The Convergence of Supply Chain Integration and Electronic Commerce

by

Jason B. Slibeck

B.S., Political Science (1992)
Massachusetts Institute of Technology

Submitted to the Department of Civil and Environmental Engineering in Partial Fulfillment of the Requirements for the Degree of

MASTER OF ENGINEERING IN LOGISTICS

at the

Massachusetts Institute of Technology

June 1999

Signature of Author _____
Department of Civil and Environmental Engineering
May 28, 1999

Certified by _____
George A. Kocur
Visiting Senior Lecturer
Department of Civil and Environmental Engineering
Thesis Advisor

Accepted by _____
Andrew J. Whittle
Chairman, Departmental Committee on Graduate Studies

The Convergence of Supply Chain Integration and Electronic Commerce

by

Jason B. Slibeck

Submitted to the Department of Civil and Environmental Engineering on May 28,1999, in partial fulfillment of the requirements for the Degree of Master of Engineering in Logistics.

ABSTRACT

Three forces drive the information economy: competition, customers, and change. Supply chain integration seeks to achieve optimal performance of the supply chain as a whole by finding the appropriate balance of "focused" excellence and process coordination. Successfully overcoming the corporate silos of the industrial age requires a focus on process, systems integration, and supply chain collaboration. New models of electronic commerce were often not economically feasible before the widespread deployment of the internet, but they fundamentally rely on the same principles of marketing and supply chain integration that govern the pre-internet business world. Intelligent agents, real-time personalized marketing, and online catalogs all rely on a solid relational database model. These business models have been successfully measured by analyzing their impact on Return on Assets and cost reduction. Recent implementations of technology such as the Extensible Markup Language (XML) hold much promise. Using XML as a communication standard, and with a strong process focus, the RosettaNet organization and other leading companies are demonstrating that the key elements of a successful business model in an internet economy are process-oriented supply chain integration and customer-focused electronic commerce.

Thesis Advisor: George A. Kocur

Title: Visiting Senior Lecturer

Biographical Note

Jason Slibeck is a native of Russell, Kentucky. He received a B.S. in Political Science from MIT and a commission in the United States Navy in January 1992. After advanced training at the Navy Supply Corps School in Athens, Georgia, he served overseas as a logistics officer aboard the USS SUMTER, a tank landing ship, and the USS DETROIT, a combat logistics ship. In 1995, he left the Navy to join Interport Communications, an internet service provider in New York City, as General Manager. He has also been the Acting Director of MIT Alumni Network Services, a pioneering online community, and a Senior Consultant at Syntra Limited, a global commerce software provider. In 1998, Jason became a member of the first class of graduate students enrolled in the MIT Center for Transportation Studies Logistics Program. He is married and became a new father while trying to finish this thesis.

Table of Contents

Table of Figures

Introduction

Overview and Objective

The inevitable collision between supply chain integration and electronic commerce will take many forms over the next several years. Both fields continue to evolve, or more realistically, to mutate. Spurred on by rapidly improving technology, a prosperous global economy, and a growing acceptance of open standards, these changes reflect the increasing clockspeed of an internet-enabled economy and promise a significant return on investment in information technology. The successful convergence of the network economy and the integrated supply chain may be analogous in this information age to a chassis dragged by rope and windlass along the floor of the Highland Park, Michigan, plant in the summer of 1913 at the apex of the previous industrial age.

Barriers still exist. While most large corporations are actively involved in the deployment of enterprise resource planning systems, few are integrating these systems beyond corporate boundaries to third party suppliers, outsourcing partners or their customers. Furthermore, the deployment of enterprise resource planning systems and the information integrity they represent is just beginning to penetrate companies not in the top tier of global industrial firms. Indeed, significant market opportunities still exist in many key areas of transaction processing and execution such as international trade, warehouse operations, and transportation management. A rich marketplace is evolving to fill these gaps, but even these companies focus primarily on a single company approach. Supply chain integration has almost become a cliché in corporate marketing material, but it is still far from being a reality in most organizations.

Barriers exist below the hype of electronic commerce as well. At the March 1999 Internet Commerce Expo, not a single exhibitor had an integrated solution for linking the customer-focused front end of internet marketing to the back-end reality of order fulfillment, demand management, picking, packing, and delivery. The electronic commerce world is somewhat prepared to securely validate a credit card on a personalized web page tailored to your eclectic tastes. Yet, in the spring of 1999, most vendors still rely on a manually entered offline spreadsheet to check inventory availability and a dot matrix printed pick slip to retrieve an order from a warehouse shelf. Getting it to its proper destination using the most efficient means of transport and verifying the order status at any point in the process are still unrealized goals.

These barriers are being overcome in a variety of ways and in certain situations where front runners have recognized the power of closing the gap between electronic commerce and supply chain integration. This thesis hopes to examine some of the creative solutions, apply them to a business model, and hopefully reveal some important lessons about using the modern equivalent of a rope and windlass in an Internet world.

Outline

The first section of this thesis is the Introduction, and covers an overview of the issues and includes this outline. The next section, Supply Chain Integration, begins with a definition of a supply chain, logistics and supply chain integration. Corporate and functional silos are placed in a proper historical and modern context. Limitations to this organizational model are discussed and an alternative approach, process orientation, is introduced. An analysis of

the collaboration between Diamond Multimedia and Skyway shows the application of supply chain integration and provides some information on the success of this partnership venture.

Electronic Commerce is the next section. Different business models that take advantage of the unique marketplace on the internet are introduced. Intelligent agents are explored through a look at MortgageQuotes.com. Customer loyalty programs at eBay.com and at Levi.com illustrate two different approaches to personalization. Following these discussions is an in depth look at the process behind building an electronic commerce website and at the relational data model necessary to support the process at MIT Press.

Measurement and Benefit Analysis uses the DuPont Model to illuminate the broad impact supply chain integration and electronic commerce can have. This model is used to show Return on Assets as a key indicator of success or failure. Cost efficiency is a major driver behind many efforts at supply chain integration or electronic commerce. Cisco Systems is used to illustrate this point.

One of the key Convergence Technology Enablers is shared information systems. XML is introduced and its potential role in future convergence efforts is described. The architecture of the language and its current status is outlined. RosettaNet.org is used to show the detailed process that the high tech industry has created to apply XML to convergence problems and to illustrate the kind of solution that is possible.

Finally, some conclusions on the issue of convergence, supply chain integration, and electronic commerce are drawn.

Supply Chain Integration

In this section, the concept of supply chain integration is introduced and defined. Supply chain integration is a way to address the problems of corporate silos and independent systems. These problems are identified and discussed. One way to grasp the complexities of supply chain integration is the use of process maps to clarify the relationships between corporate activities and to clearly establish the flow of goods, information, and cash. Supply chain integration is dependent upon advances in information systems and system integration for effective collaboration to occur. Finally, process mapping and systems integration are linked together to illustrate how supply chain integration enables collaboration between organizations such as third party logistics providers and manufacturers.

Supply Chain Defined; Concept Introduction

In 1985, John Houlihan introduced the term "supply chain management" in an article in the *International Journal of Physical Distribution & Materials Management.*[1] The definition of a supply chain has evolved significantly as more and more companies recognize the need to move beyond functional excellence and toward a process-oriented business approach that includes the entire scope of activity from the extraction of raw material to the satisfaction of end user demand. In 1995, the Council for Logistics Management defined logistics as the process of planning, implementing, and controlling the efficient, effective flow and storage of goods, services, and related information from point of origin to point of consumption for the purpose of conforming to customer requirements. This view of logistics accompanied by an

---

[1] Houlihan, John B., "International Supply Chain Management," *International Journal of Physical Distribution & Materials Management 17:2 (1985): 51-56.*

understanding of the manufacturing process leads to an understanding of the broad scope that a supply chain approach can have.

