Human Aspects of Scheduling: a Case Study

by

Yishai Boasson

M.Eng. Logistics
Massachusetts Institute of Technology

B.A. Statistics
Hebrew University, Jerusalem

Submitted to the Engineering Systems Division in Partial Fulfillment of the Requirements for Master of Science in Engineering Systems degree

at the

Massachusetts Institute of Technology
[September 2006]

Signature of Author ....
Engineering Systems Division
January 19th, 2007

Certified by ....
Prof. John S. Carroll
Professor of Behavioral and Policy Sciences, Sloan School of Management
Thesis Supervisor

Accepted by ....
Prof. Richard de Neufville
Chair, Engineering Systems Division Education Committee
Professor of Engineering Systems

Human Aspects of Scheduling: a Case Study

by

Yishai Boasson

Submitted to the Engineering Systems Division
on January 19th 2007 in Partial Fulfillment of the
Requirements for the Degree of Master of Science in
Engineering Systems

Abstract

This work presents a look at real-life production-floor scheduling, comparing and contrasting it to both normative OR theory and Cognitive Psychology theory. Relevant literature in OR, scheduling and psychology is reviewed, and gaps in theory are pointed out, calling for observation of real-life scheduling and for modeling of the cognitive processes underlying such activities. While normative theory and cognitive psychology theory suggest certain behaviors should be observed, a case study conducted with a large manufacturing company reveals real-life scheduling to be different from behavior expected by OR as well as by cognitive psychology. Future research is suggested, which may enable better modeling of human schedulers.

Thesis Supervisor: Prof. John S. Carroll
Title: Professor of Behavioral and Policy Sciences, Sloan School of Management

Acknowledgements

I would like to thank Prof. John Carroll for his support and guidance. This has been a great learning experience. I would also like to thank Prof. Yossi Sheffi for his generous support and honest criticism. I would like to thank Prof. Chris Magee for his support and clearheaded reality checks. I would like to thank the people of FunChem for their cooperation, insightfulness and friendliness. Lastly, I would like to thank Prof. Joel Cutcher-Gershenfeld. It is a journey and nothing is more valuable on this journey than a true friend and mentor.

Dedication

This thesis is dedicated to my family who stood by my through the trials and tribulations of MIT. My deepest gratitude and admiration goes to my wife Ayelet. Her unwavering love and courage are a lighthouse in a sea of uncertainty.

Table of Contents

1 Introduction

Scheduling involves, among many other factors, temporal arrangements of jobs to be performed, under certain constraints such as limited resources, deadlines, rewards and penalties. Another way of viewing scheduling is as allocation of resources to tasks. Scheduling has been researched extensively using Operations Research (OR) tools and methods (Graves 1981). These research efforts have yielded a vast array of state-of-the-art algorithms used by Decision Support Systems. Decision Support Systems (DSS) are computer applications designed to assist decision making. Computerized Decision Support Systems have become ubiquitous as planning and scheduling tools. Sometimes referred to as Executive Information Systems (EIS), Advanced Planning Systems (APS), or a myriad other three-letter-abbreviations, these systems utilize the relatively massive data-processing power wielded by modern-day computers to make possible the timely incorporation of large amounts of data into decisions, based on optimization algorithms and heuristics.

The cost of buying a decision support system could be as high as millions of dollars, while the cost of maintaining it and feeding it with all the required information possibly amounts to additional hundreds of thousands, if not millions of dollars per year (based on an interview with a Supply Network Optimization Business Lead of a large north American company). Still, experienced users often choose to override these systems' recommendations. This phenomenon is known as "user non-compliance." Studies have shown the average rate of non-compliance to be between 30% and 40% (Powell, Towns et al. 2000), while an exploratory interview I conducted with a senior analyst from large company (in charge of specifying requirements for a company-wide resource scheduling system) indicated it to be around 10%. The reasons for this non-compliance are not well understood.

Non-compliance with the systems' recommendations is believed by the companies owning these systems to cause considerable financial losses, or the forfeiture of potential profits. This is of course due to these companies' belief that, unlike the recommendations made by the system, people's decisions are sub-optimal (with "optimality" being profit-maximization, cost-minimization, or any other pre-defined goal). It is believed that decision optimality cannot be reached by people when faced with very complex systems, as they lack the computational faculties required to solve such problems, as well as a rigorous methodology for selecting a heuristic or algorithm for solving such problems (Newell and Simon 1972; Simon 1978; Einhorn and Hogarth 1981; Lipman 1999).

In addition, as far as the decision support system is concerned, non-compliance can not be predicted, and is therefore an **uncertainty introduced by users**. In fact, the systems' inability to incorporate this non-compliance into its set of considerations renders its recommendations sub-optimal (Powell, Towns et al. 2000). Furthermore, under certain conditions, non-compliance renders the systems' recommendations so sub-optimal that it would actually be better for the users to follow a myopic[1] strategy rather than the now-rendered-suboptimal recommendation. Through simulation, Powell et al. show that post-optimization user non-compliance de-optimizes the initial solution, to the point where a user who suspects some non-compliance might occur can benefit by not considering the overall schedule suggested by the decision support system as optimal, and actually devising her own myopic schedule which would not be too inferior (and in some instances, even superior) to the initial optimized schedule.

For these reasons, companies (as indicated by the interviewed pilot company) have been spending much time and effort in an attempt to better understand and reduce non-compliance. If indeed non-compliance is a cause for losses, understanding it and attempting to curb it would be desirable.

However, it is also possible that non-compliance actually serves other purposes and it is even possible that it is saving companies from bigger losses, in which case it should be viewed in a broader context and evaluated accordingly.

As the interaction of users with DSS becomes better understood, better systems can be built. From a practical perspective, this could mean better user interfaces that promote compliance, better algorithms that anticipate non-compliance and better work procedures aimed at curbing undesired non-compliance.

Furthermore, fundamental understanding of the causes for user non-compliance from a purely scientific perspective would fill a gap in current knowledge. It is entirely possible that this knowledge, once gained, would prove human decision mechanisms to

[1] Myopic: giving extra weight to short term benefits over overall performance, considering only a subset of available data in making a decision.

be not only the cause of non-compliance, but also surprisingly simple yet extremely effective. In such a case, it might be possible to use these insights to create better (fast, efficient, frugal, etc.) algorithms, or at least reevaluate the idea that non-compliance is always just a manifestation of sub-optimality. This view is supported by a group of researchers led by Gigerenzer (Gigerenzer, Todd et al. 1999; Gigerenzer and Selten 2001), who believe human decision making techniques, while very simple, are well suited for the way information is presented in the real world, and in fact might even be superior to optimization techniques, once the costs of decision making are included.

It is this combination of practical application and fundamental science that gives research into scheduler decision making much of its value, in addition to its interdisciplinary nature, addressing a real issue on the borderline between operations research and cognitive science. The need for a better understanding of the role humans have in scheduling systems was pointed out by Graves (1985): "*To the extent that scheduling is not a totally automated activity, then we need to understand how the scheduling system should interface with a human scheduler. What is the proper role of a human vis-à-vis a model in the scheduling activity? What is the best information to present to the human, and how should it be presented?*"

A few researchers have looked at this issue of user non-compliance, in an attempt to document and better understand it. These research efforts have yielded a body of literature concerned with the differences between schedulers' behavior and decisions and the way OR approaches and algorithms would have solved problems similar to those faced by the studied schedulers. However, these research efforts were mostly focused on understanding what it is that schedulers do (using more qualitative tools to phenomenologically describe the schedulers' work) as well as on benchmarking schedulers' performance against OR generated algorithms (Sanderson 1989; Wiers 1996; Wiers 1997; Crawford, MacCarthy et al. 1999; MacCarthy and Wilson 2001; Powell, Marar et al. 2002; MacCarthy 2006).

Other researchers have done similar work in an attempt to advance the use of decision support systems of differing degrees of interactivity (Haider, Moodie et al. 1981; Higgins 1994; Higgins and Wirth 1997; Mackenzie, Pidd et al. 2006). However, although these research efforts have contributed to the understanding of the nature of the scheduler's job, very little has been done in the way of understanding how schedulers make decisions. **The cognitive aspects of schedule generation are a virtually uncharted territory**.

This study aims at gaining a better understanding of some of the reasons why users of Decision Support Systems often choose to ignore or override these systems' recommendations. More precisely, I hope to better understand the difference between the way people think about scheduling and the way OR approaches the subject. This research may thus contribute to finding approaches for improving decision support systems so that they are more useful to users. In addition, in case of a system malfunction, shutdown, or attack, people might have to take over decision-making, at least until the system functionality is restored. In such an event it might be beneficial if the computer system, once up and running again, could anticipate (at least to some degree) the state in which the engineering system (such as logistics or job-shop setting) might be, even before precise system-state data is collected. This would allow faster recovery time, as well as more efficient allocation of computational resources in attempting to restore overall system optimality. Finally, findings may contribute to better modeling of systems where human decision making plays a role. More generally, there is promise of possible scientific insights about the way people perceive time and perform scheduling tasks.

Quoting Simon (Simon, Dantzig et al. 1987): "*There are no more promising or important targets for basic scientific research than understanding how human minds, with and without the help of computers, solve problems and make decisions effectively, and improving our problem-solving and decision-making capabilities. ... A substantially enlarged program of empirical studies, involving direct observation of behavior at the level of the individual and the organization, and including both laboratory and field experiments, will be essential in sifting the wheat from the chaff in the large body of theory that now exists and in giving direction to the development of new theory.*"

2 Literature Review

The literature reviewed in the next section does not reside within one single discipline or scholarly branch. It is rather an amalgam, touching on several fields of knowledge. This is not coincidental. It is my intention to show with this literature review the progress that led up to my proposed research. I will show that the study of scheduling has stemmed from a recognized need for an analytic approach to solve a real operations need. As such, it was pursued within the Operations Research discipline, with a focus on normative/prescriptive issues. I will then compare these OR heuristics to some cognitive heuristics offered by cognitive psychology and show a high degree of similarity between the two, prompting a research question regarding whether or not schedulers actually exhibit behavior suggesting the use of such heuristics.

As more and more data accumulated about the actual application of these OR prescriptions, a descriptive approach started to emerge, which was geared towards answering emerging questions about the assumptions made by OR theory as well as the actual nature of the scheduler's work. In addition, researchers started to look at the benefits of using computer scheduling systems vs. human schedulers, as well as at possible ways to combine both to exploit their strengths. By then, a need for understanding the cognitive mechanisms behind scheduling was beginning to be recognized. I will conclude by suggesting two approaches to examining scheduling behavior in its real-life context. Reviewing those, I will show some deficiencies in their current methodology and suggest a possible way to combine them to support a holistic look at the scheduling phenomenon as well as at the environment in which it takes place.

