# Chapter 1

# THE PATH OF SERENDIPITY

This book is about people who operate, maintain, design, research, and manage complex systems, ranging from air traffic control systems, process control plants and manufacturing facilities to industrial enterprises, government agencies and universities. The focus is on the nature of the work these types of people perform, as well as the human abilities and limitations that usually enable and sometimes hinder their work. In particular, this book addresses how to best enhance abilities and overcome limitations, as well as foster acceptance of the means to these ends.

This book is also about serendipity and how unforeseen connections and distinctions enable innovative approaches to problems as well as solution concepts. The serendipitous connections and distinctions discussed here stem from perspectives that cut across domains and disciplines, and sometimes cultures. This enables, for example, using understanding of the operations of library networks to develop concepts for maintaining aircraft engines. The key, by the way, is the networked nature of relationships among system elements in both domains.

The possibility of serendipity can be enhanced by a willingness to ask questions such as, "How is a NASCAR pit crew like a Waffle House cooking and serving team?" Pursuit of answers to these types of questions is facilitated by adopting a transdisciplinary perspective that borrows from behavioral and social science, management, engineering, computing and so on to create cross-cutting frameworks for understanding questions and formulating answers.

A wealth of ideas about people and organizations are discussed in this book. These ideas appear, at least in retrospect, to follow a fairly orderly path from individuals to teams to organizations to enterprises. Yet, as illustrated by the many stories and vignettes in the chapters that follow, the path is only smooth when looking backward. Looking forward, the ideas emerged from serendipitous insights and matured independently at first, only coming together as a coherent whole more recently.

## EARLY SERENDIPITY

The picture in Figure 1 symbolizes this process. While the circle in the center of this picture may appear to be a UFO landing site, it is actually a driver-training track. When I was 14 years old, I eagerly looked forward to obtaining my driver training permit when I turned 15. In preparation, I bought a 1952 Plymouth for $35. I wanted to teach myself the basics of driving before I hit that all-important 15th birthday.

There was a field behind our house, a bit of pasture owned by Mrs. Cook, who lived across the street. I asked her if I could drive my car in this field. She agreed and, over the next few months, the track shown in the picture emerged. I achieved my goal of being ready for my permit, although my experience was limited to first and second gears, as I could never get going fast enough to get into third.

**Figure 1.** One Path of Serendipity

While all this was happening, an aerial survey was being taken of our town (Portsmouth, RI) in preparation for a new highway. The survey captured my track and my aunt, Nancy Lantz, a newspaper reporter, got me a copy of the picture. Only then did I fully realize the pattern I had created. Prior to that recognition, I simply viewed the field as the place I practiced driving.

The pattern of the track helped me to understand the significance my first car played in my earliest research ventures – this picture now hangs prominently in my office to remind me. For instance, I was able to figure out how to drive downtown without going on any public roads, since 14 year olds weren't supposed to be driving. It was a very bumpy route through fields and along abandoned streetcar tracks.

All the bumping resulted in many needed repairs. I accomplished some of these repairs using parts from my Erector Set because I could not afford to buy spare parts. I learned a lot from that 1952 Plymouth because it kept on breaking. Within six months, a 1949 Chevrolet, also purchased for $35, replaced the Plymouth, and my lessons continued. This was serendipity at its best, and certainly was one of the reasons I later pursued engineering. More recently, these experiences have contributed to the great fun we are having studying the automobile industry, as discussed in Chapter 11.

## ROLE OF SERENDIPITY

This research discussed in the following chapters clearly illustrates the substantial role serendipity has played in this work. This was not fully apparent to me until I began to study strategy and planning more formally in the 1980s, continuing into the 1990s. In *Best Laid Plans* (1994), and later *Don't Jump to Solutions* (1998), I concluded that planning is a process of placing yourself in the path of serendipity.