Therefore, a supply chain is the set of processes that plan, source, produce and deliver products or services to customers through the management of the flow of materials, information, and cash. The notion of a single chain however can be misleading since modern organizations are more often a network of processing nodes; the analogy of a chain directly highlights the dependency of each link on every other link. The motivation behind supply chain integration is the recognition that traditionally "focused" behavior both within and between companies results in sub-optimal performance of the supply chain. Supply chain integration seeks to achieve optimal performance of the supply chain as a whole by finding the appropriate balance of "focused" excellence and process coordination.[2]

Corporate Silos and Independent Systems

Corporate silos and independent systems are the greatest achievements of the industrial revolution. This should not be a surprising statement if one remembers that dinosaurs were ideally adapted to the Jurassic period. In Wealth of Nations, Adam Smith analyzed the workings of a pin factory and introduced the world to the principle of division of labor.[3] This division was possible because the pin making process could be deconstructed into specialized tasks. Advances in industrial technology facilitated labor and allowed one worker to do the work of many. Henry Ford refined the concept even further and the modern industrial age was nearly complete. However, it wasn't until Alfred Sloan carried the division

---

[2] Franciose, Michelle M., Supply Chain Integration: Analysis Framework and Review of Recent Literature. Thesis. MIT Master of Science in Transportation, M.S. C.E. (1995), 14.

of labor concept from the factory floor to the management bureaucracy that the subdivision of a process into functionally efficient units was complete. The creation of this industrial management system made the expansion of General Motors possible and ushered in an era of management by financial accounting goals and marketing objectives for functional business units with unique goals and performance metrics.[4]

The era of unprecedented growth, which fueled this migration to easily scalable hierarchical organizational pyramids, was not able to withstand the information age and the three forces of the modern economy: customers, competition, and change.[5] The national mass markets of the 1950s, the 1960s, and 1970s have been replaced by an empowered set of business and individual customers. These customers know what they want, what they want to pay, and which service level will accommodate their needs. This strategic shift in power is in part a result of easier access to an enormous amount of information for price comparison, quality judgement, and producer reliability.

The national market forces also now face stiff competition as global trade becomes a reality. International differences in factor costs such as labor rates have created enormous sociological changes. Competition has come from entrepreneurial innovators as well. Entire sectors of the freight forwarding and distribution industry are now trying to expand the scope of their operations to justify a continued existence as more nimble logistics providers force the question, "How do you add value for the customer?"

---

[3] Smith, Adam. An Inquiry Into the Nature and Causes of the Wealth of Nations. (1776), available at http://www.geolib.pair.com/essays/smith.adam/woncont.html (May, 1999).
[4] Sloan, Alfred P., Jr. My Years With General Motors. Doubleday (1963).
[5] Hammer, Michael & James Champy. Reengineering the Corporation: A Manifesto for Business Revolution. Harper Collins (1993), page 19.

Finally, change has become a constant state of existence and it is accelerating. To keep up with nimble global competitors and to react effectively to global market inputs, companies have accelerated the product development lifecycle from years down to months. With only a limited time frame to react, accurate forecasting and flawless execution have become critically important. Supply chain integration requires a re-examination of these industrial revolution assumptions, and a move beyond functional excellence.

Beyond Corporate Silos

Process Mapping

One of the key tools used by companies when they begin to address the issues of supply chain integration is process mapping. Processes in a company correspond to natural business activities, but they are often fragmented and obscured by the organizational structures.[6] Fred Hewitt defines a business process as "a set of logically related tasks performed to achieve a defined business outcome." Or, a business process could be equated to a transformation function that uses energy to convert inputs to outputs. Process efficiency then can be measured as the reciprocal of the energy lost during the transformation process, usually expressed as time or cost. Effectiveness then follows as the amount of time the output of the process conforms to specifications.[7]

---

[6] Hammer, Michael & James Champy. Reengineering the Corporation: A Manifesto for Business Revolution. Harper Collins (1993), page 118.
[7] Hewitt, Fred. "Supply Chain Redesign," *International Journal of Logistics Management* 5:2 (1994), pages 1-9.

Collaboration and System Integration

Skyway[8] and Diamond Multimedia[9] provide an example of supply chain integration and collaboration on the introduction of a new product.[10] Diamond Multimedia is a $450 million manufacturer of media devices for home and business personal computers. The young company's success can be attributed to quick identification of consumer needs coupled with an ability to rapidly bring affordable, high quality products to market ahead of the competition. Skyway is a well-established third-party logistics provider with a comprehensive, information-driven approach to demand planning, order fulfillment, and warehouse management.

Diamond's latest product, the Rio, was a portable .MP3 music player and would represent a departure from the traditional products which Diamond sold primarily through indirect channels like value-added resellers, distributors, and retailers. Since .MP3 music files are typically downloaded from the internet, it only made sense to realize the higher margins and higher volumes that Diamond could obtain by selling directly over the internet. With only six weeks before product launch, Diamond recognized that they could not effectively handle the volume of traffic, the credit card processing, warehouse picking and order management of single unit SKU's, so they turned to Skyway.

Skyway and Diamond focused on integrating the technology and information infrastructures. Diamond was maintaining the front-end web site. Information needed to flow from there to Skyway's credit card verification, and then, if approved, to a check for availability and

---

[8] http://www.skyway.com/supplychain/solutions.html (25 May 1999).
[9] http://www.diamondmm.com/products/current/rio-le.cfm (25 May 1999).
[10] Benchmarking Partners. "The Power of Business-to-Business Integration" Benchmarking Partners, Inc. (1998), page 17.

picking in the warehouse. After pick, pack and ship, Skyway contacts the merchant bank for final billing and the bank transfers funds to Diamond. Skyway also regularly feeds information directly into Diamond's enterprise resource planning system where inventory levels are monitored for replenishment. Skyway is internally configured as a sub-inventory location and in is handled as if it were a Diamond location.

By working beyond functional boundaries and beyond corporate borders, Skyway and Diamond react as a single integrated supply chain focused on delivering maximum value to the customer. Using Skyway's technical and information expertise, Diamond was able to successfully launch the Rio in a very short period of time. For Skyway, the benefits were the loyalty and retention of a long-term stable customer relationship and higher margins from value-added supply chain management services beyond the traditional role of a third party logistics provider. Besides the revenue gains of a highly profitable new product line, Diamond was able to achieve customer service levels beyond its internal capability. They also saved an estimated $250,000 to $500,000 in manual labor costs during the time it would have taken the company to convert to an automated internal system. Based on the manual labor savings alone, Diamond achieved a high return on investment, 5 times in the first six months.

Electronic Commerce

Focus

Electronic commerce on the internet has been exhibiting alternatives to standard business models. While some companies are essentially building the online equivalent of a mail-order catalog, other approaches are emerging which capitalize on the unique nature of the internet. These new models were often not economically feasible before the widespread deployment of the internet, but they fundamentally rely on the same principles of marketing and supply chain integration that govern the pre-internet business world. Intelligent agents, real-time personalized marketing and integrated supplier-customer purchasing systems are creating new models for business.

Competitive Markets

The Internet for many products is approaching the economic ideal of a frictionless market. This implies a market with no profit margin, but barriers can be erected by extending the product concept to include considerations beyond price. The online mortgage market provides a good example of both the exploitation of an internet technology, intelligent agents, and the value of branding.

Intelligent agents, known as shopping bots, are programmed search tools that allow buyers to evaluate products from a diverse set of suppliers. The current generation of bots can collect some information beyond a base cost such as shipping charges and inventory status, but the level of decision-making is still a very primitive form of comparison shopping and is

subject to manipulation. MortgageQuotes.com[11] has an internet bot that catalogs over 1,300 lenders. To ensure that none of the participants are engaging in bait-and-switch, they use staff review, consumer feedback and software algorithms that check for exceptions to a reasonable range of values based on current market offerings. To date, 44 lenders have been removed.

The gnawing worry about shopping bots is that they will undermine brands and focus all competition on price. So far that concern isn't proving true. "Most people are wary of the absolute lowest rate," says Charles Collins, director of corporate development at Microsurf Inc., the company behind MortgageQuotes.com. "Our stats show that a competitive rate, not the lowest rate, is the one that gets the click-throughs. People wind up evaluating lenders based on lots of other criteria: location, style, rapport, ease of use." Mortgage seekers want a lender they can trust. Savvy online merchants understand that building a brand is less about manipulating mass perception of your product or service than about creating an enjoyable customer experience—at the purchase stage and beyond. Sites that offer that kind of experience will convert onetime buyers referred by bots into loyal customers; they will build brand equity, and suddenly consumers may pay a small premium for a trouble-free shopping experience with a trusted merchant.[12]

---

[11] http://www.mortgagequotes.com/policy.html (25 May 1999).
[12] Kirsner, Scott. "Shopping bots are back." CNN Interactive, http://www.cnn.com/TECH/computing/9905/14/bots.ent.idg/ , (14 May 1999).