2.1 Production Scheduling: an Operations Research (OR) Overview

Scheduling activities are performed by many businesses, industries and service providers, and are considered to be a very fundamental part of their operations. The importance of scheduling activities is evident from the fact that these activities have come to be regarded as a critical aspect of operations (Sanderson 1989). The term "scheduling" is usually used to describe the process of allocating available production resources over time to meet certain goals (Graves 1981). However, the line between such activities and other activities (higher level planning, inventory management, shipping and routing, global supply chain coordination) is often blurred, leading some researchers to concede that the term is not precisely defined. For example, Crawford et al. (1999) mention some definitions of what scheduling is (shop-floor allocation of resources to perform tasks decided on by upper management levels etc.), while pointing to the assumptions (which they challenge) that have to be made when distinguishing scheduling from planning, as well as the lack of reference to or consideration of actual schedulers' abilities and actual performance. Generally, the Operations Research approach to scheduling is mostly concerned with working under constrained inputs and resources, which distinguishes it from higher level planning.

The realm of production scheduling has been extensively researched by Operations Researchers in the past 60 years. While Gantt developed the Gantt chart to aid scheduling and compare schedules with their execution during WWI, it was in the 1940's that many modern scheduling concepts and classifications were formed. This ramp-up in operations research was perhaps mainly due to the advent of computers, allowing computationally heavy techniques to be developed and utilized. Since WWII, constant advances in algorithm development were (and still are) made, as computing power became cheaper and ubiquitous. By the 1960's and 1970's, review and survey papers (which are cited by Graves) have appeared. Some later review papers of the state of the art in scheduling algorithms are also mentioned by Wiers (1997).

Although some of the scheduling problems described in Graves' paper (like the one-stage, one processor problem) are relatively easy to solve, with the solution depending on the selected cost/performance criteria, many problems tackled by OR are either NP-hard or NP-complete. This means that the computational requirements for solving such problems explode exponentially very quickly, depending on the size and characteristics of the problem. Basically, the only way to truly solve a "hard" problem is to enumerate all possible solutions and outcomes, go over each and every one of them and select the best one. Therefore much of the OR-developed methods for solving scheduling problems have been heuristics for sub-optimally solving such problems, with a sub-optimality measure that provides some indication of the worst-case scenario performance of those heuristics. As these heuristics are very sensitive to the assumptions made about the scheduling problem, there is currently a huge number of such heuristics, each suited for a very specific problem. However, to make the problems manageable, a great many simplifying assumption are made by OR theorists, which possibly render these problems very stylized, perhaps even too far removed from real life problems (Tanaev, Gordon et al. 1994; Maimon, Khmelnitsky et al. 1998; T'kindt and Billaut 2002;

Wezel, Jorna et al. 2006). For example, most of the OR work deals with static problems, while real life environments are dynamic, with the consequences of the scheduler's decisions being fed back to affect the scheduler. This gap between scheduling theory and scheduling reality is also recognized by Graves, who calls for additional research on the reality of scheduling rather than solely on the theory.

As detailed above, the Operations Research approach to scheduling involves formulation of problems, understanding whether an optimal solution is obtainable in a finite time, and construction of a solution algorithm, which might lead to an optimal solution (in very simple problems) or which might be a "heuristic," designed to systematically drive towards an acceptable solution that is the best feasible in a finite operation, while providing some estimate about the quality of such a solution, measured against a theoretical solution that is optimal on given criteria, for example on minimizing the maximum job tardiness, the average job tardiness, the number of late jobs, etc.

However, given that the actual scheduler is faced with pretty much the same problem every time he has to build a schedule, while the system constraints and properties are well known to him, it stands to reason the scheduler himself would not follow this OR development process each and every time he has to build the schedule. Rather, one might expect the scheduler to use some of the heuristics or algorithms already well known to OR, choosing among these tools the one most suitable to accomplish a workable schedule, or switching between several of those tools as he progresses through the schedule building process.

As the purpose of this thesis was to understand what it is that the scheduler does, rather than to prescribe the "correct" way in which he should go about his work, some of the most used OR-generated scheduling algorithms were surveyed, so their use could be recognized once observed. While different researchers mention different algorithms, there are some algorithms that are the staple of scheduling (Fox and Kriebel 1967; Haider, Moodie et al. 1981; McKay, Safayeni et al. 1988; Nakamura and Salvendy 1988; Sanderson 1991; Higgins 1992; 1994; 1994; 1995; 1995; 1995; 1996; 1996; 1996; Wiers 1996; 1997; Wiers 1997; 1998; Baek, Oh et al. 1999; 1999; MacCarthy and Wilson 2001; MacCarthy 2001; 2003; 2003; McKay and Wiers 2004; MacCarthy 2006), including:

- Shortest Processing Time (SPT): the job with the shortest processing time required for completion is processed. If two jobs have the same processing time, a coin is flipped or another job attribute is considered to decide between the two.
- Longest Processing Time (LPT): the job with the longest processing time required for completion is processed. If two jobs have the same processing time, a coin is flipped or another job attribute is considered to decide between the two.
- Earliest Due Date (EDD): the job with the earliest due date is processed. If two jobs have the same processing time, a coin is flipped or another job attribute is considered to decide between the two.
- First-In-First-Out (FIFO): jobs are processed by order of arrival.
- Last-In-First-Out (LIFO): the most recent job to come in is processed first.

It should be noted that all of these OR heuristics have a common feature. In all cases, all the jobs are ordered according to a single deciding feature, for example processing time, due date, or arrival date, and are processed in that order as much as possible. There is some flexibility in more advanced versions of these heuristics, for example when the last two or three jobs to be processed can be modified to optimize the schedule.

It is this aforementioned ordering of jobs by a single attribute that provides the strongest connection with some heuristics for decision making suggested by cognitive psychologists. As psychologists look at and analyze decision making, it is often viewed as choice making. In particular, the task of constructing a schedule could be viewed as choice making, where the alternatives are the unscheduled jobs and the choice is which of them to schedule next. Specifically, the use of heuristics enables the decision maker to arrive at a final decision without considering tradeoffs. Even though this means giving up on a possibly optimal solution, it does allow a hopefully good enough decision to be made within a finite time.

One school of thought in particular has been advocating the view that real-life decision making is more about simple yet powerful and effective decision rules, rather than about optimization. Specifically, the works of Gerd Gigerenzer and the ABC group show how people use a decision strategy they call Take The Best (TTB), in which only a subset of the available information about a choice (typically just one attribute) is being used by the decision maker to pick between options (Gigerenzer, Todd et al. 1999; Gigerenzer 2000; Gigerenzer and Selten 2001). As a choice strategy in itself, TTB seems to make sense in the scheduling decision situation for two main reasons. For one, the scheduler has limited time and resources to make his schedule, so to create a working schedule, suboptimal as it may be, he is most likely to use some sort of heuristic, rather than try to optimize a computationally hard problem. In addition, it is this match between the TTB heuristic and the OR prescription that would prompt the researcher to look for instances where the scheduler uses the OR prescriptions in a mechanical manner.

This expectation that the scheduler will use TTB tactics is also supported by two other sources. In his work, Peter Higgins (Higgins 1994; Higgins 1994; Higgins 1995; Higgins 1995; Higgins and Wirth 1995; Higgins 1996; Higgins 1996; Higgins 1996; Higgins and Wirth 1997; Higgins 1998; Higgins 1999; Higgins 2003; Higgins 2003) proposes an interactive human-computer scheduling system, in which the human scheduler would have the option of designating a group of jobs-to-be-scheduled, to which the computer would apply a TTB scheduling strategy picked by the scheduler, so as to save the scheduler the rote task of applying the rule, leaving him with the task of determining which rule should be used by the computer. This would suggest an overall scheduling strategy in which "islands" of TTB can be observed, or even a scheduling strategy entirely made up of such islands.

A specific demonstration of the way people might make decision is also given by Payne, Bettmen and Johnson (1988). In their work, Payne et al. devised basic information search heuristics that yield approximately the same results as normative (OR prescribed) search methods, using a computer simulation. They then proceeded to examine whether these heuristics would in fact be similar to the search methods employed by real people faced with an information searching task. They found people were in fact using heuristics similar to the expected ones and were also able to switch between several different heuristics, based on the task characteristics. At the expectation-formation stage of the research, there was no reason to believe these principles were not generalizable to other decision making domains such as scheduling.

To summarize, reviewing both OR and cognitive psychology literature we can arrive at a hypothesis that people might use heuristics similar to those prescribed by OR in constructing schedules unaided by computer systems, possibly switching among several such heuristics. Therefore the research question emerges: can we find evidence of this expected behavior by schedulers in real life?

2.2 What Do People Really Do? How Good Are They?

The 1980's saw a proliferation of OR tools designed to assist scheduling, as well as cheap and reliable computers that penetrated almost every aspect of industry. This meant that relatively sophisticated scheduling tools could theoretically be used by anyone who would benefit from them. In addition, the OR research that had up until the 60's been centered around very simplified problems with relatively low-computing requirements, had shifted gears, modeling problems that were more similar to those encountered in real life, counting on computational power to be available to schedulers so they could utilize these new, more sophisticated algorithms (for example, integer and dynamic programming). However, during the 1980's and 1990's, it became clear that people were not using all available OR tools, or that these tools were not that well suited to the real world. It is during those two decades that more and more people started to look at possible differences between the prescriptions of OR and the reality of scheduling, in a way echoing Graves' call. These research efforts took several different forms. First, a question had to be asked: "What is real life scheduling anyway?"

2.2.1 Describing Scheduler Work

To answer that question, several field studies were conducted which yielded a description of the scheduler's work in its context (McKay, Safayeni et al. 1988; Liu, Fuld et al. 1993; Wiers 1996; McKay and Wiers 1999; McKay, Pinedo et al. 2002) as well as similar or adjacent issues, such as routing/dispatching (Roth, Malsch et al. 2001). These studies provided much valuable insight into the realities of shop-floor scheduling. According to these studies, the complexity shaping the schedulers' work originates with other parts of the organization in which the schedulers work. The authors of these studies emphasized time and again the complexity and constant shifting of the scheduling environment, as well as the wide variance in scheduling styles among different schedulers. However, due to the observational nature of such studies (coupled with the complex environment), they can offer limited insight into the effect specific factors might have on a scheduler's behavior. Noteworthy is the work of Wiers (Wiers 1996; Wiers 1997), who not only observed scheduler behavior, but also suggested a quantitative model to analyze factors affecting behavior (such as performance, action and disturbance variables).

However, these performance, action and disturbance variables are high level and should be distinguished from lower level variables which are cognitive and therefore internal to the decision making that might affect decision making and strategy selection (Payne, Bettman et al. 1988) when trying to figure out the basic mechanisms of human scheduling ability. Looking at lower level variables rather than at the more complex social interactions and environment issues also makes sense as it puts such research on par with similar work done in cognitive psychology on economic issues such as decision making under uncertainty (Simon 1979; Einhorn and Hogarth 1981; Kahneman, Slovic et al. 1982; Simon, Dantzig et al. 1987; Bettman, Johnson et al. 1990; Gigerenzer and Selten 2001; Newell, Weston et al. 2003).

