CITY KNOWLEDGE:

An Emergent Information Infrastructure for Sustainable Urban Maintenance, Management and Planning

by

FABIO CARRERA

B.S. Electrical Engineering
Worcester Polytechnic Institute, 1984

M.S. Computer Science
Worcester Polytechnic Institute, 1995

Submitted to the department of Urban Studies and Planning
in partial fulfillment of the requirements for the degree of

DOCTOR of PHILOSOPHY in URBAN INFORMATION SYSTEMS and PLANNING
at the
MASSACHUSETTS INSTITUTE of TECHNOLOGY

September 2004

Signature of Author: ______________________________
Department of Urban Studies and Planning

Certified by : ______________________________
Joseph Ferreira, Jr.
Professor of Urban Planning and Operations Research
Thesis supervisor

Accepted by : ______________________________
Frank Levy
Chair, Ph.D. Committee

ACKNOWLEDGMENTS

I would like to thank Jackie, Nick, Wilma, Cino, Barbara, Alberto and many others too numerous to mention – and whose names I am likely to misspell anyhow – for bearing with me through this challenge. Without them, I probably would have led a completely different life and would have probably written a totally different dissertation (if at all) and would have ended up thanking somebody else for their help instead.

I will put more time into thoughtful acknowledgements for each of the personalized copies of this tome that I will donate to unsuspecting friends and colleagues.

I hope everyone will enjoy the reading.

Meanwhile, thank you all indiscriminately. You included. Really.

Thanks.

City Knowledge: An Emergent Information Infrastructure for Sustainable Urban Maintenance, Management and Planning

by

Fabio Carrera

Submitted to the department of Urban Studies and Planning on September 14, 2004
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Urban Information Systems and Planning

ABSTRACT

Recent advances in geo-spatial technologies, together with a steady decline in their cost, have inspired many spontaneous bottom-up municipal GIS initiatives aimed at improving many aspects of urban maintenance, management or planning. Some communities have institutionalized top-down citywide urban information systems with limited results, due to many organizational and institutional factors. Despite some encouraging progress, comprehensive urban information systems are still not commonplace and planners and decision makers still struggle to acquire the rich information that they need to conduct in-depth analyses and to make important decisions.

This dissertation suggests a plausible strategy and several practical, tactical solutions to set municipalities on a trajectory leading to City Knowledge. The concept of City Knowledge is introduced by first presenting numerous case studies ranging from the maintenance of the canals in Venice, Italy, to tree management in Cambridge to planning for Worcester, Massachusetts.

Each of the cases reveals some lessons about City Knowledge, contributing to the identification of fourteen desirable qualities and consequently to the distillation of the six foundations of City Knowledge: (1) the "middle-out" approach; (2) informational jurisdictions; (3) fine-grained, distributed data management; (4) sustainable updates; (5) information sharing and (6) interagency coordination.

The middle-out approach combines the virtues of top-down rigor and reliability with the bottom-up qualities of energy and creativity. Being an emergent system, City Knowledge leverages the dominant plan-demanded mode of data acquisition to gradually and inexpensively accumulate high-return data and to ensure sustainable, low-cost updates. It produces plan-ready information, by exploiting the self-serving and opportunistic pursuit of instant return-on-investment by frontline offices. Thanks to its emergent qualities, City Knowledge engenders unexpected plan-demanding situations, where the ability to conduct second-order analyses leads to deeper knowledge of our cities.

In the end, this dissertation proposes a paradigmatic shift by recommending that information be considered as a *bona fide* infrastructure and be consequently treated with the same attention that cities reserve to other infrastructures such as utilities and roads. It proposes that communicative planners become catalysts of this transformation away from the "hunting-and-gathering" of urban data and toward the "farming" of municipal information.

THESIS SUPERVISOR:
Joseph Ferreira Jr.
Professor of Urban Planning and Operations Research

THESIS READER:
John de Monchaux.
Professor of Architecture and Planning

THESIS READER:
Lorlene M. Hoyt.
Assistant Professor of Technology and Planning

FOREWORD

This document represents the synthesis of over fifteen years of work for the transformation of communities on both sides of the Atlantic. In writing this treatise, I have finally had a chance to be a "reflective practitioner" and look back at the lessons I learnt while trying to save my hometown of Venice, one project at a time.

This thesis proposes a middle-out approach to the accumulation and sustainable upkeep of city knowledge, through a spatially-based, self-organizing framework produced semi-independently by individual municipal departments. This emergent knowledge infrastructure will permit an efficient collection, organization, integration, distribution, use and re-use of urban facts by government agencies and will thus guarantee the constant availability of information for municipal decision-makers.

After identifying what elements of the urban realm are typically the target of knowledge-based interventions by city agencies, this research explores the feasibility and desirability of the creation of a reliable, permanent, updatable, maintainable, reusable and sharable knowledge infrastructure to support municipal operations at all levels, from the ordinary upkeep of basic city properties, to the complex envisioning of alternative futures through sustainable collaborative planning paradigms.

As part of this study, I analyze the technical, institutional and logistical obstacles that complicate the development of such an infrastructure, and I demonstrate a possible path for the gradual accrual of city knowledge. As an alternative to the prevailing top-down and bottom-up approaches to the organization of municipal information, this dissertation explores a "middle-out", parallel, distributed, emergent and self-organizing approach that promises to gradually produce a flexible, multi-purpose knowledge infrastructure on which day-to-day operations as well as long-range planning decisions can be based.

John Quincy Adams once said that he was a soldier so that his son could be a farmer, so that his grandson could be a poet[1]. As soon as civic authorities begin to treat city knowledge as they treat any other infrastructure element of the urban realm, we should witness a shift from the current "plan-demanded" mode of data collection – where we have to "fight" for each piece of data – to a more "plan-ready" approach to information management whereby we slowly grow our place-knowledge. The constant availability of information may then lead to "plan-demanding" situations in which the mere existence of well-organized urban knowledge inspires plans that may have never been envisioned otherwise.

The paradigm shift I am proposing will emancipate city officials – municipal authorities, political decision makers, urban researchers and citizens – from "hunter-gatherers" of urban data to "farmers" of urban information and ultimately to "poets" of city knowledge[2], inventing alternative futures for the transformation of communities.

1 *" I must study politics and war, that our sons may have liberty to study mathematics and philosophy. Our sons ought to study mathematics and philosophy, geography, natural history and naval architecture, navigation, commerce and agriculture in order to give their children a right to study painting, poetry, music, architecture, statuary, tapestry and porcelain."*. John Quincy Adams, 1780. I am indebted to Prof. Mark Schuster of MIT for this quote.

2 I refer here to the "strong poet" envisioned by Harold Bloom (*The Anxiety of Influence*) as discussed by Richard Rorty in *Contingency, Irony and Solidarity*, p. 24 ff..

TABLE OF CONTENTS

(page left intentionally blank)

"Beyond six rivers and three mountain ranges rises Zora, a city that no one, having seen it, can forget. But not because, like other memorable cities, it leaves an unusual image in your recollections. Zora has the quality of remaining in your memory point by point, in its succession of streets, of houses along the streets, and of doors and windows in the houses, though nothing in them possesses a special beauty or rarity. Zora's secret lies in the way your gaze runs over patterns following one another as in a musical score where not a note can be altered or displaced. The man who knows by heart how Zora is made, if he is unable to sleep at night, can imagine he is walking along the streets and he remembers the order by which the copper clock follows the barber's striped awning, then the fountain with the nine jets, the astronomer's glass tower, the melon vendor's kiosk, the statue of the hermit and the lion, the Turkish bath, the café at the corner, the alley that leads to the arbor. This city which cannot be expunged from the mind is like an armature, a honeycomb in whose cells each of us can place the things he wants to remember: names of famous men, virtues, numbers, vegetable and mineral classifications, dates of battles, constellations, parts of speech. Between each idea and each point in the itinerary an affinity or a contrast can be established, serving as an immediate aid to memory. So the world's most learned men are those who have memorized Zora. But in vain I set out to visit the city: forced to remain motionless and always the same, in order to be more easily remembered,
Zora has languished, disintegrated, disappeared.
The earth has forgotten her.

Italo Calvino, 1972
"Invisible Cities", p.15

PART I:

THE CONTEXT OF CITY KNOWLEDGE

INTRODUCTION

importance of information

Information is a fundamental ingredient in all decisions. Deeper knowledge of an issue allows one to better ponder the options and select what seems to be the wisest path to a resolution[3]. Decisions based on "gut feeling" or instinct can turn out to be just as wise, but one would generally prefer to deliberate from an "informed" position, especially when major policy or planning decisions are made, which may entail considerable expenditure of human and financial resources[4].

[informed decisions]

municipal use of information

Every single day, municipal governments make maintenance, management, planning and policy decisions that affect the inhabitants of the city or town as well as its coffers[5]. Invariably, to support these decisions, a great deal of time is spent gathering information by scouring the archives of the various departments and by leveraging personal contacts with those who are the "institutional memory" of the department[6]. Meanwhile, administrative data are gathered by city offices incessantly for specific purposes, most often connected with revenue-generation (taxes, fees, etc.) or regulatory compliance (permits, licenses, etc.)[7]. Yet these data are more often treated as "documentation" supporting a specific act or deliberation rather than as "information" that can be reused over and over in other contexts to support other municipal tasks[8].

[information vs. documentation]

urban information technology

People who are engaged in urban maintenance, management or planning use information daily, thus they have been rather receptive toward the adoption of computers to organize municipal information ever since the early days of personal computers[9]. After the first commercial Geographic Information Systems (GIS) appeared in the mid-eighties, through the nineties and until today (2004), there has been a steady increase in the use of GIS in various city departments[10]. More and more, geographic information systems are delivering on promises[11] and their effectiveness in municipal

[increased adoption of GIS]

---

3 Logic and rationality have always been part of human development, at least since historic times. For instance, Adam Smith's "invisible hand" is predicated on the unimpeded circulation of information leading to rational choices about economic options.

4 The "rational" approach to decision-making was the dominating paradigm until the mid-1950's (Caron and Bédard, p. 19), and later reincarnated as "bounded-rationalism" (Simon, 1960). Today, although communicative paradigms (Innes, 1998) are reassessing the role of information in planning, nobody denies its usefulness. See also Hammond *et al.*, 1980, 1991 and Keeney and Raiffa, 1976 and Linstone, 1984.

5 Nedović-Budić *et al.*, 2004, p. 333.

6 See, for instance,. Budić, 1994, p. 244; Nedović-Budić, 2000, p. 82 (see also reference to Arbeit, 1993); Yeh, 1999; Ghose and Huxhold, 2002, p. 5; Ferreira, 1998.

7 ICMA, *Electronic Government*, 2002. See also Ferreira, 2002.

8 Although a tad dated (1984), see the interesting conclusions that Masser and Wilson arrive at when looking at the different attitudes that are connected to the choice between what they call "hard" (quantitative information and data) and "soft" (qualitative documentation) approaches to information management. Apparently, hard approaches are more often restricted to limited domains, whereas soft approaches are more comprehensive and systematic.

9 Nedović-Budić, 2000, p. 81.

10 Masser and Wilson, 1984 ; Budić, 1994; ICMA, 2002..

11 McFall, *ENR*, New York: February 16, 2004. Also, Budić, 1994: abstract p. 244.

practices is perceived to be increasing[12]. The organizational implications of GIS adoption are often underestimated[13], though – perhaps counter-intuitively – it tends to increase the demands on staff[14] as they inevitably changes the internal dynamics of a department[15].

Geographic Information Systems

The diminishing cost of computer hardware and software[16] has led to a proliferation of homegrown GIS initiatives that address specific needs of municipalities[17]. Increasingly, there have been attempts to harness the richness and diversity of such independent activities to reduce wasteful redundancy and duplication and to maximize the synergistic potential of a coordinated approach to geospatial information management[18]. Top-down initiatives emanating from the national level[19] have led to the creation of Spatial Data Infrastructures (SDIs)[20] which in turn frequently include provisions for a core set of so-called "framework data"[21]. Simultaneously, bottom-up efforts are looking at the role of neighborhoods in the development of a finer-grained spatial data infrastructure[22] through Community Statistical Systems (CSS)[23] and the like. Citizen groups, with the aid of academia, are producing the neighborhood-level equivalent of national framework data[24], through such efforts as the National Neighborhood Indicator Partnership (NNIP)[25] that tries to consolidate indicators of urban well-being using public administrative data sources[26]. Meanwhile, technical standards that allow the exchange of spatial data are also being developed primarily by the Federal Geographic Data Committee (FGDC) of the National Spatial Data Infrastructure (NDIS)[27].

[Spatial Data Infrastructure)]

[framework data]

[Community Statistical Systems]

[Neighborhood Indicator Partnership]

[FGDC and NDIS]

Planning Support Systems

Despite all this positive ferment, the development of Planning Support Systems has been lackadaisical[28], perhaps because of planners' own inability to take full advantage of the technology[29], due to organizational,

---

12 Budić, 1994, p. 251.

13 Innes and Simpson, 1993; Campbell and Masser, 1995.

14 At least initially, as discussed in Masser and Wilson, 1984. See also ICMA survey, question 7.

15 Budić, 1994, Note 1.

16 *Idem.* Though the lower costs of technology are offset by the "hidden" costs of human resources (see Note 1).

17 Witness the innumerable cases illustrated in *URISA Proceedings* and in trade magazines such as those printed by ESRI and Mapinfo, to name but a few. See also www.nexpri.nl.

18 Like that in the "Digital Earth" initiative (http://www.digitalearth.gov/); see also Klosterman, 2001, pp. 11-13;

19 National Research Council, 1993.

20 Nedović-Budić *et al.*, 2004. See also, Masser and Wilson, 1984.

21 Tulloch and Fuld, 2001. See also the FDGC web site at http://www.fdgc.gov.

22 Talen, 1999.

23 Ferreira, 2002.

24 At an even finer-grain than that discussed in Tulloch and Fuld, 2001, but with similar upward scalability.

25 Sawicki and Flynn, 1996. See also Neighborhood Knowledge Los Angeles. (http://www.nkla.org) and connected efforts.

26 For a complete listing of administrative data sources see also Coulton, Nelson and Tatian, 1997.

27 See http://www.fdgc.gov

28 Budić, 1994, p. 245; Innes and Simpson, 1993; Nedović-Budić, 2000; Geertman and Stillwell, 2004, p. 307.

29 "The most important impediment to the implementation of GIS in planning may be the planners themselves", Innes and Simpson, 1993, p. 232.

institutional, sociocultural issues[30], or perhaps simply because planners are too preoccupied with gathering useful data[31] for the plan at hand[32] to have time to dedicate to the development of tools beyond the mere computerization of manual tasks[33]. In fact, with the exception of a small number of well-funded and established comprehensive top-down efforts[34], and despite the operational benefits in terms of effectiveness that GIS technology affords[35], progress has been slow and "the application of PSS is currently still in its infancy"[36]. The puzzling lack of effective[37] penetration of GIS technology in local planning operations at the municipal level can be attributed to many factors, but "people" issues are always the main culprits[38], followed by issues of technological complexity and cost. Another major stumbling block are data[39]. Planners are voracious consumers of information, but they produce very little new information themselves[40]. "The use of systematic models to bring knowledge to bear on assessing the probable outcomes of planned actions is not a general practice in planning[41]". Comprehensive systems to bring together multipurpose geographic information systems for second-order spatial analyses are needed but not quite available to planners[42].

local GIS trategies

Current trends indicate a move toward the development of local geographic information strategies[43] to capture the finer grain of urban data that community statistical systems require[44]. There begins to be also a discussion about the importance of supporting the development and maintenance of local databases[45] so that a distributed municipal information systems can be assembled without redundancies from a series of networked systems, connected via the World Wide Web[46], and developed in a coordinated manner[47], based on agreed-upon structures, processes and

---

30 Campagna and Deplano, 2004, p. 35.

31 Nedović-Budić, 2000, p. 82.

32 Masser and Wilson, 1984 (see in particular Table 8 on p. 421).

33 In essence, there is little of Zuboff's (1991) "informating", most efforts going toward "automating".

34 Like those in Singapore (Arun and Yap, 2000) and in Vienna (Wilmersdorf, 2003 – municipal site in German at www.wien.gv.at/wiengrafik/suche.htm) to name a couple.

35 Budić, 1994, p. 251.

36 Geertman and Stillwell, 2004, p. 307.

37 As measured according to Budić's (1994) indicators of "Operational Effectiveness" and "Decision-Making Effectiveness".

38 Nedović-Budić, 2000, p. 82; Nedović-Budić and Pinto, 1999, p. 60.

39 Yeh, 1999.

40 I think Yeh and Webster (2004) meant to say the same thing, though their statement that "urban planning is a process that *generates* a lot of information" (italics added) may be easily misread.

41 Harris, 1999, p. 324.

42 As far back as Innes and Simpson, 1993, p. 232 and as recently as Brail and Klosterman, 2001, "the always imminent revolution" (pp. 2-3). Most recently Geertman and Stillwell, 2003, p. 8 and again Geertman and Stillwell, 2004.

43 Craglia and Signoretta, 2000; Tulloch and Fuld, 2001.

44 Ferreira, 2002.

45 Nedović-Budić, 2000, p. 87.

46 Kelly and Tuxen, 2003.

47 Following a Geospatial Information Management Plan, as suggested by Keating *et al.*, 2003.

policies[48] for the creation, archival, maintenance, updating, removal and sharing of urban data[49].

the challenges

As Keating *et al.* put it, a first challenge "lies in striking a balance in the degree of centralization of data storage, administration, and procedural control while serving the needs of the community […]" [50]. Moreover, according to the University Consortium for Geographic Information Science (UCGIS),

> *"As the variety of geospatial information and data resources increases each year, the demand for understanding and building sustainable information and knowledge structures remains a critical research challenge for the geo-spatial information community."*[51]

So the problem today is not the availability or capability of technology for planning, but rather the availability of "good" fine-grained, up-to-date data[52]. What's also missing is an active pursuit of the creation of systematic storehouses of urban knowledge. According to the established "Communicative Action" paradigm[53], the way forward is to embed in the planning community – and in municipal administration in general – an innate appreciation for the value and importance of information at any level of urban maintenance, management and planning. This shift of mindset would enable a sea-change to take place in how cities collect and organize information.

the solution: City Knowledge

This dissertation, as its title implies, addresses directly the aforementioned research priority of the UCGIS. It specifically focuses on the organizational, institutional, technical, logistical and financial mechanisms whereby appropriate local authorities could systematically build up a comprehensive set of *framework* (and later also *thematic*) datasets and map layers. The gradual, but systematic compilation of all the disparate datasets accumulated by a wide variety of government and non-government organizations is what I term the *City Knowledge* approach. I propose to "grow" this knowledge from the middle-out, integrating a top-down approach to standardization[54], with a bottom-up approach to neighborhood-scale ("atomic") data accrual[55].

Comforted by the latest developments in the definition of "core" framework datasets[56], and by the current trends among researchers toward a

---

48 Nedović-Budić and Pinto, 1999, pp. 56-59.

49 Nedović-Budić and Pinto, 1999 and 2000; For more information about these matters, consult the FGDC web site at www.fgdc.gov . See also the Geographic Information Resource Management short term research priority for the UCGIS at www.ucgis.org .

50 Keating *et al.*, 2003, p. 35.

51 Shuler, 2003. *UCGIS Research Priority*. Last accessed 8/20/04 at http://www.ucgis.org/priorities/research/2002researchPDF/shortterm/g_resource_management.pdf.

52 See Budić, 1994, p. 252, Table 4, under Operational Effectiveness indicators.

53 See Innes, 1998.

54 As in Craglia and Signoretta, 2000 and Nedović-Budić and Pinto, 2000.

55 As in Talen, 1999 and Ferreira, 2002.

56 Tulloch and Fuld, 2001. Harris, 1999, p. 330; Nedović-Budić *et al.*, 2004; See also FDGC web site at http://www.fgdc.gov/framework/framework.html. Last accessed 8/20/04.

"local" approach to the creation of comprehensive municipal information systems[57], using smaller and smaller jurisdictions to allocate responsibility over specific data layers[58], I am confident that my City Knowledge approach will prove to be a valid contribution to the creation of emergent, comprehensive, updatable municipal information infrastructures.

Research Question

This dissertation will attempt to answer the following question:

What gradual and unobtrusive approach can cities adopt to create a fine-grained, sustainable, affordable, updatable, comprehensive information infrastructure on which to base decisions about urban maintenance, management and planning?

Methodology

I have grounded my approach on the Case Study methodology[59] based on my own substantial experience in planning-related data collections and spatial analyses in Venice, Italy and in the greater-Boston area. In fact, the primary method I have applied is the "reflective practitioner" approach[60], which allows me to use my own cases as the foundations for the lessons I extract, without being accused of bias. Although I can cite dozens of cases to support my claims, this dissertation suffers from the typical limitations of case-study research[61]. While case-study methods do not allow simple generalizations, they do provide the benefit of richer and deeper datasets, which are not to be taken as "samples of one", but can nonetheless allow "analytical" (as opposed to statistical) generalizations[62], such as the ones I will present hereafter.

My arguments will be based on direct experiences, supported by appropriate references to similar cases in the literature (when available) and framed around logical propositions based on factual evidence. In the spirit of the emerging paradigm of planning as communicative action[63], I put forth this thesis as the synthesis of almost twenty years of *praxis* to contribute to a process that I hope will make "city knowledge" a household phrase when information[64] is treated as any other infrastructure on which our towns can count to support all other municipal activities.

57 Craglia and Signoretta, 2000.

58 Nedović-Budić and Pinto, 2000, p. 468.

59 Yin, 1994.

60 Schön, 1983.

61 Nedović-Budić and Pinto, 2000, p. 471.

62 Caron and Bédard, 2002, p. 22, citing also Yin (1989), p. 21.

63 Innes, 1998, p. 52 and especially Notes 1 and 2.

64 I refer here to what Innes (1998) calls "scientific" , "technical", or "formal" information and knowledge. It is possible, in theory, that City Knowledge could embrace the other types of information that Innes discusses: personal experience, stories, images, representations and intuition, though these will not be discussed herein.

Intellectual Contribution

To put it in Nietzsche's terms, and paraphrasing philosopher Richard Rorty[65], in this dissertation I am simply putting forth a set of contingent metaphoric redescriptions that I hope will contribute to the process of municipal knowledge-making in ways that go beyond the specific suggestions, techniques and methods I propound herein. I hope the latter will be useful in the immediate and may be inspiration for transformation of communities around the world. However, beyond that, I also hope that this paper may in some measure penetrate into the planning consciousness in the subtle, yet powerful way that the communicative planning paradigm champions. I would be gratified if only a fraction of the "new descriptions[66]" I put forth herein could "also strike the next generation as inevitable" and thus become woven into the fabric of what municipal agencies do. Indeed, as Innes put it "*when information is most influential, it is also most invisible*"[67].

In short, I hope some day City Knowledge will be utterly undetectable, not because it has failed miserably and has been long forgotten, but because it will be completely ingrained into the collective consciousness of those who are in charge of urban maintenance, management and planning.

Problems

lack of infrastructure

My personal experience in urban studies and planning in my hometown of Venice, Italy, as well as more recent forays in Massachusetts, have lead me to realize that cities often lack a comprehensive and systematic "knowledge infrastructure" on which to base planning decisions, from the grand urban design projects to the more mundane municipal maintenance tasks[68]. What I found to be the prevalent mode of functioning of the various branches and departments of a modern city is a form of "ad hoc-ism" whereby data are collected for specific purposes and then quickly forgotten or stored in inaccessible places, unbeknownst to any other department or even to other personnel in the same department[69].

ad-hoc data collection

lack of dissemination and re-use

Although some systematic data collection takes place, mostly for regulatory or revenue-generating purposes (such as for permits, licenses, property assessments, and the like), even these data are often hard to obtain or utilize, both internally – by the rest of the municipality – and even less so externally – by academic scholars, independent researchers or planners. Frequently, access to important information is made possible only through personal connections and by means of "under the counter" transfers which bypass the official channels that otherwise would render the dissemination of data virtually impossible.

redundancy and waste

Due to the dearth of coordination between departments, I have personally encountered several situations whereby the <u>same</u> data were being

65 Rorty, 1989. See particularly p.27 ff. According to Nietzsche truth is "a mobile army of metaphors".

66 *Ibid.*, p. 29.

67 Innes, 1998, p. 54.

68 Confirmed by Geertman and Stillwell, 2004 and Brail and Klosterman, 2001.

69 Evans and Ferreira, 1995; Ferreira, 1998.

collected by different departments almost simultaneously. Moreover, I have also personally witnessed how a single department will pay external consultants to collect the same type of data multiple times, over the course of a few years[70]. Redundancy and waste seem to be endemic when it comes to municipal data collection, due to the apparent lack of long-term planning and standardization. The pursuit of a systematic accumulation of comprehensive urban data has not generally been part of municipal policy for a variety of reasons and this has inevitably led to a number of inefficiencies along the way[71].

lack of a systematic approach

low effectiveness of GIS

So, despite the diminishing cost and increased adoption of GIS technology in municipal government, the impact of these technologies is still very minor[72], since applications have been thus far limited to first-order tasks, where geospatial information technology (GIT) simply replaces procedures that were previously carried out by hand[73]. As Caron and Bédard (2002) surmised, GIT projects tend to follow a "garbage-can" type of development, whereby solutions are "technology-driven and looking for problems to solve, and an *a priori* identification of possible solutions seems to determine the problem formulation"[74].

[technology-driven]

However, as is often the case, these shortcomings in the development of comprehensive municipal information systems can also be seen as opportunities waiting to be seized, as described below.

Opportunities

taming the urban infoscape

The perceived vastness and complexity of the urban infoscape can be tamed with today's technology using a "divide and conquer" approach, as my own experience and that of many urban planners around the world demonstrates. Practically all individual tasks or functions that a municipality performs have been automated and/or informated[75] somewhere by now. In short, we have the technology and we know what we want to do with it.

framework data

neighborhood indicators

Meanwhile, the "framework[76]" and "indicator[77]" movements are clarifying what urban information is useful to collect and what public datasets are available to infer the needed information from[78]. Community Statistical Systems[79] as well as municipal information systems can harness these datasets and make them widely available. In short, we know what data we need and where to get it.

getting it together

It is possible today to envision being able to gather, organize, maintain, use and re-use the datasets that would feed a comprehensive

70 See for example Flynn *et al.*, 2003.

71 This can be explained if we consider municipalities as "multidivisional" or "M-form" firms, wherein rewards and decision-making exist within isolated, uncrossable departmental boundaries, which in turn creates a dis-incentive for the adoption of coordinated or "enterprise" information systems (Singh, 2004). Caron and Bédard, 2002, arrive at similar conclusions based on empirical data (p. 33).

72 Budić, 1994.

73 Innes and Simpson, 1993, p. 233; Reeve and Petch, 1999.

74 Caron and Bédard, 2002, p. 32.

75 Zuboff, 1991.

76 Tulloch and Fuld, 2001; Nedović-Budić *et al.*, 2004. FDGC at http://www.fgdc.gov/framework/.

77 Sawicki and Flynn, 1996.

78 Coulton *et al.*, 1997.

79 Ferreira, 2002.

[comprehensive municipal information system]

municipal information system. It remains a tall task, but it is no longer an insurmountable one in this day and age. The transaction costs and complexities associated with geospatial data collection continue to decline and the impacts of technological[80] and organizational change[81] have been understood and can be factored into any economic calculation of benefits vs. costs. In short, we know how we could assemble the system and how much it will cost us. We may even be able to afford it.

So, if all the pieces seem to be in place, why aren't cities initiating the process that will get them a comprehensive municipal information system?

the elusive solution

What remains elusive is how to put together and finance a well-oiled and organized machinery that will keep all of the possible datasets organized and up-to-date so that ever-improving applications can run each aspect of a municipal operation in an efficient and cost-effective manner that can be maintained over time, migrated across all of the foreseeable technological advances and in the face of all of the changes that the city will undergo in the future.

This dissertation addresses precisely these hurdles that are preventing the full implementation of City Knowledge systems in municipalities.

capturing permanent features

[slow change]

[systematic accumulation]

Fortunately, a majority of the characteristics that make up the physical city change very slowly, if at all, and are thus amenable to a gradual and systematic collection effort, the bulk of which would only have to be conducted once. Until now, the apparent complexity of the gargantuan task of collecting and organizing such a multidimensional body of information has discouraged a wholesale approach to the accumulation of city knowledge. Today, however, technical tools that can facilitate the recording and archiving of most, if not all, of the idiosyncratic features of the urban landscape have finally become widely available and affordable, making it possible to realistically envision how cities could begin to accumulate this wealth of information about themselves. This paper proposes a possible pathway toward this end[82].

low-cost technology

lower transaction costs

The marginal returns one can obtain by systematically (or even opportunistically) collecting and archiving finer-grained urban data are beginning to outweigh the transaction costs that such refined data collections would entail. The case studies I present herein showcase the actual extent of the first- and second-order returns that we were able to exact from our progressively accumulated knowledgebase, when we used our data for a specific immediate purpose and then later re-utilized these same data for new research on a different topic. This is an important aspect of my dissertation since it demonstrates the obvious gains attainable by "automation", as well as the hidden and somewhat unforeseeable advantages

80 Evans and Ferreira, 1995; Nedović-Budić and Pinto, 2000, p. 466.

81 Nedović-Budić and Pinto, 2000, p. 467.

82 The pragmatic process I propose – starting from low-hanging fruits and expanding gradually as needed – agrees both with theoretical approaches proposed by Innes and Simpson (1993) and Nedović-Budić and Pinto (1999), but also with the empirical conclusions drawn by Keating *et al.* (2003) based on the Cerro Grande fire.

one could obtain by "informating" through the development of a City Knowledge system[83]. These case studies will make evident both the immediate "value" of a specific GIS application, but also the "value-added" bonus that one often reaps from sowing seeds of city knowledge around town at every opportunity.

recent technological advances

This thesis takes advantage of the window of opportunity created by recent advances in Geographic Information Systems (GIS) , the World Wide Web (WWW) and other information and communication technologies (ICTs) to propose a specific institutional and organizational, as well as technical, approach that will enable cities to gradually and non-traumatically accrue and maintain an exhaustive, comprehensive, flexible, reliable, multipurpose and sharable knowledge-base.

knowledge accrual

Such a distributed, grass-roots, self-organizing system could effectively bring into the world of urban planning what some of the most creative and brilliant scholars in other disciplines have also been trying to harness: the power of *emergence*[84]. One of the core elements of my approach is the reliance on a middle-out tactic for the gradual accrual and permanent upkeep of urban information. In essence, I do not propose an all-encompassing system, but simply suggest mechanisms that will lead to such a system eventually and organically, in true emergent fashion.

emergence

Goals

inventory information needs

With this research, I make the case that City Knowledge is a resource to be fostered and maintained as any other infrastructure of the city is. Through numerous examples, I illustrate the variety of information that municipal departments need in order to fulfill on their obligations toward the citizens of the city or town. Investments in City Knowledge should therefore be viewed as capital outlays and steps should be taken to ensure that taxpayer money is well spent. Just as a city would not dream of rebuilding its sewer system over and over again (or to create two parallel sewer systems), so too each municipality should ensure that information about important urban elements is not lost, inaccessible, underutilized or repeatedly and redundantly collected.

establish importance

address obstacles

This dissertation proposes ways in which cities can overcome the most common obstacles – be they organizational, financial, technical, psychological or logistical – that are hindering the institutionalization of a knowledge infrastructure.

identify desired qualities

Using real-life case studies and recent literature, I identify the desired qualities of a City Knowledge system, which include being affordable, permanent, sustainable and reusable. These qualities are ideals, but also guidelines. Only well-established, full-fledged, comprehensive City Knowledge systems will be well positioned on all fourteen of the quality dimensions I discuss herein. In fact, achieving at least some measure of

83 Zuboff, 1991.

84 Johnson, 2001.

success on some of these factors should also facilitate the implementation of City Knowledge[85].

propose approach

Then, I propose unobtrusive mechanisms that can be put in place at the ground level (or "frontline"), to gradually, but systematically, build up a body of knowledge about the city in an emergent fashion. Using this infrastructure, plans can be devised when needed, but more importantly, with this information readily available, municipal agencies can better serve their constituent taxpayers who are footing the bills and are living, day in and day out, with the positive or negative consequences of the urban management practices adopted by their town.

support with case studies

The feasibility of the institutionalization of a self-organizing city knowledge infrastructure, built from the ground up through a middle-out approach at the level of municipal departments, is supported by numerous examples from personal observation of not-entirely-foreseen outcomes that spontaneously emerged after an initial organization of routine operations in cities as far apart, culturally and geographically, as Venice, Italy and Worcester, Massachusetts. These prototype applications have enabled the target city departments to, first of all, organize the information that is needed for their day-to-day operation, but have also demonstrated true emergent qualities, when the self-serving micro-behaviors that have improved the individual departments' operations have later combined to produce unexpected macro-benefits that have exceeded the mere sum of the parts.

organize routine operations

demonstrate emergent qualities

assess usefulness to planners

When possible, I have investigated how such solid, fine-grained and rich datasets of fungible information can be used by city planners, who most frequently need to gather a variety of information from disparate sources, across department boundaries[86].

distill and generalize

From my case studies, I have distilled a generalizable approach that I then applied in Massachusetts – Cambridge, Newton, Boston, Quincy and Worcester. I have also identified the sectors of municipal operations and services that are more amenable to these tactics in the short- and medium-term. All along, I have focused specifically on the advantages and opportunities that exist in the areas of urban maintenance, management and planning.

aintenance, management and planning

explore planning implications

My dissertation also attempts to demonstrate how the self-emergent process can be enhanced by adopting a "planning mindset" in order to enrich the datasets with teleologic parameters aimed at maximizing the information content of collected data for little or no added cost. This "enrichment" of the datasets in turn facilitates sharing and also guarantees the cost-effectiveness of field campaigns both for the collection of baseline data as well as for subsequent periodic updates.

spontaneous emergence of plans

At the same time, this research also aims to illustrate with real examples that, under certain circumstances, the availability of rich city knowledge can and will promote the spontaneous surfacing of the need for

85 A good approach would be to give precedence to the qualities that best map onto Rogers' (1983) five principles for success in innovation: simplicity, observable benefits, relative advantage, ability to make small trials and compatibility (in Innes and Simpson, 1993, p. 232).

86 Yeh, 1999; Innes and Simpson, 1993; Brail and Klosterman, 2001.

new plans that are dictated by the mere existence of such information, in true emergent fashion.

Document Structure

This document is divided into five main parts. At the beginning of each part, a quote from some famous planning personality sets the stage for the topics covered therein.

I: The Context of City Knowledge

The rest of this introductory section (Part I - *The Context of City Knowledge*) is dedicated to the exposition of background materials useful to contextualize the rest of the paper.

II: The Path to City Knowledge

Part II, entitled *The Path to City Knowledge* chronicles a selected sample of my personal experiences in Venice, Italy, as the founder and director of the *Venice Project Center*, where every year since 1988 dozens of students from the Worcester Polytechnic Institute (WPI) have completed research projects exploring different aspects of the functioning of my hometown. These are the core "cases" in which my students and I were forced to deal with many practical, as well as logistical issues pertaining to the collection, archival, manipulation, analysis, and presentation of municipal information in the most diverse areas of urban maintenance, management and planning. What we learnt through these experiences in Venice is highlighted at the end of each section, in order to build support for the conclusions that are later summarized in Part IV. The last chapter in this part showcases how City Knowledge can permit advanced analyses – in this case about the effects of flooding to the physical integrity of the city – that would be unthinkable without an acquired storehouse of information. After years of data accumulation, the availability of "plan-ready" information paid off handsomely in the example illustrated and in several other studies where we were able to quickly get to second-order thinking, leveraging our repository of knowledge for advanced analytical tasks.

III: The Devolution of City Knowledge

The Devolution of City Knowledge (Part III) presents some brief excerpts about my experiences on U.S. soil, primarily in Cambridge and Worcester, which demonstrate how some of the concepts developed in Venice were exportable to "normal" cities – with cars instead of boats. Other examples from Massachusetts are also mentioned throughout the paper in a number of footnotes and references. It is thanks to these instances of successful city knowledge transplants that I began to consider a "generalized" approach to city knowledge. This section also allows me to present a few more "cases" that highlight some additional significant "lessons" about City Knowledge that had not yet emerged in the Venice section.

IV: The Essence of City Knowledge

In Part IV – *The Essence of City Knowledge* – I step back and reflect on the lessons learnt throughout and I formulate my theory of City Knowledge. I begin by distinguishing the containers – the physical structures that compose the tangible city – from the contents – the ever-changing activities that people engage in within the material containers. The second chapter introduces the premises on which the City Knowledge approach is based, and the third chapter acknowledges and analyzes the *obstacles* that continue to prevent city offices from adopting the City Knowledge approach The fourth chapter on the *qualities* of City Knowledge serves as guidance for the features that we would like to see in a future City Knowledge system. The core of my theoretical contribution, culled from my *praxis*, is in the last chapter dedicated to the six pillars that form the *foundations* of City Knowledge: the middle-out approach, informational jurisdictions, distributed knowledge, sustainable updates, interagency coordination, and information sharing.

V: The Emergence of City Knowledge

Part V – *The Emergence of City Knowledge* – discusses the *emergent* qualities of the approach. The second outlines a middle-out pathway to begin building municipal knowledge infrastructures. Finally, in the last chapter of this dissertation, I look at the future and discuss ways to mainstream these concepts into day-to-day municipal operations, how to spread them to other municipalities, as well as how to extend them and expand them to be even more useful.

MUNICIPAL MAINTENANCE, MANAGEMENT AND PLANNING

This chapter is devoted to introducing what municipalities do as institutions, how they structure themselves to provide services to citizens, how they pay for these services and how exactly the services are performed either by civil servants or by outside agencies, contractors or consultants. All of these city activities are based on the availability of some information and knowledge which are the topic of later chapters.

MUNICIPAL ADMINISTRATION

Whether they are located in Africa or Asia, in Italy or in the U.S., in Massachusetts or Iowa, at the most fundamental level all municipal governments, in order to be viable and to provide a civilized environment for their inhabitants, will always attempt to offer very similar services to their citizens. Of course, some municipalities carry out these duties more diligently than others, and some are endowed with more human and financial resources to provide the needed services than others are. This section explores how cities and towns function both administratively and operationally, as they provide fundamental services for their inhabitants. Although many considerations are applicable to municipalities anywhere in the world, this chapter focuses exclusively on the U.S. situation, except where otherwise noted.

what cities do

Cities[87] are administrative entities set up to provide community services to people who live within their boundaries (and occasionally to outsiders too). Numerous guides, handbooks and manuals[88] directed to city officials are published every year, by organizations like the *International City/County Management Association*[89] and others. Frequently, the chapter breakdown inside these publications follows the structure of an archetypal city government[90]. After all, there are only so many functions that a city government is called upon to perform. Roughly speaking, one could divide the main administrative areas of a city into:

- Political and executive branches
- Internal Services
- Public Health and Safety
- Culture and Leisure
- Education
- Physical Services

[87] Hereafter, whenever the term "city" is used, it should be interpreted to stand for *municipality* and thus to also include towns or villages; the difference being merely the size of the local unit of government. Smaller units will not necessarily have all the departments and divisions that a larger municipality may have, but the fundamental services and functions should remain unchanged. Neighborhoods can also be real administrative entities that provide services that a municipality provides, albeit at a more local level.

[88] Banovetz *et al.*, 1994; Hoch *et al.*, 2000.

[89] See http://www.icma.org. Last accessed 8/20/04.

[90] See, for example Kemp, 1998.

political and executive branch

The political and executive branch includes, for example, the City Council, City Manager and City Clerk who are selected by the citizens to administer the city and make everything else work. Numerous combinations of political and executive bodies exist within the U.S. and in other countries. The bottom line is that the citizens always run their own city either directly or by proxy. In many New England municipalities, town meetings are still used to let citizens express their views on policies by direct vote, without any elected intermediaries. Most of these towns elect selectmen to ensure the town is run according to popular sentiment, with a "first" selectman usually heading the board of selectmen. Many New England towns up to about 20,000 population are run with these "amateur" systems[91], where the pseudo-mayors are frequently local business owners who run the towns part-time, barely compensated for the time they put to the service of the community. These interesting forms of local democracy are common in New England and rare anywhere else. Similar direct systems of local government are en vogue in "tribal" communities in developing nations where councils of elders decide on major communal issues. The Afghan "loya jerga", albeit a patriarchal and conservative system with limited democratic access, represents the equivalent of a New England town council of yore.

[town meetings]

[selectmen]

[first selectman]

In the U.S., cities may have an elected political mayor who heads the city council and runs the city through appointed political department heads. European cities by and large follow this system, which in the U.S. is sometimes termed the "strong mayor" model. In the U.S., Boston, Mass. and Providence, R.I. adopt this model, as do New York City, and most major metropolitan cities like Los Angeles, Chicago, Miami and Atlanta.

[strong mayor]

[city manager]

Another model, involves the figure of the "City Manager", who is hired to run a city without being "politicized". A "weak" elected mayor usually heads the city council and interprets the will of the citizenry to the city manager, who translates the political will into concrete programs and actions. The city of Worcester, Massachusetts has this sort of a structure with a political mayor with no first-hand executive power, but with an electoral mandate to set the direction and strategies for the city, paired with a hired city manager who is the chief tactician as the principal executive implementer of the city's strategies[92]. "Strong City Manager" systems are in place in Cambridge, Massachusetts, and in other mid-sized cities that have elected mayors as well.

bottom-up flow of information

Regardless of the precise form of city government, the city's top administrators are fed knowledge distilled from information put together by the lower echelons, who in turn rely on data collected by the city departments and divisions who deal with urban issues on a daily basis and have a better sense of the "pulse" of the city at any point in time[93].

internal services

Internal services include Legal services, Finance, Human Resources, and other functions that ensure that everything the executive branch

[91] Nelson, 2002, p. 2, Table 2.

[92] Interestingly, there is a movement afoot to change Worcester's executive branch into a "strong mayor" system.

[93] Reeve and Petch, 1999, p. 25.

deliberates actually has the wherewithal to be implemented successfully. In many ways, these are standard functions found in most businesses and organizations. Just as is true in the world of commerce, city operations need to be carried out inside the bounds of the law, within the capabilities of the available workforce, and under the tight confines of budgets, which in U.S. cities are almost entirely funded by local real estate taxes.

revenues

Revenue generation through taxes, fees and fines is a very important function for any municipality or township. Some State and Federal monies are passed down to local governments through a variety of programs, like the Community Development Block Grants (CDBG's), but – at least in the U.S. – cities need to supplement an increasingly larger portion of their operating budgets through independent revenue, generated for the most part by local real estate and business taxes and a variety of fees (and fines) attached to administrative acts (and infractions) as well as to specific essential services, such as garbage removal and sewage collection and treatment.

public health and safety

Public health and safety, which comprises Fire, Police, Health and Social services represents the safety net that citizens need in order to deal with "bad stuff" that happens to the more unfortunate people, even in the best of communities. In the U.S., this area of municipal operations often represents the second largest expense a city incurs (after education), mostly due to the salaries paid to the policemen, as well as the high cost of the purchase and maintenance of equipment, such as cruisers, ambulances (often donated by a local charity, such as the Lions or Rotary Clubs), and fire engines (also frequently paid for by donations and fund-drives). In large cities, the police chief is frequently one of the highest paid officials on the payroll of the municipality. Firemen, on the other hand, are frequently volunteers, as are often the ambulance drivers as well, especially in smaller cities and towns. In the U.S., fire squads of professional, paid firemen are maintained only in large metropolitan areas of more than 50,000 inhabitants[94]. Only the larger cities have a municipal hospital, which represents another major expense for the city. Despite heavy subsidies from the state and federal governments, city-owned hospitals can be quite a drain on the city coffers. Other social services, such as eldercare or provisions for people with physical and mental disabilities may also be financed by local revenue, depending on the size of the needy population. Public housing is another form of social service that is made available to needy families and individuals to guarantee decent shelter to the working poor and to other disadvantaged people. Community Development Corporations (CDCs) are often created to manage public services (primarily public housing) at the neighborhood level.

culture and leisure

Culture and leisure for most cities include such operations as the running of the municipal library and the management of the various sports leagues in the local park system. Events, such as parades, concerts, fireworks, commemorations, celebrations and other ephemeral public

---

[94] In Italy, only the alpine autonomous, German-speaking region of *Süd Tyrol* (*Alto Adige*) maintains a tradition of volunteer *Feurwehr* troops in each mountain village, much akin to the tradition that is still widespread in smaller New England communities.

displays of community spirit[95] are often funded both by volunteer donations and by a financial contribution by the town or city.

education

Education is a municipal responsibility from the financial standpoint, though the actual running of school districts is delegated in the U.S. to elected boards of education and to elected or hired superintendents of schools. Sizeable sums of money trickle down from the State government to aid the needier municipalities (i.e. the ones with inadequate tax bases). Education is always the biggest drain of local funds, mostly because of the large number of teachers and employees needed to operate schools from Kindergarten, to Middle Schools (or Intermediate schools[96]), Junior High and High School. Teaching the "three R's" is an expensive activity that falls mostly on the lap of town treasurers. Real estate taxes – somewhat indiscriminately, as libertarians would point out – fetch the needed funds from all homeowners, whether they have children in school or not.

physical services

The rest of the city administration consists of such services as Parks and Recreation, Planning, Facilities Maintenance, Transportation and Public Works that simply make the "physical" city perform smoothly. These are perhaps the most "visible" services the city renders, as evidenced by the ubiquitous yellow trucks that plow snow, collect leaves, sweep streets, fix potholes, mow lawns and generally keep basic city operations in order year after year.

Despite some overlaps, dividing the functions of a municipality into these six areas offers a simple and useful starting point for a discussion of what a city does and how information is used in every area.
This dissertation dwells primarily on the types of information that are specific to city maintenance, maintenance and planning, hence the first two areas above – political and executive activities and internal services – are not analyzed in great detail, especially when many of their information needs are not too different from those of private business enterprises. However, the former area (politics and executive administration) plays a key role in dictating the type of information that needs to be distilled into knowledge to allow decision makers to make wise choices. The latter area (finances, legal and HR) is important to the degree that its departments generate a flow of administrative accounting paper trails that may be tapped into to feed the data streams that make up the foundation of City Knowledge that is the topic of this dissertation.

[95] Discussed in Carrera (1998) as a reaction to early drafts of Schuster in *Imaging the City* (2000).

[96] Some school systems, such as in Spencer, MA, where I live, have four levels, whereby Elementary school goes only up to third grade, followed by Intermediate that caters to 4th, 5th and 6th graders, followed by Junior High reserved for 7th and 8th and finally High School for the last four grades (9-12th). Every new building, grade and class, requires additional resources.

DATA, INFORMATION, KNOWLEDGE AND ACTION

data, information and knowledge

Whereas in the previous chapters I have used the terms "data" and "information" synonymously, there seems to be some consensus on a hierarchy of "types of information", from data to information to knowledge (some scholars, such as Klosterman, even add a fourth level of intelligence)[97]. In this context, *data* would refer to raw facts, both quantitative and qualitative, *information* would pertain to data manipulated and organized in a meaningful form, and *knowledge* relates to "understanding based on information, experience and study".[98] *Intelligence*, a term which agencies such as the C.I.A. frequently use to refer to "top secret" information, is sometimes considered to be the application of knowledge to guide behavior[99]. In this paper, the terms data, information and knowledge are only occasionally used interchangeably, as synonyms, especially in quoted references from third parties[100]. For the most part, though, I have attempted to use them appropriately throughout this dissertation.

from knowledge to action

The figure below reflects the aforementioned taxonomy and also summarizes the typical processes involved in the decision-making that leads to action on the part of municipal agencies. Knowledge is supported by hard facts, based on data, which are organized into information, but it also taps into the "fuzzier" realm of "experience", whereby social, institutional, legal, tactical, methodological and theoretical factors play a role in the shaping of final decisions leading to actions.

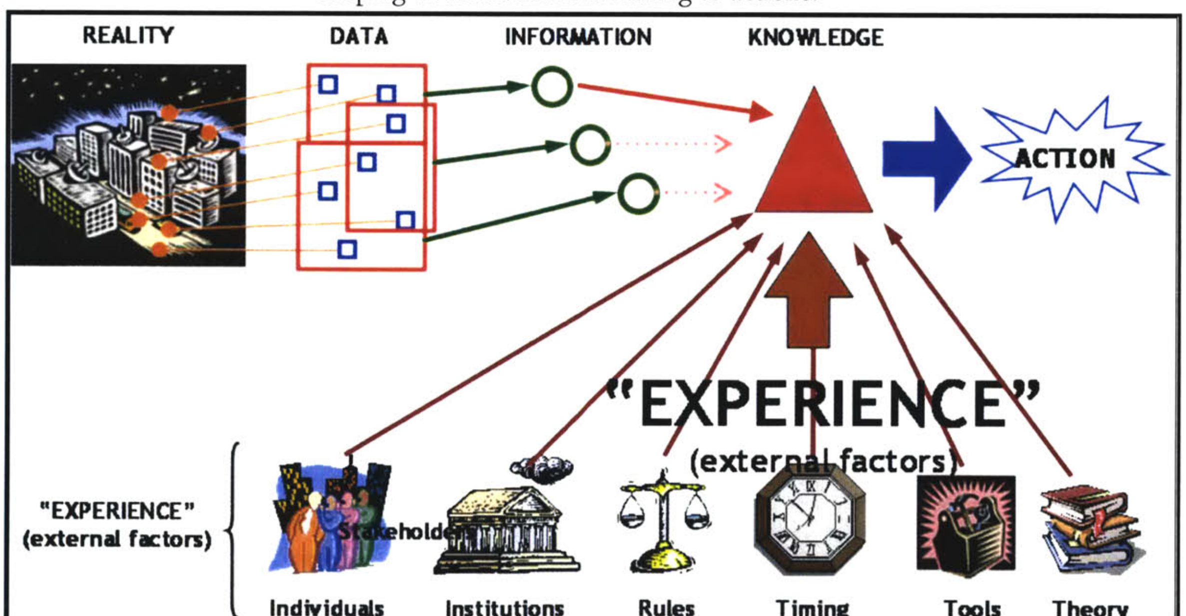

[97] Klosterman, Richard E.. 2000. "Planning in the Information Age", p. 42.

[98] *Idem.* See also Laurini, *Information Systems for Urban Planning*, pp. 41-42 and p. 149.

[99] *Idem.* I am not convinced by this fourth level, since it seems to me that it basically labels a goal which is implicit in the gathering of knowledge more than a "type" of information…

lack of data (hard facts)

The "hard facts" are unfortunately not as available as one would imagine them to be. In fact, many distinguished planners of the past (such as Olmsted[101] and Geddes[102], for example) as well as many contemporary observers of urban affairs[103] clearly point out that we are not doing a really good job of knowing our cities. In summary, as Yeh succinctly put it: "[t]oday, the main constraints on the use of GIS in urban planning are not technical issues, but the availability of data, organizational change, and staffing"[104].

My personal experience confirms these views. Despite the relative "permanence" and "immutability" of the physical elements composing our urban realms, knowledge of our cities is not as developed as one would hope it to be. Although data are gathered daily for a variety of reasons, information is not necessarily obtained as a consequence and knowledge is therefore hardly augmented in the process. The seemingly subtle differences between data, information and knowledge are quite apparent in the fields of urban maintenance, management and planning. Data are all too frequently collected to satisfy very specific needs. They are mostly treated as "documentation" and therefore they are rarely organized into information that can be used for other purposes. Thus, they hardly ever contribute to the creation of knowledge on which decisions and actions can be fruitfully based.

information vs. documentation

Planning Professionals need to navigate the inherent "political realities" of the cities in which they operate since the balance of power within the municipal hierarchy affects a planner's access to information and knowledge[105]. Planners, by the nature of their trade, have to learn how to collect, use and share knowledge using a variety of "rational", "scientific" or "formal" rhetorical forms[106]. They must be "intelligent" in their use of such knowledge and sensitive to the circumstances in which they operate[107]. Most importantly they need to bridge the gap between politicians and citizens[108] while they themselves straddle the fine line between comprehensive rationalistic theories, good on paper but highly impractical, and more realistic muddling-through[109] by incremental empiricism and communicative action[110]. Compromises between the textbook and the ad-hoc approaches have also been proposed, such as the "mixed-scanning" method put forth by Etzioni[111], but the business of planning is and will

communicative planning

100 Often, the authors I quote will use "information" to mean "data, information and knowledge".

101 Olmsted, 1913.

102 Geddes, 1911.

103 For example Geertman and Stillwell, 2004.

104 Yeh, 1999, p. 887.

105 Forester, 1989.

106 See particularly, Mandelbaum *et al*, Part IV (1996).

107 Wyatt, 1989.

108 Meyerson in Faludi, 1973.

109 The classic here is Lindblom's *The Science of Muddling Through*, 1959 (also in Faludi, 1973).

110 Habermas, 1984, 1987; Innes, 1998.

111 In Faludi 1973.

reflective practice

always remain more of an art than a science, as Don Schön clearly emphasized in his *Reflective Practitioner*[112].

It is precisely because of these multifaceted socio-political, psychological and behavioral issues that are so closely connected to the planning profession that I have become such a fervent proponent of the development of a solid, cohesive, coherent and unified knowledgebase of the urban makeup. The fact that I remain convinced that a lot of planning is an art, based on unquantifiable, instinctual, emotional and interpersonal gut feelings should not be construed as being in contradiction to my clamoring for a systematic and continuing accumulation of city knowledge.

I do not advocate total rational planning at all. I consider it utterly impossible and undesirable. What I do think, on the other hand, is that planning is so difficult – hindered as it is by the quagmires of personalities, power-struggles, ambitions and fears – that it would be highly beneficial to be able to rely on a solid foundation of factual knowledge as an anchored platform from which to deal with the capricious rollercoaster of public hearings, municipal commissions and citizen activists. In my view, a comprehensive and well-maintained urban knowledge infrastructure is not only a useful concept, but could really be a valuable reality. It could be, in my mind, entirely feasible as long as its development was approached methodically and modularly, without grandiose schemes and high-tech flights of fancy. My own experience is a living testament to such an assertion.

from information to knowledge

The aforementioned "artistry" with which a good planner has to be equipped has a lot to do with how he or she deals with institutions and organizations, both in the public and in the private domain. Information is useless if it is not accepted, appreciated or welcomed by the powers-that-be. Often, information can be threatening to bureaucrats or politicians and it is not unusual to encounter attempts at "burying" potentially damaging knowledge. Delay tactics are often employed by city officials to thwart the adoption of new data that may offset some real or perceived political balances. Patience is a virtue that planners must develop in order to weather the ups and downs of political whim. I think I have experienced the whole gamut of institutional tactics that are deployed from the arsenal of entrenched "apparatchiks". I was privileged to be exempt from the frustrations of city planners working inside government agencies, whose proposals are often mothballed after months of careful analysis and painstaking data collection. Being a freelancer, with academic credentials, and armed with altruistic motives thanks to the job-security of my American paycheck, I could outwait the long stalemates and eventually I was able to re-propose old solutions when the moment was more propitious.

inspiration and discovery

All the while, I often found myself acting as a go-between, connecting individuals and institutions which probably would not have been interacting so freely or willingly under normal circumstances. I did all this based on instinct and on my growing knowledge of city issues, as well as– quite frankly – thanks to a large dose of naïveté. My motivations were those of a

112 Schön, 1983.

self-serving altruist. I was offering my "free" advice to benefit my own lot as a citizen of Venice. My perspective was informed by personal experiences in daily living in my city. My concerns, I am sure, were shared by many fellow citizens. I was luckier than most Venetians in that I had at my disposal the tools and resources to do something about my concerns. I employed my armies of undergraduate students to comb the city, studying whatever seemed appropriate and timely. The big discovery for me was that, despite the fact that our data were collected by young troops, – rather short in experience but long in motivation and stamina – what we assembled was an unprecedented collection of city knowledge, unparalleled both in scope and depth and, most importantly, organized systematically for easy access, unlike the piecemeal assemblages that lay scattered in file cabinets and drawers throughout the various offices and agencies that were staffed by "real" professionals.

CITY DATA

City data, in my restricted definition for this thesis, include everything that is of interest to municipal maintenance, management or planning. In a broader sense, however, city data comprises everything within a city that is of interest to anybody for any reason whatsoever. Of course, there will be low-hanging fruits wherein the benefits of acquiring the extra data and organizing them properly according to City Knowledge principles would far-outweigh the additional transaction costs.

low-hanging fruit

Ideally, the ongoing consensus-building efforts will arrive at a finite number of fundamental datasets that can constitute the core data for a developing local spatial data infrastructure framework[113].

local spatial data infrastructure

"PLAN-DEMANDED" DATA

Indeed, the planning process is predicated on the availability of a myriad of data, but information is almost never available as a consequence of a systematic data-collection strategy by government agencies. Rather, "[t]o develop new land-use plans and proposals (or to form opinions as new opportunities and proposals surface), all of these agencies typically spend considerable energy researching and analyzing land use and ownership in the neighborhoods surrounding the sites that are targeted in the plans."[114] Urban Planning is largely based on ad-hoc collections of data, gathered on an "as needed" basis in what I term a "plan-demanded" mode of operation. Every time a plan is envisioned or proposed, "we need to integrate, and reinterpret many data sources now dispersed among agencies and groups that are administratively isolated and focused on different issues and goals"[115].

Automation plays a certain role in this process, in that some planning data are collected fairly rigorously by some government agencies, but the tendency toward automation in this field has been limited, for the most part, to areas that are under strict regulatory control (like land use) or that generate municipal revenue (like parcel ownership). Record keeping in such

113 See the FDGC efforts in this sense at at http://www.fgdc.gov/framework/. Last accessed 8/20/04.

114 Ferreira, 2000, second page of chapter 7.

115 *Idem.*

instances has always been necessary to the proper functioning of civil society, so the introduction of Information Technologies (IT) has been merely a convenient way to make the process faster and smoother.

Generally speaking, though, the representation of space in many municipal computerization efforts has been shortchanged. At best, locations are represented by address, with all of the standardization and referencing problems that such an approach entails. A systematic approach to the acquisition of fine-grained city knowledge is still considered too cumbersome, even after the introduction of the first G.I.S. tools in the late 80's. Unfortunately, without a reliable, shared knowledgebase of urban information, the "speed-up" effect brought about by traditional automation "may not make much of a dent in the considerable amount of time that our prototypical neighborhood planner must spend studying land use and ownership"[116].

plan-demanded examples

Examples of "plan demanded" data gathering abound. In fact, most data gathering outside of the realms of regulatory or revenue-generating operations probably fits in this category. In my own personal experience, I have actually participated in several plan-demanded campaigns of data collection, both in Italy and in the U.S. For instance, I have led teams of students from the Worcester Polytechnic Institute (WPI) in campaigns to collect information about numerous aspects related to the canals of Venice. The data are currently used by Insula S.p.A. (a private-public company in charge of the maintenance of the Venetian canals) to actually conduct dredging and restoration projects on the Venetian waterways. The data we collected included: measurements of the physical dimensions of the canals, including the water depth and sediment levels at the bottom; a catalog of all sewer outlets and wall damage along canal banks; measures of the water currents in the canals; counts of the boat traffic in the canals and quantification of the wakes produced by passing motorboats; an inventory of all bridges spanning the canals and an assessment of their state of (dis-) repair; a similar census of all boat docks and their condition; a series of campaigns to quantify the amount of cargo delivered to each island in the city; and several other more specific studies.

Similarly, I coordinated programs for the systematic inventory of: trees in the city of Venice and in the City of Cambridge, Massachusetts; parking facilities in downtown Boston, Cambridge and Newton; underground storage tanks for the Boston Fire Department; brownfields for the Boston Redevelopment Authority; historic monuments and landmarks for the Boston Landmarks Commission and for a variety of Venetian organizations; archeological sites both in Boston and Venice; etc. Overall, I have participated in over 100 such projects on both sides of the Atlantic.

All of these projects were completed to fill informational lacunae and were used by the sponsoring agencies to carry out specific actions related to urban maintenance, management and planning that required immediate attention. All of them are examples of "plan-demanded" data gathering.

---

116 *Ibid.*, p. 4 of chapter 7.

City Information

communicative planning

The use of information in planning is discussed, among others[117], by Judith Innes (De Neufville) who promotes a new paradigmatic approach of interactive communicative action, based in part on Habermas' fundamental theories[118], which inspired many other authors, including Forester[119] and Faludi[120]. Many of these authors belong to a camp that is in striking contrast to the "systematic thinking" approach of previous academics[121]. The claim is that communicative and interactive action will help close the gap between theory and practice that has characterized the discipline until recently (Innes 1985). More importantly, this literature constitutes a scholarly attempt at understanding "how, and under what conditions, this information makes a difference" and begins to provide "better conceptual tools to see more clearly how public action is shaped by information"[122] above and beyond the usual mode of planning as *analysis for decision-makers*. This new communicative approach, which is gaining more and more currency in the planning world, is very much in tune with my personal experience in the field.

Judith Innes' own 1992 paper on Geographic Information Systems (GIS), albeit a little dated, is one example of the bridge between the more theoretical assessment of the role of information and the more practical implementation of these theories using modern instruments and current technologies. Although some authors were already looking *Beyond Geographical Information Systems* even before the 1990's[123] , it seems that GIS technologies, together with Database Management Systems (DBMS), are really the primary tools at the disposal of planners. A lot remains to be done to make the use of these technologies more widespread and the information they encapsulate more broadly shared among all the actors[124]. My personal practice has relied very heavily on the application of GIS and DBMS to the systematic accumulation of City Knowledge ever since 1988.

The main trend these days, in technical circles, is the push towards "enterprise GIS"[125] and on the use of web-based internet tools to provide shared platforms for the exchange of spatial urban information[126], pioneered, among others, by our own Planning Support Systems group at MIT[127]. While I appreciate these state-of-the-art concerns, which will indeed constitute the next hurdle in the development of comprehensive Urban Information Systems, I find that this sophisticated know-how is only useful, in the "real world", as a teleological guide for active practitioners like me.

---

117 For a recent text on the topic of information for planners, see Dandekar, 2003.

118 Habermas, 1984, 1987; Innes, 1998.

119 Forester, 1989, 1993.

120 Faludi, 1986. Although Faludi concentrates on Carl Popper's theories more than Habermas'.

121 Innes 1995.

122 Innes 1996, p. 4.

123 Harris, 1999.

124 Geertman and Stillwell, 2004; Brail and Klosterman, 2001.

125 Keating *et al.*, 2003. Azad, 1998; Singh, 2004.

126 Mancuso, 2003; Kelly and Tuxen, 2003; see also page 103.

127 Schiffer, 1992; Evans, 1997; Singh, 2004.

Being cognizant of the leading-edge trends can help us design the building blocks of much less ambitious systems in such a way as to ensure their future upgradability to these high-tech new systems that, so far, have found very little application in day-to-day municipal operations[128]. In other words, I find it much more useful, at this point in time, to be actively engaged in the promulgation of the "basic" concepts of urban information archival in the many municipalities that have yet little or no direct experience with these computer tools.

Oddly enough, my own experience has provided some disconcerting surprises as to the low level of adoption of GIS and even Databases in the actual workings of many departments, even in such "progressive" cities as Boston and Cambridge[129]. Just like in the society at large, there are definitely major discrepancies between the "haves" and the "have nots" in terms of computerization of operations in different cities and even among different departments within the same municipal governments.

"Plan-Ready" Information

While involved in the aforementioned "plan-demanded" projects, I immediately realized that such Herculean efforts would be much more effective if they not only contributed to the pressing needs of the agencies that commissioned the studies, but also contributed to the long-term creation of a knowledge infrastructure that could be re-used in other contexts. While Insula S.p.A. could use the canal data for its immediate necessities, certainly such permanent and immutable features as the canal lengths and widths could come in handy for some other purpose at a later date. More importantly, the canal coding scheme that we developed, which assigned unique identifiers to each segment of the water network, could certainly be useful for posterity. If all future data-gathering utilized the same scheme, it would be possible to compare and correlate datasets referring to the same canal segment at any time. The benefits of such a standard referencing system and of any permanent records connected to it were indisputable. So, when we collected our data on all of the various elements of the urban realm, we always did so with an eye to this fundamental infrastructure of knowledge that could be reusable by other researchers or government agencies for years to come. We were able to do so, partly because of intuitive insights due to our background in the rigors of engineering, but also because we based our referential system on spatial features, which were part of the unchanging (or slowly changing) urban world. We were fortunate to have been pioneers in the use of Geographical Information Systems (GIS) as early as 1987, which opened up the possibility of actually being able to georeference our data to their real-world locations.

standard referencing system

What is still often lacking in today's municipal agencies, in fact, is a decentralized "informating"[130] strategy that properly accounts for the spatial

128 Geertman and Stillwell, 2004.

129 See Part III.

130 Zuboff, S, 1991. In *the Age of the Smart Machine,* Zuboff coins the verb "informate", which means essentially to be able to archive and organize data as they are produced through automated processes, so

spatial georeferencing

dimension of urban features and makes these and other data available to those who need them. To remedy these shortcomings, I am proposing to introduce a space-based representation of the urban realm based on the fundamental, quasi-permanent physical elements that are already the object of regular municipal attention for maintenance or management. While this may not be a novel idea in itself, the innovation I am suggesting would lie primarily in the manner in which these data could be systematically collected, and especially updated, by capturing transaction data and even some low- to no-cost snapshots, starting from a few key areas that are especially relevant to planning. An important aspect of my approach is to focus first and foremost on the permanent and immutable features of the urban world, which, once recorded and organized should require very little upkeep, thus eliminating any redundant effort to collect the same data for a variety of different purposes.

gradual accrual

low hanging branches

low-hanging fruit

The representation I propose can be gradually and systematically "grown" into a reliable, flexible, multi-purpose and shareable knowledge base of the urban landscape, beginning from the "low-hanging" branches of the hierarchy of municipal agencies, which are most directly interacting with the "real world" of the city and would benefit the most from a structured approach to the representation and computerization of the urban features that are already under their jurisdiction. It is at the level of these "low hanging fruits" that the systematic approach I propose can be most effectively overlaid on ordinary municipal operations where the tradeoffs between maintenance necessities and the added requirements of the encoding of city knowledge are most advantageous.

atomic recordkeeping

Whereas traditional recordkeeping methods for these "atomic" elements of the urban realm are generally ill-suited to planning, because their level and method of representation is usually inadequate for higher-order manipulations, the cumulative process discussed herein can not only produce usable information for both the front-line operators of the municipal departments directly in charge of each set of urban elements, but it can also generate solid, fine-grained and rich datasets of usable information that planners and decision-makers can tap into for the formulation of government actions that address more complex urban conditions. In short, the approach that is going to be explored in this dissertation promises to produce "plan-ready" information and may even lead to the inductive development of plans and actions that may be demanded by the preponderance of evidence produced in the process.

planning "automation"

My own approach to the development of "plan-ready" city knowledge is, in a sense, an attempt to bring more "automation" into the planning process, so that the "informating" will be based on reliable, systematically collected, up-to-date and easy-to-update data. This approach espouses Zuboff's argument, though it is applied to fields (city maintenance, management and planning) where informating requirements are already evident and implicit (or "tacit"[131]) knowledge is already used empirically.

that this information can be utilized for higher-order management and control activities that go beyond the original intent of the mere automation of routine operations.

131 Choo, *Information Management for the Intelligent Organization*, p. 11.

The difference between the more traditional manufacturing, and data processing applications studied by Zuboff and the urban disciplines that I am interested in, is that while information about many aspects of urban life is somehow available to city managers and planners – on demand and with substantial effort – there is little or no automation to feed the demand for such information. Whereas the traditional industries in Zuboff's case studies followed the straightforward path of technological development from a manual management and control of operations to a computer-assisted, automated version of the same tasks, many areas of urban management and planning do not have any automation in place at all. Yet, the power of information, which was only gradually realized as an afterthought of automation in Zuboff's companies, is an ever present reality in the urban management and planning arena, where the need for informating actually predates the need for automating. In this dissertation, I try to show that plan-ready information can indeed emerge from a sort of "automation" of the front lines of municipal operations.

ambulance dispatching system

Over the years, I have accumulated a number of personal cases in which data that were collected for one purpose were later used for a completely different reason, without the need to go back in the field. For example, in 1997, we were asked to develop a prototype of an ambulance dispatching system for the Venice general hospital[132]. Ambulances in Venice are of course boats that need to contend with the same obstacles as their land cousins, such as traffic congestions and the like, but also have to deal with the vagaries of tidal fluctuations which may make some routes impassable, either because of high tides – which make some bridges too low to pass under – or because of low tides – which make it impossible to navigate where the sediment build up has made the canals too shallow.

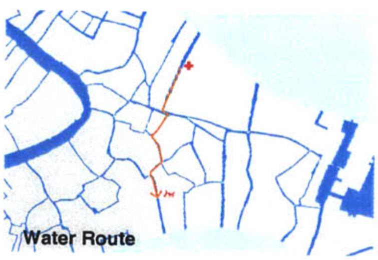

To tackle this problem, we actually borrowed a software package (called TransCAD[133]) from the Planning Support Systems group of the Department of Urban Studies and Planning (DUSP) at MIT. This application figures out the shortest route (by time or by distance) and is capable of taking into account "delays" due to a variety of causes. Thanks to our extensive knowledge of the canal system, we were able to re-utilize our information about the depth of canals and the height of bridges to insert appropriate delays along the routes, depending on the tide levels.

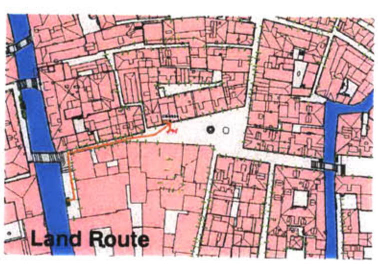

More importantly, we were able to create the canal network graph needed to run the whole system, basing it on the canal centerlines that we had already determined in the course of another project[134]. Moreover, we had to tweak the system to allow for two separate routes, a sea route to a specific dock location, and a land route, from the dock to the emergency address. We already had all of the dock locations from another project and all of the address locations as well. The only thing we had to do from scratch was develop a land network graph and we were already in business...

132 Caporale *et al.*, 1997.

133 Which incidentally is produced by Caliper Corp., a company that is an off-shoot of MIT.

134 We had created the centerlines to "measure" the canal lengths.

the power of plan-ready

This project exemplifies the power of "plan-ready" information, when data are collected with an eye to the potential re-utilization of the information in a context that is different from the one for which the original data collection was conducted. In this case, we actually tapped into at least four previous projects, which were completed long before this study was even envisioned. The utilization of all these datasets in this particular application was never even remotely imagined at the time when the data were collected. What made all of this possible was the fact that each dataset was connected to spatial features of the canal network, through standardized reference identifiers of the canal segment codes in all of the records.

City Knowledge

from knowledge to decisions

The lesson is clear. If you gradually and methodically collect urban data and systematically organize it in computer databases and GIS maps, the data can turn into information and eventually into knowledge of concrete utility to a variety of city agencies… This may not seem like an earth-shattering discovery, but – alas – it is not a widespread practice anywhere in the world, at the fine-grain that I am envisioning[135]. While I am not sure that my personal situation is typical, I do believe that the moral of the story may be broadly applicable and exportable to other urban contexts. This dissertation addresses these issues in detail.

I was very encouraged by my discovery that many of my seat-of-the-pants approaches to the institutionalization of urban knowledge were pretty much "standard" practices in the field, although there is certainly a lot of latitude in this arena, where nothing is ever quite black and white and a lot of variations on themes are employed in the management toolkit[136].

"Plan-Demanding" Knowledge

from knowledge to action

Although the mere availability of "plan-ready" information would already be a major step forward in the management of urban affairs, there is an even more intriguing byproduct of the approach that I am propounding. It seems plausible that, once enough plan-ready information is available to a variety of municipal agencies, the data may begin to "suggest" the need for plans that would otherwise go unnoticed. Patterns may emerge from the data repositories that require attention and this in turn may lead to the spontaneous emergence of the need for solutions in the form of actions or plans.

This is not a far-fetched concept, since it is really how many city plans are born anyhow. Except when real-estate development pressures force a certain "reaction" on the part of city planners, many plans are "proactive" and reflect the overall "direction" in which the city would like to go. If negative trends are somehow noticed, or if citizen groups voice concerns or demands for a certain correction in the way the city is moving,

135 Cf. Geertman and Stillwell, 2004. This is not to say that there are no examples of comprehensive efforts in the world like the national Census efforts, or the national geological surveys and many others. The emphasis here is on the word "widespread" and on the scale – and consequent "grain" – of such efforts.

136 I am indebted to Prof. Wanda Orlikowski of MIT for her guidance in the enormous body of literature on organizations. A great summary of organizational theories can be found in Vibert, 2004.

plans are usually generated to counteract these negative tendencies. If a particular pattern of development is perceived as positive in one part of a city (or in a nearby community), plans are created to try to emulate these success stories.

Having a serious critical mass of city knowledge at one's disposal, would probably invite exploratory analyses that would most likely generate ideas for corrective plans (in the case of negative patterns) possibly based on positive trends noticed elsewhere thanks to the same body of knowledge. Sometimes, planners may suddenly realize that they had actually misdiagnosed a certain urban condition or they may find the real cause of a problem whereas they had been attacking only the symptoms for years[137].

plan-demanded: cargo deliveries data

I have at least one example from my personal experience of a real case of "plan-demanding" knowledge, again from a real experience in Venice, Italy. It all started in 1995, when, under the auspices of UNESCO, we first began to study the quantities of deliveries made to each of the 125 islands that compose the city of Venice. This plan-demanded study was meant to quantify the amount of cargo that was unloaded onto each island from each of its perimeter docks[138]. The purpose was to determine how much disruption would be caused to the cargo delivery system if and when a particular canal around an island was closed for maintenance, thus eliminating a number of docks normally used for cargo operations. The study purported to pinpoint critical areas that needed to be taken into account when the canal closures were planned, to avoid isolating an island completely, thus forcing delivery personnel to surmount a number of bridges with loaded carts in order to make a delivery. The best sequencing of canal closures was thus arrived at, taking into consideration not only the effects on deliveries to local businesses, but also the amount of traffic that would be diverted to nearby canals when a specific canal was shut down for maintenance, which would in turn impede the flow of traffic and hence make deliveries on other perimeter canals more difficult. This was a successful project that led to some follow up studies in other areas of the city, commissioned by Insula S.p.A., which is the company in charge of making the canal closing decisions. Thus, plan-demanded data were turned into plan-ready information.

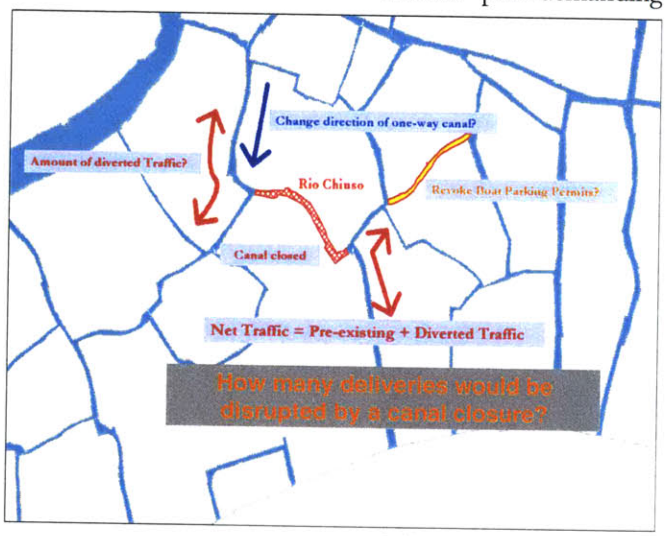

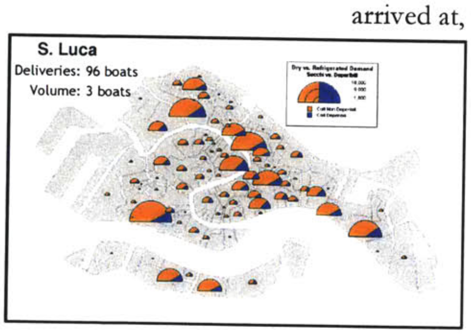

137 See the discussion about the origin of canal wall damage at page 116.

138 Doherty *et al.*, 1995.

plan-ready: canal closure analysis

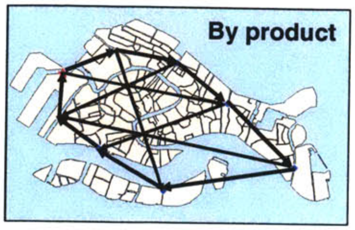

The plan-ready information that was generated by these projects led to the realization that there were in fact many cargo boats that were affected by these inevitable maintenance works. For instance, a specific island near the Rialto bridge was visited by 96 cargo boats every day. What was even more interesting though was the realization that these 100 boats carried cargo that, by volume, would have fit easily in only 3 fully-loaded boats… The sequencing of canal closures was therefore made much more difficult than it had to be by some absurd inefficiency in the actual cargo delivery system. The problem, we quickly discovered, is that cargo in Venice is not delivered "by destination" but "by product". The "water boat" drops off cases of water in each of the islands, as do the "wine boat" and the "beer boat", and the "toilet paper boat" and so on. Each boat only drops off a few boxes or cases, and each boat visits many many islands every day. No wonder there was such a glut of boats around these islands!

plan-demanding: cargo re-engineering

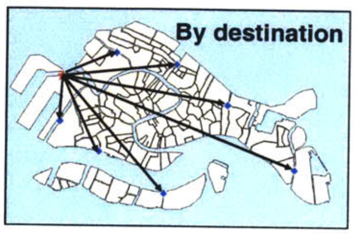

The discovery of the exact extent of the wastefulness at play in what was already known to be a sub-optimal system, transformed the plan-ready information into plan-demanding knowledge when we proposed an in-depth study to the local boat "teamster" union (*Consorzio Trasportatori Veneziani Riuniti – CTVR*) to explore ways to improve the system and eliminate these gross inefficiencies. The study[139], conducted in the summer of 2001, resulted in a proposal for a central warehouse where cargo would be sorted by destination and delivered by only a few fully-loaded boats to only a handful of contiguous islands, in one of the 16 zones in which the city was divided. The project was awarded the WPI President's Prize as the best project of the year 2001. Since then, the City of Venice has embraced the plan and intends to implement the proposal by 2005. Once operational, the re-engineered system of deliveries will reduce overall cargo boat traffic and the consequent wake damage in Venice by over 90%[140].

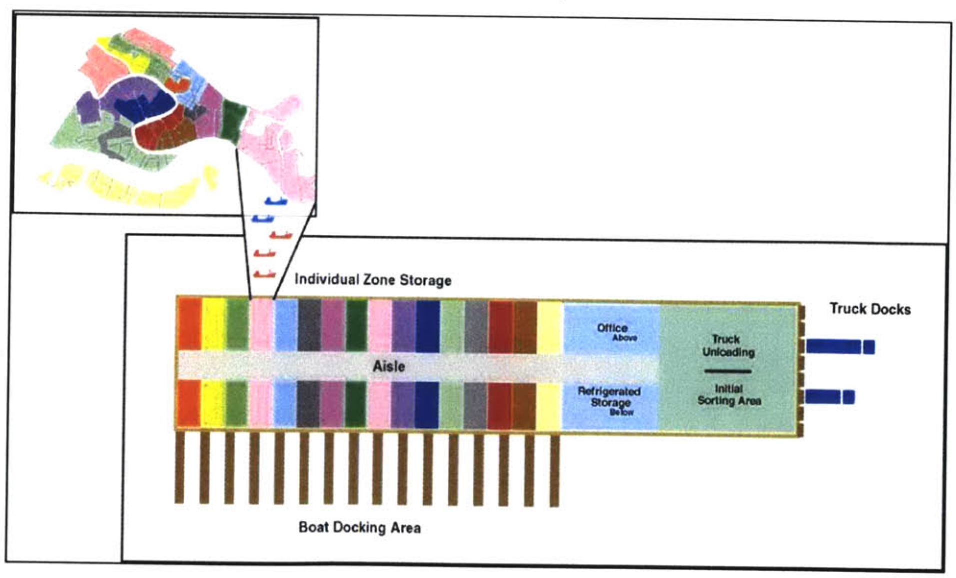

139 Duffy *et al.*, 2001.

140 More on this topic starting on page 117.

(page left intentionally blank)

"Data on the appearance of the environment must be gathered in order to prepare designs and take action. Professional uncertainty as to what is relevant at the city scale, and how it may be best organized for analysis and manipulation, has made it difficult to include visual considerations in city designs. The nature of the data, and the language in which it is recorded, always have a profound impact on the nature of proposals.