Customer Loyalty

Companies do best when they define their target market carefully and prepare a tailored marketing program.[13] The key is to fully understand the customers' needs. eBay.com[14], an online person-to-person auction, handles 400,000 auctions daily. eBay's typical user spends 1 ½ hours each month on their site. To understand these registered users, eBay regularly analyzes the history of all online transactions. Based on this analysis of user demand, service is tuned to rapidly respond to changes in the market. For instance, when eBay noticed that growing numbers of bidders hailed from Japan and Australia, it decided to launch versions of its auction service in those areas. Customer analysis also changes the product mix, with eBay adding new categories as interest in different types of collectibles grows. The number of categories is now up to 1,000. eBay relies on a simple user interface, so that even the least internet savvy individual can just jump in and use it.[15] eBay has chosen to overbuild hardware and engineering capacity to handle traffic and changing market requirements. This strategy allows the broadest possible customer base and minimizes customer technical support requirements.

Other e-commerce sites use a different technology approach to achieve customer loyalty and improve the sell-through rate. At Levi.com[16], shoppers take a quick survey that ranks their likes and dislikes in music, fashion and fun activities. Based on the answers to a few questions, the Levi site displays a personalized apparel catalog. Teenage ravers would be shown baggy club clothes, for instance, while older Levi customers would be offered more

---

[13] Kotler, Philip. Marketing Management: Analysis, Planning, Implementation, and Control. 8th Edition, Prentice Hall (1994), 19.
[14] http://pages.ebay.com/aw/nut-start1.html (25 May 1999).
[15] Messmer, Ellen. "Secrets of the e-commerce stars," CNN Interactive, http://www.cnn.com/TECH/computing/9905/25/ecomm.idg/ , (25 May 1999).
[16] http://store.us.levi.com:80/store/home.asp (26 May 1999).

classic styles. The technology for the Levi Style Finder comes from San Francisco-based Andromedia Inc.[17]

"By incorporating personalization into our web site, we are able to show products to consumers that they are interested in buying," said P.J. Santora, Levi's director of database marketing, adding that shoppers who use Style Finder visit the site more often and buy more merchandise when they do. To further reach out to web-savvy customers, Levi lets shoppers collect, save and share their favorite Levi items with friends. For instance, a shopper can pick out a Levi skirt, save it in her personal collection, and then e-mail a friend to get her feedback.[18]

Andromedia produces a personalization web server product called LikeMinds. LikeMinds collaborative filtering technology uses individual purchase history or preferences to find people with similar tastes, and predicts in real-time what each person will like, based on the choices made by other like-minded individuals. The system is a middleware layer between a web server and a relational database management system.

---

[17] http://www.andromedia.com/products/likeminds/benefits.html (27 May 1999).
[18] Feuerstein, Adam. "E-commerce gets personal," San Francisco Business Times (19 April 1999). Available at http://www.amcity.com/sanfrancisco/stories/1999/04/19/story2.html?h=Andromedia .

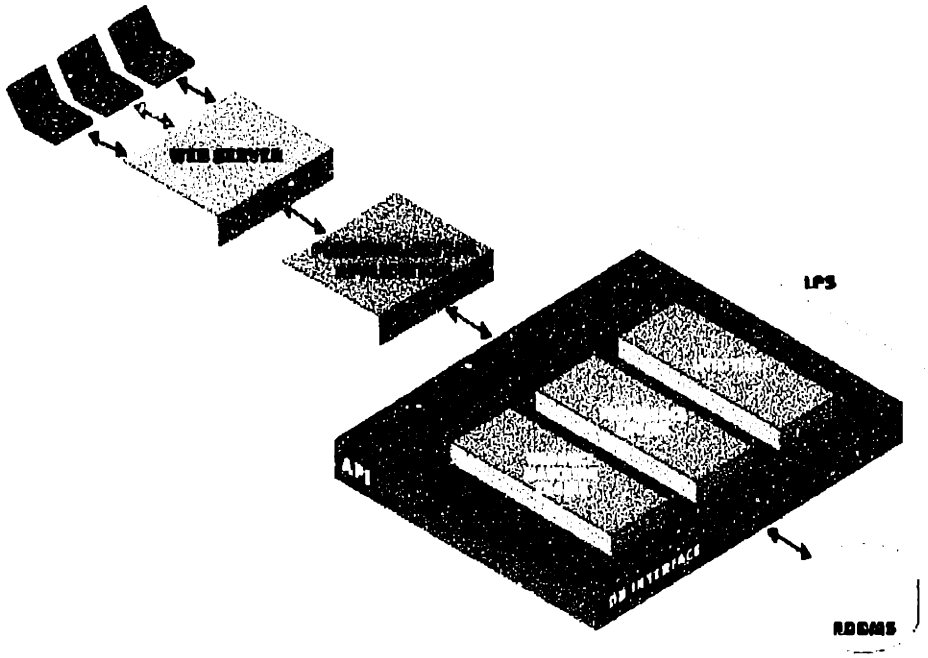

Figure 1 Architecture of Andromedia's LikeMinds Personalization Software[19]

E-commerce Architecture: MITPress Online[20]

MIT Press is a publisher of books and academic journals. Since 1993, they had been putting progressively more catalog info and content on their Web site, which was developed and maintained by a part-time webmaster. By 1996, the job had expanded to a full-time position. MITPress approached Philip Greenspun, a founder of ArsDigita, a Web development and hosting firm that supports open-source toolkits for building community and electronic commerce websites, with the following goals for the Web site, http://mitpress.mit.edu :

- Serve up a catalog of 6,000 books and 40 journals in a consistent fashion; the existing collection of static files was becoming too cumbersome to maintain.
- Collect comments on the catalog items.
- Collect names and contact information for readers who wished to be on various mailing lists.
- Collect orders for products.
- Allow MIT Press personnel to spam groups of readers with announcements, e.g., "those readers who've signed up to the cognitive science mailing list".

---

[19] http://www.andromedia.com/products/likeminds/architecture.html (24 May 1999).
[20] Excerpted and edited from Greenspun, Philip. Philip and Alex's Guide to Web Publishing. Maorgan Kaufman Publishers Inc. (1999).

MIT Press maintained overlapping catalog data in several separate systems. Their accounting/ordering system was a turnkey software package for academic publishers. The production staff maintained a FoxPro database with information on current and past projects. The legacy Web site had reasonable history information but these data were merely arranged in Unix filesystem files. The management of the Press was not ready to consider fundamental changes to the rest of the business, so a decision was made to maintain independent e-commerce and financial systems with a batch file sent daily for reconciliation and credit card verification.

A primary purpose of the Web site was to facilitate on-line shopping. An invisible shopping cart where a user would click on a product and get an order form was used. After submitting it with shipping address and credit card number, a magic cookie would be written back to the user's browser. Orders initiated the previous day would be transmitted in a batch from the RDBMS to the legacy system every morning at 7:00 am. That meant a user had between 7 and 31 hours to complete an order using the invisible cart.

AOLserver connected to the Illustra database was selected for the webserver and RDBMS respectively. MIT Press ran the system on a SPARC 5 Web server and had not budgeted for a larger machine. They expected to serve about 200,000 hits a day from this computer, some of them requiring RDBMS queries.

Using International Standard Book Number (ISBN) alleviated many thorny data modeling issues. A text file dump of the data in the accounting system was converted into SQL with

some Perl scripts and fed into the Illustra's Unix shell client. AOLserver Tcl scripts were written to semi-automatically transfer over book jacket images and long descriptions from files on the old Web server. A decision was made not to store book and journal images in the RDBMS. Instead, a Unix file system directory was created for each catalog item. For books, the AOLserver Tcl script would look in the directory for cover-sm.gif or cover-sm.jpg. If present, it would be displayed on the book catalog page.[21] If not, a page with a slightly different design is served.