The basis of this conclusion was my own experiences as well as numerous discussions with very successful people and organizations. Success seldom emerges exactly as planned. I have often heard that technologies were most successful in ways other than expected, often in markets that were not anticipated. The stories of Post-Its and Super Glue are familiar examples. There are countless stories such as these (Burke, 1996; Evans, 2004).

This conclusion has frequently been countered with questions such as, "Why bother to plan at all if it's all serendipity?" The answer is that plans are needed to be able to recognize serendipity – to assure "smart" luck rather than having to accept "dumb" luck. Prepared with a clear set of intentions and at least notional plans, you'll recognize serendipity when it happens. As a result, you'll pursue your plans and achieve success in ways not previously imagined.

The people I have talked with about this phenomenon were usually pursuing clear goals with great energy when it struck them that there was a better way of achieving success, typically with slightly altered goals and refined plans. Perhaps the unfortunate fact that the glue will not permanently stick is actually a benefit – hence, Post-Its. Perhaps the inopportune fact that the gunk happens to create an almost unbreakable bond is actually a benefit – hence, Super Glue.

This book is focused much more on research to understand people and organizations than it is about product innovations, although see Chapters 6, 8, and 9. Serendipity in research involves recognizing new phenomena, seeing old phenomena in new ways, and identifying new concepts, principles, methods, and tools for addressing phenomena. In essence, serendipity in research involves the unexpected emergence of new ways of thinking about something.

There are many sources of serendipity and ways to place yourself in the path of serendipity. One way is to work across domains. My domain experiences range from the aerospace and electronics industries, to consumer products and retailing companies, to government agencies and non-profit organizations. The contrasts across this range of domains have yielded many serendipitous insights. For example, the heterogeneity of the constituencies in a domain – the stakeholders -- seems to strongly impact decision-making processes and how these processes can best be supported.

Another way to foster serendipity is working across disciplines. By disciplines, I mean engineering, computing, management, behavioral and social sciences, law, medicine, architecture, humanities, arts, and so on. Different disciplines often bring very different perspectives to problem solving and decision making, both in terms of formulation and solution. Understanding these differences can lead to new, transdisciplinary formulations. For instance, our ongoing research on modeling work and workflow has benefited from the contrasting approaches of management, engineering, computing, and architecture, and the ways they address business processes and flows of physical objects, information, and people, respectively.

A third way to foster serendipity is to cross cultures. The research discussed in this book has been pursued in the contexts of more than 30 countries, usually with colleagues from these countries. The contrasts between Western and Eastern cultures, as well as the contrasts between developed and developing countries, have provided many useful insights. For example, it seems that the role of the family in a culture affects entrepreneurship in particular and risk taking in general. Thus, in one initiative, we found that simply providing venture capital was not sufficient to overcome risk averse attitudes – support was needed to help would-be entrepreneurs to better understand and manage risks.

Most of the serendipitous insights discussed in this book emerged from crossing domains (e.g., libraries and electronics), disciplines (e.g., technology and arts), and cultures (e.g., developed and developing countries). These "border crossings" placed us in the path of serendipity and new insights were found on this path. An overarching theme of this book is the value of such border crossings.

## HUMAN-CENTERED DESIGN

The primary theme of this book is human-centered design (Rouse, 1991, 2001). Human-centered design is a process of assuring that the concerns, values, and perceptions of all stakeholders in a design effort are considered and balanced. Human-centered design can be contrasted with user-centered design (Billings, 1996; Booher, 2003; Card, Moran & Newell, 1983; Norman & Draper, 1986). The user is a very important stakeholder in design, often the primary stakeholder. However, the success of a product or service is usually strongly influenced by other players in the process of design, development, fielding, and ongoing use of products and services. Human-centered design is concerned with the full range of stakeholders.