Kevin Lynch, 1968
"City Design and City Appearance"
in *City Sense and City Design*, p. 475

PART II:

THE PATH TO CITY KNOWLEDGE

TRAILBLAZING

I had only graduated three years earlier with a Bachelor Degree in Electrical Engineering[141] when I hatched the idea of creating the *Venice Project Center* in 1987. Following a stint as a "coop" student, I was appreciated enough by my company – BTU/Bruce Engineering of N. Billerica, MA – to be retained on a three-year H-1 visa (I was – and still am – a citizen of Italy). Taking advantage of my fringe-benefits, I applied to the Masters program in Computer Science at WPI and enrolled in two CS grad courses per semester, attending evening classes paid for by my company. By the time my visa expired, I had almost enough credits to graduate, so I slipped into a Research Assistant job at WPI with the Intelligent Machines Laboratory[142]. I had just started my Masters thesis when I received the news that my dad had stomach cancer and had already been operated on before I even knew of the illness.

THE VENICE PROJECT CENTER

There was no point in hurrying back, but I decided that I needed to move back to Venice as soon as possible to be near my family in this time of crisis. Accordingly, I began to make arrangements to complete my Masters thesis in Italy and simultaneously invented the Venice Project Center (VPC) as the "bridge" that would keep me connected to WPI while living in Italy.

creating a bridge

What amazes me about the creation of the VPC is not how I possibly conceived that the WPI administration would let me – a mere graduate student – start a center where dozens (eventually hundreds) of WPI students would spend a term conducting a serious interdisciplinary project to complete a major degree requirement[143]. Nothing works better that the naïve drive of an inspired 27-year-old for such audacious gamble. On the contrary, what truly amazes me is that a mature, distinguished professor of English and Associate Provost of WPI was actually convinced by the passionate grad student to embark in a transoceanic adventure of such magnitude. Granted, the plan seemed innocuous when I proposed to simply conduct a preliminary "bootstrap" project to determine the feasibility of creating the center in the first place.

bootstrapping

The fateful "feasibility study" took place from mid-October to Christmas of 1988. The center proved "feasible" indeed as witnessed by the over 500 students who have spent a term there since then. The rest – as they say – is history.

141 My senior project had to do with "machine vision", which later got me into Robotics and Artificial Intelligence.

142 Where we worked with autonomous vehicles controlled by an "activation framework" of intelligent agents communicating with each other and acting locally using heuristics and stored knowledge. These were prototypical "emergent systems" except the name had not been coined yet.

143 WPI operates about 20 centers around the world, from Melbourne to Copenhagen, from San José de Costarica to Bangkok, where every year over 400 students complete some part of their college curriculum for a Bachelors Degree in Science, Technology or Engineering.

the early years

Meanwhile, I had miraculously[144] got a job at the University of Venice, to run the fledgling Laboratory for Humanistic Informatics. After the pioneers in the winter of 1988, I had a couple of WPI teams in Venice, in the summer of 1989, but the first big year was 1990. Students came in the Winter, Spring and Summer terms, a total of 11 teams and almost 40 students. That year, a project on lagoon archeology received the Best-Project-of-the-Year Award (known at WPI as the "president's IQP Award"). It was the first of many such awards, but served as early recognition of the quality of the center, where I conceived and advised all projects, together with a faculty colleague as co-advisor.

286 laptops

hard drives!

Due to my computer engineering background, all projects relied strongly on the use of personal computers. When we first started, we were operating with top-of-the-line 286 laptops, with two floppy drives (and no hard drive as I recall). Technology was progressing fast and we quickly moved up to 386's with 10, then 20 Megabyte hard drives. Wow! We felt like we were at the forefront of high-tech (and in fact we were), although the machines we used had fewer resources than even the simplest Personal Digital Assistant (PDA) of today.

GIS

mapinfo for DOS

unavailable maps

Geographic Information Systems (GIS) were barely making their debut in the late eighties, yet somehow we jumped into the GIS bandwagon from day one. This was a serendipitous twist of fate that augured well for the future of the VPC. I remember vividly taking a long drive to Troy, New York, some time in 1987, to go visit a company called Mapinfo, which at the time consisted of two people in a small office in a warehouse-style building. We bought one of the first licenses of Mapinfo that year and were very happy with it, even though it was operating under DOS and there were no maps to use with the software in Venice[145].

making our own

In order to make use of this powerful tool, we spent a great deal of time, for the first couple of years of VPC operation, digitizing Venice maps from a printed book called *Atlante di Venezia*[146], which actually showcased early GIS maps and orthophotos that were not – unfortunately – available to the public in electronic form. We threw away our homemade maps a few years later when the first electronic maps of Venice became available to us through our contacts with the City of Venice and the Consorzio Venezia Nuova.

DataBases

Dbase III

Dbase IV

Similarly, we also pioneered the use of real databases, as opposed to the spreadsheets that were very much *en vogue* at the time (remember Lotus 123?). The best that DOS had to offer in the mid-eighties was Ashton-Tate's Dbase III and III+ which appeared around the time of my return to Venice. Around 1990, Dbase IV would make its appearance on our 386

144 It's a long story that involves calling Venice historian Frederick Lane out-of-the-blue, after realizing he lived in an adjoining town, which was listed after his name at the end of the foreword to his book *Venice A Maritime Republic*. Another naïve, spontaneous action that led to unbelievable consequences.

145 As it turned out, the Italian projection system was not supported either, but we were so unaware of the finer nuances of GIS that we didn't even notice this drawback…

146 Published by Marsilio Editore in 1990.

the database "habit"

machines. Despite the numerous limitations of these early relational databases, we began to put away the results of our first studies in electronic formats that proved to be rather long-lasting in retrospect. More importantly, we acquired a "habit" of structuring our data in ways that made it amenable to analytical computation and to cartographic representation through GIS.

PRELUDE TO CITY KNOWLEDGE

the grand scheme

I recall quite clearly having dinner with the parents of one of my first six students in the early winter of 1988. The father of the student surprised me a bit by asking me: "What grand scheme do you have in mind for these student projects"? He thought he could detect some sort of an ulterior motive in the fanatical manner in which we treated data even during that very first project in Venice. I was taken aback by the question since I had never publicly admitted to anyone that I indeed had a grand scheme in mind.

City Knowledge by any other name

I guess I was not very good at keeping a secret since that man read right through me and saw that I intended to gradually build up a storehouse of knowledge about my hometown. Although we all appreciated the ambitious nature of an endeavor that aimed at gathering information about the "totality" of Venice's reality, none of us had any idea about how difficult that would be, as we sat around that dinner table in 1988. Well, it took almost twenty years and a lot of effort to produce this dissertation, but the scheme I hatched in 1988 now has a name. It's called City Knowledge. And this document is an attempt to define it as I describe just how it evolved from 1988 until today.

THE REST OF PART II

the Venice Inner Canals project

The rest of Part II is composed of two chapters. The chapter on the *Venice Inner Canals Project* contains a sample of the numerous aspects of the canals that we studied for over a decade, most of them in the framework of a United Nations Educational, Scientific and Cultural Organization (UNESCO) project – the *Venice Inner Canals* project – sponsored by the Italian Ministry of University and Research in Science and Technology (MURST). The chapter does not cover all of the aspects of canals we studied in the 1990's[147], but introduces only the those that exemplify some of the lessons that were most valuable in the formulation of my City Knowledge approach, which is detailed in Part V of this dissertation.

Protecting Venice from the Tides

The last chapter in this Part[148] exemplifies many of the tenets of City Knowledge. It shows how it is possible to tap into "plan-ready" information produced by a variety of past projects in order to postulate a method for the quantification of the damage that high tides can inflict on the material fabric of the city of Venice. Without the incremental accumulation of urban information that I adopted since 1988, the difficult task of ascertaining the cost of *acqua alta* would be neigh impossible. More importantly, without some of the acquired knowledge to guide such an

[147] For complete details see http://www.unesco.ve.it/. See also note 151 and Carrera, 1996, 1999b, 1999c, 1999d.

[148] Adapted from a paper currently being published (Carrera, 2004).

effort, we may in fact arrive at erroneous conclusions and miscalculate the consequences of human and natural acts that affect the behavior of the waters of the Lagoon and of the inner canals, which in turn affect the stones of Venice.

All of the rich examples in the rest of this Part II will help to gradually build a case for City Knowledge, which will be further enriched in Part III, dissected in Part IV, and finally summed up in Part V.

THE "VENICE INNER CANALS" PROJECT

a fateful suggestion

The canals of Venice have been a defining feature of the city since Charlemagne's son Pippin forced the Doge to relocate to the islands of the *Rialto* in AD 815. When a Venetian scientist[149] suggested that we study the hydrodynamics of the Venetian waterways in 1989, I thought that it would be an interesting project for my students which may provide some interesting information to contrast against past studies on the same subject. A comparative study of the hydrodynamics of the inner canals could, for example, provide support for the hypothesis that the creation of the artificial *Canale dei Petroli* in the central lagoon had dramatically altered the flow of water in the city. Such a study could also demonstrate how some canals became more stagnant after some adjoining canals had been filled-in and turned into *rii terà*. When the first team of students began studying the Venice inner canals in the *Dorsoduro* section of town in the winter of 1990,[150] we were looking forward to these before-and-after comparisons. Generally, canals exhibit a bi-directional current, flowing in one direction when the tide is coming in from the sea, and in the opposite direction once the tide retreats back out to sea. Instead, our study showed that the four canals we monitored always flowed in the same direction, regardless of whether the overall tide was rising or falling, which was a rather unexpected result. The most surprising discovery, however, was the fact that these thousand-year-old canals had never been scientifically studied before. In their millenary history, nobody had ever bothered to do what these four twenty-year-olds from America had done in just a few weeks. This was an amazing discovery that led to a decade of in-depth studies of the entire canal system, making the Venice Project Center the foremost repository of knowledge about the inner canals in the history of Venice. What had started as a comparative study of water currents, soon turned into a thorough multi-year baseline survey of many different aspects of the target canals.

the first scientific study of the canals

UNESCO
VENICE OFFICE

UNESCO

MURST

Our solo beginnings caught the attention of a local official of the United Nations Educational Scientific and Cultural Organization (UNESCO) who enlisted me to write a major funding proposal to the Italian Ministry of University and Research in Science and Technology (MURST). The resulting project, based on our initial studies, was officially named the *Venice Inner Canals* project and it attracted over one million dollars in funding. From 1992 to 2002, the Venice Project Center played a major role in carrying out the planned research, which can be perused on the web at www.unesco.ve.it.[151]

149 The man's name is Gianpietro Zucchetta, a Venetian chemist who at the time was working with the National Council of Research in Venice, and who had just published a book entitled *I Rii di Venezia.*

150 Ciacciarelli *et al.*, 1990.

151 The synthetic reports in English and the web site design were part of the final contract that the author had with UNESCO. The web site may have been relocated since this writing, but my personal web site should redirect an interest reader to the new site (see www.wpi.edu/~carrera).

Insula S.p.A.

the Venice Inner Canals project

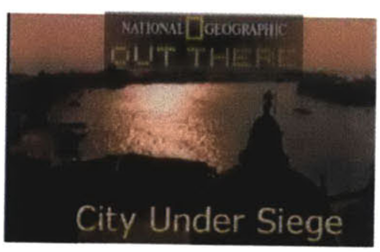

In 1997, the City of Venice, spurred in part by the UNESCO-MURST project, created a public-private company, named Insula S.p.A. to carry out urban maintenance activities, primarily on the inner canals. In 1998, Insula and UNESCO entered into a cooperation agreement whereby all of the data collected by the UNESCO teams (including the WPI Venice Project Center teams) were transferred to the fledgling company to get the initial maintenance activities off the ground as rapidly as possible. The cooperation gave Insula the jumpstart it needed to quickly begin canal dredging and repairs and also resulted in the publication of a joint UNESCO-Insula book entitled *Venezia la Città dei Rii* (Venice the city of canals)[152] which I co-edited.

In this chapter, I retell the story of the Venice Inner Canals project as a way of introducing many of the essential elements of City Knowledge that were first identified in the course of these studies and later became permanent features of our *modus operandi* in other areas of research. In particular, this chapter will describe our early efforts at a systematic standardization of the spatial reference frameworks along "reasonable" atomic elements (the so-called "canal segments"). Even though most of our studies were "plan demanded" in the early years, the "scaffolding" that we created quickly became the source of "plan ready" knowledge that later was re-used for a variety of different purposes.

Our work received lots of recognition from the media. In addition to the aforementioned book, we were also featured in numerous magazines, such as National Geographic[153], the Smithsonian[154], Wired[155] and New Scientist[156]. I was also personally featured as the main character of a National Geographic documentary entitled *Venice: City under Siege*[157].

Most importantly, though, our success can really be measured by the impact that our projects have had and continue to have on the city of Venice, where many of our suggestions have been implemented and where much of our data are used daily to make the city function, from the canal dredging that Insula is performing, to the management of boat docks by the city's public services department, from traffic modeling to sewage disposal. City Knowledge is in fact working in Venice, albeit in embryonic form and through isolated cases that still lack overall coordination and integration. Yet, the examples shown in this chapter will demonstrate that the dice has been cast and I think it's only a matter of time before the first full-fledged City Knowledge system is developed in its entirety.

152 Caniato *et al.*, 1999.
153 National Geographic magazine, August 1999.
154 Smithsonian magazine, September 2002.
155 Wired magazine, August 2003.
156 New Scientist, September 2003.
157 National Geographic Video, "City Under Siege", series: *Out There*, aired worldwide 2002-2004.

the moral of the story

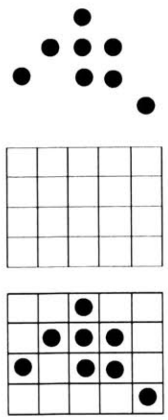

The last few paragraphs of each section, entitled *Lessons Learnt* (or some variation thereof), recap and distill the "lessons" that the section is meant to convey. One of the overall lessons of this particular chapter is that there are immediate benefits to the systematic standardization of a city knowledge framework. These instant advantages ought to be enough to justify the standardization effort by themselves. Breaking up the canal network into small bite size chunks allowed us to reorganize these atoms into any combination we needed when more advanced issues were tackled later in the research effort. Here the moral is that if you just look at the "trees" you may lose track of the big picture, but seeing the urban environment as groups of objects or elements affords great freedom to recombine the atoms into any configuration, as long as there is an overarching scaffold to support the organization of these pieces. Once the framework is in place, the benefits extend far beyond the immediate.

The Canal Network

double duty

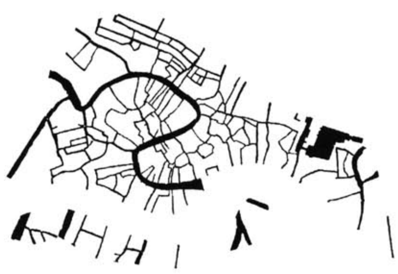

The Venetian canal network performs a double duty for the City of Venice. It is at the same time both a transportation and a sewage disposal network. Boats ply the network to transport all manner of goods and people, as trucks and buses do in "normal" cities. Everything that moves on wheels on the mainland moves on water in Venice. Taxis are boats, as are buses, ambulances, police cruisers and fire engines. The uniqueness of Venice is that the transportation grid never crosses the pedestrian routes at grade. LeCorbusier (among others) considered this an ideal system for urban mobility. Similarly, up until the late 1800's, most people looked at the Venetian canals as the ideal sewage removal conduits. Before the advent of modern sewer systems,[158] large cities were infested by untreated sewage. Human and animal[159] excrements piled up along streets and in back lots until rains mercifully washed off the noxious refuse. Not so in Venice, thanks to its canals. Unfortunately though, to this day (2004) a majority of the residential sewage still ends up in the canals, to be flushed out to sea by the daily tides. What was once a marvel in the eyes of XIX century visitors, has become an eyesore to contemporary travelers and a source of embarrassment and occasional discomfort to today's Venetians.

lack of canal studies

Given the importance of these dual roles, one would think that, over the centuries, Venetian authorities would have committed significant resources to the study of these bodies of water. Surprisingly, such was not the case. For example, the hydrodynamic behavior of the canals, which is of considerable importance in the dispersal of the sewage outflows and can also affect boat traffic[160], was never really quantified until WPI students carried out their first measurements in the 1990's. The earliest record of a qualitative investigation dated back to 1900[161] and the only limited quantitative campaign was carried out in 1966[162]. In all, even counting studies of dubious scientific value and studies that were only vaguely referenced in the literature, as of 1990, only about nine publications existed about the topic of the hydrodynamics of the inner canals and none of these were truly scientific, systematic or representative of typical behavior.

lack of standard referencing

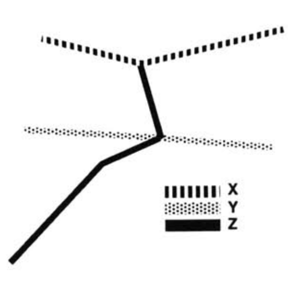

Even when previous data existed, it was nearly impossible to compare them to what we were measuring, since the location of the past measurements was identified just by a canal name, without any specification of where, along its length, the measurement was taken. This issue was non-trivial since traditional canal names often referred to waterways that didn't change name through several intersections, just like many roads in American cities don't change name every time a side street intersects with them. Since water currents are affected by the flows of intersecting canals, not knowing exactly where the measurement was taken vis à vis these intersections meant

158 Ancient Rome had the "cloacae", so sewers are not really a modern invention.

159 Let's not forget that horses and mules were the primary means of cargo transportation, which compounded the problem. Cf. for example, Rodolphe el-Khoury, "Polish and Deodorize: Paving the City in Eighteen Century France".

160 Especially when oars were the only means of propulsion.

161 Paluello,1900.

162 Dorigo, L., 1966.

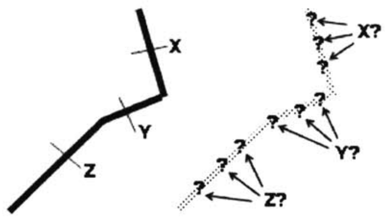

that there was no reliable way to make an apple-to-apple comparison with later measurements.

The problem was that, whereas we knew exactly where we took our measurements, the reference systems of past studies did not allow for the pinpointing of their measurement locations with the same level of accuracy, rendering such comparisons impossible. In order to change this state of affairs and to make possible future comparisons with our data, we devised a system for the unequivocal identification of each tract of canal, through a process of segmentation.

The next section explains how such a system was developed.

Representing the Network

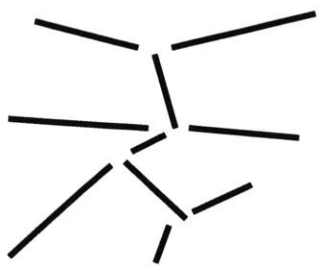

canal segment

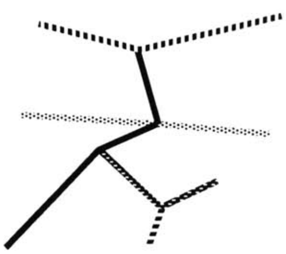

islands

In order to define the indivisible, fundamental components of the canal network that would make hydrodynamic comparisons feasible, the discriminant was the presence of an intersection. The "atom" was defined as a tract of canal between two intersections and was labeled a "canal segment". This simple concept was complicated a tad by the fact that in Venice there exist several filled-in canals (called *rii terà*) that were turned into pedestrian streets. In many cases, these *rii terà* maintain a subterranean conduit that is overlaid with pavement, thus preserving the hydrodynamic function of the former canals. Therefore, for our segmentation, we had to also consider as intersections the points where *rii terà* joined regular canals. In fact, since nobody had ever mapped out the canal network on a Geographical Information System (GIS), we were faced with the task of determining the shapes of the canal segments as we identified their boundary intersections.

Of course, we could have created "stick canals" or "stick segments" by simply representing the water network as one-dimensional centerlines, as is still frequently done in many GIS applications when dealing with roads[163]. Back in 1990, working on Mapinfo® for DOS, lines were certainly more appealing than full two-dimensional shapes or regions to represent the canal segments. However, the linear representation was deemed inadequate early on for a variety of reasons. For example, it didn't allow to tap into the geospatial dimensioning features of GIS, such as the ability to automatically determine surface areas which are needed to calculate water and sediment volumes. Also, the arc representation of canals prevented the measurement of widths, which are crucial for canal navigation. Conversely, shapes do not directly provide canal length information, which lines would provide instantly. What aided our decision was the consideration that, once regions are defined, it would be somewhat easier to derive the respective centerlines than to transform lines into shapes in the opposite direction. But, before we even began the creation of two-dimensional canal regions, we needed to first define the most fundamental units of space that gave shape to the canals, i.e. the islands. Once we had defined the perimeters of all the islands, everything between these islands would belong to the canals layer[164].

In 1990, when we had started this project, G.I.S. maps were still treated as rare commodities and access to the few existing layers was essentially barred to anyone but a few insiders. Thus, we began to make our own maps by tracing on a digitizing tablet the printed outlines of islands that were published in a marvelous publication that had just been issued at that time, called the *Atlante di Venezia*[165]. A few years later, once we obtained electronic maps from the City, we repeated the process and produced a final "official" version of the layers for islands, rii terà, intersections, canals and canal segments, all of which have been adopted in day-to-day use by all

163 For example, the Census Bureau's *Tiger Files* are just simple line arcs.

164 Using a similar approach with regular city blocks may be a good way to define road shapes in "regular" cities, since blocks are probably fairly stable features in urban settings.

165 *Atlante di Venezia.* 1990.

departments of the City of Venice and also by Insula S.p.A. and other governmental and non-governmental organizations[166].

rii terà

The other fundamental layer needed before segmentation could be applied to the whole network was an *intersections* layer, representing the connections between two canal segments and also between a canal segment and a *rio terà*. Thus the second layer that we developed was a *rii terà* layer. This layer was extracted primarily from existing publications on the topic of rii terà[167]. These former canals were classified as vaulted or filled, based on the available information, but both types were equally used to determine valid intersections for the intersection layer.

Once all of the 125 islands and the 50 rii terà had been defined, it was possible to identify and map out all of the intersections that bracketed all of the segments in the city. Intersections could come in two flavors: one type of intersection occupies physical space and has a definite surface area; another type is a symbolic line that represents a potential location for an intersection, but does not physically occupy space. The latter are generally connections between real canals and filled-in *rii terà*, but can also represent junctions between small canals and very large ones like the Grand Canal. "Real" intersections, which occupy physical space, occur at intersections between segments and include tracts of water that do not belong to any of the intersecting segments, which were terminated flush with the end of the two islands that formed the sides of the segment.

intersections

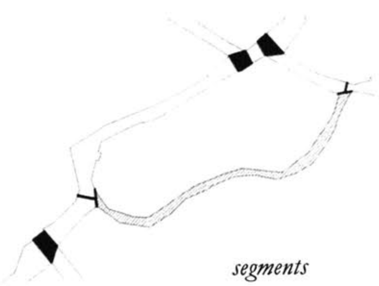

segments

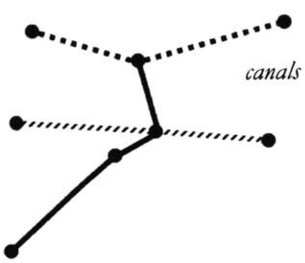

canals

After all 307 intersections were identified and represented, we could define all of the 367 segments in the network. By definition, each segment could only connect to exactly two intersections. Concatenations of two or more segments could then be used to reconstruct as completely as possible the traditional canals that are still the "unit" that is understood and referred to in common parlance by today's Venetians. Even these traditional canals have been nonetheless "formalized" and must begin and end on an intersection, since they are "molecules" composed of the "atoms" of canal segments. In all, we defined 182 neo-traditional canals with official names, and distinct beginnings and ends.

Lesson 1: atomize

The main lesson we learnt from trying to deal with an unstructured framework of traditional canal names was to divide the territory into the smallest units (or atoms) that made sense at that time. The guiding principles in determining when to stop the breakdown were technical and logical. First of all, we would only be able to keep track of elements that were adequately representable and measurable using current technology. Moreover, we would only divide reality into units that "made sense" in terms of what our purposes were for the practical utilization of our territorial datasets[168]. The units of analysis that we foresaw using for urban maintenance, management or planning dictated the units of measurement that we set up.

166 Personally, I feel that the naming and coding of canals and canal segments is one of the most rewarding and lasting consequences of my work.

167 Zucchetta, Gianpietro. 1992. *Un'altra Venezia.*

168 See Evan's (1997) discussion of river reaches for more on this subject.

I guess the generalizable lesson is that one ought to make the units of measurement as small as technologically feasible and as big as practically appropriate in terms of transaction cost vs. analytical benefit.

STANDARD REFERENCE SYSTEM

In order to make all references to these atomic segments uniquely distinguishable, each segment was assigned an alphanumeric code as the primary identifier. Each of the 182 newly created canals was assigned a four-letter code representing the most significant letters of the prevailing name of the traditional canal whose course most closely approximates the course of the new concatenated one. For example, the canal commonly known as *Rio de San Salvador* was assigned the code SALV. Once canal codes were assigned, based on the prevailing traditional name, the latter became the "official" name of the canal and was permanently associated with that canal.

Many Venetian canals however are known by more than one name, so an arbitrary decision had to be made about which name to assign to the canal as the official appellation. To allow multiple names to be connected to a canal, we created an "alias" list in a database table, so that each canal code could be associated to a variable number of names.
Any segment that is part of this newly defined canal would subsequently be assigned a unique identifier according to the syntax XXXXnn, where XXXX is the code of the longer, neo-traditional canal that encompasses that segment, and "nn" is a two-digit consecutive numeric index that sequentially numbers the segments from north to south, starting at 1. So, for instance, if a traditional canal is called *Rio dei Scoacamini*, the code for the neo-traditional canal that most closely approximates the course of that ancient waterway would be assigned a canal code of SCOA, and the two segments that make up the new canal would be labeled, from north to south, SCOA1 and SCOA2. If a new canal was made up of only one segment, then such segment would not have a numeric suffix, thus the canal code and segment code would be identical, as in the case of the aforementioned SALV.

Islands had already been numbered, though not completely, by the City of Venice, so we adopted the existing enumeration, and numerically labeled the remaining islands. In addition, being a fervent admirer of alpha codes, because of their inherently higher information content, I also insisted on a four-letter code for each island. The four letters would represent the most significant letters of the most prominent landmark on that island, be it a church, a palace or even a famous street or square. Thus, each island would have a dual link to outside data, through the pre-existing numeric code and through the new alpha code. For example, the island where Saint Mark's square is located is doubly labeled with the mnemonic character label MARC, as well as with the more compact, but less identifiable number 92.

In spite of my personal bias toward the more explicit character labels, intersections were simply numbered sequentially from northernmost to southernmost, according to the Y coordinate (latitude) of the centroid. The increment was set at 5 units instead of 1 to allow for potential insertions of additional intersections in future years. Thus, the northernmost intersection was numbered 5, the next one to the south 10,

canal codes

canal names and aliases

segment codes

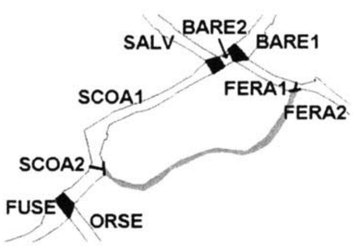

island codes and names

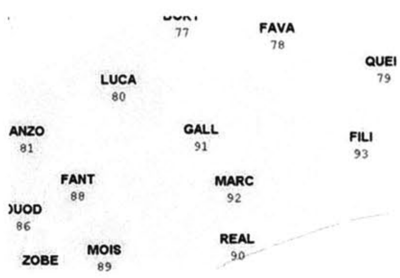

intersection codes

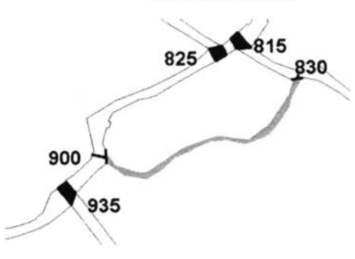

then 15, until the 309th intersection that became number 1,545[169]. This north-to-south ordering implied that the two intersections that bracketed a specific segment would frequently have widely differing intersection numbers, but the northernmost node would always have a lower number than the southernmost one[170].

rii terà codes

Rii terà were not assigned text codes, but were simply numbered sequentially. After their important initial role in the determination of the intersections to use for the segmentation of the network, filled-in canals were not associated with specific datasets until a later date, when I wrote a small article on *rii terà* for the book on the canals that I co-authored in the year 2000[171]. On that occasion, the date of the conversion of a canal to a pedestrian street and the government[172] that commissioned the work were recorded and a preliminary analysis was conducted to determine what public administration was most culpable for these *ante litteram* "urban renewal" activities.

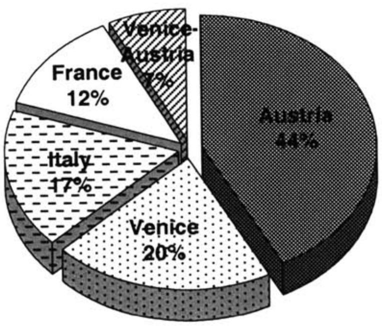

Aside from the standard codes that uniquely identify each element, the fundamental layers discussed so far – namely the islands, rii terà, intersections, segments and canals – each had additional pieces of information that were permanently embedded into the structure of the layer.

legacy identifiers

First of all, individual objects in most layers are frequently associated with a number of additional reference codes and IDs that are part of the heritage that these objects carry with them. Pre-existing IDs should always be included when possible[173], to allow for some retroactive linking to legacy databases that may have been in use up to the present.

embedded permanent parameters

My personal approach is to limit the amount of data embedded in the layer itself and instead keep all related data in separate external databases. Nevertheless, a modicum of information is permanently affixed to each element in each layer, in addition to all of the necessary identifying codes and labels discussed above. The permanent parameters that may be typically embedded within each element of a layer are:

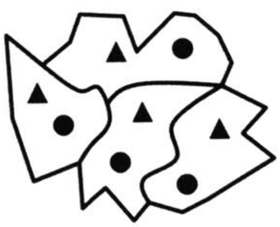

- immutable physical features of the geographical object, such as its surface area or its perimeter[174], and/or
- topological aspects, such as codes of objects that are connected to the element in question, and/or

169 Since the original numbering, two intersections have been deleted, bringing the total from 309 to 307 and creating two 10-unit gaps in the sequence at number 400 (395-405) and 1,515 (1510-1520). Despite the deletions, the north-to-south sequencing remains intact and proves the resilience of the method adopted.

170 This built-in feature proved really useful at a later date, when Insula S.p.A. requested that we include, inside the segment layer, the codes of the north and south nodes attached to each segment.

171 Caniato, Carrera, Giannotti and Pypaert, 2000, *Venezia la Città dei Rii.* Cierre ed., pp. 216-221.

172 Namely the French empire, the Italian Kingdom, the Austrian empire or the Venetian republic itself. Some projects were combined Venice+Austria efforts, since they were started before the fall of the republic and completed by the Austrians.

173 One must always be very cautious with existing IDs to ensure that the old and new objects are exactly the same or sufficiently comparable for the purposes for which they are used.

174 Even though any GIS system could easily extract these features from the geometry, it is usually convenient to have these parameters ready without additional effort, especially if they are to be used frequently.

- categorical classifications, such as typologies (e.g. whether a rio terà is vaulted or not), and/or
- information about the object's creation history (such as what government filled-in the canal, to stay with the above example) or other metadata, where appropriate.

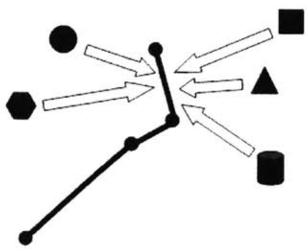

The information attached to each object in a layer should be homogeneous and, to be truly useful, it should be as complete as possible across all objects in the layer.

The unique coding of each element of each of the fundamental layers had to be done manually, but once this was done, any piece of data subsequently collected on any aspect or phenomenon related to any of these objects could be automatically linked to its location in space, thanks to these spatial units of reference. This spatial framework made it possible for all of the data collected in later years by dozens of WPI teams to be forever referenceable and hence re-usable *ad infinitum*, as will be shown in detail in later chapters. Once defined and labeled, a typical segment (like the aforementioned SCOA1), could very easily be linked at any later time to additional information such as: physical dimensions, boat traffic counts, hydrodynamic current velocities, the extent of canal wall damage, the quality of the water, estimates of sewage discharge and many other pieces of city knowledge. This is exactly what we did in the rest of the 1990's.

reusability

The next several sections describe in some detail many of the varied pieces of knowledge that we collected following the creation of the reference system, showing how we were able to instantly link several datasets to the corresponding segments via the aforementioned standard codes.

Lesson 2: codify

Alongside the atomization of the spatial objects that compose the city, it is important to assign unique codes to each of these atoms, in order to make further references possible. Once each object has a clear nametag, many disparate sets of attributes can be connected to it by a variety of different organizations or individuals to suit different needs. This simple concept, which has been around for years in the Relational DataBase Management System (RDBMS) community, is not often applied wholesale to all mappable objects that make up the municipal territory.

Later on, in Part IV, I suggest ways in which these identifiers can be assigned uniquely and reliably by specific public agencies who have jurisdiction over the "birth" and "death" of these physical objects in reality.

Physical Characteristics

The first pieces of city knowledge that were attached to the newly defined and identified canal segments were related to physical aspects of the segments, specifically length, width, surface area, depth and volumes of both water and sediment in each segment. All of the dimensional parameters except for the depth (and consequently the volumes) were obtained directly from the GIS maps. The measurement of canal depth (also known as bathymetry) took several years from 1990 until 1994, since students were only in Venice for a few months every year.

surface area

The easiest dimension to calculate for each canal segment was its surface area, since this measure was automatically implicit in the segment's geographical shape. GIS systems will instantly provide the surface area of an object once it has been spatially defined.

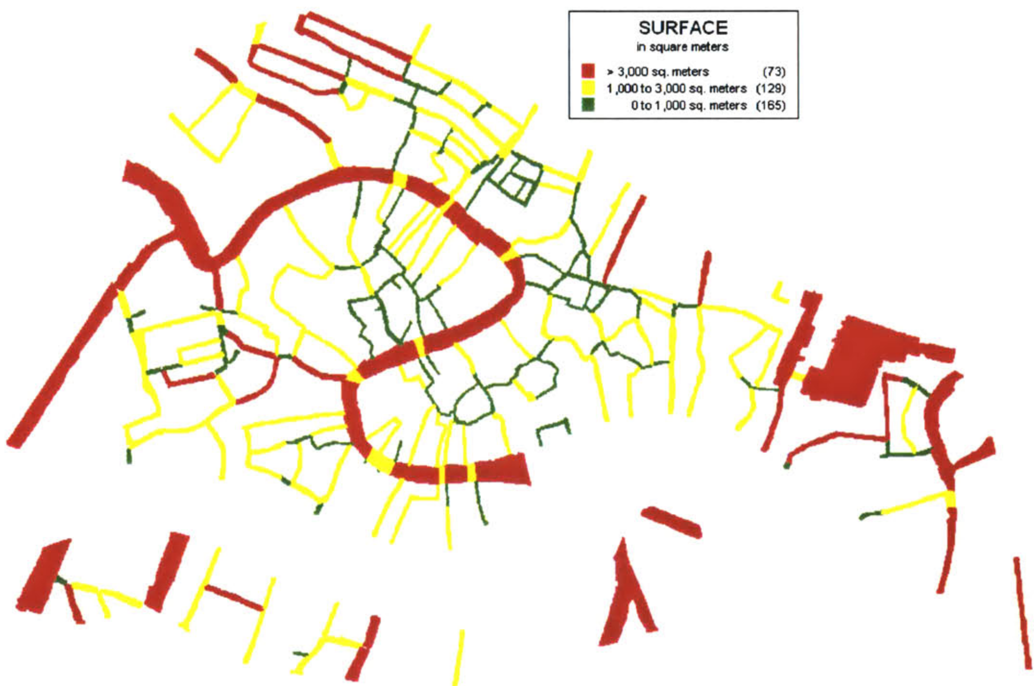

length

Now that segments had a clear beginning and end at an intersection, it was possible, for the first time ever, to truly determine the length of canal segments and of neo-traditional canals in a definitive way. It may seem patently obvious, but one cannot measure the length of a stretch of transportation network – be it a road segment or a canal segment – if there is no clear definition of the extremities of the stretch being measured. Knowing the length of a transportation arc is useful for a myriad of mundane maintenance and management tasks. Contracts and payments are often based, for example, on lengths of road plowed, or miles paved, or – in Venice – on the meters of canal dredged. It is amazing that Venetian authorities had operated for so long without an official listing of canal lengths. Once again, we were the first to produce a fundamental piece of

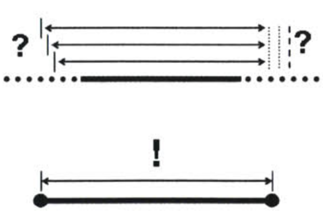

city knowledge – one that will conceivably remain unaltered for ever and will be used for years to come.

centerlines

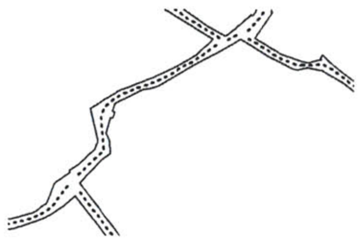

Determining the length of a canal segment entailed first and foremost the definition of *centerlines*. These were constructed manually for all canal segments, by drawing lines along the segment's midpoint from one extremity of the segment to the other, keeping the centerline equidistant from both sides at all times. Intersections received a similar treatment. The centerline of a neo-traditional canal was simply formed by the concatenation of the centerlines of all of its component segments together with the centerlines of all the intervening intersections. Once the centerlines were drawn, segment and canal lengths were computed instantly through standard GIS functions[175].

width

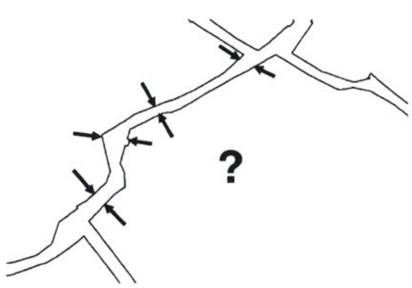

The width of a canal is an elusive measure, due to the irregular shape of the Venetian waterways. Even though typical city roads are more regularly shaped, any single measure for the width of a transportation link is often misleading. One can use an average or resort to other statistics, such as the width at the widest and narrowest points, but it will be rarely possible to define a single width as easily as one defines a single length or single surface area. Measuring the narrowest and widest points is also non-trivial since this determination implies the laborious ranking of several

175 Since canal lengths were expressed in meters, we settled for one decimal digit, representing 10ths of meters as the highest precision that made sense, though meter-precision is probably even more realistic.

measurements. With highly irregular segment shapes, it is even difficult to determine exactly what point on the opposite side to measure to.

average width

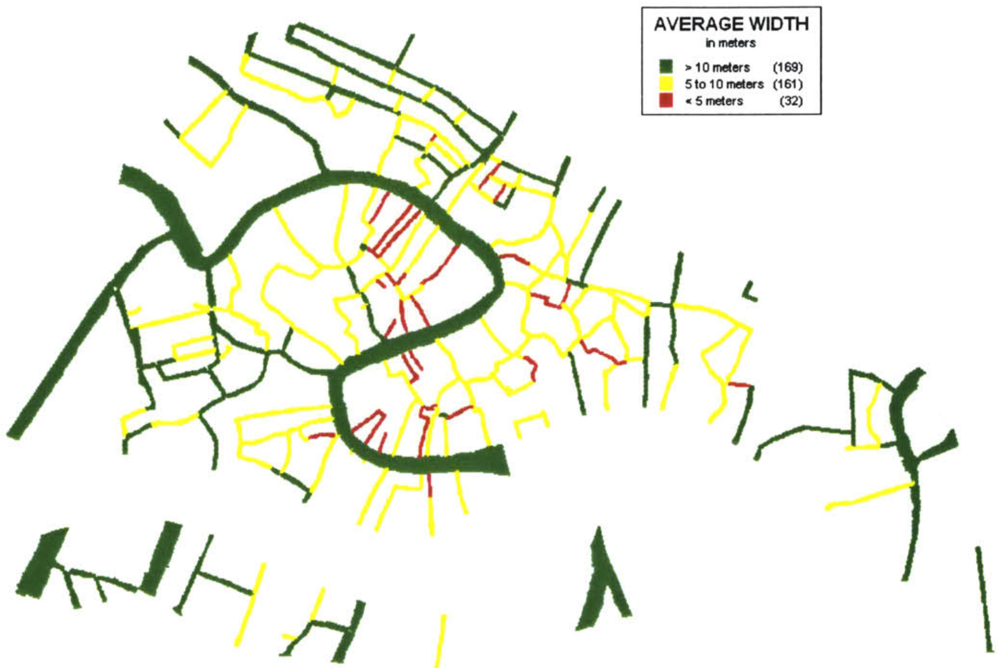

Fortunately, the average width of a segment can be computed by simply dividing the surface area by the segment length. Thus, we were able to determine this dimension for <u>all</u> canal segments in Venice.

modeling canal geometry

Average widths were very useful for the creation of the geometry of the hydrodynamic model discussed in later sections. In modeling applications, the complex network is reduced to an abstract chain of parallelepipeds characterized by the length, average width and average depth of the segments that they represent. For practical purposes, the longer and the more irregular the network arc, the less useful a single measure like the average width will be. For instance, a cargo boat driver looking at the thematic map of the average canal widths (above) could detect some potentially troublesome canals that one should probably avoid traversing with a wide boat (the red canals in the above map), but there is no guarantee that there are no additional choke-points in the other canals that would still impede navigation. The narrowest width of each segment is extremely useful for practical purposes, and the author's company (*Forma Urbis*) has just completed all such measurements for the development of a boat traffic model for the City[176].

narrowest point

176 See p. 82.

depth or bathymetry

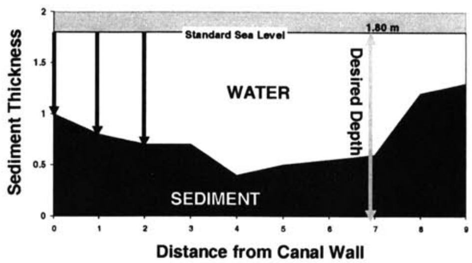

Measuring the depth of the canal segments was a non-trivial endeavor. First of all, depth is a three-dimensional quantity that is hard to represent succinctly. Ideally, depth could be determined with some sort of a sonar or radar device that would produce a topographical map of the bottom of a canal. Unfortunately, since canals are only about 1-2 meter deep on average – depending on the tide – the error that many commercial electronic sounding devices produce is often greater that the entity being measured, thus making such an automatic approach not feasible. Instead, we resorted to manual soundings of the canal bottom using a weighted tape measure that was reeled down from the surface to the bottom of the canal at intervals of one meter across the canal's width, starting with a measurement along the northernmost wall and ending with a measurement along the southernmost wall[177].

normalizing to the absolute zero

To take into account tide fluctuations, we had to normalize all our soundings to the so-called "mareographic zero" of 1897 – a standard reference level for all tide-related measurements in Venice. To do this, we had to monitor tide levels in parallel with our bathymetries, in order to measure the difference of each of our depth soundings from the standard zero. These parallel tide measurements were conducted near the location of the bathymetric sounding by a team member who measured the height of the water from a sidewalk of know altimetry (or elevation).

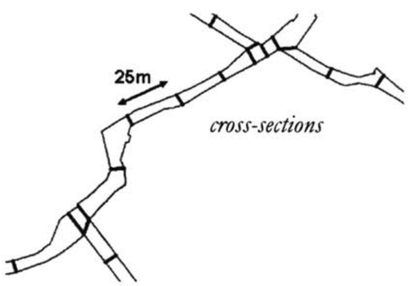

cross-sections

Cross-sectional soundings were repeated at intervals of 25 meters along the length of the segment. In some ways, this could be interpreted as a "sub-atomic" segmentation of the atomic segments. Each section was uniquely identified by the code of the segment in which it occurred plus the distance of the section, in meters, from the northernmost end of the segment along the centerline. This type of sub-segmentation is similar to the so-called "dynamic segmentation" used in "linear referencing systems" to, for instance, pinpoint accident locations on road segments[178]. The Venetian bathymetric cross-sections prove that whatever unit of space is chosen as the "fundamental atomic particle", there will invariably be cases in which sub-atomic divisions will be necessary, therefore all we can really do is to settle for a "reasonable" happy medium and adopt the breakdown that seems most appropriate as the unit of analysis for the majority of users.

dynamic segmentation

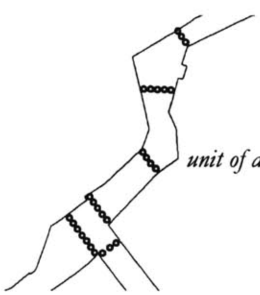

unit of analysis

Whereas we could determine the definitive lengths, surface areas and average widths for <u>all</u> segments, we only obtained reliable depth measurements for 130 of the 367 segments[179].

177 This north-to-south order was not always maintained, but an arrowhead on the section lines in the corresponding GIS layer indicates the exact direction of measurement for each section.

178 Some GIS products, like ArcGIS, support dynamic segmentation on-the-fly. See also Fletcher *et al.*, 1998.

179 The Insula company has since contracted additional bathymetric campaigns to determine the depths of the remaining canals.

soundings

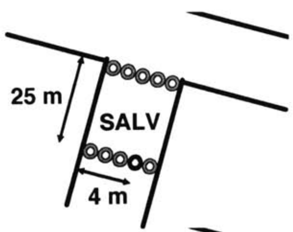

average depth

In the end, the bathymetric dataset was represented by 7,768 individual measurements along 850 section lines. We chose not to represent each individual measurement graphically as a dot in its exact location along the section[180], though this approach may have been useful to determine navigation paths, as explained below. The tedious complexity of such a task made us opt instead for a sort of recursive dynamic sub-segmentation, achieved by including the section's identifiers (i.e. the segment code and the distance of the section from the northern end of the segment) with each bathymetric sounding in the database. In the linked database, each depth measurement was labeled by the code of the segment where the measurement was taken (e.g. SALV), plus the distance of the sounding from the northern end (e.g. 25) and the distance of the sounding from the northern wall of the segment (e.g. 4).

Similar to widths, depths are also difficult to express uniquely for each segment. Moreover, unlike widths, average depths cannot be easily computed from other GIS-derived dimensions, but require instead extensive field campaigns. As mentioned, almost 8,000 measurements[181] had to be conducted – all at night to avoid interference from boat traffic – to quantify canal depths. On average, 55 soundings were performed in each canal segment. The average of all the depth measurements in one segment was used as the overall average depth of the entire segment[182].

modeling canal depths

Average depths, like average widths, proved useful for modeling purposes, but not too practical for everyday life. However, there is another useful utilization of the average depth – in conjunction with the surface area – namely for the calculation of sediment volumes, which are used, among other things, for estimating dredging costs. Being the first to systematically determine these dimensions, we were also the first to be able to produce reliable estimates of the volumes of sediment that needed to be removed from each canal segment. Such volume calculations obviously depend on a definition of where the bottom of a segment lies.

sediment volume

180 Each sounding location could in theory be determined algorithmically and mapped with a small GIS application utilizing the information already available.

181 In reality, well over 10,000 measurements were conducted, but only 7,768 were validated and transferred to Insula in the framework of the UNESCO-Insula accord. In all, the bathymetric campaigns took over 12,000 hours of work to complete.

182 In fact, in order to make the average depth as realistic as possible, the average was computed omitting measurements along the canal walls on both sides of each section.

canal bottom

Being semi-natural watercourses, canal segments do not have an easily identifiable bottom. For practical purposes, the city of Venice has arbitrarily defined the so-called *fondo di progetto* (project bottom) of each segment as the desired depth that the segment should be maintained at to allow smooth navigation of emergency boats even during low tides. Such a bottom is generally set between 1.8 and 2 meters below mean sea level[183]. Once the target bottom depth is defined, sediment thickness is simply derived by subtraction.

water volume

Complementary to the notion of sediment volume is that of water volume. In a perfectly dredged canal, all of the volume, dictated by the surface area and the desired depth, would be occupied by water. As sediment begins to accumulate, the total capacity of the canal segment is shared by a combination of water and sediment. Water volume is indirectly useful in modeling applications, both for hydrodynamics and for transport of suspended matter in water, as will be discussed later.

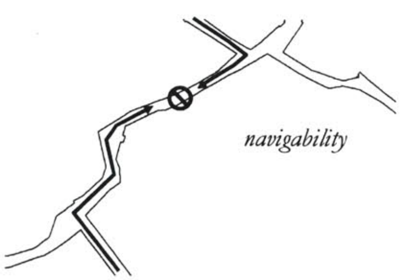

Depending on the draft of one's boat and on the tide conditions, one would probably tend to avoid canals with low average depths (red in map above) no matter what. But, once again, there would be no guarantee of

[183] I have publicly argued whether this depth is to be taken from the standard mareographic zero of 1897 or from today's mean sea level, which is 23 cm higher than the standard. I even calculated the difference in cost, due to the additional sediment one would remove if measurements were based on the absolute zero of 1897, which adds up to several million dollars. My view is that nature (or global warming) has done us a favor in this case by rising the mean sea level thus sparing us from dredging about a quarter of a meter of sediment throughout the entire network. We should take advantage of this bonus and not squander resources on unnecessary excavations.

being able to navigate the segments with medium average depths (yellow above) or paradoxically even the ones with the highest average depths (green). One shallow point in an otherwise deeper channel would be sufficient to completely impede passage, even though that canal segment may be very deep on average as a whole.

shallowest bottleneck

Looking at just the data, one may be tempted to identify the shallowest data point as the bottleneck. This would be fallacious, since the absolute shallowest points are almost always along one of the canal walls. Knowing that such points are the shallowest would be useless for all intents and purposes. No boat could possibly travel along the canal walls anyway. In fact, a boater would always attempt to tread a route through the deepest part of the canal, so that, at any moment, the boat would be passing over the deeper point of a particular cross-section, staying well clear of the shallowest. Therefore, what would be truly useful, is to know where these deeper points are across any section. We did precisely that, with a two-step manipulation of the fundamental soundings dataset. First of all, we identified the deepest points of each section, and then we selected the shallowest of these to truly determine the navigability bottlenecks. In our research, we defined what we called a *navigability axis* as the line connecting the deepest points of each of the bathymetric sections. Thus, even though we could not vouch for the depth of any water between sections, we could determine the point along this navigation axis where the water was shallowest, hence providing a more useful indicator of navigability to local boaters. This "shallowest of the deepest" concept is one demonstration about the difference between data and information.

partial navigability

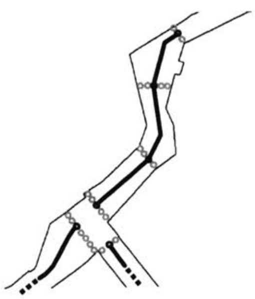

navigation aids

As a further example of the usefulness of dynamic segmentation, one could imagine an intelligent application that could determine, for a boat with a given draft traveling during a given tide phase, the impassable sections in all canal segments. Depending on the exact destination of the boat's travel, one could even imagine that a system could allow travel part of the way down a segment with a known shallow obstruction, as long as the boat did not have to go through such a point, but was going to dock before reaching the impasse. Although we didn't use such level of sophistication, the ambulance dispatching application discussed earlier[184] could potentially benefit from such dynamic segmentation of the canal network.

data longevity

Unlike the other physical dimensions, the determination of a canal depth is ephemeral in nature. With the passage of time, additional sediment will accumulate and canals will get shallower, rendering our measurements obsolete. Also, as Insula proceeds with its dredging program, canal segments will be deepened all the way to their *fondo di progetto* and sedimentation will start over from a *tabula rasa.* The action of human beings and nature combine to make these measurements short-lived, but not necessarily futile. In fact, our data collection jump started the operations of Insula when it was a fledgling company in 1997. Since then, several canals have been dredged and Insula has maintained the information up-to-date by resetting the bathymetries to the *fondo di progetto* for those segments. More

data maintenance

184 Page 34.

complex is the issue of maintenance of the information for segments subjected to natural sedimentation. One way to monitor the silting process is to conduct periodic measurement campaigns similar to the ones we pioneered. However, this would be a rather expensive proposition. An alternative that we have proposed to Insula, through UNESCO, is the creation of a sedimentation model to simulate the natural processes and estimate sediment levels over time, to plan preventive maintenance on a regular basis. This model is discussed in more detail later. Our data, even though it is obsolete by now, is still useful as a baseline to compare with subsequent measurements. The "growth" of sediment over repeated measurements, using our initial efforts as a baseline, will help determine sedimentation rates and thus inform the development of the sedimentation model.

sedimentation model

lessons from our physical studies

Having set up the geographical framework made up of coded spatial objects, this aspect of our canal studies, despite its simplicity, taught us a few additional lessons.

Lesson one was that many of the physical characteristics of real-world objects are easily managed by GIS as long as the representation of the real objects is "literal", i.e. unfiltered and un-simplified. If we represent a canal segment as a "region", i.e. a polygon whose shape reflects as accurately as possible the shape of the actual segment, then we can rely on the GIS representation to tell us the basic physical dimensions (like area and perimeter) automatically and for free. Simplifying the canal network into a series of centerline segments may bring some efficiency savings in some arenas (like calculating the segment length and enabling transportation modeling of links), but it also produces a net loss of usefulness in many other areas. The fact is that these two representations are not mutually exclusive. What worked for us was to use the geometric representation that best suited our analytical needs.

[do not oversimplify reality]

[represent objects as they are]

[create geometry to suit needs]

In general, this lesson suggests to create GIS objects that reflect the true shape of the objects in real life, which means to use "regions" (or polygons, a.k.a. "shapes") whenever an object with a real surface area is being represented. More specifically, it seems appropriate to represent road-like elements (like our canal segments) using both a region/shape geometry as well as linear objects (centerlines).

[multiple geometries can be appropriate]

Lesson two was that most permanent physical traits can be measured and archived once and for all, as soon as we have settled on an appropriate atomization of the objects. GIS provides us with basic measurements for free, but we first have to define the exact extent of each object using the atomization that best suits our needs[185]. The choice of how much detail to include and where to draw the line while divvying up reality is not so simple[186]. In fact, just when we thought we had identified the "indivisible" atoms of the canal network – namely the segments – we immediately ran into "sub-atomic particles" when we had to deal with depth

[capture permanent physical dimensions]

[prepare for sub-segmentation]

185 Cf. Evans, 1997.

186 See lesson 1 on page 50.

measurements at cross-sections. The bottom line is that we will use whatever segmentation makes sense at the time. Some of these partition decisions may have staying power and prove useful for an overwhelming majority of uses and some may not. But we should not refrain from making these choices lest no progress will ever be made in representing our urban features[187].

[2-D vs. 3-D]

[update ephemeral dimensions]

[modeling predictable processes]

[modeling vs. re-measuring]

The natural process of sedimentation immediately emphasized the shortcomings of a 2-D GIS platform. Depth is an inherently 3-D measure, so we couldn't quite visualize it in our Mapinfo system[188]. This process is also dynamic, so our measurements are ephemeral and bound to become quickly obsolete. Yet, the sedimentation process is gradual and somewhat predictable, so our third lesson was that it makes sense to develop a sedimentation model to predict the silting up that will happen over time. The cost of such a model, if successful, would certainly be cheaper, in the long run, than repeating the bathymetric measurements periodically for ever and ever.

187 Fletcher *et al.*, 1998.

188 For an example of how we dealt with this shortcoming, see the 3-D graphing functionality we added to our *SmartInsula* application (page 102 ff.).

HYDRODYNAMICS

tides

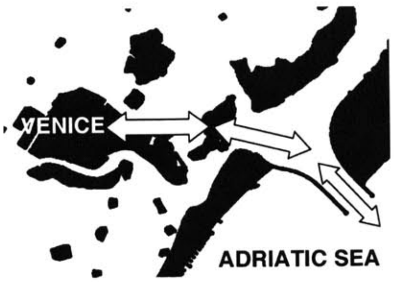

acqua alta (flooding)

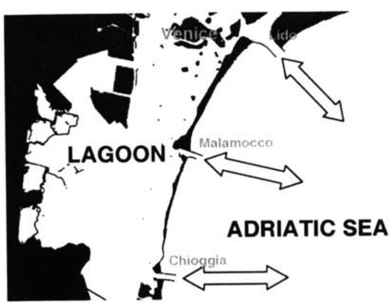

currents

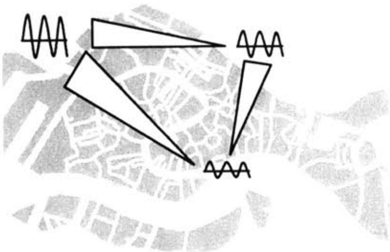

hydrodynamic behavior

Tidal currents in the canals are responsible for flushing the urban sewage out to sea, redistributing sediment along the canal bottoms in the process. Hydrodynamic currents are created predictably by the cycles of the moon and less so by meteorological phenomena such as barometric pressure and winds. Stagnant canals will not benefit from the cleansing action of tides and will thus become malodorous pools that negatively affect the lives of residents and visitors alike. Stimulating good hydrodynamic activity in all canals has always been a goal of government administrations in Venice for centuries. The worst of the stagnant canals were the first to be converted to *rii terà*[189].

Tides in Venice are infamous primarily because of the phenomenon of *Acqua Alta* (high water), which has grabbed most of the headlines since the record-breaking event of November 4, 1966, when 1.94 meters of water submerged the entire city and all of the islands of the lagoon [190]. Since then, numerous national and international initiatives have been undertaken to safeguard Venice from floods. Most notorious is the long-lasting controversy about the construction of flood gates at the three lagoon openings (a.k.a. *bocche di porto*) of S.Nicolò (Lido), Malamocco and Chioggia. The biblically-titled *MOSE* project, developed by a consortium of large construction companies called the *Consorzio Venezia Nuova* (CVN) has been quite appropriately in the works for a biblically long time (since about 1984), with an infinite litany of approvals and rejections by a wide variety of agencies, both Italian and international, that have monopolized the attention of public opinion away from some of the more pressing issues affecting Venetian citizens, some of which are described herein. In fact, tides are not a "bad" thing at all, as long as they keep below the threshold of 1-1.2 meters above standard sea level (1897). Unless and until Venice is endowed with a modern sewage system, tides will continue to play a crucial role in the hygienic health of the city.

The ups and downs of tide levels, together with the complex topologies of the lagoon channels and of the internal canals, create height differentials in the water levels in different parts of the city and of the lagoon such that downhill gradients are created that dictate the direction and velocity of flow of water from the higher points to the lower points. In most cases these differentials are ever-so-slight and therefore produce very minimal flows, but especially in conditions of full or new moons, during so-called *spring tides* (*sizigie* in Italian), the influx of tidal waves from the Adriatic sea into the lagoon produces significant flows in many parts of the canal network, with peak currents of more than 2 Km/hr in places.

From 1990 until 1999, several teams of WPI students collaborated on the systematic recording of the hydrodynamic behavior of all of the segments that make up the Venetian water network[191]. The various campaigns generally entailed the simultaneous measurement of as many segments as the number of members of the research team allowed. Three to

189 Cf. page 50.

190 Florence was subjected to an even worse flooding by the *Arno* river on the same day.

191 We even engaged the local schools in two days of simultaneous current measurements involving over 1,000 children from all middle schools in Venice and in the lagoon.

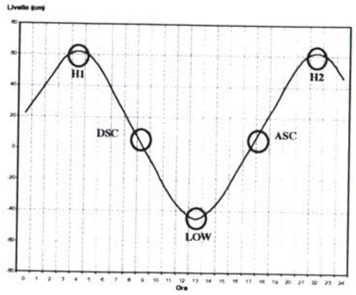

six sets of segments were measured by the team in each campaign both during *spring tides* (full or new moons) as well as during *neap tides* (half moons) for about two months at a time. Each day of campaign entailed 5 measurements at the predicted High tide time (H1), the Low tide time (LOW) and the subsequent high tide (H2), plus the two highest-velocity moments half-way between high and low (DSC) and then half-way between low and the next high tide (ASC). At the end of the study, the hydrodynamic behavior of practically all the inner canals was determined and mapped, both for incoming (bottom) and for outgoing tides (next page)[192]. The arrows in these pictures indicate the prevailing direction of flow during that type of tide. Red arrows indicate a general flow toward the north and west, whereas the green ones indicate a general tendency to flow south and east.

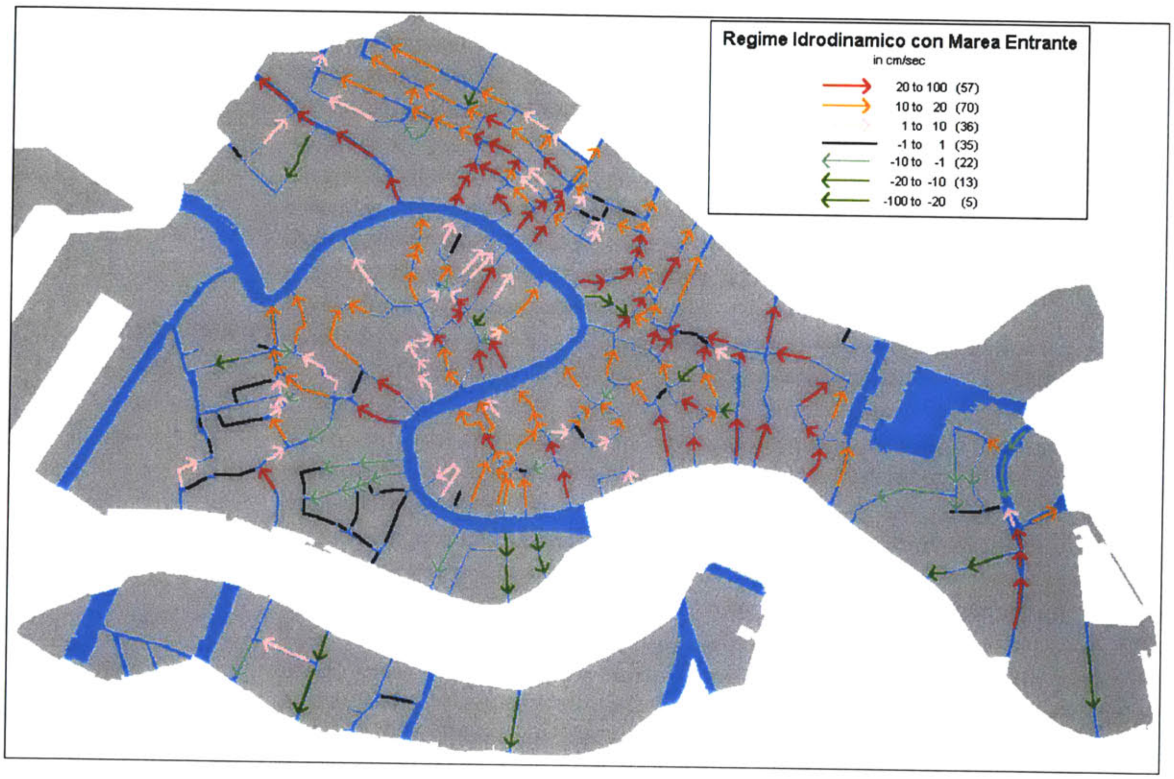

Mapping the currents in this manner had to wait until some advanced features became available in Mapinfo® that allowed to create thematic maps using bidirectional arrows. The arrows are based on plain line segments, drawn with a south-to-north direction along the centerlines. Black arrowless segments are used to represent the absence of flow, i.e. stagnant canals.

192 The maps shown in this section come from Carrera, 1999b.

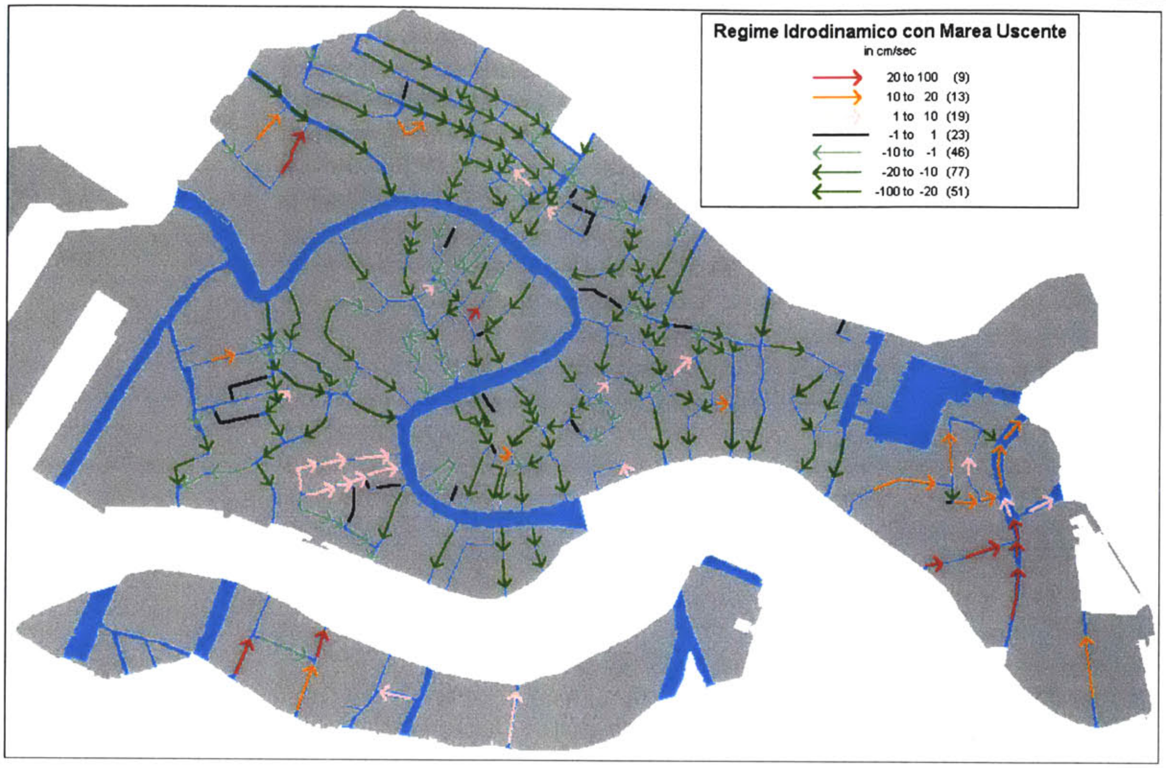

cyclical regularity

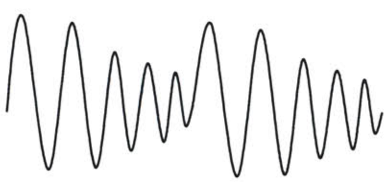

In the previous section, most of the physical characteristics of canal segments, with the exception of canal depths, were immutable and permanent. An effort to characterize parameters such as surface, length and average width would produce a definitive body of knowledge that would be good forever and would only need to be collected once. The hydrodynamics of a body of water like the lagoon of Venice are constantly changing from day to day, due to the moon cycle and to weather patterns. Nevertheless, despite its cyclical nature, hydrodynamic behavior is also rather repetitive and displays a high degree of regularity from year to year. Unless major topological changes are made to the underlying physical environment through which tides travel, one would expect to get fairly constant peak current velocities and unwavering current directions on the same segment under the same lunar conditions, after the "noise" due to weather has been factored out.

hydrodynamic model

Such steady behavior makes investment into the development of a hydrodynamic model more appealing, since such a model would also only have to be developed once (and possibly adjusted or tweaked once in a while). In fact, our pioneering work on hydrodynamics led to the development of the first hydrodynamic models of the circulation in the inner canals of Venice.

reutilization of datasets

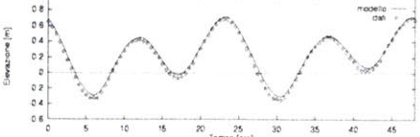

Starting in 1995, two different hydrodynamic models of the inner canals were created in the framework of the UNESCO-MURST Venice Inner Canals project. Both of these models made use of the fine-grained datasets collected by us in our hydrodynamic campaigns. Curiously, even though we focused primarily on the typical hydrodynamic behavior during *coupling GIS and models*

spring tides, which we arrived at by discarding unusual readings with high standard deviations, the hydrodynamic modelers instead used whatever data were available for specific measurement days in their entirety. This was an example of the reusable nature of our fine-grained, spatially referenced datasets, even though the modelers used separate systems (and FORTRAN!) to run their simulations on our data.

The issue of coupling GIS with models remains important and worthy of a separate discussion. The UNESCO models were developed at UC Berkeley[193], the International Institute for Applied Systems Analysis (IIASA)[194], Georgia Tech[195] and the National Council of Research in Venice[196].

The hydrodynamic model of the inner canals was interconnected with existing models of the Lagoon to allow simulations to start with the forcing functions of the tides at the Lido inlet. In the final model, the inner canals were represented by one-dimensional segments and these in turn were connected with a 2-dimensional finite element model of the surrounding lagoon channels, as shown below[197].

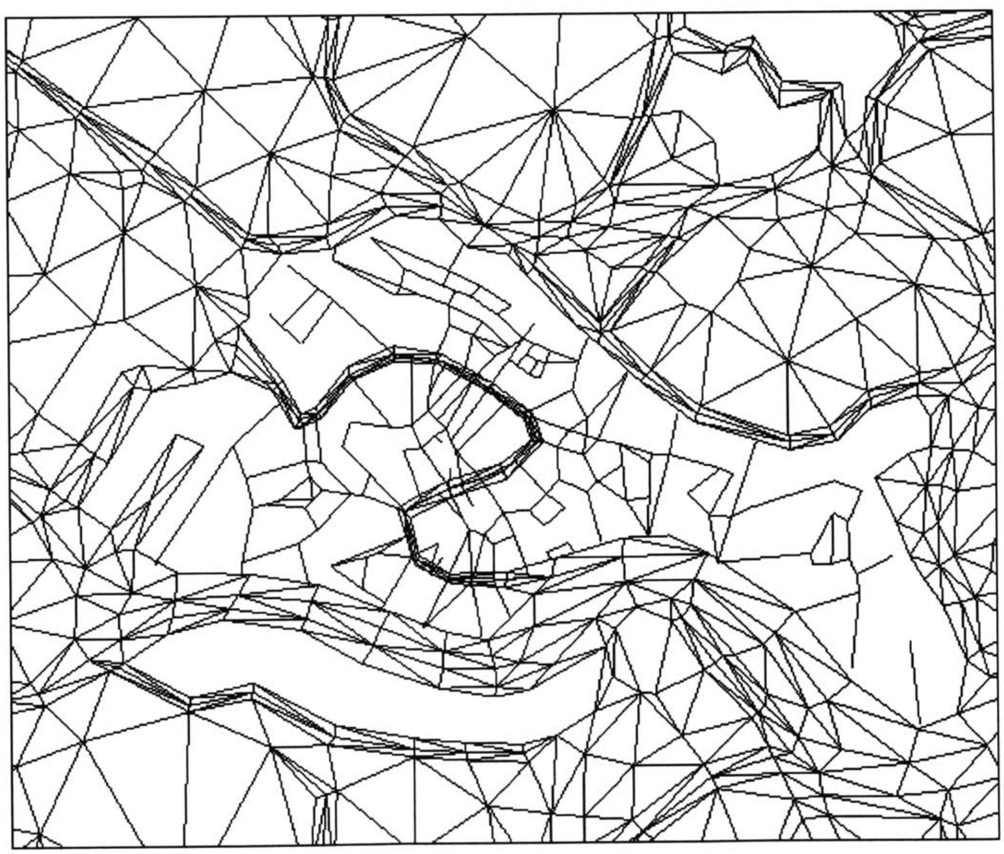

193 De Marchi, 1993.

194 De Marchi, 1996.

195 De Marchi, 1997.

196 Umgiesser and Zampato, 1999.

Hydrodynamic models could be used to forecast the effects of local interventions (like the dredging of a canal) on surrounding areas; to determine the most effective dredging sequence; to simulate the reopening of filled-in canals; to analyze the possible utilization of pumps to stimulate circulation in stagnant areas and also to act as a foundation of more complex models of sediment transport and deposition[198]. This sophisticated sedimentation model, that will be discussed in later chapters, will need, among other things, to take into account the re-suspension of sediment due to boat traffic, which is the topic of the next section.

lessons from our hydrodynamic studies

Hydrodynamic behavior is predominantly guided by the gravitational pull of the moon and the sun, coupled with the topology of the paths that the water has to travel along. The physical dimensions of the segments that make up the canal network and their topological connections are thus essential to understanding hydrodynamics. The cycles of the moon and the sun are equally important. Hydrodynamic behavior displays a cyclical regularity and as such it is a candidate for modeling. Knowing the physical dimensions of the network, together with measurements of tide levels and current velocities allowed the calibration and validation of the model.

[cyclical regularity and modeling]

As we saw here, fine-grained hydrodynamic measurements are eminently re-usable. With some statistical manipulation, they allowed us to characterize the "permanent" behavior during the highest tides (*spring* tides), giving us the predictable current direction and velocity in each segment at the time of fastest flows with both incoming and outgoing tides. The same data – this time taken in its raw, untreated form – was later re-used as input to the model calibration and validation. After mapping our "statistical", maximum current scenarios we could have thrown the data away, as people often do after achieving the aim at hand. Thankfully we didn't, so modeling was that much easier to carry out and that much cheaper too, since there was no need to re-measure the currents and levels all over again.

[re-usability of fine-grained data]

simulating and forecasting

The lesson therefore is twofold: (1) collect data at the finest grain possible (within reason) and (2) do not throw it away.

[use finest grain]

[keep data in raw form]

A more sophisticated lesson here was that the canal segments that were used in the model were already "right" for hydrodynamics since we had already thought ahead to the possibility of hydrodynamic modeling when we conducted our segmentation. Hence, as you may recall, we had decided to use intersections with filled-in canals (some of which retained hydrodynamic flows under them) as nodes of our canal network. Had we not acted teleologically, we may have had to create a "new" segmentation just for this modeling exercise[199]. So the third lesson here is: think ahead! And structure your data accordingly.

[think ahead!]

[enrich maps and data teleologically]

[re-use model itself]

A final lesson that further compounds the benefits of a City Knowledge approach is that the hydrodynamic model, which brought

197 Can you detect the outline of the city of Venice in the figure?

198 More on the sedimentation model starting at page 92.

199 As we have had to do for our traffic model (see page 82 and ff.) to characterize pseudo-nodes corresponding to docks and bridges.

[hierarchical re-utilizations]

together plan-ready data about the physical canal network, and the fine-grained hydrodynamic measurements to determine the "permanent" behavior of the canal currents, is itself, as a whole, a component of an even bigger model – the sediment transport model – of which it forms the backbone. So, here we have a multiplicity of re-utilizations at different hierarchical levels of complexity, which further demonstrates the power of City Knowledge.

BOAT TRAFFIC

motorboats

moto ondoso

Row boats and sail boats have traveled the canals of Venice ever since the fall of the Roman Empire, but motorboats have only been around since after WWII. In the few decades since their widespread diffusion in the 60's, motorboats have become one of the major problems the City is facing today. Boat traffic in Venice not only generates the problem of congestion that most mainland cities are gridlocked by, but it is actually a much more insidious phenomenon than its land cousin. Unlike automobiles, boats have actually the ability to physically destroy the foundations of buildings along their path. The passage of cars on a street cannot dismantle a city block piece by piece like water turbulence and wakes do when boats ply the inner canals[200].

Despite the massive exodus of Venetians away from their city of birth, boat traffic continues to increase in order to cater to the needs of a growing tourist industry. According to our own calculations, traffic has almost doubled in the last 25 years[201], as population declined by 50% in the same period. People who live along the canals are literally seeing their dwellings crumble under the constant pounding of boat waves. In desperation, many property owners have begun a public protest against the daily assault on their homes. "*Stop Moto Ondoso*" (loosely translatable as "stop making waves") is a common phrase that even tourists are beginning to recognize after seeing it posted on sheets hanging from windows all over the city.

The *moto ondoso* problem is particularly intense along the primary arteries. In fact, despite the apparent intricacy of the web of canals, the entire water network can be schematically simplified to just a few primary routes where most of the boat traffic concentrates. Although there are 367 segments and 182 inner canals in the city, for a total of 47.5 kilometers of waterways, the main thoroughfares can be counted on the fingers of two hands.

---

200 The problem is so big that the mayor has been nominated *Commissario al Moto Onsoso* ("Commissioner of Boat Wakes") by the national government in 2002 to deal with this issue once and for all. As commissioner, the mayor has absolute powers over regulations related to boat traffic, above and beyond any existing law.

201 Fiorin and Carrera, 1999, pp. 143-159.

inner and outer ring

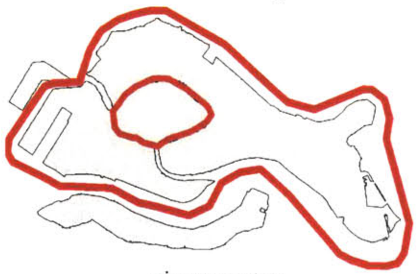

One could schematically represent the important traffic arteries in and around Venice starting from two concentric rings[202] (red rings in figure on left). The outer ring includes the *Canale delle Fondamente Nuove*, *Colombola*, *Tronchetto*, *Giudecca*, *Bacino di S.Marco*, and the *Canale delle Navi*. The inner ring is made up of the central part of the Grand Canal and the entirety of the *Rio Nuovo*.

ring connectors

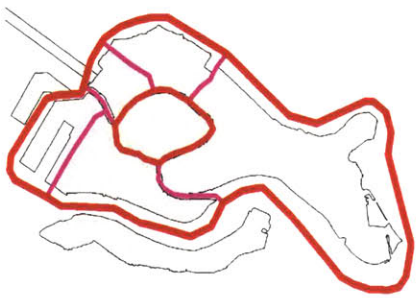

The next level of arteries include the *Rio de Noal*, *Canale di Cannaregio*, the northernmost and southernmost tails of the Grand Canal and the *Scomenzera*. These are all canals that connect the inner and outer ring (purple). The canals that are part of the ring and the ones that connect the rings are among the most traveled in Venice, especially by taxi boats and large cargo vessels. Public transportation boats from the ACTV (*Azienda del Consorzio Trasporti Venezia*) travel all of these canals, except the *Rio de Noal*.

secondary arteries

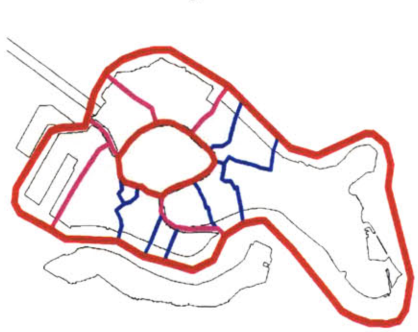

Beyond the two rings and the ring connectors, there are two more types of arteries that play an important role in the flow of traffic in Venice, the so-called *secondary arteries* (blue in the side diagram) and the *bypass* canals (yellow). The secondary arteries include the *Rio de S.Sebastian, Carmini, Briati,* and *Tre Ponti.* connecting the *Canale della Giudecca* to the *Rio Novo*. The blue canals also include the two parallel canals in Dorsoduro that allow travel between the *Canale della Giudecca* and the *Canal Grande* (*Rio de San Trovaso* and *de San Vio*), and the two parallel canals that connect outer and inner ring across the borough of San Marco (*Rio de S. Moisè* and *Rio de la Canonica, S.Zulian, de la Fava* and *Rio del Fontego dei Tedeschi*). The list of secondary arteries is completed by two canals that connect inner and outer ring to the north and East, namely the *Rio dei Santi Apostoli* and the *Rio de Santa Marina* (which actually connects the inner ring to one of the bypass canals described below).

bypasses

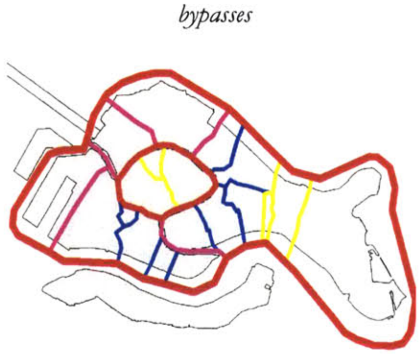

The final major elements of the water traffic network in Venice are canals that *bypass* parts of the network and create shortcuts (yellow) between two parts of the inner and outer rings. For instance, the *Rio de San Polo* and the connecting *Rio Marin* e *Rio de San Zandegolà* together create a "Y-shaped" bypass that cuts right through the heart of the inner ring inside the boroughs of S.Polo and Santa Croce. Similarly, the *Rio de Santa Giustina* and the *Rio dei Greci* and *Sant'Antonin* form an inverted Y that connects two parts of the outer ring in the Castello region, as does the *Rio de l'Arsenale* a little farther East.