MIT Press brought Ben Williams[22] aboard the project team to do the graphic design. Each book page had the following computational structure:

1.Query the books table for basic title and description for the ISBN key in the URL.
2.Stream out the top part of the catalog page.
3.Query the reader_comments table to pull out reader comments for display right on the page.
4.Stream out the reader comments, if any.
5.Query the books table again to find related items, such as a paperbound version of the same title.
6.Stream out the related items (as hyperlinks).
7.Do an expensive LIKE query against the books table to find other books by the same author(s).
8.If found, stream out the list of books.

To have a site that felt responsive, the site was designed with stacked tables. The first table contained the information from the first three RDBMS queries and the last contained the "books by the same author" content plus the page footer.

[21] http://mitpress.mit.edu/book-home.tcl?isbn=0262522349
[22] www.blueperiod.com

The AOLserver/Illustra combination requires minimal maintenance. About four person-weeks of programming went into creating the site, yet it has the same functional capabilities as amazon.com[23] circa 1997 (amazon had more than 10 full-time programmers). Because packaged software was not used, lots of shortcuts were programmed into the admin side to facilitate ease of use. An annotated data model file is included as Appendix A.

[23] http://www.amazon.com/exec/obidos/ASIN/0139397299/qid%3D927839756/002-1399425-5052640 (28 May 1999).

Measurement and Benefit Analysis

Return on Assets (ROA)

Supply chain integration and its convergence with electronic commerce can be understood and illustrated using the DuPont model devised by F. Donaldson Brown around 1914 at the apex of the industrial revolution.[24] The DuPont model in the following figure shows the impact of a supply chain order-to-cash process and links into key financial indicators from the corporate balance sheet and income statement. The model can also be used to show that a focus on electronic commerce enabled supply chains directly impacts return on assets. The focus can be on increasing sales through personalization and customer loyalty, reducing fixed expense and fixed assets by moving from traditional bricks and mortar distribution to electronic storefronts, or reducing inventory and the cost of goods sold by focusing on improved forecast reliability and integrated purchasing systems.

[24] Chandler, Alfred D., Jr., Thomas McCraw, and Richard Tedlow. Management Past and Present: A Casebook on the History of American Business. South-Western College Publishing (1996), page 3-73.

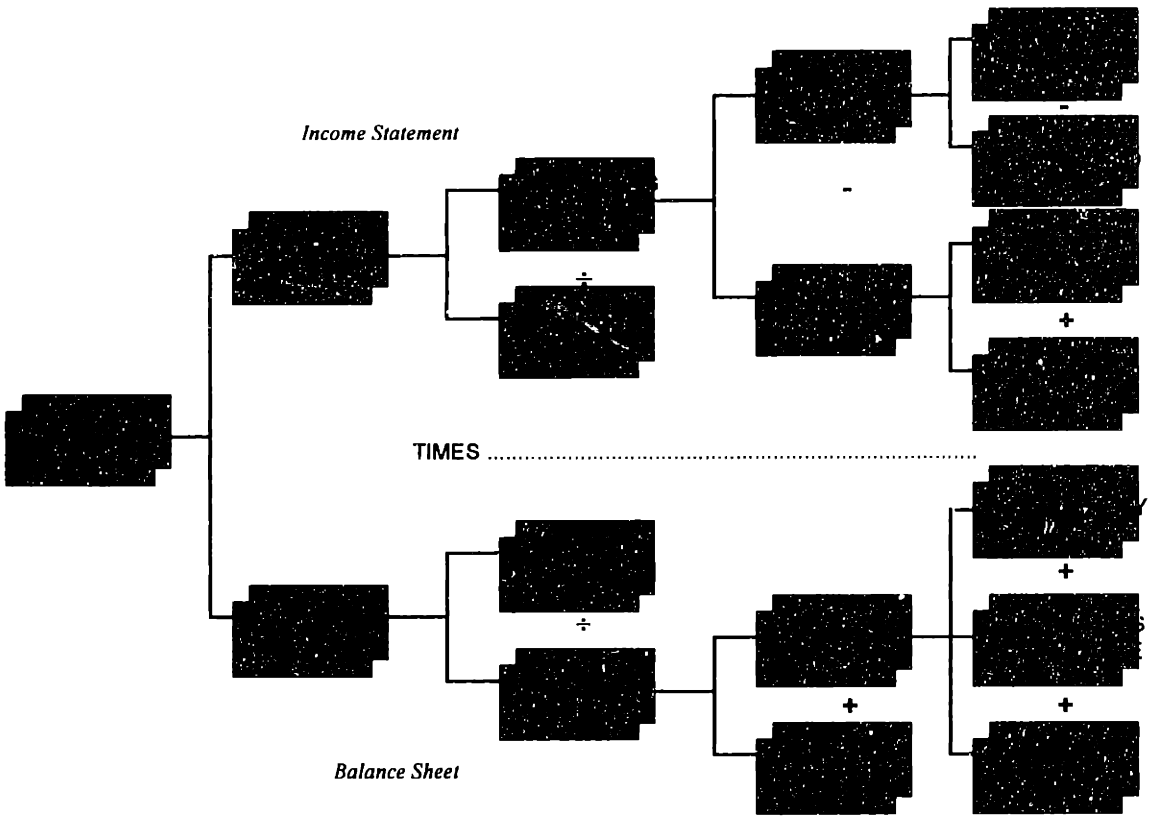

Figure 2 The DuPont Model

Cost Reduction

Cisco Systems[25] is a leading provider of internet routers, switches, telephony support products and other network hardware and services. 75% of its total equipment orders arrive through its Web site; that's 2,000 purchase orders for a total of $23 million worth of business each day.[26] When the company began using the Web for business transactions almost three years ago, Cisco hooked cisco.com into its in-house Oracle enterprise resource planning (ERP) system. From the beginning, every order went directly into the order-management process without having to be keyed in again.

---

[25] http://www.cisco.com/ (28 May 1999).
[26] Messmer, Ellen. "Secrets of the e-commerce stars," CNN Interactive, http://www.cnn.com/TECH/computing/9905/25/ecomm.idg/ , (25 May 1999).

This saves considerable time and promotes the accuracy of the order. Accuracy is also safeguarded by the Web site's "configuration engine," based on a product from Internet Commerce Connection[27]. The engine looks over the order to spot commonly found errors. The Web engine automatically checks customer account information and purchase details against a database of technical information. By eliminating errors in order entry and by tying the information directly into their ERP system, costs for handling exceptions and error control can be drastically reduced.

[27] http://www.inetcc.com/Default.asp (28 May 1999).

Convergence Technology Enablers: XML

Shared information is widely recognized as one of the most important enablers for supply chain integration. Even within a well-organized supply chain, disruptions and inefficiencies can result from inadequate information sharing. Both information technology and integrated information systems play significant roles in coordination of the supply chain. Integrated information systems are needed to coordinate the movement of information throughout the supply chain.[28] One of these emerging enablers is the Extensible Markup Language.

XML Architecture[29]

Extensible Markup Language (XML) is a simple, very flexible text format based on SGML (ISO 8879). Designed to meet the challenges of large-scale electronic publishing, XML will also play an increasingly important role in the exchange of a wide variety of data on the Web.

XML will:

- Enable internationalized media-independent electronic publishing
- Allow industries to define platform-independent protocols for the exchange of data, especially the data of electronic commerce
- Deliver information to user agents in a form that allows automatic processing after receipt
- Make it easy for people to process data using inexpensive software
- Allow people to display information the way they want it

---

[28] Franciose, Michelle M., Supply Chain Integration: Analysis Framework and Review of Recent Literature. Thesis. MIT Master of Science in Transportation, M.S. C.E. (1995), 57.

[29] Excerted and edited from Extensible Markup Language (XML) Activity. W3C Architecture Domain Activity Statement. http://www.w3.org/XML/Activity.html (1 May 1999).