Stoop and Wiers (1996) provide an overview of practical scheduling complexity (with a review of models and relations to planning), followed by improvement suggestions. After discussing the problems of measuring performance, the authors provide a case study (of successful implementation) as an example for how complex actual scheduling is, the kinds of disturbances that affect it, methods used and performance assessment. The case describes the implementation of a scheduling system and the reasons for its success, but does not analyze the scheduling task in the way Cognitive Task Analysis (Vicente 1999) would, as we will discuss later[j1].

Insights gained by these studies (in my opinion) pertain to what the scheduler's job actually is, rather than to the purely cognitive aspect of the scheduler's work. It appears that schedulers are not only responsible for generating schedules, but also for making sure these schedules are actually followed, as well as for mitigating other disturbances as much as possible. This puts schedulers at a crucial junction in the supply chain, and dramatically increases the number of factors they routinely consider. In addition, these factors often include hard-to-measure things like company politics and organizational issues, in addition to increased complexity caused by the very existence of "supply chains" or networks (such as being sensitive to weather across the globe).

The reviewed research, although not necessarily unified by a coherent research agenda, provides much insight into schedulers' work and the differences between it and the OR view of scheduling problems. However, it provides little to no understanding of the mental mechanisms used by schedulers to perform their jobs, or for that matter, used by all to perform everyday scheduling.

2.2.2 Comparing Scheduler Performance to OR Models

The case studies and direct observations are but one side of a continuum, between the descriptive and the prescriptive. When moving along this continuum towards the latter, we see works such as that by Powell (Powell, Marar et al. 2002). In his work, Powell tried to build a new truck routing system and implement it alongside the system already in use by actual schedulers (or dispatchers). His aim was to compare his system's performance with that of the dispatchers. His hopes did not materialize, as the dispatchers seemed not to enter all information available to them into Powell's system. He thus concluded the information was too intricate and complex to be effectively entered into a system. This study however serves as an example of what is on the other side of the continuum, namely work attempting to compare human performance with that of computers, with varying degrees of control, from field oriented to complete control in a laboratory setting.

The term "laboratory" is often used (in the reviewed studies) to denote a controlled attempt at observing schedulers' behavior outside their natural production-floor environment, coupled with benchmarking efforts designed to compare such behaviors under different circumstances to that of a computer system, as well as to the performance of interactive human/computer systems. This is done either because the system on which participants are asked to work can not be implemented in a natural setting to form a field experiment, or because the nature of the production floor environment shifts and changes too rapidly to be properly controlled (Sanderson and Moray 1990).

Pretty much all laboratory experiments reviewed were not done to gain a fundamental understanding of the heuristics used by schedulers, but rather as a means for benchmarking. This point should be emphasized, since laboratory experiments closely control environments where only a few parameters are being manipulated, to enable the testing of hypotheses. Such hypotheses might include theories about the underlying mechanisms that govern human scheduling abilities etc. However, in the reviewed papers, laboratory experiments were not used to test a theory, but rather to compare performance of human schedulers to normative models, with the goal of determining "who is better," not to gain a deep understanding of the way scheduling is done by humans.

Broadly viewing these papers, the laboratory experiments all focus on having schedulers use a computerized tool (decision aid) that facilitates the generation and manipulation of schedules through an interface. Naturally, the interface varies according to the time in which each of these papers was published and ranges from the very basic (Haider, Moodie et al. 1981), using only alphanumeric data displayed on a computer screen, to the intermediate (Nakamura and Salvendy 1988) text interface to convey graphic information, to completely graphic interface with different degrees of sophistication (Higgins and Wirth 1995; Baek, Oh et al. 1999). All of the aforementioned studies give subjects a scheduling task to perform, either on their own, or using a computer system. Their performance is measured on a given criterion (maximum tardiness, asset utilization, etc.) and compared to the performance of other pre-defined scheduling heuristics. In addition to this benchmarking, interviews are sometimes conducted with the test subjects (Crawford, MacCarthy et al. 1999), and their thought process is sometimes recorded through real-time or after-the-fact reporting (Nakamura and Salvendy 1988). A conclusion reached by all researchers seems to be that in relatively complex (that is in non-easy problems) or dynamic settings, human schedulers perform markedly better than fixed algorithms, although the variability in the human performance and methods used by different schedulers (or "style") is high and is not easy to explain. High variability also means that while some human schedulers perform better than the fixed algorithms, others do not. In addition, performance can be further improved by using computerized aids, which enable "look-aheads", simulations, graphic representation of data etc., prompting the researchers to advocate the use of "interactive" man/machine scheduling tools that might enable better performance than that of human or computer agents alone.

While Crawford et al. (1999) review several lab experiments and recognize the value of information and insight gained through these experimental studies, they also offer some criticism on the suitability of the methodology of experimentation as a valid tool for examining the practice of scheduling. They argue that the controlled experiments do not sufficiently convey the complexity, uncertainty and dynamic nature of the environment in which real-life schedulers do their work. They also cite the danger of recognizing variables as important in an experimental setting (for example variables that define the scheduling problem in the OR sense, such as number of machines, nature of jobs, etc.), while in the real world it is other variables that are the important ones (for example, they list motivation, environmental disturbances and interpersonal relationships). Although this claim makes sense, especially in light of the very high complexity in real scheduling environments, the argument seems flawed. If indeed the complexity, uncertainty and dynamic nature of the environment are so important in shaping scheduling behavior, then their effects should also be investigated in a laboratory setting to determine their importance as is done in the emerging field of behavioral operations management (Sterman 1989).

Another concern they raise is with the fact that laboratory studies define scheduling according to the OR view, which (as mentioned before) might not agree with what scheduling actually is in practice. However, the relevance of laboratory studies should be judged in relation to the aim of the study. While their claim might be true for studies geared towards finding out "what is industrial scheduling practice," it would not hold for studies that seek to measure an effect of a specific parameter on scheduling behavior, or "how is scheduling done," as this type of problem might be decomposed into smaller problems that can be used to experimentally compare actual performance with optimal (OR). Therefore, a quantitative (or mixed methods) experimental approach could be suitable for looking at such problems.

2.3 Can We Build Better Systems?

Some of the aforementioned studies were done as part of an effort to design or market new and better scheduling tools. Others were more exploratory, aiming to understand the realities of shop-floor scheduling. In any case, more and more data accumulated indicating a need (or an opportunity) to create better interactive scheduling tools that would seek to take advantage of both human and computer strengths, letting the two work together effectively. Most of the work can be summed up as attempts at redesigning the interface (Higgins 1994; Higgins 1994; Higgins 1996; Higgins 1996; Baek, Oh et al. 1999), or at allowing greater interactivity (Haider, Moodie et al. 1981; Higgins 1992; Higgins 1995; Higgins and Wirth 1995; Higgins 1996; Higgins and Wirth 1997). Such interactivity could mean for example letting the human scheduler manipulate schedules generated by the computer, control display of information, or choose sub-groups of tasks that can be scheduled by the computer, while manually scheduling others. These attempts probably did yield better scheduling tools, as indicated by some test results presented by the authors. In any case, the aim was to design a **tool** (cf. Powell, Marar et al. 2002), rather than designing the entire engineering system. Hence, these efforts did not entail a fresh look at the scheduling task itself, but rather just the design of better tools, based on more knowledge in human factors for example. In fact, there had been so much interest in the human factors domain, that Sanderson identified a possible need to define scheduling as its own human factors domain, since it involves dealing with discrete tasks, rather than the continuous control human factors was traditionally concerned with (Sanderson 1989)[j2].

Other researchers saw the design possibilities in a broader perspective. These researchers identified an opportunity to re-examine the design of the entire engineering system in which the scheduler operated, to include organizational, operational and work related aspects. This was justified by the availability of new interactive scheduling tools. While this connected back with the aforementioned descriptive field inquiries into the nature of the schedulers' work, it was not the same.

The fact that schedulers were very involved with the goings-on in the factory and were enforcers rather than just planners, plus the very dynamic nature of their work environment and the tight coupling of their functions with the many other functions in the enterprise (such as suppliers and customers), have led to the realization that schedulers were a part of very large and complex socio-technical systems, which we call engineering systems (Moses 2004). To better understand the schedulers' function and performance, and to enable the design of better procedures and systems, the whole system needed to be considered. This realization was compounded by the fact that the growing body of knowledge about human scheduling was also exhibiting differing (and sometimes contradictory) claims about human behavior and abilities. For example, some researchers found that performance was markedly better when people generated the initial schedule fed into an interactive scheduling system, as compared to when the initial schedule was generated by a computer (Baek, Oh et al. 1999). Other researchers found that people monitoring their own performance are both biased against declaring their own non-optimality, as well as less sensitive to it (Liu, Fuld et al. 1993). Researchers agree however that a deeper understanding of human behavior and abilities is needed. The implications of such findings are clearly on the system-level, as the decision whether to use people or computers and in what ways goes beyond the design of scheduling tools and interfaces.

Other indications that schedulers were influenced by the behavior of other people in the enterprise, as well as that workers are influenced in their decisions by what they think other workers might do, is given by Powell, Towns et al. (2000), who develop a routing and scheduling model through computer simulation and analyze its performance under different uncertainty conditions. A very interesting insight emerged that under certain levels of user non-compliance, greedy or locally optimal strategies[2][j3] significantly outperform optimization (even myopic optimization). In addition, under high levels of uncertainty (other than user non-compliance), myopic optimizing is only slightly better than greedy sub-optimal strategies.

These insights have far-reaching implications. In a complex manufacturing environment where scheduled jobs are to be performed by people and not only by machines, the scheduler can never be sure the schedule will be adhered to. If a scheduler can have a sense of the expected rates of non-compliance, that might change the scheduling strategy as well as the planning horizon. This would be true for a human scheduler as well as a computerized one. This would also explain an observation made in many field studies that schedulers do not just put schedules together, but also enforce them. As for the operators themselves and their decision whether to follow a schedule or not, if they estimate the chances of non-compliance in other parts of the process to be high, they might actually be improving overall performance by non-complying with a long-range plan and following a greedy strategy instead.

[2] Greedy Strategy: using an algorithm (or a heuristic) of always making a locally-optimal choice, with the hope of eventually getting to (or close enough to) the global optimal solution. Greedy strategies are a subset of myopic strategies, as they consider a subset of information. However, a myopic strategy also assigns different weights to different pieces of information.

An overall framework for thinking about the scheduler's role (and its possible re-design) in the enterprise can be adapted from Cognitive Work Analysis (CWA), based on understanding constraints and the way these constraints shape the way work has to be done (1999). Constraints are analyzed on different levels of abstraction (by work domain, the task to be performed, strategies that can be used, social-organizational factors and worker competencies, for example). Vicente offers a set of tools for analyzing each of these abstraction levels. These tools include an abstraction-decomposition framework for understanding the work domain structure (what influences what, what is made of what, etc.), Rasmussen's (1986) decision ladder for understanding control tasks (a graphic representation of a linear sequence of information processing tasks, including possible shortcuts and leaps that allow exploitation of previous knowledge, recognition and environment/information structure), and information flow maps and skills analysis.