This broad view of human-centered design emerged for me in the late 1980s as our research on intelligent interfaces for aircraft pilots (see Chapter 4) matured and we pursued plans for integrating this concept into production aircraft. Our intelligent interface concept was clearly pilot centered. However, we soon discovered that "pilots may fly 'em, but they don't build 'em or buy 'em." We needed to pay much more attention to the aircraft manufacturers, the airlines, and various regulatory bodies if we wanted to get our idea on airplanes. This led us to identify and understand the various stakeholders who affect the design, development, manufacture, operation, and regulation of airplanes and their operations.

## Design Objectives

There are three primary objectives within human-centered design. These objectives should drive much of designers' thinking, particularly in the earlier stages of design. Many discussions in later chapters illustrate the substantial impact of focusing on these three objectives.

The first objective of human-centered design is that it should *enhance human abilities*. This dictates that humans' abilities in the roles of interest be identified, understood, and cultivated. For example, people tend to have excellent pattern recognition abilities. Design should take advantage of these abilities, for instance, by using displays of information that enable users to respond on a pattern recognition basis rather than requiring more analytical evaluation of the information.

The second objective is that human-centered design should help *overcome human limitations*. This requires that limitations be identified and appropriate compensatory mechanisms be devised. A good illustration of a human limitation is the proclivity to make errors. Humans are fairly flexible information processors -- an important ability -- but this flexibility can lead to "innovations" that are erroneous in the sense that undesirable consequences are likely to occur.

One way of dealing with this problem is to eliminate innovations, perhaps via interlocks and rigid procedures. However, this is akin to throwing out the baby with the bath water. Instead, mechanisms are needed to compensate for undesirable consequences without precluding innovations. Such mechanisms represent a human-centered approach to overcoming the human limitation of occasional erroneous performance.

The third objective of human-centered design is that it should *foster human acceptance*. This dictates that stakeholders' preferences and concerns be explicitly considered in the design process. While users are certainly key stakeholders, there are other people who are central to the process of designing, developing, and operating a system. For example, purchasers or customers are important stakeholders who often are not users. The interests of these stakeholders also have to be considered to foster acceptance by all the humans involved.

## Design Framework

We soon found that human-centered thinking could be applied in a wide range of domains. The process was formalized in terms of the framework shown in Figure 2 involving four phases and specific issues, methods, and tools (Rouse, 1991, 2001). These phases support the above design objectives by first focusing on the full range of humans involved in the success of the solution being designed, and then emphasizing how abilities can be enhanced, limitations overcome, and acceptance fostered.

The Naturalist Phase involves understanding the domains and tasks of stakeholders from the perspective of individuals, the organization, and the environment. This understanding not only includes stakeholders' activities, but also prevalent values and attitudes relative to productivity, technology, and change in general.

The Marketing Phase – or market research phase – builds upon the understanding of the domains and tasks of current and potential stakeholders to conceptualize alternative products or services to support these people. Product or service concepts can be used for initial marketing to determine whether or not people perceive a product or service concept as solving an important problem, solving it in an acceptable way, and solving it at a reasonable cost.