- Provide metadata -- data about information -- that will help people find information and help information producers and consumers find each other

Simple Example of XML Usage

The best way to appreciate what XML documents look like is with a simple example. Imagine a company sells products on-line. Marketing descriptions of the products are written in HTML, but names and addresses of customers, and also prices and discounts are formatted with XML. Here is the information describing a customer:

```
<customer-details id="AcPharm39156">
  <name>Acme Pharmaceuticals Co.</name>
  <address country="US">
    <street>7301 Smokey Boulevard</street>
    <city>Smallville</city>
    <state>Indiana</state>
    <postal>94571</postal>
  </address>
</customer-details>
```

The XML syntax uses matching start and end tags, such as <name> and </name>, to mark up information. A piece of information marked by the presence of tags is called an element; elements may be further enriched by attaching name-value pairs (for example, country="US" in the example above) called attributes. Its simple syntax is easy to process by machine, and has the attraction of remaining understandable to humans. XML is based on SGML, and is familiar in look and feel to those accustomed to HTML.

Building Applications with XML

XML is a low-level syntax for representing structured data. This simple syntax can be used to support a wide variety of applications. This idea is demonstrated in a simplistic way in the diagram below, which shows how XML now underpins a number of Web mark-up languages and applications.

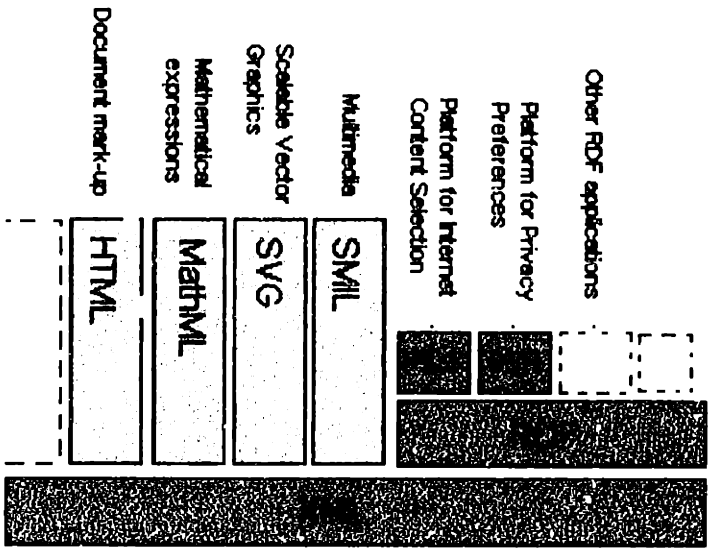

Figure 3 XML Application Architecture

The World Wide Web Consortium[30] (W3C) are redeveloping HTML as a suite of XML tag sets so that, although documents will still be marked up using HTML, this will conform to the rules of XML. Presumably, other domain-specific XML-based tag sets will become candidates for inclusion in HTML documents.

[30] http://www.w3.org/MarkUp/#future (28 May 1999).

Outside of W3C, many organizations are defining new formats for information exchange. Dell Computer[31] has announced that it will be using XML as an interface between its supplier's enterprise resource planning systems and its own online order management system.[32] Sun Microsystems and Netscape Communications have formed an alliance with a dozen other charter members to launch the Netscape/Sun Supplier Network to provide XML product feeds to Netscape procurement applications.[33] Sterling Commerce has also announced plans to integrate XML with EDI systems for enterprise applications.[34] As with many new technologies, competing XML specifications are emerging though. But, two of the dominating players, Ariba Inc. and Microsoft, have agreed to work together on a common tag set for e-commerce.[35] An examination of RosettaNet.org, a high tech industry trade group, illustrates the application of XML and the process an industry must go through to create a new standard.

Application- RosettaNet.org[36]

The lack of electronic business interfaces in the IT supply chain puts a huge burden on manufacturers, distributors, resellers, and end-users, ultimately creating tremendous inefficiencies and ultimately inhibiting the ability of the industry to leverage the Internet as a business-to-business commerce tool. Here are a few examples:

---

[31] http://www.dell.com/corporate/index.htm (27 May 1999).
[32] Booker, Ellis. "XML Applications Stand Up To EDI" Internetweek. http://www.internetwk.com/story/INW19990415S0007 (15 April 1999).
[33] Karpinski, Richard. "Sun-Netscape Alliance Outlines E-Commerce Plans" Internetweek. http://www.internetwk.com/story/INW19990426S0002 (26 April 1999).
[34] Karpinski, Richard. "Sterling Delivers E-Commerce Capabilities" Internetweek. http://www.internetwk.com/story/INW19990514S0003 (14 May 1999).
[35] Karpinski, Richard. "E-Commerce Tools, Standards Alleviate" Internetweek. http://www.internetwk.com/story/INW19990513S0004 (13 May 1999).
[36] Excerpted and edited from http://www.RosettaNet.org/general/overview.html (May 1999).

Manufacturers today utilize complex processes to all but guess inventory levels and locations across the supply chain at any point in time. This is because there is no agreement on something as simple as how a part number is defined or how inventory queries can be made through a standard interface. This significantly impacts production planning, channel allocation, and the cost of returns.

Distributors, who provide pre- and post-sale technical support to their resellers on tens of thousands of SKU's, must grapple with disparate forms of product information collected from hundreds of manufacturers with no common taxonomy. The lack of product information standards makes the current aggregation and dissemination of such content an expensive and inefficient proposition - an effort duplicated by each distributor in the channel. This problem is further compounded by the content's explosive rate of change.

Resellers must learn and maintain different ordering/return procedures and system interfaces to each distributor and direct manufacturer with whom they trade, causing them to spend valuable resources in back-office operations (50% by some estimates), which they could otherwise use to make new sales or service their customers.

End-users have no mechanism enabling effective procurement through uniform templates, which can be contextually linked to government authorized schedules. This often causes a nonsensical lengthening of the purchasing cycle whereby most PC orders are old technology by the time they make it through this inefficient cycle and onto the desk of the requisitioner.

Electronic business is fundamentally changing how IT products and services are bought and sold. The Internet is fueling the growth of electronic business and driving all members of the supply chain to find news ways to strengthen customer relationships, to identify and capitalize on new revenue opportunities, and to create operational efficiencies. Much research points to the economic efficiencies derived from standard setting in any industry. Similar common business interfaces will undoubtedly result in measurable cost savings and productivity gains throughout the IT supply chain.

The ubiquity of the Internet, which dwarfs the coverage of EDI VANs, is precipitating the dire need for such common business interfaces, as more commerce takes place electronically requiring powerful search features, cross-channel inventory access, and deep content - all of which necessitate strong industry-wide agreements on common formats, processes, content, and system interfaces. Moreover, the Internet provides a unique opportunity to effectively establish and maintain these interfaces.

While EDI and bar code standards present a formidable foundation for electronic transactions across many industries, RosettaNet.org was created to complement them with deeper IT industry-specific electronic business interfaces which not only address basic transactional data, but encompass all aspects of commercial relationships.

The need for such common electronic business interfaces is depicted in the following diagram, showing the parallel between a human-to-human business exchange and a server-to-server electronic business exchange. In order to communicate in a human-to-human business exchange, humans must be able to produce and hear sound. Further, they must

then agree on a common alphabet, used to create individual words. Grammatical rules are then applied to the words to create a dialog. That dialog forms the business process, which is conducted (or transmitted) through an instrument such as a telephone.

The fundamental system of exchanging sounds in a human-to-human business exchange can be compared to the Internet, which enables two servers to exchange information during a server-to-server electronic business exchange. HTML/XML functions as the "alphabet" of this electronic exchange. And, presently electronic commerce applications serve as the instrument by which an electronic business process is transmitted.

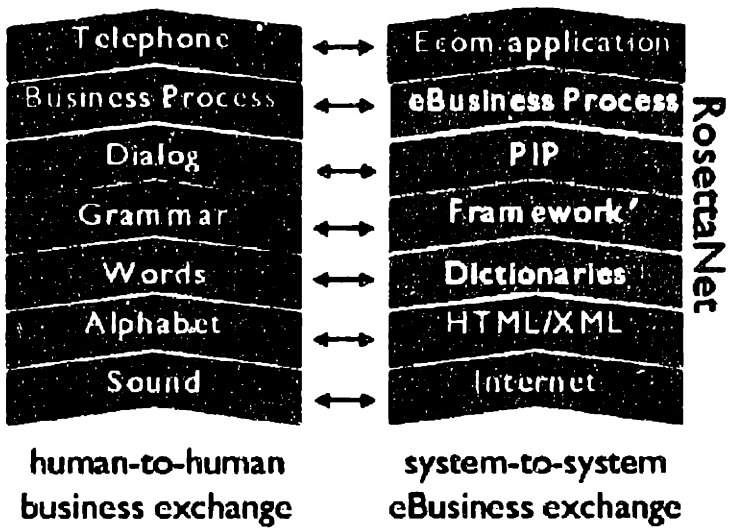

Figure 4 RosettaNet.org Technology Communication Analogy

What is missing in order to scale eBusiness are the "dictionaries," the "framework," the "Partner Interface Processes - PIPs" and the "eBusiness processes." RosettaNet fills this existing gap by focusing on building a master dictionary to define properties for products, partners, and business transactions. This master dictionary, coupled with an established

implementation framework (exchange protocols), is used to support the eBusiness dialog known as the Partner Interface Process or PIP. RosettaNet PIPs create new areas of alignment within the overall IT supply-chain eBusiness processes, allowing IT supply-chain partners to scale eBusiness, and to fully leverage e-commerce applications and the Internet as a business-to-business commerce tool.