Analysis is linked across levels. Insights reached on one level are used as inputs on other levels (e.g. an understanding of constraints imposed by skills is an input to the analysis of possible strategies). Moreover, concepts introduced by models used in one level are used to tie the other abstraction levels together (e.g. Rasmussen's notion of Skill Based Behavior, Rule Based Behavior and Knowledge Based Behavior connect an understanding of the work domain which is knowledge based to strategy selection which is rule based to actual physical actions performed in day-to-day operations that are skill based). It is the analysis and enumeration of these constraints that Vicente claims should serve as a guide for designing work environments, tasks, machinery etc. CWA emphasizes design based not only on current practices, but rather on what new practices should (or can) be designed and how they might evolve through the lifetime of the system. This approach is rooted in the view that potentially risky situations arise when workers are faced with a novel situation which they haven't encountered before, and do not have a correct mental model of their work environment to support correct decision making, or are otherwise limited by tools (be it physical or procedures) that were not designed to accommodate such unexpected situations. It is in such cases that workers should have the freedom to make use of their understanding of the system (which should be correct) and generate new behaviors that would accomplish the task safely and effectively. More importantly, CWA is concerned with understanding **what** needs to be done (or cannot be done), expecting an understanding of how it could be done and by whom to emerge where the constraints have narrowed down the set of possibilities. Specifically, CWA tries to avoid taking a position on the how and who issues.

While Vicente does discuss the applicability of this framework to the study of the scheduling environment, he does not elaborate on the subject. However, other researchers have used parts of this framework to look at what the scheduling task might be made of. In a widely cited paper, Sanderson (1991) uses a slightly modified decision ladder as a platform for what she deems as necessary mechanisms in performing a scheduling task. This includes production rules that allow (or generate) situation analysis, strategy selection and follow-through action based on selected strategy. These production rules are backed by an estimated "calculation time" for each of them, based on "the model human processor" (Card, Moran et al. 1986). Another example for the use of the decision ladder is Hettenbach's work (1989), which didn't look at expert performance but rather used a more classical experimental psychology design with students in a laboratory setting, again using Rasmussen's decision ladder. Both these works drew on models of human information processing abilities constructed in a laboratory setting. This addresses one of the weaknesses or misalignments in Vicente's framework, as Vicente recognizes the importance of environmental constraints on decision making, but uses data and models of decision making that could have been generated without consideration of those factors and their complexity (i.e. in the lab), instead insisting on a holistic view.

Other researchers emphasized the need to consider human scheduling abilities in their natural environment. For example, MacCarthy (2006) outlines a high-level framework including the organizational, systems and human issues in scheduling. The framework intentionally does not explicitly consider the cognitive aspects of the issue, because an understanding of the scheduler's internal representation of the scheduling system as well as the mental models used to do scheduling is currently unavailable. MacCarthy argues for the use of Naturalistic Decision Making frameworks (NDM), which take into account (or at least acknowledge) the way real world decisions are affected by context and environment. However, it is not clear whether NDM has the capacity to methodically look at system-wide issues pertaining to scheduling and draw generalizable theories from those observations (Klein 1993; Zsambok and Klein 1997; Klein 1998; Salas and Klein 2001). This is due to the need not only to observe and understand "scheduling in the wild", but also to compare these observation to normative theory across many different instances and circumstances.

2.4 Summary of Literature Review

The literature review followed the historical course of research of scheduling, from its inception as a branch of OR, to the realization that there are gaps between OR's views of scheduling and evidence from the real world, which led to qualitative descriptions of what scheduling actually is, as well as more concrete attempts at improving scheduling practice. These research efforts have yielded important insights about the makeup and complexity of scheduling issues, scheduling environment, etc. In addition, an emergent conclusion from many of the newer papers is that in today's complex manufacturing environments, it is often still up to a specific person to determine a schedule and enforce it. It is for this reason that many of those studies emphasize the importance of understanding the way people perform scheduling tasks, considering issues such as interaction with computer scheduling systems, as well as people's abilities and strategies when scheduling. This is identified as one of the major missing pieces that should be better researched if better scheduling tools that make use of human abilities are to be effectively developed, as well as an essential piece of information required to better understand the performance of schedulers in real life situations (Fox and Kriebel 1967; Graves, Abraham et al. 1985; Dunkler, Mitchell et al. 1988; McKay, Safayeni et al. 1988; Sanderson and Moray 1990; Dessouky, Moray et al. 1995; Higgins and Wirth 1995; Wiers 1996; Wiers 1997; Baek, Oh et al. 1999; Crawford, MacCarthy et al. 1999; MacCarthy and Wilson 2001; McKay and Wiers 2003; MacCarthy 2006).

Pointing out the similarity between the scheduling heuristics prescribed by OR and those documented by cognitive psychology in other decision making situations, we are left with the question of whether or not schedulers indeed use heuristics similar to those suggested by normative scheduling theory in their day to day scheduling.

3 Methods

In the next chapter, I will be describing the results of a pilot study done with a scheduler at a leading manufacturer of liquid solutions. The first part of this chapter will describe the setting in which the participating scheduler works, as well as the details of the scheduler's work. In the latter part of the chapter, a description of the data collection methodology will be provided.

3.1 The Setting

The settings in which the scheduler (who I call Bob) was observed can be described from two aspects: (1) the physical settings and the nature of the plant and business in which the scheduler works, as well as (2) the scheduler's job description and the nature of his working routine. Information was gathered through direct observation of the physical settings, explanations and information given by Bob about the plant and the company, as well as verbal protocol collection as Bob was "thinking aloud" while scheduling.

3.1.1 The Plant

The participating scheduler, Bob, works at a regional manufacturing facility for FunChem[3], a leading manufacturer of liquid chemical solutions. The manufacturing plant is located just outside Washington DC, and is responsible for satisfying demand for the entire Mid-Atlantic region. The plant has four production lines. Two of the lines (called Line I and Line II) make the same product in different package types (plastic and metal), while the third line (called Line III) makes a different product in all package sizes (plastic only). A fourth line (Line IV) is for making concentrated chemical compounds, to be shipped out and used off-facility in making the end product at the point of sale (one package type, different labels that are printed on site).

The plant operates continuously, with an on-site inventory as a buffer. Incoming raw materials include water (from city mains), which is purified as necessary in the plant, chemicals, gases, and packages. Packages arrive pre-printed from an outside vendor and therefore serve as a constraint on manufacturing, as the packages are inserted into the production line and can not be changed at a later date. Chemicals and gases do not normally serve as a constraint, as they take up relatively small amounts of storage space and are not hard to obtain.

While the plant has its own warehouse, holding both raw materials and finished goods, it also serves many regional warehouses supplying their respective regions with the finished products that are made by the plant. In addition, the plant communicates continuously with other manufacturing plants responsible for other geographic regions. This allows the different plants to ask for assistance in case of raw-material shortage, malfunction, or high demand. The finished product is shipped to regional warehouses and to some local wholesale customers who make bulk orders by truck. As the chemical solutions are water-based, truck weight limit serves as a constraint on the amount of product that can be shipped in each truck load.

---

[3] Company name cannot be revealed, for confidentiality reasons.

While FunChem has several such plants, Bob's plant performs markedly better on all metrics. Due to confidentiality reasons, the metrics and their values can not be quoted.

3.1.2 The Work of the Scheduler

Regarding the work of the scheduler himself, a distinction should be made between the scheduler's job description along with the procedures he's supposed to follow (to be described next) and the actual work performed by the scheduler, to be described later in the Observation section.

Overall, the scheduler has two main official functions: Planning and Scheduling. As this thesis focuses on the scheduling aspect of his job, it was important to ascertain that the scheduler does in fact make a conceptual distinction between the two, as well as that the two can be functionally separated for analysis purposes. In a preliminary phone interview with the scheduler, conducted about two weeks prior to the site visit, the scheduler made it clear that not only does he understand the difference between the two functions (although that difference is not universally agreed-upon, as mentioned earlier), but also that there is in fact a temporal distinction between the two jobs, with the scheduling part being done only after the planning phase was over. During his workweek, the scheduler gathers information about incoming orders (from warehouses and large bulk customers), demand forecasts from within the company, and data about the availability of raw materials, as well as about the timing and quantities of incoming shipments. In addition, any other information such as production line downtime, holidays etc. is communicated to the scheduler. The scheduler should then enter this information into his computerized planning and scheduling tool. The outcome of the system's processing of this data is an overall demand, presented in a graphic format for the scheduler to view, with markings indicating the timing of orders that require filling. Having such a complete and final display finalizes the scheduler's planning. Once a week (Wednesday afternoon), the scheduler should finalize the weekly demand and feed this demand data into the scheduling software. He should then generate a production schedule based on that demand. This one-evening activity is the scheduling part of the scheduler's job. The outcome of the scheduling process is a detailed plan of what compounds are to be manufactured and in what quantities, in what order they should be manufactured and on which production line.

After the schedule is generated and checked for feasibility, a schedule committee convenes. The committee includes the scheduler himself, the plant manager, the production manager, the QA manager, the maintenance manager and the warehouse manager, as well as the person in charge of raw materials. Once the committee reviews and approves the schedule, it is passed down to the production floor, where it is adhered to.

3.1.3 The Scheduler

The participating scheduler, Bob, was hired by FunChem fresh out of college, where he completed a program focusing on manufacturing and operations. He has been with FunChem for two years, going through several positions within the company as training for his role as a scheduler. These positions included manufacturing, shipping/warehouse, as well as procurement. He attributes his success at his job to his thorough understanding of the processes and constraints involved in each of these functions.

3.2 Methodology

To obtain as complete a picture as possible of the scheduler's work, three rounds of interviews were conducted with the scheduler. The first interview was done over the telephone. In the first interview, the research was presented to the scheduler and the benefits of insights to be gained clarified. The scheduler was asked to describe his job in general terms, in order to make sure he was indeed a good fit for this research. To be a good fit, the scheduler needed to demonstrate that he uses to a certain degree a computerized system to assist in the construction of a daily/weekly work schedule for the plant's production floor. However, these requirements were not made known to the scheduler during the interview, as it was necessary to obtain his own views of what it is that he does.

During the interview, Bob elaborated on the nature of his job, and explained that it was made of three parts: (1) generating, collecting and aggregating demand data and entering it into his system, (2) schedule generation based on data entered, and (3) administrative duties/problem solving activities. Bob also made a clear distinction between functions (1) and (2), which were separate both functionally and temporally (described using terms such as "totally different activities" and "...when I'm finished doing the data entering, I can start with building the schedule..."). More precisely, Bob explained that he was performing function (1) during the entire week, pausing on Wednesday to do the schedule construction based on the data in the system at that point. Schedule construction was done on a weekly basis for the coming week. Based on this information, it became clear that the scheduler can distinguish between the data gathering and generation and the schedule construction, meaning that scheduling was a completely separate activity which could be studied uninterrupted by other activities. At this point, a time for the second interview was established, coinciding with the time he was to do his next round of scheduling.