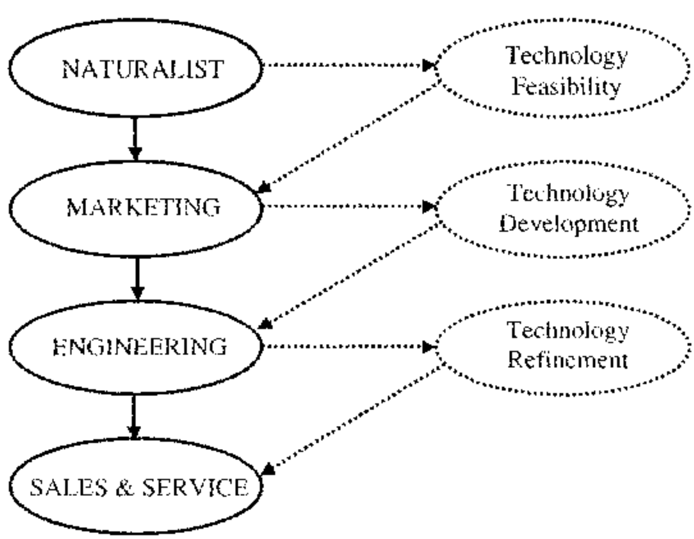


**Figure 2.** Framework for Human-Centered Design

The Engineering Phase addresses tradeoffs between desired conceptual functionality and technological reality. As indicated in Figure 2, technology development will usually have been pursued prior to and in parallel with the Naturalist and Marketing Phases. This will have at least partially assured that the concepts shown to stakeholders were not technologically or economically ridiculous. However, one now must be very specific about how desired functionality is to be provided, what performance is possible, and the time and dollars necessary to provide it.

As the Sales and Service Phase begins, the product or service should have successfully been developed and evaluated, that is, proven to be a "verified" solution. The focus is now on the extent to which it solves the targeted problem (i.e., is a valid solution), solves it an acceptable way, and does so at reasonable cost (i.e., is a viable solution). It is also at this point that one assures that implementation conditions are consistent with the assumptions underlying the design basis of the product or service.

It is important to indicate the role of technology in the human-centered design process. As depicted in Figure 2, technology is pursued in parallel with the four phases of the design process. In fact, technology feasibility, development, and refinement usually consume the lion's share of the resources in a product or service design effort. However, technology should not drive the design process. The aforementioned human-centered design objectives should drive the process and technology should support these objectives.

## Design for Success

As the human-centered design methodology matured and was applied in a wide range of domains (see Chapter 8), we conducted numerous workshops. The workshop notes evolved into a book manuscript that was to be titled "Human-Centered Design." Upon concluding a workshop at NASA where this material was used, several participants approached me and said that they thought the tentative book title was a poor choice. Further, they argued that the human-centered methodology could be used to design virtually anything, as all products and services have stakeholders, even those products that have no real "users" in the sense of anyone interacting with them. They suggested that the book be titled "Design for Success." When I returned from this trip, I called my editor at John Wiley and he readily agreed that this new title would be much more descriptive – and catchy. In this way, a serendipitous exchange in the waning moments of a workshop fundamentally repositioned a product (the book) in a broader market.

## OVERVIEW OF BOOK

As shown below, Chapters 2-11 each address what I term an essential phenomena and a central question:

2. Estimation: How to assess what is happening, has happened, will happen?

3. Queueing: How to allocate resources to provide efficient service?

4. Control: How to determine what task should be performed next?

5. Diagnosis: How to determine the cause of observed symptoms?

6. Design: How to synthesize human-centered solutions?

7. Information: How to identify sources and retrieve valuable information?

8. Stakeholders: How to consider and balance stakeholders' interests?

9. Future: How to consider future uncertainties?

10. Challenges: How to address essential management challenges?

11. Transformation: How to address and pursue fundamental change?

These essential phenomena played central motivational roles as the research on these phenomena was pursued. I can recall discussions of how estimation (e.g., predicting future states of a system) is a central aspect of being human rather than any other species. I can remember similar discussions of control (deciding what to do next) and diagnosis (deciding what's wrong). The notion of an essential phenomena implied that our research was central to understanding the essence of being human. This made the research very compelling.

Table 1 describes the essential phenomena of Chapters 2-11 in terms of the three human-centered design objectives discussed earlier. Examples of human abilities and limitations are indicated, as are possible approaches to support, that is, enhancing abilities and overcoming limitations. In the following subsections, abilities. limitations, and support are summarized across the ten essential phenomena.

## Human Abilities

The human abilities column of Table 1 includes three broad types of abilities. First. people are good at recognizing familiar patterns and mapping from these familiar patterns to past successful responses. Consequently, they tend to be good at running an "as is" system – for example, vehicle, process, or enterprise -- to achieve familiar objectives. People tend to have a wealth of "common sense" that is invaluable in situations where this knowledge and these skills apply.