202 *Ibid.*, p. 144-145.

previous studies

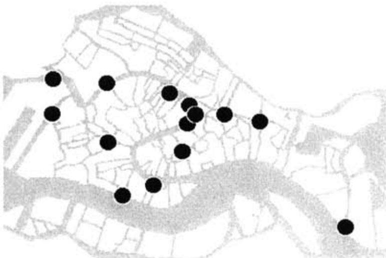

The canals listed in the above discussion could be identified as "important" transportation links by anyone who has experienced Venetian boat traffic first hand, but the relative levels of traffic in each of those arteries, despite the obvious importance of such knowledge, were not systematically quantified until WPI students began to record traffic flows in 1992. In truth, the City of Venice had conducted three studies before WPI students began their research in the framework of the UNESCO *Venice Inner Canals* study. The 1978 City study was never published, but was referenced in the following 1986 City Traffic Report[203]. Both the 1986 and 1987[204] City studies attempted to give a snapshot of the traffic levels at key points in the canal network through the monitoring of "maneuvers" that resembled what automotive traffic engineers call "turning movements".

turning movement counts (TMC)

Unlike professional Turning Movement Counts (TMCs), though, the studies conducted by the City of Venice in 1986 and 1987 did not exhaustively monitor all of the "maneuvers" that a boat could make at any given location. In fact, at some counting locations, the counts resembled more the "volume counts" that mainland traffic engineers conduct along traffic routes using "automatic traffic recorders" (ATR)[205], i.e. total number of transits in each of the two opposing directions along a single network segment.

automatic traffic recorders (ATR)

WPI-UNESCO traffic studies

Starting in 1991, our canal studies (not yet under the auspices of UNESCO) began to analyze traffic experimentally[206]. Our early forays led to a full blown traffic study, in the framework of the UNESCO-MURST project *Venice Inner Canals.* Between 1992 and 1994, six two-month campaigns were conducted in Winter, Spring and Summer, recording almost 60,000 transits in almost 400 hours of counting as part of an effort to which 21 WPI students dedicated a total of about 15,000 thesis-hours. These efforts were summarized in a final report to UNESCO in 1996.[207]

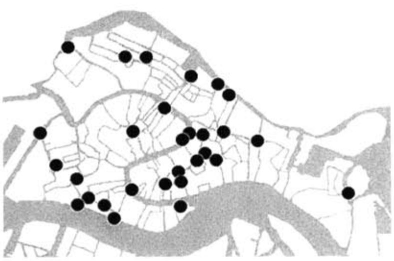

stations

Our traffic campaigns monitored traffic from a total of 29 stations (intersections) in the city of Venice, concentrating exclusively on the inner canals and only marginally touching the main thoroughfares of the City, like the Grand Canal, and the canals of the *Giudecca*, the *Fondamente Nuove* and the *Scomenzera*. Each traffic counting station was uniquely labeled with an ID that reflected the borough (*sestiere*) where the station was located (e.g. San Marco stations were prefixed by the letters "SM") followed by a numerical suffix (e.g. SM-4). Counting stations were selected in such a way as to maximize the number of canal segments that could be observed and they were placed along the main thoroughfares that were described earlier in this chapter.

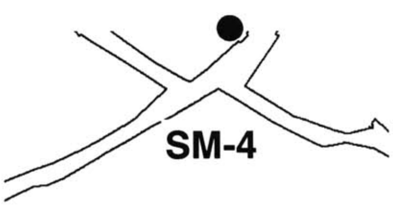

turning movements

At each location, a number of "maneuvers" were monitored, that captured all of the possible turning movements available to a boat at that intersection[208]. For any number *n* of segments that connect at an

203 Comune di Venezia, 1986.

204 Comune di Venezia, 1987.

205 These are the rubber strips laid across roadways that we have all noticed at some point as we ran over them with our cars.

206 Carvajal *et al.*, 1991.

207 Carrera, 1996. An English summary is available at www.unesco.ve.it.

208 Excluding U-turns.

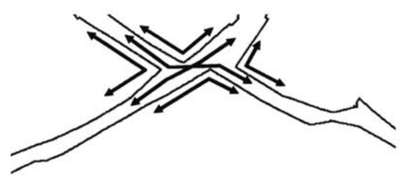

maneuvers

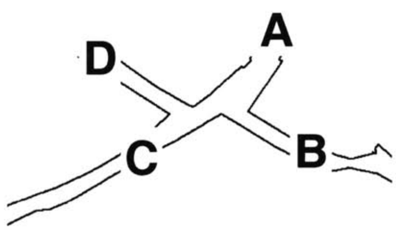

from turns to volumes

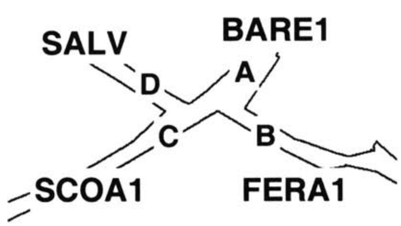

from traffic data to traffic information

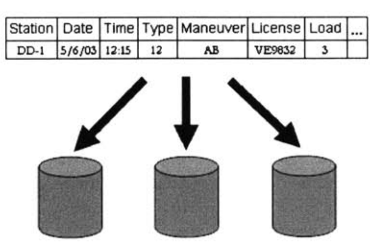

intersection, the number of available moves turns out to be $\binom{n}{2}$ which is read as "n choose two" and translates to a mathematical formula of $\frac{n!}{(n-2)!}$, which in turn simplifies down to $n\,(n-1)$. This simply means that a three-way intersection will allow a total of 3 × 2 = 6 turning movements, a four-way will permit 4 × 3 = 12 moves, a 5-way 20 moves and so on.

To capture all of the traffic at each station, we labeled each of the approaches to the intersection with consecutive letters of the alphabet, starting from the northernmost approach and moving clockwise around the intersection. In this way, any particular boat would be recorded as making maneuver AB or BA or AC, and so on. Contrary to what is commonly done in mainland traffic counts, we decided to actually record every single boat that went by, identifying its type, load (number of people or amount of cargo, in ⅛ increments) and even its license plate[209]. Each passage was time-stamped in 15 minute intervals. The number of records in the database is exactly equal to the number of transits recorded. Each record in the database is uniquely connected to a specific turning movement (labeled according to the "from" and "to" approach letters) at a specific station (labeled with a station ID). Although we came up with this method on our own, it isn't very different from what traffic engineers commonly use around the world to record turning movements. The main difference is that, for automotive traffic, movements are usually recorded as Left, Right or Straight.

Using appropriate ancillary tables, each lettered approach at each station was related to the codes of specific segments touched by a boat making that particular turn. Total traffic volume in any specific segment was simply equivalent to the total count of boats that showed the corresponding letter of the alphabet in either the <u>from</u> or the <u>to</u> maneuver fields. So, for example, if one wanted to know the total number of boats that traveled up or down the segment coded as FERA1, a query would be set up to count all of the boats that either came from or went to the segment labeled B, namely AB+DB+CB+BA+BD+BC.

The structure of the boat traffic data we collected was rather simple. As mentioned, each boat that was observed passing by a counting location was recorded as a single record in our field forms and in our databases. Each record was characterized by a station code, a date and time stamp, the turning movement (maneuver) the boat performed, the type of water craft observed, its license plate, name and permit number, and the load it was carrying. Such a simple raw dataset allows numerous and complex manipulations that can produce higher-order information, by reconfiguring the data along any of the characterizing dimensions. For example, one could want to explore the temporal distribution of taxi boats in the course of a day, to find out when the peak times of taxi traffic occur. Or one could

209 License plates were not present on each boat at the time of our studies, though a 2003 ordinance has finally made them mandatory for every vessel. They were only recorded starting with our study for the Provincial government and COSES in 1997 (Carrera, 1997). The data we recorded in the framework of the UNESCO project did not include license plates.

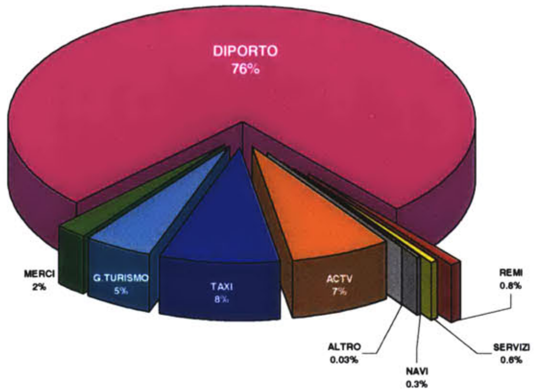

be interested in the number and/or percentage of boats of each type that travel a certain canal over the entire day. Or else it may be of interest for someone to know the average number of motor boats per hour that make a certain turn at an intersection on weekends. All of these are plausible requests that would satisfy specific informational needs. Each inquiry can be resolved by crafting an SQL query (or a sequence of queries) that will address the issue. Custom-tailored answers to these pointed questions are always possible, as long as a technical expert is at hand who can deftly interact with the database.

Similarly, a GIS expert could manipulate the information to produce geospatial renditions of the information, creating thematic maps to accentuate one particular pattern of behavior or another. In GIS, however, the temporal component would be much harder to depict, although some techniques do exist to portray time in GIS, at least in part[210].

depicting traffic data

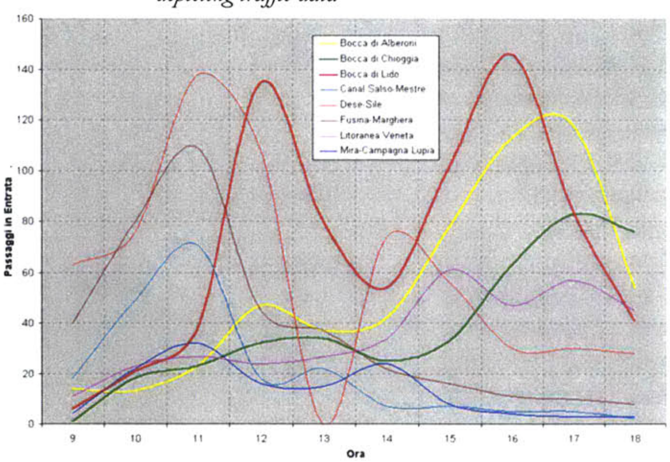

Nevertheless, if one had to produce a final report on a traffic monitoring campaign, as I had to do on several occasions[211], then the issue becomes one of depicting a complex phenomenon in a digestible manner. Instead of focusing on specific areas, a general report would have to provide the "big picture". Graphs of various types come in handy to display some of the trends and behaviors observed, but these inevitably entail pretty massive spatial aggregation. Bar graphs, pie charts and line graphs are useful when looking at the entirety of traffic in the whole city either as a total or as averages or percentages. If dozens of canal segments were monitored, it would be impossible to show all of the segment codes on the X-axis of a graph while presenting the traffic levels in each segment on the Y-axis.

thematic maps

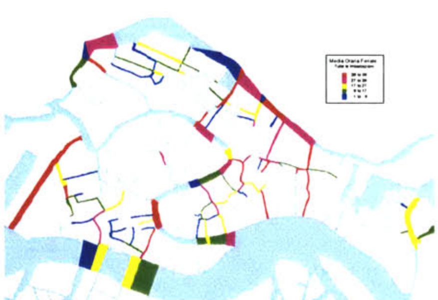

GIS tools allow the depiction of spatially-specific information in its entirety, as long as the area to be displayed is not too vast in comparison to the individual elements being depicted. The city of Venice is right about as big an area as is feasible to thematically map, when the units of analysis are the narrow canal segments. When it comes to thematic mapping, traffic data lends itself to three primary visualizations: by maneuver, by counting station and by segment. These three units of spatial analysis are not just appropriate in Venice, but also in mainland cities, as I have personally experienced in the cities of Worcester, Cambridge, and Quincy in Massachusetts[212].

210 See bottom of p. 80.

211 See Carrera, 1996, 1997, 1999a.

212 See part III.

traffic by maneuver

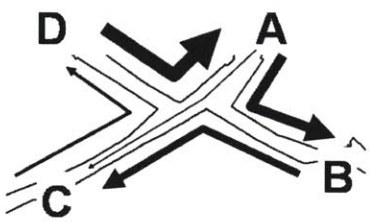

The type of spatial visualization that most closely approximates the fine grain at which traffic data are collected in the field is the representation by maneuver. In this depiction, each turning movement would be represented by an arrow and the color or size of each arrow would be representative of the number of boats making that turn. The quantities could represent the total number of turns in a day of counting, or an average per hour, or any other time scale. This finer-grained visualization allows a more in-depth analysis of driver behavior and can give clues to the preferred pathways across the city.

The problem with a presentation by maneuver is that it is difficult to view many stations in a single map, since the arrows tend to blur into each other as the view is zoomed out more and more. Nevertheless, careful placement of the arrows will allow a decent number of stations to be viewed simultaneously. With some effort, the entire city of Venice was visualized in a single map (shown below are total daily transits by maneuver).

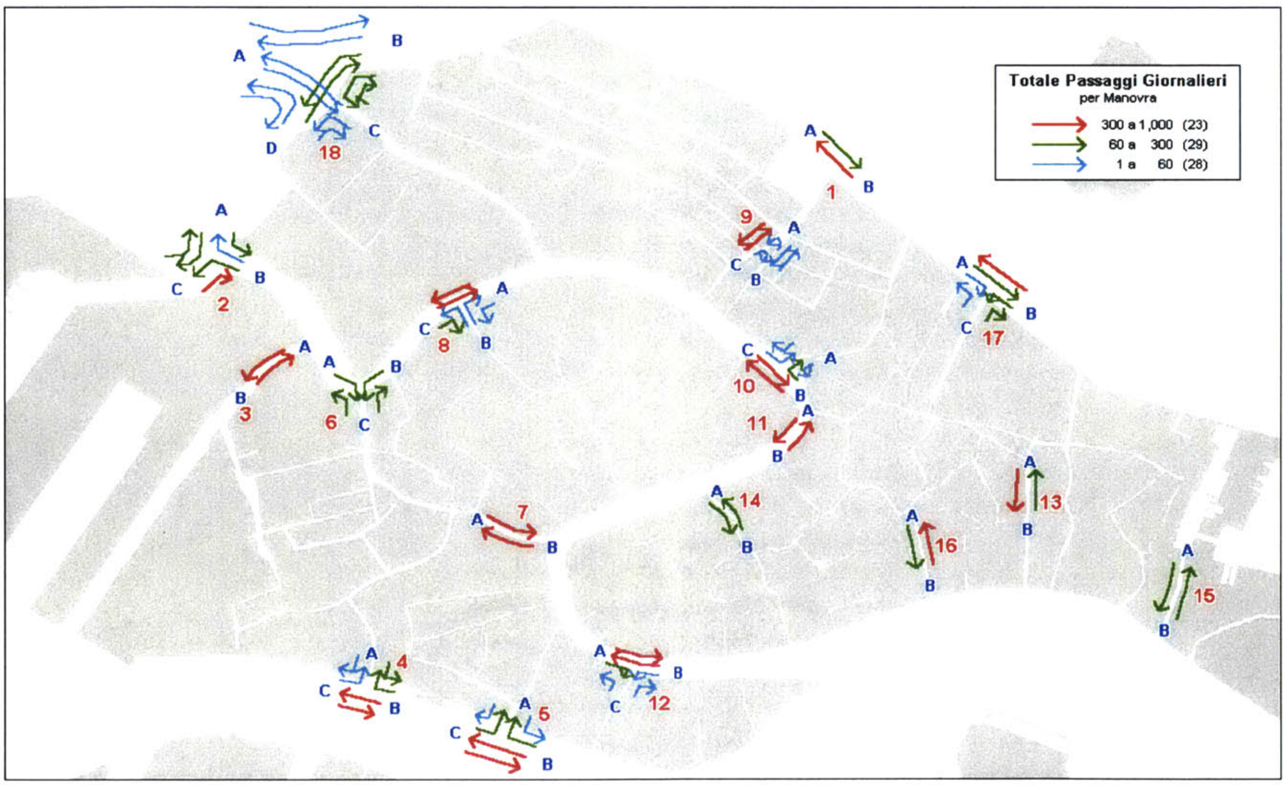

traffic by station

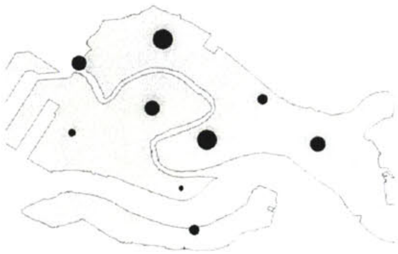

The representation of traffic counts by station (or counting location) is the simplest of all. It entails aggregating all traffic counted at each counting location, regardless of the turning movements. These data may be visualized as an overall total for a whole day of counting, which would be applicable only when the counting was conducted uniformly for the same amount of time, on the same day, at each of the stations. Alternatively, the data could be averaged per hour of counting, thus removing many of the possible idiosyncrasies, though uniformity of field methodology would still be necessary. These totals or averages could be easily depicted by graduated symbols, whose size is proportional to the value being depicted (the total or average/hr).

Additionally, the type of boats that were counted at each station over the whole time period of choice could be displayed by replacing each dot with a pie chart showcasing the percentage of each boat type that contributed to the total. The size of each pie would still convey the overall volume of traffic at that station, whereas the size of each slice would give a proportion per boat type.

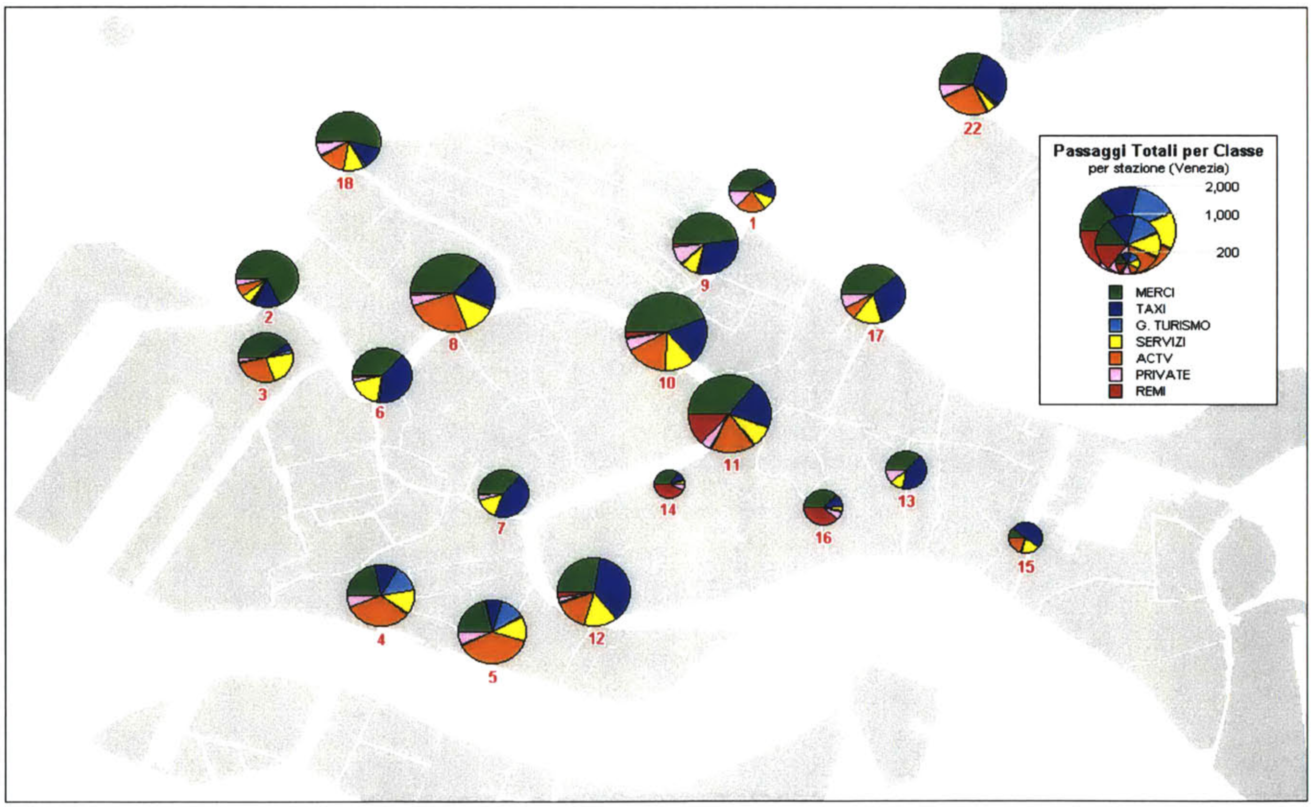

critical sums

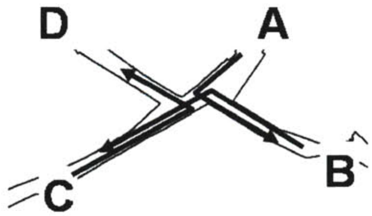

The usefulness to traffic engineers of a representation of overall traffic volumes at an intersection, however, is rather minor. A more valuable visualization that is also based on each intersection (or counting location) is what traffic engineers call "critical sums", which represents the sum total of all crossover movements, which interrupt straight through traffic. Critical sums are not too difficult to obtain starting from turning movement counts, but they would be impossible to obtain from typical automatic traffic recorders.

traffic by segment

ATR vs. TMC volumes

The third and final means of display is that by segment. In this representation, the volume of traffic in each segment can either be calculated from turning movements or can be obtained by automatic traffic recorders. The volumes based on TMCs are more nuanced than the simpler ATR volumes, but the latter can be much cheaper toobtain for much longer periods of time. Despite their resource-intensive nature, the breakdown of TMCs into segment volumes holds the potential to allow the extrapolation of ATR counts to adjoining segments, through a statistical knowledge of the preferential turns made by cars at the intervening nodes.

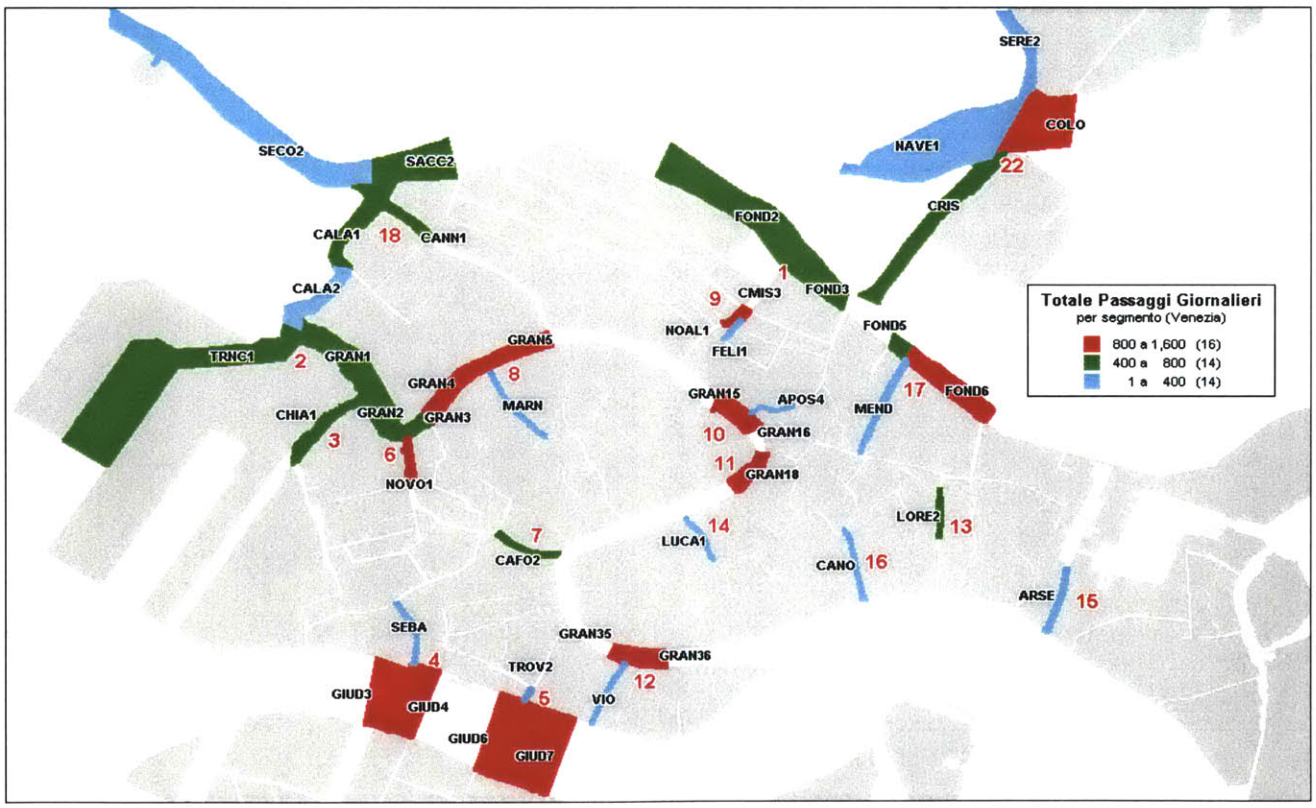

In this type of visualization, each segment is colored thematically depending on the total (or average) hourly traffic it sustained during a temporal interval of choice (usually hour or day). This type of presentation facilitates the identification of the main arteries that carry most of the traffic in the city. Segment volumes can be used to determine the carrying capacity of a certain pathway, as well as to forecast the number of boats displaced when a segment is closed for upkeep. Displaced boats will tend to be diverted to adjoining canals, thus compounding themselves upon the traffic volumes already existing in those nearby segments.

segment traffic and moto ondoso

Segment-specific issues like wear and tear, congestion and noise will be directly related to traffic volumes regardless of whether the city streets are made of water or asphalt. In Venice, total volumes of traffic in each segment are also useful to know because they relate to the total amount of *moto ondoso* that is produced by motorboats traveling along the segment. Direction of motion is immaterial in all of these contexts, whereas it may be very significant in other types of analyses.

direction of travel

An example of where direction of travel is important, is in the analysis of the effectiveness of directional regulations. Venetian canals, like roads in mainland cities, are regulated by one-ways. Some directions of travel are prohibited in some canals, although going the wrong way on a one-way canal is not usually as dangerous as the equivalent action on the mainland. In fact, wrong-way travel is rather frequent in Venice. To capture the extent of this phenomenon, we wanted to separate the traffic volumes in one segment into two separate volumes in the two directions of travel.

extracting and encoding direction

In a database, separating the traffic flows into two separate directions is fairly trivial, though there are some standardization issues that need to be addressed. For instance, how should the two directions be labeled? Our choice was to separate traffic into North and South directions based on whether a boat was moving towards the northernmost or southernmost node at the two ends of a segment. The northernmost node would be the one with the higher Y-coordinate (or latitude). Even in the more challenging situations, when a segment is oriented East-to-West, one of the two nodes would always be ever-so-slightly farther to the North than the other[213].

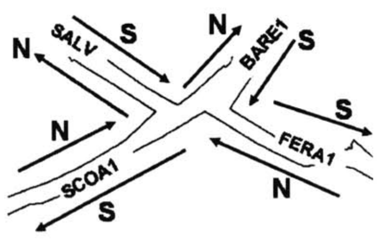

Keeping the traffic information separated into the two directions of travel along each segment became useful later, when we wanted to explore the effects of the creation of a new one-way canal using the MANTA traffic model[214]. This level of disaggregation will also provide a more sophisticated appraisal of the local, direction-specific effects of canal closures on boat movements.

depicting directional traffic

One of the practical issues that we had to deal with, when separating segment traffic into two directions, was the issue of how to depict the two flows (geo)graphically. One possibility is to create two overlapping layers of polygons, one representing the Northern segments and one the Southern ones, so that thematic coloring of the regions could be possible, based on the volumes recorded in each of the directions. The problem with this approach is that only one direction could be shown at a time. Another possibility is to attach two arrows to each segment, a N arrow and a S arrow

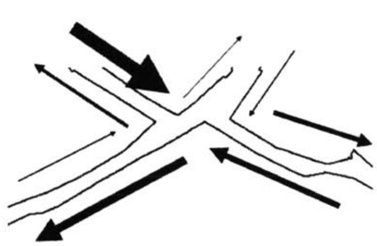

213 In fact, the intersections (nodes) in Venice were coded with numeric labels proceeding from North to South (again based on the Y coordinate of the centroid) so that implicitly the Northernmost is always the node labeled with the smaller numeric ID.

214 See page 82.

so that the traffic volumes for each direction could be thematically depicted by coloring (or varying the thickness of) these arrows appropriately. The drawback here, as explained above for the representation of maneuvers, would be once again the difficulty in visualizing large areas of the city on a single map, since the arrows would tend to tangle up above a certain zoom level.

analyzing traffic data

Traffic data lends itself to a myriad of analyses. Information could be necessary for very focused decisions about regulating traffic in a specific canal, or for the development of very broad planning guidelines for the entire city. A long time period may be of the essence in one determination, whereas a very narrow spatial extent may be highly significant in another.

analysis by vehicle type

classes and categories

The typologies of boats add another dimension to the combinatorial explosion of the innumerable cross-tabulations that are possible and desirable. To bring down the complexity of dealing with the 18 boat types, many of the analyses in my reports are limited to just the 9 boat classes (with 6 classes being the truly important ones), or to the 6 use categories, all the way down to the 4 main boat classifications of Cargo, Recreation, People and Other.

temporal analysis

Another way of looking at the data is on a temporal scale. Peak hours may vary from day to day and from boat type to boat type. Overall traffic volumes at each hour will identify times of congestion, whereas a type-by-type analysis of the hourly fluctuations may stimulate reflections on the underlying causes for the different peak times, which in turn may suggest actions to limit the overlapping of peaks, by creating temporal restrictions for some canals at specific times of the day for specific types of boats. For instance, gondolas could be barred from some inner canals until 11am to give cargo boats space to complete their deliveries.

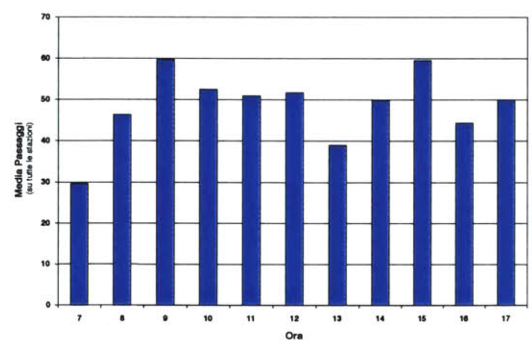

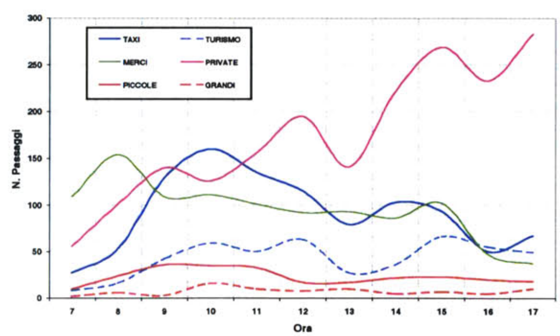

seasonal variations

Due to the influence of tourism, traffic in Venice exhibits some distinct seasonal trends, with an increase in overall volumes in the summer months. The nice weather also encourages private boat owners to use their vessels for recreation, further compounding the problem, especially out in the Venice lagoon, as evidenced by our studies conducted in the entire lagoon on behalf of the Venice Provincial Government[215].

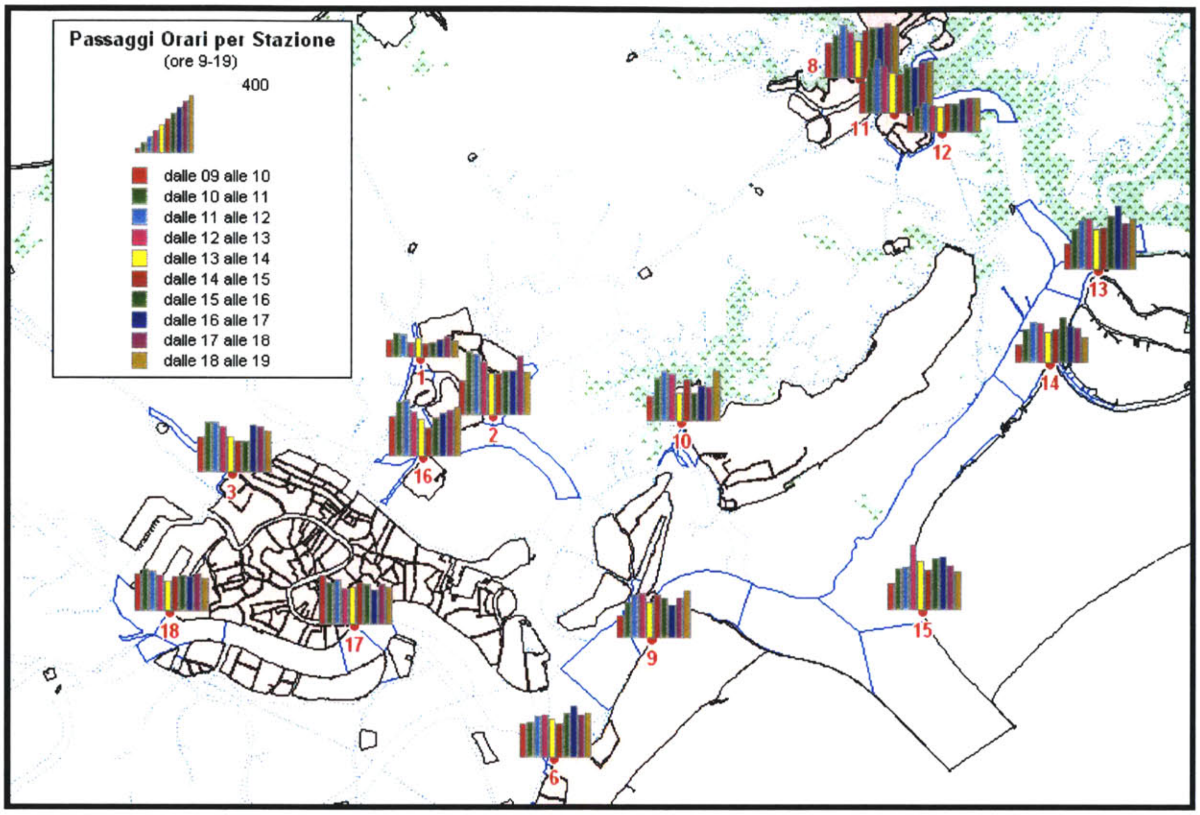

traffic decision support systems

The intricacies of the analyses of traffic data are reflected in the commensurate complexity I encountered when I designed a preliminary prototype of a Decision Support System (DSS) for Venetian traffic engineers. The DSS needed to first of all provide instant information about boat traffic on any of the Venetian canals, predicated of course on the availability of such data. Even this simple spec turns out to be much more complicated than one would expect at first glance. Exactly what data does the user want to extract from a specific segment? The total traffic? The average hourly? For a weekday or weekend? For just one type of boat or for all (or some)? It is easy to see how quickly things can get really messy, even when dealing with a single segment and one single counting campaign (the latest one of course). Imagine the difficulty of allowing the user to select a set of segments or permitting the analysis of longitudinal trends over several datasets!

215 Carrera, 1997, 1998.

multimedia interface

The prototype I proposed would allow the user to visualize data by segment, maneuver or station, plus it allowed one to focus on any category, class or type of boat. The interface conforms to a "look and feel" that I have used in other information systems. It allows the user to simultaneously see data, photographs (of the station location), maps (at varying zoom levels) and even graphs (with a variety of choices of representation). Although it was never fully implemented, this DSS would have allowed multiple segments to be aggregated together, and would have allowed the user to look longitudinally at past campaigns on the same segments.

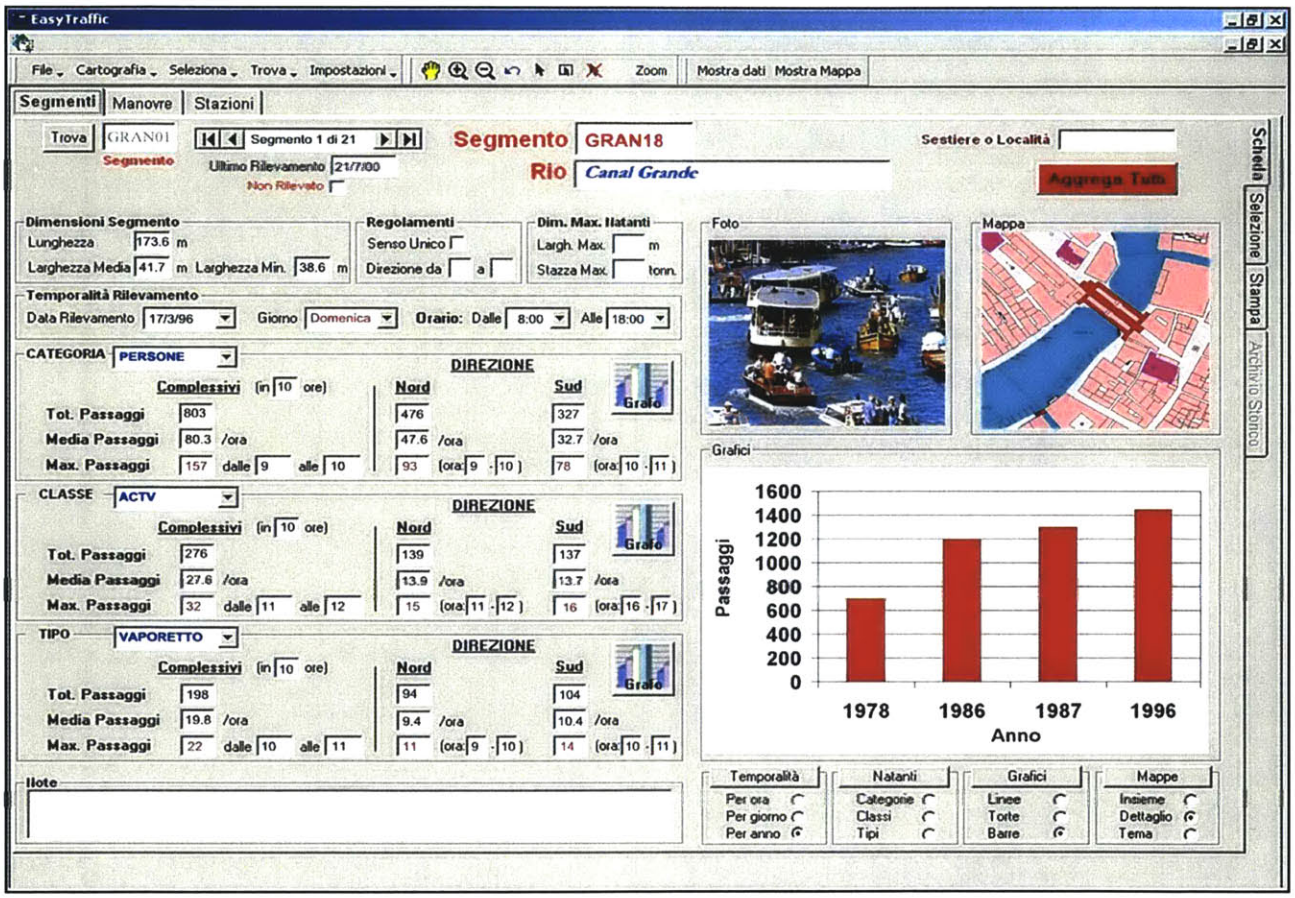

traffic modeling

MANTA

origin and destination

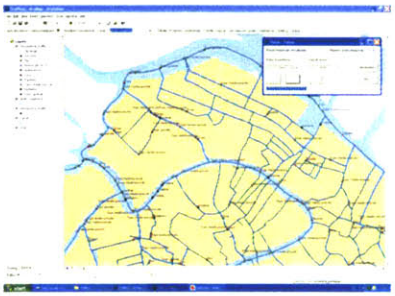

The traffic DSS is a stepping stone towards a more complex traffic model. Late in 2003, I was asked to contribute to the development of just such a traffic model for Venice[216]. This traffic model differs substantially from *terraferma* models in that Origin-and-Destination (O&D) information is only marginally useful in Venice, where nobody commutes to work on his/her own boat. In Venice, rush hours are due exclusively to commercial traffic, like cargo or garbage, as well as people transport, like taxis, gondolas and public transportation (ACTV). O&D studies would have to be replaced or integrated with very localized "detour studies" to determine the local O&D within a sub-network, and thus forecast the most likely diversion a boat of a certain type would take when a canal is closed or precluded *ex legis* to navigation.

Despite its idiosyncrasies, the Venice traffic model needs to answer the same questions that "normal" land traffic models typically address. It should allow city officials to forecast the effects of canal closures and to simulate the consequences of changes to regulations, like the institution of new one-ways, or the temporal restriction of access to some canals for some boats at some time of the day, or the effect of alterations to the routes of such regular services as garbage pick up and other realistic scenarios in which the city can purposely attempt to modify water transportation for the better. The development of the model has benefited tremendously from existing, plan-ready city knowledge.

moto ondoso index

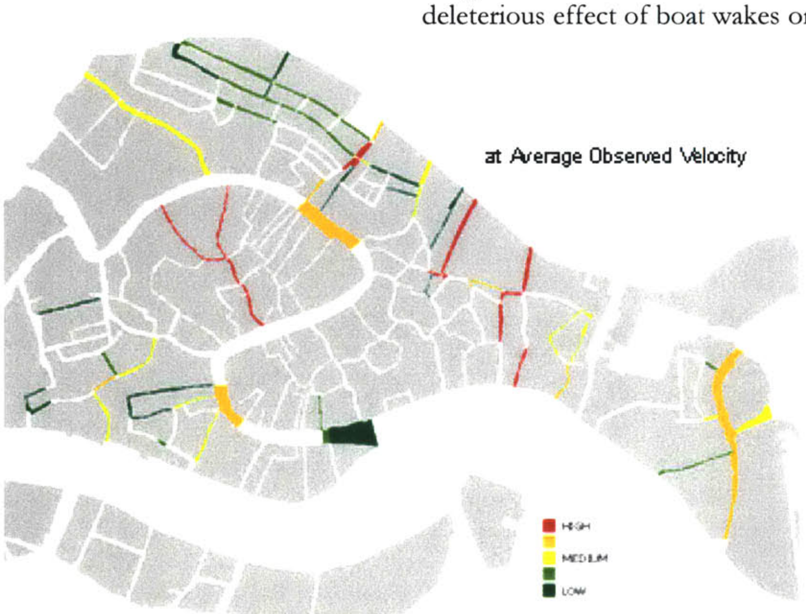

Since the real problem with traffic in Venice is not so much the congestion that occurs in the canals during rush hours, but rather the deleterious effect of boat wakes on the foundations of buildings that flank the canals, in the summer of 2002, we developed a metric dubbed the *moto ondoso* index that would parameterize the wake-production potential of each boat type[217]. In this framework, boats that produce lower wakes than other boats, while traveling at the same speed, would be assigned a lower "moto ondoso coefficient". The *moto ondoso* index would then be the product of these coefficients (related to boat type and velocity) and the traffic volume for each specific boat type in each of the canals. At any arbitrary maximum speed limit, the overall index for any segment would be the sum of all of the indices of each of the boat types that had been recorded in that segment during the traffic

216 The *Modello per l'ANalisi del Traffico Acqueo* (MANTA) was completed in the winter of 2004 and presented to the public in the summer of 2004.

217 Chiu *et al.*, 2002.

monitoring campaigns. Once again, our plan-ready city knowledge made the development and quantification of this higher-order concept fairly expeditious.

turbulent energy discharges

While the *moto ondoso index* was being developed, we immediately realized that the observable waves created by boats are only the more visible symptom of the problem. Another major factor in the destruction of canal walls would certainly be the turbulence created by propellers undergoing sudden acceleration or – even more significantly – deceleration, during breaking or docking maneuvers. To measure the effect of turbulent energy discharges in the canal network, I supervised a WPI, on-campus project for the creation of an instrument[218] that enabled us to create a map of turbulent discharges in the canals. The instrument was tested in the summer of 2003 and the plan is to duplicate it and mount it onto a dozen boats of various typologies for an extended period of data collection.

updating traffic data

Traffic data are ephemeral by its very nature. Today's data are of little use to us next year, except to evaluate change and to analyze trends. Unlike the more permanent features of the physical make-up of the city, traffic is a dynamic human activity that requires periodic updating. To make the collection of traffic data easier, mainland traffic engineers make use of devices called Automatic Traffic Recorders (ATRs) that normally consist of pneumatic hoses connected to electromechanical boxes, usually chained to a post or guardrail along a road. Recent developments in this arena have seen the adoption of buried inductive spires that sense the passage of cars due to electromagnetic disturbances. Radar devices are also used in some applications, despite their high cost. Although the most recent ATRs are capable of distinguishing several types of vehicles, from trucks to small sport cars, they can only be used to collect volume information along single segments of a road network. Moreover, despite their otherwise impressive capabilities, inductive pickup devices cannot reliably count motorcycles, scooters or bicycles. These devices are also more intrusive in their installation, making them more suitable to being installed as permanent implants, frequently associated with traffic light controls.

automatic traffic recorders

turning movement counters

Turning movements need to be collected with different devices, like Turning Movement Counters (TMCs), that are generally still manually operated, by pushing a button for each vehicle making a specific turn. These simple devices usually have upwards of 4 or 5 buttons for each possible turn, one for each of 4 or 5 types of vehicles. These devices can be used to count any type of vehicle, and they are often used also to count pedestrians (or *peds* as they are called in the lingo of traffic engineers).

218 Chiu *et al.*, 2004. See page 115 for a picture of the instrument.

RFID boat monitoring device

For obvious reasons, standard ATRs cannot be used in Venice. Radar devices, on the other hand, have great potential applicability in Venice's narrow canals, as proven by an experimental instrument that we developed at WPI, which was capable of measuring boat wakes and could recognize an approaching boat using a Radio Frequency (RF) tag, similar to those that are used to automate the payment at tollbooths (in Massachusetts it's called Fastlane™)[219]. The problem with this device was that it assumed that every boat would be "tagged" with a Radio Frequency ID (RFID). The benefit of this device, despite its intrusive nature, would be the exact identification of each individual boat, which would allow the ability to tap into very precise information about the boat, through the ID code. Unfortunately, the fear that the ID would be linked to enforcement and ticketing of traffic infractions made this device politically unpopular.

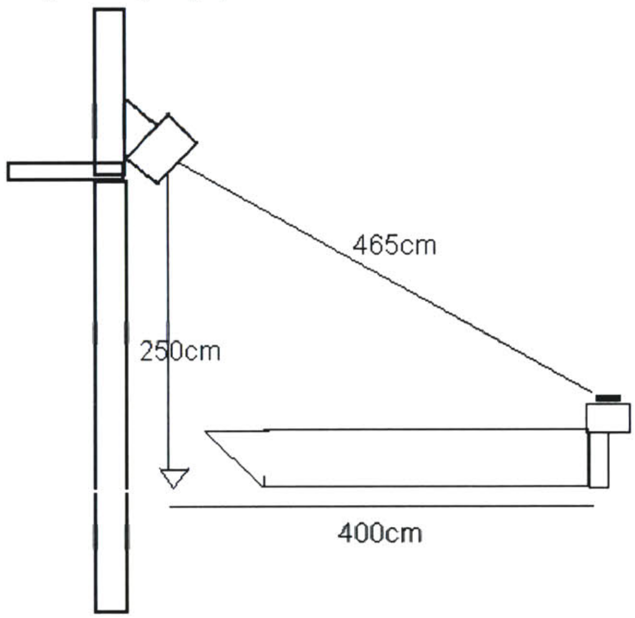

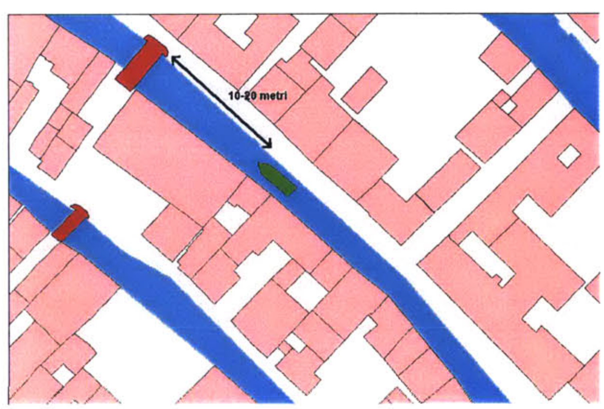

219 Blomberg *et al.*, 1999.

GPS boat monitoring device

Although the RFID solution was ultimately rejected on "political" grounds, the alternative that was subsequently pursued by the city of Venice was a GPS system installed in all licensed boats (taxis as a start, then licensed cargo boats at a later date). The oddity is that the objections that made the RFID undesirable were not really done away with in a GPS system. Each unit is still tagged and identified and its whereabouts are known with even more precision. The danger of these devices being used for purposes beyond mere knowledge and statistics, i.e. the possibility of the GPS being used to detect speed limit violations and the like is exactly the same as with the RFID system, yet this system was installed on the entire fleet of Venetian taxis in late 2003.

ultrasound boat monitoring device

The limitation of the GPS device, as it is being used in Venice, is that it will only provide a partial picture of the traffic in the city since only "licensed" boats will be forced to carry the devices. All other boats will not be detectable through these gadgets. To allow a more generalized automatic ability to monitor all types of boats, we have been developing a second-generation instrument at WPI, based on ultrasound technology, that will allow the automatic monitoring of boat traffic, without imposing the installation of an RF tag or a GPS upon every boat owner[220]. The next version of the instrument will count boats, determine their velocity, measure their wakes and noise levels, and quantify the pressures exerted on the canal walls by each wave.

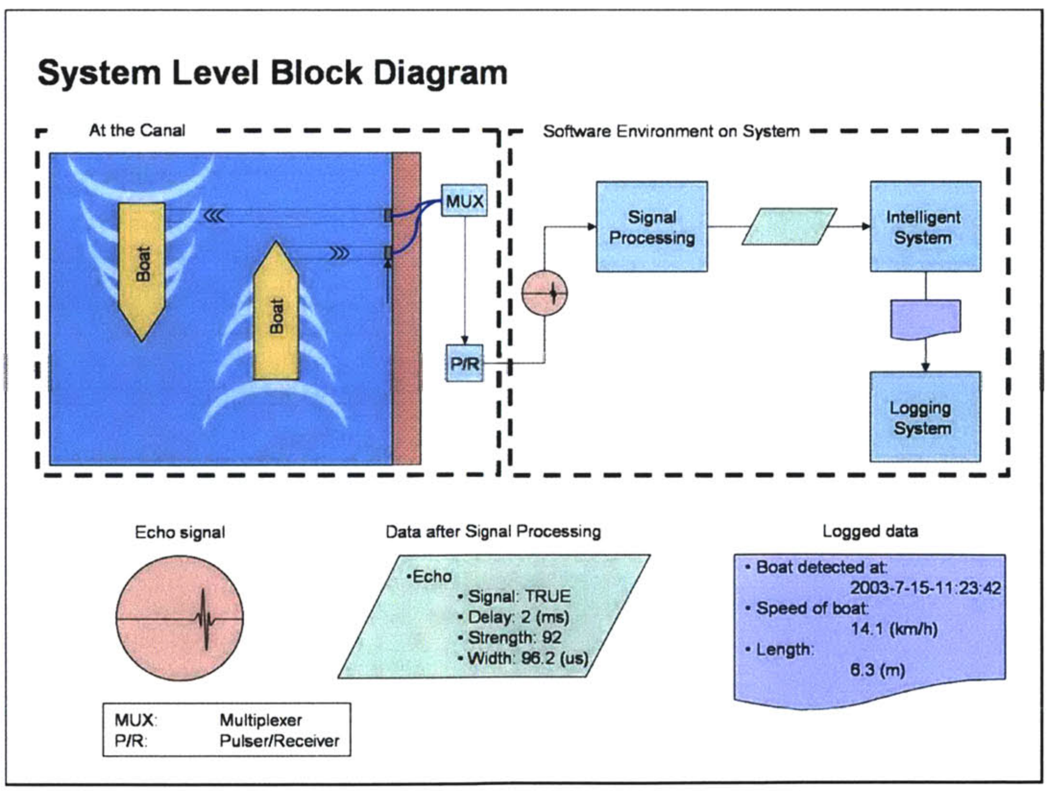

220 Johnson *et al.*, 2003 and Gomperts *et al.*, 2004.

automating traffic volume collection

Our ultrasound devices are being designed to eventually be placed at key locations in the canal network where they will collect traffic volumes autonomously. Each boat will be archived as a record that will be tagged with a time-and-date stamp, plus with an indication of the boat type (detected through the acoustic signature of the engine noise) and with the speed and direction of travel, plus the height of the wake produced and the pressure the wake exerted on the canal wall (measured through a series of pressure transducers). These devices will collect data continuously for months and years, and will exceed what Automatic Traffic Recorders (ATRs) do for vehicular traffic.

automating turning movement counts

Unlike ATRs, Turning Movement Counters (TMCs) are potentially usable in Venice, though we never actually used them, preferring to employ paper field forms which allowed us to collect a much richer qualitative dataset than the merely quantitative data produced by the electronic TMCs. In general, any automated device will allow longer time periods to be monitored with little or no human intervention, but these benefits will be offset by the fact that the data produced will be much less rich in nuance and detail. TMC's do not automate the collection anyhow, since human beings are still needed to push the buttons. Basically TMCs constitute a shortcut in the data entry process, since they remove the need to transcribe paper forms and input the data into a computer database.

manual traffic data collection

Despite the existence of technological solutions to monitor traffic, good old fashioned manual counts will always retain an important role in the periodic updating of the more nuanced and possibly more useful traffic data. Origin and Destination, for one, would be never possible using technological devices unless every car is somehow tagged and identified automatically. It is difficult to foresee the day when sophisticated traffic analyses such as trip-chaining would be possible without a substantial investment of human resources to collect and organize the necessary underlying information.

Since periodic traffic counts are inevitable, cities can be creative in how to make these cyclical updates as effortless and inexpensive as possible. In Cambridge, MA, we experimented with modifications to existing city regulations that force developers of large sites to produce traffic studies in order to gain approval for these sizeable construction projects[221]. Since these are already accepted requirements, all we had to do is figure out how to parlay these professional reports into a system that would "grow" and accept all studies into a single framework, which would thus continually provide the best available traffic information for all areas of the city for which data were available.

221 See Gage *et al.*, 2003.

the infrastructure of traffic knowledge

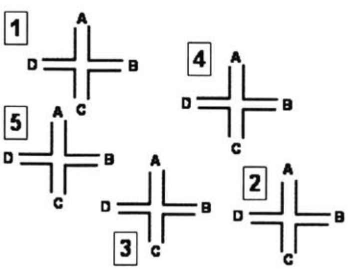

structure of traffic GIS

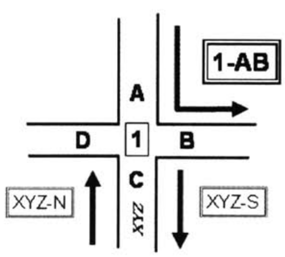

structure of traffic databases

In fact the creation of an informational scaffolding into which to structure traffic data is a *condition sine qua non* for the creation of any serious traffic information system[222]. Three basic spatial objects need to be defined, positioned and labeled on a GIS platform. These same labels will then be reflected in the associated relational databases that will contain the bulk of the actual traffic data. The three spatial entities that need to be firmed up are the aforementioned: stations (or intersections), maneuvers and segments. The latter is assumed to be already in existence if a proper road (or canal) network has already been structured for other purposes. Intersections are beginning to be standardized nationally, through an effort of the Census Bureau to include them into their *Tiger* maps[223]. Maneuvers are thus the only traffic-specific GIS layers that need to be mapped appropriately. Maneuvers are usually labeled with a unique identifier that combines the intersection code as well as the maneuver code, to simplify the linking to the associated database of traffic counts.

In addition, a layer of bidirectional traffic arrows would also be useful to attach directional information to each road segment, to allow even more nuanced types of traffic analyses and simulations, like those that would accompany the creation of new one-way streets.

In all, though there can be considerable room for personalized referencing standards from one town to the next, the items that require a definitive representation on a traffic GIS and a clear and unique labeling are:

- the Intersections (or *nodes*)
- the Maneuvers (or *turning movements*)
- the Segments (or *approaches*)
- the North and South flows on each approach

My personal recommendation would be to create an exhaustive layer of all of the possible intersections, turning movements, approaches and directional (N-S) flow arrows for the entire road network[224]. This needs to be done only once and requires updating only when a new road/intersection is created. Once everything is drawn and labeled, the data that are produced either with ATRs or with TMCs, either internally by city workers, or externally by developers' consultants can all be connected to these spatial representations that can in turn be immediately useful to city managers and traffic engineers.

In the traffic databases that are associated with these GIS elements, each record will typically represent the counts within a specific 15 minute interval on a specific date across a specific maneuver at a specific intersection. Additionally, *n* fields of each record could represent the counts for *n* types of vehicles (autos, small trucks/SUVs, large trucks, mopeds, motorcycles, bicycles, etc.). An alternative, simpler and more flexible

222 Butler and Dueker (2001) present an in-depth approach to the creation of such an infrastructure for automotive traffic.

223 See www.census.gov.

224 Butler and Duecker (2201) are of similar persuasion, though their GIS-T is more geared toward geometric modeling than actual management of traffic data.

format for the traffic data would only comprise the date, time, location, maneuver, vehicle type, and count, so that there would be more than one record per 15' interval, namely one record for each vehicle type that was counted in that period. The database structure in this latter example would be the same regardless of the number of vehicle types one decided to monitor, thus making it truly universal.

lessons from our boat traffic studies

As already discussed in previous sections, dynamic process require periodic updates. However, unlike gradual natural processes – like sedimentation – or cyclically predictable phenomena – like hydrodynamics – traffic contains a more difficult variable that makes it harder to model it reliably over time: the human factor. Traffic is a human activity and as such it is subject to the vagaries of human volition and behavior. Modeling traffic is useful for short-term predictions – *which way will traffic go if we close this canal?* – or simulations – *how much traffic will we get if we change this canal to one-way only?* – yet the longer term can only be vaguely guessed by extrapolating trends that may or may not reflect future motivations for individual or group choices and the attendant consequences on transportation of people and goods.

[the "human factor"]

Because of the fickle nature of traffic, we focused on developing a solid traffic-counting methodology and we also worked on the development of automatic traffic-counting devices. So lesson one in this section is that if the data are deemed to be truly useful and if the dynamic phenomenon lends itself to automation, then automatic devices should be employed to collect the needed data because of the cost-savings.

[rigorous data-collection methods]

[automatic instruments]

If automation is not possible, then a rigorous methodology for data collection and a strict schedule of updates ought to be put in place, as long as there is a clear commitment to acquiring these data. Lesson two is that these "manual" updates can be farmed out to third parties, by modifying existing regulations or administrative processes, so that they happen more or less free-of-charge[225].

[scheduled updates]

[free updates]

Although activities generally imply a dynamic "movement" through space, it is indeed quite possible to harness traffic data into a spatial geo-scaffolding, as was shown. By exhaustively representing and coding road intersections, as well as turning movement arrows and bi-directional flow vectors, each traffic datum can be made to fall into its proper spatial pigeonhole[226]. The generic lesson is to look for spatial "hooks" onto which dynamic data can be geocoded. This geospatial infrastructure for activities also makes the collection of "free" data updates that much easier, as will be demonstrated later[227].

[spatial geo-scaffolding for activities]

225 See also Note 317 and Gage *et al.*, 2003.

226 See, for example, Butler and Dueker, 2001.

227 See Note 225.

Environmental Conditions

The *Venice Inner Canals* project included the analysis of environmental conditions in the waterways of Venice, wherein, after all, most of the urban sewage is still unloaded to this day. Due to the complexity of some of the bio-chemical analyses that are needed to provide an accurate assessment of the environmental quality of the waters, our role in this aspect of the project was limited to the mapping of point-source discharge locations and to the collection of samples for other scientific teams that participated in the UNESCO-MURST project.

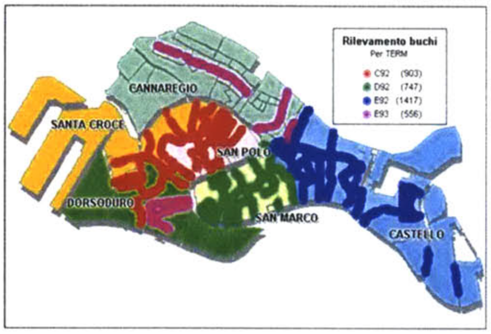

sewer outlets

An important informational foundation to these studies was the identification and inventory of all sewer outlets through which domestic sewage is still dumped daily into the inner canals so that it can be carried out to sea by the circadian tidal cycles. Our inventory process entailed visually inspecting both sides of each canal segment at low tide and mapping and photographing each outlet[228], while recording some key pieces of information, like its condition and activity level. The cyclical nature of the astronomical tides dictated that our activities had to take place in the wee hours of the morning, from 1am to about 8am for the most part, since that's when the most "usable" low tides occur during the spring and summer months when our students were operating in Venice. Collection during the day was made impossible by the constant boat traffic in the canals and only on some occasions were late afternoon and evening windows of opportunity available.

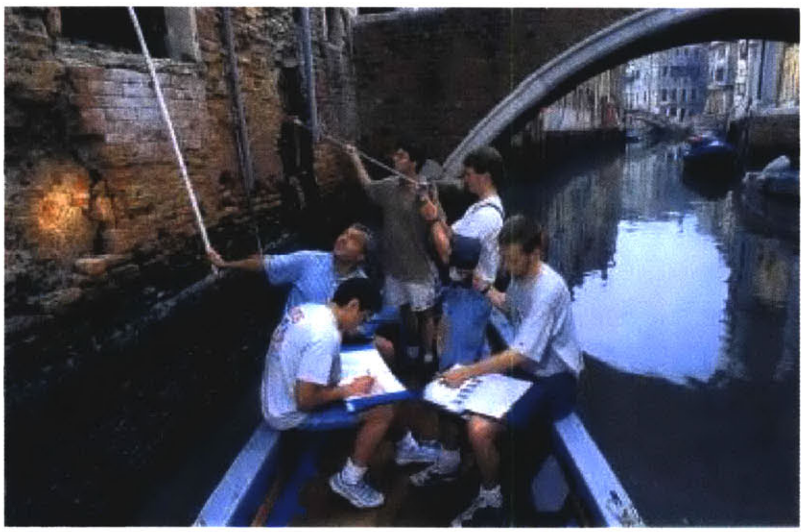

[row boats and gondoliers]

We needed to work during extreme low tides so that the lowest of the sewer outlets would be revealed, but at the same time we also needed enough water in the canal to allow our row boat to float[229]. We used a row boat to carry out our data collection so that we could reduce our draft to the bare minimum, but also because we wanted to conduct our measurements as quietly as possible, given the ungodly hours during which we were forced to operate. I personally manned the oar on many of these nightly excursions, but we also were helped by a group of Venetian gondoliers who were concerned with the deteriorating conditions of the canal walls. In fact, the inventory of sewer

228 The photo on this page, depicting me with four WPI students was published in the August 1999 issue of *National Geographic Magazine*, p. 100 (reproduced with permission).

229 In those years (early 1990's) the amount of sediment that had accumulated in the canals during over 30 years of neglect made it quite possible for a canal to be completely dry during the lowest tides, so that not all low tides were really usable for our studies.

outlets was conducted together with an inventory of wall damage along the same canals.

photometrics

Each photo that we took included in its frame a scale bar, which we used to subsequently measure the dimensions of the outlets. Photometric techniques were employed for this purpose, based on the knowledge of the length of the yardstick included in the picture (50cm in our case).

mapping sewer outlets

Keeping track of these point-sources on a GIS was complicated by the dearth, in the existing basemaps, of sufficiently recognizable features along the canal walls at a scale that would permit a precise pinpointing of the outlets' locations. Nevertheless, we were able to create the first detailed maps of these openings, which later led to several follow-up studies of great importance for the City. The mapping process began in the field, where paper maps were used as initial records of the identified locations. Each outflow was labeled on the map and a corresponding record on a field form was tagged with the same label. On the same field form we also recorded additional information about the outlet, such as its shape, its condition (damaged or undamaged), and its level of activity (active or inactive).

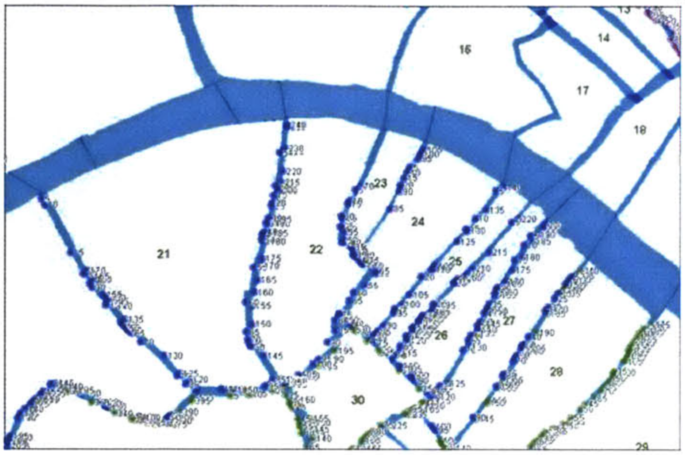

coding sewer outlets

Each outflow was uniquely tagged with a composite numeric code that included the island number and a sequential enumerator to distinguish among the outlets found along the perimeter of each island. These sequences started from the northernmost discharge location, which was number 05, then proceeded clockwise along the perimeter of each island, in intervals of 5, so that the next outflows were numbered 10, 15, 20 etc. These gaps of 5 units in the sequence of codes allows for the addition of intermediate outlets at a later date, as necessary, while preserving a modicum of order to the system.

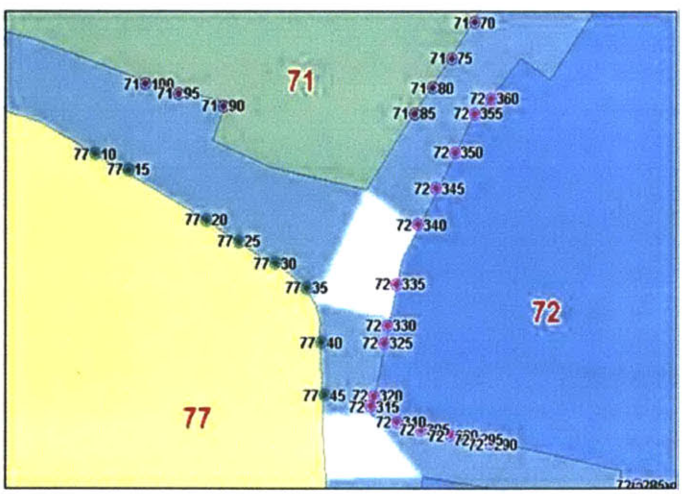

In all, we recoded and validated a total of 2,725 sewage release points in all of Venice's *sestieri*[230]. Though thousands more had been inventoried, they were never confirmed by subsequent inspections. Although we never attained 100% coverage of all canals, this study still represents one of the most complete reconnaissance of sewer outlets ever conducted in Venice.

public health

Knowing where all the sewage is being unloaded has great implications for studies of the effects of canal waters on public health. In fact, one must remember that every day there are hundreds of people who

---

230 Gallo, 1999.

come in contact with these waters in the course of their jobs. Therefore, one of the questions that the UNESCO project addressed was whether there was any cause for concern for those of us who come into contact with these waters on a regular basis.

environmental health

A complementary question was whether – or rather to what degree – the environmental health of the canal network was compromised and what effect the urban discharges have on the surrounding lagoon ecosystem.

To answer both of these questions, several teams of scientists from a variety of scientific and health organizations, local and international, participated in extensive campaigns to ascertain the health of the Venice inner canals. Each campaign was able to furnish insights into the level of pollution in the canals' waters and in the bottom sediments. The main concern was organic pollution deriving from the discharge of excrements from residences, offices and hotels. The microbiology of these discharges was studied[231] to look for pathogens, viruses and bacteria, that could cause harm to humans. Indeed some highly dangerous pathogens were found, though their ability to spread disease was called into question, due to their inability to survive for long in the hostile saline environment of the canals.

[organic pollution]

[pathogens]

Most of the organic pollution that could be harmful to humans is frequently beneficial to plant and animal life in water. Fecal excrements are used as fertilizers everywhere in the world, so many of the Venetian discharges are playing the role of manure for life forms in the canal network. More problematic for the environment are chemical pollutants, which were also studied in the course of the UNESCO-MURST project[232].

[fertilizing]

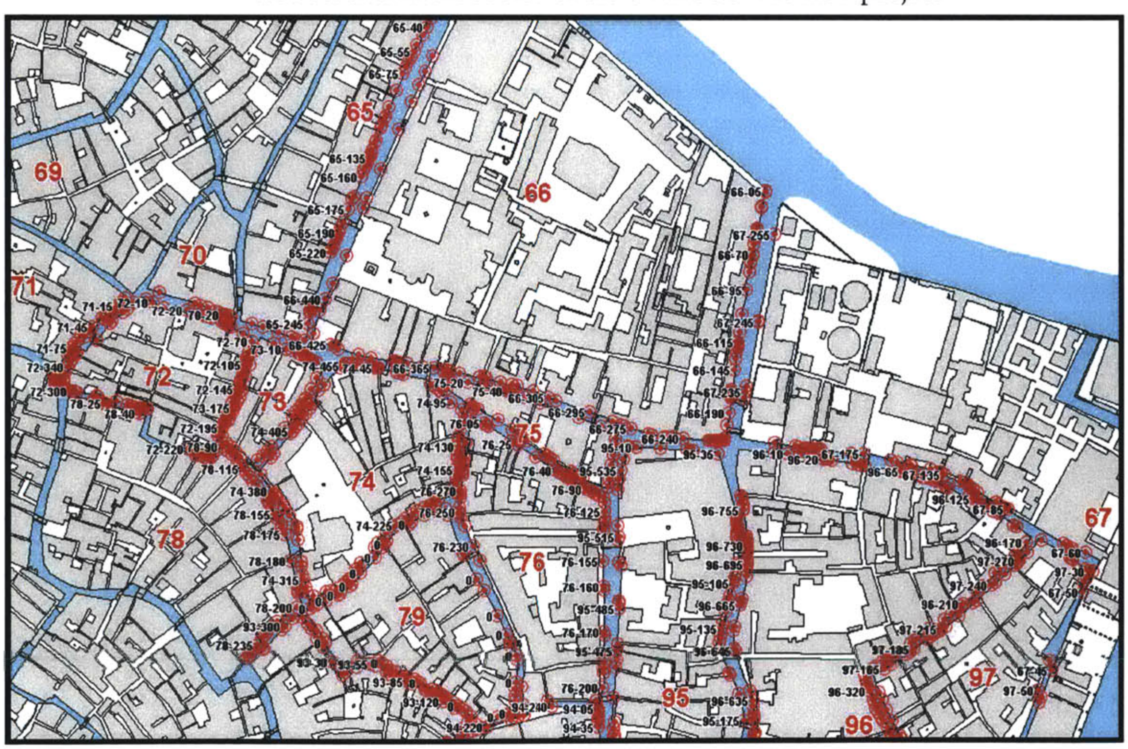

231 Aimo *et al.*, 1999.

water quality

In all, in addition to the high presence of nutrients (nitrogen and phosphorus primarily), the studies found that the city waters contained chemical pollutants in concentrations that were two orders of magnitude greater than the levels found in the lone lagoon sampling site, although the exact effect of this pollution on the lagoon ecosystem was not ascertained, though we do know that the city only contributes 4% of the nitrogen and 15% of the phosphorous that ends up in the lagoon[233].

[salinity]

Meanwhile, scientists have been using an interesting proxy measurement to ascertain the organic pollution in the canals due to sewer discharges. Salinity, as it turns out, is lowered by the freshwater that is flushed down toilets and drains, so measuring the salinity level is tantamount to measuring the concentration of sewage in a canal.

[water consumption]

Along the same lines of thinking, knowing how much water is consumed in a household allows us to estimate how much organic pollution is released by that household into one of the canals near the house. We used this simple proxy as a way to estimate the total organic pollution released into the *rii* by the homes and businesses on each island. This in turn allowed us to estimate the solid waste that would be contained in the sewage, which lead to our first estimate of the total sediment being deposited in the canals by the precipitation of solids in the waste stream.

sedimentation

Sedimentation is a never ending process which leads to the periodic need for dredging to remove the excess mud that accumulates on the bottom of canals. The sediment does not just precipitate to the bottom as soon as it is unloaded into the canal, but rather it floats to the bottom with a certain rate of precipitation which depends on what the suspended solids are made of. If the canal has strong currents, a lot of the solids will be carried elsewhere before they have a chance to settle on the bottom. Boat propellers will also re-suspend the sediment even after it has finally settled down. However, due to the topology of the network, most canals are just as likely to receive stuff from adjoining ones as they are to export their own stuff onto others. The bottom line is that much of the sediment gets trapped in the tortuous waterways and very little floats of it out to sea with the liquid waste.

[sediment transport]

232 *Idem.*

233 *Il Gazzettino*, October 15, 2003, p. III.

[sediment sources]

A specific study we conducted in 1999 explored the sediment sources that contribute to the sedimentation in Venice[234]. We estimated that the majority of sediment actually comes in from the lagoon and gets trapped during the daily tide cycles in the sieve that is the inner canals network. According to our preliminary study, more than 80% of the deposit originates this way. About two thirds of the remaining muck (i.e. 11% of the total found in the canals) derive from the precipitation of solids in the sewage inflow. The last 6% of the sediment was estimated to be generated by the slow crumbling of plasters and by other debris originated by urban materials.

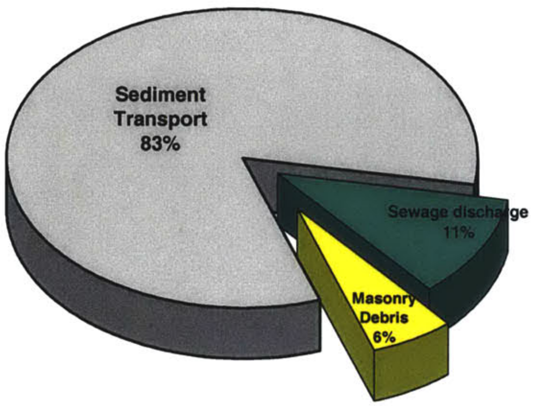

outlet clogging

Over time, the accumulation of residue on the bottom of canals, if neglected as it was from the 1960's until the 1990's, leads to the clogging of sewer outlets – traditionally placed at 4 Venetian feet (approx. 1.2 meters) below mean sea level, to ensure that they would always be submerged. The congestion of the outlets caused internal raptures within the old sewer pipes and forced the relocation of the outlets to a higher elevation, to avoid blockage. The height of the sewage outlets is therefore an important parameter to keep track of since it could (or even should) have an impact on the scheduling of periodic dredging maintenance[235].

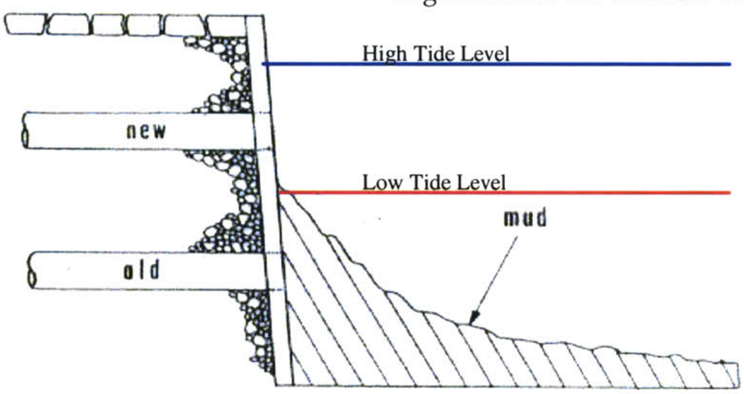

sediment quality

All of the aforementioned UNESCO studies considered both pollutants floating in the canal waters, as well as those immobilized in the canal sediments and slowly released into the system. Some of the worse chemical pollution from heavy metals dates back to the 1960's and 70's when the factories in nearby Marghera were releasing huge quantities of toxic materials into the lagoon without any significant public oversight. This negative environmental legacy is still "archived" in the deeper sediment layers, which trapped some of the worst heavy metals (like mercury, cadmium and copper) for posterity.

current situation

The current situation in Venice has greatly improved thanks in large part to the stricter controls on the industries in Marghera and also in part thanks to the installation of septic tanks for 1200 hotels, 450 businesses, and 2700 residences[236]. Nevertheless, the lack of a "modern" sewer system in

234 Borrelli *et al.*, 1999.

235 See also page 115 and following for more details about the clogging process and it effect on canal wall damage.

236 *Idem.*

the historical center of the city of Venice remains a source of pollution, and thus a cause of concern as well as embarrassment for Venetians.

vacuum sewers for Venice

To address this major shortcoming in Venice, we proposed a vacuum sewer system in 1997, with an award-winning project that exploited a US patent held by a WPI alumnus[237]. Several proposals for the first implementation of a HI-FLO™ system were prepared since 1997, the latest of which is being considered for the collection of sewage from the neighborhood of S.Elena in Venice. To some degree, this is also an example of *plan-demanding* knowledge because the preponderance of data at our disposal suggested this project to our sponsor and to us.

sons from our environmental studies

[urban environmental information]

From the point of view of City Knowledge, the collection and organization of urban environmental information can once again be separated into two distinct aspects: the inventory of permanent contributors to environmental degradation and the monitoring of dynamic conditions that measure the health of the environment.