Since Bob described the schedule construction process as well defined in time and scope, happening every Wednesday afternoon, the second meeting was scheduled to take place during this activity. The "day-in-the-life" observation and interview took place in Bob's office at his plant outside Washington D.C. Bob was asked to go through the schedule construction process as he would normally do, while thinking aloud and describing in detail his thoughts, the information he was considering while scheduling, any heuristics he might be using, etc. This "think aloud" verbal protocol method was used in previous studies of scheduling behavior (Nakamura and Salvendy 1988) and is also advocated by other researchers (Newell and Simon 1972). While as a method it certainly does not provide a complete and absolutely reliable insight into the scheduler's decision making (Bainbridge 1979), it does provide information that is very difficult to get any other way, for example the scheduler's consideration of other people in the organization. It was made clear to Bob that not only are the scheduling heuristics of interest to this research, but also other reasons for decisions, such as organizational considerations (for example, certain people's expected behavior and performance might be taken into account while making decisions). As his thoughts about individual performance and abilities were also of interest, it was decided that there will be no electronic recording of his comments, only note taking, taking care not to identify any of the individuals he might be mentioning while talking. Key phrases that were indicative of his scheduling heuristic were noted during the interview. These phrases included (but were not limited to) "considering these options", "I'm looking for", "I need to have", "because", etc.

In addition to collecting the verbal protocol from Bob, his work on his computer was also visible the entire time, and he explained all the data available on his screen and the way he interacted with the software. This allowed the comparison of his thoughts and action descriptions with the actual work being done on his computer. As will be described later, the scheduling process was often interrupted with Bob having to take care of problems around the plant, occasionally leaving the office to do so. When this happened, I accompanied Bob and a description of these extra activities was obtained as well. The entire second interview took about seven hours, from 16:00 to 23:00.

After the second interview was finished, a list of questions for the third interview was compiled, based on a review of all the data collected and initial insights gained during this interview. The third interview took place at Bob's office the very next day and lasted about 4 hours. In this interview, he was asked to verify some of the observations and insights gained through the second interview. In addition, Bob was asked to reflect on his work, explain some of the differences in scheduling tactics between the different product lines, elaborate on scheduling methods used by his colleagues, raise concerns about the software he's using, and explain his relationships with his peers and managers, as well as the way these relationships might affect his decision making process. The interview concluded with a tour of the plant, which gave the rather abstract product-line notions developed in the second interview a concrete meaning, grounding insights in the visible scale and complexity of manufacturing.

During the second and third interviews, questions were raised both by the researcher and by Bob about the availability and use of certain features of the software, as well as about the way the software incorporates (or doesn't) some of the constraints that have to be considered during schedule construction. These questions were later discussed through email correspondence with the person in charge of the next software version at FunChem.

3.3 Limitations

While verbal protocols are used extensively in other studies of decision making in general and schedulers in the field in particular (Simon, Dantzig et al. 1987; Nakamura and Salvendy 1988; Liu, Fuld et al. 1993; Klein 1998; Crawford, MacCarthy et al. 1999), the method has some inherent potential drawbacks. A few of these potential drawbacks would be the possible unavailability (or unverbalizability) of the underlying decision making processes to the decision maker, possible incentives for inaccuracy (if the decision maker doesn't want the researcher to know some things, like personal reasons for certain decisions or ulterior motives not aligned with company goals), as well as a possible change in behavior caused by the additional reporting. However, in this case there is some evidence mitigating fears of these drawbacks. The scheduler was observed scheduling for six hours, during which he never once had to be reminded to think out loud. In addition, there wasn't a single instance where the verbal information relayed by the scheduler did not match the observed behavior. In fact, the scheduler commented that this "talking to yourself" (as he called it) was in fact a natural part of the way he normally does scheduling, used to make sure every scheduling decision he makes has a clear rationale behind it and could therefore be justified later in the scheduling committee meeting. However, he did perform the reporting twice, first while scheduling and the second time when reviewing the schedule. The second repetition was the mental justification process, as he described it, an "afterthought" to the scheduling process, rather than the driving force behind the scheduling process.

4 Observations and Discussion

In this chapter I will contrast the observed behavior with the prescribed OR behavior, as well as with some existing cognitive psychology and human factors literature, to show why the observed behavior is surprising yet effective. I will conclude by suggesting possible reasons for this observed behavior, as well as possible future research directions.

The literature (reviewed earlier) describes the scheduler has having three functions: scheduling, problem-solving and enforcing. I present my observations in that order.

4.1 Building the Schedule

After entering all the demand and available resources data into his system, Bob builds the weekly schedule. As mentioned earlier, his display includes all manufactured products and their current days of inventory. In addition, days of inventory are color coded into the future, with different colors for safe inventory, unsafe inventory and critically unsafe days of inventory. On each product timeline, orders to-be-filled are marked as a small square, without an indication of the order size. However, these orders are already incorporated into the available days-of-supply inventory he sees. To build his schedule, Bob can decide how much of each product to make on the relevant product line, constrained by the line's capacity and the availability of raw materials. While this may seem straightforward, Bob does some of the scheduling and spreading out of demand during the demand data-entering process. As Bob is aware of shipping constraints, he bundles up demand into feasible shipments (as the weight constrains possible shipping configurations), allowing him to spread out demand as well as to pre-plan production based on these constraints, using his extensive experience. Shipping constraints are not only derived from pallet size, but are also further complicated by the weight differences between the different products. Although all products are water-based with the same type of packaging, the miniscule differences in the chemical characteristics of the products affect pallet weight as well. Bob makes these adjustments himself because many such constraints are not available in his system, which will be elaborated on later.

When Bob is ready to start constructing the schedule, he first orders all the products by safe inventory level, from lowest to highest. This sorting corresponds very closely with the Earliest Due Date (EDD) heuristic. As our initial expectation was for Bob's scheduling heuristic to be similar to an EDD type heuristic, Bob was clearly not following closely the order in which the jobs were arranged. In fact, in one of the production lines, he ordered jobs by EDD on his screen, while selecting candidate jobs for scheduling based on their chemical makeup. As he did so, it became clear from his thinking aloud (and later confirmed in a follow-up session) that he was using this sorting of jobs as a purely visual aid, facilitating the location of a specific job among the rest (quote: "This is just so I know where the jobs are. Other people sort them out by flavor"). He mentioned that other schedulers prefer other sorting methods, for example by product type or simply by product name. Overall, the emerging scheduling tactic seems to be:

- Sorting of jobs by safe inventory levels
- Picking a specific job and reviewing all the constraints associated with that job
- Reviewing the daily schedule for constraints
- Shifting scheduled jobs as required in order to satisfy constraints
- Reviewing overall weekly schedule
- Shifting scheduled jobs as required

Bob described his scheduling tactic as one continuous process as he performs his scheduling. As notes were being taken, the activity was broken down into the aforementioned sub-activities as distinctly perceived by the researcher (sorting, picking, iterative reviewing and shifting, reviewing, concluding). After the scheduling was over, the scheduler was asked to break down the process into sub-activities. As it turned out, the scheduler's thoughts matched perfectly the notes taken.

4.2 Contrast with Expectations

This way of scheduling differs significantly from our expectations and is perhaps the most exciting phenomenon observed in these interviews. While we expected Bob to adopt a set of Take The Best (TBB) heuristics that would correspond to OR normative heuristics, it seems he uses a Satisficing method instead. Satisficing was described by Herbert Simon to be a form of decision making strategy where options are ordered and searched, and the first option to meet the criteria is selected (stopping rule) (Simon 1982). The search for the optimal or best choice does not continue, and the decision maker is satisfied with the "good enough" option selected. This strategy seems to fit perfectly with

Bob's, as he examines the constraints attributed to each potential job, and schedules the first job in which all the constraints work with the already existing schedule. While doing this, Bob described his thoughts and actions with phrases like "I just take this first one here... and I check that I can make it... because of what I made before... and that it's ok...". While this scheduling strategy clearly differs from the use of heuristics that are optimization oriented, it does offer a fast and frugal way of obtaining a good enough (meaning workable) schedule.

While Bob is using his scheduling software purely as a graphic "clipboard" of sorts, the software does have strong scheduling capabilities that go virtually unused by Bob. This is partly due to the fact that very few of the constraints considered by Bob are available in the system. These include constraints on shipping, precedence, production line operators' breaks, number of changeovers allowed in one day, excess inventory, capacity, etc. These constraints and business rules can be documented, formulated and entered into the system, either by Bob or by the corporate level function in charge of the scheduling application specification and tailoring. However, at the time Bob was working with his software, none of them was. While the reasons for that are not well understood at this point and warrant further investigation, some insights have been gained through communication with the person responsible for the requirements of the next scheduling software version at FunChem. It appears that FunChem considers some of the constraints to be network-optimization related (for example excess inventory) and not scheduling-related, although the scheduler currently considers them as part of the problem he has to solve. In addition, other constraints such as number of shifts and changeovers are considered by the corporate head office as issues that should be resolved before the scheduling starts by doing cost/benefit analyses, while the scheduler identifies those as things that can be discussed in real time. This indicates Bob is operating in a problem space which differs than the one corporate headquarters believes or expects him to work in.

Another issue raised by previous field studies is that of scheduling "style" and the quality of the schedules (Wiers 1997; Powell, Marar et al. 2002). It was noted in these studies that different schedulers have very different ways of scheduling. When asked about the way other schedulers in his company do their job, Bob said that other schedulers do their initial sorting of jobs to be scheduled by different parameters, such as product name, while selecting candidate jobs for scheduling based on chemical makeup, popularity, or just a predetermined order with which they are comfortable and which has yielded good results in the past. However, while other research (Wiers 1996) describes schedulers as comparing their schedules on performance (based on measured metrics) and elegance (a term that is ill defined), Bob indicated that all schedulers are considered to produce schedules of similar quality, and that schedules were never judged on their elegance. When asked to define "quality" of schedule, Bob admitted that the term was not well defined and depended on subjective impressions of people working with the schedule. Specifically, the schedulers' output was not benchmarked against one another and certainly not against an optimal OR generated schedule.

As was mentioned in the literature review, much of the work that has previously focused on the real-life scheduler was concerned with describing what it was that the scheduler was doing, rather that how things were done or why. In that respect, this work is not much different, as it describes the observed behavior of Bob, while offering some insight into the reasons behind it and only a little insight into the actual details of the cognitive mechanisms that generate his behavior, work that will be left for future endeavors. Nevertheless, these observations do lead to some thoughts about the way Bob views his schedules' sub-optimality (the "Why" question), as well as about the way in which Bob's modus operandi might serve organizational purposes.