Second, when they need to consider alternatives – for instance, when recognition-primed decision making (Klein. 1998, 2002) is not sufficient – they are often good at articulating their interests, defining the attributes of these interests, and specifying the relative importance of interests and attributes. However, their abilities tend to be limited to making pair wise comparisons of alternatives, as opposed to being able to compare many alternatives simultaneously. Thus, they tend to satisfice (Simon, 1969) and make reasonable rather than optimal choices.

Third, when new alternatives are needed, people are often quite good at envisioning the possibilities, formulating new ideas, and imagining the consequences of these ideas. They are also often good at articulating a vision, and leading people to pursue this vision. Sometimes, but not always, they are good at recognizing when the vision needs to change to better fit the evolving environment, for example, production needs, competitors, or the weather.

These three types of abilities are among the primary reasons why humans are included as elements of complex systems. Common sense, abilities to abstract, inventiveness, and communication abilities are almost always very important in complex domains. Thus, human abilities are usually essential. However, the information processors (humans) with these abilities also have limitations that tend to be central to the "costs" of enjoying the benefits of human information processing.

| Chapter | Essential Phenomena | Human Abilities | Human Limitations | Approaches to Support |
|---|---|---|---|---|
| 2 | Estimation: How to assess what is happening, has happened, or will happen? | Good at recognizing familiar patterns and mapping to action | Inaccurate mental models and perceptions of the state of the process | Stochastic forecasting models and displays of filtered, smoothed & predicted states |
| 3 | Queueing: How to allocate resources to provide efficient service? | Good at specifying attributes and making pair wise comparisons | Difficult to trade off multiple attributes in designing service | Stochastic process models and multi-attribute decision analysis methods |
| 4 | Control: How to determine what task should be performed next? | Good at satisficing and making reasonable choices | Poor at switching among tasks and performing several tasks at once | Adaptive aiding to perform tasks when task load becomes excessive |
| 5 | Diagnosis: How to determine the cause of observed symptoms? | Good at recognizing familiar patterns and mapping to past actions | Difficult to identify the best test and determine implications of test results | Aiding to see implications of available information to choose most diagnostic tests |
| 6 | Design: How to synthesize human-centered solutions? | Good at envisioning & formulating ideas for new solutions | Difficult to consider and balance numerous design attributes | Frameworks, methods & tools for systematically identifying & addressing issues |
| 7 | Information: How to identify sources and retrieve valuable information? | Good at recognizing valuable information once retrieved | Difficult to specify the attributes of valuable information | Aiding in tailoring and use of search tools, as well as compilation of results |
| 8 | Stakeholders: How to consider and balance stakeholders' interests? | Good at specifying interests and importance of associated attributes | Difficult to deal with stakeholders' differing and conflicting interests | Multi-stakeholder, multi-attribute models that enable tradeoffs and decisions |
| 9 | Future: How to consider future uncertainties? | Good at imagining alternative futures and possible consequences | Difficult to consider future contingencies and specify long-term returns | Decision models that provide economic assessments of the value of contingencies |
| 10 | Challenges: How to address essential management challenges? | Good at running the "as is" business to achieve familiar objectives | Tendency to be tactical rather than strategic & too focused to see situation | Toolkits that enable systematic addressing & pursuit of the essential challenges |
| 11 | Transformation: How to address and pursue fundamental change? | Good at articulating a vision and leading people in pursuing this vision | Difficult to recognize forces for change and then commit to change | Methods that address value deficiencies, work processes, decisions & social networks |

**Table 1.** Phenomena, Questions, Abilities, Limitations, and Approaches to Support

## Human Limitations

Humans have difficulty perceiving variables accurately – that's why we provide pilots instruments. However, in general, they tend to have inaccurate perceptions of system states, including past, current, and future states. This is due, in part, to limited "mental models" of the phenomena of interest in terms of both how things work and how to influence things. Consequently, people have difficulty determining the full implications of what is known, as well as considering future contingencies for potential systems states and the long-term value of addressing these contingencies. The implications for supporting humans are discussed in the next subsection.