[mapping point sources]

We have seen how it is possible, albeit arduous, to identify and map all of the point source sewage outflows that are the gateways through which urban refuse flows into the canals, negatively affecting the water quality in Venice. So the lesson again[238] is to look for the "permanent" spatial objects that relate to the dynamic processes (in this case environmental pollution), so that space can once again act as the "glue" to keep information organized.

[proxy measurements]

We have also seen how proxy measurements of salinity – a much easier and cheaper measurement than biochemical analyses – can be used to measure the presence of fresh water flushed down from municipal drains. The lesson is therefore to consider proxy measures that adequately reflect

237 Felices *et al.*, 1997. The alumnus' name is Alan (Chip) Hassett.

238 Similar to the last lesson in the traffic section, on page 88.

the phenomenon we want to monitor, whenever the "ideal" measurements are impractical or too expensive to conduct on a scheduled basis.

[administrative data]

Finally, we have seen how administrative information – namely the water consumed by residents and businesses – can also be used to determine the total amount of sewage discharged, which in turn can lead to the estimation of the quantities of "black sewage" and thus allow us to gauge the solid waste that will be released into the water to contribute to the sedimentation problem. This lesson extends the last one in that it suggests that some of the proxy data that we could use for our urban maintenance, management or planning may be already collected by some public agency (or even by some private company) and thus could be tapped into fairly easily and reliably, as long as some agreement was made with the owners of these existing administrative data sources[239].

Instead of continuing with the plethora of canal examples that I have in store, we will now switch to a more "derivative" application of city knowledge in Venice – one related to Venice's most publicized ailment: *acqua alta* (high water) and its effects on the stones of Venice.

[239] Singh, 2004.

PROTECTING VENICE FROM THE TIDES[240]

flooding in Venice

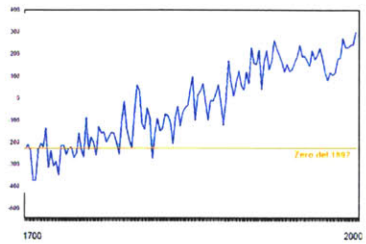

[urban features affected by flooding]

This chapter looks at a particular set of issues that are important for urban maintenance, management and planning in the city of Venice, namely the effects that high tides and flooding have on the physical upkeep of the city[241]. It provides a very concrete set of examples of the benefits of City Knowledge and begins to illuminate the difference between ad-hoc data collection and a systematic approach to information accrual, which leads to returns that are unforeseen, yet very tangible, as the following sections demonstrate.

Flooding affects the physical infrastructure in Venice in a variety of ways, most of which have been recognized since the early days of the city's existence. Manuscripts in the Venice archives chronicle the frequent requests for maintenance along the inner canals of the city ever since records were kept[242]. Natural erosion due to fluctuating tidal water levels is a fact of life in a place like Venice[243]. Construction materials are gradually weakened by the constant wet-dry cycles and by the natural salts and unnatural pollutants contained in the tidal waters.

A simplified taxonomy of the primary elements of the city that are negatively affected by exposure to the salt water of the lagoon and canals includes:

- Above-ground Infrastructure
    - Streets and Pavements
    - Bridges
    - Docks
- Underground Infrastructure
    - Utilities
    - Sewers
- Private and Public Buildings
- Artistic and Architectural Heritage
    - Churches
    - Palaces
    - Convents
    - Public Art
    - Street Furniture
- Canal Walls

While there is no doubt that these elements of the built environment are to some degree harmfully impacted by the gradual, but

240 This chapter is adapted from a refereed paper entitled *City Knowledge as key to understanding the relation between waters and stones in Venice* that was recently accepted for publication (Carrera, 2004).

241 The choice of focus was not mine, but that of the organizers of the Cambridge conferences on Venice in September 2002 and 2003 (Cambridge University, 2003).

242 Caniato, 1999; Dorigo, 1999; Piasentini, 1999.

243 So much so that we very recently collaborated on an interdisciplinary research project to explore the fluctuations of sea levels over the centuries, using archeological, biological, artistic (Canaletto paintings)and semiotic indicators (Angelini *et al.*, 2004). The pictures on this page come from that study.

incessant, assault of the waters of the lagoon, the precise extent of the overall impact of flooding on the state of conservation of today's architectural and urban structures is hard to measure and the costs nearly impossible to quantify.

What is clear is that there are several concurrent factors at play in the undermining of Venice's built environment. Perhaps flooding is not the most destructive of all of the forces participating in the constant interplay between the liquid and solid components of the city. Many see "*Moto Ondoso*" (motor boat wake) as a major player in this arena[244]. Another potential – though perhaps unexpected – culprit is sedimentation, which is accused of engendering damage on canal walls through the clogging of underwater sewer outlets, leading to underground ruptures and thus to seepage and weakening of the mortars that bond together the bricks and stones of the canal walls[245].

"moto ondoso" in Venice

the case for City Knowledge

Without aspiring to actually quantify the damage done to these elements of the urban landscape, this chapter will instead use these pressing issues to delineate a path that would make these and other analyses much more feasible in the long run. In this manner, this chapter will attempt to demonstrate that Venice, like any other city in the world, would benefit from espousing the systematic approach I call City Knowledge.

The complexity of the relationship between the stones of Venice and the waters of the Lagoon is a great demonstration of the types of issues that could be better understood once a plan-ready City Knowledge infrastructure becomes available. The rest of this chapter briefly explores the nature of the interactions between natural and human-caused phenomena vis-à-vis the maintenance of the aforementioned elements of the Venetian built environment. The entire discussion will then focus on the indispensable foundation of information that is needed whenever we try to explore complicated relationships such as the one between stones and water in Venice. This chapter is therefore designed to highlight the knowledge infrastructure that can support such complex analyses.

As mentioned, since 1988, the Venice Project Center of the Worcester Polytechnic Institute (WPI)[246] has been at the forefront of the exploration of the causal relations that cumulatively produce the physical damage that is visible everywhere in Venice. In collaboration with UNESCO and other agencies, the Center has systematically collected a wealth of information about the phenomena connected with architectural damage and decay, both along the canals and elsewhere. Though none of these data connect flooding to structural damage *per se*, numerous correlations were tested out and verified, by relating a variety of independent and dependent variables that link the "waters" with the "stones". These include: traffic levels, boat wake-loading, sedimentation, hydrodynamics, construction materials and maintenance.

244 Pulliero, 1987; Mencini, 2000, Carrera, 2001.

245 Carrera, 1994. See also previous chapter, page 93.

246 See http://www.wpi.edu.

[no "true" City Knowledge system]

The numerous databases and Geographic Information System layers that have been created since 1988 make it possible to test many assertions and draw useful conclusions, but, despite their sophistication, they only hint at what could be possible if a true City Knowledge infrastructure were created and maintained in Venice. This chapter now sketches out the type of information that one would need to have accrued in order to measure the "before and after" of flooding situations and thus begin to hypothesize about the "cause and effect" of phenomena that can – at best – be treated as "natural experiments" over which there is very little design control.

DAMAGE TO PUBLIC INFRASTRUCTURE

If one considers the "public domain" only, then the above-ground infrastructure in Venice that is affected by flooding can be succinctly defined as: streets, bridges and docks. Moreover, *Moto Ondoso* (motorboat wakes) also damages all three elements, although streets are only marginally affected. The causal nexus between flooding, *moto ondoso*, and the corresponding damage to these artifacts is extremely complex to isolate, but one can start by simply knowing as much as possible about these three elements of the public realm that are subjected to the destructive force of water. The Venice Project Center has teamed up with a Venetian company called *Forma Urbis*[247] to complete: (1) the mapping of over 1 million square meters of Venetian street pavements for Insula (top left); (2) the inventory of all of the 472 bridges in the city also for Insula[248] (center left); and (3) the catalogue of all 1321 public docks in the city for the city's department of public services[249] (bottom left).

These three major physical inventories create the backbone of a possible study of the effects that tides and *moto ondoso* have on them. By adding the digital elevation maps[250] that the *Magistrato alle Acque* has developed, through the *Consorzio Venezia Nuova*[251] (1988), it is possible to produce maps of the flooding of the streets of Venice at a succession of tide levels (bottom right).

247 This is the for-profit company that I founded in 1997 to carry out some of these City Knowledge activities in Venice. See more at page 189.

248 After an initial project by Bahn et al. 1998.

249 Felices et al., 1994; Doherty et al., 1995.

250 All elevations in Venice are from the "mareographic zero" of 1897, a.k.a. the "Punta Salute" datum.

251 This is the consortium of construction companies that is in charge of the "floodgate" project.

Knowing how frequently these various tides occur leads to the identification of what portions of the city's streets, bridges and docks get wet and how often. So, with only a modicum of approximation and extrapolation, it would be possible to arrive at reasonable estimates of the damage that can be caused by mere flooding.

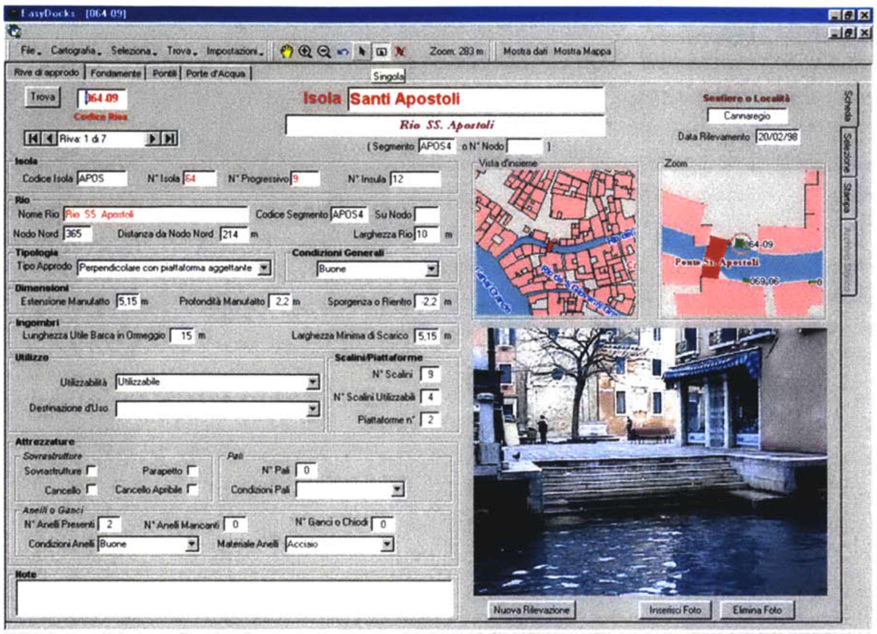

It must be noted that these data-gathering efforts have not been aimed at simply producing these maps, but also entailed painstaking surveys of the artifacts inventoried. Each of the Forma Urbis catalogues has resulted in an information system application that brings together the permanent (left) and ephemeral (bottom) characteristics of both the bridges (bottom) and the docks (left). In fact, whereas the state of conservation of an artifact will change and thus will need to be collected repeatedly over time (allowing the measurement of "damage"), the permanent aspects will never have to be recorded again, which is one of the advantages of a City Knowledge system.

These user-friendly, multimedia information systems make it possible for the dynamic data to be maintained up-to-date through the creation of a new, time-stamped, state-of-conservation assessment every time an intervention (or mere inspection) takes place. One can begin to see how useful such "plan-ready" information would be if it were accessible to scientists and decision makers as well as to the frontline users who are in charge of the upkeep of the physical artifacts.

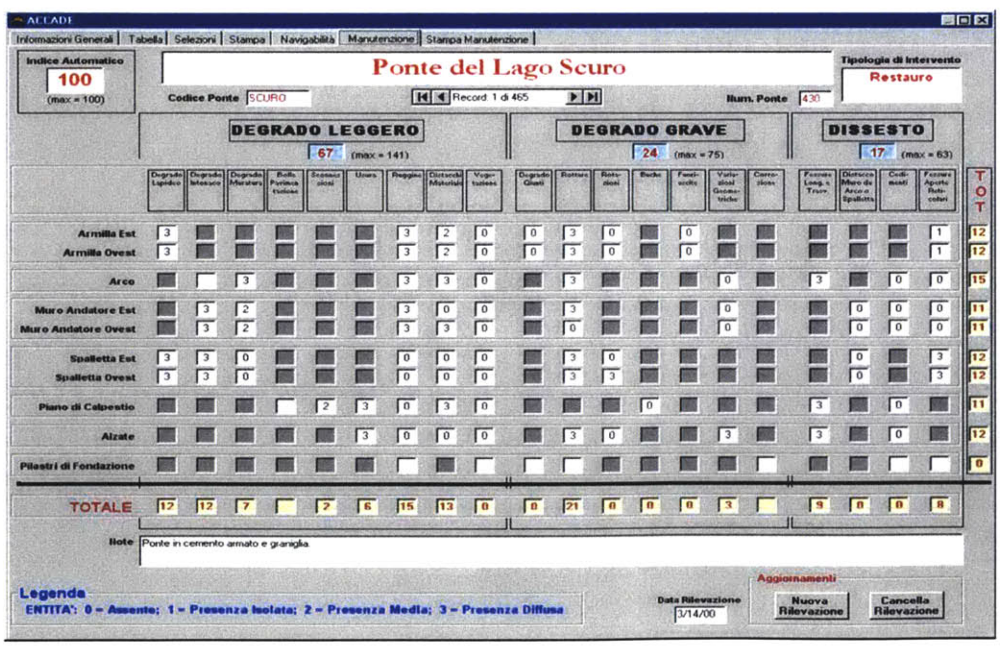

This section shows how plan-

lessons from infrastructure analysis

[second-order re-use]

ready information which is already useful in and of itself for the maintenance of various infrastructure elements can also become useful for second-order analyses such as this one about flood damage. The "emergent qualities" of City Knowledge are starting to emerge.

[separating permanent and ephemeral]

The section also shows how useful it is to separate the "permanent" aspects of our maintainable municipal assets from the "ephemeral" ones, such as condition assessments. In fact, the condition assessments and the log of physical maintenance interventions to ameliorate the assets' conditions can be themselves useful to measure damage over time and possibly correlate it with potential causes, such as flooding. A true City Knowledge system should always separate the slowly-changing from the fast-changing features of municipal assets.

[include update functions]

Moreover, these systems should also be constructed with provisions for updating of the data at whatever interval is appropriate. In this case, our systems can be kept up-to-date as a consequence of repairs, which entails that some of the conditions will actually improve accordingly. However the same system would also allow periodic check-ups of these structures to detect worsening conditions, as long as scheduled inspections were arranged by those in charge of the upkeep of these structures.

[citizen reports]

Citizen reports and complaints would be greatly facilitated by the existence of a well-defined information infrastructure which attributes clear, unique labels to each object. By exploiting the standardized reference system, these reports could be logged and used as rough "condition reports", thus feeding "free" information into the City Knowledge system[252].

---

252 See also page 186. For an interesting and potentially useful way to use cell phones to tap deeper into the world of citizen reports, see www.yellowarrow.org.

Damage to Underground Infrastructure

The next category of items that can be damaged by water includes objects under ground, namely the infrastructure for electrical, water, gas, telephone, street lighting and sewage services, composed of pipes, ducts, valves, manholes, inspection wells, cables, etc. Insula S.p.A. has worked on these sub-systems since it was originally chartered by the four main utility companies (electricity, phone, gas and water) who own 48% of the shares – the remaining 52% being owned by the City[253].

managing the underground infrastructure

One of the goals of Insula is to eventually produce a maintainable map of the underground infrastructure in Venice. The effects of flooding on subterranean utilities could be much more predictable with this sort of city knowledge at one's fingertips. Work below street level could also be coordinated in such a way as to avoid waste and redundancy. Once in place, these systems would represent practical applications of city knowledge principles. The first major system to showcase these principles was SmartInsula, the pioneering and award-winning application which formed the backbone of Insula's sophisticated information system that was developed by a UNESCO team in 1997[254] and has since evolved way beyond that initial application.

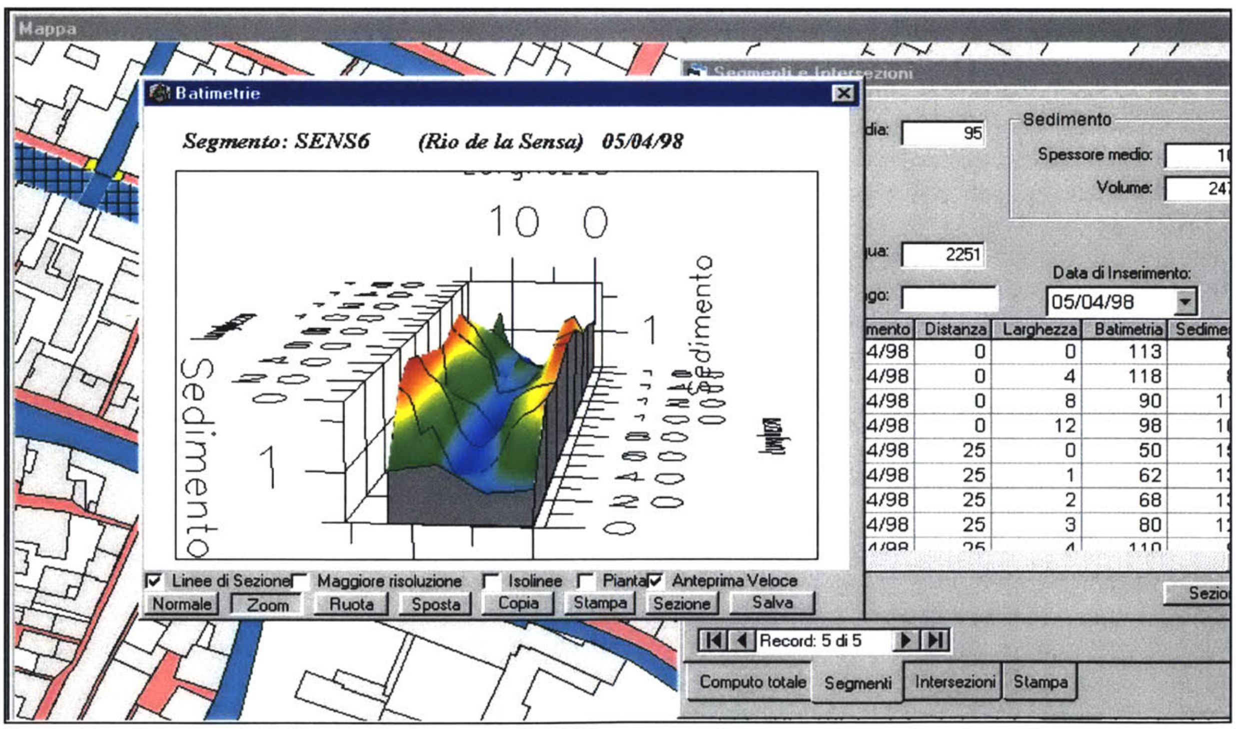

253 As is discussed elsewhere in this dissertation (see fottnote 486), the original charter did not clearly specify whether sewers would be under Insula's jurisdiction, which has led to some friction with the department of public works from which many of the Insula staff were recruited.

254 I was personally involved in the initial design, later carried to completion by Cozzutto, 1997.

updating the information

The challenge that Insula is now experiencing lies in keeping such rich systems up to date. Insula needs to constantly update bathymetric measurements, bridge conditions, work progress and much much more. Truly emergent, self-sustaining, city knowledge systems delegate data entry to the most appropriate external user, or to beneficiaries of the work – in other words to the end users of the system. In fact, the major technical difficulty for Insula so far has been in decentralizing the data entry, while reconciling accounting systems with technical or engineering systems (frequently CAD-based), and integrating them with geographical information systems (now web-GIS). Many difficulties are being eliminated by forcing compliance with a desired file format as part of a contractual stipulation with outside contractors. Internal adherence to this standard tends to be harder to enforce. As the entire GIS system is ported to the internet, web-based applications to assist contractors in submitting the appropriate digital documents and files are beginning to relieve internal staff of data entry tasks[255].

Another web-GIS that deals with the underground and embodies some of these principles has been developed by Dr. Enrico De Polignol and Dr. Lapo Cozzutto for the Environment Department of the City of Venice. The *Sistema Informativo del Suolo* (S.I.S.) was initially dedicated to the self-reporting of core-sample analyses about contaminated sites in Porto Marghera. Private companies are entering all of the data into this system, through a password-protected internet browser and the data are later analyzed and mapped semi-automatically by the system[256]. The system has recently been incorporated into a more ambitious *Sistema Informativo Ambientale* (S.I.A.) which is a web-GIS system that will also include information about electromagnetic pollution and green amenities[257].

lessons from underground data management

Even though we did not collect data about underground utilities ourselves, the paragraphs above have illustrated how custom-made systems to manage urban information of that sort can be created and can be allowed to later "grow" on their own. The difficulty becomes how to keep these systems up-to-date. A true City Knowledge system should not only provide the technical means to capture periodic updates in computer databases (as discussed in the previous section), but should also include logistical and administrative mechanisms to make these updates actually happen. When possible, the task of keeping the information up-to-date should be left to "customers" of the system, through web-based self-reporting mechanisms, as was done in the S.I.S. system described above.

[logistical and administrative updates]

What is still missing, to allow a sustainable use and re-use of these data repositories for more complex, higher-order analytical studies, is a clear definition of "informational jurisdictions" and a mutually beneficial agreement to share the information between different agencies – two basic tenets of the City Knowledge concept.

[informational jurisdictions]

255 Todaro, personal communication.

256 http://www.ambiente.venezia.it/sis/. Last accessed 8/19/04.

257 http://www.ambiente.venezia.it/web_sit.asp. Last accessed 8/19/04.

Damage to Private and Public Buildings

damage to buildings

The next big category of physical objects that could be impacted by frequent floods includes all buildings: public and private. By implication, this category also includes all of the stores, shops, restaurants and all other businesses housed in these buildings. Flooding affects all buildings and businesses in its path. *Moto ondoso*, on the other hand, only affects buildings along the canals. Using the information systems developed, it is possible to know how many buildings are affected by floods at any tide level. For example, during a tide that reaches 130 cm., 9,124 buildings come into direct contact with the *acqua alta* out of a total of 15,486 buildings in the entire city (including the Giudecca). The system does not allow the prediction of whether or not the interior of each individual building actually gets inundated with tide water, although a specific inventory on the *piani terra* was conducted by the city in 1999 to answer just such a question[258]. Insula S.p.A. has been actively working to increase the elevation of public streets to around 1.2 meters – a process called *rialzi* – in order to reduce the number of buildings reached by high tides.

[58% of buildings flood with 130cm tide]

[local anti-flood measures]

The owners of private buildings and commercial establishments affected by *acqua alta* are doing whatever they can to protect their property from floods: using small barriers, impermeable membranes and *vasche*[259] to seal out the water, and even installing sophisticated drainage systems to direct the water to sump pumps that expel it from the inside of the building. Public buildings are similarly protected and access to many of them is guaranteed even during high floods, by the installation of wooden walking planks[260].

[costing out anti-flood measures]

Quantifying the damage that floods do to buildings and stores may be difficult, but the expenditure related to the local prevention of flooding building-by-building and business-by-business should be somewhat easier to calculate, by inventorying and estimating the cost of all the measures that have been put in place to either protect private and public property from floods or to make them accessible during floods. Moreover, in addition to tallying the cost of preventive measures, one could also account for all of the restorations and repairs that were caused by particularly severe floods. It may be rather difficult to do so, however, unless the government was involved in disbursing emergency relief funds for such activities and thus records were kept of the repair costs incurred.

damage to stores

[67% of stores flood with 130 cm tide]

As discussed above, City Knowledge would help us with such difficult estimates, by telling us where all the buildings are with respect to the flood lines. A map of all stores that was produced in 2001[261] shows that 2,862 shops (out of 4,263) would get flooded by a 130 cm tide. Together, these two figures, about the number of flooded buildings and number of

258 COSES, 1999

259 Literally "tubs", which are basically concrete underground chambers.

260 By VESTA, the local public-private Water and Sanitation authority.

261 Duffy, 2001.

flooded shops, represent a necessary place to start in an estimate of flood damage to private and public property.

[anti-flood permits]

Since permits are necessary to install local barriers or to raise the ground floors by adding a step or two to the entrance, an estimate of the overall citywide cost of localized flood prevention measures would be possible if a system was put in place to geographically archive permits that affect the external built environment. The City Knowledge framework recommends that such mechanisms for capturing administrative transactions be put in place to guarantee that the information systems are maintained up-to-date as piecemeal changes to the urban fabric are allowed through the permitting process. Together with Dr. Pypaert (UNESCO) I have been actively promoting such a self-sustaining system, by bringing together data from the former *Assessorato all'Urbanistica*, that keeps track of zoning and land use; the *Legge Speciale* department, that is in charge of disbursing restoration funds, based on specific work estimates; the *Edilizia Privata* department, which administers permits; and the *Soprintendenza*, which updates the "*vincoli*" that restrict modifications to registered historic properties[262]. All these organizations provide information for the benefit of the *Commissione di Salvaguardia*. This institution has the final word on all major modifications to buildings in the historic city and would greatly benefit from such contextual knowledge at its fingertips when making important decisions[263].

coordinating building restoration efforts

Being one of the institutions with a representative in the *Salvaguardia* commission, UNESCO had a vested interest in bringing about such a confluence of information from all of these sources, but so do all Venetians, who are in the end affected by the permanent changes that are approved by this regulatory body.

lessons about building data management

Dealing with this aspect of flooding would have been much easier if the data collected by a myriad organizations were already coordinated and interconnected. If that was the case, we would not only be able to know exactly what effects a flood would have on the *piani terra*, but we could also begin to estimate the cost of local flood protection by tallying up the permits to install miniature flood gates at one's doorstep or to raise the entrance one step higher, etc.

[self-serving interests first]

The first lesson for City Knowledge we glean from this section is that each organization that is in charge of some aspect of building construction, restoration, preservation or permitting should get its house in order first and foremost, as Dr. Pypaert and I have been trying to propose, so far unsuccessfully[264]. If the various departments listed above had operational information systems in place within the boundaries of their respective informational jurisdictions[265], changes to the city of Venice would

262 Halloran *et al.*, 2002.

263 See also Mancuso, 2003.

264 Halloran *et al.*, 2002.

265 Mancuso (2003), using primarily VPC datasets and layers, put together a web-GIS that encompassed all of these jurisdictions, which can be found at http://www.intelligencesoftware.it/unesco/venezia/. Last accessed 8/19/04.

happen more cogently, thanks to the ability of these agencies to deliberate about permit applications while having all of the appropriate contextual information in front of them. By serving their own interests, these organizations would make their own job easier and would provide a better service to the citizens that foot their bills.

[second-order uses emerge later]

[coordination adds value]

The other lesson, which repeats a refrain that already appeared in previous sections is that after the public agencies have independently taken care of their own affairs, they can coordinate with each other to provide value-added benefits to each other and to the city, such as, in this case, the ability to estimate costs incurred for the local protection from high tides.

DAMAGE TO PALACES, CHURCHES AND CONVENTS

Some of the buildings which are part of Venice's more prestigious architectural heritage, such as palaces, churches and convents, are also touched by high tides. Being more important than others, they have received more attention from public authorities and from philanthropic organizations such as the so-called "private committees".

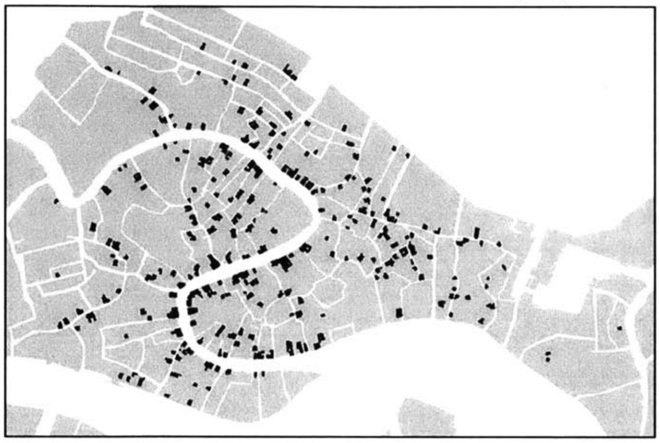

Venice owes much of its fame to its aqueous *forma urbis* and to the art and architecture it contains. In 1987, the whole city was inducted by UNESCO into the list of World Heritage Sites as a treasure that belongs to all humanity; the first city to receive such an honor in its entirety. When it comes to damage due to floods, the parts of Venice's heritage that stand to suffer the most are palaces and churches, which tend to have ground floors containing more elaborate craftsmanship and more precious materials.

cultural heritage catalogs

Right after the 1966 flood, UNESCO funded the creation of catalogues of Venetian Palaces, Churches and Convents. In the three decades since then, these catalogues have proved invaluable as a knowledgebase supporting the relentless efforts for the restoration of the priceless treasures of art and architecture first inventoried in the late 1960's. Starting in 1999, teams of WPI students began the task of revisiting and computerizing the UNESCO catalogs[266]. These efforts allow us to say, for instance, that out of the 383 *palazzi* in Venice, 308 get wet with a 130cm tide (top map at left), as do 46 out of 59 convents (center map).

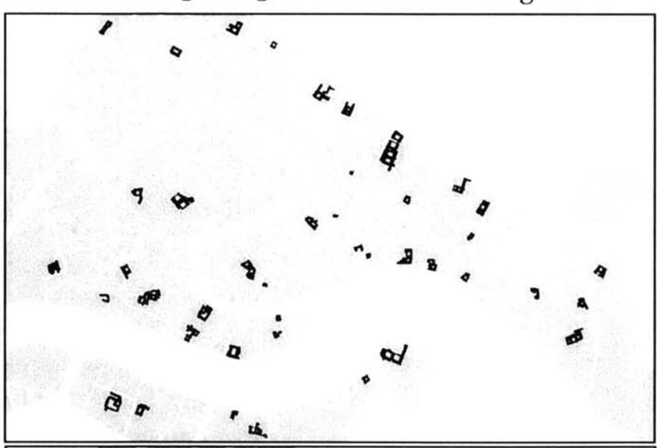

Similarly, we can identify all of the churches that would get flooded with the same 130cm tide. Out of a total of 113 churches in Venice (including Giudecca), 86 are affected by this *acqua alta* (bottom map).

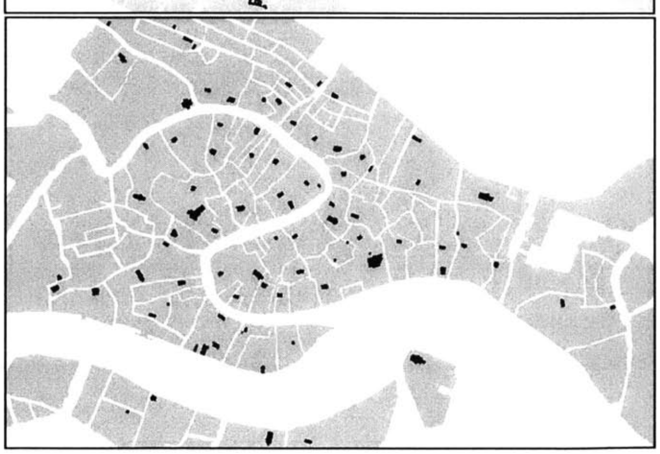

Nevertheless, in order to convert the knowledge of what gets wet at the various tide levels into a more useful estimate of the damage incurred by these artifacts when they get flooded, it is necessary to know a lot more about what's inside these historic buildings.

266 Donnelly et al., 1999; Halloran et al., 2002; Marchetti et al., 2003

church floor artifacts

Starting in 2002, we began recording the tombstones, inscriptions and artifacts that are embedded in church floors[267] (see photo at left). In addition to being frequently at lower elevations due to the age of the original foundation, churches have the added handicap of being vulnerable to flooding through their floors, which are riddled with tombstones. The underground cavities where the entombments took place are a conduit through which high tides can quickly reach the artifacts on the floor's surface. This process is abetted by the high permeability of the bottoms of the tombs, which were constructed in such a way as to purposely allow tide waters into them, so that the mortal remains could be rapidly washed away and the tomb could be promptly recycled and reused.

The photo below clearly shows the huge gaps purposely left between the planks laid at the bottom of a tomb recently excavated under the church of San Samuele[268].

Once this additional city knowledge catalogue of church floor artifacts is finished (over 80% of the surveys have been completed), a more accurate assessment of the potential damage inflicted upon church floors by frequent floods will be possible. Arriving at a similar inventory for the ground floors of palaces and convents would also be useful in this regard.

The overall impact of flooding on churches, convents and palaces can thus include a better estimate of the damage to their floors, but should also include the deleterious effects of salt water on any other artifact that may be touched by tidal waters in the interior of these historic structures. Appropriate monitoring of the decay by the *Curia* and *Soprintendenza* could help maintain these catalogues up-to-date and thus prevent catastrophic damage to these important artifacts.

267 Delaive et al., 2002; Hayes et al., 2003;

268 Courtesy of Dr. Luigi Fozzati, *Soprintendenza Archeologica* and Dr. Marco Bortoletto, archeologist.

lessons from palaces, churches and convents

[truly permanent catalogs]

[never-ending condition assessments]

[be thorough and complete]

[share results to avoid duplication]

This topic does not yield many new lessons, but is more of a "refresher course". It reminded us of the permanent/dynamic dichotomy and it reaffirmed the usefulness and re-usability of systematic catalogs.

Nowhere is the advantage of separating the "permanent" from the "dynamic" more clear than when dealing with cultural heritage. These historic artifacts have been around for centuries and such ancient relics only change by subtraction, not by addition. No new artifacts are ever added, yet some may disappear or dissolve to dust. Cataloguing historic assets is therefore a finite process, but monitoring their state of conservation is not.

The biggest lesson we learnt with these and similar projects is that it really pays to be thorough and complete when cataloguing antique artifacts. In this day and age, such a process could indeed be done once and for all. The key to avoid duplication of effort in the future is to disseminate the results widely and openly. Today's internet technologies can help eliminate redundancy as long as people are amenable to sharing their inventories freely and willingly. The avoidance of duplication should free up some time that can be put instead into periodic checks of the conditions of these objects to detect degenerative processes that may lead to their complete loss.

DAMAGE TO OTHER CULTURAL HERITAGE

After the flood of 1966, most, if not all, precious paintings in Venice have been moved up and out of the reach of even the highest high tides. Practically all damage to heritage would now be limited to fixed and immovable structures, such as floors, bases, pedestals, columns, steps and other artifacts within a 2-meter band from ground level (which translates to over 3 meters above sea level)[269].

endangered artwork

With the exception of our aforementioned recent work on church floors[270], there are no systematic assessments of the artistic or historic heritage contained in this "danger zone" in the entire city. Common sense suggests that everything that could be moved away from this perilous band should have been already moved, though it is quite possible that some artwork might still be in a vulnerable location to this day. Estimating just how many works of art still remain within the "danger band" is arduous at best, whereas a fully-developed (utopian?) City Knowledge system could provide the answer to this enigma in a few seconds. Under such a system, the various authorities in charge of heritage collections (like the municipal, provincial and regional governments, the Curia, the *Soprintendenze*, the *Archivio di Stato* and the two main libraries – Querini and Marciana) would have already catalogued all of the objects that they are respectively in charge of, namely the buildings, properties, church floor artifacts, paintings, mosaics, manuscripts, parchments, books etc..

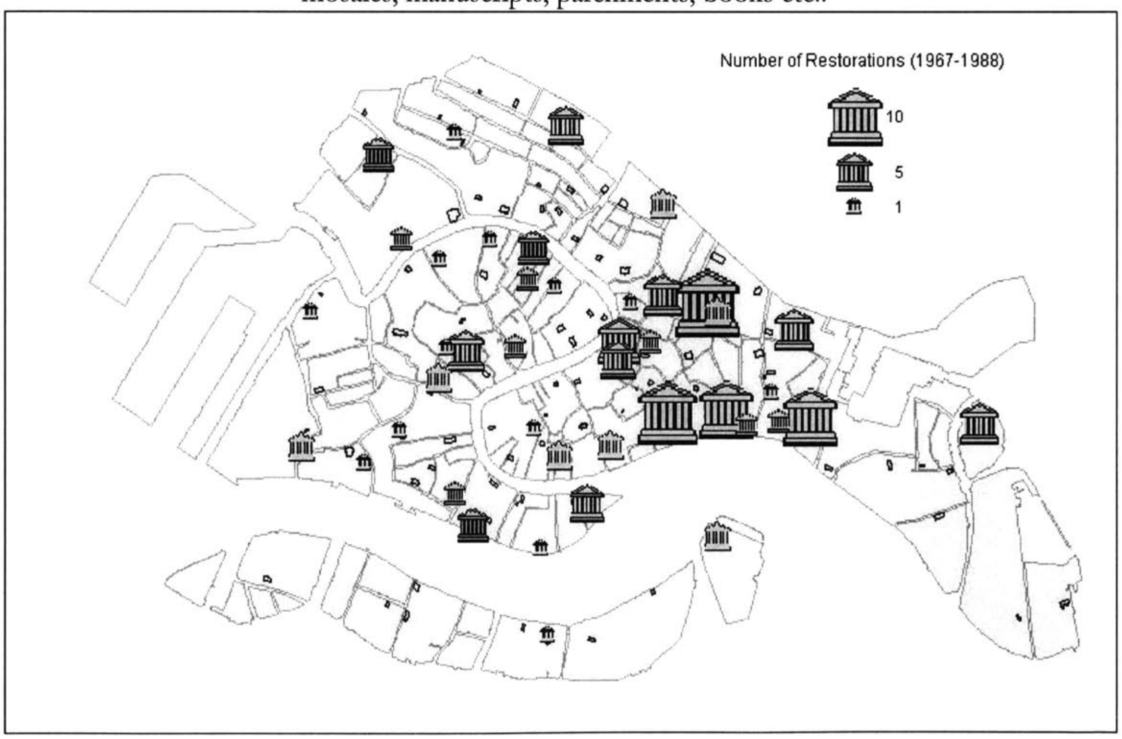

269 This measure has been picked somewhat arbitrarily to reflect the approximate height of the theoretical maximum tidal surge. A one meter threshold would reflect the actual street flooding during the historical maximum level of 1.92, recorded in 1966. The exact dimensions of this "danger zone" are irrelevant to this discussion.

emergent transcriptions

The information contained in the Venice archives and in the historic libraries would be even better protected if electronic transcriptions of the manuscripts were produced using the Emergent Transcription Assistant System[271] (left) and the Ultraviolet Scanner[272] that we have been concurrently developing at WPI[273].
With some foresight, these computerized catalogues could have included a field for the height of the artifact from the floor, which in turn would allow us to simply select all objects whose distance from the floor was less than 2 m.

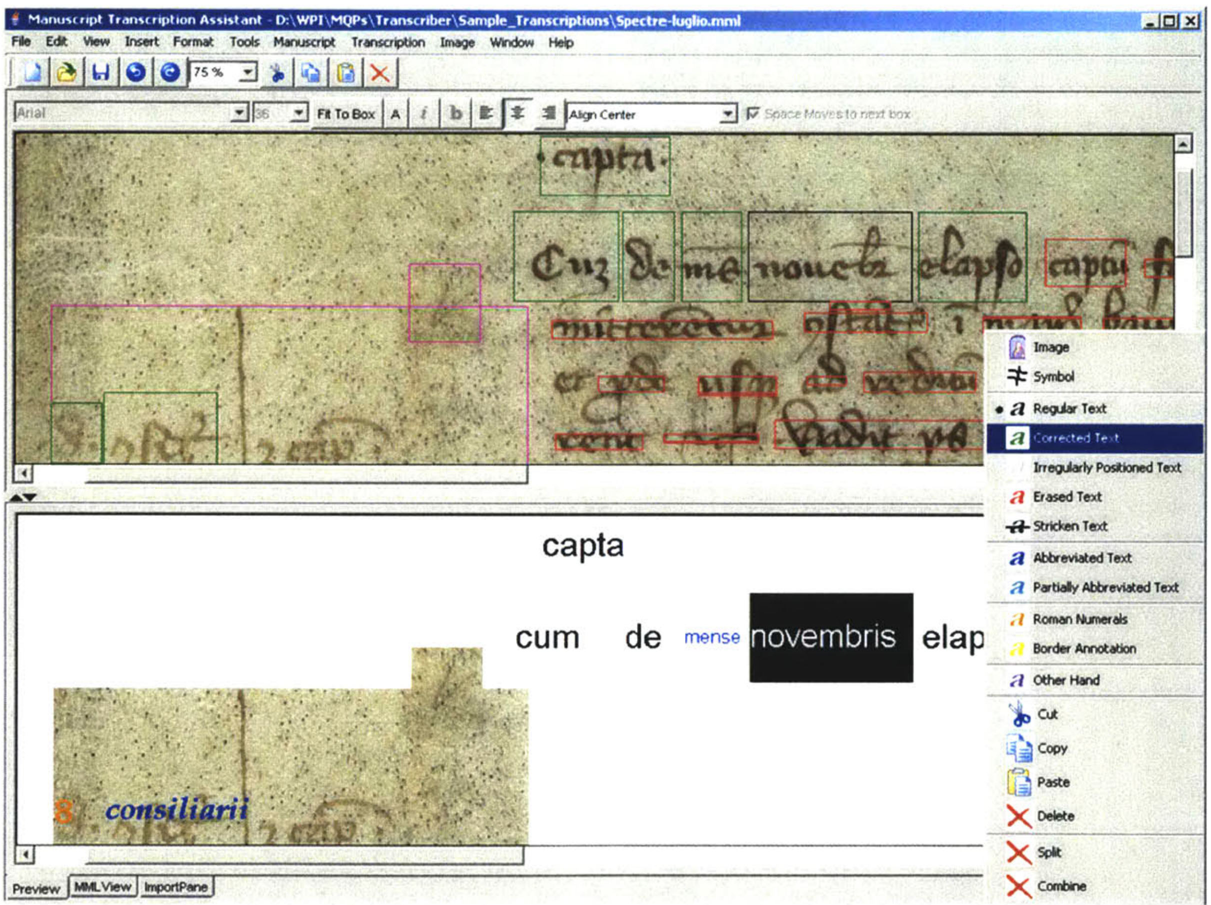

270 Delaive et al., 2002; Hayes et al., 2003.
271 Ho *et al.*, 2003; Calhoun *et al.*, 2004.
272 Dehri *et al.*, 2004.
273 More information at www.wpi.edu/~carrera, under WPI research.

public art

As partial demonstration of the benefits of having city knowledge systems in place, it has been possible to select, from the 2,930 pieces of outdoor sculpture that have been catalogued over the last decade (figure at left), the ones within 2 meters of ground level[274]. This search instantly reports that 69 artworks are on public display at a height of 2m or less.

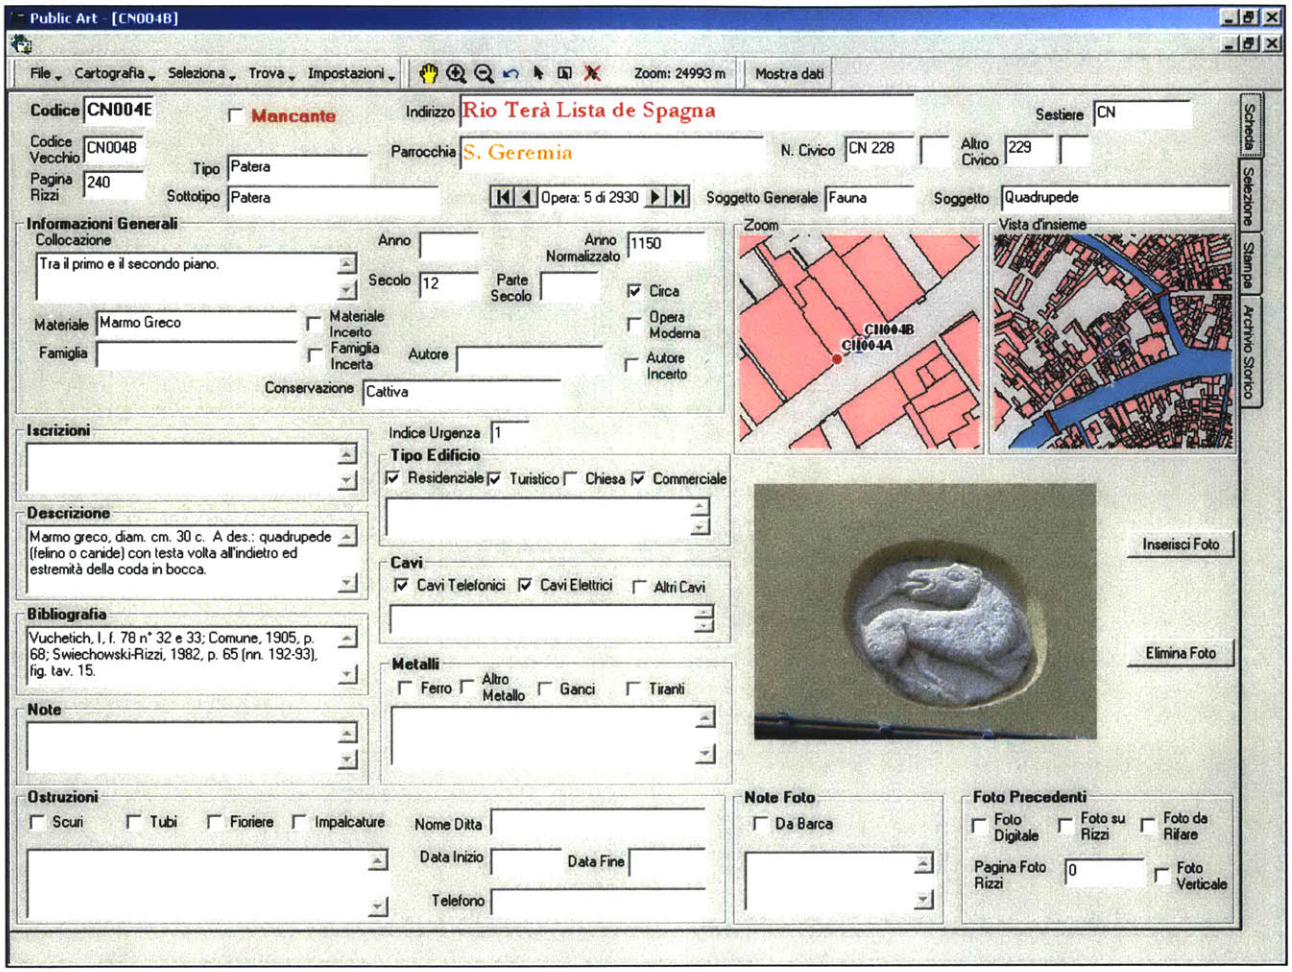

274 The reason for including the height from the ground in the database in the first place was to allow us to calculate the cost of scaffolding. It should be noted that the height from the ground does not correspond to the actual elevation above the standard sea level of 1897.

wellheads and flagpoles

Similarly, one could also include in the 2-metre band all 232 wellheads from the wellhead catalogue (bottom of page), since they all sit at ground level, as do all 22 historic flagpole holders that dot the city (left). More specifically, with the same 130 cm flood used as an example, 122 of the 232 wellheads would get wet (detail below).

Fortunately, though, most of the 4.500+ pieces of public art and street furniture that have been inventoried in the *calli* and *campi* of the city are outside the 2m danger-zone.

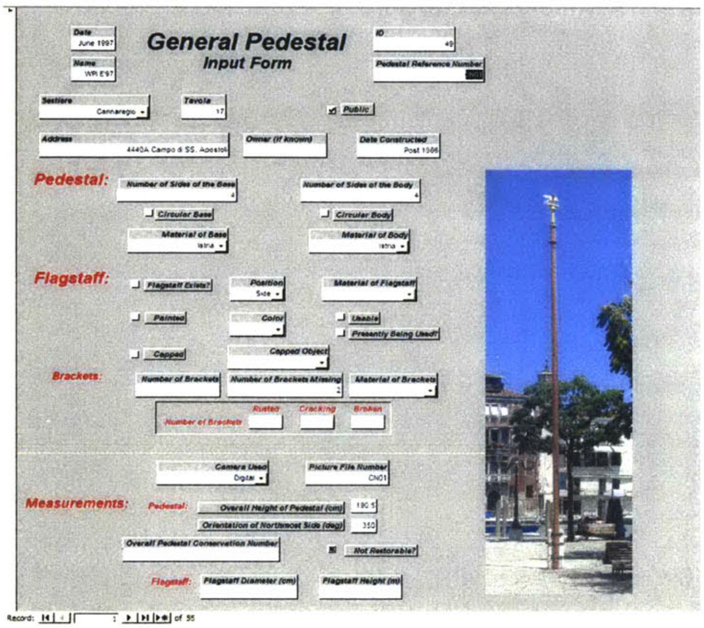

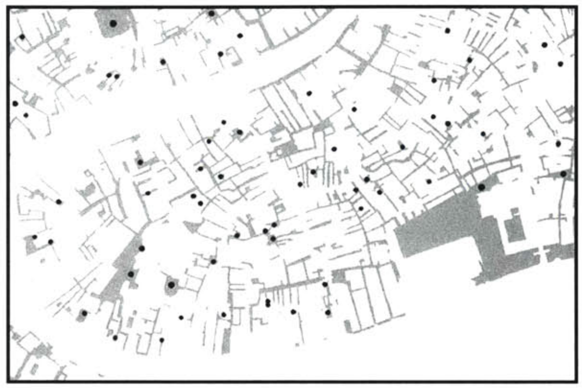

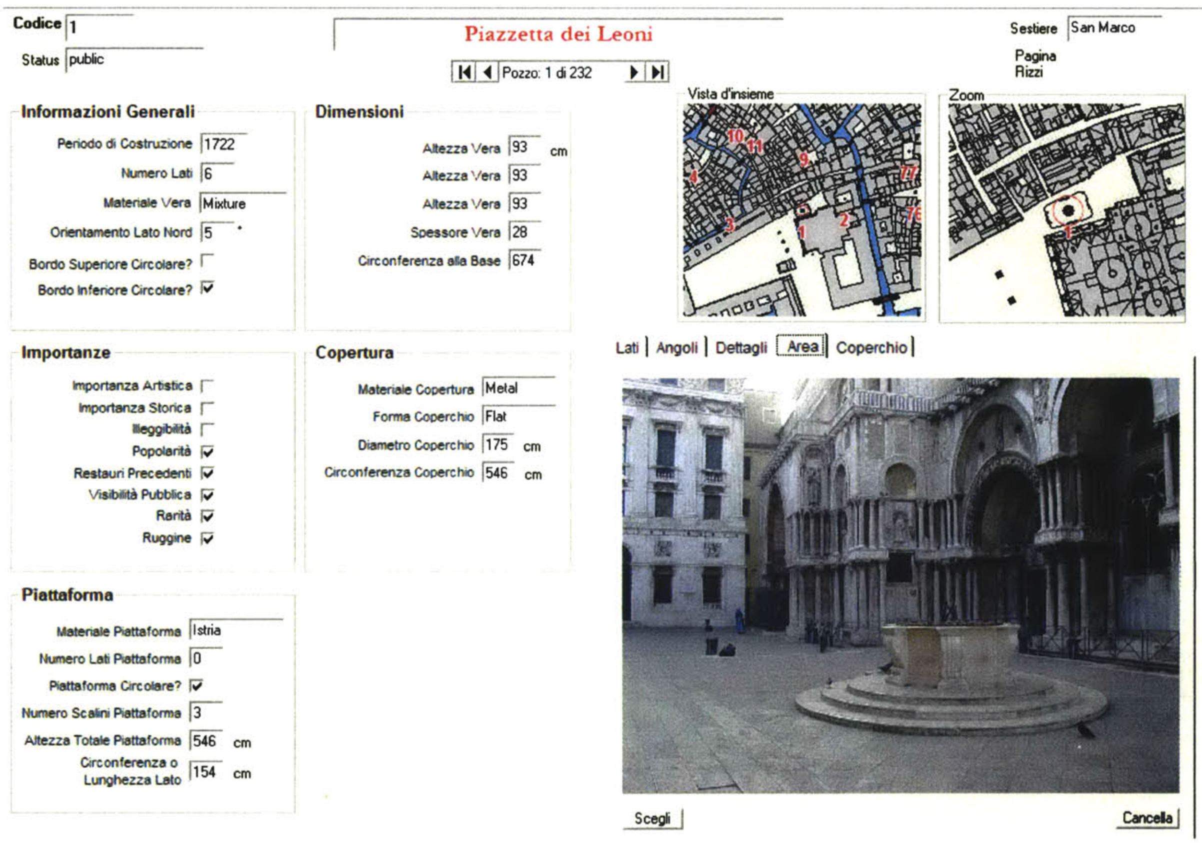

lessons from cultural heritage

[inventory within jurisdiction]

Once again, the availability of a City Knowledge system would make many of these preliminary assessments instantly possible, as long as each institution in charge of artistic treasures had developed its own catalogue according to the aforementioned informational jurisdictions. It is also obvious that a disaster relief agency, such as the Italian *Protezione Civile*, would greatly benefit from a distributed, yet interconnected City Knowledge system that was able to direct emergency crews to the exact locations where works of art were in danger of being flooded during an *acqua alta*[275].

[no centralized know-it-all system]

[emergent distributed network]

[no data management mediators]

The key here is not to focus on creating a centralized know-it-all system, but to foster the emergence of a distributed network of smaller (and more manageable and maintainable) systems, through a process that I chose to call "middle-out[276]", to differentiate it from "top-down" or "bottom-up", both of which have demonstrated severe shortcomings. It may be that the information could be actually housed in a single server somewhere[277], but each different agency will be managing the data updates independently, without an overarching entity mediating the data management.

[minimize in-house data input]

Data input, on the other hand, could be farmed out to a central agency (like an MIS department) to achieve some cost-savings, especially when specific skills are required, such as for GIS mapping. It is important to remember, however, that the goal of an efficient City Knowledge system is to outsource most data input to "customers" or contractors in order to basically get it done for free, or at least to minimize the costs of in-house data input.

---

275 In fact, we began to work with the *Protezione Civile* in the summer of 2004 precisely to prepare for these types of catastrophic events.

276 This term has been used before by many to mean several different things. I personally first heard the phrase from Prof. Joseph Ferreira who used it in a paper (Ferreira, 1998) but in a different context and with a somewhat different meaning.

277 As was the case for Mancuso's (2003) web-GIS.

Damage to Canal Walls

In concluding this brief excursus into the various points of contact between waters and stones in Venice, this section discusses what is perhaps the most critical interface between these two fundamental elements: canal walls.

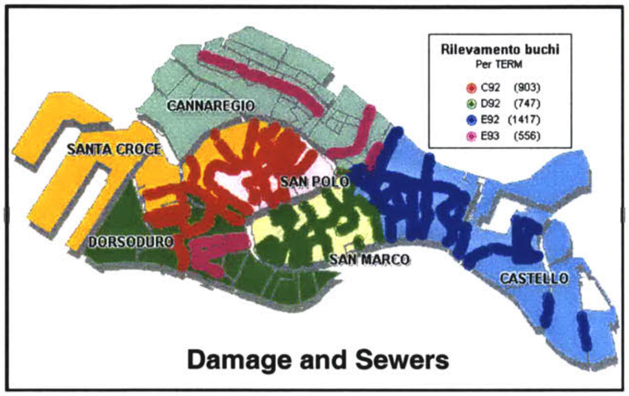

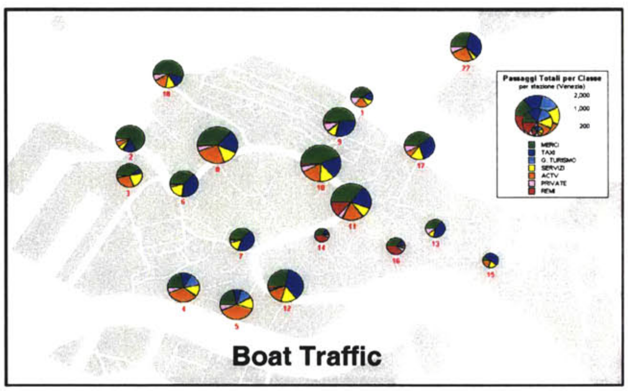

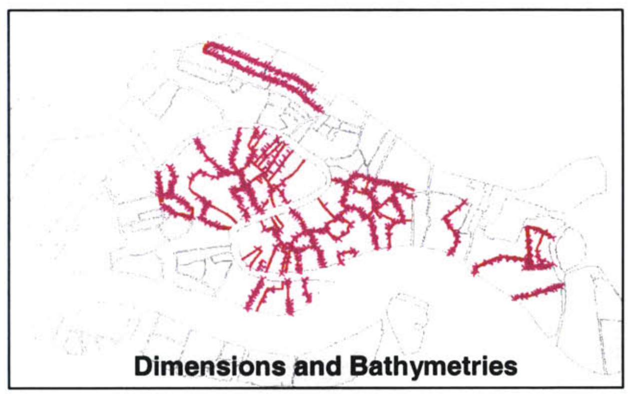

"deep knowledge"

This section illustrates how a systematic and methodical approach to the accumulation of urban information can yield "deep knowledge" about the causal nexus between phenomena.

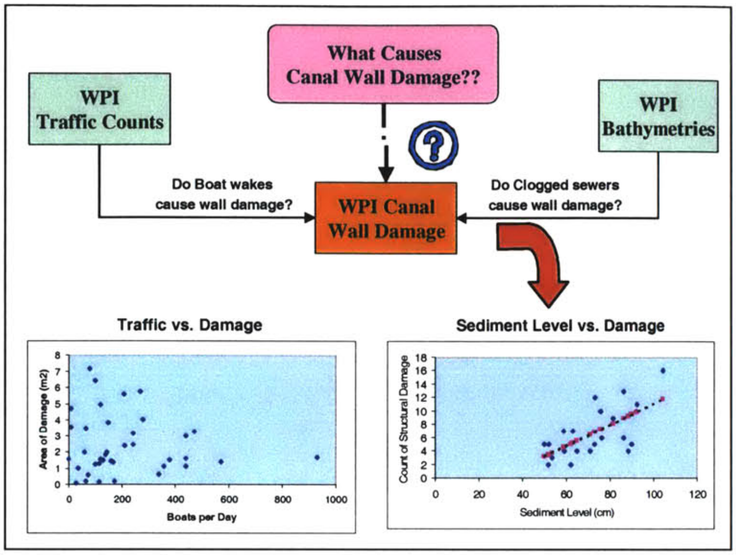

Back in 1998, Babic *et al.*, using the accumulated storehouse of knowledge on traffic, wall damage and bathymetries, concluded that the root cause of wall damage is lack of dredging, which is only later compounded and exponentially accelerated by traffic and wakes[278]."

Based on this knowledge, in the year 1999 a further study was conducted[279] to quantify the relative and absolute contributions for a variety of possible sediment sources, including the debris produced by crumbling masonry, to the accumulation of sediment at the bottom of canals. This study led to the proposal for a Sedimentation Model to predict the rate and location of mud accumulation at the bottom of canals, which would allow a more effective and efficient scheduled maintenance before more wall damage was generated by the clogging of sewer outlets.

[278] Babic *et al.*, 1998.

[279] Borrelli *et al.*, 1999.

wake-loading

In 2002, the concept of a "moto ondoso index" was developed, that translates levels of boat traffic (i.e. number of boats) to levels of "wake-loading" (how much wake energy is discharged in the canal), which helps to better correlate traffic to damage[280].

turbulent discharges

More recently, WPI students[281] designed and successfully tested an instrument that maps the locations where motorboats discharge energy into the canals when maneuvering to make turns, or otherwise stopping abruptly by shifting into reverse when an approaching boat threatens a collision, or even when moving back and forth near a dock to tie up the boat and unload people or cargo. This custom device is equipped with a differential GPS, a triaxial accelerometer and an RPM meter and will produce the first ever map of "turbulent discharges" in the inner canal network, further facilitating the prediction of future damage along canal walls.[282]

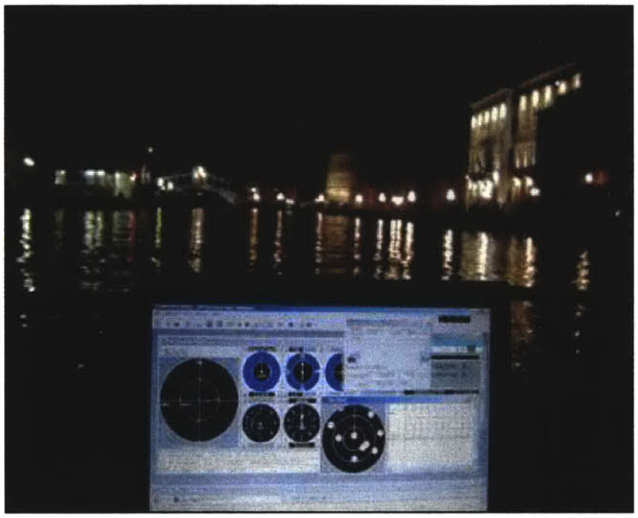

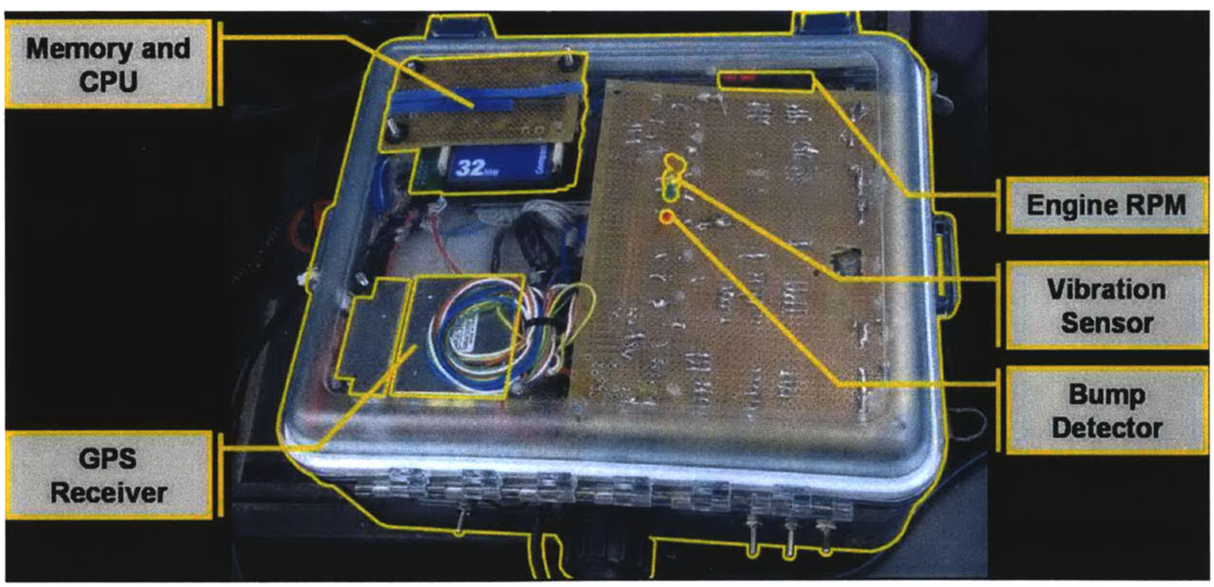

280 Chiu *et al.*, 2002; see also page 82.
281 Chiu *et al.*, 2004.
282 See page 83 for a map produced by this instrument.

"interscambio merci"

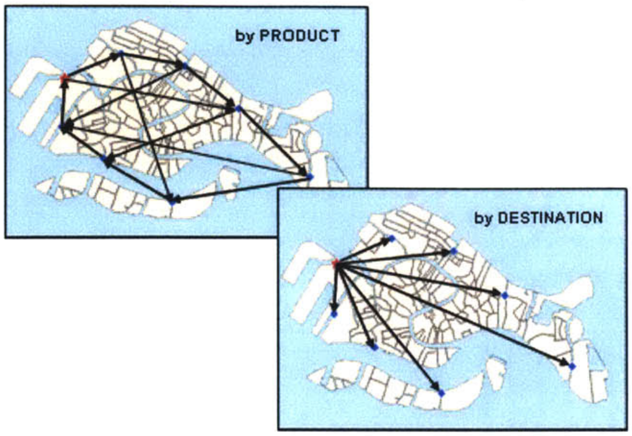

The application of city knowledge principles has paid off dramatically in another project related to *moto ondoso* entitled *Re-Engineering the City of Venice's Cargo System*[283]. The project has demonstrated the plausibility of "plan-demanding" knowledge as a consequence of "plan-ready" information, in opposition to the traditional modus operandi of "plan-demanded" data collection.

Here, work on the optimization of canal closures[284], which produced plan-ready information on the amount of cargo delivered to each Venetian island, has led directly to the spontaneous emergence of the need to develop a plan (hence the term "plan-demanding") to restructure the way deliveries are conducted in Venice. The 2001 award-winning project was conducted with and for the former *Consorzio Trasportatori Veneziani Riuniti* (a group representing about 70% of all cargo boat drivers in 2001) and resulted in a proposal to redistribute merchandise "by destination" instead of "by product".

We estimated cargo demand by inventorying all shops and stores in Venice and by surveying representative samples of each typology of commercial establishment to quantify the number of parcels delivered to each store of that type everywhere in the city.

283 Duffy *et al.*, 2001. This is the same project mentioned in the "Plan-Demanding" discussion on page 37 ff.

This effort led to the aggregate depiction of the total demand of dry (yellow) and refrigerated (blue) cargo for each island in the city.

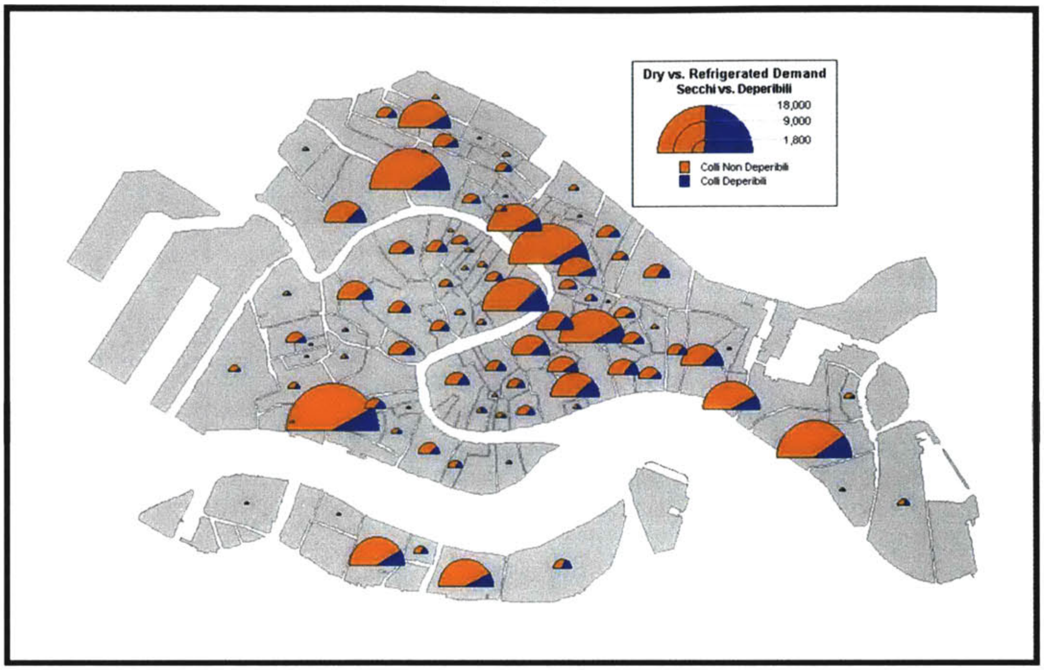

The research also proposed the creation of a cargo warehouse (*interscambio merci*) in the Tronchetto area to allow for the sorting of the merchandise arriving by truck. We envisioned that the city would be divided into 16 zones with commensurate delivery demands. The warehouse would mirror such a division, reserving a loading bay for each zone.

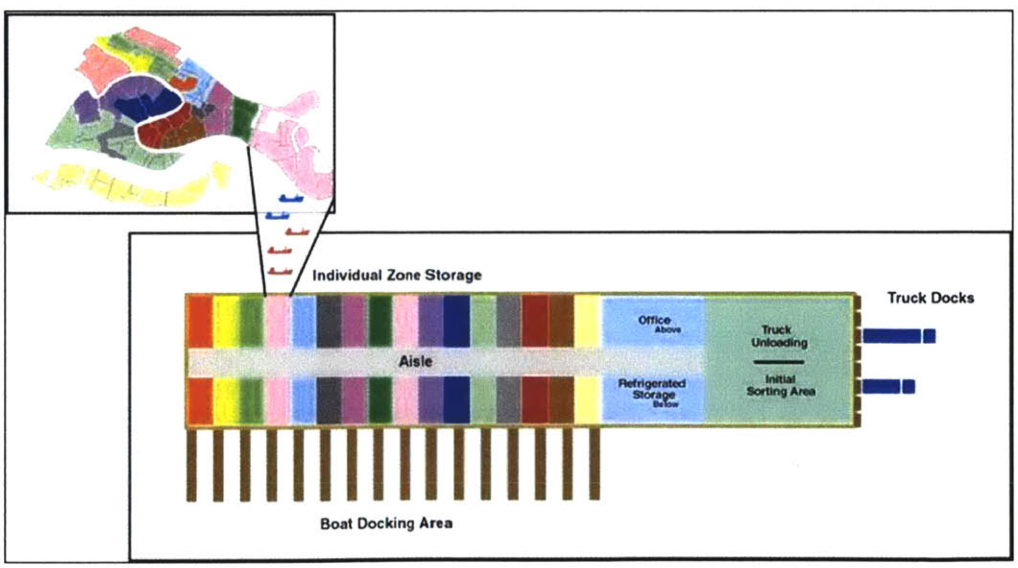

284 Amlaw *et al.*, 1997.

In the end, it was possible to estimate the total number of boats (for dry and refrigerated goods) that would be necessary to deliver all necessary merchandise to all the stores in each zone.

reducing wall damage with knowledge

Boat drivers were repeatedly engaged in public meetings as well as surveyed and interviewed during the development of this plan. This participatory process created consensus over a plan that would make all licensed drivers shareholders of a single, large logistics and distribution company delivering most of the goods in the city. Boat drivers would be able to deliver to their destination dock at a set time and without interference from other cargo boats. Vacations and sick days would now be possible without the fear of losing clients. The great majority of boat drivers that participated in the public meetings found the benefits of such a system advantageous.

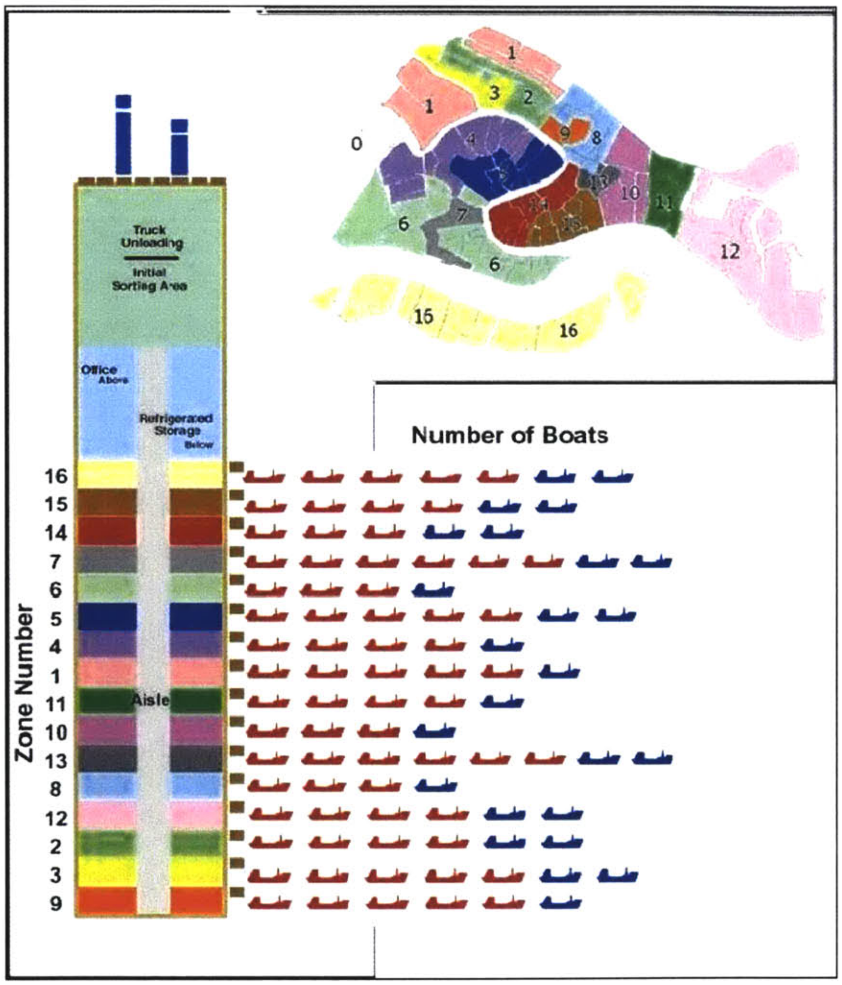

Presently, the municipality of Venice is moving forward with the actual implementation of this plan. This proposal will reduce cargo-generated *moto ondoso* by a substantial amount while preserving non-seasonal jobs. Point-to-point cargo journeys will be reduced to one or two per boat, instead of the dozen or more segments traveled each day under the current system. A reduction of the order of 90% of the wall damage induced by cargo boats is therefore not implausible[285].

In the summer and Autumn of 2001, the proposal was presented to the Mayor and the Vice Mayor who espoused it. The project, which is currently well on its way to being realized, was a triumph of City Knowledge principles, demonstrating the "plan-demanding" potential that such an approach entails. A plan-demanded study dedicated to minimizing disruption to deliveries during canal closings for maintenance and dredging operations led to the accumulation of enough plan-ready information about the inefficiencies of the system of deliveries "by product" to spur the inception of a follow-up study to explore the boat-drivers' perspective on the revolutionary approach of deliveries "by destination" that

285 Carrera, 2002.

was proposed. The project went full-circle, from plan-demanded to plan-demanding.

Although we have "deep knowledge" that tells us that timely dredging is more important than reducing wake-loading, nevertheless, whatever damage motor boat wakes do induce upon the canal walls will be reduced by a fraction that is proportional to the wake-loading that the number of cargo boat transits being eliminated would have caused as a percent of the total wake-loading induced by all boat types in each canal segment.

lessons from "moto ondoso" analyses

[value-added]

This section manifests the ultimate power of City Knowledge, i.e. the ability to create wholes that are bigger than the sum of their parts. Here is the ammunition for the "value-added" argument that should provide the final clinching proof about the efficiency, effectiveness and efficacy of City Knowledge.

[efficiency]

The efficiency of City Knowledge is demonstrated here by the lack of redundancy in the datasets and the smooth collation of information needed for the plan-demanded study, which mirrored the idealized situation of non-overlapping informational jurisdictions envisioned by the City Knowledge approach. Of course, the fact that the VPC functioned as a single overarching entity and not as a distributed cadre of information-producers may be taken as evidence of the possible primacy of the centralized approach versus an emergent decentralized one. The fact is that the data that were used in these examples came from a variety of sources, administrative, academic, municipal, as well as our own data collection. The upkeep of the fundamental datasets for each one of the aspects we dealt with in this example could be easily attributed to one agency or another, based on informational jurisdictions.

[effectiveness]

We could measure the "operational effectiveness" of the example discussed in this section according to Budić's indicators[286]. She selected a sample of GIS practitioners in 125 county and municipal governments in four south-eastern states. According to that sample, our Cargo Project would be very appreciated by 63.6% of GIS professionals because of the instant accessibility and availability of the data we used to arrive at our re-engineering proposal. Validating the accuracy of the data would make 27.3% additional users very happy. A well-crafted City Knowledge system should therefore be very effective operationally.

As far as "decision-making effectiveness", the plan-demanded example discussed herein should have made clear the power of the GIS tools to communicate the information pertaining to the cargo delivery system. This alone would have been highly desirable to half of the sample of GIS practitioners surveyed in Budić's study. It could also be argued that the example discussed herein showcased both the ability to aid in the identification of conflicts and also in the confidence in the analyses that produced our radically re-designed delivery system, making it even more effective for decision makers.

286 Budić, 1994, p. 252, Table 4.

[efficacy]

The efficacy of our framework in terms of the measurable consequences of our proposals in the "real world" cannot be quantified quite yet since the proposed delivery system has not been put in place as of this writing. Measures of efficacy may include the financial benefits or quality-of-life gains of the citizens of Venice after the inauguration of the system.

THE PROMISE OF CITY KNOWLEDGE

This paper argues that not only should Venice entertain the notion of a 'central bank'[287] guaranteeing the free flow of data, to be open to all and transparent, but that the city should more importantly employ a sustainable methodology to allow such a bank to emerge from the middle out (not from the top down or from the bottom up) as the sum total of the data produced by a whole variety of contributors distributed in the territory. For such an endeavor to really have staying power and to take on a life of its own, it will be necessary to forego the old-fashioned notion of a "central bank" and replace it with the principles of City Knowledge introduced in this paper.

middle-out from each jurisdiction

The basic tenets of City Knowledge which are aligned with some of the most recent trends in Geographic Information Systems, suggest the adoption of a middle-out approach based on clear "informational jurisdictions" assigned to the producers of urban information, starting with the plethora of municipal offices which, through the approval of permits or the assignment of licenses or other administrative acts actually cause – or more frequently allow – the city to change ever so slowly.

intercept administrative transactions

Intercepting these administrative transactions will enable City Knowledge systems to maintain the information up-to-date more or less automatically. Self-interest is the fuel that will make the City Knowledge system thrive, since all offices will have self-serving incentives for making their operations smoother and more citizen-friendly. Once a number of departments in the City have embarked in the full life-cycle analysis of the information flows that guide their actions, leading to the identification of their specific jurisdictions, and once the backlog of existing information is captured in databases and GIS, scholars and scientists, planners and decision-makers, as well as citizens at large, will be able to enjoy the benefits of the synergic, emergent properties of a connected and shared City Knowledge system[288].

creating a virtual central bank

Although the path to a full City Knowledge system will suffer from typical implementation woes, such an approach would create a virtual 'central bank' on everyone's desktop and would make Venice a model for sustainable municipal information systems all around the world.

287 This was proposed as a conference recommendation (Cambridge Conference, 2003).

288 As partially demonstrated by Hart et al., 2004, who analyzed the information flows in the Environment Department of the City of Boston, specifically with respect to the Boston Landmarks Commission and the Boston Conservation Commission.

PART III:

THE DEVOLUTION OF CITY KNOWLEDGE

"In truth, because of the nature of the work to be done, almost all city planning is concerned with relatively small and specific acts done here and done there, in specific streets, neighborhoods and districts. To know whether it is done well or ill - to know what should be done at all - it is more important to know that specific locality [...]. No other expertise can substitute for locality knowledge in planning, whether the planning is creative, coordinating or predictive.

The invention required is not a device for coordination at the generalized top, but rather an invention to make coordination possible where the need is most acute - in specific and unique localities."

Jane Jacobs, 1961
"The Death and Life of Great American Cities", p. 544

EXPORTING THE LESSONS

the unreflective practitioner

The indisputable uniqueness of Venice may have discouraged any type of generalization of what we learnt there since 1988. As intriguing as my insights were, it wasn't immediately clear that my work there might contain some exportable lessons for other urban professionals to learn from. I hadn't really had a chance to reflect upon my practice as Don Schön would have encouraged me to do[289]. I was too busy walking the walk to even find the time to talk the talk (or more appropriately to "think the thought"). However, as fate would have it, in 1998 I was asked by my department head at WPI to take over as director of the nascent Boston Project Center, so, quite serendipitously, I was thrust into a new context and began to advise projects for the City of Boston and for other municipalities in the greater Boston area, like Cambridge, Newton and Quincy.

serendipity

"Unencumbered by the thought process"[290], I went ahead and began to develop projects in the same vein as I had been doing in Venice for a decade, in order to help frontline city officials to find practical solutions to day-to-day issues confronting the municipality. The issues were unlike those in Venice, but not drastically so. The approaches we developed seemed also different at first sight, but upon further inspection, they revealed some common principles not too dissimilar from those we employed in Venice.