The process used to create RosettaNet's common Partner Interface Processes (PIPs) is set forth below:

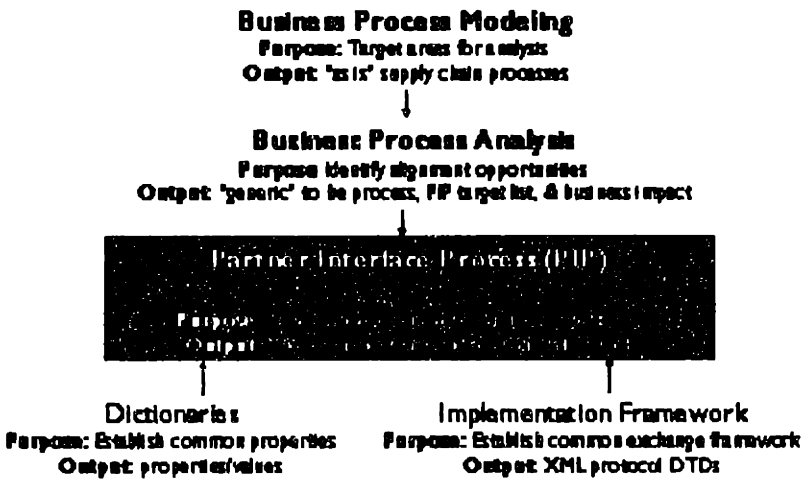

Figure 5 RosettaNet.org Partner Interface Process Creation Framework

Business Process Modeling

Business Process Modeling is used to identify and quantify the individual elements of a business process, creating a clearly defined model of the supply chain partner interfaces as they exist today. This is called "as is" modeling. This model reflects the results of extensive

research at every level of the supply chain and is then analyzed to identify any misalignments or inefficiencies.

Business Process Analysis

Through the analysis of the detailed "as is" model, a "generic to-be" process emerges, showing the opportunities for re-alignment in the form of a Partner Interface Process (PIP) target list, and estimating the business impact of implementing the resulting PIPs (savings as a function of time and money).

PIP Development

The purpose of each PIP is to provide common business/data models and documents enabling system developers to implement RosettaNet eBusiness interfaces. Each includes a) XML document based on Implementation Framework DTDs, specifying PIP Service, Transactions, and Messages which include dictionary Properties; b) Class and sequence diagrams in UML; c) Validation tool; and d) Implementation guide.

Dictionaries

As part of RosettaNet's foundational projects, two data dictionaries are being developed to provide a common set of properties required by PIPs. The first is a technical properties dictionary (technical specifications for all product categories), and the second is a business properties dictionary which includes catalog properties, partner properties (attributes used to describe supply chain partner companies) and business transaction properties. These

dictionaries, coupled with the RosettaNet Implementation Framework (exchange protocol), form the basis for each RosettaNet Partner Interface Processes (PIPs).

RosettaNet.org PIP Model[37]

Partner Interface Process (PIP) models identify the activities performed by each partner role involved in the PIP, identify and define the business properties required by each role interaction, define the volume and timeout requirements, security requirements, audit requirements, and exception handling requirements. All of these components make up what is known as the "PIP Blueprint." The following is an explanation of each of these components.

PIP Models

PIP models are created to document each process. PIP models are unique in that they depict the activities of two or more Partner Roles in a single process. Following is the PIP model pertaining to the Price Change Announcement PIP.

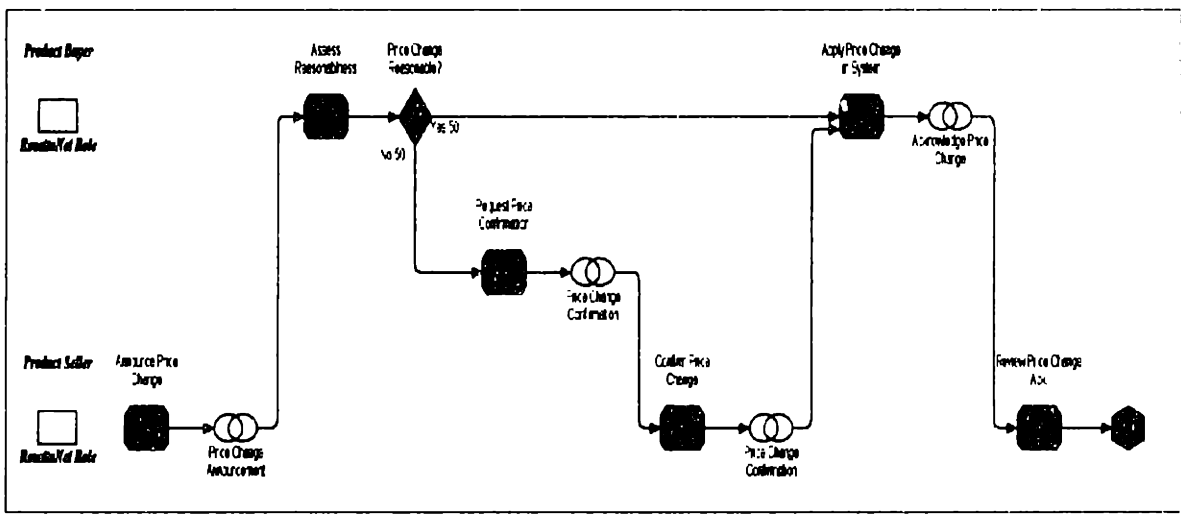

Interpreting PIP Models

The proceeding information is a legend that will aid in the review of the PIP models. The most important thing to remember when reading PIP models is that there are horizontal "swim lanes" in which the activities of a certain Partner Role are displayed. Additionally, for ease of readability, all activities have been color coded to differentiate between the Partner Role that performs the activity.

The following objects are displayed in a PIP model:

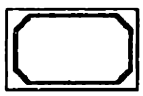 Activities. Activities are represented by rectangular objects with the corners cut off. Activities represent the most granular level of operation that a human or machine can perform. Within PIP diagrams activities are placed in horizontal "swim lanes" according to the Partner Role that performs the activity. The Partner Role is defined by the Role object in the left margin of the diagram. Therefore, all activities appearing in the "swim lanes" denoted by the Role object are performed by that Role. Additionally, activities are color coded to differentiate between the roles.

 Choices. Choices are represented by a very small octagon. Choices exist only when associated with a decision that has multiple possible outcomes. Thus, a choice represents one possible outcome of a decision.

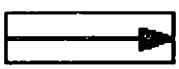 Connectors. Connectors are represented by a blue line with an arrow pointing toward the right. Connectors determine the flow.

---

[37] RosettaNet PIP Blueprints. Cluster: Inventory Managememnt. Segment: Price Protection. (31 March 1999).

 Decisions. Decisions are represented by a diamond. There are two kinds of decisions; binary decisions and multiple decisions. Binary decisions are those that have only two possible outcomes; yes and no. Multiple decisions are those that have more than two possible outcomes.

Go To. Go To objects may appear as a "star" shaped object as follows:

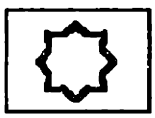 Or they may appear as triangles the either "point" forward in the model or back ward in the the model. Go Tos that are shaped as "stars" merely connect in with the same go to object that occurs either before or after the Go To object in the model. Go To's that involve re-work, which is defined as re-doing previously completed activities, are triangular and "point" in the direction of the Go To.