When there is a need to consider multiple alternatives, people have difficulty trading off multiple attributes across the multiple alternatives. They tend to want to think in terms of the direct attributes of the alternatives rather than the attributes of the value of the alternatives. These difficulties often cause people to satisfice rather than optimize, in part because the nature of the phenomena (e.g., highly nonlinear) may make optimization difficult, but also because people find it difficult to consider many things simultaneously. They also tend to be poor at multi-tasking, particularly for high demand situations such as emergencies and other types of crises.

For domains where there are multiple types of stakeholders with differing and perhaps conflicting interests, people often have difficulty addressing these differences and conflicts. They tend to focus tactically rather than strategically and not see broader situations that could inform the resolution of tradeoffs and conflicts. People often have difficulty seeing forces for change, as opposed to needs to restore the status quo. Consequently, they may have difficulty committing to change and sustaining such commitments.

The above "costs" of including humans as key elements of complex systems might lead one to pursue approaches to automating human abilities to avoid these types of costs. Alternatively, one can pursue means of supporting humans to overcome these limitations and, thereby, enjoy the benefits of humans' abilities without the costs of human limitations. The next subsection considers this possibility.

## Supporting Humans

An obvious way to support humans is to provide them information – hence, the aforementioned pilot's instruments. Information can be provided at varying levels of aggregation and abstraction (Rasmussen, 1986, 1994), for instance, temperature measurements at regular points along the plant's pipes vs. energy flow throughout the power plant. One might display estimated future states, for example, using "predictor" displays for air traffic control. Another example is a display of information search results, including relationships among results and their implications.

At a deeper level, one might display underlying processes and employ various models to infer underlying states and the bases of these states. Examples include time series forecasting models, queueing process models, business process models, and value stream representations. Network models might be used to portray, for example, queueing networks or perhaps social networks.

Such models could be used to "drive" displays of abstractions of evolving systems phenomena. The models themselves might also be useful to overcoming human limitations of not being able to understand the underlying complexity of a system without appropriate abstractions. For example, libraries are not simply networks of queues of people and materials, but portraying libraries in this way can facilitate understanding and managing the complexity of libraries' business processes (see Chapter 3).

Another type of support involves helping people decide what to do – what choice to make and how to allocate resources. There is a wide range of models for addressing economic uncertainty, as well as related models for dealing with multi-stakeholder, multi-attribute decision making situations. There are both prescriptive and descriptive models for these types of problems, addressing what people should do and what people naturally tend to do, respectively. Such models can be used directly or incorporated into decision aids that embody these models.

Yet another type of support addresses task performance. Aiding can be employed to perform part of a task, to sometimes perform the whole task (adaptive aiding), or to always perform the whole task (automation). One can think in terms of a range of levels of automation, from none to complete with several levels in between (Sheridan, 1992, 2002). One often needs to provide aiding in use of aiding, for example, to assure that search tools provide the greatest value. Chapters 4 and 7 address task performance aids.

Humans can also be supported by methodologies that provide frameworks, methods, and tools in domains such as research, design, and management. Use of these methodologies can be facilitated via toolkits that provide alternative models, methods, and tools for targeting specific problems. Methodologies and toolkits can greatly facilitate the "data-driven decision making" discussed in Chapters 8-10.

For all of these types of support, there is another dimension – aiding vs. training. Aiding involves directly augmenting human performance while training is focused on enhancing the potential to perform. Given any particular knowledge or skill, the question arises of whether the humans in question need to know it or simply know how to do it. Perhaps they do not need to understand nuclear engineering, but do need to be able to operate nuclear plants, for instance. This question is pursed in more detail in Chapter 6.