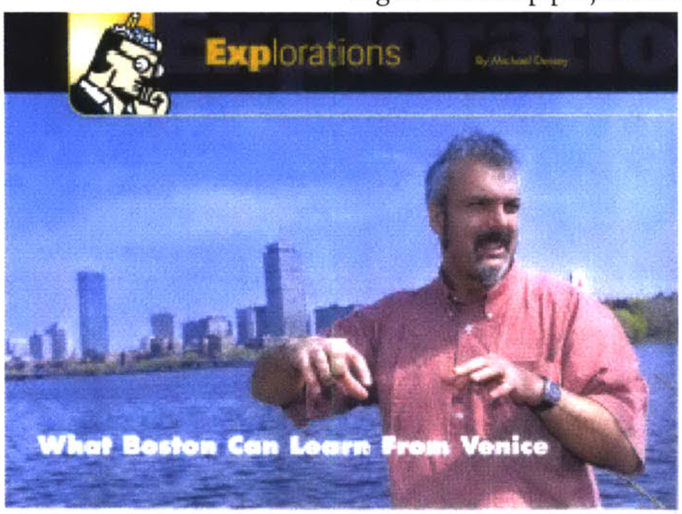

In the two chapters that follow, I introduce a small sample of the projects we have conducted in Massachusetts. Together with the over 30 other projects we have completed on this side of the Atlantic using the approach developed in Venice, this sample demonstrates – if nothing else – that the City Knowledge approach can work in widely different social, cultural and administrative contexts.

After each section, as was done in Part II, I summarize the *lessons* learnt through the experiences described therein. All of these lessons will come together into Part IV which represents the culmination of the dissertation.

289 Cf. Schön, 1983, *The Reflective Practitioner.*

290 I couldn't resist the opportunity to quote the "Tappet Bothers", fellow MIT graduates from "our fair city" of Cambridge, heard weekly on National Public Radio's *Car Talk.*

CAMBRIDGE CITY KNOWLEDGE

WPI Projects in Cambridge

The City of Cambridge was one of the towns founded in the 1630's by the earliest pilgrims. It is home to two of the most prestigious universities in the world, Harvard and the Massachusetts Institute of Technology (MIT). As of the 2001 census, Cambridge had a population over 100,000. The municipal government in Cambridge is composed of an elected mayor and a hired city manager that effectively runs the city in the day-to-day sense. The city has been a leader in innovative bottom-up planning initiatives and has even instituted a form of "middle-out" strategy, orchestrated by the MIS department, to distribute GIS capabilities into many of the front-line departments. Simultaneously, there has also been a top-down push to institute uniform accounting practices such as the management of work-orders, with a system by a company called Hansen Information Technologies[291].

Through a fortuitous connection, I arranged the first WPI project in Cambridge during the spring of 2001. This was the third year of operations for the WPI *Boston Project Center* and it marked the first time we actually stepped outside of the boundaries of the city of Boston proper.

Cambridge Trees

"the Works"

tree maintenance

I came in contact with Mr. Larry Acosta, City Arborist in the department of Public Works of the City of Cambridge[292]. Despite the conspicuous resources allotted for the maintenance of the urban forest, the City did not have a detailed, computerized inventory of its trees. From year to year, tree pruning was carried out "by zone" instead of more selectively. Basically, in a given year, a fixed amount of money was set aside to take care of any maintenance need for all trees within a certain neighborhood. The work was contracted out to private firms who were expected to provide detailed records of the work carried out, so that payment could be made appropriately. Staffers from "the Works" inspected the work done and confirmed that the job had been properly executed. Despite all of the checks and balances, Mr. Acosta felt that the scheme he inherited when he moved to Cambridge from California was still inadequate vis-à-vis the millions of taxpayer dollars spent on tree maintenance. He envisioned a computerized system that would enhance the overall accountability and would thus also optimize the allocation of funds from year to year.

tree-induced energy savings

In parallel, the same department was also interested in quantifying the benefits in term of energy savings that the "urban forest" brought to the citizens of Cambridge. Trees shade buildings in the summer, thus reducing the need for air conditioning, and they also protect buildings from chilly winds in the winter, thus reducing the heating bills. The quantification of the energy savings is based on the size of the canopies as well as on some topological relations between the trees and the buildings in the city. The public works department had acquired a special software to help in this

291 See www.hansen.com.

292 Officially nicknamed "the Works" in its logo.

exercise[293]. The approach the city adopted, which was dictated by the structure of the simulation software, entailed using aerial photographs to measure the total canopy in order to enter the necessary figures into the software application, which was then able to produce estimates for energy savings based on the information received.

My instinctive recommendation in reaction to the issues presented to me was to suggest scrapping the software package and to invest resources in a thorough inventory of all city trees, which would serve the dual purpose of supporting the maintenance activities as well as to create the platform for the subsequent canopy calculations that could be translated into estimates of energy savings, without the need to pass through another software program.

tree inventory

The WPI project was therefore conceived to develop and field-test a methodology for the inventory of individual trees in Cambridge[294]. The catalog included all of the parameters that are useful in order to maintain the tree in good physical health, free of diseases and deformities. The

[tree codes]

parameters included were also deemed sufficient to derive the energy saving estimates that were of interest to some sectors of the municipal government. In concert with Mr. Acosta, the WPI students defined the syntax of a "tree code" that would identify each tree uniquely and proposed a list of useful parameters that should be collected about each tree while in the field.

[tree condition assessment]

Defects and disease symptoms were codified and a rubric for their evaluation was prepared by the student team who then took the methodology into the field to test its validity and efficacy.

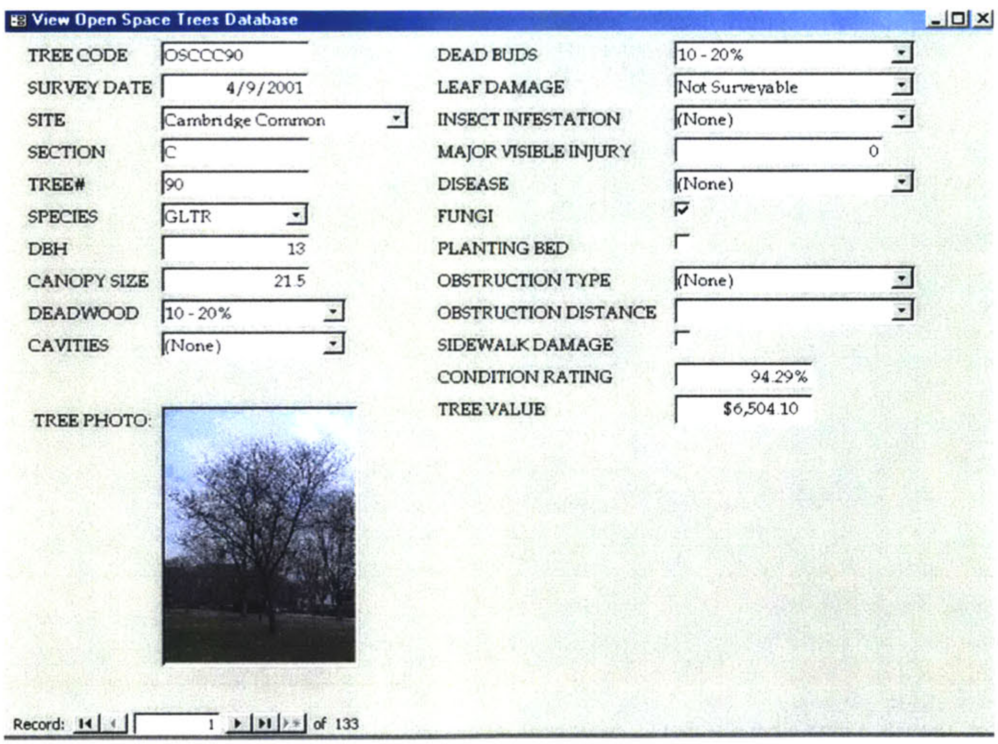

293 See for instance http://www.americanforests.org/productsandpubs/citygreen/.

294 For something similar and with the addition of a web-GIS, see Kelly and Tuxen, 2003.

The location and species of a tree, together with its unique code, constitute the fundamental items of information associated with each plant. Locations were first pinpointed using a laser-guided GPS. This system was plagued by systemic errors[295] which required extensive manual corrections. In this case, as we had already realized during the measurement of canal depths in Venice, the manual process of positioning the tree on a map using local reference landmarks and tape measures proved to be more effective and expeditious than the high-tech solution involving GPS[296]. Fortunately, the position of the tree may be corrected repeatedly as better and better geopositioning tools become available, but the object itself will not need to be modified after its initial mapping. The codes and basic information associated with each tree symbol are good forever, irrespective of minor positional adjustments that may follow.

[low-tech beats high-tech]

trees and energy savings

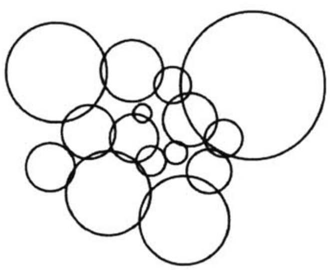

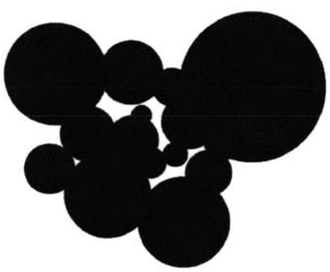

Energy savings were calculated by adding the canopies of all individual trees in the catalog. First of all, the "atomized" information about canopy radii enabled us to calculate the area of each tree's canopy[297]. The sum of individual canopy areas would approximate the overall canopy coverage, which in turn was converted into an estimate of energy savings through proven formulas obtained from scientific literature. If one had started by measuring canopy areas as amorphous "blobs", the energy calculations could have been made just the same, but the human and financial costs incurred would have been useful only and exclusively for this effort. No useful information about single trees could be possibly derived by attempting to disaggregate the "blobs" into individual trees. On the other hand, once all the trees in a city have been immortalized in a GIS and recorded in a relational database, the catalog will lend itself to a variety of uses, making the long-term endeavor I propose well worth the protracted effort.

Some may argue that the widespread availability of aerial or satellite imaging, coupled with remote sensing techniques, make the top-down analysis of tree canopies more and more affordable and feasible. My counter-argument would be that the tree-by-tree approach I propose is going to provide a more sustainable, updatable and re-usable foundation for multiple applications – including the tree canopy energy estimates – for years to come, provided an appropriate update mechanism is adopted.

update mechanisms

To ensure that tree information is always as current as possible, we proposed that the update mechanisms for tree information ought to be tied to the physical maintenance of the plant. In essence, we suggested that the

295 Even though this project took place after the US government loosened the anachronistic, national-security restrictions on the accuracy of GPS, our GPS results were not satisfactory. Prior to May 2, 2000, the US government inserted errors into the system on purpose, through a process termed "selective availability". The error went from about 45 meters to about 6 meters overnight (http://www.igeb.gov/sa/diagram.shtml). However, a 6 m error is still unacceptable for pinpointing a tree's location.

296 The discrepancies we encountered when using our laser-guided GPS may have been at least partially due to inaccuracies in the basemaps we were using at the time and not only due the the selective availability limitations of the GPS itself.

297 Assuming, of course, that a circle would sufficiently approximate the shape of a tree's canopy.

[informational returns by contract]

[lower cost of maintenance-based updates]

tree-pruning squads measure a new canopy radius and a new trunk diameter during routine maintenance operations, providing an informational return together with the horticultural services. In addition to these physical measurements, we recommended also that the maintenance crews provide feedback about the health of the plant, by noting telltale signs of disease or damage, based on an evaluative rubric that we developed and tested for efficacy. Finally, we proposed that the field teams should also note the impact of the tree on surrounding assets, like power lines, sidewalks and storm-drains by using objective metrics that could be applied consistently by different staffers or contractors.

Our deconstructionist method of parameterizing tree characteristics in such a way as to allow non-expert individuals, with a modicum of training, to consistently measure, evaluate and assess the collection of trees in a city[298], translates into tangible savings for the city coffers because:

[no waste]

(1) **Superfluous maintenance is eliminated**. Trees are maintained selectively based on solid factual information, instead of applying blanket procedures to entire sections of town based on brute-force, information-blind methods;

[no travel]

(2) **The time spent to reach a tree in order to collect updated information is eliminated**. Information can be kept up-to-date by appropriately-trained ground crews who are already sent out to the field anyhow, thus eliminating one of the major cost components of information retrieval, namely the travel time to and from the assets to be monitored;

[blue-collar updates]

(3) **The cost of retrieving updated information is greatly reduced**. Ground personnel is less expensive than a trained botanist, even if we factor in the limited training needed to ensure adequate data quality and consistency[299].

maintenance-based cataloguing

Of course, if we are willing to be very patient, we could use the very same maintenance-based approach I just described to simply initiate the process of gradual data collection and mapping of each tree as it becomes the object of maintenance. The only difference here (which may be actually a major issue in terms of cost and training) is that the ground crews would have to also be able to recognize the species of the tree and would need to be trained in surveying techniques, either with GPS or otherwise, to produce accurate positional coordinates to generate the correct GIS maps. These added tasks would demand a considerable amount of time from the maintenance crews, though they would not necessarily cost significantly more, thanks to the widespread availability of low-cost technology in this day and age. Despite the ever-diminishing cost of technology, using ground crews to create the catalog from scratch will always be much less cost-effective than employing them to simply update existing data.

298 Following Hammond *et al.*, 1980, 1991.

299 It could be argued that collecting updated information all the time, even if done by blue-collar workers, may in the end cost more than hiring a botanist once in a while. However, the sheer magnitude of an urban forest, with tens of thousands of trees, makes it unrealistic to consider using highly-trained individuals to conduct a thorough survey of all plants owned by the city.

Nevertheless, a maintenance-based approach to cataloguing could work for simpler types of point-like city assets – like street lights or fire hydrants – which would only require adequate GPS locating skills (not very complicated to teach) and would not bog down the maintenance crew with too much additional work above and beyond the actual maintenance.

On the other hand, this promising approach would not work well for region-like assets – like lawns, or flower-beds – which require a more concerted initial effort to map out, just as it wouldn't work well for more complicated point-like objects that present complex challenges, like the need to determine the botanical species of a tree. My personal bias here is that the initial catalogue should be done by trained professionals if financially possible, to guarantee that the whole database and GIS infrastructure is constructed on very solid and reliable foundations, leaving just the updating (which will be the bulk of the work in the long run) to less-skilled individuals.

[start catalog with professionals]

Aside from the aforementioned "instant savings", achieved through our City Knowledge approach, the fine-grained nature of our deconstructed datasets engenders numerous – sometimes unforeseen – opportunities for creative re-use of the data to produce information above and beyond what was intended originally.

leveraging plan-ready information

For example, our tree canopy area calculations were not only useful for determining energy savings, but were also leveraged to calculate the impact of trees on air quality. Depending on the species and the canopy dimensions, we could estimate the quantities of major pollutants (like CO_2 or SO_2) removed from the air by our trees, thus providing an additional measure of their "value" to society, who might therefore be more forthcoming with the funds to support the efforts to catch up with the backlog of tens of thousands of trees in each city that are already out in the streets and parks and need to be recorded into GIS and Databases to begin applying these cost-saving measures when it comes to their upkeep.

[trees and air quality]

separating facts and values

Tree limbs constitute ever-present threats to public safety. Every year, dozens of people die because of tree accidents, exposing municipalities to potential liability lawsuits. A well-designed updating protocol would try to capture potentially dangerous situations and nip these hazards in the bud, within reasonable limits. The inclination of a tree, for instance, could constitute a factor in determining its level of threat. As tempting as it may be to simply assign "danger levels" based on visual inspections, we have always tried to stay away from these "value judgments" that are prone to be inconsistent and subjective, preferring to focus on the ingredients that factor into such judgments, and trying to objectively quantify these factual components before applying more value-based weights to each factor and before applying judgment-based techniques to combine the weighted factors into a final "danger score"[300].

[trees and safety]

[danger]

[disease]

Similarly, we devised a way to separate visual symptoms of disease ("facts") from the determination of the level of health of the plant ("value").

[300] Hammond *et al.*, 1980, 1991.

This way, after minimal training, ground crews would be able to detect early warning signs of illness.

maximizing informational returns

Another source of potential harm, as well as costly damage, is the interference of tree branches with power lines. Utility companies routinely hire tree companies to trim down trees that are threatening to interfere with the power grid. Trees are a major source of blackouts in the world[301]. Our approach in Cambridge took these issues into consideration and a special parameter was set up to record the closest type of obstruction to the tree, which often was a power line. The distance was estimated to the closest foot. In a similar vein, we also made provisions for the assessment of the status of the root system vis á vis the sidewalks and roads.

[trees and power lines]

[sidewalks]

The project produced a very successful demonstration of a system that would allow the city arborist to:

- plan future maintenance based on full knowledge of the tree inventory, including tree locations and species
- detect spatial disease patterns thanks to GIS mapping. This in turn would dictate preventive measures to deal with epidemics.
- receive early warning signals about diseases, defects, interference with power lines, disruption of sidewalks, and other useful diagnostics.
- estimate energy savings and pollution abatement based on the cumulative sum of all canopies, calculated according to sophisticated spatial analyses.
- maintain a constant watch over the "urban forest" and react accordingly in a timely fashion.

[301] Fifty million people across the entire country of Italy were left without electricity on September 29, 2003 after a tree branch fell on a power line in Switzerland. This was one of the largest blackouts ever and was the cause of at least three deaths (AP, 9/29/03, 11:10am).

Lessons from Cambridge

transferability of City Knowledge

The tree management application described herein allowed us to first of all demonstrate that basic City Knowledge tenets apply in Cambridge as well as in Venice. Of course, the way in which we go about pursuing City Knowledge is skewed by the fact that we are outsiders coming in laterally with an army of skilled students. This fact may account for much of the similarities we encountered between Boston and Venice.

[transatlantic cross-cultural transferability]

Interestingly, the Cambridge tree project gave birth to a sequel in Venice, bringing the process full-circle. In turn, the Venice project lead to my involvement in the development of a professional tree-management web-GIS that is currently in the process of being brought on line at VESTA, the public/private company that manages garbage, water and parks for the city of Venice.

[transferring from academia to municipality]

The main challenge remains how to set up virtuous mechanisms that will self-generate similar successes – without relying on free academic studies – but tapping into a variety of resources and using a variety of schemes to move the process forward in a financially sustainable way. The real transfer does not have to take place from Italy to the US or vice versa. It needs to take place between academia and municipal administration.

transaction vs. snapshot

In this project, the demarcation between systematic tree-by-tree cataloging and the broad-brushed canopy assessment made clear that one of the hurdles that City Knowledge systems need to address is the cost/benefit dilemma between transaction-based and snapshot-based systems . The one-tree-at-a-time approach which I propose can serve multiple purposes and will completely eliminate redundancy and waste. The snapshot approach of doing the canopy analysis and move on seems utterly wasteful. In a sense, the tension is really the familiar one between short-term gratification versus long-term commitment. City Knowledge relies on the transaction model to capture the changes that occur to "permanent", administratively-controlled elements of the city. The transactional approach could potentially be fed by tapping into administrative databases, but could also extract greater informational returns from outside contractors and internal staff to slowly but surely build the storehouse of urban information that City Knowledge calls for.

[transaction approach for permanent features]

[snapshot approach for dynamic features]

Nevertheless, for dynamic anthropogenic activities such as traffic, as well as natural phenomena such as hydrodynamics, there are no transactions to be captured, so a snapshot model is the only one applicable.

maintenance-based updates and cataloging

Hybrid approaches can also be advantageous. Tree assessments consist of a the tree-by-tree snapshot of current dimensions, health and contextual impact. This snapshot-type of information gathering can be turned into a combined transactional/snapshot hybrid by attaching the data collection activities to routine maintenance operations through contractual obligations.

separating facts and values

One way to ensure that tree maintenance crews can be competent participants in the proposed distributed, maintenance-based data collection system is to devise a fool-proof methodology for the data-gathering. The key to a successful field methodology is to limit the number of decisions that the field worker needs to make while gathering data. Ideally, field forms should aim at collecting the most objective type of data, without any judgments or interpretations needed. With some training, crews can be taught not only how to make simple measurements, but also how recognize diseases and how to assess dangerous or dubious situations.

cost savings

By removing expert-judgment and allowing lower-paid employees to conduct the data-collection, this approach can greatly reduce the costs of information management. The information can be gathered *ex-novo* or simply updated using lower-skill personnel, thus achieving a net cost saving. Additional savings come from requiring the collection of up-to-date data concurrently with regular maintenance activities, such as pruning. It is hard to put a precise dollar value on the savings, but logic suggests that the solutions I propose would indeed provide sufficient cost-savings to justify pursuing a rigorous and systematic knowledge acquisition process.

WORCESTER CITY KNOWLEDGE

Despite being the home of our Institute, the city of Worcester was initially neglected by the WPI Global Perspectives Program. As our sights were set to global horizons, we passed up the challenges and opportunities that were in our own backyard. Nevertheless, it is not surprising that we were slow to recognize the desirability of offering the same "full-immersion" project experience to students who were not as interested in exotic locales, but nonetheless equally eager to make a difference in the community in which they lived.

THE WORCESTER COMMUNITY PROJECT CENTER

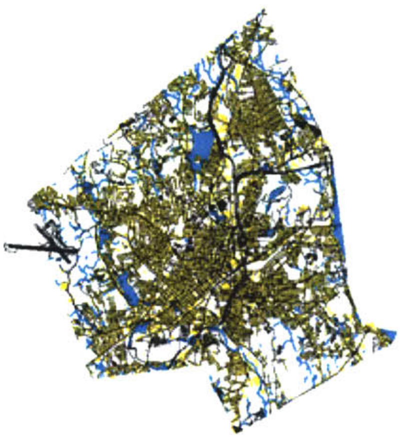

To address this shortcoming, WPI created the Worcester Community Project Center (WCPC) in the year 1999, with support from local foundations. After the initial stewardship of Associate Provost Lance Schacterle, Prof. Rob Krueger was nominated director of the center in 2001. Worcester is the second largest city in Massachusetts and third largest in New England, just slightly behind Providence, Rhode Island in terms of population[302]. After playing a major role in the War of Independence and in the Industrial Revolution, Worcester underwent the same boom-and-bust cycles of many other manufacturing towns that have had to reinvent themselves in the post-WWII period. Worcester has the advantage of having been traditionally the home to several colleges and universities, including Clark, Holy Cross, Assumption, Worcester State, Quinsigammond Community, Becker Junior College and UMass Medical, in addition to WPI.

The colleges have been an economic staple of the city and have contributed to the establishment of a strong biomedical and biotechnological sector in recent years, which has in part offset the continued losses in the manufacturing base.

IMBY – in my back yard

Since its founding in 1865, WPI has contributed greatly to the development of the city's industrial heritage. WPI intends to play a renewed role of economic stimulus in its own backyard by leveraging its expertise in science and technology, as witnessed by the recent creation of the Metal Processing Institute, the Bio Engineering Institute and the plans to incubate new high-tech businesses through the Gateway Park redevelopment.

a city as a test bed

From my perspective as a student of urban issues, I have also belatedly awakened to the opportunities that Worcester provides to be a real-life test bed for my municipal information theories. In that vein, starting in 2002, I began to become engaged in the WCPC to explore how City Knowledge principles could be exported to a city of significant size, yet amenable to new ideas.

The first opportunity to play an active role in the functioning of the city of Worcester was provided to me by my participation in the creation of

[302] In the year 2000, the population in Worcester was 172,648, versus a total of 173,618 in Providence. (Census Bureau, http://quickfacts.census.gov/qfd/)

a Community Development Plan (CDP) for Worcester, as part of a statewide planning initiative called Executive Order 418.

Executive Order 418

EXECUTIVE OFFICE of ENVIRONMENTAL AFFAIRS

dhcd Massachusetts

Executive Order 418 (EO418) [303] financed a state-wide effort for the development of "Community Development Plans" to "help communities plan for new housing opportunities while balancing economic development, transportation infrastructure improvements and open space preservation"[304]. The upshot of this innovative program was to create a widespread planning consciousness at the local level, where change really happens. EO418 addresses four main areas:

1. Housing
2. Economic Development
3. Transportation
4. Open Space and Resource Protection

regional planning agencies

The primary focus of EO418 was to assist "communities in addressing the housing shortage", so the Department of Housing and Community Development was the leading agency in the effort, under the auspices of the Executive Office of Environmental Affairs. Each of the 351 communities of the Commonwealth of Massachusetts received $30,000 to complete the EO418 process. Regional Planning Agencies (RPA) would act as intermediaries between the towns and the State. The RPAs would conduct the initial buildout studies and support the development of the CDP. The general purpose behind this very progressive initiative was to engender many local conversations about the lot of each town in order to inspire citizens to consider if the future that was on the horizon was truly what they envisioned as a community. If the "probable, almost certain, future[305]" of a community does not match the collective "vision" for the development of the town, then each city and town should fulfill on the expectations of the population through the implementation of the Community Development Plan.

community development plan

build-out

In order to raise the consciousness of the citizens of Massachusetts, the first step in the process entailed a form of "shock therapy" to jolt the apathy out of their systems. The best way to make people pay attention is to show them how bad things could get if they don't. The buildout analysis that forms the foundation of EO418 is supposed to do just that, by simply demonstrating the extent of construction that the current zoning regulations would allow if each zone was developed to its maximum allowed envelope, taking into consideration other physical or environmental limitations. This is what EO418 calls the "Default Future Land Use Map"[306].

community visioning

Once citizens were awakened to the real impact that current zoning could have on the shaping of our urban realm, EO418 foresees tapping into the collective imagination in order to extract an alternative "vision" that

[303] Cellucci and Swift administration, 2000 (http://www.massdhcd.com/eo418/homepage2.htm).

[304] Commonwealth of Massachusetts, *Building Vibrant Communities*, p. 1.

[305] I am borrowing this phrase from *Landmark Education*.

[306] EO418 Technical Assistance Bulletin #3, p. 2.

could be used to reinvent the zoning regulations and to creatively apply the other four tools of government action[307] to realign the future growth of the city with the will of the people. Such a process would entail conducting surveys, focus groups and other participatory interactions with the inhabitants of the town, representing all categories of people and all demographic segments.

assets and liabilities inventory

Each of the four main areas is then scrutinized carefully and thoroughly, by using available information and, when possible, by collecting new data about useful indicators of the state of affairs in each area. Although the various EO418 documents do not put a lot of emphasis on this aspect, the underlying inventory of existing assets (and liabilities), such as the housing supply inventory or the inventory of environmental resources and others are essential foundations for the subsequent analyses. City Knowledge approaches could be applied to the accrual of these inventories, as will be described below.

deliverable GIS maps

The main deliverables of all these four sector studies would be the following maps:

1. Open Space Suitability Map
2. Housing Opportunities Map
3. Economic Development Opportunities Map
4. Transportation Issue Map
5. Community Development Plan Final Land Use Map

Of course, the production of each map would in turn demand the collection, organization and manipulation of several datasets – some easily available in the public domain, some not. Suitability analysis entails the establishment of a rank-order classification of uses that would be most suited to each parcel of the city based on the result of the visioning sessions discussed above. The documentation that accompanies EO418[308] is not very specific about the exact nature of these suitability analyses, except to require that they result in suitability maps that display the "best use" that could be located in each part of town, depending on physical features (such as slopes), environmental characteristics (like the presence of wetlands), current zoning regulations (like Commercial, Residential, etc.), as well as a host of other parameters that are deemed important by the community. All suitability analyses really need to be framed by a careful assessment of current and future demand for each of the possible land uses, as well as supply of resources to meet all or some of the predicted needs in the years to come.

[supply and demand]

review of past plans

A useful preamble to the suitability analyses is a through review of past plans. I a sense, plans represent an expression of demand. Some plans embody a collective vision and some only the will of an individual. Some may express powerful social demands, others venal yearnings for financial

---

[307] Schuster *et al.*, 1997. The five tools of government action are: (1) Ownership and Operation; (2) Regulation; (3) Incentives and Disincentives; (4) Information and Education; and (5) Rights (the most confusing and controversial). I would recommend thinking of something like "quid pro quo" or "barter" as a better name for the 5th tool.

[308] http://www.massdhcd.com/eo418/homepage2.htm.

gain. Plans offer important snapshots of the local conditions at a certain moment in history. They synthesize a multitude of city information and often try to predict the future development of the area, with or without the implementation of the proposed plan, so as to provide a "before and after" comparison of alternative futures.

open space protection

Open space suitability would entail the determination of what parcels should be protected from development and which would be best suited to be used as parks of various shapes and sizes, from the small neighborhood Playground, to the larger City Park (which may or may not include sport fields), to the more protective Conservation Land, which is often fairly extensive in size. Citizens ought to have a say about the criteria that combine to make a parcel suited to be a park. The various types of parks will abide by different criteria:. Moreover, different criteria will also carry different weights in the overall suitability determination. Again, the local population should be dictating how much importance (or weight) to attribute to each suitability factor.

economic development

The economic development component of a CDP would start with the specification of current and future economic profiles, followed by a goal statement that would lead to an implementation strategy that would include an assessment of where different types of economic development could and/or should take place. Parcels can be more or less suited to house economic activities, depending on several factors, such as contiguity to population centers, accessibility to highways, presence of rails, etc. A suitability map in this arena would showcase what type of business would "make sense" in different parcels of the city.

housing suitability

In the area of housing, the CDP would first include the inventory of the current housing stock, and an assessment of demand for different types of housing to determine the housing needs of the town. After setting community goals, a housing suitability map would be produced to show how suitable each parcel is for housing units of various types (single-family, multifamily, etc.) and whether the housing should be specifically dedicated to the elderly or to residents with special needs. Once again, a public participation process would be used to identify what criteria should play a role in the determination of suitability for housing as well as how much weight these criteria should carry in the overall suitability.

transportation issues

EO418 is rather vague of how to deal with mobility despite its being a major determinant of development, for good and for bad. Economic development is frequently dictated by accessibility as is housing. Conversely, housing and businesses also generate the need for additional transportation capacity when the infrastructure that supported the initial development becomes inadequate to support the subsequent growth. In a nutshell, though, the transportation element of a CDP should identify transportation improvements that can be implemented within one to five years.

WORCESTER E.O. 418

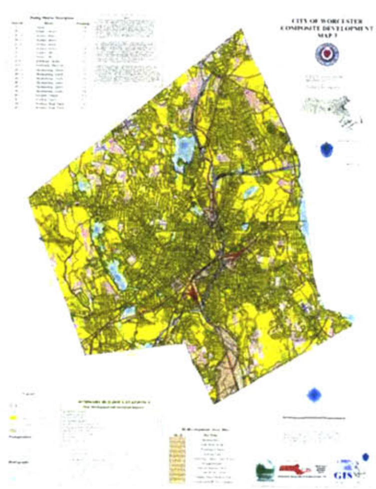

Thanks to the resourcefulness of WCPC director, Prof. Rob Krueger, we were asked by the City of Worcester to act as consultants on the EO418 project through the Sustainable Cities Research Group (SCRG) at WPI, which included. beside Prof. Krueger and myself, a local professional Architect and Planner, named Daniel Benoit. The SCRG received the $30,000 state grant, which was administered through the local regional planning commission, the Central Massachusetts Regional Planning Commission (CMRPC). By the time we were hired by the city, the CMRPC had already conducted some of the preliminary work and had produced a buildout analysis by land use zone.

Part of the required asset inventories were available to us thanks to previous WCPC projects, such as a "blue-green amenities[309]" inventory and a brownfield inventory[310]. In addition to the existing datasets from previous projects, to fulfill our mission we conducted four additional undergraduate interdisciplinary studies, through the WCPC, starting in the spring of 2003[311]. The first of these EO418 projects consisted in a review of existing plans and resulted in a database of plans that made it possible to search and retrieve the principal plan documents[312].

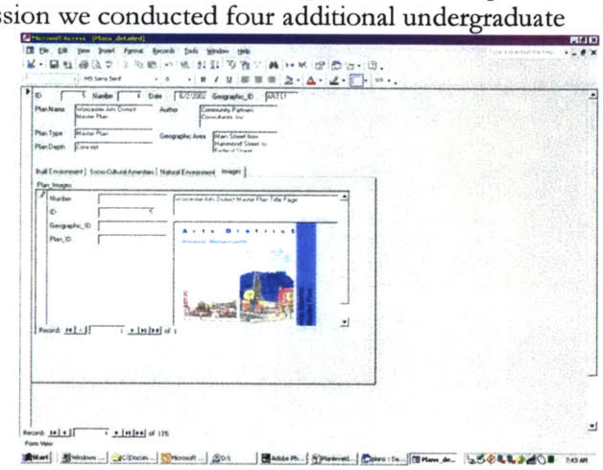

REVIEW OF EXISTING PLANS

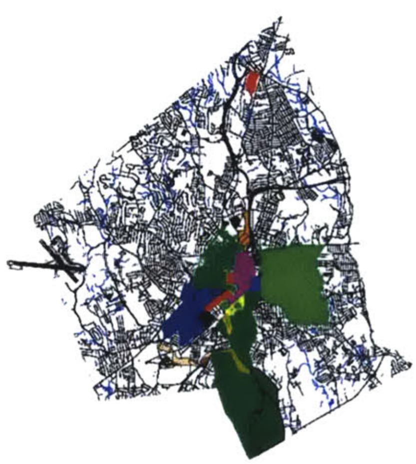

More importantly, our plan review identified ways to capture the considerable amounts of information that plans contain. There is a lot of data about demographics, markets and traffic that is included in many plans, be they at the level of a master plan or at the scale of a site plan. A strategic plan, such as the Community Development Plan required by the EO418 process, naturally results in the proposal of a series of actions based on the five tools described by Schuster *et al.*[313]. The more concrete and specific site plans are instead negotiated between the developer and the town, through a planning board (or the like) that cajoles the proponents into adapting their plans to the dictates of the strategic master plan or the visionary Community Development Plan.

What became clear in the course of this study – but also in similar experiences I had in Cambridge[314] and elsewhere – was that plans represent a missed opportunity for a city to acquire and <u>retain</u> useful urban

309 Dow *et al.*, 2003.

310 Bassa *et al.*, 2003.

311 Brown and Groeli., 2003; Farmer *et al.*, 2003; Hamir *et al.*, 2003; Jajosky *et al.*, 2003.

312 Brown and Groeli., 2003.

313 Schuster *et al.*, 1997.

314 Gage *et al.*, 2003.

information that can be reused by itself outside of the scope of the plan, or in the evaluation of a single plan over time, or in the comparative evaluation of the net effect of multiple concurrent plans. Plans will almost always contain up-to-date information about land use, building use, traffic, jobs and other economic indicators of supply and demand.

[parameterizing plans]

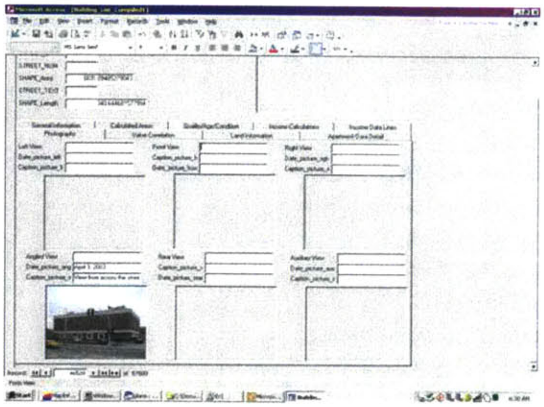

To be able to make use of the data included in plans, these data would need to be parameterized into atomic chunks that could be reutilized, aggregated and analyzed in any combination from then on. In particular, the current use of each floor of each building affected by the plan would be extremely useful, whether the plan became a reality or not. The main upshot of any plan is that buildings will be created, destroyed, modified and/or converted to different use. In turn, the net square footage dedicated to each use will dictate the types of people who will frequent these buildings and hence the type and quantity of traffic that the new uses will generate, the number of jobs that will be created, and the municipal services that will be required.

[before and after building use]

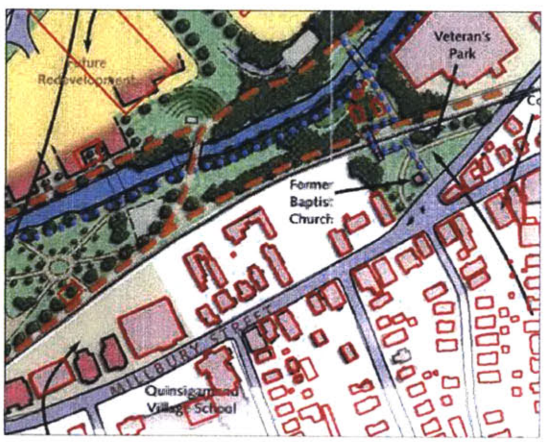

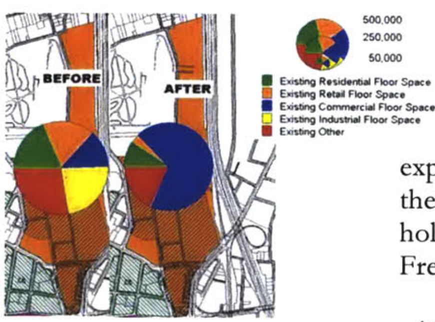

With a proper method of parameterization, each plan could produce the "before and after" in terms of square footage of each building's floor dedicated to each of the main uses (retail, office, industry and residential for example). Using SIC (or NAICS) codes will provide the finest grain of business use, whereas MacConnell codes would specify the uses for parcels of land. An information-aware planning ordinance could force all developers to produce a detailed, building-by-building inventory of current uses for all edifices within the boundaries of the plan area.

The uses forecasted by the plan for each of the "new" buildings would be thus comparable to the pre-existing uses and a net square footage could be calculated for each of the uses contemplated by the plan. Using standard metrics, it would be thus fairly straightforward to estimate the trips generated by these uses, as well as the jobs that would be created by the businesses that the plan expects to move into the newly created spaces. These new spaces will satisfy some "demand" for the various uses (like housing, retail, open space, etc.) that will thus diminish the overall demand for them vis a vis other plans elsewhere who are also expecting to satisfy those same unmet needs of the local population. Thus the synergies between different plans will come to the fore and a more holistic approach to town planning will become possible, as envisioned by Frederick Law Olmsted Jr. in 1913[315].

Two possible outcomes could result from the proposal of the plan with all of the accompanying "before use" data: (a) the plan is implemented and some or all of the forecasted uses become a reality, or (b) the plan is

315 See quote on p. 159.

[lasting benefits of "before" data]

never put into action and target area undergoes the "natural" evolution that the place was destined to experience in the absence of planning efforts. In either case, the before data would be an extremely useful addition to the requirements for submission. If the plan is implemented, then the before data can be compared to the "after" a few months or years down the line to see if the forecasted benefits of the development have been achieved as planned. If the plan was rejected, then the before data would still be useful for any other town activity that will be planned for the same geographical area in the immediate future. Building use, in particular, would be of great interest to many types of municipal activities and some form of use recordkeeping would be highly desirable.

In addition to uses, the before-and-after datasets could include traffic (volumes and turns), jobs (by occupation code and NAICS), open space (acres per 1000 inhabitants) and several others parameters that could help evaluate the "success" of a development as well as its impacts on the surrounding areas and on the outcome of adjoining plans.

[synergy between plans]

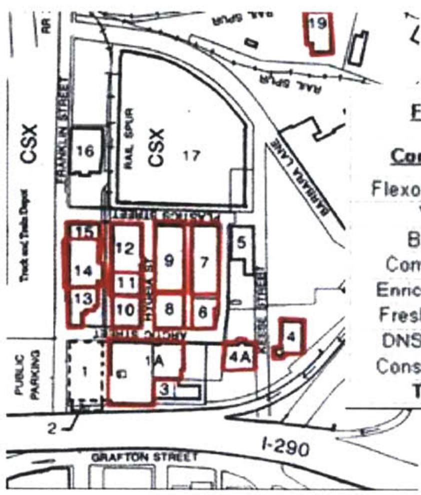

Franklin Science Park

Company	Area (sq. ft.)
Flexographics	14659
Visi	14659
Burton	14659
Come Play	91977
Enrich Foods	16000
Fresh Foods	14350
DNS Printer	14350
Consolidated	8180
Total	**188834**

As an example of the possible <u>positive</u> synergy that could result from such detailed and uniform knowledge across multiple plans, we explored the feasibility of relocating some industrial and commercial enterprises from the site of the proposed Franklin Science Park to aid in the proposed revitalization of the South Worcester Industrial Park. Due to the proposed demolition of 283,185 sq. ft. of industrial space at the Franklin site, eight businesses would need to find a new home in the South Worcester site where 218,500 sq. ft. of industrial space are planned to be created. Detailed knowledge of floor space usage building by building enabled us to map these eight businesses to four new buildings in just a few minutes.

On the flip side, one new development could engender additional traffic that ought to be taken into account in the planning of a nearby complex. A successful project will instantly affect other plans that are contiguous to it, but it may also affect sites that are on the opposite side of town, if the target audience of those plans are the same. If one development is supposed to capture the demand for new high-end retail shops, another similar development anywhere else in town would be deleterious to the financial viability of such an enterprise no matter how far.

South Worcester Industrial Park

Building	Area (sq. ft.)
Building 1	82500
Building 2	64000
Building 3	37000
Building 4	35000
Total	**218500**

The approaches to the parameterization of plan data that we experimented with in this project will enable city planners to engage in "concurrent holistic planning" instead of being trapped into the more typical, and almost inevitable, discrete mode of planning that is the norm in most cities and towns.

lessons form plans

The revelation during this project was that, since plans are catalysts of information, they could be treated as vehicles for the accrual of city knowledge. A parametric approach to the structuring of plan data would allow the city to accumulate useful data from each plan submitted for approval. Even if the plan was eventually botched, the data it furnished would be nonetheless acquired by the city for re-use in any of its activities[316].

[building use by floor]

[free traffic data]

[supply and demand basins]

Plans could produce gradual inventories of building use, potentially even floor-by-floor. They could provide current traffic data for affected intersections[317]. They could even provide useful macro-economic regional information about unmet "demand" as well as exiting "supply" for specific industries, services, labor or housing.

316 This approach would fit well with the Planning Analysis and Modeling Markup Language (PAMML) proposed by Singh (2004).

317 As was proposed by our project for the City of Cambridge(Gage *et al.*, 2003) which suggested to require developers to submit traffic impact studies in formats that would lend themselves to instant geo-coding and acquisition.

Buildout Analysis

parcel-level build-out

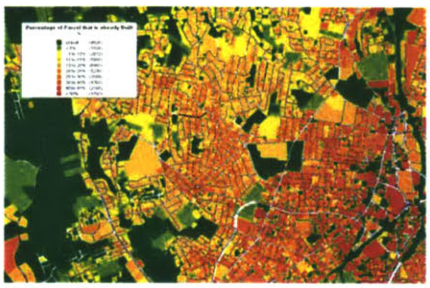

Being believers in the "atomization" of information, we conducted a parcel-level buildout analysis that calculated the maximum buildable area in each parcel, after removing all regulatory and physical limitations. The total surface area of a parcel was reduced first by the amount taken up by water bodies, inclusive of the mandatory 30' buffer. Then, a fixed 15' ring of land was removed to approximate the setback requirements for the lot. Any area that was already built (e.g. map to the left) was also deducted from the total to give the net buildable surface.

The resulting "number of developable square feet" did not account for sites where construction would be difficult or nearly impossible due to steep slopes. Since there is no law to prevent construction on such rough terrain, we treated sloped parcels separately. In essence, by ignoring slopes we could visualize a maximum buildout scenario, with the largest amount of land open for development. We then produced a more realistic case where slopes were allowed to preclude some construction. We calculated a sort of "maximum growth index" by calculating the percent of expansion that each parcel could undergo from the current situation to the maximum possible buildout. Note that this analysis does not take current zoning into account, except for the provision of a generic setback buffer around the perimeter of each property.

The map below shows in redder tones the areas where more growth is physically possible due to the morphology of the terrain and in the absence of zoning regulations to keep development in check.

The gray areas above (and left) indicate existing parks or bodies

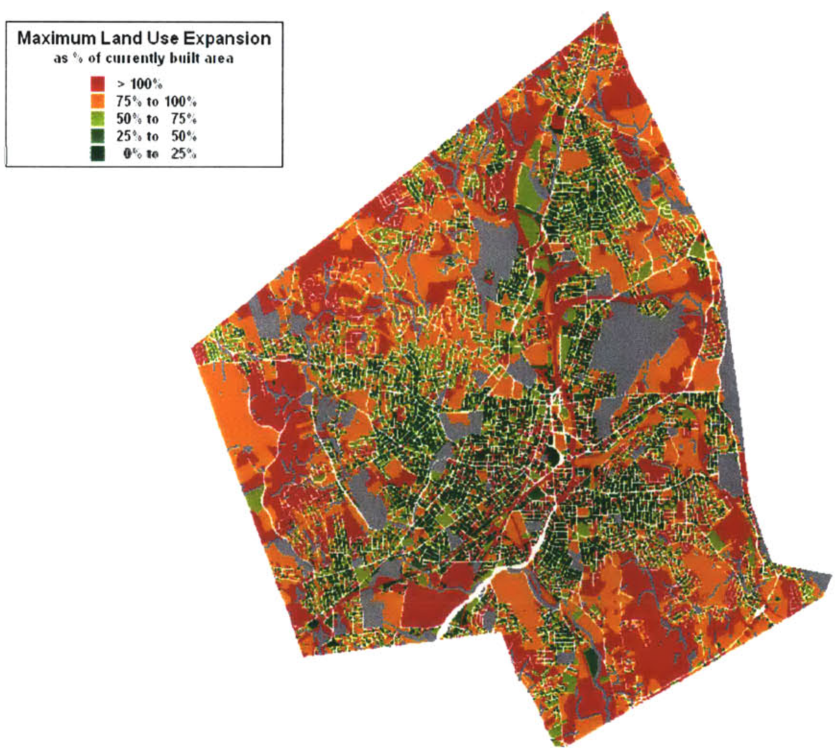

of water that are not going to be developed at all. When one zooms into this map, it becomes possible to distinguish the parcel-level detail that this approach affords. With a complete City Knowledge system, a more sophisticated differentiation among parcels would become possible, thus allowing this analysis to focus only on the truly realistically developable tracts.

Such an approach opens up the possibility of introducing zoning as a variable in this analysis to experiment with "what-if" scenarios for exploring how much development would be allowed under different zoning regimes.

[parametric zoning]

A parametric approach to zoning would make the process of recalculating offsets, FARs, maximum footprints and the like quite effortless. Reclassification of a parcel from one zone to another would simply result in a recalculation of the amount (and type) of development that would become possible after such a change. Parcel level knowledge of the terrain (slope and water primarily) may even suggest the possibility of making zoning a dependent variable, determined by suitability, which is discussed later in this chapter, and by accessibility to transportation, which is the topic of the section that follows.

lessons from buildout analysis

[pseudo-City Knowledge situation]

In order to carry out the four main tasks that follow, as required by EO418, we relied on existing studies and data that formed an adequate factual foundation for our analyses. Our successful use of less-than-ideal datasets reinforces how useful a fully-developed City Knowledge system of plan-ready information would be. Thus, one of the principal recommendations that emerged from our work in Worcester was that the city should invest into the creation of virtuous mechanisms that will produce "plan-ready" information, using a middle-out tactic according to the tenets of City Knowledge. In other words, the availability of a more complete and up-to-date foundation of knowledge would have allowed us to produce an even finer analysis of the challenges and opportunities facing Worcester in the years to come.

[exploratory vs. optimal]

Parameterization of the zoning regulations would further enhance the capabilities of a comprehensive municipal information system. The confidence and ease with which one can re-run exploratory analyses, reduces the need for obtaining "optimal" solutions.

TRANSPORTATION ACCESSIBILITY[318]

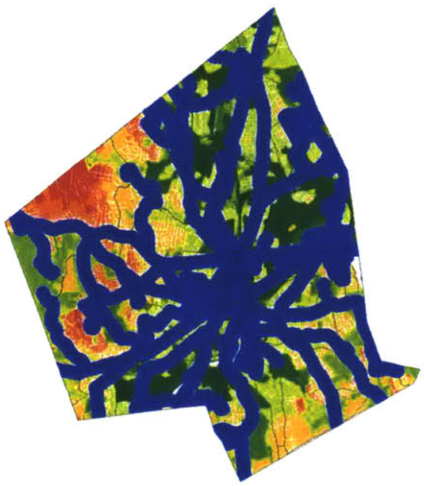

The purpose of the transportation section of EO418 was to determine where the transportation infrastructure should be improved. We found that the best way to achieve this goal was to simply determine how accessible each parcel is. To do this we used existing traffic volumes to identify the major arteries that crisscross the city. We then attributed higher accessibility to parcels that were closer to highways (up to 2 miles) and major arteries (up to ½ mile away).

Separately, we also analyzed accessibility with public transportation, attributing higher scores to parcels within a ½ mile of a bus route (blue buffers at left) and an even higher score if they were within ½ mile radius of a bus stop. Each parcel would thus receive a cumulative accessibility score that represented how easy (green in map below), or hard (red) it was to reach it with either private or public means of transportation.

This accessibility information was used in all further suitability analyses, since transportation plays a major role in the choice of land use, be it for residential, commercial, industrial or recreational purposes. All sections that follow incorporate, in some measure, the accessibility scores at the parcel level.

lessons from traffic accessibility

The main lesson from this section was that transportation data ought to be more easily accessible and usable. We had to resort to manual data input from paper maps just to get volume information for major arteries in the city. Obviously the cost of our manual data entry – or even worse the cost of field data collection – would be spared if the data were acquired through mandatory impact study submission requirements attached to developers' plan proposals[319].

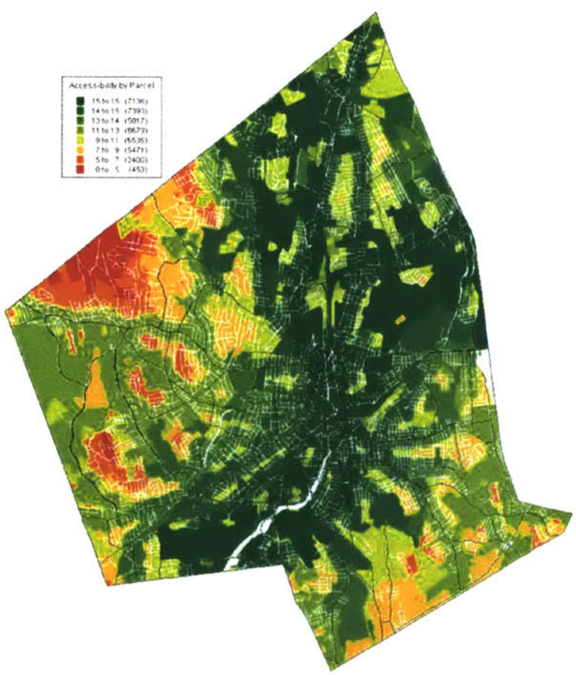

---

318 This section and the next (on open-space) are based on Farmer *et al.*, 2004. All figures come from that report.

319 As proposed by Gage *et al*, 2003 and implemented by the City of Cambridge and as discussed earlier in this section on page 140.

OPEN SPACE ANALYSIS

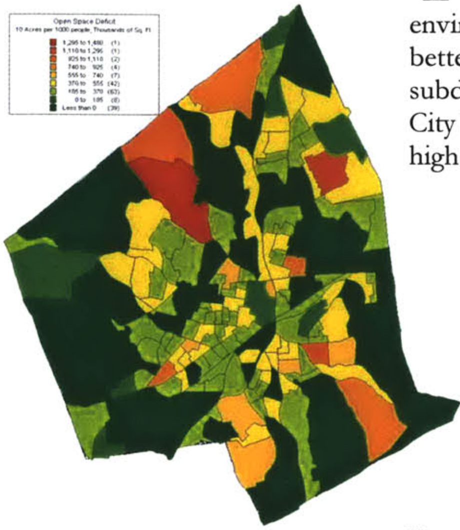

The next step in the EO418 process deals with environmental preservation. The goal here is to identify what parcels are environmentally precious and ought to be preserved and what parcels are better suited to be parks instead of being turned to other uses. We subdivided all open space into three main categories: Conservation Land, City Parks, and Neighborhood Playgrounds. The criteria used to assign a higher or lower open space suitability to a parcel were the following:

- slope (not desirable for playgrounds, but good for conservation)
- presence of water (desirable except around children in playgrounds)
- deficit (whether a particular part of town was lacking open space)
- proximity to other OS (conservation land is better in large swaths)
- ownership (public is easier to convert to OS than private)
- size (playgrounds small, conservation big)
- value (the cheaper the better in general)

The open space deficit was calculated based on industry benchmarks like the de-facto standard metric of 10 acres/1000 inhabitants[320]. The map on the left shows in red the areas where there is a more pronounced need to

[open space suitability map]

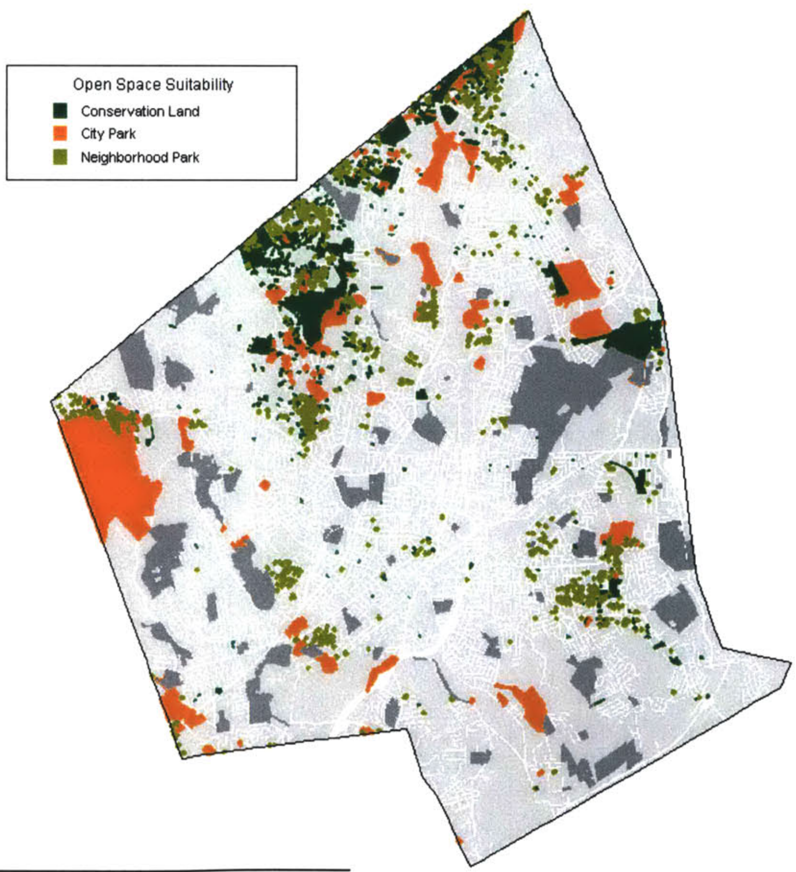

320 Recommended by the National Recreation and Parks Association (NARPA).

offer green amenities to the population. Not surprisingly, in the results shown below the most suitable OS parcels are concentrated around these areas of higher deficit.

feasibility and desirability

By including information about parcel value and ownership, these calculations reflect not only desirability but also feasibility. Nevertheless, these fact-based rankings and classifications represent only the first step toward protecting open space. The next step would be to act on these indications and begin the more painstaking process of turning them into reality. The five tools of government action mentioned elsewhere in this document[321] would suggest a combination of the following actions:

applying the 5 tools

1. Conversion of highly-suited, publicly-owned properties to OS;
2. Give incentives to owners of valuable parcels to convert them to public use;
3. Regulate the use of these parcels through zoning and other restrictions;
4. Educate and inform the public about OS conservation;
5. Transfer property rights of some crucial parcels of land to some other less precious parcels (better if the latter are publicly owned);

Once the most "actionable" parcels are identified, a visual survey of the sites ought to be made to ascertain their true potential as open space.

lessons form open space analysis

[exploratory filters vs. perfect match]

In the end, if done well, the McHargian analysis[322] of Open Space suitability can be a great time-saver and represent an excellent filter that can quickly lead to the most desirable and feasible parcels to target for protection and/or recreation. Once again the ease with which rough filters can be applied removes the need to more precisely identify the "most suitable" parcel up front. Instead the tool gets us close enough to allow more nuanced decisions to take place.

321 Schuster *et al.*, 1997.

322 McHarg, 1969.

HOUSING ANALYSIS[323]

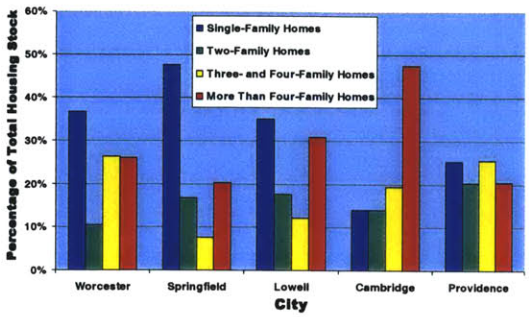

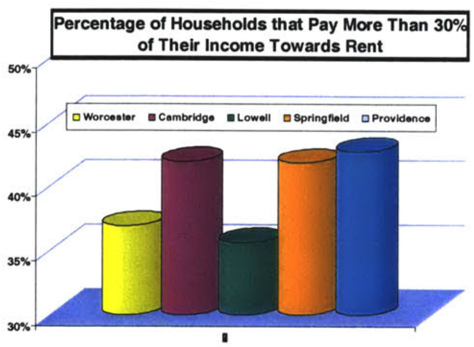

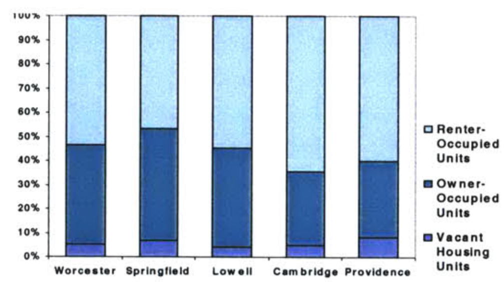

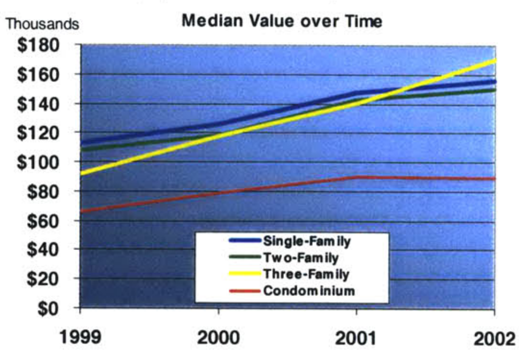

The primary stated purpose of EO418 is to assist communities in addressing the housing shortage, so the housing suitability map is an essential deliverable of this process. The housing analysis is meant to start from a thorough housing inventory, which was not available for the city of Worcester. We were nonetheless able to develop a housing market profile based on a recent study conducted by a local consulting firm on behalf of the city[324]. Using the MacConnell coding scheme[325], one could try to roughly classify the housing typologies by parcel. Unfortunately, the Worcester Assessor's data does not explicitly describe land use, though it contains a useful description of the building (*blddesc*). In all, the Worcester parcel dataset categorizes buildings with 103 different descriptors. Single family homes are the most numerous (23,075), followed by triple-deckers (5,046), then by two-family homes (3,780). Compared to other major New England cities, Worcester seems to have more single-family homes and more triple-deckers, with relatively few duplexes and a moderate amount of multifamily dwellings.

Some of the building categories betray a lingering bias toward taxation (e.g. *exempt* classifications). It would be worthwhile to make a renewed effort to standardize land use codes using an improved version of the MacConnell codes and/or MassGIS's Land Use Zone crosswalk table[326]. Housing inventories would be instantly available at any moment if a set of building use codes was developed and attached to the spatial object through a unique building code. In parallel, a land use coding scheme should also be used to account for differences between unbuilt parcels.

Despite these difficulties in use classification, we were able to analyze the housing suitability for two fundamental types of dwellings: single family and multi-family. Moreover, we also considered the suitability of parcels as sites for housing for the elderly or for people with special needs.

From other sources[327], we had already surmised that Worcester homes are more affordable and that there is a decent demand for rental units. Triple deckers recently surpassed single-family homes as the dwelling with the greatest median resale value, which suggests that high-density living is still a desirable quality in this city.

323 This section and the next (on open-space) are based on Hamir *et al.*, 2004. All figures come from that report.

324 RKG Associates, Inc., *Housing Market Study*, 2002.

325 See for instance http://www.mass.gov/mgis/landuse_stats.htm.

326 *Idem.*

327 RKG, *op. cit.*.

opportunistic choice of criteria

To simplify the visualization of the data, the suitability analysis for housing was divided into three sub-categories:

1. Single-Family
2. Multifamily
3. Special Needs and Elderly Housing

The principal criteria that were brought to bear on the housing suitability calculations were:

elderly housing

- lot size (smallest and largest more suited for multi)
- accessibility (multifamily and special needs require better access)
- deficit (for elderly housing only)
- proximity to OS (more important for high density multifamily)

As was the case for the housing and transportation analyses, the choice of criteria is based in part on the availability of data to support the inclusion of each criterion. Spatial metrics that can be directly extracted from GIS – like size and proximity – are always available with a modicum of manipulation. Other, more complex proxies – such as accessibility – may require considerable computational efforts[328]. Deficit was probably the most sophisticated indicator that we were able to calculate based on available datasets. We did not pursue even more complex suitability factors, such as the homogeneity of the housing stock in the neighborhood, since such measures are impractical, if not utterly impossible to even approximate with the typically available urban datasets.

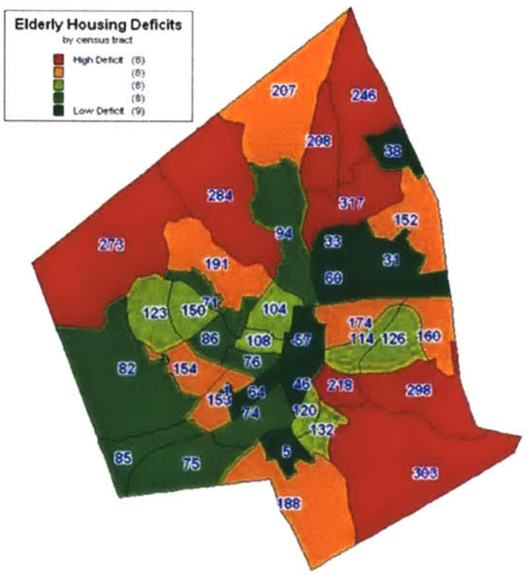

Elderly housing was allotted according to the deficit that was calculated as the actual number of people over 65 in each of Worcester's census tracts minus the number of beds in that tract that are already dedicated to elderly people[329]. One of the complications was that these data were based on census tracts instead of parcels. The scalability of the units of analysis is thus very important when performing these analyses.

When making decisions about how to allocate suitable uses to the various parcels based on the results of these analyses, housing would probably be given a lower priority than Open Space and Economic Development, as will be discussed below.

328 As briefly discussed on page 143.

329 From the Massachusetts Housing Finance Agency *Housing List*.

suitability analysis as a tool

Another thing to keep in mind – and this applies to all suitability maps – is that these maps provide a "portfolio of choices" and are not meant to point to the optimally best choices. The colors show only the highest suitable use. The second highest use may be right up there, but only slightly lower than the topmost. On the other hand, the unsuitable uses should be fairly clear, so these tools are primarily useful to identify sets of suitable uses versus unsuitable ones. Once a particular parcel is assigned one

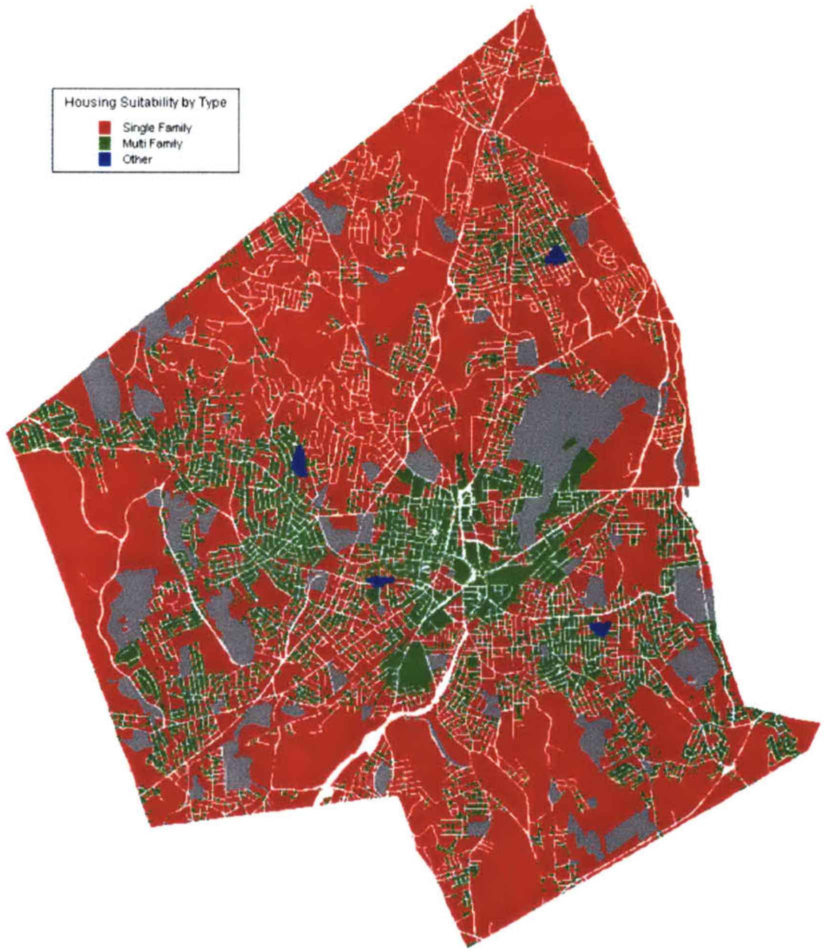

of the more suitable uses, surrounding parcels may become suddenly less suited for their topmost use. For instance, if a particular parcel was selected to be the site of a new elderly housing project, any additional parcels in the same census tract that had been earmarked for potential elderly housing will need to be re-assigned to the next highest suitable use.

Far from being the silver bullet that will make planning a largely automated enterprise, suitability analysis does provide a powerful tool for screening potential uses in the early stages of planning. It also provides enough intelligence to guide a possible rezoning effort and it helps to focus our limited resources on important situations that may be the target of incentive programs.

housing suitability map

In the bigger scheme of things, if we leave current zoning out of the picture, the suitability map shows multifamily houses (green color in map below) clustered around downtown and along major arteries, with single-family homes filling the in-between spaces (red color).

lessons from housing analysis

This section showed how much more useful it would be to have a true City Knowledge system in place so that we would not be forced to resort to opportunistic choices for our suitability criteria. It also reinforced the benefits of adopting suitability analyses as precursors to planning decisions, more as a way to weed out unsuitable locations than to truly zero in on the absolute best options.

ECONOMIC DEVELOPMENT ANALYSIS[330]

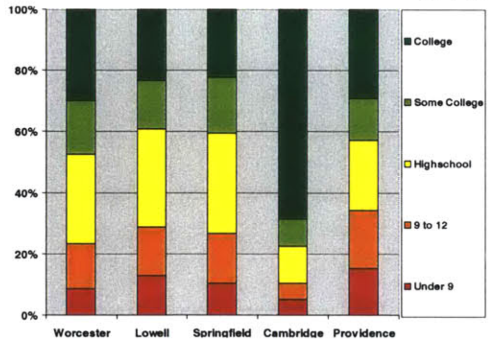

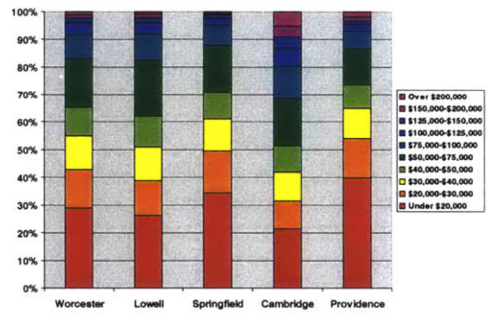

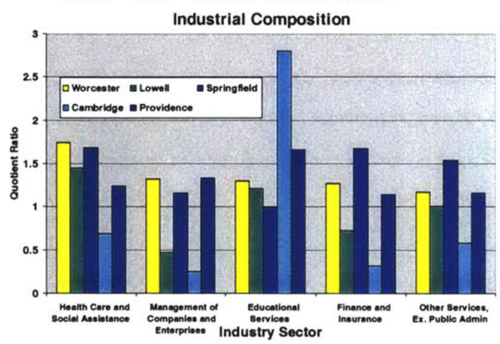

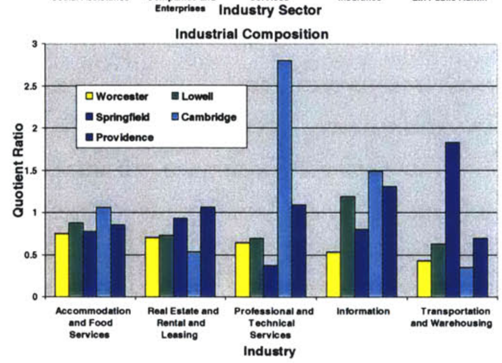

The final aspect covered by the EO418 process is Economic Development. Our students compared Worcester to the Massachusetts cities of Lowell, Springfield and Cambridge as well as with the capital of Rhode Island, Providence. Basic workforce statistics were used to determine the human resource pools that these cities are able to dip into. Worcester has the second highest level of education after Cambridge, which is very encouraging for future business development.

The income distribution of the populations of these cities was also examined and once again Worcester faired comparatively well. Cambridge is undoubtedly the leading city in this comparison group, but Worcester is well equipped too.

Our students then conducted a quotient analysis[331] to bring out the strengths and weaknesses that Worcester exhibits in its business sectors, as compared to other Massachusetts cities.

Based on the NAICS codes[332], the leading industries in Worcester turned out to be Health Care, Management of Companies, Educational Services, Finance and Insurance, and Other Services (not public administration). Yet, despite the fact that these are the best sectors of Worcester's economy, the city only leads all others in health care and in the management of companies. All of its colleges notwithstanding, Worcester is only third in the field of educational services after Providence and Cambridge, which far outdoes all others.

The trailing industries in Worcester are Accommodation and Food, Real Estate Rental and Leasing, Professional Technical Services, Information and Transportation and Warehousing. The fact that Cambridge and Providence, who are leaders in educational services just ahead of Worcester, are also leaders in professional and technical services – where Worcester lags even behind Lowell – suggests that Worcester may be missing an opportunity in this area. Even more acute is the shortcoming in the realm of Information, where Worcester sits at the bottom of the heap despite its educated workforce and despite WPI itself.

In the end, a detailed economic profile based on 21 industry sectors (2-digit NAICS) suggested that Worcester should focus its economic development efforts on declining industries that are complementary to the strong sectors. For instance, the declining manufacturing base in Worcester could be reconstituted in part by providing incentives for the development of biomedical manufacturing plants that could ride on the coattails of the very successful biomedical and health sectors that are thriving in the city.

330 This section and the next (on open-space) are based on Jajosky *et al.*, 2004. All figures come from that report.

331 Jajosky *et al.*, 2003.

332 US Census Bureau. *2001 MSA Business Patterns* (NAICS).

In developing these insightful recommendations, we were yet again limited by the availability of good local data. Business activity, like many other phenomena, is poorly represented at the local level, especially at the scale of parcels. Even when datasets containing addresses are obtained, they only provide a partial and potentially misleading picture of the actual economy of a place. Most data are available only at the census-tract level, if not at just the overall municipal scale.

economic development suitability map

Thus, when it came to analyzing the suitability of locating different industries in different parts of town, the data had to be extrapolated from the tract-level sets available[333]. Despite the difficulties of spreading these data across individual parcels, the resulting suitability map seems to capture the essence of a rational siting strategy, which is based on the following main parameters:

- accessibility to rail (important especially for industry)
- accessibility to major arteries(important for retail and commercial)

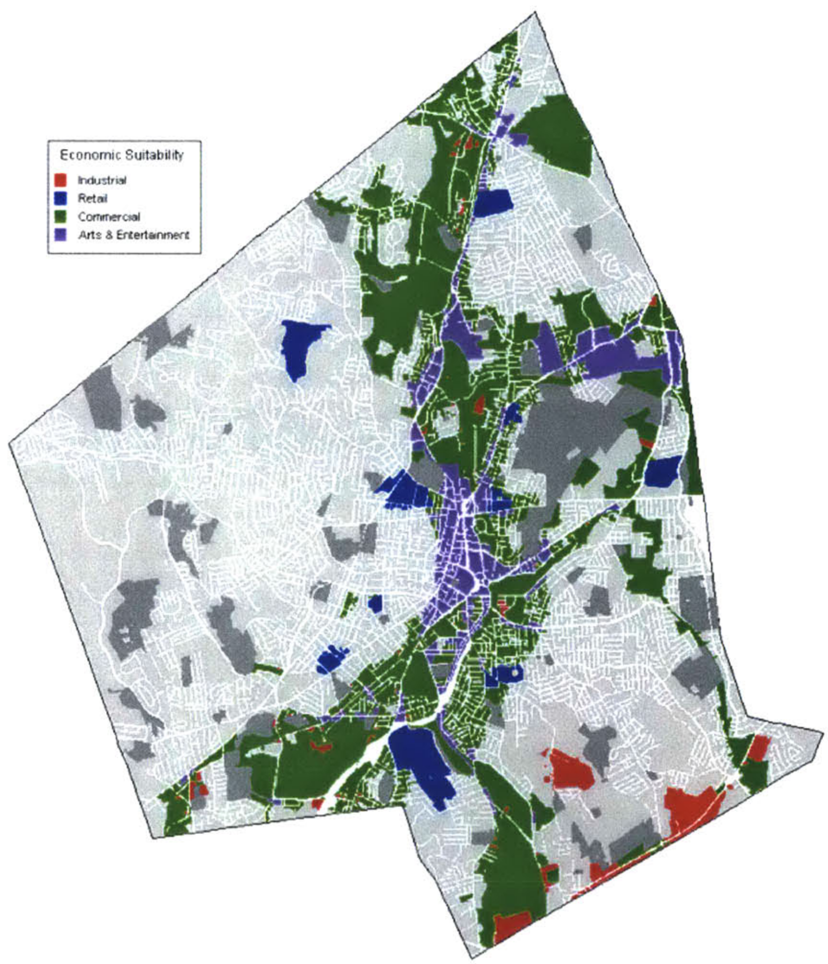

[333] Massachusetts Division of Employment and Training. 2003. ES-202 data.

Suitable industrial sites (in red) are strung along the main roadways and railroad tracks, especially in the Southeast part of the city, where Route 20, 146 and the Massachusetts Turnpike are located and through which the main north-south rail trunk between Worcester and Providence, along what used to be the Blackstone Canal that connected the two cities. Retail (in blue) is also laid out in big chunks, not far from the principal roadways, but with a more equitable distribution in the various neighborhoods, where shoppers live. Parcels that are highly suited for commerce (green) create a rather thick buffer around the main arteries whence goods and customers would reach the businesses.

[arts district prediction]

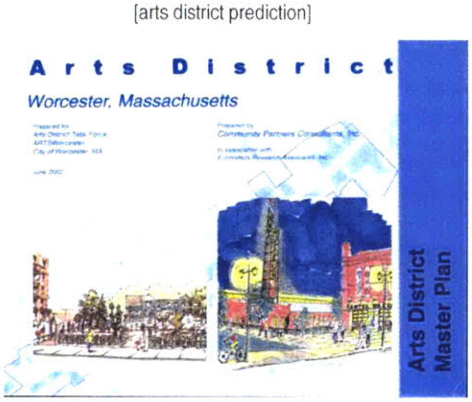

lessons from economic analysis

As an example of the siting potential for specific NAICS sectors, the economic suitability maps show the best locations where Arts and Entertainment businesses could be established (purple in maps above and below). It was rather amazing to see how this computational analysis ended up mimicking real life in an astonishing manner, by picking as the prime locality for this type of activity precisely the neighborhood where the "Arts District" is being planned, yet doing so based on a purely numeric algorithm predicated on objective physical parameters.

The big issue here was not so much the siting suitability for the various types of businesses, but just which businesses it would be preferable to try to attract to Worcester. Our quotient analysis achieved just that by using census data. Of course, a full-fledged City Knowledge system may afford us much more latitude in the locational choices, yet the problem here is much more regional than local, and it is not so much spatial as economical.

[regional planning]

So the lesson here is that a uniform treatment of datasets across town boundaries would enable Regional Planning Authorities (RPAs) to coordinate macro-level economic development decisions.

COMMUNITY VISIONING

The relative weight of the various factors that participated in our suitability calculations can be adjusted to account for differing perspectives on what is more important or more appropriate. The EO418 process encourages public participation in the formation of the sets of important criteria and in the assignment of relative weights between all criteria in each of the three suitability categories. As the overall suitability across the three domains is ascertained, once again the community could be given the chance to voice its opinion and sway the importance of the three aspects in one direction or the other. As part of our Worcester EO418 study, we conducted three visioning exercises involving a cross-section of individuals representing a broad spectrum of views. Before the EO418 process even began, Prof. Krueger had already conducted an international workshop entitled "Envisioning Worcester's Future" that provided a foundation for our subsequent work.

A Participatory Planning Tool

Having spent many hours to develop the initial sets of suitability maps, we quickly realized that we needed some way of streamlining the suitability calculations in order to rapidly incorporate citizen comments in the iterative process of successive refinements that is needed to create a consensus over the most appropriate land use for the future of the city. To do this, we developed a computerized tool that would interact with the three databases we had developed for Housing, Open Space and Economic Development and allow a dynamic modification of the relative importance of each set of criteria associated with each of the three main areas of study. These suitability databases were directly linked to GIS maps to instantly produce a visual rendition of the suitability produced by the citizen groups.

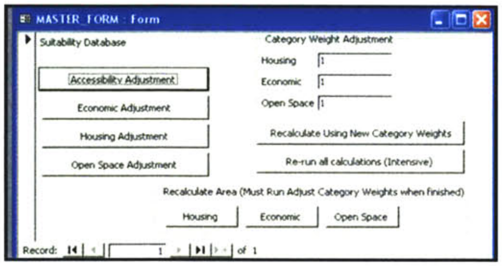

As was summarily discussed above, we identified a list of criteria culled from planning texts and interviews with professionals, but filtered by how feasible it was to obtain the data necessary to adequately "measure" each criterion, directly or by proxy. After the criteria were selected and filtered, we were left with 7 factors for the analysis related to the 3 main types of housing considered, 7 factors for the assessment of economic development suitability across the 4 broad business segments we analyzed, and 8 factors for the identification of ideal sites for the 3 types of open space.