Numbers. Objects that appear in models with numbers following their name indicate that the object is the same object that appears elsewhere in the model. For example, assume we have an activity named "Review Claim" that is performed more than once in a model. The first time the activity appears as "Review Claim." The second time it will appear as "Review Claim 01" and so on. Each instance of the same object will be sequentially numbered to indicate it' reuse.

 Role. Roles are represented by a rectangle object. Role objects relate directly to RosettaNet Roles. Role objects are always placed in the left-hand margin of a PIP diagram as they establish ownership of the activities appearing in the horizontal "swim lanes" associated with that Role.

 Role Interactions. Role Interactions are represented by two interlocking ovals. They represent the data exchange between two Supply Chain partners. Role Interactions contain data.

Stops. Stops are fairly self-explanatory. They represent points in the model where the path within the process comes to an end. A process may have multiple paths, with some coming to an end while others continue. Ultimately, each process comes to a stopping point.

On April 15, 1999, RosettaNet announced that it had completed a successful test of its first PIP, an update to product specifications and part numbers in an online catalog shared between Microsoft and IBM. RosettaNet plans to test more than 100 additional PIPs and XML schema before disbanding in June 2000.[38]

[38] Booker, Ellis. "XML Applications Stand Up To EDI" Internetweek. http://www.internetwk.com/story/INW19990415S0007 (15 April 1999).

Conclusion

In his book, Clockspeed: Winning Industry Control in the Age of Temporary Advantage, Charles Fine describes what he calls the first and second "laws" of supply chain dynamics.[39] The first law is the law of volatility amplification, better known as the "bullwhip effect." The bullwhip effect says that the volatility of demand and inventory is amplified the further away from the end user one is in the supply chain. The second law is that clockspeeds tend to be amplified the closer one is to the end customer. How do these two "laws" relate to supply chain integration and electronic commerce, and does the convergence of these two concepts provide any particular strategic advantage to an organization?

Supply chain integration provides a framework for collaboration. When information is shared simultaneously across multiple systems in the value chain the opportunity for demand amplification is minimized. Companies can discover real value in operating from collaboratively developed forecasts and in response to true end-user demand. Additionally, by integrating systems with key suppliers and partners, like Skyway and Diamond Multimedia, organizations develop trust and gain competitive advantage based on the capabilities of the entire supply chain. This non-zero-sum game creates real value for the customer. In a marketplace driven by competition, customers, and change, supply chain integration is the key to success.

When the amplification occurs at the customer end of the supply chain, the ability to respond rapidly is critical. Fortunately, companies like eBay and Levi.com have shown how

the new business models of electronic commerce can provide organizations with the ability to react quickly to changing customer needs and to use this immediate feedback to fine tune marketing efforts. The rewards can be substantial. The MIT Press electronic catalog demonstrates that a solid understanding of strategic goals and data modeling can effectively exploit mature technology to achieve outstanding returns on investment.

Biologists study Drosophilia, fruit flies, because their rapid lifecycle allows one to reach quicker conclusions in the field of genetics. Companies like MortgageQuote.com are at the leading edge of industry change and exhibit extremely high clockspeeds. Fine suggests that these companies can provide the business researcher with the organizational equivalent of a Drosophilia species.[40] What can we learn from these fruitflies of modern industry? One need look no further than the efforts of the high tech members of RosettaNet to discover an answer.

Informational technology such as the Extensible Markup Language (XML) holds much promise. Using XML as a communication standard, and with a strong process focus, the RosettaNet organization and other leading companies are demonstrating that the key elements of a successful business model in an internet economy are process-oriented supply chain integration and customer-focused electronic commerce.

---

[39] Fine, Charles H. Clockspeed: Winning Industry Control in the Age of Temporary Advantage. Perseus Books (1998), page 89.
[40] Fine, Charles H. Clockspeed: Winning Industry Control in the Age of Temporary Advantage. Perseus Books (1998), page 5.

Appendix A: MITPress Annotated Data Model[41]

```
--
-- data-model.sql
--
-- created by philg@mit.edu on July 6, 1996
--

-- Illustra doesn't have built-in sequence generators so you
-- have to use transactions on a one-row table to generate unique
-- keys

create table id_numbers (
      reader            integer,
      reader_order      integer,
      reader_address    integer,
        ...
);

create table journals (
      issn        text not null primary key,
      title       text,
      title_prefix      text,
      -- a Unix directory name, usually sort of like the title
      -- this is where we'll expect to find static files
      dirname     text,
      -- stuff we need for presentation of journal characteristics
      frequency   text,
      ...
      -- extra stuff for the PLS Blade to catch
      search_keywords         text,
      -- prices
      individual_price  numeric(7,2),
      institutional_price     numeric(7,2),
      student_price           numeric(7,2),
      -- for journals like LEA, where subscribers to Leonardo
      -- get a price break
      special_price                 numeric(7,2),     -- usually NULL
      special_price_explanation     text,
      foreign_shipping  numeric(7,2),
      -- individual issue prices
      current_issue_price     numeric(7,2),
      individual_back_issue_price   numeric(7,2),
      institutional_back_issue_price      numeric(7,2),
      individual_back_issue_double_price  numeric(7,2),
      institutional_back_issue_double_price     numeric(7,2),
      foreign_shipping_single_issue       numeric(7,2)
);

-- build a full-text index using the fabulous PLS Blade
create index journals_pls_index on journals using pls
```