## Levels of Understanding

This book is about people and organizations in terms of how to understand them and how to support them. At one level, understanding can be expressed in terms of data about phenomena, for instance, human perception of visual displays in particular circumstances. This level of understanding may enable predicting human performance in similar circumstances. To some extent, you predict by reviewing the tabulated results and, hopefully, interpolating to match the conditions of interest.

Another level of understanding enables control of the phenomenon of interest. Rather than predict what humans will do, you cause the outcomes that you desire via particular inputs and environmental conditions. I have long felt that the ability to control tends to represent a higher level of understanding than the ability to predict. Put simply, if you cannot influence a phenomenon, you do not fully understand it.

Yet a higher level of understanding involves the ability to design the phenomenon to behave as you desire. While controlling an aircraft's trajectory represents a higher level of understanding than predicting its trajectory, designing an aircraft to be controllable represents yet a higher level of understanding. The engineer in me does not feel that you really understand a phenomenon until you can design a means of enhancing that phenomenon.

Consequently, this book not only addresses the essential phenomena listed in Table 1. The lion's share of the attention is devoted to discussion and illustration of how to enhance human abilities and overcome human limitations via a wide range of support mechanisms that my colleagues and I have designed, developed, deployed – either as contract deliverables or off-the-shelf products sold in the marketplace – and supported in use by hundreds of enterprises and thousands of users.

As articulated earlier, an overarching theme of this book is the human-centered design of these systems, products, and services. Throughout the following chapters, there is much discussion of the basic and applied research that supported these design efforts, including summaries of what we learned about people and organizations. It is important to keep in mind, however, that our overall goal was not just prediction – we also wanted to control and, ultimately, design.

## SERENDIPITY REVISITED

I have portrayed human-centered design as the systematic pursuit of understanding of human abilities and limitations, as well as development and deployment of means for enhancing abilities and overcoming limitations. From this perspective, human-centered design may seem to be a very predictable process that inevitably leads to success. To some extent, this is true. However, the results of this process often are surprising.

The reason for such surprises is that serendipity intervenes to take you in another direction, offer an alternative solution, or possibly provide a novel interpretation. From this perspective, this book is a story of how one thing leads to another, in this case, in the context of understanding and supporting people and organizations.

For instance, our studies of library networks (see Chapter 3) led us quite naturally to drawing network diagrams, often on large white boards in offices or meeting rooms. These diagrams served as the basis for discussions of finding optimal routes for particular types of library services. It occurred to us that the properties of the network might not be static. For example, a node might be "up" or "down" at any particular point in time. This led to the question, "How would we know?"

At that moment, our research thrust in fault diagnosis emerged (see Chapter 5), although we did not realize this for some time. We next generated many random networks and studied how to determine a node was down. The process of generating and displaying networks was very slow as the computer searched for optimal routes. I commented on this to a colleague in electrical engineering as we met at the coffee urn. He said that he was not surprised as we were trying to solve the general circuit routing problem.

My inclination was to run off and take a look at that literature. However, it struck me that the structure of our "circuits" (as we came to often call them) was such that the optimal routes could only have five possible shapes. Rather than searching all possible routes, we could just try variations of these five shapes and pick the best one. Programming this approach led to almost-instantaneous generation and display of networks that supported a long series of studies.

Serendipity emerged by crossing domains (library networks and technical networks) and disciplines (industrial and electrical engineering). The result was a new line of research and an idea that was key to making the research feasible. The path of serendipity, in this case, was in the Coordinated Science Laboratory at the University of Illinois, a place where many disciplines mingled at coffee urns and seminars. The characteristics of such places are revisited in Chapter 12.

Another example emerged from an invitation to give a lecture at Los Alamos National Laboratory. I was CEO of Search Technology at the time and one of our largest initiatives was the intelligent interface for the Pilot's Associate, an artificially intelligent co-pilot that we were developing as a member of Lockheed's team (see Chapter 4). Another much smaller effort was focused on understanding fundamental limits of modeling behavioral phenomena (see Chapter 2).