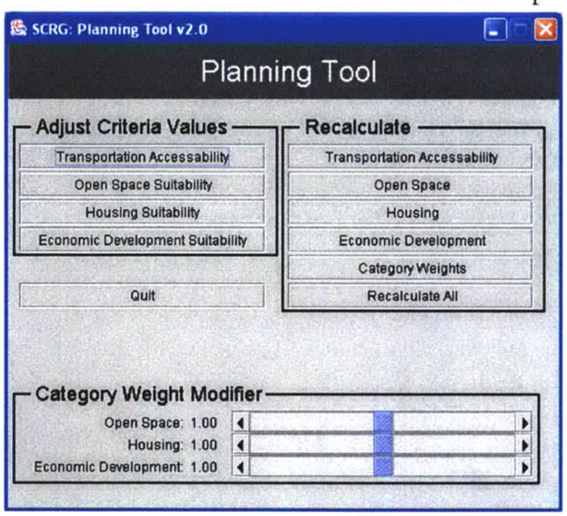

In essence, we set up three independently-modifiable lookup tables wherein we recorded the scores that the focus group participants assigned to each of the relevant criteria. The first completed version of the tool was built using Microsoft Access® for both the calculations and the data storage. To calculate the suitability after modifying the data, a number of macros would run a sequence of queries that would cumulatively produce the end result. The databases include data that have been pre-processed to reflect some of the useful measures dictated by the analyses. For example, a number of spatial queries involving buffers of varying widths were used to assign a proximity value to each parcel vis a vis parks. These spatial queries would typically be at our disposal without much additional information, since they are based on the geographical characteristics of the GIS objects.

The next generation of the system was written in Java™ to add more power and flexibility. Adding a new criterion is as simple as adding another row in the criteria lookup table, and separately providing the data necessary. Some criteria, however, may be based on knowledge collected from government databases – such as the US census, or the city's real estate tax assessments – so these would need to be manipulated to suit the desired analysis. Even though our data are stored using Microsoft Access®, this program is very flexible and is capable of working with any other standard database type that supports SQL.

lessons from public participation

The interactive visualization tool allowed us to conduct our visioning sessions in a more constructive and efficient manner, which in turn allowed us to progressively refine the delicate balances of weights and scores that led to the development of our final composite suitability map.

The visioning sessions made clear the instant value of the suitability analyses that would make them immediately cost-effective with today's tools and knowledge. Moreover, these sessions added credence to the potential of City Knowledge as a foundation for citizen empowerment through education and information dissemination.

COMPOSITE SUITABILITY MAP

After each of the four main areas of study have been analyzed, a composite suitability map was produced to visualize the net effect of all of these suitabilities taken together. In creating this combined map, the relative importance of each land use category needs to be decided whenever different uses are competing for the same parcel. During the visioning sessions such tacit hierarchies were made explicit. In the map below, we placed the open space parcels suitabilities on top of the other layers, giving them priority over all uses. Next we placed the economic development set of suitable parcels and finally, at the bottom of the land use food chain, we relegated the housing suitabilities, which would fill in all the land gaps between the higher priority parcels that would be better suited for recreation or for business.

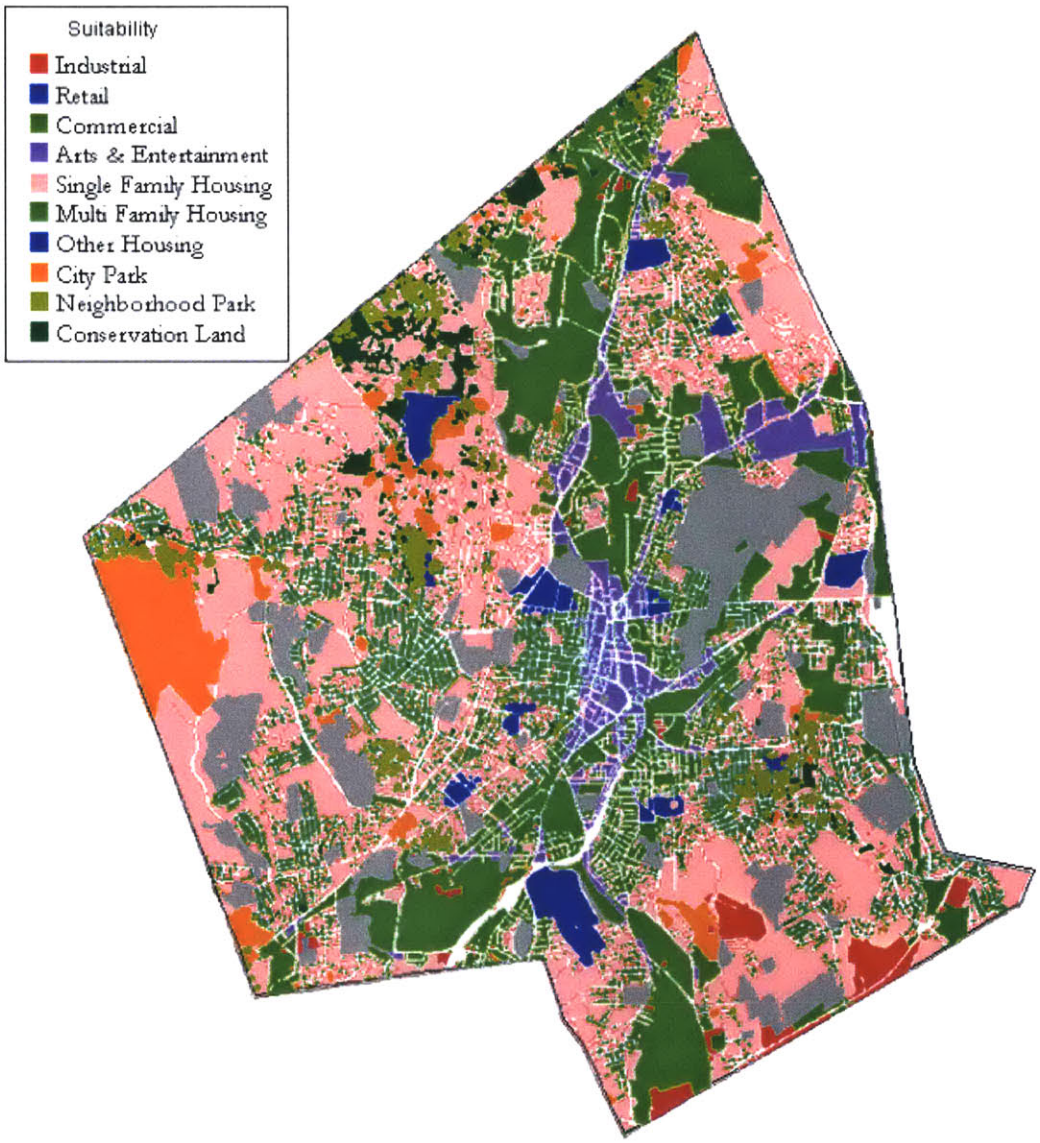

LESSONS FROM MY OWN BACKYARD

Working on the Worcester EO418 project enabled me to specifically focus my attention on the impact of City Knowledge on municipal planning. I awakened to the fact that plans are untapped repositories of up-to-date information, but unfortunately they are still treated as nothing more than "documentation", made up of paper reports and paper maps. As is the case in other departments, the data that are received by the planning department as part of a plan are still considered like accessory documentation instead of being treated as crucial, up-to-date, free-of-charge, retainable and reusable information.

plans as information

Our efforts in Worcester[334] and Cambridge[335] have shown how, by simply revising existing plan submission requirements, cities can begin to amass a considerable amount of current information on urban data such as traffic, land use and business activity. I expect more cities will revise their guidelines for the submission of plans to maximize the informational return of these routine planning processes.

buildings as atoms

[building IDs]

Another lesson that was emphasized in Worcester, though it had already become evident in Venice and elsewhere, was the desirability of acquiring and managing information at the level of the individual building. This type of information would have been very useful for our buildout, economic development and housing analyses and it would also have permitted the apple-to-apple comparison of adjoining or competing plans. The start of this effort would be the creation of a unique building ID to identify each single edifice in the city.

floors as atoms

[building typology and use by floor]

Once building IDs were assigned, we could begin to keep track of building typology (including architectural styles), as well as building use, which I recommend to be collected using the individual floor as the unit of measurement, as discussed earlier in this chapter[336]. It would be good to adhere to standards for the classification of uses according to either the NAICS (for business uses), or the MacConnell use codes (especially for undeveloped parcels)[337].

free traffic data

[volumes and capacities by segment]

A similar lesson was learnt with respect to our transportation analysis. If the road network had been appropriately subdivided and organized to provide segment-by-segment information about traffic volumes and capacities, and if a method similar to that adopted in Cambridge[338] was put in place to obtain "free" traffic data from developers' impact studies, the accessibility map would have been much more representative of the real level of access that is available to each part of town.

plan-ready facts and interactive values

When we tried to adopt the "separation of facts from values" approach[339], we gained the insights that lead to the development of the suitability maps and subsequently to the interactive visualization tool. It would be great to be able to dig deeper into the variety of additional criteria

334 Brown and Groeli, 2003.
335 Gage *et al.*, 2003.
336 Page 138.
337 See for instance http://www.mass.gov/mgis/landuse_stats.htm. Last accessed 8/20/04.
338 Gage *et al.*, 2003.
339 Hammond *et al.*, 1980, 1991.

[suitability criteria and values]

that may affect land use choices and to arrive at a comprehensive review of the datasets that could be used to inform the factual component of such decisions[340], leaving the "value" side open for consensual modification through the participatory process, augmented by the interactive tool that allows instant visual feedback and iterative refinement.

demonstration of value of plan-ready

The Worcester case highlighted the potential returns we could get if we were able to reliably tap into administrative databases or into other existing municipal knowledge bases. Partially developed versions of the plan-ready exhaustive municipal information system that I envision were used fairly successfully in the example illustrated herein. This bodes well for the full-fledged systems I propose.

[340] Tapping into such rich literature as that produced by the Urban Institute, with its *Catalog of Administrative Data Sources* and the like (Coulton *et al.*, 1997).

(page left intentionally blank)

PART IV:

"The enormous importance of such a City Plan Office as we have been discussing, with its elaborate, active and obviously costly human machinery for systematically recording these live ideas which form the real city plan, for interpreting them and for deliberately amending them, lies in the fact that without such machinery these functions are performed unsystematically, intermittently and very imperfectly by people whose principal interests and duties lie in other directions. Without it the actual set of ideas and purposes concerning probable future improvements and conditions which are really kept in mind in such a way as to have practical influence upon current decisions, is dependent upon the memory and personal equation of scores of different individuals, no one of whom has opportunities to be cognizant of the whole field or to keep in touch with all the other people."

Frederick Law Olmsted Jr., 1913
Proceedings of the
Fifth National Conference on City Planning, Boston
http://www.library.cornell.edu/Reps/DOCS/olmst_13.htm

THE ESSENCE OF CITY KNOWLEDGE

A TALE OF TWO CITIES: STRUCTURES AND ACTIVITIES

Kevin Lynch – a professor in the same MIT department where this dissertation was developed – who remains, despite his premature death, a role model for those of us who are interested in the phenomenological aspects of a city in relation to planning, once said:

> "The Fundamental problem is to decide what the form of a human settlement consists of: solely the inert physical things? Or the living organisms too? The actions people engage in? The social structure? The economic system? The ecological system? The control of the space and its meaning? The way it presents itself to the senses? Its daily and seasonal rhythms? Its secular changes?
> Like any important phenomenon, the city extends out into every other phenomenon, and the choice of where to make the cut is not an easy one".[341]

His resolution of this dilemma was simple and straightforward. As he put it:

> "[...] the chosen ground is the spatiotemporal distribution of human actions and the physical things which are the context of those actions [...]".[342]

My shorthand way of rephrasing Kevin Lynch's dichotomy is simply that cities are made up of two components: *structures* and *activities*[343]. Italo Calvino eloquently captured the complementarity of permanent structures and ephemeral activities by splitting in half the city of *Sophronia*, my favorite of his *invisible cities* (left).

" The city of Sophronia is made up of two half-cities. In one there is the great roller coaster with its steep bumps, the carousel with its chain spokes, the Ferris wheel of spinning cages, the death ride with crouching motorcyclists, the big top with the clump of trapezes hanging in the middle. The other half-city is of stone and marble and cement, with the bank, the factories, the palaces, the slaughterhouse, the school, and all the rest. One of the half-cities is permanent, the other is temporary, and when the period of its sojourn is over, they uproot it, dismantle it, and take it off, transplanting it to the vacant lots of another half-city. And so every year the day comes when the workmen remove the marble pediments, lower the stone walls, the cement pylons, take down the Ministry, the monument, the docks, the petroleum refinery, the hospital, load them on trailers, to follow from stand to stand their annual itinerary. Here remains the half-Sophronia of the shooting galleries and the carousels, the shout suspended from the cart of the headlong roller coaster, and it begins to count the months, the days it must wait before the caravan returns and a complete life can begin again...

Italo Calvino
"Invisible Cities", p. 63

structures

Structures include all of the "containers", environments, spaces and places that make up the physical, material city. These concrete components of the urban realm include buildings, parks, rivers, roads, trees, fire hydrants and everything else that's "out there" in our cityscapes. Structures lend themselves to "permanent" inventories since they change ever so slowly and can thus be captured once and for all through an initial cataloguing effort, only to be occasionally updated by intercepting administrative acts that signal the changes that do occur – however seldom – in the physical make-up of the city.

activities

My dissertation deals predominantly with the physical structural elements of municipalities, though *activities* are also discussed at length. These dynamic phenomena are more difficult to track because of their ever-changing nature[344], which makes it impossible to capture the information

341 Lynch, *Good City Form*, p. 48.
342 *Idem.*
343 Hopkins, 1999, p. 335 offers essentially the same breakdown in Figure 1.
344 See however Longley and Harris, 1999.

once and for all, and then only deal with small changes from that moment on – as is instead the case with our urban structures. Whereas information about *structures* can be maintained through a piecemeal "transactional" approach, which is less costly and more manageable on a day-to-day basis, *activities* need to be monitored periodically and regularly, thus they require more resources to produce "snapshots" of their status quo at a particular moment in time.

Having made clear the major *distinction* between permanent, physical "structures" and ephemeral, dynamic "activities", in the chapters that follow, I try to condense and organize the lessons that have been discussed thus far, to propose a provisional framework for what I call City Knowledge.

In the next chapter, I establish the *premises* for City Knowledge, that revolve around a paradigmatic shift of perspective on municipal information awareness. In essence, I propose the adoption of an "information-aware *modus operandi*" so that towns can begin to treat information as an infrastructural element, as essential as roads, sewers and electricity.

The third chapter introduces the *obstacles* that have hindered the spontaneous emergence of comprehensive municipal information systems around the world. In that chapter, I make a case for why now is the right time to overcome these obstacles and move toward full-fledged City Knowledge.

I then enunciate in the fourth chapter of this Part IV the *qualities* that would be embedded in a comprehensive City Knowledge system in chapter four. Each quality is discussed and examples of how to attain it are provided.

The last chapter sums everything up and distils the six pillars that constitute the *foundations* of City Knowledge, namely the *middle-out* approach, informational jurisdictions, distributed knowledge, sustainable updates, citywide standards, and information sharing.

PREMISES OF CITY KNOWLEDGE

Based on my *praxis* as described so far in parts II and III, in this final Part IV, I draw my conclusions and provide a generalized, theoretical synthesis of the lessons we learnt, in order to propose a reasonable pathway that municipalities could embark on to acquire and maintain comprehensive City Knowledge systems.

In the chapters that follow, I introduce the foundations of a methodology aimed at fostering the emergence of a municipal knowledge infrastructure that can be constructed gradually, with a systematic process, without undue stress for local officials and municipal workers, and without breaking the bank. As utopian as this concept may seem, my personal experience has convinced me that this type of "plan-ready" knowledgebase is not only desirable but also quite feasible and sustainable.

single-use, disposable information

Information is already used in cities and towns on a daily basis. Unfortunately though, most of the documentation that is acquired is only used for a single purpose. For instance, the city's building inspector receives an application for a construction permit, with attached drawings and plans that describe an addition to a home. These attachments are only treated as supporting evidence, used exclusively to grant or deny the permit, and then they are shelved and forgotten forever. When the inspector visits the construction site and gives final approval to the finished work, updated drawings are generally not filed away in the city's archives to permanently record the change to the form of the city that just occurred. Unfortunately, in the long run these seemingly insignificant piecemeal changes to individual properties add up to irreversible transformations of our cities and towns.

failure to record change

Moreover, failing to record these modifications as they happen has more immediate consequences on the efficiency of the municipal machine. For example, the assessor's department may never receive notice about the increased footprint of the building, therefore real estate taxes for that property will remain unchanged until the next round of appraisals is done to bring the assessed values up to date with market values[345].

loss of tax revenue

When the tax assessors for the town where I live (Spencer, MA) came to visit my house in 2002, the official assessor's map they were carrying was still missing an addition done in the mid 1990's, as well as the more recent addition completed in 2001[346]. I wouldn't personally complain about the fact that they undertaxed my property for more than a decade, but one can see how this lack of attention to information leads to gross inefficiencies that can result – among other things – in loss of revenue.

345 In actuality, some automated reporting between the building inspector and the assessor does take place, but the inspector does not exchange information with the assessor above and beyond the simple signaling the completion of a renovation or construction. This alert will in turn generate a visit by the assessor. The process is still sub-optimal since the assessor does not receive any concrete information about the work, but merely a notice of completion. In fact, typically no city office receives final drawings or any information worth retaining from the owner or contractor, which represents a missed opportunity for "free" accrual of city knowledge from the ground up.

346 Together, these additions had quadrupled the living space and more than doubled the footprint.

missed opportunities

Above and beyond these blatant examples, inattention to the value of information is so widespread to be mostly unnoticeable. When a contractor is hired by a city to trim branches off the trees in a certain part of town, nobody is going to notice the missed opportunity to extract an informational return from these services. Indeed, the opportunity for a city worker to be face-to-face with an individual tree is so rare that these periodic trimmings are probably the only chances cities have to get a report on how the tree is doing. As we proposed in Cambridge, MA and Venice, Italy[347], in addition to pruning the branches, these crews could collect numerous pieces of useful, yet simple, information from a brief visual inspection of the trees, and a few quick measurements. For instance, the diameter at breast height and the canopy radius could be measured, so that the city arborist could have an idea of whether the growth of the tree is being stunted or proceeding normally. With minor training, crews could be taught to identify telltale signs of the main diseases, so that a botanist could be sent to the plant for a rapid follow-up to make treatment available as quickly as possible. The workers could also report on the condition of the sidewalks, curbs and storm drains vis-à-vis the trees' root systems, as well as on the distance of the branches from the closest houses, telephone poles and electrical power lines.

informational return

maintenance with a plus

added informational value

These examples of value-added informational extensions to typical city services exemplify the basic tenets of a sustainable municipal knowledge infrastructure. If cities make a conscious decision to extract informational returns from every single activity, City Knowledge will naturally emerge as a byproduct, with only minor additional efforts above and beyond current procedures. As soon as a city adopts an information-aware *modus operandi*, slowly but surely the accumulation of city knowledge will become routine and the shift to plan-ready information will occur almost effortlessly.

information-aware modus operandi

office-wide information assessment

Once the switch to this information-conscious approach is made, in order to be able to adopt this tactic, the various departments of a city ought to reassess their standard operating procedures (SOP) to identify:

- the exact type and form of the information needed to fulfill the various responsibilities of the office;
- the sources of the data to fulfill those informational needs;
- the modifications to the current SOP that could enable the acquisition of information that is long lasting, updatable and reusable[348]
- the extensions that could be made to the current procedures, forms, interactions, and related activities, in order to gather richer information that could lend itself to multiple uses

technical and organizational issues

Of course, there will be technical and organizational issues to be dealt with, but the bottom line is that the transformation of municipal maintenance, management and planning operations will happen if and only

347 See for instance page 128.

348 We carried out just such an assessment in the spring of 2004, on behalf of the Boston Environment Department. Cf. Hart *et al.*, 2004.

if city knowledge principles are embraced at least by one city department as a start. Once the collection and organization of information becomes ingrained in any municipal activity, it will be only a matter of time before a thorough and complete knowledge is accrued that will be resilient, reusable, sharable and updateable in perpetuity.

actions to bring about City Knowledge

Once the decision has been made to treat information with the importance it deserves here are two basic activities that need to be undertaken:

1. Collect and organize information about all of the physical elements and human activities that already exist within the municipal boundaries and are either maintained or managed by the city. This is what I will refer to as "the backlog".
2. Develop mechanisms to capture future changes as they happen.

As needed, the update mechanisms could be implemented concurrently with the data gathering, especially since catching up with the backlog may take months or even years to complete. Nevertheless, the backlog will never be taken care of unless some action is taken now. The first and most important move simply consists in consciously deciding to make information as important to the city as other, more traditional infrastructures like roads, sewers and water already are. Such a move would be tantamount to what Thomas Kuhn calls a paradigm shift[349]. When this cultural revolution takes place, the city will be well on its way towards the creation of a sustainable municipal knowledge infrastructure.

information as an infrastructure

a paradigm shift

This new approach to urban ontology seems so natural and obvious that one has to wonder why this paradigm shift hasn't taken place already. The next section explores the financial, technical, logistical and organizational hurdles that have prevented the emergence of city knowledge until now.

349 Kuhn, 1962.

OBSTACLES TO CITY KNOWLEDGE

The paradigm shift I am invoking is not exactly new nor unique. In fact, what I am proposing may be so commonsensical that it may appear borderline trivial. Certainly, many city pundits before me have proposed bits and pieces of the overall approach I propose – from Geddes[350] to Olmsted[351], from Jacobs[352] to Lynch[353] – and it would be surprising if city workers, who frequently need to make quick decisions in contexts fraught with uncertainty and lack of information, had never wished to have what I call "plan-ready" information at their fingertips. So what has prevented these simple concepts from taking hold sooner? What were the obstacles that have made the cumulative collection of city knowledge infeasible until now?

I think there are at least five main reasons to explain why City Knowledge has not been practical until very recently:

1. the perceived high cost of such an enterprise, both in terms of the initial cataloguing effort and in terms of the subsequent upkeep of the information;
2. the difficulty in cross-referencing different data archives, especially before the advent of the PC and desktop databases;
3. the intricacies of relating information to specific locations in space, even when street addresses are used;
4. the complexity of coordinating and synchronizing data within and across agencies;
5. the frustration that many municipal officers have experienced when trying to keep up with constantly changing technology, even after the introduction of computer tools into municipal operations.

The five "problems" are closely interlinked and the solution of the spatial referencing dilemma thanks to Geographic Information Systems (GIS), together with the ease of cross-referencing brought about by relational database management systems (RDBMS) have combined to greatly lower the cost barrier, making the whole City Knowledge proposition feasible and affordable in the late 1990's.

The biggest problem remains a "people" problem, due to the all-to-human reluctance to accept and adapt to changing situations and technologies[354]. "The people and organizations designing and managing GIS often are uninterested in such comprehensive systems"[355] so, after a discussion of the five problems listed above, I dedicate a final section to this overarching problem of the adequacy of the human skill sets for the tasks

---

350 Geddes, 1911.

351 Olmsted, 1913. See quote on page 159.

352 Jacobs, 1961. See quote on page 123.

353 Lynch, 1968. See quote on page 39.

354 Reeve and Petch, 1999, p. 5.

355 Innes and Simpson, 1993, p. 232.

required and the consequent resistance to change that such endeavors invariably encounter.

THE COST OF CITY KNOWLEDGE

Inventorying the physical urban infrastructure already in existence is an imposing task. Before the advent of personal computers and relational databases, systematic and exhaustive tracking of city assets was very cumbersome and consequently not very flexible. Paper records were maintained (and often still are) in file cabinets, using a variety of ad-hoc indexing schemes[356] suited to the mission of the office where the records resided. Re-indexing and cross-referencing were simply not available options if one wanted to re-utilize an existing archive for practical or analytical reasons that differed from the original intended purpose of the documentation. Enriching the archive with complementary information that augmented the core collection of indispensable data was not even contemplated, given how unwieldy these paper stores were, even when the bare minimum of necessary information was retained.

lower cost of computerization

Nowadays, the widespread adoption of computerized databases for many municipal operations has greatly reduced the cost of keeping the records organized, and of making them accessible for multiple purposes. In fact, some of the bigger towns have even begun to use the web as the vehicle for making the information more accessible to citizens, though a full two-way interaction is still not commonplace even in the more advanced e-government systems in operation[357].

As mentioned earlier[358], there are two principal tasks a municipal office needs to address once the paradigm shift has taken place toward an information-aware modus operandi: (1) create computerized inventories of the pre-existing city structures and activities already "out there" and (2) make sure future change is captured as it happens.

creatively catching up with the backlog

The cost of computerizing the "backlog" of information that is already in our municipal archives – or simply already in existence in the "real world" – may appear to be prohibitive for some communities, but there are numerous ingenious ways to make the process affordable even for cash-strapped municipalities. Some of these creative approaches may include:

1. Making the task of computerizing past records part of the daily routine of some of the municipal <u>staff</u> already on payroll;
2. Leveraging inexpensive (or free) <u>volunteers</u> such as interns, summer workers, university students or high-school students to do the bulk of the data entry;

---

356 Like "place-over-time" as Bryan Glascock of the Boston Environment Department describes the typical filing system where permits and other paperwork are first of all filed in folders organized by address ("place"). Within each place-indexed folder, one would then find the documents organized chronologically ("over time").

357 Hart *et al.*, 2004.

358 Page 164.

3. Outsourcing the computerization of the pre-existing situation, through contractual obligations, to contractors that are hired for routine maintenance tasks;
4. Actively pursuing "instant gratification" in the form of a rapid financial return by investing on applications that promise to yield immediate economic benefits;
5. Setting aside some funds (5-10%) from the ordinary budget of each department to gradually computerize the backlog of paper records already in the municipal archives. These funds could either pay overtime for regular staff members or could go toward paying some outside consultants to do the job.

information-conscious job descriptions

Perhaps the most cost-effective manner to make City Knowledge principles a reality in any municipal office is to simply decide that information is utterly important to city operations. Once such a momentous decision is made, the city can review the job description of every civil servant in town with the intent of extracting the maximum informational return out of every worker. So, the next time crews from the Department of Public Works (DPW) go out to unclog a drain after a storm, the "new" job description will force them to map (however approximately) the location of the drain with respect to a nearby landmark. The drain will thus not only be unclogged but it will also be recorded and codified so that future interventions will be able to refer to it by its code-name and will also be able to detect patterns of clogging that may have gone unnoticed until information was accrued with such specific place references.

maintenance-based data collection

However slow this process of information recording on a "need-to" basis may appear, it is guaranteed to eventually produce a complete map of all features that are under the jurisdiction of each department. As an added bonus, the data collection will also be done on a priority basis, leaving the less trouble-prone elements for last, as logic would suggest. The drawback of such an approach is that it would require crews to be always on the alert and ready with the appropriate field forms and data-collection equipment. The marginal returns may not be worthwhile once the occurrence of an unrecorded element becomes more the exception than the rule.

free or inexpensive labor

Another low-cost approach to the computerization of the backlog is to enlist the help of *pro bono* volunteers or inexpensive summer interns to do the data collection. In many ways, this is the approach I have personally undertaken through the hundreds of WPI undergraduates and Earthwatch volunteers that I guided through the data collection campaigns I described in previous chapters[359]. Towns already enlist the help of interns for several purposes, so the difference I am proposing would be simply in the sustainability and "staying power" of the work conducted by these volunteers. As with anything else, the ultimate impact of the work carried out free-of-charge by volunteers is only as good as our ability to follow up and use the work – hopefully more than once. Therefore, this money-saving option would entail managing the volunteers in a way that will produce continuous growth of the knowledge database. Making good use of free or

[359] Page 39 and following.

inexpensive labor requires careful management and adequate attention lest it become another exercise in futility. Keeping interns "busy" is not a guarantee that city knowledge will accrue incrementally. Data collection, if carried out by volunteers, is only as good as the methodology employed in the process and is only as usable as one's ability to turn data into information.

contractual obligations

Another surefire way to minimize costs while knowledge is being gradually and systematically augmented is to revise all outsourcing contracts to include informational returns as part of the services rendered. It will be up to the contractors to equip themselves to fulfill the new knowledge-focused contracts, which require updated information in digital form. Thus, the city could demand that the company that is doing the usual pruning of the city trees will also measure the circumference of the tree trunk, estimate the distance from power lines to the nearest tree branch, assess the condition of the sidewalk vis à vis the tree's root system, and conduct a thorough check of visible signs of potential disease. In this vein, the City of Cambridge, based on our suggestion, has obliged its contractor (Lockheed) to keep track of parking ticket information in a manner that will allow analyses to be made to resolve problematic situations[360].

low-hanging fruits

A really convincing way to dispel any doubts about the economic viability of the construction of a City Knowledge system is to look for "low hanging fruit" in the form of City Knowledge projects that will yield instant financial returns, amply justifying any up front outlay of funds. An example of low hanging fruit would be the creation of a catalog of parking meters and curb regulations that would enable parking control officers to be more efficient in their routes[361]. Improvements in efficiency would greatly increase the ability in detecting infractions, which would in turn yield immediate and permanent returns quantifiable in tens of thousands of dollars.

budgeting for city knowledge

My experience suggests that espousing City Knowledge principles saves money – if nothing else for the demonstrated reusability of the datasets. Once a sufficient number of City Knowledge projects will be operational, a more exhaustive economic analysis of the advantages of these systems should provide conclusive evidence to prove or confute this inductive assumption. Eliminating redundancy should at the very least free up lots of time to focus on more fun endeavors than the hunting down of datasets to analyze. A careful weighing of the pros and cons may move a city to simply decide to invest in the creation of an emergent, self-sustaining, once-and-for-all city knowledge system. A yearly amount budgeted for such a cause would initiate the emergence of City Knowledge and would be well worth the investment even if the value-added synergies that I described in previous chapters should fail to materialize. Having well-organized information at one's fingertips will make all operational decisions more

---

360 Flynn *et al.*, 2003.

361 As in Flynn (2003) when our team was able to suggest practical procedural changes to the City of Cambridge, that enabled parking control officers to warn parking meter collection crews about jammed meters on a daily basis. Hundreds of thousands of dollars could be saved with this simple procedure.

effective and efficient[362]. Eliminating redundancy is the minimal result one is going to obtain from embarking in a City Knowledge path and such an outcome is incontrovertible and guaranteed.

the real cost of city knowledge

Inventorying all 472 Venetian bridges cost Insula about 74,000 dollars (€61,000), i.e. $ 156 per bridge. Similarly, our survey of 1627 docks in the entire lagoon cost the city $250,000 (€205,300), which translates to $153 per dock, which is amazingly identical to the per-bridge cost[363]. The savings derived from these multimedia catalogs are hard to quantify, but the cost per object translates into roughly 10 person-hours of work[364]. Would more than 10 hours be spent by someone gathering these same data from now until the end of time? If so, how soon would someone have to come along to collect such data? How often would such a redundant activity take place? These questions are hard to answer definitively. We spend much more time and money on finding, requesting and obtaining data than on manipulating the data and analyzing results[365]. So, I personally think that a city or town that was prepared to invest in City Knowledge would be making a smart choice. If the will to purposely budget for City Knowledge is not there, then one could still resort to the other low-cost or no-cost solutions described earlier.

the cost of knowledge maintenance

While the backlog is being whittled away, one needs to worry about the maintenance of the knowledge being acquired. Knowledge updates always represent a cost for someone – either in terms of money or in terms of time. From the perspective of a municipality, the trick is to extract as much information as possible from the private sector without significant cash disbursements. As long as we restrict the discussion to the maintenance of information about "things", it is possible to imagine how maintenance of the physical object could be coupled – by design – to activities aimed at revising and/or verifying the underlying dataset. As mentioned, the company that is contracted to prune all of the city trees would also be charged with measuring the circumference of the trunk, the height of the tree, the shortest (and more "dangerous") distance between branches and power lines and the canopy radius. These scheduled maintenance activities would thus contribute to a periodic updating of the underlying knowledge-base as well. Eventually, the wealth of up-to-date knowledge available to the city arborist would perhaps dictate a different scheduling or sequencing of the physical maintenance operations, allowing for an optimization that is far from possible today.

self-reported updates

One "free" updating strategy that worked for the Boston Air Pollution Control Commission (APCC) was to let the "customer" (parking facility owner), do the updating of much of the essentials, by instituting a periodic renewal of the permit or license that required the submission of

---

362 Budić, 1994, p. 252.

363 This amazing agreement between such disparate projects is worthy of further investigation, but if provides a wonderful heuristic for ball-park estimates.

364 The about $15.00 per hour used here is an approximate average of the hourly cost for the more skilled tasks such as GIS and database manipulation, together with the less demanding tasks of field measurement and surveying.

365 Budić, 1994; Nedović-Budić, 2000, p. 82.

updated statistics[366]. In general, self-reported data need to be validated, but linking the submission of information to fees may, in some cases, provide a highly reliable means for information maintenance, while also providing funds for the upkeep of the datasets[367].

administrative updates

In general, a focus on informational returns has yielded many surprisingly simple ways to acquire updates with minimum effort, as was amply discussed earlier in this document. By far, the most promising way to keep the knowledgebase up-to-par – insofar as change is produced by human acts and not by natural dynamics – is to actively work to intercept and process the administrative paperwork that accompanies such change, since almost all anthropogenic modifications to the world we live in are decided, requested, required, approved or authorized by some level of government.

Disconnected Datasets

Until the advent of the personal computer in the 1980's, it was inconceivable for city officials to even imagine how different archives could be cross-referenced to produce augmented information from mere documentary data. Only a handful of hard working scholars would ever attempt to do something like that even in a very limited research domain since such a project entailed literally consulting hundreds of paper files, trying to reconcile them with each other to glean at some hidden pattern that explained some interesting phenomenon. It was hard enough to thumb though a single paper archive, never mind two or more. Ironically, it was easier to do such arduous research on very ancient materials than on very recent ones. The paucity of antique records made the task more manageable than it could ever be in the presence of miles and miles of massive shelves brimming with modern paper records.

relational databases vs. flat-files

The arrival of the computer age in the 1970's did little to improve this situation, since the early tools were first and foremost geared towards the "keeping" of the records and not so much for their analysis. Eventually, in the following decade, PCs began to make their appearance on the desktop of researchers and scholars and the first personal database tools became available to the masses under the novel operating system nicknamed DOS created by an upstart computer company called Microsoft. Ashton-Tate's Dbase III was the first database that gathered a certain following in the DOS community. The "relational" capabilities of such database tools were hardly ever tapped, however, and even today "flat-files" seem to prevail, which prevent datasets from being connected into wholes that are bigger (or at least more informative) than the sum of their parts.

disconnected vs. distributed

Paradoxically, this "insular" mentality was further aided by the subsequent mass diffusion of the personal computer in the 1990's, which was not accompanied by a parallel dissemination of networking hardware

---

366 As we suggested in Allard *et al.*, 2001.

367 In the case of the APCC, a yearly fee per parking spot was instituted so the owner would immediately report spots that were no longer available for public parking. Of course, this system could be vulnerable to gross underreporting, but at least we eliminated the opposite problem of over-estimating parking availability. See also Eichelberger, 2004.

and software. Stand-alone applications held sway until the end of the nineties when finally the usefulness of networking became more apparent to everyone and connecting PCs became cheaper and easier, as the World Wide Web became an overnight sensation since its modest beginnings in the early-nineties[368].

Creating an emergent system that would grow knowledge from many different-yet-connected systems, managed and maintained by different agencies having jurisdiction over different urban domains, only became truly possible very recently. Until the end of the last century (and millennium) such a system was a mere utopia, though the technical difficulties involved in such an endeavor did not stop some visionaries from predicting the day when such a networked, distributed intelligence would be commonplace[369]. I was fortunate to be there during such a momentous period of human history and I have kept in my personal archives documents that I wrote in the 1980's that hint at such a system. In my own way, I was one of those visionaries…

organizational disconnects

Today, technology has finally made good on the promises of those heady days[370] and the difficulty in cross-referencing different data archives is organizational and no longer technical. It is now people, departments, organizations and agencies that create the barriers between datasets that make a truly emergent city knowledge system difficult to implement[371]. This dissertation is my contribution towards the breaking of these artificial barriers that prevent such a distributed knowledgebase from being widely used to maintain our urban infrastructures, manage our civic activities and plan the future of our cities.

Vague Spatial References

Another major obstacle to the widespread adoption of the intuitive mechanisms that constitute the City Knowledge approach has been the inadequate manner in which our municipal recordkeeping has dealt with references to geographic locations. Once cities began to keep track of data using computers, their primary purpose was to manage financial records, such as payroll, taxes, fees and fines. To this date, spatial references – when present at all – are limited to traditional street addresses, which, while a step forward, have amply demonstrated their inadequacy for analytical and management purposes.

Even in today's most advanced GIS efforts, buildings rarely referenced by unique IDs[372] and, except for Venice, I have yet to see a city which has coded each individual doorway with a unique code to replace (or

368 For a good on-line recounting of the evolution of the web and other related technologies, see for example, http://www.netvalley.com/intval_intr.html (accessed 6/27/04).

369 See for instance Budić, 1994.

370 McFall, *ENR*, New York: February 16, 2004. Also, Budić, **1994: abstract p. 244.**

371 Nedović-Budić, 2000, p. 82; Nedović-Budić and Pinto, 1999, p. 60.

372 The British TOIDs are a step in the right direction, though I find them a bit too arbitrary and not sufficiently mnemonic to be used successfully by humans when needed, altough Building IDs may be one of the few codes that might as well be numeric sequences since it would be hard to come up with a mnemonic identifier.

at least augment) the inadequate yet typical "street and number" addressing scheme[373].

GIS and spatial references

This dearth of geographic references in city-owned datasets has made it difficult to integrate the information available in different departments and has therefore impeded the development of a distributed City Knowledge system. The advent of GIS in the eighties began to change things a little bit, though real progress has been slow since the power for spatial analysis that GIS provide has been misunderstood and underused in favor of more mundane uses of these powerful tools as glorified mapping and plotting applications serving the needs of planning commissions and the like[374]. This underutilization of GIS is a pity, but it has had the beneficial effect of at least making the tool a household word in most mid-to-large-size communities in the developed world[375].

space as the glue for the urban puzzle

Just as networking has only recently come of age, GIS has also reached a critical mass in terms of its widespread adoption in municipalities worldwide, making this decade ripe for the next step, namely the final tapping of GIS's real power as a tool that will enable disparate pieces of the municipal information puzzle to be glued together through the space that they share. A better appreciation of the "power of space" and a better and more educated application of City Knowledge principles for the unique coding of objects in the real world and the linking of maps to databases will enable the emergent qualities of what I am proposing to be unleashed so that the benefits of city knowledge can become evident to all.

Asynchronous Co-Dependence

Once we put in place a distributed system that exploits spatial relations and employs well-designed codes to connect pieces of information under different departmental jurisdictions, to create a whole that is superior to a mere compendium of multiple databases, we run the risk of creating dependencies between datasets (and hence between departments[376]) that may spell the ruin of our distributed system. If one piece of the puzzle were to fail, it may take the rest of the system down with it. The fear of such dependency has prevented even the best-intentioned municipalities from embarking in the creation of networked systems, linking the different departments along functional lines. The biggest of these dependencies occurs when the entire municipality is expected to connect into a top-down mega-system that encompasses all of the different departments that feed into a huge central repository. If the big, all-powerful central system fails, nothing works. Fortunately, these monstrous systems are a thing of the past since they have demonstrated weaknesses that have prevented them from ever being implemented in full, thus avoiding the risk of catastrophic failure by simply failing to come on line in the first place[377].

top-down dependencies

---

373 Though I am sure that there are quite a few other cities that must have done the same. The overwhelming majority probably has not.

374 Budić, 1994.

375 ICMA survey, 2002.

376 See for example Nedović-Budić and Pinto, 1999, p. 56.

377 Keating *et al.*, 2003. See also Reeve and Petch, 1999, p. 5.

bottom-up independencies

The alternative until recently has been a bottom-up trend toward stand-alone computing. The various municipal departments that relied on a number of disconnected PCs and applications to support their work quickly discovered the need to connect, and the corresponding difficulties in coordinating and synchronizing the efforts of different offices[378]. Such apparent complexity discouraged most cities from even attempting to coordinate the various divisions and even internal interactions within a single department were not so common. Such had been the history of "distributed" computing in city governments until the 1990's and the coming of age of the internet.

the interconnectedness of a webbed world

With the advent of the web, a culture of interconnectedness and a certain familiarity with the concept of sharing through a distributed network of independent computers have created the right mindset upon which the City Knowledge concept of "middle-out" can now be grafted. Middle-out entails that each department will first and foremost take care of its needs, so that the primary functions that the department or office performs will be invariably performed with or without the connection to the outside world. With proper safeguards, each branch office would be capable of functioning on its own, regardless of the state of other offices in the city. Nevertheless, if one department requires knowledge of some aspect of the city that falls under another department's jurisdiction, a City Knowledge system would expect that such knowledge would be shared and that the information would be kept up to date by the department in charge. In a worst case scenario, old-fashioned means of communicating information between departments could be employed and the last-best-version of a dataset could be used if the absolute latest is somehow unavailable at the time.

middle-out to the rescue

CK on the information superhighway

Barring the occasional server that goes down, the City Knowledge system discussed herein would rely on normal internet technology that has reached a high level of reliability and resiliency, so the distributed, interconnected City Knowledge infrastructure would be no more vulnerable to co-dependency than our email system is, upon which we already rely rather heavily to conduct our daily business. If something happened that disrupted these systems in a major and long-lasting way, we would be probably facing problems that are much bigger than the mere malfunctioning of our City Knowledge system.

If there are no hardware problems, the real hurdle will remain the difficulty in coordinating the efforts of different departments who need to share some of their information. This problem will not go away magically thanks to City Knowledge, but I think the gradual self-generated transitions that are envisioned in City Knowledge will make these interconnections more likely to be successful than with any imposed-from-above solution.

378 *Idem.*

TECHNOLOGICAL OBSOLESCENCE

A final concern that has traditionally been cause of a certain reluctance to adopt new technologies in general, and has thus prevented City Knowledge from emerging as a municipal *modus operandi* until now, is the fear of technological obsolescence. As computers became commonplace in the last two decades of the XX century, people have witnessed also a relentless escalation and evolution of both hardware and software. Every year a new improved model or version is released in a never ending race to the ultimate computing power according to Moore's law[379]. In this wild race, many people, companies and municipalities have got burnt at least once by a technology that did not deliver what it promised and resulted in an irretrievable waste of time and money. These bad experiences were so common that terms such as "vaporware" were coined to capture some of the disappointment people experienced when products, companies and entire technologies came and went at lightning speed, leaving behind lots of disgruntled victims, who were thus turned from enthusiastic adopters to cynical neo-luddites. Even the more savvy and fortunate had to endure repeated migrations of their data from one platform to another, having to reinvest time and money at regular intervals to avoid losing years of work by being painted into a technological corner from which there was no escape.

getting burnt by vaporware

growing pains

When we started the project center in Venice in 1988, we used 8086 PC's with 5¼" floppies. Our database was Dbase III; we used Lotus 123 for graphs and – being way at the forefront of technology – we even had Mapinfo for DOS version 1, a GIS program that had just been developed a year or so before by a small company that I visited when both of its workers shared a small cubicle in Troy, New York. We were hugely proud of what we were able to accomplish with those tools. It was already a big step up from the Commodore 64 on which I wrote my undergraduate thesis. Since then, our database migrated to Dbase III+, then IV, then FoxPro, then the early versions of Access up until the current version, and even up to SQLServer and Oracle. Some of the data we use today, though, is still the same we collected way back when. It was not easy, and it was frequently unpleasant and frustrating, but we lived through these transitions with only a few scars to show. We wasted thousands of dollars in bogus hardware and software or on products that served us for only one season before being discontinued. But I would do it all over again because, despite all that, we inched forward and finally blossomed in the mid-nineties, tempered by our harrowing experiences and all the better for them.

future evolution through migrations

I would be delusional if I were to suggest that the internet and GIS and databases as we know them today are the ultimate tools that a city will need to use in order to maintain its data for posterity. Many amazing technologies will come and go in the years to come and what today seems utterly unbelievable and totally awesome, will some day (not too long from now) seem silly, childish and banal. Nevertheless, the mechanisms, the

---

379 In 1966, Gordon Moore suggested that the number of transistors in a microchip would double every 18 months. This original quote was subsequently extended to cover the doubling of computing power and halving of price every 18 moths (see http://firstmonday.org/issues/issue7_11/tuomi/).

structures, the procedures, processes and codes, the jurisdictional partitions and the overall City Knowledge approach that I propose should have staying power way beyond any technology that we use today.

What counts here is not the tool but the approach. Data can always be migrated to a new format as the last 15+ years of my personal experience can prove. However, badly-structured data will remain bad, whereas a well structured dataset and an accurate map layer will remain. If the data and layers relate to permanent (or very slowly changing) features of our urban landscape, then we can be sure that whatever effort we put in today will not go to waste because of the vagaries of technological advancement.

Resistance to Change

The biggest obstacle for the development of tools to deal with significant planning tasks that "require comprehensive, multipurpose, and multiuser geographic information systems" have been people[380]. Resistance to change is in some measure due to self-perceived inadequate skills and lack of training programs to ameliorate them. Another part of the resistance that planners experience is probably imputable to the separation between planners and the repositories of data[381]. The "fear of losing autonomy, control over information sources, independence, and organizational power is widely acknowledged"[382]. There seems to be a consensus among researchers that planners have a "limited vision of the potential of GIS[383]" and this has resulted essentially in a stagnation in the development of Planning Support Systems despite the great technological advances of the last decade[384].

Although, "the most important impediment to the implementation of GIS in planning may be planners themselves[385]", in this paper I propose a distributed system of data accrual and sharing that may allow planners to skip completely the issues of data collection that take up so much of their time[386] and thus be able to focus on the more challenging issues that planners are supposed to concentrate on. The key to overcoming the "issues of organizational inertia, mistrust, and "turf[387]", according to Innes and Simpson[388] is to follow an implementation path that displays the following traits: simplicity, observable benefits, relative advantage, ability to make small trials, and compatibility. I think that my City Knowledge approach incorporates all five of these principles.

The preeminence of this obstacle is why I think that the most important step toward City Knowledge is to accept information as a core component of the city' infrastructure, on par with water, sewer, roads and

380 Innes and Simpson, 1993, p. 232; see also Reeve and Petch, 1999, p. *xiv* for ex.

381 Klosterman, 2001, p. 4; Reeve and Petch, 1999, p. *xiii* and p. 6.

382 Nedović-Budić and Pinto, 1999, p. 54.

383 Innes and Simpson, *idem.*

384 Klosterman, 2001; Geertman and Stillwell, 2003.

385 Innes and Simpson, *idem.*

386 Nedović-Budić, 200, p. 82.

387 Nedović-Budić and Pinto, 1999, p. 60.

388 Innes and Simpson, *idem.*

electricity, and to begin treating it as such in all aspects of municipal operations. This paradigmatic shift alone will generate the rest of the transformations needed to gradually bring City Knowledge to be embedded in the municipal *modus-operandi.*

QUALITIES OF CITY KNOWLEDGE

In this chapter, I present a sort of "wish list" that contains all of the positive qualities that a City Knowledge system should display. After explaining why each quality is desirable, I point to specific instances from my personal experience or from the literature that demonstrate that such a quality is indeed achievable in a municipal information system and I then discuss how that quality can be achieved by a City Knowledge system. I this chapter, I will try to address each of the obstacles from the previous chapter, while also referring back to the lessons sprinkled around Parts II and III to support my arguments.

In my opinion, the distinguishing qualities that a comprehensive City Knowledge system should try to achieve are[389]:

- Affordable and Easy-to-assemble
- Gradual and Systematic
- Permanent and Exhaustive
- Sustainable and Up-to-date
- Rich and Reliable
- Flexible and Re-usable
- Shareable and Secure

AFFORDABLE

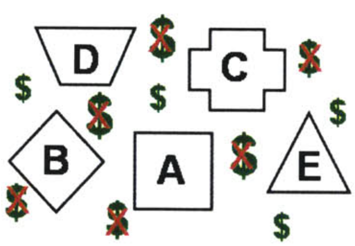

lower transaction costs

Capturing the information for all of the structures and activities that are already "out there" in an urban environment may seem like a daunting task. Catching up with the backlog will have a cost associated with it, but the expense can be amortized over a long period. As shown in previous chapters, most of the data we collected in Venice was gathered by students working *pro bono* (actually *pro grade*). In a typical year, we only had 24 students in Venice, for a period of only two months, yet we were able to acquire an impressive collection of datasets on a variety of different aspects of the city. Though there are costs associated with these endeavors, they can be defrayed in a variety of creative ways, as explained in the previous chapter.

The archival of city knowledge that I propose makes practical and economic sense today due to the declining costs of such activities, thanks to the technological evolution that has brought databases and geographic information systems (GIS) into the mainstream of municipal operations, even in smaller towns. Once the existing state of things is recorded and organized in databases and GIS layers, the task will then be to intercept change on a day-to-day basis, so that there will never be the need to catch up with backlogs again in the future. If done carefully, this constant upkeep of information should cost very little additional money. The difference between the current procedures and those that will be put in place in the

[389] I think that all of the five principles for an effective implementation of GIS that Innes and Simpson (1993) list (see page 175) can be mapped onto the qualities discussed here.

future is really minimal in terms of resources. The main difference is one of focus, as discussed in the previous chapter[390].

EASY-TO-ASSEMBLE

As already mentioned, we are at a peculiar juncture in the advancement of information and communication technologies (ICT) that enables us to take for granted computer tools that simply were not widely available as recently as the late nineties. Each individual element of city knowledge is typically not that complicated to capture in databases and/or geographic information systems. In fact, one of the contentions that is made in a later section is that the fundamental elements of the urban realm – from trees to streetlights, from roads to traffic lights, from park benches to public art – are rather common across the international municipal landscape.

uniform characteristics

The universal nature of the components that make up a city creates the possibility of the development of standards that parameterize the uniform characteristics of typical city assets. Such a solid foundation will then enable supplementary customizations to suit the peculiar needs each specific township. The relative simplicity of the parameters that characterize each category of objects (or of actions, as in the case of traffic) makes it possible to train staffers, volunteers and contractors in the procedures necessary to consistently collect the information in the field[391]. A fundamental tenet of our approach is to always atomize the parameters so that each aspect is gathered in a manner that is as objective as possible. Our data collection always relies on visual inspection and simple measurements, with the occasional support of more sophisticated instruments when necessary. Once again, the proof of the fact that these activities are easy to conduct lies in the fact that our massive city knowledge effort in Venice was carried out by twenty-year-old students, who did not even speak the local language, yet were capable of gathering all of the necessary data even about items that they had never heard of before their projects began.

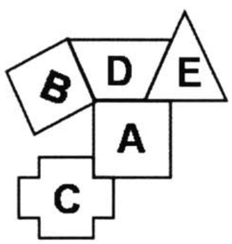

GRADUAL

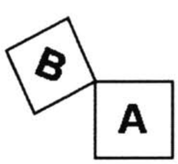

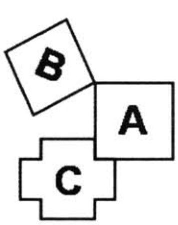

It took us more than a decade to complete some of our largest databases in Venice. As mentioned, tens of thousands of student-hours went into the creation of our public art and canals information systems. Yet, in any given year only a handful of students devoted a maximum of just two months to each of these undertakings. The trick for us was to tackle one borough at a time, so that at any one point we would always have some parts of the city completely done. The difficulty, from the academic and pedagogical perspective, was to propose ever-challenging projects, even though the topics may have been the same as the prior year's, albeit in a different part of town.

The way around this conundrum was, on the one hand, to focus on the analytical aspects of the projects, challenging the teams to higher-order and more complex analyses of the data, thus making the data collection just an incidental part of a more sophisticated study that tested the critical-thinking abilities of our students. On the other hand, once all of the challenging analytical angles were exhausted, as a responsible educator I had

[390] See page 166 ff.

[391] See for example our tree projects in Cambridge (Creps *et al.*, 2001) and Venice (Bennett *et al.*, 2001).

to stop proposing projects for my WPI students on that specific topic. I made up for this by continuing the more mundane data collection work using volunteers. Much of the public art catalog was completed by dozens of pro bono assistants under my guidance as Principal Investigator of research projects funded by an organization called *Earthwatch*[392]. Volunteers, interns, gradeschoolers and university students all can play a role in the gradual accrual of city knowledge, as they did for us in Venice.

[better late than never]

The key here is to accept the fact that since these objects have been "out there" for decades (centuries in the case of Venice) without being thoroughly investigated and inventoried, it won't be a problem if we take our time cataloguing them now. Even if it takes a few years, they're not really going anywhere and a slow progress is better than none. Conversely though, if we do not begin the process now, we are guaranteed to never see it done.

All of the structures and activities that make up our urban reality seem so intricate, complicated and innumerable that it is hard to fathom how we could hope to actually keep track of all this complexity. The principle of *graduality* is one way to deal with the apparent immensity of the task. Biting just what one can chew is a good way to "divide and conquer" the apparently insurmountable hurdles.

392 www.earthwatch.org

SYSTEMATIC

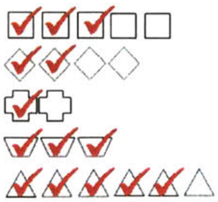

To guarantee that the final product is complete, though, one needs to be *systematic* about the way in which data are collected. In dealing with the backlog, we need to be deliberate and methodical, because the records of past changes will become permanent features of our municipal information infrastructure. Since we waited so long to get this "city stuff" organized, we might as well take our time gathering the information so that we don't miss anything along the way. In Venice, some of our projects benefited from the ancient city's idiosyncratic addressing system, which allowed our students to systematically visit every single address in each borough, in sequence, to make sure that nothing was missed along the way[393]. Similarly, in Boston and Cambridge, target streets were combed thoroughly by our teams to ensure full coverage of the particular aspect of city knowledge being studied.

The modular partitioning of the city into manageable neighborhood units that were analyzed one at a time is another way to ensure steady progress toward the completion of the full inventory of pre-existing urban elements in a *systematic* manner.

PERMANENT

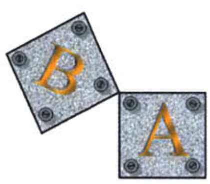

Cities may be seemingly intricate and unwieldy, but their physical make up changes really slowly. If we go back to Kevin Lynch's simple separation between structures and activities – the container and the contained – we can see that much of what makes us think of cities as very dynamic and ever-changing is due to the frenetic pace of activities that take place in their streets and sidewalks. Despite all of the observable vitality, the structures that provide the backdrop for those activities remain practically unaltered from day to day. The elements that make up the concrete physiognomy of our hometowns are finite, enumerable and thus eminently recordable once and for all (or *una tantum* as we would say in Italy). It is my contention – borne out of my experience – that it is possible to capture all of the pre-existing features of our material urban environments in a gradual and systematic way and thus to create permanent records that will make it unnecessary to go back to collect any more information about these objects ever again[394]. Realizing the immanence of the tangible city makes any effort at permanently recording city knowledge worthwhile. While it is hard to quantify the exact financial benefits of the efforts I propose, it can be logically argued that a one-time-only, in-depth, systematic and gradual campaign to organize the information about municipal assets for perpetuity will save money in the long run, vis-à-vis the costly and often redundant consultant studies that are commissioned year after year to collect data demanded by the "plan *du jour*".

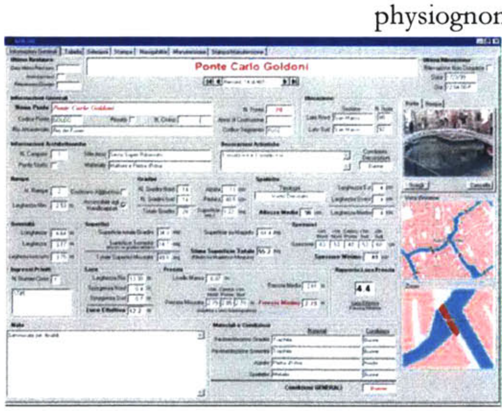

permanent vs. dynamic datasets

It must be remembered, however, that even the most immutable features of a city will always be subject to changing conditions, which is why all of our databases are always separated into permanent and dynamic

393 In the summer of 2004, for instance, a team of students followed the entire address space systematically to map all of the storefronts in Venice.

394 This statement is not entirely correct since even "permanent" data will change ever so slowly, but for the purpose of this discussion, such splitting of hair has been purposely omitted.

components for each asset. For instance, bridges in Venice have been there for almost a thousand years, so one database was dedicated to their invariable characteristics, like the span, number of steps, height of rail, clearance, materials and other such features. According to the classic "entity-relation" model for Relational Data Base Management System (RDBMS), alongside this permanent dataset we created a dynamic database – linked to the former via the unique bridge code – that captured the physical conditions of each bridge at the time of our inventory, listing such things as the damage to the steps, the physical integrity of the arch, the state of conservation of the pavement and many other aspects affected by wear and tear over time. Insula, the company that commissioned our bridge inventory, later proceeded to create additional linked databases to keep track of maintenance work conducted on each bridge, which in turn lead to the updating of the time-stamped condition assessment database records.

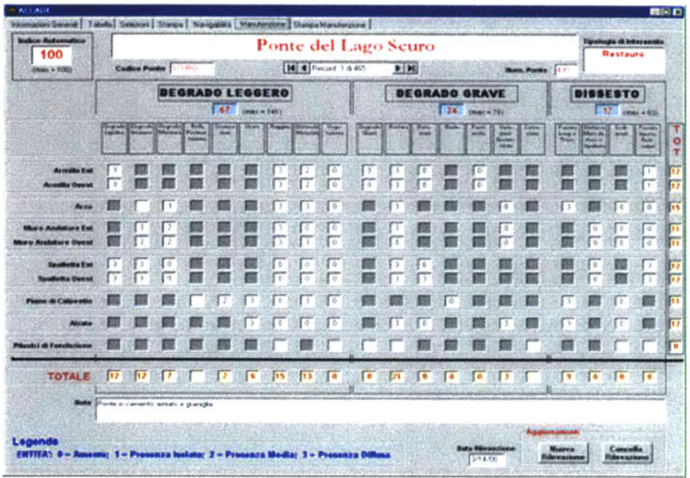

EXHAUSTIVE

informational equity

A gradual and systematic approach to the collection and organization of permanent urban features will eventually result in a comprehensive municipal knowledge infrastructure. Unlike the happenstance consultant reports that focus on this or that part of town as the need arises, the approach I propose will not leave any neighborhood out of the picture. The availability of information about specific areas of a city will not depend on the vagaries of past studies, but will be guaranteed to be *exhaustive* and complete for every single borough, block and street, with no exceptions[395].

In some ways, this feature of my proposed municipal knowledge infrastructure will address issues of social equity that are often hidden in the confusing piecemeal approach to urban information that currently prevails in cities all over the world. It may be argued that the current state of affairs is so fragmented and disorganized that there is an odd form of equity at play in that every part of town and every social class is equally subjected to the inefficiencies that are ubiquitous in our municipal governments. City knowledge, as I see it, will correct this oddity and ensure that the right kind of informational equity is attained through plan-ready information.

SUSTAINABLE

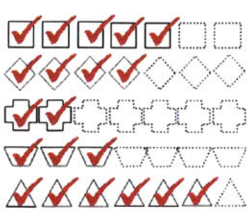

Cities may be slow to change, but they do change a little bit every day. Indeed, it may be a slight misnomer to call some of our urban information "permanent" when in fact it is just "changing very slowly". At any rate, complete city knowledge is a moving target that requires constant upkeep. For a municipal knowledge infrastructure to be truly sustainable, it needs to be constructed in such a way as to facilitate updates in the most cost-effective, transparent and effortless manner. To avoid cost overruns,

[395] Our completed Public Art catalog in Venice (see picture on page 110), with over 4,000 records, for example, is exhaustive of all public art of each type.

the data that are collected to feed this urban knowledgebase must be consistent with what is typically acquired by cities in their normal operations[396].

Maintaining the data should be a "natural" extension of typical current practices. The updating effort should be commensurate with the usefulness of the information. To be true to the tenets of sustainability, we should not go out of our way to willy-nilly gather data that expends resources that we are borrowing from future generations of citizens, particularly when the effort is costly and the returns are only marginal. A sustainable municipal knowledge infrastructure does not live beyond its means nor does it waste precious informational capital.

informational capital

The resolution of the data included in a City Knowledge system today needs to only address today's needs and can only be commensurate to the current technological capabilities. A sustainable City Knowledge system will adapt to changing circumstances[397] and to evolving technologies, so it needs to be resilient. All of the qualities listed herein contribute together to making these systems adaptable to varying degrees of precision and data quality.

UP-TO-DATE

updating information about "structures"

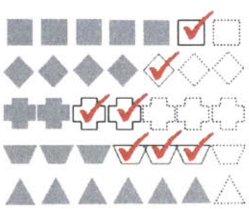

[transactions]

Updating the portion of city knowledge connected with the physical structures that change only sporadically should be simply a matter of intercepting administrative transactions as they happen within the already established procedures that are part of standard municipal operations. This means that from now on, when an inspector gives the "thumbs up" to a newly constructed addition, data from the corresponding drawings are entered into a "buildings database" and, on the GIS side, the "buildings layer" is updated to reflect any change of footprint. The more sophisticated municipal systems may even keep track of internal volumetric changes, by linking 3-D CAD drawings to the building records[398].

These procedures could be institutionalized in such a way as to cost next to nothing to the local community, by simply making the owner/contractors submit these attachments according to predetermined formats[399]. Notice of these changes could be cascaded down to the assessor's office immediately to spark a new appraisal that would in turn be parlayed into an updated tax bill. As unpleasant as such efficiency may sound (especially for those of us who own a home and pay real estate taxes), these revenue-generating examples are meant to emphasize the measurable instant returns a town could obtain on rather modest investments aimed at streamlining current operations.

updating information about "conditions"

Even though the fundamental character of urban structures does not change dramatically over time, their *condition* or *state of conservation* changes constantly since most of these items are outdoors, exposed to the elements.

---

396 See Innes and Simpson, 1993.

397 As we had to do when we constructed the boat traffic model in Venice, which forced us to add pseudo-nodes to the canal network wherever there was a dock or a bridge.

398 I have advised a Masters thesis in Civil Engineering at WPI on precisely this topic (Samdadia, 2004). Michael Batty in the UK and many others have also been pursuing the integration of 3D CAD with GIS.

399 As we proposed in Cambridge (Gage, 2003).

For some urban elements – like city blocks or parking garages – talking about state of conservation does not make much sense. For some other city assets, condition assessments may need to be conducted only occasionally – as would be the case for traffic signs, for instance. Some categories of physical objects, however, may require frequent updates as do trees for instance, which may become infested by parasites if neglected for too long, or may endanger people or property with their limbs, especially in bad weather. Anything that may constitute a hazard for public safety usually gets priority status and is monitored more closely. Worsening conditions usually entail physical decay and potentially dangerous static deterioration of the object, with perilous consequences for people and properties near the object.

[transactional snapshots]

[citizen complaints]

Fortunately, the worse states of decay are usually reported to city hall by concerned citizens who thus perform this surveillance duty for free in order to protect themselves, their properties or their loved ones. Properly managing citizen complaints can be an effective way to inexpensively monitor the most delicate and treacherous deteriorating conditions around the city. Provided these reports are followed by some corrective action, some form of citizen *vigilantism* may be cost effective for the municipal treasury as well as empowering for neighborhood communities. It is worth remembering that a well maintained city is respected by its citizens, as mayor Giuliani's policies proved in New York city in the 1990's. Enforcing maintenance standards in public and private property can create virtuous cycles of overall improvement of the quality of life in a community, dramatically contributing to the reduction of petty crimes against property and even diminishing the frequency of more serious felonies in the long run[400].

updating information about "activities"

Keeping current with the information about the "activities" within a city's boundaries requires an approach that is different from that used for tracking permanent physical characteristics and is also somewhat different from the methods one can adopt to monitor changes in the state of conservation of these material elements of the urban environment. Dynamic processes, like traffic flows, or business vitality, or demographic change are harder to maintain up-to-date by simply tracking the administrative "red tape".

[snapshots]

[administrative proxies]

Keeping track of new car sales or new car registrations may provide proxies for some of the dynamics that urban managers and planners are interested in, but they will never tell us where a car customarily travels to, or at what time or with what frequency. Toll-road accounting systems and other automated devices can fill the gap, though privacy safeguards prevent us from using the data at the finest grain that is technically available, as we experienced with the Venice boat-monitoring projects described earlier[401]. Although there is some unexplored potential for a more careful analysis of how existing record-keeping systems may help us in developing a framework for the monitoring of these activities, the bottom line is that these dynamic

[automatic devices]

400 Gladwell, 2000, pp. 144-146.

401 Starting on page 86.

practices need periodic monitoring to provide the needed information to administrators and decision-makers.

Automatic devices, such as the ones described in earlier sections[402], can be brought into the system to collect reliable data 24/7, greatly reducing the costs of data collection, after an initial investment in the required hardware and software.

[grain and privacy]

Different agencies customarily collect information about ephemeral activities with reasonable frequency[403], though these informational snapshots often present shortcomings of one type or another. Opportunities exist to leverage existing procedures to rein in some of the possible "free" updates that may be available to municipalities, as was discussed in our experience with traffic data in Cambridge, Massachusetts[404]. Another creative and educational way to conduct updates is to involve local schools. For instance, in Venice we enlisted the help of local gradeschoolers to "keep an eye" on the public art collection by promoting yearly "treasure hunts" for the school children, to ensure that the artwork was still there and in good condition year after year. Beyond that, city planners should make sure that adequate funds are available to collect whatever datasets are deemed essential about these more short-lived processes that take place within our cities and towns.

RICH

While we need to be sustainable in our practices and limit our data collection to the realm of established procedures and to facets of urban life that merit attention and are already acknowledged as important to city maintenance, management and planning, we also need to make sure that we do not flatten our data accrual in such a way as to make it impossible to turn our datasets into a re-usable and sharable information infrastructure. Our base data need to be *rich* enough to allow both horizontal sharing and vertical aggregation. If the fundamental datasets are too plain and/or too specific to a particular task, they will not lend themselves to multiple uses, thus they will not allow municipalities to enjoy the economies of scale and savings that could be obtained when the same data are reutilized in a different context without the need for additional data collection[405].

horizontal sharing

vertical aggregation

The richness of a dataset cannot rely on fortuitous coincidences that somehow make it possible to take data from one particular realm and use it in an unforeseen way to analyze another facet of urban operations. These happenstances may give us useful insights, but we need more than that to create a sustainable knowledge infrastructure. In my work in Venice, I relied on intuition to enrich the data that my students were collecting, by suggesting the inclusion of parameters and measurements that were not immediately useful to the task at hand, but were collected "just in case" to

fortuitous vs. intuitive vs. teleologic

402 See page 83 ff.

403 See, for example, the efforts by the National Neighborhood Indicator Partnership (NNIP, http://www.urban.org/nnip) and by Neighborhood Knowledge Los Angeles (NKLA, http://www.nkla.org), to name a couple.

404 See footnote 317.

405 Specific cost-savings are reported in the literature by Nedović-Budić and Pinto, 1999, and others. The re-usability of the data can be gleaned from the example sin Parts II and III, like for example on page 66.

make possible some ulterior use of the dataset at a future date in some foreseeable situation down the line. This intuitive approach needs to be translated into a more cognizant teleological method of data enrichment that makes the foreseeable future applications much more overt and hence makes the dataset extensions more explicit and codified[406].

limited "face time"

One needs to remember that "face time" with the physical objects (or activities) out in the urban domain is a rare commodity and these infrequent opportunities need to be exploited for the maximum benefit[407]. When a city worker is sent out to change a light bulb on a streetlight, this ought to be seen as an occasion to get an update on the condition of the entire light post (Does it need painting? Is it damaged? Are the cables hanging low? Are there branches occluding the light?) as well as on the conditions of other lights along the same street. Enriching the datasets collected in the field may seem to add too much extra work in the context of the specific operation at hand (in this case the simple changing of a light bulb), but the marginal added cost of gathering the few additional pieces of information could be recouped in the long run since a fuller, richer and more up-to-date picture of the city is thus made available to municipal administrators, who will therefore be in a position to make maintenance and management decisions in a more informed manner, without guesswork.

[maintenance-based updates]

RELIABLE

For city knowledge to be a powerful tool in urban maintenance, management and planning, the underlying datasets need to be eminently *reliable.* People need to trust that the information they are using is accurate and up-to-date. If data are perceived to be inaccurate or unreliable, they simply will not be used. Moreover, the data need to be reliably available, and not here today and gone tomorrow. The sustained existence and proven reliability of municipal datasets will create a constituency of users that will come to rely upon them for daily operations, thus perpetuating and reinforcing the need for such information.[408]

continual reliability

[data quality]

reliable sources

Depending on the source of the data, people may be inclined to rely on them with more or less confidence. What happened to me in Italy was that city administrators were loath to accept data produced by students as sufficiently dependable to be incorporated into their activities. I was forced to create my own corporation (Forma Urbis s.a.s.) to validate and integrate the student work and thus put my company's seal of approval on it, taking full responsibility for data accuracy. This process of certification was not only accepted but actually encouraged by Venetian administrators who needed a "fall guy" in case problems with the datasets were later discovered.

FORMA URBIS

[reliable or liable]

To expedite the whittling away at the backlog, cities may choose to enlist trusted contractors and consultants to carry out some of the work. In the long run, however, I envision a self-supporting municipal infrastructure wherein each department will be responsible for collecting and validating all

departmental certification

---

406 The process described mirrors my own personal metamorphosis from empirical intuition to a more theoretical reflection, which in turn parallels the evolution of planning theory itself, as described by Peter Hall in his *Cities of Tomorrow*, pp. 322 ff.

407 As in our tree projects (Creps *et al*, 2001 and Bennett *et al.*, 2001).

408 See Budić, 1994, p. 252, Table 4, under Operational Effectiveness indicators.

[implicit quality control]

of its own data, thus ensuring their reliability since the department's own operations will depend upon the accuracy of those datasets. This implicit quality control will in turn guarantee to all departments that the data they exchange with each other will be intrinsically reliable, or at least as accurate and up-to-date as possible at that time[409]. This virtuous cycle should gradually improve data reliability, accuracy and precision for the overall benefit of all municipal operations.

FLEXIBLE

rich and flexible

standard and flexible

predictable flexibility

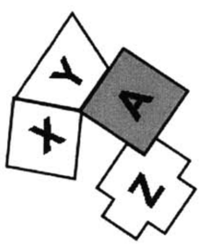

atomized flexibility

[reconfigurability]

atomized resilience

The richness of the municipal datasets that I envision implies a degree of malleability that opens up possible avenues for reutilization of the data in disparate contexts. City knowledge is *flexible* because of its richness, and ironically also because it is fairly rigidly structured along standard coding and reference schemes. It is counterintuitive and almost paradoxical that the predictable and rigid backbone of our data structures enables the great flexibility and adaptability of city knowledge to changing circumstances. Yet the oxymoron of "predictable flexibility" is implicit in the teleological enrichment of our datasets discussed earlier. This apparent contradiction is also inherent in any situation where standards are widely adopted. The rigid abidance to the GSM telephone standard gives me the possibility to travel back and forth between Massachusetts and Italy without having to change cell phone. The universal power supply in my laptop functions in both countries regardless of whether the line voltage is 110 (US) or 220 volts (Italy). Yet life would be even easier (and more flexible) if every country strictly adhered to the same voltage and the same electrical plug configurations.

The predictably standardized nature of the data structures that encapsulate city knowledge is a prerequisite for flexible reutilization, but so is also the nature of the data that are stored in those data structures. Flexibility in our city data starts at the moment of data collection and initial data archival. Atomized data, collected and archived in the most disaggregated and fragmented – yet logical and organized – manner that is reasonably achievable will always be more adaptable and reusable than data that are aggregated, manipulated or pre-digested before storage[410]. Census data on a tract level are less flexible than those based on the block (or block group). If I want to study a specific neighborhood that straddles two tracts, I will do a better job if I can reconfigure my data using the blocks than I could possibly do by using tracts[411]. Atomized data may also offer a degree of resilience[412]. Aggregate data can often be produced even if datapoints were corrupted or missing altogether. The choice of how to make up for these flaws is left to the analyst. Aggregate data, on the other hand, frequently hide these lacunae and leave the end user no choice as to how to

---

409 See for instance Craglia *et al.*, 2004 and Tulloch and Fuld, 2001.

410 "Data warehousing" is one of the pillars of the Urban Institute's National Neighborhood Indicators Partnership (NNIP). Their current attitude is to "keep the whole file at the ready so you can respond quickly as new data needs are expressed" (*Building and Operating Neighborhood Indicator Systems*, p. 36).

411 See, for instance Tufte, 1997, p. 35.

412 Resilience is one of the topics addressed by Kathi Beratan in her lecture on November 21, 2003, entitled *Managing Complexity, or the Information Needs of Adaptive Co-Management: The Durham NC air quality Case*,, in the framework of the MIT E-planning seminar, Fall 2003.

deal with them. It is better to obtain an imperfect set of fine grained data points than to get an aggregate set that appears to be whole and complete. The fact is that real world data are rarely perfect.

[fine grained = adaptable]

The spatial "unit of measurement" is therefore a key predictor of the flexibility of city data. In fact, the degree to which data can be converted into information depends heavily on the level of spatial disaggregation.

deconstructing city knowledge

Deconstruction of city data will always yield richer possibilities for later information building. This deconstructivist approach need not be applied only to spatial units of measurement either.

I have struggled over the years to dissuade experts in various fields from using the prevailing "analytical" approach to data collection that entails, for instance, the evaluation of the condition of a piece of public art as a single synthetic "expert opinion" expressed symbolically as *poor, mediocre, acceptable, good, excellent*, or on a pseudo-quantitative Likert-scale from 1 to 5.

[non-spatial deconstruction]

The conceptual jump that experts subconsciously make when translating visual clues into their final "grading" of the state of conservation of an artifact is tantamount to gathering data at the block level but then only storing it away as a summary by tract[413]. The visible evidence that is mentally factored into the evaluation is not recorded and hence it is lost to posterity. Yet these atomic indicators could be useful for other purposes or for different analytical summations based on alternative evaluation criteria.

Observable traces of tree disease or architectural damage are fairly easy to distinguish, as discussed in earlier chapters[414]. Non-experts, like students or staffers, can be easily trained to recognize and record these clues

[visual inspection rubrics]

as long as a rubric is created to facilitate the process of visual inspection and detection. So, not only is an atomic approach to the collection of the parameters that characterize an urban element more flexible in the long run,

[atomizing is more flexible and cheaper]

but it is also – in the short term – less costly, since pricey expertise is not necessary during the time-consuming field work. Only after the evidence is collected by inexpensive staff, can experts be called in to conduct quality control spot checks on the validity of the collected data and also to perform more in-depth follow-ups with the objects that appear to be in the worst condition.

Reusable

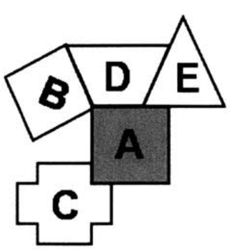

One of the main claims of the municipal knowledge infrastructure I am proposing is that it will eliminate waste and redundancy by allowing accumulated knowledge to be brought to bear as needed, time and time again, without the extra cost of collecting data repeatedly upon demand. This shift from plan-demanded data to plan-ready information is predicated upon a concerted effort to amass the necessary data using a standardized framework that allows the piecemeal addition of new records in a cumulative fashion[415]. The underpinning of this reusability lies in the

413 This represents a commingling of "facts and values" that goes against some well-respected decision-making approaches (such as Hammonds's 1980 and 1991).

414 Page 129.

415 In previous chapters, I have illustrated several examples that demonstrate the benefits of plan-ready information.

rigorous structure for the labeling and referencing of data items which are spatially linked through their geographic location.

units of measurement

The atomization of the "units of measurement" and their geocoding onto standard basemaps creates the foundation for the reusability of the datasets at a later date, and/or under a different set of circumstances altogether. For instance, I was able to compare lagoon boat traffic data collected on two weekdays in 1997 and 1998 to determine changes from one year to the next, simply by re-using the datasets as they were and extracting just the records that were collected from the same locations in both campaigns[416]. The same traffic data are also being re-used in a higher-order project for the calibration and validation of the boat traffic model commissioned by the *Commissario al Moto Ondoso* in 2003[417].

units of analysis

Moreover, not only can the information prove its plan-readiness simply by being reusable without modification as in the examples above, but it can also be re-used and re-aggregated according to completely different "units of analysis", sometimes in ways that would have been hard to imagine *a priori*. For instance, our cargo delivery data, originally collected from individual docks, were later aggregated "by island" to determine the overall "demand for deliveries" for each isle. This lead to the plan-demanding cargo re-engineering project that has revolutionized the delivery of goods in Venice[418].

SHARABLE

intra-departmental sharing

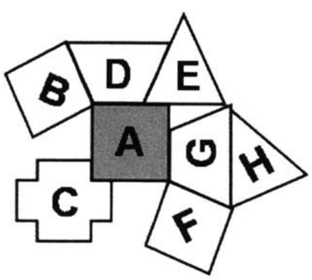

longitudinal legacy sharing

Re-usability, richness, flexibility and reliability, empower us to share information with others. First and foremost, the data could be shared internally <u>within</u> a division or department. Even such a straightforward form of sharing is not quite the norm today, despite its apparent simplicity. For instance, a WPI team proposed to institute a simple intra-departmental form of communication between two divisions within the Traffic and Parking Department in Cambridge, which saved the city upwards of $300,000 a year that would have been lost in meter jams[419]. Obviously, such patently advantageous forms of sharing are the easiest to put in place, since they instantly pay for themselves. Without such venal incentives, WPI teams practiced a form of intra-departmental sharing all along, both in Venice and in Boston, by exchanging and bequeathing our internal legacy databases from one team to the next over the years, parlaying previous successes into stepping stones that would catapult new projects into more ambitious undertakings. Thus we were able to gradually and systematically build upon past projects toward bigger and better results that are truly changing the way things are done in Venice. Thanks to our internal sharing, we can truly claim that we are "leaving Venice better than we found it", and

---

416 Carrera 1997, 1999.

417 *Forma Urbis* was contracted by the *Consorzio Venezia Ricerche* to create the informational foundation for the model. Phase one was completed in March of 2004.

418 as discussed in detail on page 37 and page 118. Our plan was featured in the September 27 issue of *New Scientist* and was also the subject of my interview with the BBC World Service Radio on October 2, 2003.

419 See Footnote number 361.

we are well on our way to do the same in the greater Boston metropolitan area and in Worcester as well.

[the framework for sharing]

We also shared our outcomes and methods with numerous agencies and organizations in Venice and in the US who are all benefiting from our work. More importantly we have created the premises for further sharing across agencies by, for example, creating a *de facto* standard for the coding and labeling of Venetian canals. Insula S.p.A. and various departments in the city of Venice, as well as offices in the Provincial and Regional government have all shared information in an effortless manner thanks to the standard reference system for canal nomenclature that I developed in the 1990's. Similar forms of sharing are occurring in many other areas as well.

inter-departmental sharing

Once the municipal information framework is created, sharing is possible among different departments as well. Such arrangements could first of all follow "paths of least resistance" by concentrating on operations where inter-departmental sharing is already a reality due to institutional mandates as is the case between the Boston Inspectional Services Department and the Historic Districts Commission whenever a building that requested a particular permit is a registered historic property[420]. This is an example of institutional sharing that is already in place and is mandated by law wherein City Knowledge solutions could be easily introduced with instant benefits and without revolutionizing standard practices[421]. After the "low hanging fruits" have been addressed, sharing could be treated as an additional instrument at our disposal to maximize operational efficiency. In an advanced sharing framework, common resources could be mainstreamed into municipal operations in such a way as to exploit synergies that are completely untapped today due to the disconnects that exist among departments.