4.3 Problem Solving

Evidence for the importance of problem solving in Bob's work was easily observed and was driven by the organizational context in which Bob operated. In addition, this was supported by Bob's experience and training, which included a rotation through different functions related to Bob's current role (production, warehouse, procurement, etc.) designed to prepare him for his role as a scheduler.

4.3.1 Observations

The reviewed case studies point out that schedulers were acting as problem solvers, constantly dealing with constraints and trying to find ways around them, as well as making sure the schedule is adhered to. These findings were corroborated by our study as well, with some additional explanations provided by the scheduler himself.

As soon as one entered Bob's room, it was obvious he was not just sitting at his computer devising the schedule. Apart from the standard telephone and computer, Bob has a fax machine, two cell phones and a two-way radio. These suggested that although he supposedly has all the required information he needs to build his schedule, he still needs to be in constant contact with various functions around the plant to support them in their work. In addition, through the entire interview process (which lasted about 10 hours), Bob would constantly leave the room every 5-10 minutes to take care of one thing or another. When asked to describe the reasons for this, Bob used the term "Problem Solving." When asked to elaborate about it, as well as about the distinction between problem solving and enforcement, Bob re-iterated that none of these unplanned activities were in fact to enforce the schedule. However, in regards to the problem solving aspect, he attributed it to his background more than to his job description. In fact, according to Bob, almost none of the problem solving is in his job description, but is rather "attracted" to him since he has diverse and deep understanding of the different aspects of the plant, including manufacturing, procurement, shipping and QA.

To give an example of that, Bob told us that the plant's second-shift manager had left two weeks earlier, and that much of his work was left for Bob. For example, as Bob's office overlooks the front gate, he would get telephone calls from the security guard regarding the comings and goings of trucks. In addition, Bob carried a flashlight on his belt, used to verify the content of trucks, as well as for verifying truck serial numbers in the dark.

4.3.2 Organizational Factors

In the second day session, Bob was asked to explain some of the behavior observed during the first day. As mentioned before, Bob has the role of a problem solver, due to his extensive training that encompasses many functions across the organization, as well as due to the nature of his job, that is, to bridge between the manufacturing, shipping and purchasing functions. As such, Bob actually sees much value in the fact that he does not get all the information required for schedule construction directly off his computer screen, but that he actually has to talk to many different people in the organization and collect the information. In addition, once the schedule has been devised, it is not transferred electronically down to the production floor, but is rather printed out and hand-delivered, giving him a chance to talk to more people about it. As Bob explained, this actually keeps him informed in the goings-on of the factory while maintaining close personal relationships with many people he interacts with. This updated information about the state of the system and the people in it becomes invaluable once he is called upon to solve a problem. These views agree with statements made in both Naturalistic Decision Making and Cognitive Work Analysis about the importance of maintaining a coherent, complete and up-to-date mental image of the system state by the experts operating the system (Klein 1993; Zsambok and Klein 1997; Klein 1998; Vicente 1999; Salas and Klein 2001).

4.4 Schedule Sub-Optimality

Bob does not produce an optimal schedule, at least not as far as the normative OR definition of optimality goes. Two main questions arise when considering this sub-optimality. Can Bob optimize? Should Bob optimize?

As far as Bob's ability to optimize goes, it is obvious that under the current conditions, where he schedules manually, he can not optimize. As detailed earlier, schedule optimization is extremely difficult computationally, clearly beyond the capabilities of a normal human being. In addition, in order for Bob to know what an optimal scheduling strategy (or heuristic) might be, he has to first define his environment and the scheduling problem in mathematical terms and then follow the procedure an experienced OR specialist would for finding a possible solution. While Bob does not have the required OR background to do so, it might not have mattered much even if he did, as real life setting are much more complex than the OR models of them, with many of the assumptions made by OR models rendering a solution only theoretically optimal, while in reality it is those subtle differences between the model and the actual setting that again cause schedule sub optimality (Pinedo and Chao 1998; Pinedo 2002; Pinedo 2005; Wezel, Jorna et al. 2006).

Some of these unrealistic assumptions can be illustrated in Bob's behavior. First, OR models often ignore the ever changing work environment and data which might cause a need to redefine the scheduling problem and consequently its solution every hour (Sanderson 1989). In Bob's case, the DSS assumes that the entire schedule can be generated by the system as one unit, without consideration for the uncertainty inherent in the later parts of the schedule. The way Bob sees the schedule is different. For Bob, the near future (the next couple of days) will probably go according to plan, while the last couple of days in the work week will probably be different than planned. Second, OR models often ignore the fact that a human scheduler can in fact react faster to surprises by making local changes in a schedule rather that re-doing the entire schedule. In case a change in plans happens (while interviewing Bob, a request for help came in from another plant that got stuck without packaging materials), Bob can modify parts of the schedule to accommodate the required change, while the software needs to go through the entire scheduling process, delivering a new schedule that is potentially entirely different than the original schedule. Third, OR models ignore the fact that a human scheduler does not have the required cognitive computational resources to analyze a scheduling problem in a normative way, handle the data in a computationally optimal fashion, or the time to do so.

However, when talking to Bob about these issues, it was not clear that normative optimality is something to be sought in Bob's work. As previously mentioned, other research on schedulers all note the very dynamic nature of the environment in which the scheduler operates. Some studies have suggested that for scheduling in a typical dynamic environment, there would be a shift in goals that need satisfying or constraints every half-hour to an hour (Sanderson 1989; Sanderson and Moray 1990). However, since Bob schedules once-weekly, the circumstances in which Bob works do not change as he schedules, necessitating a reevaluation of his scheduling heuristic. Still, Bob recognizes the fact that circumstances and requirements will change after the schedule has been constructed, requiring some changes in the schedule to be made later on during the week. In addition, Bob does not need to satisfy just one metric (e.g. minimize inventory or minimize unfulfilled orders), but rather many (often conflicting) goals, some of which are hard to quantify (e.g. inventory levels, work shifts, maintenance schedule, as well as happiness of production-floor staff, ability to help other plants out, etc). Hence, as Bob uses one tool, he recognizes its sub-optimality, but acknowledges that this scheduling heuristic works under those different circumstances, based on his own past experience.

When discussing the schedule sub-optimality, Bob claimed it to be advantageous. As Bob describes it, having a tight schedule reduces the room for play should schedule changes be required later on. Such changes might be necessitated by a change in incoming orders or unavailability of raw material, as well as by requests for assistance from other plants. Were his schedule to be optimal, such a change would require the redrafting of the entire schedule, while a sub-optimal schedule usually allows him to change just a day or two out of the entire weekly schedule and is therefore much more flexible in accommodating last minute changes in a less disruptive manner. This view actually resonates with some cutting-edge thinking about building resiliency and flexibility into operations (specifically supply-chain related operations) by giving up optimality and acknowledging the inevitable uncertainty inherent to the supply chain (Sheffi 2005).

4.5 Enforcing

Previous case studies describe the scheduler as constantly enforcing his schedule on those charged with executing it, so as to maintain its optimality. Specifically, Sanderson and Wiers (Sanderson 1989; Wiers 1996) point out the importance of this activity in the scheduler's routine, while Power, Towns et al (2000) point to the value of this activity from a pure OR perspective. However, as mentioned earlier, in this plant the participating scheduler does not engage in any enforcement activities, as the people on the production floor do not deviate from the schedule once it has been approved by the schedule committee.

5 Further Study

Although the overarching research goal is to understand the reasons why people choose to override DSS (or not even use it in the first place), a first step in the research would be to narrow the focus down to a set of research questions. In this last chapter I will briefly summarize my findings and conclusions, point to gaps in the current way scheduling is perceived, and offer possible direction for future research.

5.1 Conclusions

This case study was intended to compare and contrast the behavior of Bob the scheduler with existing theory in different domains: Operations research, Scheduling and Cognitive Psychology. The data collected corroborated some of this existing theory, especially in the field of real-life scheduling. However, Bob's behavior stood in contradiction to both the prescriptions of OR and the expectations of cognitive psychology.

5.1.1 The Role of the Scheduler

As discussed in detail in chapter 4, Bob's behavior and working environment conformed to previous descriptions of similar environments. In particular, the scheduler was observed not only as a schedule-builder, but also as a problem solver that sits in a critical junction within the organization (Sanderson 1989; Wiers 1996; Crawford, MacCarthy et al. 1999; MacCarthy 2006). In addition, the scheduling environment was observed to change rapidly and constantly, requiring the scheduler not only to adapt his schedule to these surprises in real-time, but also to try and mitigate such surprises in advance (Sanderson 1989; Sanderson and Moray 1990).

5.1.2 Operations Research and Cognitive Psychology

According to normative Operations Research scheduling theory, scheduling should be done by the use of heuristics that take into account only a subset of the available information to deliver a schedule as optimal as possible. Such scheduling heuristics can include ordering jobs by due date, processing time, etc. (Graves 1981; Pinedo 2002; Pinedo 2005). On the other hand, Cognitive Psychology theory suggests people do in fact use such heuristics in their decision making, particularly when making choices (Gigerenzer, Todd et al. 1999; Gigerenzer 2000; Gigerenzer and Selten 2001), possibly switching between such heuristics as they perform the decision making (Payne, Bettman et al. 1988).

However, it is here where the most surprising findings were, as Bob's behavior did not seem to conform to either OR's prescriptions or to Cognitive Psychology's expectations. Instead, Bob seemed to be choosing by Satisficing, as he orders jobs arbitrarily (or using a system that is irrelevant to schedule optimality), examining each job's attributes against his set of constraints, making sure the job (and consequently the entire schedule) is **good enough** (Simon 1955; Simon 1978; Simon 1982; Simon, Dantzig et al. 1987).

5.2 Gaps in Scheduler's Model

As previously detailed, several gaps in the existing theory became apparent during this study: gaps in understanding the context in which the scheduler works as well as gaps in understanding the scheduler's mental activity.

5.2.1 Understanding the Context

There seems to be a disconnect between the way OR perceives scheduling and the way scheduling is done in practice, as exhibited by the fact that schedulers often choose to override DSS generated schedules.

There could be many reasons people would choose to override the system, including "systemic/knowledge", social/organizational, and cognitive.

- Systemic/knowledge: users might be aware of data of which the system is unaware, enabling them to outperform the system. For example, people might be aware of future surprises ahead of the system (Kleinmuntz 1990). In addition, people's mistrust of the system (people think they know better than the system) due to a lack of understanding of the system's inner workings (lack of transparency) might cause people to ignore the system's recommendations. In Bob's case, there is indication of scheduling constraints that the system is unaware of, which might be part of the reason why Bob does his scheduling manually and does not trust the system to do it for him.
- Social/organizational: the system might instruct people to perform actions that would cause inconvenience to themselves, or to other people in their organization. In such cases, people might prefer having to explain their choice to override the

system to their supervisor rather than suffer the consequences of following the system's instructions. In fact, if users can "advertise" their non-compliance (and the derision duly suffered), non-compliance could be used as a currency, or as a favor to be repaid. For example, Bob sees value in talking to the people around him as an integral part of the scheduling process. In addition, Bob intentionally leaves slack in the schedule that would allow him to later choose whether to allocate production resources to requests from other plants. Bob indicates he does this not only because he sees the importance of schedule flexibility, but also because he sees social advantage in his ability to choose whether to accommodate requests, etc.