[41] http://photo.net/wtr/thebook/mitpress-data-model.txt (6 July 1996).

```
( issn, title, editor_name, editor_bio, editorial_board,
submission_guidelines, blurb );

-- a journal can be in more than one category, so we need a separate
-- table to express the relation
create table journal_category_map (
      issn        text references journals,
      category    text
);

-- let's make it fast to find out either which journals are in
-- a category OR which categories a journal is in

create index journal_category_map_by_issn
on journal_category_map using btree ( issn );

create index journal_category_map_by_category
on journal_category_map using btree ( category );

create table books (
      isbn  text not null primary key,
      -- to make sorting work, we do NOT include leading article, e.g.,
"The"
      title       text,
      subtitle    text,
      title_prefix      text, -- put an article here, with no trailing
space
      volume_number     integer,    -- can only have 1, 2, 3
      series            text,
      -- a Unix directory name
      -- this is where we'll expect to find static files
      -- we construct this with the first three chars of the
      -- author name, then the first char of the title
      -- then P or H for binding, all uppercase
      -- e.g., ABESH for Abelson & Sussman, Structure & Interp
      -- of Computer Programs, hardcover
      dirname     text,
      -- some books we don't want to present in category home pages
      -- for example, if we have a multivolume set, we don't want
      -- to see Foobar Vol 1, Foobar Vol 2, Foobar Vol 1 & Vol 2
      -- in the Engineering section because the page for
      -- Foobar Vol 1 & Vol 2 will have all three ordering options
      -- the solution is to set the PRESENT_P to false on the
      -- individual volumes but then make them have a related_to_type
      -- of 'book' and a related_to_key that is the ISBN of the
      -- multi-volume set
      present_p   boolean default 't',
      -- this is the type of the thing to which I'm related
      related_to_type   text, -- 'book','video','cd','journal'
      related_to_key    text, -- ISBN or ISSN
      out_of_print_p    boolean default 'f',
      publication_date  date,
      -- this column is computed from the publication_date column (see
below)
      new_p       boolean virtual,
      -- we want to know if a book is a paperback original
      -- in which case we present it with "new" books despite its
```

```
      -- binding_type of 'paper'
      paperback_original_p    boolean default 'f',
      blurb       text, -- HTML, main sales stuff
      description text, -- e.g., '144 pp., 78 illus, 8 in color'
      table_of_contents large_text, -- HTML
      author_name text, -- could be several people
      author_bio  text, -- HTML, describes author(s) in one paragraph
      author_web_site   text, -- url
      author_email      text,
      search_keywords text,   -- arbitrary stuff that Terry and Marney
add for PLS
      -- prices
      intro_price numeric(7,2),
      intro_price_limit_date  date, -- last day you can get the limit
price
      price       numeric(7,2),
      discount_code     text,       -- 'S', 'T', 'X' (only used by
mega-bookstore-chains)
      -- we're going to add semantics to binding_type to handle videos
      -- and software with binding_type = 'video' and 'software' and
'audiocd'
      binding_type      text
);

-- for the virtual new_p column

create function new_p(like books)
      returns boolean
      as return $1.publication_date + interval '180' day(3) >
current_date;

-- this makes it fast to find all the books related to Book X

create index books_by_related_to_key on books
      using btree(related_to_key);

-- a PLS index on a bunch of columns

create index books_pls_index on books using pls
( isbn, title, subtitle, author_name, binding_type, description,
search_keywords );


--
-- a book can be in more than one category
-- so we have a map table
--

create table book_category_map (
      isbn        text,
      category    text
);
create index book_category_map_by_isbn on book_category_map
      using btree(isbn);
create index book_category_map_by_category on book_category_map
      using btree(category);
```

```
--
-- textbooks are categorized separately as well
--

create table textbook_category_map (
      isbn        text,
      -- is this primarily a textbook?
      primary_p   boolean default 'f',
      category    text
);
create index textbook_category_map_by_isbn on textbook_category_map
      using btree(isbn);
create index textbook_category_map_by_category on textbook_category_map
      using btree(category);

-- requests for textbooks for evaluation by professors we only allow
-- each prof to request each book once in the table; they have to send
-- email explaining why they want a second copy (probably they are in
-- fact complaining that they didn't get the first copy they requested)

create table textbook_requests (
      reader_id          integer not null references readers,
      isbn               text not null references books,
      -- bunch of stuff typed in by the requesting professor
      course                   text,
      enrollment         text,
      date_commences           text,
      level              text,
      current_textbook   text,
      institution        text,
      department         text,
      shipping_address   integer not null references
reader_addresses(reader_address_id),
      request_time             timestamp(0),
      -- here's the integrity constraint that keeps prof from
      -- requesting the same book again
      primary key(reader_id,isbn)
);

-- professional, published reviews of books or journals.  We can store
-- either the complete review or just a URL from a well-known, stable
-- site.

create table reviews (
      product_key text, -- ISBN for a book, ISSN for a journal
      product_type      text, -- 'book', 'journal'
      review_date date,
      url         text, -- URL of a review (if blank, we assume to
store the text)
      review_text large_text
);


--
-- Readers
--
```

```
-- Names of people who've either bought books
-- or joined a spam list.
create table readers (
      reader_id   integer not null primary key,
      email       text, -- we use email, first_names, last_names as a
key
      reader_type text, -- probably one of 'individual', 'library',
'wholesaler', 'retailer/bookseller'
      first_names text,
      last_name   text,
      birthyear   integer,
      sex         char(1),    -- 'M' or 'F'
      day_phone   text,
      eve_phone   text
);
create index readers_index on readers using btree(reader_id);

-- this should dramatically speed looking up to see if a person
-- is already in the READERS table (except it doesn't
-- because we need to do case-insensitive search)
create index readers_email_index on readers using btree(email);

-- this functional index actually does speed the lookup, but
-- sadly it locks us into Illustra or Informix since Oracle doesn't
-- give us functional indices.
create index readers_upper_email_index on readers using
btree(upper(email));

-- we use this to figure out if a nickname or something was typed
create function substring_p(text,text) returns boolean
as
return position($1 in $2) > 0 or position($2 in $1) > 0;

-- takes email, first_names, last_name and tries to find a match
-- returns the reader_id or NULL

-- must be used like this
-- select * from lookup_reader('philg@mit.edu','Philip
G.','Greenspun');

-- this was wicked slow until we added the functional index above

create function lookup_reader(text,text,text) returns setof(integer)
as
select reader_id from readers where upper(email) = upper($1) and
substring_p(upper(first_names),upper($2)) and
substring_p(upper(last_name),upper($3));

create table reader_addresses (
      reader_address_id integer not null,
      reader_id integer not null references readers,
      address_type      text,   -- bill, ship, spam_list
      line2 text,
      line1 text,
      postal_code text,
      state text,
      city  text,
```

```
      country     character(2) references country_codes(iso_code) --
ISO country code
);

-- args are reader_id ($1), line1($2), line2($3), postal_code($4)
create function lookup_reader_address(integer,text,text,text) returns
setof(integer)
as
select reader_address_id from reader_addresses where
reader_id = $1 and line1 = $2 and ((line2 = $3) or (line2 is null and
$3 is null)) and postal_code = $4;


create table reader_comments (
      reference   text, -- 'ISBN0262123451' for a book, 'ISSN903412'
for a journal
      email       text,
      name        text,
      message           text,
      posting_time      timestamp(0)
);
create index reader_comments_index on reader_comments using
btree(reference);

--
-- we record readers who are interested in books or journals
--

-- we allow duplicate submissions because
-- we are going to do a DISTINCT for all queries
-- anyway

-- we record interest in specific books or journals though
-- we tell people that they're joining "the Anthropology list"
-- and send out email accordingly

create table spam_list (
      domain            text, -- for random lists
      issn        text, -- if a journal
      isbn        text, -- if a book
      reader_id   integer references readers
);


--
-- Orders
--
-- the business end of the system
--

-- a table like this is useful for generating user interface
-- but also to maintain database integrity and make sure that bogus
-- country codes can't be entered

create table country_codes (
      full_name   text not null,
      iso_code    char(2)
```

```
);

--
-- if multiple items are ordered, then that puts multiple rows
-- in this table, but all of those rows will have the same
-- order ID
--

-- we don't store credit card numbers in the database; we encrypt them
-- with a public key and then transmit them to the legacy system where
-- they are decrypted with a corresponding private key

create table reader_orders (
      reader_order_id   integer not null,
      reader_id   integer not null, -- who ordered it
      product_key text, -- ISBN for a book, ISSN for a journal
      -- these have the values that go into the PPSB report
      product_type      text, -- 'book', 'journal'
      subscriber_class  text, -- relevant mostly for journals, e.g.,
'individual', 'current_issue'
      renewal_p   boolean default 'f',
      account_number    text, -- relevant only for journal renewals
      order_time  timestamp(0),
      quantity    integer default 1,
      payment_method    text, -- 'creditcard', 'purchaseorder', maybe
'ecash'
      credit_card_type char(1),     -- M, V, or A
      -- holds output of a PGP run with both cc# and exp date
      credit_card_encrypted   text,
      purchase_order_number   text,
      gift_name   text, -- NULL if not a gift
      transmitted_p     boolean default 'f',    -- set to true when
order xmitted to legacy system
      fulfilled_p boolean default 'f',    -- set to true when order is
shipped
      paid_p            boolean default 'f',    -- set to true when
First Virtual is confirmed, for example
      cancelled_p boolean default 'f',    -- set to true if user
cancels subscription or returns physical product
      -- this is stuff that we really ought to be able to get from
other
      -- tables but it is safer to insert them here (since the price
      -- in the books table might change but in a report you want an
order
      -- from 1996 to have the price paid at the time
      price       decimal(7,2),
      shipping    decimal(7,2),
      tax         decimal(7,2),     -- also use this field for 20%
international tariff
      -- arbitrary user-entered comment
      comment           text
) archive;
create index reader_orders_index_0 on reader_orders using btree
(reader_order_id);
create index reader_orders_index_1 on reader_orders using btree
(reader_id);
```

```
create index reader_orders_index_2 on reader_orders using btree
(product_key);


---
--- Reporting
---

--- many of these are kludges to work around Illustra's inability to
--- GROUP BY anything other than a raw column name

create function book_0_or_1 (text) returns integer
as
return position ($1 in 'book');

create function journal_0_or_1 (text) returns integer
as
return position ($1 in 'journal');

create function video_0_or_1 (text) returns integer
as
return position ($1 in 'video');

create view orders_for_report (reader_order_id, year, month, book,
journal, video) as
select reader_order_id,
extract(year from order_time),
extract(month from order_time),
book_0_or_1(product_type),
journal_0_or_1(product_type),
video_0_or_1(product_type)
from reader_orders;

create view all_products (title, product_key, product_type)
as
select title, isbn, 'book' from books
union
select title, issn, 'journal' from journals;

create index all_products_by_product_key on all_products using btree
(product_key);
```