[paths of least resistance]

[low hanging fruits]

sharing beyond the firewall

Beyond inter-departmental sharing, one can foresee the possibility of making some of the information available to interested parties outside of the municipal firewall. Some information could, and possibly should, be made available to citizen groups, both in its raw original formats and in pre-digested versions for public consumption, as is done in the flourishing community statistical efforts like "Neighborhood Knowledge" in California[422] and the National Neighborhood Indicators Partnership (NNIP)[423]. We experimented with this type of sharing with the citizens of the island of Pellestrina (Venice)[424] and, through the Environmental Protection Agency (EPA), with the Chelsea Creek communities of East Boston and Chelsea, Massachusetts[425]. Beyond public constituencies, the other major users of municipal data are probably academics, especially in the fields of public policy and urban studies and planning, who are always scouting around for information to support their research interests, as I personally continue to do to this day.

[neighborhood data]

[sharing with academia]

420 Hart *et al.*, 2004.

421 *Idem.*

422 www.nkca.ucla.edu

423 www.urban.org/nnip/

424 Battocchi *et al.*, 2003.

425 Desmond *et al.*, 2002.

emergent system

Sharing will allow the distributed data producers in all of the municipal branches to connect together to form a virtual intelligence that is greater than the sum of the parts and that could surface progressively and spontaneously in whatever way is most suitable at any particular time, gradually taking on the connotations of a true *emergent* system with pleasantly unpredictable, yet promising prospects for improved municipal services that will benefit individuals as well as government and non-government organizations.

SECURE

[dig-safe]

[homeland security]

The possibility and desirability of allowing departments to share information with each other brings up the specter of the information falling into the wrong hands. Having a detailed layout of all of the city's gas mains could greatly facilitate the work of road crews and minimize possible "dig safe" hazards, but the same map can also become a weapon in the hands of local disgruntled Americans or of foreign disgruntled anti-Americans alike. In some ways, true homeland security demands the complete knowledge infrastructure that I am proposing, but at the same time all this knowledge could also become a bane if misused by terrorists or criminals. In my view, ignorance is not an option. So instead of refusing to embrace City Knowledge because it could be a dangerous tool in the hands of al-Qaeda, we should just make sure that adequate safety precautions are taken to ensure the security of the data cannot be compromised.

[passwords, certificates and authentications]

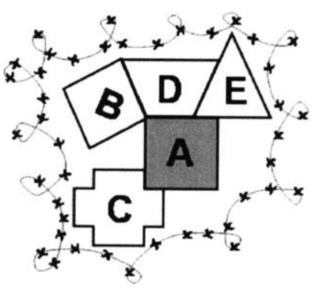

knowledge and safety

As is commonplace in ICT systems, sharing of City Knowledge will be controlled by a system of passwords, certificates and authentications, and the same safeguards that protect the thousands of other sensitive internet sites should be adopted to protect vulnerable urban information repositories. Varying degrees of protection can be placed upon municipal data, depending on its nature, so as not to unnecessarily burden the system with too much security when it is not warranted. When it is appropriate, the data should be protected in the most suitable a cost-effective way. If adequate defenses cannot be put up, it may be better to remove the data from any system that is accessible through internet connections until a safer system can be put in place.

In the end, though, the possible costs associated with protecting city knowledge should not keep us from developing that knowledge in the first place. In fact, it is my hope that the growing concern for the safety and security of urban populations will spur a well-funded effort to embrace city-knowledge principles all around the world[426]. If a natural disaster or criminal incident should occur, emergency crews will want to know exactly everything there is to know about the location where the crisis is occurring. For example, we helped the Boston Fire Department keep track of all Underground Storage Tanks (USTs)[427] and all tall buildings[428] so that fire crews could know what to expect if an emergency was called at a site where either a UST or a tall building was located. In case of evacuations, for

426 Witness our involvement with the Italian *Protezione Civile*, as mentioned in footnote n. 275.
427 O'Donnell *et al.*, 2002.
428 Gaewsky *et al.*, 2003.

instance, fire crews would like to have an approximate estimate of how many people are expected to be in a particular building, and they would also want to know whether and where hazardous materials may be stored within that building, so that the proper countermeasures can be taken.

privacy

Safe and *secure* city knowledge is a prerequisite for the system to work smoothly even in the absence of any terrorist or malicious threat. Simple protection of privacy, for instance, mandates that safeguards be put in place to prevent the release of personal information. All in all, a gradual transition from stand-alone systems, disconnected from the internet, to intranets with impassable firewalls (to share information internally), to limited internet accessibility with proper certification and authentication (to let selected outsiders get access to your information) will smoothly get each department's portion of the overall municipal information system to a proper level of security without risking violations of privacy or malevolent attacks.

[gradual development path]

onward

Attainment of the qualities discussed in this chapter will help ensure that our municipal knowledge infrastructure does not become another exercise in futility, but it actually stands a real chance of becoming an irreplaceable tool for urban maintenance, management and planning. Most municipal efforts will experience difficulties in one or another of the areas highlighted above, but the important thing is to set the system in motion and allow emergence to take place gradually at its own pace. The following chapters provide additional guidance on how to steer the process into a fruitful direction.

FOUNDATIONS OF CITY KNOWLEDGE

information-conscious modus operandi

The qualities discussed in the preceding chapter define the dimensions through which we can evaluate the performance of the knowledge infrastructure that I am proposing. Varying degrees of success can be expected in the pursuit of each of those desirable characteristics of a City Knowledge system, but we have identified practical mechanisms that can at least facilitate the attainment of such positive traits. Once the primary paradigm shift has taken place toward an information-conscious *modus operandi* in all municipal activities, City Knowledge can begin to emerge, based on the following foundations:

- Distinct informational jurisdictions
- Distributed, atomic data acquisition and organization
- Sustainable update mechanisms
- Institutional and/or voluntary sharing of information
- Interagency coordination

"glued" together by a

- Middle-out approach to the development of City Knowledge

[gradual building of foundations]

In the sections that follow, these fundamental principles are each described in detail. Although these elements must all be in place in order to extract the totality of the benefits that City Knowledge can offer, they can also be attained partially and sequentially in a gradual progression toward the ultimate goal of a full-fledged City Knowledge system[429].

[no City Knowledge systems in existence]

As the most recent literature indicates[430], such a comprehensive system does not seem to exist in its entirety anywhere in the world, and even where I was personally engaged with the inner workings of municipal departments – namely in Venice and in various Massachusetts cities as described in parts II and III – these principles have not been completely adopted – not even unofficially. This chapter is therefore purely conjectural, although it pulls together pieces of the puzzle gathered from each of the various cases presented earlier, as well as from comparable efforts discussed in the urban studies and planning literature.

[conjectural hypotheses]

[obstacles]

Numerous obstacles will impede the emergence of a full-fledged City Knowledge system, as described earlier[431]. There is ample literature in the field of Management Information Systems (MIS) that describes the difficulties in the diffusion of information systems in organizations[432]. There is also a growing body of academic literature specifically dedicated to

429 The sequencing is somewhat flexible. Some forward-looking municipalities may start by setting standards. Most practical-minded towns have generally started some sort of computerized bottom-up data collection and GIS layers (ICMA survey, 2002).

430 Laurini, 2001; Brail and Klosterman, 2001; Geertman and Stillwell, 2004.

431 Starting on page 165.

432 For example: Laudon and Laudon, 1996; Marchewka, 2003, pp. 4-5;

the institutionalization of GIS in organizations[433]. Even after the adoption of such an approach has been agreed upon, the devil is always in the implementation details. It is still unclear "how the management in a public agency can move the organization/jurisdiction through the process of [...] GIS implementation"[434]. The consensus opinion indicts the "neglect of the human and organizational aspects" as the primary culprit in the failures of the past[435]. In the municipal domain where my dissertation is focused, the "technology-push" from the top has failed to penetrate deeply enough into the frontline offices where fine-grained data are the daily currency of operational decisions for maintenance and management activities. My middle-out strategy proposes to switch instead to a "demand-pull" approach[436], which centers on the needs of the periphery first.

[where City Knowledge will not work]

City Knowledge is not a panacea and it may be better-suited for some municipal governments than for others. Small towns may not have the resources to even consider this approach[437]. Big city government may be too fractured to be conducive to it[438]. Since space plays a key role in pulling together my City Knowledge strategy, there is an implicit reliance on technology (GIS and RDBMS in particular), which may be utterly inappropriate in some contexts, such as poverty-stricken locales, where basic survival needs overwhelm any other municipal service. City Knowledge may be a luxury that only communities in the more affluent parts of the world can really afford[439]. In some other situations, the concept of a comprehensive municipal information system may be deemed undesirable for political or cultural reasons. With these disclaimers in mind, in the sections that follow I focus on how City Knowledge could be embraced as a strategy and a *modus operandi* applicable only wherever the circumstances permit this concept to be viable and desirable.

[pioneering department]

In this dissertation[440], I submit that it may be possible for one municipal department to set the process in motion by single-handedly beginning to systematically collect data that unequivocally fall under its

433 For example: Budić, 1994; Campbell and Masser, 1995; Reeve and Petch, 1999; Azad (1998) provides a thorough review of the process of managing GIS implementation.

434 Azad, 1998, p. 18.

435 Reeve and Petch, 1999, p. 5. Evans and Ferreira (1995) however suggest that technology is also culpable.

436 *Idem.*

437 Although I have begun to develop a project proposal to connect small towns (like my hometown of Spencer, MA) with Regional Planning Authorities (like the Central Massachusetts Regional Planning Commission – CMRPC) through web-based applications created and supported on the RPA's central server that will automate routine tasks such as construction permitting at town hall while "informating" planning at the regional level.

438 Although, when taken one department at a time, even the biggest bureaucracy can be tamed, as my cases showed. The biggest obstacle may come when dealing with standardization and jurisdictional coordination, which have not been institutionalized explicitly even where I worked directly with willing departments.

439 Yet, if we consider information as an infrastructure as I am proposing, we may want to put in place information systems at the same time as we lay down other infrastructural essentials, such as water, electricity and sewers. If nothing else, we could at least keep track of these physical systems more efficiently, thus minimizing maintenance costs in the long run.

440 Grounded on some recent literature surrounding the "tipping point" principle (Gladwell, 2000 and Godin, 2001).

[capture backlog]

domain of control until it has organized its own backlog of existing information in an exhaustive and systematic manner[441]. Having caught up with the pre-existing *status quo*, this pioneering department can then begin to implement procedures and invent mechanisms through which any new pertinent information that comes into being can be captured as soon as possible and with the smallest expenditure of human and financial resources.

[intercept change]

spreading city knowledge

After setting the example by creating a working operational model of how to manage its own slice of city knowledge, this innovative department could then set in motion a chain reaction[442] that could "infect" other departments[443]. Whether this chain reaction happens or not, my claim is that the mere systematic organization of a single aspect of a department's duties would justify the effort in and of itself, so any additional benefit would be a value-added "cherry on the cake".

[spontaneous coordination]

Beyond that, if two departments could demonstrate the viability of a full-fledged departmental City Knowledge system, the need or desire to coordinate a bare minimum of citywide standards would emerge naturally. Such a harmonization would in turn ensure that, as the various departmental systems come on line, they can begin to share information on a "need to" basis, starting from the existing mandatory bilateral collaborations that already bind them into institutional relations with each other.

[institutional sharing]

[full-fledged emergent system]

Over time, this gradual growth of city knowledge from one department to the next could eventually encompass the entire municipality and enlist government offices, civic associations, academic institutions, non-government organizations and even private citizens and businesses in the continuous upkeep of urban information[444].

the rest of this chapter

The following sections expand on these fundamental processes that in my view will foster the emergence of sustainable City Knowledge. At the end of each of the six sections that follow, I critically assess how my particular approach – culled from specific personal experiences I had in Venice or in Massachusetts – compares with past efforts of similar nature. I try to identify what makes my approach different and/or better than alternative methods and try to distill what aspects of current practice would have to change (and how) in order to move a municipality in the direction I propose.

[empirical basis]

[similar to other efforts?]

[different/better?]

[change what and how?]

My numerous experiences in real world implementations have convinced me that truly useful, complete and reliable urban information systems are indeed within reach of a typical municipal government in a medium-sized, "western" city like Venice, Cambridge or Worcester. Often,

441 This process may be achieved with or without the aid of an MIS department. Most medium to small towns will never have the benefit of such a department, yet I have personally experimented with several alternative ways for such small towns to embark in the systematic accumulation of city knowledge. See footnote 437.

442 Like Malcom Gladwell's *Tipping Point* (2000).

443 Like Seth Godin's *ideavirus* (2001).

444 Of course, there is a pretty huge leap between simple bilateral sharing and citywide sharing. Coordinating these efforts gets progressively hard as n>2. My argument is that ifbottom-fed middle-out initiatives can overcome the complex hurdles of multiagency coordination then the resulting approach will be necessarily long-lasting and resilient.

[MIS and GIS theories]

[diffusion and implementation]

[individual aspect is proven]

what my experience has led me to propose as a foundation of City Knowledge, coincides fairly closely with existing concepts in the area of MIS diffusion, or to parts of other approaches in the GIS implementation literature. As one would expect, not a lot of what I propose is completely novel. On the other hand, my personal experiences in the field – as well as those published by many fellow urban researchers – have at least revealed the merits of each individual foundation of City Knowledge.

[ensemble not proven]

What I cannot prove, but I instead predict – based on the evidence presented herein – is that these five ingredients, selected and combined in a fairly flexible manner, can produce results that far surpass what you would get by simply mixing the constituent parts. The novelty of my overall proposal is in the nuances of each ingredient and in how the ingredients come together into the final product.

[the middle-out recipe]

Together, the foundations that are described in the rest of this chapter can elevate municipal data to urban information and thus contribute to the sort of city knowledge that would support second-order analyses and complex decisions. But first we need to introduce the technique that makes it possible to put all these ingredients together: the *middle-out* approach. This method is not exactly like a culinary recipe from a cookbook since no specific sequence of steps and carefully measured dosages are required to put the five ingredients together.

The result is also not guaranteed since – as far as I know – nobody has quite put all of these ingredients together in the same concoction quite yet.

THE "MIDDLE-OUT" APPROACH

The "middle-out" approach to City Knowledge combines the self-serving ingenuity and energy of bottom-up initiatives and the coordination and standardization of top-down approaches. Middle-out is the methodological framework which allows City Knowledge systems to emerge. It is a consciously unobtrusive strategy which guides the choice of common-sense, low-impact tactics that gradually translate the promising concept of City Knowledge into an achievable reality.

bottom-up enlightened self-interest

This hybrid approach empowers the more peripheral branches of the municipal hierarchy and gives front-line civil servants a stake over the upkeep of the information that they themselves require. The "middle-out" approach assumes that municipal officials are motivated by a desire to improve their own performance at their daily tasks related to the maintenance of city property, the management of city services, or the planning of city futures. The bottom-up component of *middle-out* exploits the all-too-human self-interest that civil servants undoubtedly harbor, like the rest of us. Everyone wants to do a good job and receive praise and recognition from peers and superiors. Enlightened self-interest will be the motor that will feed the information infrastructure that I am describing.

top-down teleological coordination

To rein in the anarchy that might ensue if each civil servant with a modicum of computer savvy were to create an information system from the ground up – as has been somewhat the case in the last decade of the XX century after the advent of personal computers and the mainstreaming of databases and geographic information systems[445] – the *middle-out* approach entails a degree of top-down coordination and management. Information systems will be indeed developed in bottom-up fashion by the front-line departments where the action really happens and where information is produced and consumed on a daily basis, but these efforts will be managed and coordinated at the departmental or divisional level and will include a teleological (forward-looking) approach to avoid unbridled duplication of effort and overlapping of jurisdictions that will hamper the sustainability of the City Knowledge enterprise in the long run.

critique of top-down and bottom-up

[shortcomings of top-down]

Both top-down and bottom-up are fundamentally unsustainable in their pure incarnations and both frequently lead to a waste of energy, time and money[446]. Pure top-down approaches to the diffusion of information systems in organizations are generally poorly received because they fail to engage the rank and file[447]. After huge initial investments, they struggle to achieve the fully-integrated coordination and seamless operation that they promised, especially when the task is not easily automated and the organizational structure is not very hierarchical[448]. Nevertheless, the military is a glaring example of where such a hierarchical approach could indeed work – possibly more reliably than the emergent middle-out system I propose.

---

445 Geertman and Stillwell, 2004; Brail and Klosterman, 2001.

446 Reeve and Petch, 1999, pp. 4-9.

447 *Idem.*

448 See Marchewka, 2003, especially pp. 4-6; Reeve and Petch, 1999, p. 76.

[shortcomings of bottom-up]

Grass-roots, bottom-up schemes work well for a while, as long as some "champions" keep the efforts going, but the resulting information systems eventually fall by the wayside because they refuse to "grow up" or are unable to connect with other systems, or simply fail to keep up with technological or organizational changes[449]. Bottom-up efforts frequently fail because of lack of sustained support by a dedicated core of users/developers and because of the related lack of adequate funds to keep the effort alive through changes in personnel, equipment, and software tools[450].

a pragmatic hybrid compromise

The middle-out method is predicated on striking a careful balance between the potentially oppressive and unwelcome rigor of a typical top-down hierarchical system and the practical ineffectiveness of isolated bottom-up efforts, which almost inevitably fail to scale up or to integrate with other similar systems[451]. It represents a pragmatic compromise that taps into the positive aspects of the two competing approaches, leveraging the energy and self-interest that power bottom-up endeavors, and endowing them with the structure, reliability and sustained financial support that accompanies top-down enterprises.

pseudo- and proto- middle-out

Some development paths being recommended today resemble this middle-out approach[452]. As mentioned[453], the City of Cambridge, Massachusetts has begun an effort of this sort, by distributing GIS specialists in the main city departments under the orchestration of the MIS department, although that endeavor is still in its infancy. This is one of many possible paths to promote the diffusion of GIS in municipal operations. The Management Information Systems (MIS) literature has a lot to teach us about the variety of methods that have been tested to favor the dissemination of information systems in organizations[454]. My middle-out approach would be more similar to the Digital Earth effort[455], but limited, at first, to individual municipal boundaries and to the agencies operating therein. Ferreira, in particular, has championed a variation of this middle-out approach when he proposed the use of lookup tables to correct on-the-fly the "stubborn" standardization errors that regularly appear in municipal datasets[456], as well as (with Evans) when he discussed a more general approach to the "messy" technical and organizational issues confronting GIS today[457].

449 Klosterman, 2001 and Singh, 2004.

450 *Idem.*

451 See the interesting "Problems in the Middle Layer" that Reeve and Petch (1999) discuss on page 25 of their book on GIS and organizations. My overall City Knowledge approach is intended to make the "theory" illustrated in figure (a) more achievable, to overcome the "reality" in figure (b).

452 See for instance Barr, 1991; Campbell, 1999 and Yeh, 1999; Talen, 1999; Ferreira, 1998 and especially Reeve and Petch, 1999, p. 50.

453 Page 125.

454 For instance Laudon and Laudon, 1996; Avison and Fitzgerald, 1992; Flynn, 1992; Reeve and Petch (1999) argue that "the development of GIS is like watching a video of the history of conventional information systems being replayed at fast forward speed" (p. 1).

455 See http://www.digitalearth.gov/ and Crockett, 1998.

456 Ferreira, 1998.

457 Evans and Ferreira, 1995.

What these new methodologies have in common is a recognition that "GIS technologies are *not* divorced from the interplay of organizational life: rather they are subject to its vagaries and power relationships"[458]. A middle-out approach should not only simplify the more technical pitfalls of pure top-down and bottom-up approaches[459], but also promises a more gradual, hence smoother, and less traumatic path for the organizational transformations needed to ensure a widespread acceptance and a successful adoption of the principles of City Knowledge in municipal agencies.

horizontalizing growth from the middle

The hybrid combination of these two extant methodologies, not only exploits the good traits of both, leaving the negative connotations behind, but it also removes the hierarchical verticality of the previous constructs, which was in itself a source of tension between management and rank and file across the entire municipality. The locus of control of information is now at the level of a division or department[460]. Within each department, managers and staff are partners in the development of information systems that improve that department's effectiveness, with no imposition from the mayor's office or the city manager or the MIS department[461]. With a proper consensus-building approach, the entire department or division can bond together around the common goal of producing effective computer-aided tools to streamline every function that the department is responsible for[462].

["slices" of City Knowledge]

The *esprit de corps* that will result from this devolved information design will give staying power to each department-based "slice" of city knowledge. These self-contained units could then connect horizontally with other departments or divisions across the municipality. Top-down coordination among all municipal branches and bilateral agreements between individual departments could thus enable the overall city knowledge "compendium" to emerge from the middle out as a combination of these self-directed, semi-autonomous bottom-up efforts[463].

[City Knowledge as a "compendium"]

the tenets of middle-out

Recapitulating, the basic tenets of the middle-out approach are:

1. front-line offices are directly vested into the collection and upkeep of their own city data, information and knowledge[464];
2. each office will systematically address each of the city elements over which it has birthrights and assess the costs and benefits of creating a comprehensive knowledgebase relating to each element;
3. once the low-hanging fruits are identified, the office will explore all feasible and reliable means for the accrual and sustainable upkeep of the data, starting from no-cost options;

458 Campbell, A. J.. 1999, pp. 621-631.

459 Keating *et al.*, 2003.

460 Which is where the reward structures would have to also be located (see Singh, 2004).

461 Although the department may be assisted by MIS in the development and implementation.

462 Some of these issues are covered in Nedović-Budić and Pinto, 1999; Campbell and Masser, 1995; Reeve and Petch, 1999, ch. 6.

463 Cf. Barr, 1991.

464 Though they may yet delegate the nitty-gritty of the data manipulation to a technical department as was done for instance in Calgary, with its "Data Utility" concept (Findlay, 2002).

4. data are collected at the finest appropriate grain and the same office is also responsible for the efficient and effective upkeep of the data;
5. the front line office participates in intra- and inter-departmental caucuses for the definition of top-level issues, such as standardization, jurisdictions and sharing;

what's new about middle-out

Middle-out is a strategic approach that relies on opportunistic tactical choices by self-interested front-line users (the "ants") who engage in self-organizing behavior to take care of their impellent needs. The middle-out methodology is more focused in the implementation of a comprehensive municipal information system than in the diffusion of the *information-aware modus operandi* that subsumes the City Knowledge philosophy.

[empirical basis]

All of my examples from Italy and the U.S. were carried out in collaboration with specific departments or agencies. We purposely shunned projects imposed from top management or politicians and we did not work with individuals except in the case of one-person institutions[465]. All arrangements were conducted with department heads or program directors. Our projects were institutionalized at the department level.

We were able to show some evidence of horizontal diffusion thanks to shared reference codes both in Venice[466] and in Cambridge[467]. We have also shown how our approach can spread horizontally across department or even municipal boundaries, as evidenced by our ability to elicit interesting and challenging projects from new sponsors by showing them what we did elsewhere and how our sponsors benefited from our studies. To some degree, we expect that the desire to embark in a systematic inventory of physical assets under a department's jurisdiction ought to emerge spontaneously in the various branches of municipal government. In any case, I posit that the diffusion of this approach could be abetted by some lobbying by "communicative planners" who have a lot to gain from the implementation of such systems. Yet the primary focus of the middle-out approach remains implementation and not diffusion[468].

[similar to other efforts?]

As amply discussed in the sections above, there have been several efforts in the recent past that contain seeds of my middle-out approach[469]. To date, none of the cities – or even departments – with which I collaborated have formally espoused my wholesale middle-out approach. In some sense, middle-out is one of the more conjectural aspects of my thesis. I am proposing that it would be beneficial for a municipal department to

465 Such as the Boston Air Pollution Control Commission (APCC) or the Boston Conservation Commission (BCC).

466 As a consequence of the canal UNESCO-MURST project (Carrera, 2001b); the transferal of the base layers to Insula (Carrera, 1999d); the traffic projects (Carrera 1996, 1999a); and of some professional follow-ups, such as *EasyDocks* and *EasyBridge* applications mentioned earlier.

467 With the initial cataloguing of parking meters in the Traffic and Parking department (Cullen *et al.*, 2002) that spread to both the meter maintenance crew and the parking control officers (Flynn *et al.*, 2003);

468 Although I find the Tipping Point (Gladwell, 2000) and Ideavirus (Godin, 2001) concepts promising agents of diffusion.

469 See footnote 452 and related sections.

consciously adopt my approach, yet I personally have no evidence that it would indeed work[470].

[different/better?]

What is new about middle-out today is that recent developments in information and communication technologies – especially spatial tools – now enable this approach to be cost-effective at the department or division level, which gives the beneficiaries direct control over the tools used to represent reality[471]. Physical objects and dynamic activities can therefore be captured and stored in their geographic location by the agency most directly engaged with them[472]. Moreover, GIS and databases have been around long enough that the need for some top-down coordination is more commonly acknowledged.

Perhaps what is novel about my approach is simply that I am introducing a new metaphor to empower municipalities to consciously foster emergent behaviors so they materialize from the front-lines out. I am packaging my concepts in new re-descriptions that I hope will strike other planning practitioners as inevitable[473] and will thus become vehicles for the creation of the sort of City Knowledge systems that my numerous examples collectively suggest.

[change what and how?]

The path to adoption of middle-out starts with the conscious decision to apply this method in the framework of the department's information strategy. The issue here is not the simple "adopt/do not adopt" choice, but a more authentic commitment to City Knowledge and its principles – described in the sections that follow – that will guide the success of the middle-out approach[474].

470 Although others do find some positive evidence using approaches that are similar to mine. See for example Keating *et al.*, 2003; Innes and Simpson (1993); and Reeve and Petch, 1999.

471 A desirable characteristic according to all recent implementation theories as described in Azad, 1998, Reeve and Petch, 1999, Nedović-Budić and Pinto, 1999, Singh, 2004.

472 Although typical implementation paths may still be adopted, such as those in Azad, 1998 and Reeve and Petch, 1999, with the technical support of the city's MIS department, the assistance of outside consultants or the intervention of regional planning authorities.

473 See "Intellectual Contribution" on page 15.

474 More on this starting on page 230.

Informational Jurisdictions

jurisdiction over structures and activities

A very fundamental tenet of City Knowledge is the attribution of clear "informational jurisdictions" over the "birth and death" of the "structures" that make up our urban environment, as well as over the "actions" that occur within it[475]. One should not confuse the concept of informational jurisdiction with some form of *exclusive* responsibility over the maintenance and management of these city elements by a single entity or department[476]. Nor, should our informational jurisdictions be mistaken for an assignment of complete "ownership" over <u>all</u> information that pertains to a specific item[477].

birth and death certification

[parental authority]

In the domain of physical structures, the definition of clear jurisdictions that we propose refers solely to the "parental authority" over the certification and formalization of the birth and death of these objects within the municipal administration.

[administrative existence]

[bureaucratic acts accompany actions]

The notion here is that the physical assets that make up the "public" part of a city – roads, buildings, trees, parks – have an "administrative existence" that begins and ends with official bureaucratic acts that always accompany operative government actions in the real world. So, for instance, the parks department normally puts out a work order for the planting of a tree somewhere in the city. The administrative red-tape (the work order) is converted into an action in the real world by a crew that physically plants the tree into the ground. The bureaucratic birth of a tree is marked by a work order (or by a contract if the service is outsourced) and the tree's physical appearance into the real world is a direct consequence of the fulfillment of the work order or contract[478]. In order to keep track of this tree for the rest of its existence, we will use a standardized informational scaffolding to record the tree's characteristics, as explained earlier[479]. The administrative birth of a tree could thus be treated in a manner that is not too dissimilar to how our birth is also "tracked" by an official administrative act – namely our birth certificate – when we are physically born into the world[480].

naming and coding

An arborist's recordkeeping starts simply with the assignment of a label or code to each tree to identify it unambiguously vis-à-vis all other trees in the city, just like our life starts by our receiving our names from our parents and by our being tagged with a bracelet so as not to be "switched at birth" in the nursery room.

birthrights

In the case of the "urban forest", I am hypothesizing that the parks department might have "informational birthrights" over trees[481] and that it would thus be its responsibility to maintain the corresponding GIS layer,

---

475 Akin to what Nedović-Budić and Pinto (1999) call "Responsibility", page 58.

476 *Idem.*

477 *Idem.*, see "Ownership".

478 Obviously this only applies to "public" trees. A canopy study like the one described starting on page 125 would thus still need to capture the canopy of "private" trees as well as those of trees in a group like in a forest (not individual trees), which may not be recorded individually at least initially.

479 See the Cambridge example at page 125 and following.

480 We have all heard the anecdotes about "ghost" people who were not considered "alive", despite their obvious physical existence, because they could not exhibit a valid birth certificate.

481 This is merely an hypothesis used for the sake of the foregoing discussion. It is irrelevant which specific department would really be in charge of trees in a specific city.

since it knows exactly where a new tree is planted and it is also in a position to give new trees a unique name or code to identify them from that moment on. The parks department may also want to attach a few other permanent pieces of information to this "tree birth certificate", like the species of the tree and the date of planting. These birth records are a necessary foundation to city knowledge although they are not sufficient, by themselves, to support all of the desirable re-use capabilities that were discussed in previous chapters.

[birth certificate]

deathrights

Despite the longevity of many plant species, even trees do die, either due to natural causes (old age and disease) or due to sudden traumatic events (lightning, tornadoes or violent impacts with trucks or cars). Not infrequently, trees are also purposely cut down for one reason or another. Some department needs to be given "deathrights" over the recordkeeping related to the removal of the tree from reality. Just as we have "death certificates" that formalize our passing from this life, some department needs to have the authority to remove the tree symbol from the tree layer in the municipal GIS, so that everyone will know that the tree no longer exists.

[death certificates]

The assignment of very clear jurisdictions over birthrights and deathrights for each element of the municipal infrastructure is a *conditio sine qua non* for this whole system to function properly over the long run. It is possible that quibbles may arise among departments concerning the assignment of jurisdictions over birth certificates, but the most clear-cut way to determine them would be to look at who in actuality is responsible for installing, creating, constructing, or giving permission to erect or establish each category of physical objects. Conversely, it may be that the jurisdiction over the certification of the "death" of that object may reside with the same department, or with whichever other department is in charge of physically removing, dismantling, eradicating, uprooting, destroying and/or disposing of the object.

[jurisdiction over birth]

[jurisdiction over death]

[parent department]

In the end, only one parent department or entity should emerge clearly as the responsible party for tracking the birth of each of the physical assets that make up a city. Similarly, the same or another department will be in charge of certifying the administrative death of objects in a specific class, thus allowing the city to track the entire life cycle of physical assets owned, operated, maintained or managed by the city[482].

[life-cycle tracking]

jurisdiction over activities

Jurisdictions are really important for city structures, because they entail the potential for keeping track of changes after the backlog of pre-existing assets has been tackled. When it comes to "activities" (like traffic, parking, crime, or economic development), jurisdictions relate more to who manages or regulates the activities, so birthrights and deathrights apply only marginally here and the assignment of responsibility for the tracking of changes to these activities relates more to the management of periodic updates than to the tracking of individual birth or death events. Some events have actual births (which fall on the day that they occur). Events also

482 The so-called GASB-34 accounting guidelines will make such recordkeeping more and more commonplace for financial purposes (http://www.gasb.org).

have surrogate parents, such as "masters" of ceremonies[483], parade organizers, and the like. In these contexts it is possible therefore to attribute the responsibility *in loco parentis* to some individual or agency. Other dynamic phenomena are unpredictable and fickle, so they need to be adopted by some agency that is put in charge of monitoring and managing the activity, from traffic to land use, to economic development.

[relational linking]

From the moment the tree is born into the municipality (e.g. from the moment it appears on a shared web-GIS system), until its "official death", different departments will be able to attach assorted pieces of information to this tree, by creating relational databases that can link to the tree through its unique code assigned at birth. The ability to link up information to an object may depend on permissions granted by those who have jurisdiction over it. Overall accords about "read" and "modify" rights might also be agreed upon by a committee that coordinates the creation of the City Knowledge system[484]. In general, linking ought to be universally permissible to anyone who wants to attach his/her datasets to publicly-owned properties.

[open-code doctrine]

What we are proposing here is that everybody ought to call each object by its "official name" and, at the bare minimum, all of the reference codes and perhaps the corresponding map layers should be completely in the public domain. This "open-code[485]" (and open-layer) doctrine will foster the independent development of several proprietary datasets developed by individuals, academics, NGOs, private companies as well as other government agencies, with the major benefit of creating an underlying capability for sharing that does not exist today. Nobody will have to invent new names for these objects, nor will there be any need for people outside of the municipality to create their own GIS maps of the objects[486].

[read, write and modify permissions]

["field-level" permissions]

Beyond the sharing of mere codes and spatial locations, selected portions of municipal datasets may also be made directly accessible to outsiders, through a system of permissions. Authorized users could be granted "field-level" permissions that would specify exactly which fields in which database are accessible to whom, for reading, writing or modifying[487].

what's new about these jurisdictions

The originality of the concept of "informational jurisdictions" that I am proposing is not so much in the insistence on the definition of boundaries between departments, which already exist and may actually be a potential detriment to the achievement of City Knowledge. The concept is

---

483 In Italy we call the main organizer of an event the *padrino* or *madrina* (depending on the gender).

484 See section on "Overarching Standards" starting at page 219.

485 The concept of "open code" is akin to that of "open source", where programming code is made transparent to users who are thus encouraged to modify it and in turn share it with others in a never-ending virtuous cycle of successive improvements. The Linux operating system is probably the best known of the open-source applications circulating in cyberspace. In this paper, the phrase "open code" actually refers to the sharing of reference codes, i.e. nametags, that uniquely label real-world objects subject to municipal maintenance, management or planning.

486 The Ordnance Survey in the UK is a leader in the development and distribution of "mastermaps" that provide the entire country with uniform base maps and unique topographic identifiers (TOIDs) as a universal foundation for the sharing of spatial information (http://www.ordnancesurvey.gov.uk). More on TOIDs on page 222.

487 See section on "Information Sharing" starting at page 214.

more subtle than that. The modifier here is key. The "informational" borders around elements of our urban landscape that I am proposing are more specific – more spatial – than the generic departmental boundaries that already exist. In some towns all information may be considered the purview of the MIS department, if there is one. In smaller towns, information may simply not be on the agenda at all, so that all transactions with the public may be limited to acquiring "documentation" to support an impending decision. I propose to delegate the informational responsibilities to front line offices that oversee the "birth" of the element if possible.

[empirical basis]

We proposed this information-centered approach to defining jurisdictions in Venice as well as in Massachusetts. In our experience, though, it has been applied not as a "top-down" agreement among agencies defining the formal areas under each agency's control. Despite the lack of coordination from the top, the intrinsic nature of departments creates *de-facto* spheres of influence that generally match the informational jurisdictions as defined herein. Some of the clearest jurisdictional boundaries in Venice are those of Insula, that was created specifically with the charter to maintain the canals and infrastructures of the city[488] and the *Commissario al Moto Ondoso* who was appointed specifically to overcome traditional jurisdictional overlaps in dealing with the traffic problem in the canals and lagoon of Venice.

In the U.S., we really focused on jurisdictions with a recent project which explored the application of City Knowledge principles in the Environment Department of the City of Boston[489]. By analyzing the information flows in the key processes of the Boston Conservation Commission and Boston Landmarks Commission, the team was able to propose streamlined on-line procedures for the management of permit applications. Although this system has not been implemented yet, it is one of the better examples of a real-world application of this and other City Knowledge principles[490].

[similar to other efforts?]

Although several scholars have recently discussed similar concepts[491], the primary novelty of our approach lies simply in the renewed focus on information and space. In this *information aware* context, we can then be more overt about assigning jurisdictions over the accrual and updating of the data. In the simplest situations, my proposal gives primacy to the spatial aspects and therefore assigns the jurisdiction to the front-line offices that are actually interacting with the physical world and are doing the hands-on installation or "creation" of the object in the real world. Next in line for bithrights – in case no municipal department is directly involved in the birth of the object – would be the department that last authorizes the

[different/better?]

488 Though there are still some gray areas especially around jurisdiction over sewers, which is causing friction between Insula and the Venice department of Public Works.

489 Hart *et al.*, 2004.

490 Another promising example is being developed in Venice (Novello and Sartori, 2004). See also footnote number 498.

491 Craglia *et al.*, 2004 (p. 61) specifically mention a concept very similar to the one presented here. See also the "temple" vs. "triangle" structure discussed in Reeve and Petch, 1999, pp. 155-156.

creation/installation of the structures, or the department that is in charge of managing the activities.

[change what and how?]

To abide by this tenet, towns would simply need to internalize the concept of informational jurisdictions so that each department can assess its own domain of control and plan how to bring its own turf under a sustainable City Knowledge regimen. The nuance here is that the "parent" office or department is not necessarily the one that ends up carrying the burden of the automating/informating tasks. This is simply the place where the "naming" and "coding" of an object takes place, coupled perhaps with the assignment of appropriate permanent attributes to the newborn.

It could well be that the actual data entry and GIS mapping will still take place at a central repository – like an MIS or GIS department – for larger towns, or even outside of the municipality – for example in regional planning authorities – in the case of smaller towns[492]. In other words, the parent department – after naming the child – may decide to relinquish its care to a "guardian". Yet the paternity would always reside in the department/office that actually oversees the modification of the world we live in. The original parent department would be ultimately responsible for the information about such modifications, not the "guardian" that oversees the actual computerization of the data.

The main departure from currently popular practices is the attribution of a special importance to the exact place of birth and death of administrative objects, so that appropriate jurisdictions can be drawn up using consensus approaches based on the locus of such administrative events. The politics of such a consensus approach would surely be quite intriguing and potentially detrimental to the success of this aspect of the City Knowledge approach, yet there will be plenty of uncontested jurisdictions that are unequivocally already under the sphere of influence of a specific office or department. Of course, we should start organizing urban data within these clear-cut, *de facto* domains before we get into the more controversial ones.

[492] See footnote number 437

Atomic Distributed Knowledge

The way in which we have proposed to assign the responsibility over the information relating to the birth and death of each object in a city to the most appropriate department, and not to some single overarching entity, betrays another fundamental tenet of City Knowledge – namely that knowledge should be acquired and organized in a distributed manner and not though a centralized effort[493]. Top-down efforts like that of the Ordnance Survey in the UK[494] are appropriate insofar as the mapping of fundamental topography is concerned, especially when there is a huge backlog of territory that is still unmapped via GIS and unrecorded in computerized databases. But, in the long run, any form of top-down, centralized control over the continuous updating and upkeeping of the information infrastructure is doomed to fail since the real data that captures fine-grained local change originates from the frontlines of local government and not from ministerial headquarters in the nation's capital.

capturing fine-grain at the frontline

Just like ants go about their daily business guided by simple rules that control their individual tasks at the local level – yet the ant colony as a whole displays a macrobehavior of uncanny intelligence – so front offices in our cities and towns micromanage piecemeal change at its finest grain within the civic sphere and in so doing almost serendipitously produce an aggregate effect that translates into macroscopic change over time. Private actions that affect the public realm are always filtered through local branches of government. Only few "big" projects that affect our physical territory are decided and implemented at the state or national level[495]. The majority of change is managed through our local city halls, which is why I propose that the way to create – and more importantly to perpetuate – City Knowledge is through a concerted effort at the municipal level[496], and more specifically at the level of each department or division within a municipality[497].

micromanage local piecemeal change

Working within the confines of the department, the informational needs of each of the internal operations should be analyzed and a priority sequence of small incremental projects could be planned to make gradual progress towards the final goal[498]. Simple informational tools will be developed, at a very gradual and sustainable pace, to address specific operational needs, while data are gradually accrued through a variety of low-cost methods that are progressively refined and optimized to yield the maximum informational return using the minimum financial and human resources[499].

iteratively maximize informational returns

[gradually accrue data]

[minimize cost]

---

493 In line with Nedović-Budić, 2000, who support the development and maintenance of local databases (p. 87).

494 See footnote 486 and page 222.

495 Interstate highway systems, water supply projects, and federal buildings come to mind, as do military reservations and coastal developments related to navy yards, as well as the national and state park systems.

496 This should be taken to include townships and counties, for the more sparse parts of the world.

497 Tulloch and Fuld, 2001. After Hart *et al.*, 2004, I have been also advising an undergraduate diploma thesis that entails a thorough analysis of "information flow" in the management of permits for the occupation of public space in Venice (Novello and Sartori, 2004).

498 This is not a new approach at all. See chapters 5 and 6 in Reeve and Petch, 1999.

499 Cost savings are reported by Budić (1994) and Nedović-Budić and Pinto (1999), p. 55.

[maximize human resources]

This decentralized approach to the accrual of city knowledge lends itself to contributions by forces outside of the municipal sector. In fact, the work of volunteers[500], students of all grades (K-20)[501], scholars[502], and professionals[503] could be harnessed and incorporated – with proper validation – into this emergent, distributed information system.

what's new about atomic distributed data

The grain of data a city collects has always been controlled by the balance between costs and benefits. The fine atomization that we were able to employ when collecting urban data at the Venice project center since 1988 would have been prohibitively costly even for wealthy western communities until very recently. The middle-out approach suggests that we devolve data upkeep to the front lines, hence implying a distributed architecture for the ensuing information system. Neither of these concepts are particularly revolutionary, but they were simply not cost-effective at the department level until very recently. The falling cost and the increased capabilities of hardware and software have made a fine-grained, atomized approach to city knowledge affordable and hence feasible. Although there will be coordination and synchronization costs, they are beginning to be offset by the benefits of this approach.

[empirical basis]

In Venice, we have demonstrated that sizeable components of the urban realm can be systematically and exhaustively collected, with patience and with proper information design and knowledge acquisition methods[504]. Our inventories have shown resilience to technological change, as we migrated them through several generations of software and hardware tools. They have also shown flexibility and re-usability as demonstrated by our plan-ready applications and as emphasized by our plan-demanding cases.

Although we have been operating in Boston only since 1999, we already have achieved considerable success in promoting meticulous comprehensive inventories of Cambridge's curb-side parking regulations and parking meters[505], as well as of Quincy's public buildings[506] and Boston's parking facilities[507].

[similar to other efforts?]

As mentioned, the closest example to this approach, in terms of the distributed and emergent nature of the system is the *Digital Earth* effort[508].

---

500 As the over 200 Earthwatch volunteers who were instrumental in the rapid completion our Public Art projects discussed partially on page 110.

501 Like the over 500 WPI students who came to Venice and Boston over the years. Or the 1000 middle-school children who measured the hydrodynamics of the entire lagoon simultaneously under my direction (Carrera, 1998).

502 As am I planning to do with the emergent transcription system discussed very briefly on page 111..

503 As we proposed to Cambridge to harness traffic reports from consultants. As another example, I think surveyors should be enticed into some submission requirement (by the county-level Registry of Deeds) to make our cadastral system sustainable.

504 See for instance the bridge, dock, public art and canal catalogs in Part II.

505 Cullen *et al.*, 2002; Flynn *et al.*, 2003.

506 Blizard *et al.*, 2004.

507 Allard *et al.*, 2001.

508 Crockett, 1998 and http://www.digitalearth.gov. For more examples of similar systems, see the list of links at http://www.digitalearth.gov/analogs.html.

Although it is focused on earth sciences at the planetary scale, it reflects all of the main tenets of city knowledge, making it a true emergent system, albeit at a different scale than my municipal approach. It has the same flavor in terms of distributed cooperation between independent agencies, but a different – much larger – grain.

The Federal Geographic Data Committee (FGDC) seems to promote more of a centralized clearinghouse concept at this time, though there is an overall distributed approach to the collection and organization of the fundamental framework datasets: geodetic control, orthoimagery, elevation, hydrography, governmental units, and cadastral information. Connected activities such as the Geospatial One-Stop and "The National Map" [509] are producing appreciable results with many local initiatives being spawned every month in local areas. Yet these approaches all hover at a scale and resolution that is lower than the urban fine grain that I propose.

Other initiatives from the bottom-up, such as the various Neighborhood Knowledge initiatives[510] and the National Neighborhood Indicators Partnership (NNIP)[511] also resemble my distributed approach, but focus on socio-economic indicators and not on the physical elements that municipalities also require information about in order to conduct routine maintenance, exercise proper management and produce sensible, well-informed plans.

[different/better?]

It is easier to find initiatives that resemble what I propose by looking at examples at the municipal level, though the documentation about the specifics of each city's implementation are hard to track down, making a comparison with my proposal difficult if not impossible. Cities like Vienna[512] or Philadelphia[513] demonstrate some of the more advanced municipal information systems and strategies, though the implementations seem to still betray a dominance of top-down approaches.

My approach is a hybrid that combines the emergent and gradual approaches of the federal efforts like "Digital Earth" and the National Spatial Data Infrastructure (NSDI), with a more bottom-up ingredient similar to the neighborhood data efforts. My focus is at the municipal level and even more specifically at the level of departmental offices. I think that my approach, now that it is technically and economically feasible because of technological developments, promises to be more sustainable since it counts on the finest grain of urban data to produce the higher level information and the second-order knowledge that many of the other initiatives already focus on.

[change what and how?]

The way forward for an interested municipal department would be to pick a low-hanging fruit and begin the process of creating a municipal framework into which to plug urban data as they are collected from now on.

509 See http://www.geodata.gov (last accessed 9/7/04).

510 See http://www.urbanstrategies.org, http://nkca.ucla.edu/, http://nkla.ucla.edu (last accessed 9/7/04).

511 See http://www.urban.org/nnip/ (last accessed 9/7/04).

512 Wilmersdorf, 2003. See also footnote 34.

513 See http://www.phila.gov/mois/index.html.

The basis of my distributed approach are the aforementioned informational jurisdictions, so an ambitious office could get the ball rolling as soon as it identified a suitable first project. The thorny issues of coordination of distributed agents, synchronization, and replication that are standard fare in the MIS and IT fields will eventually need to be resolved through the top-down coordination phase of my City Knowledge approach. Meanwhile, the bottom-up, high-resolution data collection can be started at any time, provided that a systematic and exhaustive approach is followed, in accordance to the lessons listed in parts II and III.

For those cities and towns that are already collecting and mapping urban data, the change would be more in the direction of finer grain, richer attribute sets and exhaustive and systematic data collection. For such cities, the next step would be to work on the more advanced aspects of City Knowledge that are described in the sections that follow.

[perpetual updates]

To truly bring each small project to completion according to City Knowledge principles, the information system will not simply include an inventory of all pre-existing assets up to today, but will also include a mechanism for maintaining such an inventory up to date in perpetuity.

The next chapter explains how we envision these updates could take place semi-automatically whenever a change happens in the real world.

SUSTAINABLE UPDATES

Once cities embrace City Knowledge principles and systematically collect and organize data about the various elements that make up our urban world, the next hurdle is to devise methods to keep the information up to date from that moment on[514]. The provision of perpetual mechanisms for updating the data in an information system is one of the distinguishing features of a true City Knowledge application. Such a system not only provides a user-friendly multimedia GIS interface[515] to aid the municipal end-users in their day-to-day urban maintenance, management or planning activities, but also incorporates the data-updating mechanisms in the system from its inception and not as an afterthought[516]. Some of these mechanisms do not necessarily entail purely technological solutions, but often require a combination of technology, together with appropriate changes in policies and procedures. In fact, akin to what was discussed earlier[517], there are at least five main ways to achieve a sustainable level of informational upkeep without massive financial investments:

1. Intercept administrative transactions (e.g. permits);
2. Force contractual updates (e.g. force contractors to return up-to-date information);
3. Change job descriptions to include "informational returns" (e.g. make information updates officially part of the "job" for civil servants);
4. Exploit free or inexpensive labor such as students, interns and volunteers;
5. Budget and plan for periodic updates, particularly for dynamic activities that change over time;

intercepting administrative transactions

As I pointed out at several junctures in this treatise, there are few – if any – modifications to the physical realm that we live in, that are not in some way preceded or accompanied by an administrative act that results in some sort of entry in the public record. The obvious exceptions are private modifications to one's property, when they do not require authorization[518]. The existence of such a paper trail is almost guaranteed to exist in the realm of "structures", whereas conversely it is highly unlikely to exist in the arena of "activities". Unfortunately, the availability of these transaction logs is not routinely exploited as an information source through which a City Knowledge system can be maintained up-to-date[519]. Not only are "births"

---

514 See Nedović-Budić and Pinto, 1999, p. 58, under "Responsibility".

515 Like the ones shown in Part II and Part III of this paper.

516 In Venice, our *SmartInsula*, *EasyBridge* and *EasyDocks* systems all included updating mechanisms. In fact, the proof of the success of our update mechanisms is that the data these systems contain today are different from what we originally delivered to Insula S.p.A. and the City of Venice in the late 90's.

517 The discussion on page 166 was focused exclusively on the "Catching up with the backlog" aspect of City Knowledge accrual, whereas here we are discussing the subsequent upkeep of the accumulated knowledge, hence the different slant of the section.

518 In Spencer, Massachusetts (and probably elsewhere) these projects are commonly called "ANR", meaning *Authorization Not Required*. Despite the name, many projects that claim to be ANR are still reviewed by planning boards and/or zoning boards in order to ascertain whether authorization is indeed required or not.

519 Despite efforts by researchers such as Coulton *et al.*, 1997.

and "deaths" recorded in these archives, but so are also any subsequent piecemeal modifications, corrections, adjustments and even some of the maintenance performed on the structure – particularly if it is government-owned[520]. Therefore, it seems obvious to me that intercepting existing administrative data streams[521] ought to be the primary means to extract updated information out of the documentation that is associated with acts that require government oversight and are therefore already a matter of public record stored in some municipal recordkeeping system.

Beyond tapping into these administrative records, the operating principle would otherwise be to try to shift the burden of the maintenance of information to outsiders who have an interest or an obligation in the upkeep of such knowledge. The next four sub-sections describe some of these ways.

maintenance-based updates

As mentioned before, there can be many creative ways to incorporate information updates into routine maintenance activities. The tree[522] and light bulb[523] examples are exemplary of these "maintenance-based updates" that should become standard fare in future contractual negotiations and in the language of outsourcing bids.

[contractual obligations and bids]

information-conscious job descriptions

Next the focus ought to be in modifying the Standard Operating Procedures (SOP) internally to include information maintenance. This way internal staff will be made conscious of the importance of information in all aspects of municipal functioning.

free or inexpensive labor

Lastly we ought to look at untapped *pro bono* resources as a final source of low- or no-cost sustainable updates by considering ways to harness the power of volunteers, students, scholars and professionals.

budgeted updates

In addition to these no-cost or low-cost mechanisms, one can also envision dedicating some municipal funds to support additional ways to keep city knowledge current, through focused programs that would fill-in wherever there might be informational gaps left, and validate the data, after these other inexpensive venues have been fully exploited.

Regardless of the method employed – and it may well be a hybrid of the ones above – once these principles are adopted, the imperative will be to never waste any opportunity for updating our city knowledge from that moment forward.

what's new about sustainable updates

As was the case for all of the previous foundation elements, this too is not a particularly novel principle. So-called "lifetime" models of information systems implementation have long included provisions for the updating of the underlying datasets[524].

[empirical basis]

Our own information systems for the maintenance and management of docks (*EasyDock*) and bridges (*EasyBridge*)[525] incorporated screens for the

520 Some of this record-keeping will be mandatory and may even become standardized by the GASB-34 accounting requirements. See footnote 482 and http://www.gasb.org.

521 Coulton *et al.*, 1997.

522 See page 127.

523 See page 185.

524 Reeve and Petch, 1999.

525 See Part II.

recording of conditions and for the logging of maintenance activities. In Cambridge, the parking meter collection crew now directly maintains its electronic log book of jams and it – not the Public Works department who installs the meter stand post – is in charge of updating the meter information whenever a head is installed or removed.

[similar to other efforts?]

Of course, everyone expects that data will somehow be kept up-to-date, so traditional MIS literature always includes considerations about information upkeep[526]. The interception of transactions as a means to achieve reliable updates is also not new[527]. Traditional "waterfall" models of information system development[528] always envision a data maintenance and review step at the end of the waterfall.

The neighborhood indicator programs (such as the NNIP and NKLA efforts)[529] have demonstrated that it is possible to tap into statistical or scientific data sources reliably and repeatedly. Stubborn translation issues could be taken care of with the middleware that Ferreira envisioned[530]. Local frameworks[531] could also facilitate the upkeep of the data as the grain gets finer and finer. The GASB-34 accounting mandate[532] and the spreading of asset management tools may soon make the upkeep of data about physical elements of the city more commonplace too.

[different/better?]

The sustainable updates I suggest here differ from other approaches to data upkeep primarily in focus. One view may be that the "low hanging fruit" here are the slowly changing elements of the physical environment so that a simpler system to tap into government data sources[533] may be devised and implemented. Another view may assert that the best return on investment (ROI) is more likely to come from an application where updates are frequent and the process is important to the city. In my view, data maintenance needs to be a fanatical pursuit. There is no point in developing a comprehensive municipal information system if the data are going to be obsolete the moment the system is unveiled.

Another fine distinction is my insistence in leveraging outside self-interest to keep the records up-to-date. Towns already force developers to pay for the services of a planner who will support the town in its deliberations on the developer's project. I therefore suggest that the self-interest of developers could be exploited also to delegate data entry and GIS mapping to them since they are the ones who indeed will be changing the real world out there anyhow.

One potential benefit would be the ability to collect backlog data as well as new updates using essentially the same method. This method may

---

526 See for example Laudon and Laudon, 1996.

527 See Ferreira, 1998 and 2002.

528 Reeve and Petch, 1999, chapter 3.

529 See the Introduction (pages 11 and 16) as well as later, on page 189 and elsewhere in footnotes.

530 Ferreira, 1998.

531 Tulloch and Fuld, 2001.

532 See also footnotes 482 and 520.

533 See Coulton *et al.*, 1997.

be a composite triangulation of the five paths discussed earlier in this section[534].

[change what and how?]

The course to take in order to establish a solid and sustainable system for reliable data updates needs to proceed opportunistically[535] starting from the more cost-effective updates, as evidenced by an internal assessment of information flows and information sources in standard administrative processes that entail spatial decision-making or analysis[536]. Since we believe that sustainable data maintenance mechanisms are necessary conditions for the longevity of an urban information system, a true City Knowledge system should never be conceived without making provisions for keeping the information current. By including update considerations in the department's data collection strategy it may be possible to exploit possible synergies with the collection of the backlog information so that old and new data can be collected using the same seamless procedure.

Using the five tenets described earlier, the department could revise its requirements[537], modify its forms[538] and generally shift the burden of data upkeep to interested third-parties who may not mind the extra burden as part of doing business with the town.

534 Page 210.

535 Reeve and Petch, 1999, p. 156; Barr, 1991.

536 As was done in Boston by Hart *et al.*, 2004 and repeated in Venice by Novello and Sartori, 2004.

537 As we suggested in to the Traffic dept. in Cambridge (Gage *et al.*, 2003) and to the Boston Air Pollution Control Commission (Allard *et al.*, 2001).

538 As proposed to the Boston Fire Department (O'Donnell *et al.*, 2002) and to the Boston Environment Department (Hart *et al.*, 2004).

INFORMATION SHARING

Until some degree of sharing is initiated within and among municipal departments, we will not be able to exact the powerful, value-added benefits of City Knowledge that allowed us to quickly and easily conduct many second-order analyses yielding unforeseen plan-demanding results as described in parts II and III. Nevertheless, sharing of the "meatier" datasets, beyond the mere publicizing of the "official" reference codes and the availability of the GIS layers with the objects' positions, ought not be mandatory but rather voluntary. Nobody will be forced to share data with anybody else, unless there is either an institutional mandate or a desire to do so on the part of the rightful "owner".

voluntary sharing

institutional sharing

mandatory sharing of public records

Some forms of sharing will be more or less compulsory, based on pre-existing institutional requirements. For instance, public records access laws[539] and "right to know" would constitute a mandate to share[540]. The most immediate type of institutional sharing is the one that takes place within a department or division. Intra-departmental sharing is a patently obvious form of sharing that ought to take place within municipal organizations, for very apparent reasons. Yet, in reality, the level of sharing that occurs even within small organizations is surprisingly low, despite the intuitive expectation to the contrary. I have briefly mentioned the benefits that the city of Cambridge's Traffic, Parking and Transportation department has reaped from our recommendation that the parking control officers share the daily log of parking meter jams with their colleagues (in the same department) of the meter collection crew[541]. Yet this type of intra-departmental sharing was not happening before our project.

[intra-departmental sharing]

[inter-departmental sharing]

The next level of potentially mandatory sharing could take place between two departments that must communicate information to each other as part of their institutional duties. For instance, the building inspector must consult with the conservation commission about possible restrictions on developments that are near wetlands[542]. In this case, the conservation commission is obliged to share its information with the building department[543]. Likewise, I already discussed how local landmarks commissions must notify the inspectional services (or code enforcement) department whenever a particular building becomes listed as a registered historic property, since different rules may apply in relation to building permits or codes[544].

Intra- and inter-departmental sharing represent the "low hanging fruits" wherein information exchanges could be quickly mainstreamed through a GIS-based, permission-enabled, distributed urban information system.

---

539 Like Massachusetts 950 CMR 32.00.

540 This may also be a hook that state or regional agencies can use to induce smaller agencies at the municipal level to share information with them, so that they in turn can make the appropriate information available to the public at large in accordance with the law.

541 Flynn *et al.*, 2003. See footnote 361.

542 See Hart *et al.*, 2004.

543 Also known as: "building inspector", "construction dept.", "inspectional services", "code enforcement" and others.

544 Hart *et al.*, 2004.

extra-mural sharing

[sharing with public-private companies]

[asynchronous alerts]

[two-way sharing]

The next step beyond intramural sharing, would entail sharing with "outside" institutions, in situations where contractual mandates exist, or where the mutual benefits make sharing a desirable process. For instance, the electrical utility company (whether or not it is owned by the municipality) has an interest in keeping track of the interference of tree branches with its power lines. Power companies will trim branches of municipal (and private) trees to protect their grids. The Parks Dept. may be in a position to alert the utility company when it detects, through its routine maintenance activities, dangerous situations that may negatively impact the electrical wires. An "asynchronous alert" may take the form of an updated tree record wherein the distance of the tree branches from the closest power line is simply modified as a result of a visual inspection by Parks department staff. This modified record could be shared – through the granting of appropriate field-level permissions – with the electric company which in turn would have its own (asynchronous) built-in middleware mechanism[545] for flagging the "dangerous" distances and thus would immediately recognize the problem and send out a bucket-truck to trim those branches at the first opportunity.

[informational *quid pro quos*]

The trimming of the perilous branches should then be followed by an update, this time by the utility company itself, of the distance-of-branch-closest-to-power line parameter in the shared municipal tree database, to reflect the change that just occurred through the pruning intervention. This second-round of sharing would thus bring the process full-circle, demonstrating how two-way sharing could be used to leverage these mutually beneficial barters of *quid-pro-quos*[546]. Similarly, within the municipal confines, the Public Works department could also be allowed access to the municipal tree databases to keep track of the damage to sidewalks created by the roots of the same trees.

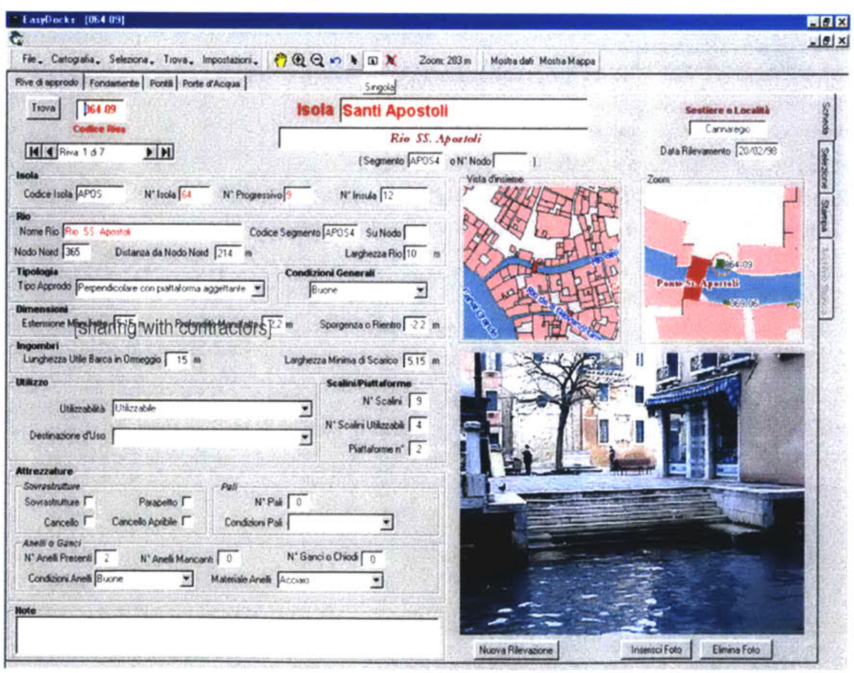

[sharing with contractors]

In Venice, our *Easy Docks* information system (below) for the management of boat docks was immediately shared between the Public Services department and a public-private company (ARTI) that was in charge of the physical maintenance of these municipal assets. The system is still in use at the time of this writing (2004). ARTI (a

545 *À la* Ferreira, 1998.

546 See Nedović-Budić and Pinto, 1999, p. 58, under "Incentives".

"contractor") was expected to update the dock's condition log as soon as a repair was made. Similarly, other outside contractors who are hired to prune trees could be required (by contract) to keep track of the maintenance done to each tree. Contractors could also be obliged to provide updates on the growth and health of the tree. All of these forms of contractual sharing could be achieved through appropriate setting of permissions in shared online databases[547].

[sharing with the public]

States have so-called Public Records Laws that mandate access to public records. The parents of all state laws are the federal Freedom of Information Act and the Privacy Act of 1974. Frequently though, instead of making raw data available to the public, cities prefer to provide pre-screened information – once-removed from the fine-grained original datasets. Thus, a filter is created that allows complete control on the interpretation of the data by precluding independent analysis of the raw facts. Whereas, in the majority of cases, converting data to information represents a major leap in sophistication – one that allows us to conduct second-order analyses that would otherwise be rarely performed due to the difficulty that is generally encountered when just trying to get the basic data together – nonetheless some analyses may be precluded if undigested data are not made available in addition to pre-digested information.

from data to information

For instance, the Environment Department may become interested in determining the energy savings that the urban forest is providing to the city[548]. An adequate energy audit could be easily achieved thanks to the accumulated knowledge about the size of each tree's canopy that would be plan-ready once all of these operations are coordinated around the shared fundamental information references represented by the tree's location (on a GIS layer) and its ID[549]. If the disaggregated, fine-grained tree data were not made available, such calculations would not be possible.

[plan-ready information]

[horizontal and vertical sharing]

The ultimate extensions of this hierarchical sharing scheme involve both horizontal sharing with other cities and towns – probably in the context of regional planning efforts – as well as forms of vertical sharing, both internally in the organization – from the front lines up the management ladder and up the chain of command to the executive branch – as well as among government agencies at different scales, i.e. metropolitan, regional, county, state and federal levels.

[the technology of sharing]

Technically, the act of sharing information among different providers and users can be supported by a variety of client-server or even peer-to-peer architectures. The possibilities run the gamut from file-servers that act as data warehouses and allow sharing through files and network applications[550], to web-GIS applications that allow interactions with shared layers through regular browser interfaces (usually supported by client-side

547 As was done for the "sudden oak death" project described in Kelly and Tuxen, 2003.

548 By allowing us to cut our AC use in the summer, thanks to their shading, trees save us money, as they do in the winter by lowering our heating bills thanks to their wind-screening ability.

549 Refer back to pages 125 and ff.

550 Like in the Citrix system in the City of Worcester.

Java applets)[551], or through web-enabled multimedia client applications that can tap into shared layers that are accessible through some internet service and make these updated layers available within a custom, client-side application.

what's new about our information sharing

Not surprisingly, this final aspect of City Knowledge is also well discussed in the literature of MIS, IT, and GIS. Full-fledged geospatial sharing mechanisms are far from becoming commonplace though. Web-GIS prototypes are more and more common, although their effectiveness in day-to-day municipal operations is not so obvious. More commonly, enterprising departments are sharing GIS and Database files through a common file-sharing system, sometimes even through network applications that allow simultaneous access and modification rights to layers and data in real time, without local copies[552].

[empirical basis]

We have dabbled only a little with institutionalized information sharing, since we have scant direct control over it, since our role as academics (and even as professional consultants) is generally that of outsiders. We come into an issue laterally and with no say in policy matters, such as institutionalized sharing. Nonetheless, we have successfully completed some prototyping of potential sharing scenarios both in Venice as well as in Boston, Cambridge and Worcester. We have also recommended specific sharing arrangements that in some cases have been implemented, at least on paper.

[similar to other efforts?]

Interoperability is a common buzzword these days. Coordination mechanisms, such as the ones described in the next section go hand in hand with the challenges of sharing municipal data across jurisdictional boundaries. Structure, process and policies regarding data, responsibility, ownership, contributions and incentives have been suggested as a conceptual framework for making progress on this front.

[different/better?]

My approach is not too different from the latest thinking in this arena, but the main focus, one again, is on leveraging the self-serving instincts of a department and favoring sharing where the quid pro quos are evident and easily achieved for instant gratification. Slowly these sort of success stories – I argue – can make more sophisticated forms of sharing less threatening to the "turf-conscious" individuals who still occupy many municipal positions.

[change what and how?]

Very similarly to what was said in the previous section, this final aspect of City Knowledge should also be left to its own evolution, only with a little bit more awareness of the benefits, possibly advertised or highlighted by self-interested individuals – such as planners and decision-makers – who need the more complicated, articulated, intermixed type of city knowledge that can only be obtained once sharing becomes common-place. Once these second-order advantages are appreciated by the higher echelons of the municipal organization, communicative action theory suggests that these

---

551 See for instance our own demo at www.intelligencesoftware.it/unesco/venezia. See also Kelly and Tuxen, 2003.

552 As can be done with the Citrix systems in Worcester and Cambridge.

fairly powerful executives and managers can become part of the lobbying group that will promote – more effectively than planners alone – the paradigm shift toward full *information awareness* by treatin City Knowledge as a true municipal infrastructure effective immediately.

The next section explores this latest assumption and other issues related to information sharing.

INTERAGENCY COORDINATION

In concert with the necessary assignment of object-level jurisdictions, a City Knowledge system requires the establishment of a modicum of interagency coordination to enable the sharing of information if and when such sharing should become desirable or necessary[553]. Standards and reference rules will be minimal and as transparent as they are in the similar infrastructure of the World Wide Web, which is the model of "emergent" behavior that this City Knowledge infrastructure hopes to emulate. These municipal standards will also guarantee the reusability of the data and will allow multiple uses by different departments, which may regulate one or more of the following aspects through memoranda of understanding, guidelines or internal bylaws:

[memoranda of understanding]

[guidelines and by-laws]

- ❖ Reference Codes
- ❖ Formats
- ❖ Methods
- ❖ Databases
- ❖ GIS layers
- ❖ Metadata
- ❖ Software applications
- ❖ Network Access

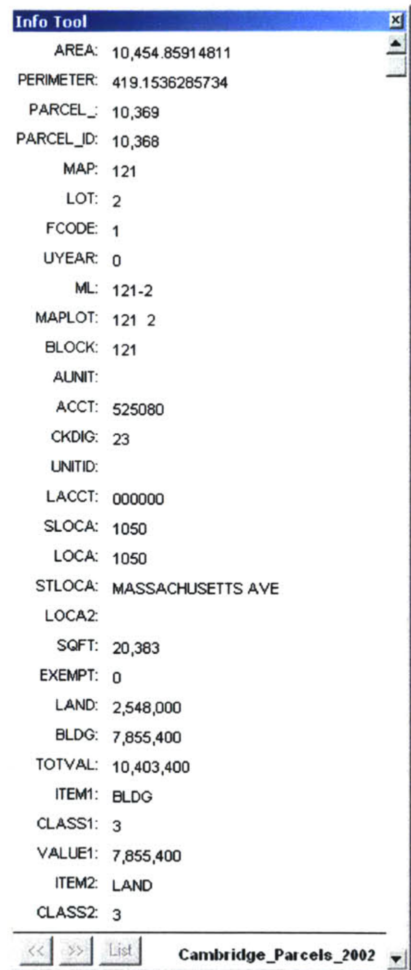

reference codes

Ample space has already been dedicated to the primary form of sharing which revolves around the reference codes, i.e. placetags, that uniquely label real-world objects subject to municipal maintenance, management or planning[554]. Beyond the assignment of jurisdictions to dictate who should be the "parent" to the newborn "city object" and hence have the right to "name the child"[555], there is an overarching citywide necessity to agree on what these identifiers ought to represent (if anything).

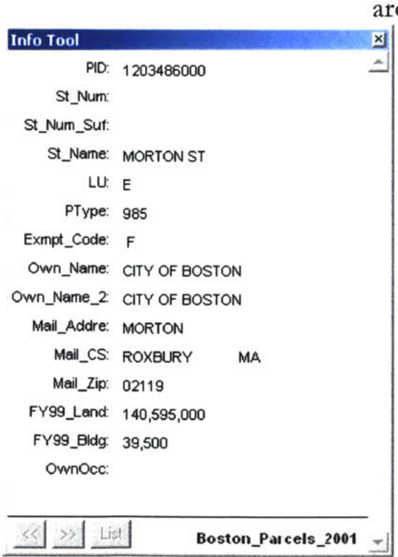

There need to be some basic agreements about

553 See Nedović-Budić and Pinto, 1999, p. 58, under "Data". According to them "interorganizational systems require standards on data models; data formats; data quality; categories of spatial data; contents of specific data layers; metadata; data dictionaries; output requirements; and data transfer."

554 See for instance page 52 and following.

555 See the "Informational Jurisdictions" section on page 201.

the criteria that the various departments ought to adopt in the process of creating standard codes. In general, all codes should be at least unique, consistent and coherent[556]. For instance, there may be an overall agreement that all codes will be always alphanumeric and the code data type will always be character (or text, or ASCII) and not integer or numeric. This uniformity will facilitate linking, when data types need to be specified or manipulated, so no conversions will be necessary on the fly. Another overall agreement may be that no code should exceed *n* characters in length (probably 20 would be an appropriate number). This simple rule would allow everyone to simply arrange for codes of 20 characters to be set aside as linkable entities in the various department databases. Yet another useful retroactive measure to take is to give each data field a proper, clear name that conveys exactly what that field contains. Indeed, it would be useful to get into the habit to do the same whenever files are named as well.

More importantly, the syntax and semantics of the fundamental codes that will be used for sharing across databases and departments ought to be agreed upon in some sort of conference committee with representatives from a variety of departments.

[legacy codes]

Whether or not new codes are introduced to do away with awkward anachronisms, it is always wise to retain all of the possible legacy reference identifiers that refer to each object, in order to maintain backward compatibility with any dataset that may reference the old codes. Perfect backward compatibility may be impossible due to spatial mismatches between the old codes and the new/updated spatial objects, nevertheless this effort should still be made to allow longitudinal analyses with archival records that, despite their antiquity, may still hold significant informational value for the establishment of long-term trends or for before-and-after comparisons with today's data. Conversely, it is equally wise to stop actively using codes whose meaning or origin is lost to current institutional memory.

[code adoption]

[parent department]

When various departments create databases that link to physical objects outside of their birthing jurisdictions, it is essential that they adopt the precise codes that the "parent" department has assigned to those objects. In this way, sharing will always be possible, regardless of whether it is currently desirable or not. Of course, this transition would need to be coordinated as do many other aspects discussed in this section.

[ancillary multimedia documentation]

[codes as file names]

In addition to linked layers and databases, advanced city knowledge systems will frequently contain ancillary multimedia information, such as photographs, videos, graphs, and sounds that provide additional information or documentation about objects, but are generally not incorporated into either the GIS or the database for reasons of efficiency. The way I prefer to deal with these items that are linked with the GIS and DB through integrated multimedia interfaces[557], is to use, as the file names of the ancillary documentation, the same exact codes that uniquely reference

556 These are standard database principles that can be found in any RDBMS textbook.

557 Such as those shown on p. 81 and 101.

each database record, using suffixes to further differentiate between multiple media attached to the same object/record[558].

[syntax]

The syntax of the various codes used as unique identifiers and/or filenames may be the object of a standardization effort within the municipality at a later date. For instance, the city may decide that characters should always be preferred to numbers (even if they are both treated as characters). A possible standard may recommend creating codes from the concatenation of a variety of fields that are already in the database to identify the object. As an example, we used the combination of island code and a sequential number to label each sewer outlet in Venice[559]. The specific syntax used within each individual application need not be agreed upon by the entire city, but generic syntactical "rules" may be part of a citywide standard nonetheless.

Furthermore, smaller committees of interested parties ought to get together to share the specific syntax of items of common interest. Even though the "parent" department has full control over the naming, it may be wise to coordinate the coding syntax with departments that are likely to interact with a specific category of objects. Outside of these object-specific syntax committees, the citywide standardization process would only address general "ways" or "criteria" to adopt in the definition of the syntax of a code. Maximum freedom will remain in the determination of the exact syntax for each specific category of objects within each individual jurisdiction.

[algorithmic code generation]

Sophisticated algorithmic procedures may be devised to generate codes automatically from some implicit or existing parameter that already contains the seed of uniqueness that is necessary to establish a successful referential framework. One such scheme would entail exploiting the singularity of centroid coordinates (for non-overlapping objects) to construct a composite, interlaced unique code that mixes the X and Y coordinates (regardless of the projection used)[560].

558 The two pictures that accompany each bridge record are labeled with the bridge code (e.g. GOLD) + a suffix of A or R to represent the "Arch" or the "Ramp" photo (GOLD-A.jpg and GOLD-R.jpg). We have used a similar system for wellhead pictures, where each well was photographed on average 10 times, from different angles. Of course, if dozens of pictures were needed for a single object, one may eventually revert to numeric sequential or time-stamped suffixes.

559 As described on page 90

560 The reason for the interweaving, as opposed to the simple concatenation of the two coordinates is so that an alphabetical ordering of the objects according to the interwoven ID would not privilege one coordinate over the other, but would combine the two and organize the objects more or less according to their proximity across both dimensions. The disadvantage is that code generation is much more complicated, requiring a "small" program versus a single concatenation operation.

[automatic codes]

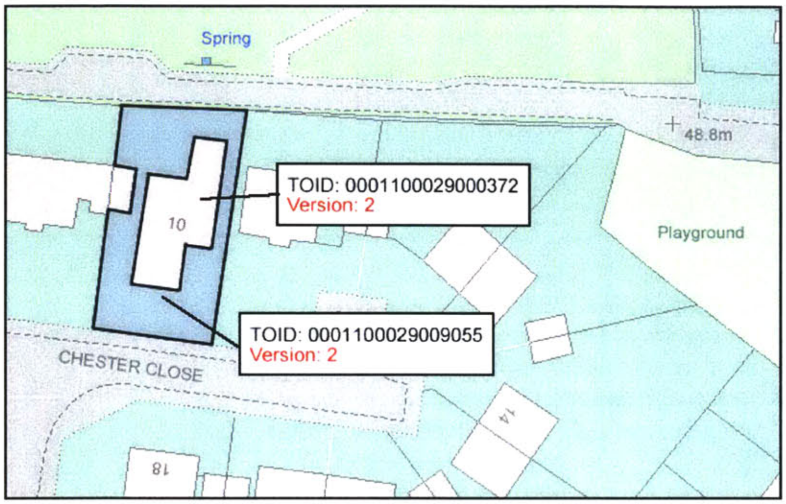

The Ordnance Survey in the United Kingdom has begun to assign automatic codes to all objects that it maps in Britain, by attaching a 16-digit TOpographic IDentifier (TOID) to each mapped object. This rather "dumb" assignment of codes in an algorithmic manner has the advantage of creating an automatic reference framework through which all other agencies can refer to the "real" elements that they are each responsible for. Although some objects may actually be composed of several TOIDs, this approach has the implicit advantage of facilitating sharing.

[code semantics]

["talking" vs. "dumb" codes]

Some of the syntactical rules adopted by a city may embody semantics of one sort or another. For instance, a suffix appended to a code may indicate a sub-partition of the object. A prefix, on the other hand, may serve as a way to identify the "group" of objects that an item belongs to. A syntactical standard accompanying such semantic breakdowns may impose the use of hyphens (instead of underscores, for instance) to separate prefixes and suffixes from the body of a code.

[numeric IDs]

Traditional RDBMS literature[561] recommends the use of numeric IDs that are much more computer-friendly since they lend themselves to effortless sequencing and other algorithmic computations like indexing and sorting. The ease of spatial operations makes "talking", spatially-explicit codes just as easy to maintain with algorithmic precision. At any rate, a compromise can probably be arrived at by providing both a computationally-efficient number and a mnemonically effective code for each object.

Finally, the suffixes that we frequently use to break down sub-units of a bigger whole are always chosen with some logic that embodies semantics. The same is true of the suffixes of file names that refer to images and other media associated with our data. The file name 101_L.jpg refers to the photo of the Lid of wellhead number 101.

At any rate, leaving aside my personal preferences, all I am proposing here is that these issues be discussed and agreed upon across the entire municipality in order to create a standardized framework for future citywide sharing.

formats

In addition to semantically mnemonic and syntactically consistent codes, the formats of the datasets, of the GIS layers, of related files and of any ancillary data ought to also be agreed upon at the city level. Agreement on formats may at first be limited to acceptable and unacceptable file types that departments should uniformly adopt and reject, respectively.

[561] Such as Ullmann and Widom, 2001.

methods

Departments may also decide to standardize the *methods* that field crews or professional consultants will adopt for the collection, archival and presentation of data, information and knowledge.

databases

A little deeper level of standardization may relate to the structure of some of these files, dictating for instance that all Access database files contain tables with names preceded by a numerical sequence number, to organize the contents of the MDB file. Database fields may also be the object of some interagency agreement. Some databases may become standards as a whole and be incorporated into high-level framework datasets.

GIS layers

Some core sets of GIS layers may be standardized, along the lines of what the FGDC is doing for "framework" data[562] so that there may be even compatibility across town boundaries – with abutting towns – even across state lines or vertically, from one level of government to the next higher (city to state to federal).

metadata

[extramural and vertical sharing]

The exact structure of the metadata used within a municipality is one of the aspects of City Knowledge that ought to be standardized on a citywide (or even state or federal) basis. The management of evolving versions of both layers and datasets could be achieved through a strict abidance to the metadata standards that are slowly emerging in the GIS industry. As of this writing, it appears that the most useful approach to this issue would be an interdepartmental agreement on what "subset" of the all-encompassing FGDC metadata standard to adopt[563]. At the very least, in the beginning, the system may rely on the simplest file-system metadata that operating systems already provide in the form of creation/modification dates, owner, permissions, and file size. Basing the municipal metadata on an existing standard will make possible the next level of sharing, beyond municipal walls, with other cities and towns, or beyond the municipal level of governance, with state and federal agencies.

software applications

It may be advantageous for some towns to force the use of specific software applications (e.g. use Mapinfo instead of ArcGIS, Oracle instead of Access) to facilitate sharing and more importantly to cut the cost of software support. Although this approach would certainly save money in the long run, it may engender resentment on the part of those who are forced to switch to an unknown package and may also preclude some "innovation" from happening. Such a level of standardization is no longer necessary thanks to today's highly interoperable software packages[564], so this standardization may not be advantageous except in terms of software licensing and support costs.

network access

Assuming that the data will at some point be shared through a network, there needs to be some coordination about network access, regarding passwords, permissions, quotas etc. These agreements could be coordinated with the department that directly manages network operations

562 Tulloch and Fuld, 2001. See also the FDGC web site at http://www.fdgc.gov.

563 As MassGIS is trying to do in Massachusetts (see http://www.mass.gov/mgis).

564 Mapinfo can read and modify native ESRI shape files, for example. Similarly, SQLServer is capable of reading Access files.

so that everyone can access the system remotely, with the appropriate level of read/write access.

what's new about interagency coordination

Standardization is a necessity that emerges naturally when multiple actors are trying to cooperate toward a common goal. My approach is hardly new or different from the myriad of examples in typical municipal MIS, GIS and RDBMS practices.