- Cognitive: the algorithms/heuristics used by DSS might be quite different than those used by people to tackle the same problems, mainly due to the limited computational ability of the human brain. As discussed in the literature review, people are known to actually perform better on certain scheduling tasks and to have innate strengths like pattern recognition. Hence, even while the same data is available to the DSS, people might come up with different solutions than those of the DSS. In such cases, people might think the system is wrong and would act out their own solutions, disregarding the system. In Bob's case, he did not disregard the system's recommendations, but rather the system's ability to make such recommendations, choosing to do the scheduling manually rather than to hand the scheduling over to the system. Bob's observed scheduling routine was based on Satisficing constraints rather than on seeking the best schedule possible.

5.2.2 Cognitive Psychology

When considering the way an actual scheduler performs his job, existing models fall short in that they do not model the scheduler's mental activity and the way the scheduler processes information, as exhibited by Sanderson's Model Human Scheduler (MHS), which assumes the scheduler selects between normative heuristics rather than rely on empirical work (Sanderson 1991). In addition, some of the tools possibly suited for examining the work of the scheduler in a real-life environment (like Cognitive Work Analysis) also build on existing knowledge in cognitive psychology rather than examine these theories first in a scheduling context and adapt/modify them accordingly, as they might not be accurate in describing the scheduling phenomenon (Vicente 1999).

5.3 Proposed Research

My own proposed future research will focus on the third bullet in section 5.2.1: the cognitive aspects of scheduling, as such an understanding is necessary if we are to predict the way schedulers might interact with DSS. By using the term *cognitive*, I refer to mental activities of individuals, as expressed through their behavior. In particular, these mental activities can be understood in terms of information processing. This view of the mind as an information processor has since the 1950's become the mainstream of modern day psychology.

5.3.1 Outline of Methods

I see scheduling, or at least the part of scheduling that interests me, as an activity that happens inside the schedulers' head. As such, and being familiar with literature on Judgment and Decision Making (JDM), I first thought of trying to gain insight into it by controlled tasks in a laboratory setting, much as was done by others looking at other JDM issues. However, one thing became very clear from the literature on human scheduling: the environment in which schedulers work is very complex. This complexity seems to be a very important factor in shaping scheduler's behavior, as exhibited by their actions which go beyond just making schedules and into controlling and modifying their environments, as well as a high degree of variability in scheduler behavior. To me it seems that these insights call for two things: looking at schedulers without eliminating the complexity that characterizes real-life scheduling, and considering the way the environment (and information) in which schedulers work is structured, as it is shaping their behavior. It is the understanding of real life behavior and the way it is connected to the structure and constraints of the environment and the information it contains that will allow a true understanding and more complete and useful description of scheduler behavior. In addition, it is the understanding of the relationships between the two and the way they might shape each other that could give us insight into the underlying mental processes and mechanisms that allow such complex behavior to be effective (Cosmides and Tooby 1994).

The case study presented in this work may serve as a first phase of a future research project. This first phase of the study focused on an exploratory study of one scheduler in a real-life environment, comparing findings with literature. Here, a Naturalistic Decision Making style approach was used in eliciting from the scheduler his own thoughts about the heuristics he uses and the very nature of his work, recognizing that the scheduler makes decisions in the real context of his factory, rather than just constructing a schedule. From this exploratory study, one could branch out and construct a theory of the mental processes involved in schedule construction. Such a theory might include the way the jobs to be scheduled are perceived by the scheduler and made into mental symbols which he manipulates, as well as a description of the rules and heuristics he uses to manipulate those symbols. These insights could then be taken into a more controlled setting and tested out. However, this testing of hypotheses will involve the actual expert schedulers rather than students, as their capabilities and ways of operating might be fundamentally different. Therefore, the participating schedulers will either be brought into the lab, or a method will be devised of controlling their work environment for a certain period of time, to the extent possible (which might be difficult). It is my hope that the insights gained through those different methodologies (field observation and laboratory experiments), coupled with cognitive psychology theory will be coherent enough to support each other and to compensate for each of their deficiencies.

Vicente's Cognitive Work Analysis (CWA) can be a great lens through which to look at the environment, be it the production floor, the supply chain, or the scheduling aids that present the scheduler with information about the former. Still, we should note that CWA generates understanding of a specific system under investigation, and can therefore be used to put specific findings from a specific system in context, rather than to generate easily generalizable theories applicable wherever people schedule. As understanding the context in which scheduling is performed and referring to it in an organized manner is crucial to understanding scheduling, such a set of frameworks is invaluable, especially since these tools were designed with the aim of supporting the design of new systems. However, when it comes to describing the human element in this environment, I believe CWA falls short. It is not enough to just use CWA down to the level of the individual's behavior and stop there. Individual decision making capabilities are fed-back into the environment, not only providing processes with input, but also actively changing it and therefore should be understood in order for the CWA model of the environment to provide a reliable dynamic description of the environment. For example, once people start working in a system, they might change their behavior or change the system, a fact that is not inherently addressed by CWA, which rather assumes a system is designed and then given to people who work in it as expected without modifying it. It is there that I see a misalignment in the CWA framework, as it recognizes the importance of environmental constraints on decision making, but uses data and models of decision making that could have been generated without consideration of those factors and their complexity (i.e. in the lab), instead of insisting on using models and data that were obtained with the holistic view in mind. Specifically, some CWA literature uses existing cognitive psychology, without making sure it really describes how people behave in the specific system that is under investigation. It should therefore be augmented by insights into scheduling abilities and behaviors gained through the work of holistic researchers, as can be found in the Naturalistic Decision Making (NDM) community.

The way NDM sees scheduling is very much in line with claims of people who have researched real-life scheduling in the past 20 years. These researchers show real-life scheduling to be much more complex and dynamic than the prescriptive OR descriptions, and consequently much harder to analyze in detail or to provide theories of underlying mental models and cognitive mechanisms. Some researchers have even used some of the techniques of NDM, like real-time verbal protocol (Nakamura and Salvendy 1988), albeit in a lab setting. It seems to me to make a lot of sense to at least consider how schedulers perceive their own decision making process, rather than just compare their performance against a (perhaps arbitrary and irrelevant) normative model, labeling deviations as "errors" on the part of the schedulers.

Much of the NDM literature is rooted in the desire to show why mainstream psychology is inadequate for describing real-life decision making. As this "attempt-at-disproof" is the driving force, anecdotal or case-study evidence serves that purpose well and is therefore widespread in the NDM literature. However, I'm interested in gaining insight into a human ability. As such, I should be able to generalize findings onto populations outside the one studied. Case studies and anecdotes should be augmented by an understanding of the context in which the behavior is observed (and later by a model, a prediction and an experiment to test that prediction). This understanding of context, especially if different contexts are encountered, has to be supported by an organized way of analyzing this context, which is not a current staple of NDM as I have encountered it. While it is possible (and indeed very likely) that each of the members in the NDM community has his or her own views of what such a framework might look like, the entire NDM school of thought has not yet endorsed or pushed for the use of such a unified framework (and perhaps it shouldn't, as it does deal with many very different circumstances). It is in this insistence on understanding and describing the environment that CWA-like approaches come back into the picture, connecting work previously done on understanding the environment in which schedulers work with future attempts at understanding what goes on inside the schedulers' minds[j4].

5.3.2 Outline of Possible Contributions

It is my hope and expectation that as more data about real-life scheduling is collected, a coherent picture of the scheduler's work will emerge. It is this coherent image that should be the basis for constructing an information-processing model of the mind that will explain the observed behavior and predict it. Such a model and its predictions can be useful in a number of ways, including an effect it might have on OR theory and models, as well as on models of the human scheduler.

For OR theory, the ability to predict where differences between OR generated schedules and human generated ones are the greatest might enable DSS to expect non-compliance in particular instances, incorporating such non-compliance into its schedule building algorithm (Powell, Towns et al. 2000). In addition, if people (unassisted by DSS) perform scheduling tasks better than expected, it might be beneficial to incorporate some of the heuristics people use into DSS, perhaps achieving schedules that are easier for schedulers to understand and manipulate, as well as schedules that are inherently more flexible and adaptable to change, relative to more optimal and "slack-free" schedules. Of course, a more complete model of the mental processes involved in scheduling can plugged into existing models of human schedulers, like that of Sanderson (Sanderson 1991). This would provide these models with greater predicting power and accuracy.

As for Cognitive Psychology theory, future work might corroborate what this case study has shown: sometimes even when normative theory aligns with cognitive psychology to suggest "Taking The Best" as the heuristic of choice, it might be the case that people work differently, with clear benefits to Satisficing in the fast and frugal flexible schedules it generates.

Bibliography

Baek, D. H., S. Y. Oh, et al. (1999). "A visualized human-computer interactive approach to job shop scheduling." International Journal of Computer Integrated Manufacturing **12**(1): 75-83.

Bainbridge, L. (1979). "Verbal reports as evidence of the process operator's knowledge." International Journal of Man-Machine Studies **11**: 411-436.

Bettman, J. R., E. J. Johnson, et al. (1990). "A componential analysis of cognitive effort in choice." Organizational Behavior and Human Decision Processes **45**(1): 111-139.

Card, S. K., T. P. Moran, et al. (1986). The Model Human Processor: An Engineering Model of Human Performance. Handbook of Perception and Human Performance. New York, NY, John Wiley and Sons: Chapter 45, 1-35.

Cosmides, L. and J. Tooby (1994). "Better Than Rational - Evolutionary Psychology and the Invisible Hand." American Economic Review **84**(2): 327-332.

Crawford, S., B. L. MacCarthy, et al. (1999). "Investigating the Work of Industrial Schedulers through Field Study." Cognition, Technology & Work **1**(2): 63-77.

Dessouky, M. I., N. Moray, et al. (1995). "Taxonomy of Scheduling Systems as a Basis for the Study of Strategic Behavior." Human Factors **37**(3): 443-472.

Dunkler, O., C. M. Mitchell, et al. (1988). "The effectiveness of supervisory control strategies in scheduling flexible manufacturing systems." Systems, Man and Cybernetics, IEEE Transactions on **18**(2): 223-237.

Einhorn, H. J. and R. M. Hogarth (1981). "Behavioral Decision Theory: Processes of Judgement and Choice." Annual Review of Psychology **32**(1): 53-88.

Fox, P. D. and C. H. Kriebel (1967). "An empirical study of scheduling decision behavior." Journal of Industrial Engineering **18**: 354-360.

Gigerenzer, G. (2000). Adaptive thinking : rationality in the real world. New York, Oxford University Press.

Gigerenzer, G. and R. Selten (2001). Bounded rationality : the adaptive toolbox. Cambridge, Mass., MIT Press.

Gigerenzer, G., P. M. Todd, et al. (1999). Simple heuristics that make us smart. New York, Oxford University Press.

Graves, S. C. (1981). "A Review of Production Scheduling." Operations Research **29**(4, Operations Management): 646-675.

Graves, S. C., C. Abraham, et al. (1985). A Research Agenda for Models to Plan and Schedule Manufacturing Systems. Cambridge, M.I.T.

Haider, S. W., C. L. Moodie, et al. (1981). "An investigation of the advantages of using a man-computer interactive scheduling for job shops." International Journal of Production Research **19**(4): 381.

Hettenbach, D. A., C. M. Mitchell, et al. (1989). Decision making in supervisory control of a flexible manufacturing system.

Higgins, P. G. (1992). Human-Computer Production Scheduling: Contribution to the Hybrid Automation Paradigm. Ergonomics Of Hybrid Automated Systems - III. P. Broedner and W. Karwowski, Elsevier: 211-216.

Higgins, P. G. (1994). A Graphical Display To Support Human-Computer Decision-Making In Production Scheduling. Advances in Agile Manufacturing: Integrating Technology, Organization and People. P. Kidd and W. Karwowski, IOS Press: 317-320.
Higgins, P. G. (1994). Graphical Features for Aiding Decision-Making in Production Scheduling. OZCHI94, Melbourne.
Higgins, P. G. (1995). Interaction in Hybrid Intelligent Production Scheduling. Proceedings of the Melbourne Intelligent Decision Support Workshop, March 20, 1995. F. Burstein, P. O'Donnell and A. Gilbert, Monash University: 55-60.
Higgins, P. G. (1995). A post-hoc contribution to the panel discussion: Finding a lucid description for IDSS. Proceedings of the Melbourne Intelligent Decision Support Workshop, March 20, 1995. F. Burstein, P. O'Donnell and A. Gilbert, Monash University: 103-105.
Higgins, P. G. (1996). "Interaction in hybrid intelligent scheduling." International Journal of Human Factors in Manufacturing **6**(3): 185-203.
Higgins, P. G. (1996). The use of Graphics to Display Messages in a Intelligent Decision Support System. International Conference on Human Aspects of Advanced Manufacturing: Agility and Hybrid Automation, Maui, Hawaii.
Higgins, P. G. (1996). Using Graphics to Display Messages in an Intelligent Decision Support System. Second Melbourne Workshop on Intelligent Decision Support Systems, IDS'96, Monash University, Monash University.
Higgins, P. G. (1998). Extending Cognitive Work Analysis to Manufacturing Scheduling. Australian Computer Human Interaction Conference, OzCHI'98, Adelaide, IEEE.
Higgins, P. G. (1999). Production scheduling: some issues relating to the location of the user in the decision-making architecture. Conference of the Computer Human Interaction Special Group of the Ergonomics Society of Australia OzChi'99, Wagga Wagga, Charles Sturt University.
Higgins, P. G. (2003). A Formal Method for Analysing Field Data and Setting the Design Requirements for Scheduling Tools. HCI International 2003, Heraklion, Crete, Lawrence Erlbaum Associates.
Higgins, P. G. (2003). A human-interaction method for analysing design requirements for custom-built decision support tools for production control. The International Conference on Manufacturing Excellence, Melbourne.
Higgins, P. G. and A. Wirth (1995). Interactive Job-Shop Scheduling: How to Combine Operations Research Heuristics with Human Abilities. 6th International Conference on Manufacturing. Melbourne.
Higgins, P. G. and A. Wirth (1997). Interactive Scheduling. APORS'97, Melbourne.
Kahneman, D., P. Slovic, et al. (1982). Judgment under uncertainty : heuristics and biases. Cambridge ; New York, Cambridge University Press.
Klein, G. A. (1993). Decision making in action : models and methods. Norwood, N.J., Ablex Pub.
Klein, G. A. (1998). Sources of power : how people make decisions. Cambridge, Mass., MIT Press.
Kleinmuntz, B. (1990). "Why we still use our heads instead of formulas: toward an integrative approach." Psychol Bull **107**(3): 296-310.

Lipman, B. L. (1999). "Decision theory without logical omniscience: Toward an axiomatic framework for bounded rationality." Review of Economic Studies **66**(2): 339-361.

Liu, Y., R. Fuld, et al. (1993). "Monitoring behavior in manual and automated scheduling systems." International Journal of Man-Machine Studies **39**(6): 1015-1029.

MacCarthy, B. and J. Wilson (2001). Human performance in planning and scheduling. New York, Taylor & Francis.

MacCarthy, B. L. (2006). Organizational, Systems and Human Issues in Production Planning, Scheduling and Control. Handbook of Production Scheduling. J. W. Herrmann, Springer.

MacCarthy, B. L., J. R. Wilson S. Crawford (2001). "Human performance in industrial scheduling: A framework for understanding." Human Factors and Ergonomics in Manufacturing **11**(4): 299-320.

Mackenzie, A., M. Pidd, et al. (2006). "Wisdom, decision support and paradigms of decision making." European Journal of Operational Research **170**(1): 156-171.

Maimon, O. Z., E. Khmelnitsky, et al. (1998). Optimal flow control in manufacturing systems : production planning and scheduling. Dordrecht ; Boston, Kluwer Academic Publishers.

McKay, K. N., M. Pinedo, et al. (2002). "Practice-focused research issues for scheduling systems." Production and Operations Management **11**(2): 249.

McKay, K. N., F. R. Safayeni, et al. (1988). "Job-Shop Scheduling Theory: What Is Relevant?" Interfaces **18**(4): 84-90.

McKay, K. N. and V. C. S. Wiers (1999). "Unifying the theory and practice of production scheduling." Journal of Manufacturing Systems **18**(4): 241.

McKay, K. N. and V. C. S. Wiers (2003). "Integrated decision support for planning, scheduling, and dispatching tasks in a focused factory." Computers in Industry **50**(1): 5-14.

McKay, K. N. and V. C. S. Wiers (2004). Practical production control : a survival guide for planners and schedulers. Boca Raton, FL, J. Ross Pub.

Moses, J. (2004). Foundational Issues in Engineering Systems: a Framing Paper. Engineering Systems Monograph. Cambridge, MIT.

Nakamura, N. and G. Salvendy (1988). "An experimental study of human decision-making in computer-based scheduling of flexible manufacturing system." International Journal of Production Research **26**(4): 567.

Newell, A. and H. A. Simon (1972). Human problem solving. Englewood Cliffs, N.J.,, Prentice-Hall.

Newell, B. R., N. J. Weston, et al. (2003). "Empirical tests of a fast-and-frugal heuristic: Not everyone "takes-the-best"." Organizational Behavior and Human Decision Processes **91**(1): 82-96.

Payne, J. W., J. R. Bettman, et al. (1988). "Adaptive Strategy Selection in Decision Making." Journal of Experimental Psychology: Learning, Memory, and Cognition **14**(3): 534-552.

Pinedo, M. (2002). Scheduling : theory, algorithms, and systems. Upper Saddle, N.J., Prentice Hall.

Pinedo, M. (2005). Planning and scheduling in manufacturing and services. New York, NY, Springer.

Pinedo, M. and X. Chao (1998). Operations scheduling with applications in manufacturing and services. Boston, Mass., Irwin/McGraw-Hill.

Powell, W. B., A. Marar, et al. (2002). "Implementing real-time optimization models: A case application from the motor carrier industry." Operations Research **50**(4): 571-581.

Powell, W. B., M. T. Towns, et al. (2000). "On the value of optimal myopic solutions for dynamic routing and scheduling problems in the presence of user noncompliance." Transportation Science **34**(1): 67-85.

Rasmussen, J. (1986). Information processing and human-machine interaction : an approach to cognitive engineering. New York, North-Holland.

Roth, E. M., N. Malsch, et al. (2001). Understanding How Train Dispatchers Manage and Control Trains: Results of a Cognitive Task Analysis, Research and Special Programs Administration, John A. Volpe National Transportation Center, U.S. Department of Transportation.

Salas, E. and G. A. Klein (2001). Linking expertise and naturalistic decision making. Mahwah, N.J., Lawrence Erlbaum Associates, Publishers.

Sanderson, P. M. (1989). "The human planning and scheduling role in advanced manufacturing systems: an emerging human factors domain " Human Factors **31**(6): 635-666.

Sanderson, P. M. (1991). "Towards the Model Human Scheduler." International Journal of Human Factors in Manufacturing **1**(3): 195-219.

Sanderson, P. M. and N. Moray (1990). The human factors of scheduling behaviour. Ergonomics of hybrid automated systems II. W. Karwowski and M. Rahimi. New York, Elsevier Science Publishers: 399-406.

Sheffi, Y. (2005). The resilient enterprise : overcoming vulnerability for competitive advantage. Cambridge, Mass., MIT Press.

Simon, H. A. (1955). "A Behavioral Model of Rational Choice." The Quarterly Journal of Economics **69**(1): 99-118.

Simon, H. A. (1978). "How to Decide What to Do." Bell Journal of Economics **9**(2): 494-507.

Simon, H. A. (1979). "Rational Decision-Making in Business Organizations." American Economic Review **69**(4): 493-513.

Simon, H. A. (1982). Models of bounded rationality. Cambridge, Mass., MIT Press.

Simon, H. A., G. B. Dantzig, et al. (1987). "Decision-Making and Problem-Solving." Interfaces **17**(5): 11-31.

Sterman, J. D. (1989). "Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment." Management Science **35**(3): 321-339.

Stoop, P. P. M. and V. C. S. Wiers (1996). "The complexity of scheduling in practice." International Journal of Operations & Production Management **16**(10): 37-53.

T'kindt, V. and J.-C. Billaut (2002). Multicriteria scheduling : theory, models and algorithms. Berlin ; New York, Springer.

Tanaev, V. i. a. S., V. S. Gordon, et al. (1994). Scheduling theory. Single-stage systems. Dordrecht ; Boston, Kluwer Academic Publishers.

Vicente, K. J. (1999). Cognitive work analysis : toward safe, productive, and healthy computer-based work. Mahwah, N.J., Lawrence Erlbaum Associates.

Wezel, W. v., R. Jorna, et al. (2006). Planning in intelligent systems : aspects, motivations, and methods. Hoboken, N.J., John Wiley & Sons.

Wiers, V. C. S. (1996). "A quantitative field study of the decision behaviour of four shopfloor schedulers." Production Planning & Control **7**(4): 383.

Wiers, V. C. S. (1997). Human-computer interaction in production scheduling. Eindhoven, Technische Universiteit Eindhoven. **Ph.D.**

Wiers, V. C. S. (1997). "A review of the applicability of OR and AI scheduling techniques in practice." Omega **25**(2): 145-153.

Zsambok, C. E. and G. A. Klein (1997). Naturalistic decision making. Mahwah, N.J., L. Erlbaum Associates.