[empirical basis]

Both in Venice and in Boston, we developed and successfully deployed infrastructural reference codes to uniquely identify key physical elements of the built environment, such as canals[565], bridges[566], docks[567], pieces of public art[568], parking facilities[569], underground storage tanks[570] and so on. On both sides of the Atlantic, we have also structured spatial frameworks for the archival of dynamic data about such activities as traffic[571], demographics[572] and economic vitality[573]. Many of our standard codes have become *de facto* standards in Venice[574] and in the U.S.[575].

We also standardized photo formats (JPG) and methods (landscape vs. portrait). We structured our Access tables internal to the MDB files in predictable ways and we even developed a set of fundamental GIS layers, similar to the local framework layers proposed elsewhere. We have not, alas, dug deep into the issue of metadata beyond the simplest file system metadata and little more. We have settled on standard software applications, specifically Mapinfo and Microsoft Access for now. We have also shared standard layers internally through passwords and a "secret" web site on WPI's server. A web-GIS prototype is being updated to improve on the promise of the interactive, password-protected system that is already available on the internet.

565 See part II, especially the chapter entitled "The Venice Inner Canals Project", starting from page 44.

566 *Idem.*, see also section starting on page 99.

567 *Idem.*

568 See page 112.

569 Allard *et al.*, 2001.

570 O'Donnell *et al.*, 2002.

571 See for instance Carrera, 1996, 1997, 1999a; Gage *et al.*, 2003; and Farmer *et al.*, 2004 to name just a few. See also page 70 and following.

572 Hamir *et al.*, 2004.

573 Jajosky *et al.*, 2004.

574 The canal, bridge and dock codes that everyone uses in Venice are essentially the ones we developed (Carrera, 1999d).

575 The curb-side regulations in Cambridge, for example, are now identified by our codes (Cullen *et al.*, 2002 and Flynn *et al.*, 2003).

[similar to other efforts?]

Our standardization efforts resemble the initiatives that are being conducted at the national level by the FGDC and Digital Earth and at the state level by the like of MassGIS in Massachusetts. Technically, the metadata standards of the FGDC and the interoperability progress being made by the Open GIS Consortium will probably create the premises for fruitful advancements in these areas[576]. At the local level, there are some movements to define standards for indicators[577] at the NNIP and NKLA. Some cities are more advanced than others at defining common practices and standards and are more effective at coordinating the mapping and data collection efforts of a variety of departments, though the approaches are often either centralistic or strongly hierarchical and implemented from the top down.

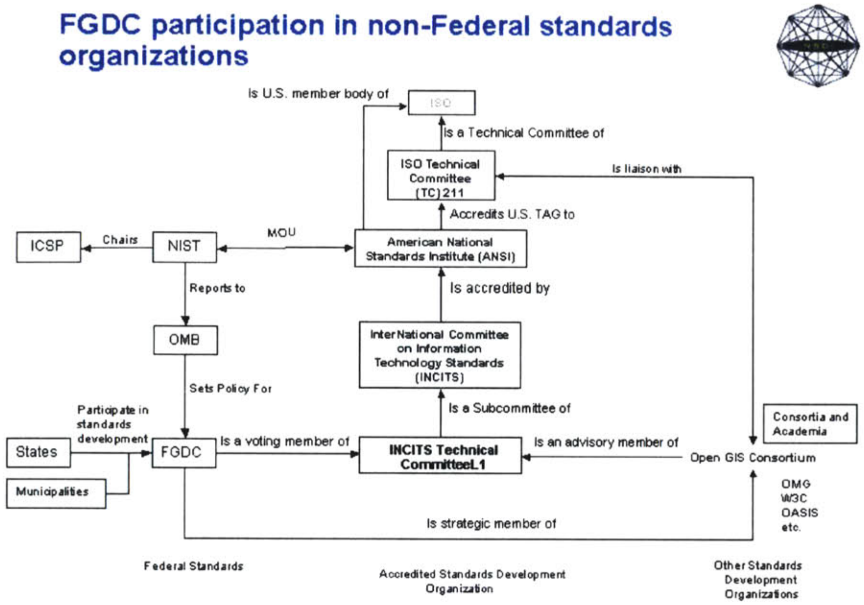

Despite all of these efforts, the systematic codification of all homogeneous GIS objects at least for the fundamental framework layers is still not standard practice even in technically savvy and fairly wealthy cities of the developed world, such as the ones discussed herein. Although the concept of unique reference codes is common knowledge and even common practice in the MIS and RDBMS fields, many GIS layers out there are still CAD-like spaghetti files that look good but are not useful for spatial data archival and analysis. Even when selectable objects are mapped in GIS (instead of the lines and points of spaghetti-maps) they often lack the all-important key ID[578]. The same goes for all of the other aspects that are amenable to standardization, such as formats, methods, databases, GIS layers, metadata, software applications and network access for data sharing.

[different/better?]

My proposal is not fundamentally different from these existing examples, but once again it focuses at a local spatial scale, where high-resolution urban data can be collected from the middle-out. Even though overarching standards are an eventual necessity in a full-fledged City Knowledge system, they do not have to be forced upon reluctant municipal offices from the top. The need to coordinate should be allowed to rise spontaneously from the natural, organic evolution of the municipal information systems and from the inevitable need to eventually interact across divisional or departmental boundaries, with other offices of the same

576 See for instance http://www.fgdc.gov/standards/related_activities.html (last accessed 9/9/04).

577 Sawicki and Flynn, 1996; Coulton *et al.*, 1997.

578 For example, most municipal buildings layers do not have meaningful identifiers that can be linked to other datasets.

or even of a different department, not to mention the possibility of vertical aggregation toward the executive level of municipal affairs or to the state or federal level of government.

Once the usefulness of standards is appreciated first-hand by the font-line offices, standardization will become standard fare itself. Coordination will become second-nature as its advantages become obvious to more and more municipal practitioners.

[change what and how?]

In keeping with the "disjointed incrementalism" of our collective "muddling through"[579] and the opportunistic[580], emergent attitude of my middle-out approach, I do not have a silver bullet to offer to municipalities who may want to engage in standardization practices across departments. Instead, I recommend that, after having embarked in the fine-grained data collection and updating mechanisms described earlier, each department should explore intra-departmental sharing necessities and subsequently inter-departmental interactions that routinely occur in daily administrative processes[581]. Once these links are established and analyzed, I forecast that overarching standards will become blatantly useful and hence will have much more of a chance of achieving "stickiness", which is one of the pre-requisites for reaching a tipping point[582].

The five foundations of City Knowledge: jurisdictions, atomized, distributed data acquisition, sustainable updates, sharing and coordination, together with the middle-out approach, have demonstrated great potential for bringing about a comprehensive City Knowledge system, as the numerous examples in Parts II and III clearly showcased. In this Part IV we have looked at the pros and cons of City Knowledge and laid these foundations.

The following Part V closes this dissertation by proposing a possible sequence of actions to get towns from point A (interest in City Knowledge) to Point B (fully implemented City Knowledge system) and beyond.

579 Lindbloom, 1959.

580 Barr, 1991.

581 As we did with Hart *et al.*, 2004 and Novello and Sartori, 2004.

582 Gladwell, 2000.

PART V:

"The indispensable prerequisite is then a real program[...] If one studies the history of a town, the development of its trade and industry, and if one carefully checks all other statistical data, taking into account local needs, he will certainly have enough information to extrapolate this data into the immediate future, which is sufficient."

Camillo Sitte, 1889
"City Planning According to Artistic Principles", p. 264

THE EMERGENCE OF CITY KNOWLEDGE

CITY KNOWLEDGE AS AN EMERGENT SYSTEM

A well-developed municipal knowledge infrastructure – as I envision it – enables city officials to not only accomplish their "normal" duties very efficiently and with less effort, through the shift from "plan-demanded" data to "plan-ready" information, but should eventually even lead to the spontaneous generation of "plan-demanding" knowledge that will inspire city officials to look into previously-unnoticed patterns or trends that may demand appropriate attention.

The needs and wants of municipalities, which have been computerizing many of their operations, and may have been dabbling with bottom-up GIS applications for the last few years, can now be met with the most recent advances in spatial management tools. Technologies such as relational databases, geographic information systems and the internet are not only affordable but also better suited than previous technologies to the new decentralized, "middle-out" approach for the spatial representation of urban features that is the focus of this paper. My own research and experience suggests that an informating approach that is neither truly top-down, nor merely bottom-up has a chance to be successful in the real world of municipal governance of urban structures and activities.

[high-resolution, fine-grained data]

Until now, the sort of fine-grained data accumulation that I propose[583] was impractical, due to the high transaction costs as compared to the perceived benefits. Today, the costs have come down and the benefits are being perceived more clearly at least by planners, who have otherwise been slow in realizing the potential of current technologies[584]. So the ratio of benefit per unit of cost is improving. Moreover, in a department-level, middle-out structure such as I propose, the primary beneficiaries of an exhaustive urban information system would be precisely those who are in charge of accumulating and updating the system's underlying datasets. The connection between cost and benefit would thus be immediately obvious within each informational jurisdiction[585].

[lower costs]

[improving benefits-to-cost ratio]

[second-order benefits]

Furthermore, the benefits are not limited to the here and now value of the mere "automation" of pen and paper processes, but actually extend into second-order, value-added realms where "informating[586]" leads to plan-demanding situations, such as the ones I presented in parts II and III.

[ripe technologies]

Current technological trends make my approach much more feasible today that it would have been as recently as the mid-nineties. In particular, the loosely-coupled nature of web-based applications, and the overall architecture of the WWW provide a natural infrastructure for my middle-out approach which is now "ripe" for real world implementations.

583 And I have applied both in Italy and in the U.S., as shown in Parts II and III.

584 Geertman and Stillwell, 2003 and 2004.

585 Singh, 2004.

586 The "automating" vs. "informating" dualism comes from Zuboff (1991). These more sophisticated benefits will require coordination to enable cross-department sharing, which would add some cost and complexity.

the emergent nature of City Knowledge

City Knowledge systems are "complex adaptive systems that display emergent behavior"[587]. Just like the use of local knowledge by individual ants informs microbehaviors that create the macrointelligence of the ant colonies, so too bottom-up, department-level data and information collected for dealing with specific aspects of municipal affairs can produce city knowledge that supports higher-order urban analyses and decisions. It is at this higher level that planners are meant to operate, instead of being forced to hunt down the data they need over and over again. Middle-out is my proposed pathway to get us there from here.

So, with costs greatly abated, and demonstrable benefits that reliably exceed the first-order advantages of computerization, I conclude that now is a good time to begin to treat information as a *bona fide* infrastructure and I therefore recommend to progressively create a City Knowledge system from the finest grain out.

analyzing emergence

Even though I have not fully explored the problems and promises of the extension of City Knowledge beyond a single department, many of the department-level information systems that I was involved with – especially in Venice – are now mature enough to be evaluated in such context. Such an evaluation would be a useful follow-up to this dissertation since it would allow us to: identify the practical impediments to the horizontal (and vertical) diffusion of City Knowledge across department boundaries; to better quantify the costs associated with the added coordination necessary to cross-reference datasets from different organizations and to keep them synchronized; and to address some of the transition issues facing municipal departments when they try to mainstream promising technologies into established work flows and entrenched bureaucratic processes. For the time being, I will conjecture what

creating emergence

In this final Part V, I suggest that instead of "dealing with" emergencies, city officials may actually purposely set up an *emergence* of their own. Instead of waiting for the arrival of a magic, all-encompassing system for the management of all municipal information, I recommend a path that promises a non-traumatic transition to City Knowledge in true *emergent* fashion. I hypothesize how cities can start from one specific aspect of municipal affairs and then sustain such a system in perpetuity, gradually spreading this approach to more and more aspects, then infecting other departments and eventually diffusing into other towns, so that this concept will become commonplace in urban communities around the world.

587 Johnson, 2001.

CREATING CITY KNOWLEDGE FROM THE MIDDLE OUT

Geographical Information Systems (GIS) and Database Management Systems (DBMS) – while rather commonplace in today's city government – are scarcely used to systematically keep track of essential urban elements, such as roads, trees, sewers, etc. in an "automatic" fashion. Ad hoc, or "implicit" knowledge is used instead to produce maps and datasets that feed the decision-making process on a need-to basis, case by case. The typical city department is generally too busy with the struggles of day-to-day emergencies to afford the "luxury" of thinking ahead and planning a systematic knowledge acquisition campaign that will last several years.

Yet, this study recommends that cities implement just such a long-range infrastructural approach to city knowledge. How can one realistically expect to accomplish such a demanding – albeit promising – task, when the average city official is burdened with myriad challenges just to perform routine activities that the administration of the city demands? This dissertation proposes to towns to embrace the piecemeal, bottom-up, opportunistic evolution of municipal information systems that results from our "muddling through", but to do so strategically and deliberately, with an eye to the overall role of each department's information in the combined big picture of City Knowledge.

In this section, we take a look at a plausible implementation path that cities and towns could adopt to harness urban data used for day-to-day maintenance into plan-ready information that can facilitate more complex management tasks and can lead to the kind of deep knowledge that can inform longer range plans for the future of our communities.

The basic steps that an office or department could follow iteratively to move toward a City Knowledge system are:

(1) commit to a City Knowledge approach to the development of a comprehensive municipal information system and accept the middle-out strategy as a possible tool to achieve it;
(2) define its informational jurisdictions by analyzing information flows;
(3) identify "low-hanging fruits" within its jurisdiction – i.e. applications that would yield the maximum benefits at the lowest cost;
(4) collect "atomic" data about the urban elements related to the low-hanging fruits;
(5) develop and apply a mechanism for sustainable data updates;
(6) share the information and data as appropriate;
(7) coordinate as needed.

In the sections that follow, we first discuss the need for a commitment to the development of City Knowledge, which can be manifested by an official acknowledgement of the infrastructural role of information. Being fully aware that this step is the most important and least trivial, we subsequently assume that this paradigm switch has taken place in order to discuss all of the rest of our strategy and tactics. The hope is that, in accordance with the theory of communicative action[588] and the "tipping point" theory[589], this dissertation and possible future pilot implementations of City Knowledge systems will make it more and more spontaneous for municipal offices to make such a commitment, which for now we will only assume.

Of course, as is often the case, the devil is and will continue to be in the details. Just where one draws the line when it comes to key decisions regarding the application of these principles will depend on the circumstances of that particular department at that particular moment. There is no one-size-fits-all silver bullet when it comes to City Knowledge. As long as there is a commitment, though, I think progress can be made by any town anywhere, with only a modicum of computer savvy and human and financial resources.

The rest of this chapter fleshes out a realistic pathway to City Knowledge using the "lessons" from parts II and III as supporting evidence and as examples of not only what worked, but also of the thornier implementation issues encountered or still to face.

The next and final chapter provides insights in how to keep these systems alive and how to make them become more common across the municipal landscape. The next section begins to address the question of exactly how to construct such comprehensive municipal knowledge.

COMMIT TO CITY KNOWLEDGE

A desire to create a comprehensive municipal information system has been around for decades[590], so for planners and municipal practitioners the concept of City Knowledge is not new. My thesis takes for granted that towns will want to use a comprehensive municipal information system.

What I think is new in this dissertation is the approach I propose for how to get to the holy grail of full city knowledge gradually and methodically as long as the town has made the simple commitment to become "information-conscious".

top-down

bottom-up

Some cities[591] have attempted to push a wholesale approach from the top down, forcing departments to comply with directives from on high. Almost everywhere, in the meantime, we have witnessed the spontaneous appearance of homegrown GIS endeavors at every branch of the municipal hierarchy in medium to large cities, primarily but not only in the "developed" world. I would like to contend that so-called "top-down" approaches have failed in their ability to percolate downwards just as bottom-up initiatives have failed to bubble upwards. In essence, it is rare (if not impossible) to find comprehensive GIS infrastructures within municipalities that are coherent and interconnected at all levels. If these two

588 Habermas, 1984, 1987; Innes, 1995.

589 Gladwell, 2000.

590 Geertman and Stillwell, 2004; Brail and Klosterman, 2001.

591 Like Singapore (Arun and Yap, 2000) and Vienna (Wilmersdorf, 2003), but also my own hometown of Venice (with the SITEA system).

popular methods for the creation of a knowledge infrastructure have largely failed, what alternative is there?

middle-out

As amply discussed in a previous chapter[592], I propose a hybrid approach to knowledge acquisition and maintenance that combines the freedom, flexibility and creativity of the bottom-up scheme with the structure, rigor and standardization of a top-down system. I call this a "middle-out" approach. This method engages offices or departments at the "front-line" of city operations, but avoids the pitfalls of the extreme, independent and individualistic idiosyncrasies associated with random homespun initiatives. In a middle-out approach, there is coherence and consistency within the front-line office; there is coordination between departments; and there is a conscious appreciation of the broader context in which information is being managed above and beyond the immediately obvious needs of a single office.

A committed city will treat information as an infrastructure and that ought to make all the difference in the world in how information is regarded in that town from that point on. Such a commitment would provide a strategic thrust for the gradual accrual of City Knowledge, which will be aided and abetted by the tactical lessons included in this chapter and throughout parts II and III.

Just what would make a town commit to this paradigm shift remains a question for further study, though the benefits derived by small-scale applications may incrementally build up the case for City Knowledge and could eventually bring about this sort of commitment from small towns as well as large cities, from Venice to Vancouver, from Peoria to Peking. Every town will have its own threshold beyond which the presumed costs are offset by the perceived benefits of such a commitment. Each department in each town, in fact – according to my middle-out strategy – will be facing this quandary and will have to "draw the line" between committing to this approach or not.

IDENTIFY JURISDICTIONS

According to our scheme, municipal information ought to be divided up among all functional divisions or departments in the city, along informational jurisdictions. We propose to "deconstruct" the urban landscape into homogeneous slices and "atomize" the data collection within each jurisdiction to achieve the smallest grain compatible with realistic needs and available financial and human resources.

As explained in detail in a previous chapter, "birthrights" and "deathrights" could be identified by analyzing information flows in typical departmental procedures[593]. It is preferable to pick an informational jurisdiction that is as unequivocal and uncontroversial as possible for the first City Knowledge application. You don't want to start by stepping on someone's toes (or turf) right off the bat.

---

592 Page 196.

593 As was done by Hart *et al.*, 2004 and Novello and Sartori, 2004.

This is the most normative of my tenets since I am convinced that this step is the most important for the long-term prospects of any municipal information system to flourish and persevere, as our *lessons from underground data management* [594] and many other positive[595] and negative[596] experiences have hinted at. This aspect will require some coordination among agencies to define the domains and "draw the lines" of demarcation between elements of the urban environment that are under each department's jurisdiction.

There will undoubtedly be turf battles over these choices and some urban elements or activities may even remain unassigned to any department, either because nobody wants them or because two or more offices lay claim to them. My view is that the latter controversies can be easily solved by letting either department – or even both – take control under strict performance expectations, with clear, yet reasonable deadlines for the delivery of shared datasets and layers that the "other" overlapping department can immediately share and use. I would let the department that lost the jurisdiction battle be the judge of the adequacy of the portion of City Knowledge that the "winner" was able to put together by the agreed deadline. If the product was deemed inadequate, then I would give the dissatisfied department a chance of doing a better job, under similar deadlines. This tension may actually produce much better results than unopposed assignments of domains.

If nobody wants to take charge of a particular aspect of the urban landscape[597], then some office may be put in charge of it *pro tempore* until the matter can be resolved by either determining that such an aspect is not worthy of special attention, or by identifying institutional reasons to attribute the onus to a specific office. If anybody ever comes looking for data about this undesirable domain, then they should be either be put in charge of it – if appropriate – or they should be probed in order to arrive at a resolution of the issue. If nobody seems to use information about the elements in this category, then there is no need to worry about this domain until somebody does.

Future research on this aspect may provide more details about the most appropriate consensus-building approach to the formation of these important jurisdictions.

594 See page 103.

595 All of our canal work fell clearly under Insula's jurisdiction, which greatly facilitate our efforts at making the data usable and re-usable. See part II.

596 The unclear boundary between the *Soprintendenza ai Beni Artistici e Storici* and the *Soprintendenza ai Beni Ambientali e Architettonici*, as well as the even more controversial private vs. public jurisdiction over public art has been the primary cause of the neglect suffered by the over 4,000 pieces of art visible from the streets of Venice. See page 112 for more on Public Art.

597 As was the case in Venice with Public Art.

PICK LOW-HANGING FRUITS

Having identified the informational jurisdictions, we recommend that each department pick a specific process within it, and target the urban structures or activities that participate in that process. The selection of the "low-hanging fruit" should be based on benefit vs. cost[598], as well as on the administrative priorities of the particular political moment. Ideally, the first project should be a "no-brainer"[599], i.e. one that actually pays for itself[600] or even makes the town money in the short term[601]. The choice of "grain" ought also to be tailored to the needs of the moment so as to make the project a 'showcase application".

Based on the lessons we learnt in Venice and in Massachusetts, each department should have at least one departmental body of information that it is already compiling, and has been for years[602]. Perhaps that ought to be the first candidate for a full-cycle City Knowledge review to see how to set up a system that would keep track of those data in perpetuity, including all of the periodic changes that might occur in the future. When dealing with cultural, historic assets[603], we have also demonstrated the benefits of acquiring plan-ready data for finite datasets of elements that are permanent, or very slowly changing. These datasets may be the lowest-hanging fruit and provide good returns on investment since they need to only be collected once and the cost of updates in minimal.

Once again, it is non-trivial to decide exactly where one draws the line between what's low-hanging and what isn't. Such decision would most likely be based on some cost-benefit calculus, which itself requires a baseline of information that may or may not be available. When it comes to inventories of municipal assets, it is always difficult to determine the cost of the complete catalog, since cities usually do not know how many items that inventory actually contains. Therefore, contracts such as those for our work on bridges and docks in Venice[604] are based on a per-item charge[605] and the city every year could just budget an amount that corresponds to *n* objects in the inventory, and so on until the job is complete.

More research could be done on the meaning of "return on investment" (ROI) in this context. The research would have to match perceptions about costs and benefits to choices with an eye to the process of selection and the definition of "low-hanging".

598 See the ample discussion of this important step in Reeve and Petch, 1999, pp. 51-81.

599 The "showcase application" mentioned in Reeve and Petch, 1999, p. 57.

600 As our Cambridge parking project (Flynn *et al.*, 2003).

601 As did our Parking Freeze fee proposal (Allard *et al.*, 2001).

602 See parts II and III.

603 See [lessons from palaces, churches and convents] on page 109 and ff.

604 See part II.

605 As you may recall (page 169) the rule-of-thumb unit cost for complex urban assets such as dock and bridges came out to about $150 per element to collect and archive attribute data as well as GIS attributes.

ATOMIZE AND CONQUER

Once the department has selected the specific municipal process to "informate", the next task is to dissect the process to arrive at the "atomic" elements about which to collect information in a systematic and perpetual way. Here is yet another situation that forces a decision on where to draw the line. Whatever our "atom" is today, it may be too big to represent tomorrow's sub-atomic particle. How fine of a grain should we use? What should be the unit of measurement? The answer, once again, is not easy[606].

atomic grain size

First of all, there are two grains to consider here: the size of the spatial units that will be managed and the detail collected for each of the spatial units. My quick choice would be: smallest area and largest attribute set. But, once again, you need to draw the line somewhere.

[spatial grain]

For the spatial dimensions, the expected unit of analysis should be one factor guiding the choice of unit of measurement – which should always be smaller than, or equal to the former. Cost considerations will dictate the grain to some degree, but the choice will be mostly practical and commonsensical, based on the physical object itself, as we learnt in Venice[607].

[attribute grain]

The choice of detail in the attributes to collect for each unit is less straightforward. As our cases showed repeatedly, there are advantages to be reaped by collecting as fine a grain as time allows[608]. This enables re-usability since finer grain can recombine more flexibly. Moreover, the attributes ought to be true to reality and not synthetic expressions of expertise, as we encountered with our tree management and public art projects. Instead of mentally calculating an overall condition rating for a tree (good, bad, OK), we collected visual clues of each factor that might play a role in the overall tree condition, such as insects, fungi, root problems, structural problems, etc. This adds many attributes, but makes the work more manageable and probably cheaper in the long run[609]. Beyond the separation of facts and values[610] in the choice of attributes to characterize the elements being inventoried, there is one last consideration one ought to make before the data structures are ready to be populated with field work or otherwise.

[attribute enrichment]

I am referring here to the teleological enrichment of the dataset through the choice of additional parameters that may not be needed right away for the task at hand, but are not too hard to collect once the crew is standing in front of the object in the field, and they may have some future use, either within the office itself or for some other maintenance, management or planning activity. This is where the science turns to an art. The line one draws here is very blurry and its shape can change dramatically depending on who is drawing it. A choice of additional fields made by someone with a planning mindset may permit a more artistic palette of re-

---

606 See the lessons we learnt when choosing between using buildings ort floors as atoms of City Knowledge in Worcester (on page 156).

607 See [Lesson 1: atomize] on page 50, and other [lessons from our physical studies] on page 62.

608 See [re-usability of fine-grained data], [use finest grain], and [keep data in raw form] on page 68 ff.

609 See page 126 and ff.

610 See the lesson entitled [separating facts and values] on page 131.

uses and more sophisticated second-order capabilities, whereas the choice made by someone with a management mindset might be more scientific and perhaps less creative and less combinatorial, but more concrete and focused on specific foreseeable future needs and wants.

locus of control

[collect fine-grained data within jurisdictions]

[sustainably maintain data at frontline]

Plan-ready information will be generated more efficiently if the work of creating and maintaining the databases, GIS layers and information systems is left to the "frontline" offices of each department, where the real action is and where municipal data are acquired and generated. Frontline offices in turn should generally attempt to delegate most data collection and updates to outside contractors or to the public and, only as a last resort, they should devote municipal staff to the task if necessary. Since the data would be collected first and foremost to suit the needs of the front-line department, their collection should be verified by field inspections by town officials to guarantee the integrity and quality of the data.

[adopt standards and methods]

[share to leverage value-added benefits]

Once enough data have been collected (using agreed standards and methods), and organized into plan-ready information, sharing with other agencies may be possible either by operating along existing interagency lines of cooperation or by crafting new virtuous networks of value-added knowledge.

data collection

Having settled on the spatial and attribute grain, and having identified the frontline offices in charge of the data collection, as well as the surrogates who will inexpensively collect the data for those offices, a town that has embraced the cause of City Knowledge faces two fundamental tasks:

- Catching up with the backlog of urban data
- Dealing with future changes in the city

catching-up with the backlog

[commitment to city knowledge]

In my middle-out approach, departments (or even divisions or offices within a department) are responsible for their own information. The fact that the department or unit takes ownership over its own knowledge guarantees that the data are *reliable, complete, pertinent* and *up-to-date*, since the job of department members rides on the quality of these data. The department needs to make a commitment to the acquisition and upkeep of its slice of City Knowledge before any of the following steps can even be envisioned.

paying for the backlog

It is undeniable that an initial collection of data will be necessary to catch up with the backlog of information that is already "out there" in the city today. Naturally, towns will explore the no-cost or low-cost solutions first.

[no-cost solutions]

The no-cost solutions entail getting data from third-parties for free. No-cost is probably an illusory concept since the mere acquisition of pre-existing data would itself entail some cost, even if the town were to receive the "complete municipal information system" as a "free-gift"[611].

611 In this section, I will call no-cost or free what is actually a fairly low cost activity, but certainly not entirely free-of-charge.

[free = self-reporting]

Data could be collected more or less for free by mandating a lot of self-reporting from the public and from businesses. This may include making information-returns part of the current requirements for existing processes (like permit applications, for instance). Similarly, contractors or professionals hired by developers submitting a plan, or hired by the town[612] to rebut the same plan, could be forced to submit information in a certain format in order to feed the growing information infrastructure[613].

[free = required information returns]

[free = secondary sources]

Free data could also come from secondary sources, be they government agencies (like the regional planning agencies, or the Census Bureau) or private or semi-private entities (like Chambers of Commerce, Real Estate professionals organizations, etc.). Acquiring these data will necessitate some formal agreement but government-to-government data exchanges ought to be free and fairly feasible with a modicum of bureaucracy and legal paperwork. Post-processing of these datasets may involve getting rid of "stubborn errors", but the guaranteed reliability of the data may justify investing in middleware such as Ferreira (1998) proposes[614].

[free = volunteers or students]

A final source of "free" primary (or secondary) data can be students, researchers or volunteers. All of my examples come from the "free" work of WPI students and even from volunteers (*Earthwatch*, grade schoolers, interns). Data validation and quality control are factors to be considered carefully here[615]. Involving local citizens in these efforts produces a multiplier effect.

[low cost= maintenance-based data collection]

Low-cost, very gradual ways of slowly chipping away at the "backlog" may include collecting data as repairs or maintenance are done. If the repair work is farmed out to contractors, a specific informational return would be included in the bid language and in the subsequent contract.

[low cost = contractual obligations]

Another way for a town to gather information gradually – yet unsystematically – is to force its contractors to produce GIS-compatible data whenever some spatially-referenceable service is rendered[616]. Towns could simply change the wording of contracts and bids to impose informational deliverables upon external contractors[617]. However, before doing so, that town needs to make sure it is ready and able to deal with such an information onslaught once it has mandated it.

[low cost = in-house]

A final recourse would be for the department to do some of the data collection with its own staff. This is not a preferred path, but could be considered, depending on the staffing levels and on how busy municipal staffers are. Some institutionalization of the process of catching up with the

612 Often at the developer's expense.

613 As we did in Gage *et al.*, 2003 and in Brown and Groeli, 2003. These information-conscious requirements may provide a net gain for the town as long as these costs are not somehow transferred to the community in some other way, like through higher real estate costs.

614 Ferreira, 1998.

615 For example, Forma Urbis ended up re-collecting all of the bridge data for Insula, despite the fact that half the bridges had already been measured by Bahn *et al.* in 1998. Generally, Forma Urbis will always collect the data again if possible. The "methodology" is therefore the primary contribution that student projects really make to the final operative information systems, in my experience.

616 For example, surveyors may be asked to do that, as well as architects, or professional planners.

617 This would not be free because it is presumed that the cost estimates would rather quickly incorporate the informating costs and pass them onto the town.

backlog may be instituted by including data-gathering as one of the requirements in the job descriptions of municipal workers.

[market cost = consultants]

Catching up with the backlog may, in some municipal domains, may be urgent enough or sensitive enough to warrant farming out to outside consultants to expedite matters along and get the program off to a good start so it can act as a model and a catalyst. Capturing "snapshots" of *activities*, such as traffic or housing patterns, may also entail the participation of professional consultants.

[information budget]

Regardless of the route a city decides to take, the town will incur some costs. These could be offset, at least in part, by devoting a fixed percentage of a department's regular budget (say 5%) to these knowledge-acquisition activities, thus making the process rather *affordable* during any FY.

Future research could focus on some of the more creative no- or low-cost data-gathering schemes mentioned herein, in which we have little or no direct experience and are not really discussed in the literature either. In particular, I plan to experiment more cogently with the *quid pro quo* bartering of information[618] and the tweaking of contracts and job descriptions in Venice and in Massachusetts.

618 Between Spencer and the Central Massachusetts Regional Planning Commission. See footnote number 437.

Update Sustainably

All municipal activities leave a paper trail that would easily lend itself to automation. Roads are regularly re-paved, cleaned and cleared of snow, so someone is issuing work-orders or stipulating contracts for these services. Similarly, trees are bought, planted, removed and trimmed and paperwork is produced to make each of these actions happen and to keep track of the corresponding expenditures. Sewers, like many other components of the urban infrastructure, are subject to similar record-keeping procedures, plus they are also regulated and licensed. The list goes on.

automating and informating

Not only should these informational opportunities be tapped into, as much as possible, to populate the city's knowledgebase as we catch up with the backlog, but they should play an even more important, primary role in keeping the information up-to-date to promote the sort of "automation" that would lead to a more rigorous approach to informating urban maintenance, management and planning[619].

[capturing transactions]

Just about all of the techniques listed in the previous section that could save a town some money in the collection of data, could also apply when it comes to data updates. The most important and most effective approach to sustainable information upkeep would be the institution of computerized mechanisms to intercept changes in the urban realm as they happen, never creating a backlog ever again. This step would entail considerable thought and planning, together with technical, legal and administrative agreements and solutions.

[update or die]

The advice to a community starting on the path to City Knowledge is to never embark in a system to keep track of a class of assets unless there is a plan to capture future updates too. Updates should be sustainable because they can be more expensive than the backlog in the long run, and they can make or break a project down the line. Some projects are born obsolete because of their disregard for the upkeep of the datasets.

[experiment with innovative updates]

The sticky points as well as the interesting research angles are about the same here as they were for the collection of the backlogged data discussed in the previous section. The next step for me in this context will be to go back to the more mature applications that I had a hand in starting in Venice and in Massachusetts and take a closer look at gaps in the updating mechanisms, so that we can experiment with some of the novel techniques that are introduced in this dissertation but are still mere conjectures and hypotheses until we put them to the test in the real world.

619 See the lesson about [maintenance-based updates and cataloging] on page 131.

SHARE APPROPRIATELY

The full power of City Knowledge and its multiplicative, emergent effects will not become tangible unless some amount of sharing takes place, to bring more than one department, agency or institution into the system.

[share to foster emergence]

As soon as a department has mastered a specific application for the maintenance, management or planning of some urban aspect over which it has jurisdiction, the information is available for sharing if so desired. This step is not necessary, nor mandatory, but it would be certainly salutary, so I would personally encourage the practice of sharing non-controversial, privacy-respecting data, in order to generate the amalgam within which a fully emergent City Knowledge system can flourish and produce unexpected, value-added benefits[620].

The technology itself, once again, is not the solution, nor the problem. But it may contribute to a more rapid transformation of the internal organizational dynamics of government agencies toward a "connected distributed" [621] *modus operandi*, that will enfranchise the citizenship as well as the front-line civil servants. This, in turn, may set the stage for a truly devolved informating "wholeness"[622], where managers and managed contribute together, "holistically"[623], to a middle-out approach for the management of urban affairs. The Internet and the WWW would certainly facilitate such a development.

All of these innovative approaches may incrementally lead to "*a truly interactive, timely planning dialogue between neighborhood planners and city agencies – as well as* [to] *a mode of interagency* [and – I would add – 'intra-agency'] *coordination that might allow agencies to keep pace with one another*"[624] and with their public constituency. Once again, the interconnectivity provided by the WWW today makes this interactive approach all the more feasible and affordable for our cash-strapped public agencies.

[share to avoid duplication]

Another – more practical – reason for sharing is to avoid duplication of efforts within the department or the municipality[625]. Of course, the avoidance of duplication implies that some method exists to inform all potential users of the existence of urban data and information for such and such an asset or activity over such a span of time and over such a spatial extent. This is what metadata is really good for. Unfortunately, the use of metadata – even in progressive GIS-intensive municipalities – remains lackadaisical so far, especially at the real front lines. Even my groups have not yet used any part of the FGDC metadata standard rigorously or methodically.

Metadata use and usefulness is an area of research that certainly deserves more attention and I plan to look at my examples through this lens as well, when I go back for the aforementioned evaluation of mature systems.

620 As was the case in the [lessons from palaces, churches and convents] on page 109.

621 Thomas W. Malone, *Is Empowerment Just a Fad?* (1997)

622 Zuboff's term (1991).

623 This term borrowed from Evans and Ferreira, 1995, p. 458.

624 Ferreira, 1998.

625 See the lesson entitled [share results to avoid duplication] on page 109.

Coordinate as Needed

Once enough sharing takes place – most likely along paths of least resistance connecting two agencies with pre-existing administrative ties – the issue of coordination will immediately come up. While this is not a mandatory step, it is very likely to be required once the number of sharing agencies becomes greater than two. In fact, many issues will become thornier once interoperability is more in demand across departments.

Coordination is costly and needs to address dependencies and problems with synchronization. Sharing also entails dealing with many other subtler intricacies such as the issues of data quality, reliability, liability, copyrights, accessibility, security and many others. All of these issues require coordination.

[grow into coordination]

I am purposely downplaying the importance of coordination because I don't want to scare people from engaging in all of the other steps listed before. In fact, I was not going to even include this step at all for that reason. Many people react negatively to too much oversight and too many rules – including me. It would be great to establish coordination committees or roundtables within a municipality, but I think that the membership into these coordinating entities should emerge slowly (like everything else in my City Knowledge approach).

[coordinate spontaneously]

I believe that individual departments need to grow their internal information capacity and begin two-way sharing until they themselves experience the self-generated impulse to coordinate, at which point they can be ready to join the interdepartmental coordinating club. Before that time, they may participate in meetings as observers to become acquainted with the issues, but the methods[626] will have no "stickiness"[627] until the urge to coordinate emerges spontaneously from within the department.

There are many ripe areas of research in this arena, but I am not interested in the technical standards per se, but in how to make them useful and used. I would like to apply some lessons from the field of emergence[628] to determine what minimal sets of standards and protocols are needed to bring out the second-order, high-level emergent qualities of City Knowledge, without overburdening the users.

626 Described starting on page 219.

627 Gladwell, 2000.

628 Johnson, 2000. I am especially intrigued by *Tipping Point* (Gladwell, 2000) and *Ideavirus* (Godin, 2001).

A SUSTAINABLE FUTURE FOR CITY KNOWLEDGE

Despite all of the great cases presented herein, and despite the continued success of our initiatives in Venice and in Massachusetts, to date none of these municipal administrations has wholeheartedly embraced City Knowledge as its official *information-aware modus operandi.* There have been comforting signals from both sides of the Atlantic[629], yet there is no proof out there that these concepts would fully work outside of the privileged position that I have put myself into – namely of having the luxury of coming at problems laterally, with considerable human and technological resources and without needing to be accountable to bosses, citizens or taxpayers.

One of the most important future developments – hopefully in the short run – will be to have some real examples of cities and towns (or at least departments) that formally espouse these principles and begin to deal with the devilish details of implementing and keeping alive such an operation. There will be innumerable lessons to learn once these experiences start in full earnest.

the tipping point of City Knowledge

Since it is an emergent system, I think that the small seeds of City Knowledge that are already germinating in Venice, Cambridge, Boston and Worcester will gradually grow to encompass entire departments and then infect the whole city. As Malcom Gladwell explains in *The Tipping Point*[630], after these innovators and "early adopters" have established the concepts in their respective organizations and communities, we'll need to wait for the "connectors", "mavens" and "salesmen" to do their subtle work before we can witness the emergence of a more sizeable "early majority". This "tipping point" will occur when the adoption of City Knowledge reaches the "moment of critical mass, the threshold [...] where the unexpected becomes expected, where radical change is more than a possibility". When or whether this will happen depends on the communicative skills of practitioners like me who can see the possibilities that plan-ready information holds for planners and for decision-makers and can articulate a strategy to achieve it[631].

value-added benefits for high-order users

value for front-line users

Planners can reap the value-added benefits of being able to jump right into second-order analytical tasks – which is what they are probably best at, in addition to being what they are paid for – without having to track down all of the necessary data first. Yet these higher order capabilities will only become available once the front line municipal workers begin to see direct first-order benefits for themselves and their job. It is the day-to-day value derived from the automation of the front-office activities that enables the back-office to harvest the deeper informating benefits. I posit that this

629 The latest was the honor of having been invited to talk about City Knowledge principles to the top echelon of the M.I.S. department of the City of Boston (on 8/7/04). Meanwhile, in Venice, Insula continues to carry the torch as the active operation that most closely resembles City Knowledge at work.

630 Gladwell, 2000.

631 For example, treating City Knowledge as an *ideavirus* (Godin, 2001) may yield some useful insights.

middle out approach

useful commingling of instant gratification with long-term capacity building can be achieved with our middle-out method that mixes the best features of bottom-up and top-down to create a sensible, low-impact approach to the accrual and sustainable maintenance of city knowledge.

cost vs. benefit

In my view, a City Knowledge system is both desirable[632] and feasible[633]. My personal cases document how the emergent power of plan-demanding knowledge provides informating benefits that are higher than the typical improvements due to mere automation. The technological costs of such an endeavor are declining every day, whereas the organizational, administrative and logistical costs can be kept under control, using the techniques that many of my cases have demonstrated in the real world. In short the balance between benefit and cost is tipping toward the former.

toward city knowledge

These favorable circumstances make finer-grain data collection conceivable and affordable, which in turn should yield more flexibility in the re-utilization, aggregation, manipulation and analysis of our urban datasets. With only a modicum of overarching coordination, different departments can set out to comprehensively capture the data that they need in order to maintain assets, manage activities and plan developments within their informational jurisdictions. As they catch up with the backlog of information already out there – starting from low-hanging fruits[634] – they can also begin to put in place self-perpetuating mechanisms for the upkeep of the fundamental framework datasets that are necessary to fulfill their own departmental needs. The resulting plan-ready systems promise to make each office perform its duties more effectively, efficiently and efficaciously.

Following the path of least resistance, these plan-ready systems can then be shared where institutionally or opportunistically appropriate. I hypothesize that sharing can give rise to the self-organizing behavior that City Knowledge has already demonstrated, allowing plan-demanding situations to emerge as we gain a deeper and deeper understanding of the urban fabric in which we live and operate.

the future as a possibility

With this dissertation, as a reflective practitioner, I have put forth my new metaphors for the redefinition of City Knowledge as a paradigmatic state-of-mind. I hope this concept will penetrate into the collective subconscious of municipal governments, so that they will begin to treat urban data, information and knowledge as essential infrastructural elements of our cities and towns.

As a communicative planner, I hope that having defined, proposed and partially demonstrated a framework for the gradual growth of City Knowledge, it will no longer be a vague holy grail that we could perhaps achieve "some day", but will instead become a clearer cause to champion for those of us who believe in its power for the transformation of communities.

632 As indicated in Marchewka, 2003, pp. 8-9.

633 As my cases in Part II and III showed.

634 i.e. the element of the urban realm that represents the best return on investment weighing the transaction costs versus the timeliness and value of the benefits.

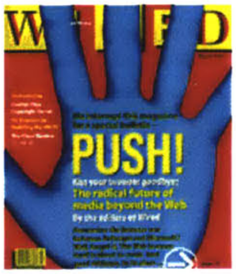

In essence I am proposing to replace the tactical process of "muddling through" with a strategic equivalent: the conscious and deliberate "middling through". I think technology will fuel this process. For example, I foresee a future in which mobile technology will not be focused exclusively on "Push"[635]. I am sure that we will soon witness the emergence of "Pull", with which we will tap into the data-collection potential of these mobile devices[636]. I think towns ought to be the first to "pull" geo-spatial data from the field using these cutting edge technologies at the very front of the front lines.

In addition to technological developments, there are many other areas of research that promise to be useful to give more body to the fledgling concepts of City Knowledge and middle-out here introduced. I look forward to deepening my understanding of these mechanisms, so that towns can really find smooth ways to acquire the data, information and knowledge that they require for urban maintenance, management and planning.

As I said in the introduction, I hope that some day City Knowledge will itself become transformed from a plan-demanded cause pushing us from the past toward a future full of emergencies, into a different type of cause – a plan-demanding cause celeb – pulling us into a future full of emergent possibilities for our cities and towns.

Ten or twenty years from now City Knowledge and middle-out may or may not be household words in municipal administrations, however, if cities and towns will be making progress towards comprehensive municipal information systems that have the qualities I listed in this dissertation, then perhaps the paradigm shift will have taken place – perhaps under another name or with a different approach – and, if things go the way I predict, these concepts may be so ingrained in the municipal psyche to be virtually unnoticeable. I hope to be around when that begins to happen.

---

635 Whereby connected users are able to receive information about everything everywhere (from stocks to location-aware advertising, news, media, etc.) through wireless PDAs and tablet PCs. Push was on the cover of Wired magazine way back in March of 1997, but there has recently been a resurgence of interest in it.

636 I am advising a project on "Pull" in the 2004-2005 academic year. As far as I know this may be my own neologism. I haven't heard it used in juxtaposition with Push.

WORKS CITED

Aimo, Emilia, Della Sala, Stefano, Stradella, Silvia, Tagliapietra, Davide, and Vazzoler, Marina. 1999. "L'inquinamento dei rii. La diffusione degli inquinanti nei canali interni di Venezia" in Caniato *et al.*, eds., *Venezia la Città dei Rii.*

Allard, Mark, Amato, Ben, Barber, Keith, Couture, Steve. 2001. *Evaluation of the Downtown Boston Parking Freeze.* Boston Project Center: WPI Interactive Qualifying Project.

Amlaw, K., Kervin, C. L., Mondine, I., Vepari, C.. 1997. *Optimization of Cargo Boat Deliveries Through the Inner Canals of Venice.* Venice Project Center: WPI Interactive Qualifying Project.

Arbeit, D. 1993. "Resolving the Data Problem: A Spatial Information Infrastructure for Planning Support". *Proceedings of Third International Conference on Computers in Urban Planning and Urban Management,* Atlanta, Georgia, July 23-25, Vol. 1, pp. 2-26.

Arun, Mahizhnan and Yap, Mui Teng. 2000. "Singapore: The Development of an Intelligent Island and Social Dividends of Information Technology", *Urban Studies*, Vol. 37, No. 10, pp. 1749-1756.

Avison, D.E. and Fitzgerald, G. 1988. *Information Systems Developments: Methodologies, Techniques, and Tools.* Oxford: Blackwell Scientific Publications.

Azad, Bijan. 1998. *Management of Enterprise-wide GIS Implementation: Lessons from Exploration of Five Case Studies*, PhD Thesis, Department of Urban Studies and Planning, MIT: Cambridge, MA.

Babic, K., Leeds G., Sidiroglou, S., Borek, M.. 1998. *Analysis of Sewer Holes and Canal Wall Damage in Venice, Italy.* WPI Interactive Qualifying Project.

Bacharach, M. and Hurley, S., eds. 1991. *Foundations of Decision Theory*, Oxford: Basil Blackwell.

Bahn, R., Deliso, A., Hubbard, S.. 1998. *The Inventory and Analysis of the Bridges and Pedestrian Traffic in Dorsoduro, San Polo, and Santa Croce Sestieri of Venice.* WPI Interactive Qualifying Project.

Banovetz, James M., Dolan, Drew A. and Swain, John W., eds. 1994. *Managing Small Cities and Counties. A Practical Guide.* Washington, DC: International City/County Management Association (ICMA).

Barr, R. 1991. "A federal approach to GIS", *Mapping Awareness*, Vol. 5, No. 6, pp. 15-19.

Bassa, Olga, Chaves Ferreira, Elisa, LaVenture-George, David. 2003. *Examining a role for the community reinvestment act (CRA) for private brownfield reinvestment in Worcester.* Worcester Community Project Center: WPI Interactive Qualifying Project.

Battocchi, Scott, Bell, Caitlin, Blair, Jeffrey, Cole, Natalie. 2003. *The Island of Pellestrina: Case Study for the Environmental Atlas of the Lagoon.* Venice Project Center: WPI Interactive Qualifying Project.

Beltran, J., Brophy, E.E., Cardenas, A. 1995. *A Computerized Catalog of Outdoor Art in Dorsoduro, Venice.* Venice Project Center: WPI Interactive Qualifying Project.

Bennett, Matt, Premo, Elizabeth, and Tavares, Jarod. 2001. *The Management and Maintenance of the Venetian 'Verde Pubblico'.* Venice: WPI Interactive Qualifying Project.

Bernard, Russell H.. 2002. *Research Methods in Anthropology. Qualitative and Quantitative Approaches.* Third Ed.. Walnut Creek, CA: Altamira Press.

Blizard, Ian, O'Connell, Edward and Schmidtberg, Erik. 2004. *Building Maintenance in Quincy.* Boston Project Center: WPI Interactive Qualifying Project.

Blomberg, Adam, Newton, Ben, Mandrila, Ana-Maria and Vacca, Tom. 1999. *Radio Frequency ID Boat Monitoring System.* Venice Project Center: WPI Major Qualifying Project.

Borrelli, A., Crawford, M., Horstick, J., Ozbas, H.. 1999. *Quantification of Sediment Sources in the Canals of the City of Venice, Italy.* WPI Interactive Qualifying Project.

Brail, Richard K. and Klosterman, Richard E., eds. 2001. *Planning Support Systems.* Redlands, CA: ESRI press.

Brown, Zachary and Groeli, Florentin. 2003. *Urban Planning in Worcester: Past, Present and Future.* Worcester Community Project Center: WPI Interactive Qualifying Project.

Budić, Zorica D.. 1994. "Effectiveness of Geographic Information Systems in Local Planning", *Journal of the American Planning Association*, Vol. 60, No. 2, Spring, pp. 244-263.

Butler, Allison J. and Dueker, Kenneth J. 2001. "Implementing the Enterprise GIS in Transportation Database Design", *URISA Journal*, Vol. 13, No. 1, Winter, pp. 17-28.

Cahan, Bruce B. 2002. "Strategic Investing in Community GIS", presentation delivered at the conference *Planning for Community Statistical Systems*, Tampa, FL, March 13, 2002.

Calhoun, Shaun, Kumar, Raja, Tiscia, Anthony, and Turner, Terrence. 2004. *Manuscript Transcription Assistant Initiative.* WPI Major Qualifying Project.

Calvino, Italo. 1972 (1993 ed.). *Le Città Invisibili*,. Milan: Mondadori.

Cambridge University. 2003. *Flooding and Environmental Challenges for Venice and its Lagoon: State of Knowledge 2003,* An International Scientific Meeting at Churchill College, Cambridge, England, 14th–17th September 2003.

Campagna, Michele and Deplano, Giancarlo. 2004. "Evaluating geographic information provision within public administration websites", *Environment and Planning B: Planning and Design*, Vol. 31, pp. 21-37.

Campbell, A. J.. 1999. "Institutional Consequences of the use of GIS" in *Geographical Information Systems.* Longley, Paul A., Goodchild, Michael F., Maguire, David J. and Rhind, David W., eds. Second Edition, 2 Vols.. New York: John Wiley and Sons.

Campbell, Heather and Masser, Ian. 1995. *GIS and Organizations*. London: Taylor and Francis.

Caniato, Giovanni, Carrera, Fabio, Giannotti, Vincenzo and Pypaert, Philippe, eds.. 1999. *Venezia la città dei rii.* Verona: Cierre edizioni.

Caniato, Giovanni. 1999. "La manutenzione dei rii in epoca moderna. Politiche e modalità d'intervento" in *Venezia la Città dei Rii, op.cit.*

Caporale, Stefan J., Cergneux, Max, Freed, William, and Matzal, Anna L. 1997. *The Development of a Prototype for an Automated Ambulance Dispatch System in Venice, Italy.* Venice Project Center: WPI Interactive Qualifying Project.

Caron, Claude and Bédard, Yvan. 2002. "Lessons Learned from Case Studies on the Implementation of Geospatial Information Technologies", *URISA Journal*, Vol. 14, No. 1, pp. 17-36.

Carrera, Fabio. 1994. "I Rii e la Qualità della Vita a Venezia" in *Coses Informazioni*, Venice.

Carrera, Fabio. 1996. *Il Traffico Acqueo nei Canali Interni di Venezia,* report to UNESCO in the framework of the *Venice Inner Canals* project. July.

Carrera, Fabio. 1997. *Il Traffico Acqueo nella Laguna di Venezia,* Report to the Provincial government on the boat traffic campaign of July 29, 1997. Venice: Provincia di Venezia, Ufficio Trasporti e Mobilità, and COSES. December 15, 1997.

Carrera, Fabio. 1999a. *Il Traffico Acqueo nei Principali Rii di Venezia e Canali Laguari,* Report to the Municipal and Provincial governments on the boat traffic campaign of March 17, 1998. Venice: COSES. January 31, 1999.

Carrera, Fabio. 1998. *Il comportamento idrodinamico dei canali interni di Venezia.* Report to UNESCO in the framework of the *Venice Inner Canals* project.

Carrera, Fabio. 1999c. "La riapertura dei rii terà" in *Venezia la Città dei Rii, op.cit.*

Carrera, Fabio. 1999d. *Geographical Information System per la Gestione dei Canali Interni di Venezia.* Report on the UNESCO-Insula agreement (phase 1).

Carrera, Fabio. 2001. in *City Under Siege.* National Geographic Video. Aired around the globe on the National Geographic Channel.

Carrera, Fabio. 2001b. *UNESCO-MURST "Venice Inner Canals project",* final report to UNESCO in the framework of the *Venice Inner Canals* project.

Carrera, Fabio. 2002. "Trasporto e Distribuzione di merci nel centro storico di Venezia: spontaneismo e riorganizzazione" in *Quaderni di Insula.* Venice: September 2002, pp. 29-33.

Carrera, Fabio. 2004. "City Knowledge as key to understanding the relation between waters and stones in Venice". Refereed paper in *Flooding and Environmental Challenges for Venice and its Lagoon.* Cambridge, UK. (accepted for publication).

Carvajal, Hernando, Federico, Manrico, Gonzalez, Carlos, Johnson, Tim, and Levesque, Jeff. 1991. *A Study of Tide Flows, Mud Buildup, Boat*

Traffic, and Structural Damage on the Cannaregio Canal Subsystem. Venice Project Center: WPI Interactive Qualifying Project.

Cellucci, Paul, Governor. 2000. *Executive Order 418. Assisting communities in addressing the housing shortage.* The Commonwealth of Massachusetts, Executive Department, issued on 1/21/2000.

Chicken, J.C. and Hayns, M.R.. *The Risk Ranking Technique in Decision Making.* Oxford: Pergamon, 1989.

Chiu, D., Jagganath, A., Nodine, E.. 2002. *The Moto Ondoso Index: Assessing the Effects of Boat Traffic in the Canals of Venice.* WPI Interactive Qualifying Project.

Chiu, D., Lacasse, C., Menard, K.. 2004. *Mapping Turbulent Motorboat Discharges in the Canals of Venice, Italy.* WPI Major Qualifying Project.

Ciacciarelli, Michael, O'Connor, Bryan, Sylvia, Barry and Testa, Steven. 1990. *A Hygenic, Dynamic, and Static Study of the Canals of Venice, Italy.* WPI Interactive Qualifying Project.

Cohen, M. D., March, J. G., and Olsen, J. P. 1972. "A Garbage Can Model of Organizational Choice", *Administrative Science Quarterly*, Vol. 17, No. 1, pp. 1-25.

Commonwealth of Massachusetts. 2000. *Building Vibrant Communities. Linking Housing, Economic Development, Transportation and the Environment.* Executive Department. Available on line at http://www.massdhcd.com/eo418/hyperlink%20index.htm. Accessed on 5/19/04.

Comune di Venezia. 1986. *Indagine sul Traffico Acqueo.*Venice.

Comune di Venezia. 1987. *Indagine sul Traffico Acqueo.*Venice.

Consorzio Venezia Nuova. 1988 *Livellazione di dettaglio del Centro Storico e delle isole.*

COSES. 1999. "Indagine sui danni provocati dagli eventi di marea sulle Unità Immobiliari pubbliche e Private ai piani terra del Centro Storico di Venezia". *Rapporto COSES n. 46/1999.*

Coulton, Claudia J. with Lisa Nelson and Peter Tatian. 1997. *Catalog of Administrative Data Sources: for neighborhood indicator systems.* National Neighborhood Indicators Partnership. The Urban Institute.

Cozzutto, Lapo. 1997. *Rapporto sull'Applicativo "SmartInsula".* UNESCO Report in the framework of the "Venice Inner Canals" project.

Craglia, Massimo and Signoretta, Paola. 2000. "From global to local: the development of local geographic information strategies in the United Kingdom", *Environment and Planning B: Planning and Design*, vol. 27, pp. 777-788.

Craglia, Massimo, Leontidou, Lila, Nuvolati, Giampaolo and Schweickart, Jürgen. 2004. "Towards the development of quality of life indicators in the 'digital' city", *Environment and Planning B: Planning and Design*, vol. 31, pp. 51-64.

Creps, Erik, Hiremath, Pallavi, Pantazelos, Robert, and Stefanik, Dan. 2001. *Re-engineering Tree Management in Cambridge: The Benefits Provided by Updated Tree Management Techniques.* Boston: WPI Interactive Qualifying Project.

Crockett, Thomas W. 1998. "Digital Earth: A New Framework for Geo-Referenced Data", *ICASE Research Quarterly*, Vol. 7, No. 4.

Cullen, Christopher, Moriarty, Michael, and Patel, Chirag. 2002. *Parking Monitoring and Management in Cambridge.* Boston Project Center: WPI Interactive Qualifying Project.

Dandekar, Hemalata C., ed. 2003. *The Planner's Use of Information.* Second ed. Chicago: Planners press.

Dehri, Genci, Laratro, Antonio, and Orlando, Michael. 2004. *Ultra Violet Scanner for the Enhancement of the Readability of Damaged Parchments.* WPI Major Qualifying Project.

Delaive, A., Kristant, E., Petrowski, C., Santos, L.. 2002. *The Church Floors in Venice, Italy: An Archeological Study and Analysis.* Venice: WPI Interactive Qualifying Project.

De Marchi, Carlo. 1993. *A hydrodynamic model of the network of canals of the city of Venice.* Master's Thesis, University of California, Davis, CA.

De Marchi Carlo. 1996. *Effetti idrodinamici di interventi sulle caratteristiche idraulichedella rete dei canali interni della città di Venezia. Rapporto sulle simulazioni matematiche delle caratteristiche idrauliche della rete dei canali Veneziani* for UNESCO-MURST project *Venice Inner Canals*; July.

De Marchi, Carlo. 1997. *Un Modello Matematico del regime idraulico della rete dei canali interni della città di Venezia, Rapporto sulla calibrazione di un modello matematico idraulico della rete di canali interni della città di Venezia* for UNESCO-MURST project *Venice Inner Canals*; November.

Desmond, Brenda, Fitzhugh, Christopher, Kheny, Vikram. 2002. *Improving Environmental Awareness for Chelsea Creek Communities.* Boston Project Center: WPI Interactive Qualifying Project.

Doherty, K., Maraia, J., Parodi, C., Souto, F.. 1995. *A Documentation and Analysis of the Traffic, Cargo Deliveries, and Docks within the Insulae of Santa Maria Formosa and Frari.* WPI Interactive Qualifying Project.

Donnelly, B., Hart, B., Pilotte, M., Scherpa, T.. 1999. *Safeguarding the Churches of Venice, Italy: a Computerized Catalogue and Restoration Analysis.* WPI Interactive Qualifying Project.

Dorigo, Livio, 1966. *Rilievi Contemporanei di velocità di marea nei canali della Giudecca, Fondamente Nuove, Canal Grande ed alcuni rii interni della città,* Ist. Ven. Sci. Let. ed Arti, Vol. III, Rapporti e Studi, Venezia.

Dorigo, Wladimiro. 1999. "I rii di Venezia nei secoli IX-XIV: un profilo storico" in *Venezia la Città dei Rii, op.cit.*

Dow, Walter, McDaniel, Matthew, Moyer, Dale, St.Pierre Jeremy. 2003. *Documenting the Blue and Green Amenities of Worcester, Massachusetts.* Worcester Community Project Center: WPI Interactive Qualifying Project.

Duffy, J., Gagliardi, J., Mirtle, K, Tucker, A.. 2001. *Re-Engineering the City of Venice's Cargo System for the Consorzio Trasportatori Veneziani Riuniti.* WPI Interactive Qualifying Project.

Dye, Richard F. and Merriman, David F., 2000. "The Effects of Tax Increment Financing on Economic Development", *Journal of Economic Development*, 47, 306-328.

Eichelberger, Pierce. 2004. "Modernizing Land Records Pays Off". *ArcUser*. July-September. Redlands, CA: ESRI corporation, pp. 16-17.

el-Khoury, Rodolphe, "Polish and Deodorize: Paving the City in Eighteen Century France," *Assemblage*, no. 31, Cambridge: MIT Press, Spring 1997.

Evans, John and Ferreira, Joseph Jr.. 1995. "Sharing Spatial Information in an Imperfect World: Interactions between Technical and Organizational Issues" in Onsrud and Rushton, eds., 1995. *Sharing Geographic Information*, op. cit., pp. 448-460.

Evans, John D. 1997. *Infrastructures for Sharing Geographic Information among Environmental Agencies*. Ph.D. Thesis. MIT-DUSP.

Faludi, Andreas, ed. 1973. *A Reader in Planning Theory*. Oxford: Pergamon Press.

Faludi, Andreas. 1986. *Critical Rationalism and Planning Methodology*. London: Pion.

Farmer, Jay, Settle, Jennifer, St. Pierre, Matt, Wall, Chris. 2003. *Transportation and Open Space Study for the City of Worcester*. Worcester Community Project Center: WPI Interactive Qualifying Project.

Felices, D., Moreno, C., Munoz, A., Smith, B.. 1994. *A Documentation and Analysis of the Docks, Cargo Deliveries, and Boat Traffic within the Santa Maria Zobenigo Insula*. WPI Interactive Qualifying Project.

Felices, Martin, Goodfellow, Lauren M., Johnson, Jay L., and Maheshwary, Sonali A. 1997. *A preliminary feasibility study of the implementation of HIFLO vacuum sewerage system within the city of Venice*. Venice Project Center: WPI Interactive Qualifying Project.

Ferreira, Joseph Jr.. 1998. "Information Technologies that Change Relationships between Low-Income Communities and the Public and Non-profit Agencies that Serve Them," Chapter 7 in *High Technology and Low-Income Communities*. Donald A. Schön, Bish Sanyal and William J. Mitchell, eds., Cambridge: MIT Press.

Ferreira, Joseph Jr. 2002. "Spatial Data Infrastructure for Economic and Community Development". Transcript of remarks presented at the Wharton Impact Conference *The Expanding Role of GIS in Business and Government*. August 21, 2002.

Findlay, Bill. 2002. "Designing a Data Utility", *GITA Proceedings*. Available on line as of 9/8/04 at http://www.gisdevelopment.net/proceedings/gita/2002/municipal/gita2002039pf.htm.

Fiorin, Franco and Carrera, Fabio. 1999. "L'evoluzione del sistema dei trasporti e del traffico acqueo" in *Venezia la Città dei Rii, op.cit.*

Fletcher, David, Espinoza, John, Mackoy, R.D., Gordon, Steven, Spear, Bruce and Vonderohe, Alan. 1998. "The Case for a Unified Linear Reference System", *URISA Journal*, Vol. 10, No. 1, pp. 7-11.

Flynn, D.J. 1992. *Information Systems Requirements: Determination and Analysis*. London: McGraw-Hill Book Company.

Flynn, Jennifer, Fydenkevez, Paul, McCarthy, Brina, O'Connell, Christine J. 2003. *Optimizing Parking Regulation Enforcement and Revenue Collection in* *the City of Cambridge.* Boston Project Center: WPI Interactive Qualifying Project.

Forester, John. 1989. *Planning in the Face of Power.* Berkeley: University of California Press.

Forester, John. 1993. *Critical Theory, Public Policy, and Planning Practice. Toward a Critical Pragmatism.* Albany: State University of New York Press.

Gaewsky, Kristopher, Gebhardt, Karl, Reanrungroch, Nicholas, Reynolds, Christine. 2003. *Safety and Reciclying in Large Buildings.* Boston Project Center: WPI Interactive Qualifying Project.

Gage, Matt, Johnstone, Kimberly, Shaver, Colleen, West, Josh. 2003. *Transportation Management in the City of Cambridge.* Boston Project Center: WPI Interactive Qualifying Project.

Gallo, Alberto. 1999. *Scarichi fognari e danni lungo le sponde dei canali interni di Venezia. Rapporto sul catalogo multimediale dei dati delle campagne '92-'93* for UNESCO-MURST project *Venice Inner Canals*; January 30.

Geddes, Patrick. 1911. *The Civic Survey of Edinburgh.* Edinburgh: Chelsea Civics department.

Geertman, Stan. 2002. "Inventory of planning support systems in planning practice: conclusions and reflections", in *Proceedings of the 5th AGILE Conference on Geographic Information Science,* Ruiz, M. and Gould, M., eds. Palma: Balearic Islands, Spain (http://www.agile-online.org).

Geertman, Stan and Stillwell, John., eds. 2003. "Planning support systems: an introduction" in *Planning Support Systems in Practice.* Berlin: Springer-Verlag.

Geertman, Stan and Stillwell, John. 2004. "Planning support systems: an inventory of current practice", *Computers, Environment and Urban Systems*, Vol. 28, No. 4, July, pp. 291-310.

Ghose, Rina and Huxhold, William. 2002. "The Role of Multi-scalar GIS-based Indicators Studies in Formulating Neighborhood Planning Policy". *URISA Journal*, pp. 5-17.

Gladwell, Malcom. 2000. *The Tipping Point.* Boston: Back Bay Books.

Godin, Seth. 2001. *Unleashing the idea virus.* New York: Hyperion.

Gomperts, Alexander, Lee, Joseph, and Tanefusa, Yusuke. 2004. *Ultrasound Boat-Monitoring System.* WPI Major Qualifying Project.

Habermas, Jürgen. 1984. *The Theory of Communicative Action.* Vol. 1., *Reason and the Rationalization of Society.* Boston: Beacon Press.

Habermas, Jürgen. 1987. *The Theory of Communicative Action.* Vol. 2, *Lifeworld and System: A Critique of Functionalist Reason.* Boston: Beacon Press.

Hall, Peter. 1988 (1998 reprint of updated 1996 edition). *Cities of Tomorrow: An Intellectual History of Urban Planning and Design in the Twentieth Century.* Oxford: Basil Blackwell.

Halloran, A., Rohleder, K., Malik, R., Fletcher, K.. 2002. *An Integrated, Multi-Agency Approach to the Preservation of Venetian Palaces.* WPI Interactive Qualifying Project.

Hamir, Akraad, Mallozzi, Nina, and Traynor, Kate. 2003. *Building a Better Future: Housing Development in Worcester.* Worcester Community Project Center: WPI Interactive Qualifying Project.

Hammond, K.R., McClelland, G.H., and Mumpower, J. 1980. *Human Judgment and Decision Making. Theories, methods and procedures,* New York: Praeger.

Hammond, K.R. , Harvey, L.O., and Hastie, R. Jr.. 1991. *Making Better Use of Scientific Knowledge: Separating Truth from Justice*, Center for Research on Judgment and Policy, Institute of Cognitive Science, July.

Hart, A., Hetrick, T., LeRay, D., LoPresti, E.. 2004. *Boston City Knowledge.* WPI Interactive Qualifying Project.

Harvey, Francis. 2001. "Constructing GIS: Actor Networks of Collaboration". *URISA Journal*, Vol. 13, No. 1, pp. 29-37.

Harris, B. 1999. "Computing in planning: professional and institutional requirements", *Environment and Planning B: Planning and Design*, vol. 26, pp. 321-331.

Hart, Adam, Hetrick, Todd, LeRay, David, LoPresti, Emily. 2004. *Applying E-Governemt Principles in the operation of the Boston Environment Department.* WPI Interactive Qualifying Project.

Hayes, H., Liu, J., Salini, C., Steinhart, A.. 2003. *Church Floors and Archeology.* WPI Interactive Qualifying Project.

Ho, Oliver, Kligman, Ricardo, Patel, Chirag, and Patel, Ravi. 2003. *Manuscript Transcription Assistant.* WPI Major Qualifying Project.

Hoch, Charles J., Dalton, Linda C., and So, Frank S., eds. 2000. *The Practice of Local Government Planning.* Third Edition. Washington, DC: International City/County Management Association (ICMA).

Hopkins, Lew D. 1999. "Structure of a planning support system for urban development". *Environment and Planning B: Planning and Design*, Vol. 26, No. 3, May, pp. 333-343.

ICMA (International City/County Management Association), 2002. *Electronic Government* survey (www.icma.org).

Innes, Judith E. and Simpson, David M.. 1993. "Implementing GIS for Planning", *Journal of the American Planning Association*, Vol. 59, No. 2, Spring, pp. 230-236.

Innes, Judith E. 1995. "Planning Theory's Emerging Paradigm: Communicative Action and Interactive Practice". *Working Papers of the Institute of Urban and Regional Development.* Working paper n. 629. Berkeley: University of California. February 95.

Innes, Judith E. 1996. "Information in Communicative Planning". *Working Papers of the Institute of Urban and Regional Development.* Working paper n. 679. Berkeley: University of California. October 96.

Innes, Judith E.. 1998. "Information in Communicative Planning", *Journal of the American Planning Association*, Vol. 64, No. 1, Winter, pp. 52-63.

Jajosky, Jessica, Moller, Christopher, Zarr, Joshua. 2003. *Economic Development Opportunities for Worcester.* Worcester Community Project Center: WPI Interactive Qualifying Project.

Johnson, Mark D., Lovisolo, Jonathan, and Okuno, Yasuhiro. 2003. *Ultrasound Boat-Monitoring System.* WPI Major Qualifying Project.

Johnson, Steven. 2001. *Emergence.* New York: Touchstone.

Keating, Gordon N., Rich, Paul M., and Witkowski, Marc S. 2003. "Challenges for Enterprise GIS", *URISA Journal*, Vol. 15, No. 2, pp. 23-36.

Keeney, R.L. and Raiffa, H. 1976. *Decisions with Multiple Objectives,* New York: John Wiley and Sons.

Kelly, Nina M. and Tuxen, Karin. 2003. "WebGIS for Monitoring 'Sudden Oak Death' in coastal California", *Computers, Environment and Urban Systems*, Vol. 27, pp. 527-547.

Kemp, Roger L. 1998. *Managing American Cities. A Handbook for Local Government Productivity.* Jefferson, NC: McFarland & Co.

Klosterman, Richard E.. 2000. "Planning in the Information Age". In *The Practice of Local Government Planning.* Hoch, Charles, J., Dalton, Linda C. and So, Frank S., eds. Washington, DC: International City/County Management Association.

Klosterman, Richard E.. 2001. "Planning Support Systems: A New Perspective on Computer-aided Planning" in *Planning Support Systems*, Brail and Klosterman, eds. Redlands, CA: ESRI press.

Kuhn, Thomas S.. 1962. *The Structure of Scientific Revolutions.* Second ed. (1970). Chicago: University of Chicago Press.

Laudon, K.C. and Laudon, J.P. 1996. *Management Information Systems: Organization and Technology.* Upper Saddle River, NJ: Prentice Hall.

Laurini, Robert. 2001. *Information Systems for Urban Planning.* New York: Taylor and Francis.

Lindbloom Charles. 1959. "The Science of "Muddling Through", *Public Administration Review*, Spring, pp. 79-88.

Linstone, H.A. 1984. *Multiple Perspectives for Decision Making*, New York: North-Holland.

Longley, P.A. and Harris, R. J. "Towards a new digital data infrastructure for urban analysis and modelling". *Environment and Planning B: Planning and Design*, Vol. 26, No. 6, November, pp. 855-878.

Malone, Thomas. Winter 1997. "Is Empowerment Just a Fad? Control, Decision Making, and IT." *Sloan Management Review.* Cambridge, MA.

Mancuso, Andrea N. 2003. *Un web gis per il centro storico di Venezia.* Doctoral thesis. Istituto Universitario di Architettura di Venezia. Venice.

Mandelbaum, Seymour J., Mazza Luigi and Burchell, Robert W., eds. 1996. *Explorations in Planning Theory.* New Jersey: Center for Urban Policy Research press.

Marchetti, P., Mazza, K., Hoey, S., Kahan, M.. 2003. *Adaptive reuse of Venetian Palaces and Convents.* Venice Project Center: WPI Interactive Qualifying Project.

Marchewka, Jack T. 2003. *Information Technology Project Management. Providing Measurable Organizational Value.* Hoboken, NJ: John Wiley and Sons.

Massachusetts Division of Employment and Training. 2003. *Employment and Wages* (ES-202). http://lmi2.detma.org/lmi/lmi_es_a.asp

Massachusetts Housing Finance Agency. 2002. *Housing List.* Boston.

Masser, Ian and Wilson, Tom. 1984. "Approaches to Information Management in County Planning Authorities in England and Wales", *Urban Studies*, Vol. 21, pp. 415-425.

McHarg, Ian L.. 1969. *Design with Nature.* Cambridge: MIT Press.

Mencini, G. 2000. *Sull'Onda Viva del Mare.* Venezia.

Mumford, Lewis. 1961 (1989 ed.). *The City in History: Its Origins, its Transformations, and its Prospects.* New York: MJF Books.

National Neighborhood Indicators Partnership, 1999. *Building and Operating Neighborhood Indicator Systems: A Guidebook.* Kingsley, G. Thomas, ed.. Washington, DC: The Urban Institute.

National Research Council. 1993. *Toward a Coordinated Spatial Data Infrastructure for the Nation.* Washington, DC: National Academy Press.

Nedović-Budić, Zorica. 2000. "Geographic Information Science Implications for Urban Regional Planning", *URISA Journal*, Vol. 12, No. 2, pp. 81-93.

Nedović-Budić, Zorica and Pinto, Jeffrey K.. 1999. "Understanding Interorganizational GIS Activities: A Conceptual Framework", *URISA Journal*, Vol. 11, No. 1, pp. 53-64.

Nedović-Budić, Zorica and Pinto, Jeffrey K.. 2000. "Information sharing in an interorganizational GIS environment", *Environment and Planning B: Planning and Design*, vol. 27, pp. 455-474.

Nedović-Budić, Zorica, Feeney, Mary-Ellen F., Rajabifard, Abbas and Williamson, Ian. 2004. "Are SDIs serving the needs of local planning? Case Study of Victoria, Australia and Illinois, USA", *Computers, Environment and Urban Systems*, Vol. 28, No. 4, July, pp. 329-351.

Nelson, Kimberly L. 2002. *Structure of American Municipal Government.* ICMA Special Data issue, No. 4. Washington, DC: International City/County Management Association.

Novello, Andrea and Sartori, Davide. 2004. *Progettazione ed avviamento di un sistema web-GIS per la gestione degli spazi ed aree pubbliche in concessione.* Tesi di I Livello. Venice: Istituto Universitario di Architettura di Venezia. Diploma in Sistemi Informativi Territoriali.

NRC. 1995. *A Data Foundation for the National Spatial Data Infrastructure*, Washington: Mapping Science Committee, National Research Council, 55 p.

O'Donnell, Malinda, Pollard, Turin, and Savain, Marvin. 2002. *Enhancing Public Safety in Boston: A Computerized System to Inventory and Track Underground Storage Tanks.* Boston Project Center: WPI Interactive Qualifying Project.

Olmsted, Frederick Law, Jr.. 1913. "A City Planning Program" in *Proceedings of the Fifth Conference on City Planning.* Boston.

Paluello, Carlo. 1900. *Note igieniche sul progettato ponte lagunare - con particolare riguardo alla malaria*, Tipolit. Ferrari:,Venezia.

Piasentini, Stefano. 1999. "Aspetti della Venezia d'acqua dalla fine del XIV alla fine del XV secolo" in *Venezia la Città dei Rii, op.cit.*

Pulliero, Augusto. 1987. *Canal Grande Mare Forza Tre.* Venezia.

Reeve, Derek and Petch, James. 1999. *GIS Organizations and People*, London: Taylor and Francis.

RKG Associates, Inc. 2002. *Housing Market Study: City of Worcester, Massachusetts.* Durham, NH.

Rorty, Richard. 1989. *Contingency, Irony and Solidarity.* Cambridge, UK: Cambridge University Press.

Samdadia, Viren. 2004. *Integration of 3D Parametric Building Model (3dPBM) and Geographic Information Systems (GIS) in Educational Facility Planning and Management.* WPI Masters Thesis in Civil Engineering.

Sawicki, D.S. and Flynn, P.. 1996. "Neighborhood Indicators", *Journal of the American Planning Association*, Vol. 62, No. 2, Spring, pp. 165-183.

Schiffer, Michael. 1992. "Towards a collaborative planning system".. *Environment and Planning B: Planning and Design*, Vol.19, pp 709-722.

Schön, Donald A. 1983. *The Reflective Practitioner. How Professionals Think in Action.* New York: Basic Books.Schuster, J. Mark, with de Monchaux, John and Riley, Charles A. II, eds. 1997. *Preserving the Built Heritage: Tools for Implementation.* Salzburg Seminar. Hanover, NH: University Press of New England.

Shuler, John A. 2003. "Geographic Information Resource Management", *UCGIS Research Priorities.* http://www.ucgis.org. Downloaded 2/24/03.

Simon, H.A. 1960. *The New Science of Management Decision.* New York: Harper & Row.

Singh, Raj R.. 2004. *Collaborative Urban Information Systems: A Web Services Approach.* PhD Thesis, Department of Urban Studies and Planning, MIT: Cambridge, MA.

Talen, E. 1999. "Constructing neighborhoods from the bottom-up: the case for resident-generated GIS", *Environment and Planning B: Planning and Design*, Vol.26, pp 533-554.

Todaro, Rudj. 2003. Personal Communication.

Tulloch, David and Fuld, Jennifer. 2001. "Exploring County-level Production of Framework Data: Analysis of the National Framework Data Survey", *URISA Journal*, Vol. 13, No. 2, pp. 11-21.

Tufte, Edward R. 1997. *Visual Explanations. Images and quantities, evidence and narrative.* Cheshire, CT: Graphic Press.

Ullmann, Jeffrey D. And Widom, Jennifer D. 2001. *A First Course in Database Systems.* New York: Prentice Hall.

Umgiesser, Georg and Zampato, Lucia. 1999. "Un modello Idrodinamico dei Canali Interni di Venezia" in *Venezia la Città dei Rii.*

U.S. Census Bureau. Census 2000. http://www.census.gov.

U.S. Census Bureau. 2001 MSA *Business Patterns..* http://censtats.census.gov/cgi-bin/msanaic/msasel.pl.

Vibert, Conor. 2004. *Theories of Macro Organizational Behavior. A Handbook of Ideas and Explanations.* Armonk, NY: M.E. Shape.

Williamson, I. and Ting, L.. 2001. "Land administration and cadastral trends – a framework for re-engineering" in *Computer, Environment and Urban Systems*, 25, 339-366.

Wilmersdorf, Erich. 2003. "Geocoded information incorporated into urban online services – the approach of the City of Vienna", *Computers, Environment and Urban Systems*, Vol. 27, pp. 609-621.

Wyatt, Ray. 1989. *Intelligent Planning. Meaningful Methods for Sensitive Situations.* London: Unwin Hyman.

Yeh, Anthony G.-O, 1999. "Urban Planning and GIS" in *Geographical Information Systems.* Longley, Paul A., Goodchild, Michael F., Maguire, David J. and Rhind, David W., eds. Second Edition, 2 Vols.. New York: John Wiley and Sons.

Yeh, Anthony G.O. and Webster, Chris. 2004. "Planning, government, information, and the internet". Guest editorial. *Environment and Planning B: Planning and Design*, Vol.31, No. 2, pp 163-165.

Yin, Robert K. 1994. *Case Study Research. Design and Methods.* Second ed.. Applied Social Research Methods Series, Volume 5. Thousand Oaks: Sage Publications.

Zuboff, Shoshana. 1991. "Informate the enterprise: an agenda for the 21 century". *National Forum* (summer).

Zucchetta, Gianpietro. 1992. *Un'altra Venezia.* Venice: Arsenale Editrice