I suggested to Chris Barrett, my Los Alamos host, that I could talk on either topic. He responded by asking me to combine both topics and talk on fundamental limits of intelligent interfaces. Having no idea what this would really lead to, I agreed and then had to invest much time into combining the two streams of thought. This effort paid off, however, as a whole new research thrust emerged (see Chapter 2). The path of serendipity in this case ran through Search Technology, Lockheed, and Los Alamos.

The chapters in this book move systematically from studies of operations and maintenance (Chapters 2-6) to studies of research and design (Chapters 7-9) to studies of strategic management (Chapters 10-11). As well planned and smooth at this path may appear, it was laced with serendipitous insights and transitions. The shift to understanding the process of research, its impact on design, and then design itself grew out of frustrations with getting the applied world to the laboratory door (Rouse, 1985).

Specifically, Ken Boff of the Air Force Research Laboratory and I, initially independently, were struck by how long it takes for research to impact practice, if it does at all. We joined forces in the early 1980s to study how best to package research results for designers. We quickly became immersed in studies of how designers work, how they seek and consume information, and the forces and motivations for their behaviors. (See Chapters 6 and 7 for discussions of these studies.)

One of our findings, not surprisingly, was that business considerations, rather than technical considerations, often drive the design issues considered and how these issues are addressed and resolved (see Chapter 8). Perhaps inevitably, the enterprises we were helping address their design processes asked us to also look at their management processes. They often noted that they were attracted to the human-centered design process we were advocating, but this process did not dovetail well with their business processes.

This eventually led to the formation of a new company – Enterprise Support Systems – focused on business strategy, market assessment, product planning, and technology strategy. Our methods and tools supported decision making by executives and senior managers. A significant level of consulting services was provided with these methods and tools, mostly for Fortune 500 companies and selected government agencies. (See Chapters 8-11 for discussions of these initiatives.)

The path of serendipity, in this case, led through the Air Force, NASA, Boeing, Motorola, 3M, Lockheed, DARPA, Raytheon, Rover, Coca-Cola, Abbott Labs, Honeywell, Hughes, Hitachi, and so on. We worked with more than 100 enterprises in the private and public sectors. Several thousand executives, senior managers, and designers continually asked us questions – and asked for help. Our dialogues, including arguments and debates, with these people led to many of the new directions, new ideas, and new methods and tools discussed in this book.

From a research perspective, the path of serendipity has led to a wide variety of studies of people and organizations, ranging from high tech to low tech, from technology to the arts. At a recent seminar where I talked about our research on teamwork in the performing arts (see Chapter 2), I was asked, "Why on earth would an engineer research such a topic?" Thinking quickly, I responded by saying that I have always found it important to have three types of research initiatives in my portfolio.

First, I like to have at least one mainstream research topic. Our recent work in investment valuation probably fits there (see Chapter 9). Second, I like to be working in an emerging area where everyone agrees it will be an important area, but it is not yet clear what form it will take. Our work in organizational simulation fits there (see Chapter 10). Third, and finally, I always like to be doing some research that causes people to say, "What the hell are you doing that for?"

The third category is often completely unsupported, both by sponsors and colleagues. Yet this work usually seems the most creative and, several years later, has provided the foundation for well-supported initiatives. The path of serendipity can be like a magnet, attracting your interest despite lack of clarity and support. My experience is that great things can and will happen – you just do not know what they are.

I began this chapter with an illustration of early serendipity – how a beat-up 1952 Plymouth needing constant repairs led to the 40 years of research discussed in this book. Another example of early serendipity was my first engineering job as an assistant engineer at Raytheon during my junior and senior years as an undergraduate. My supervisor was Alec Garnham, an English engineer who was a British pilot during World War II and a heavy smoker such as you seldom encounter any more.

Alec always helped me to see the bigger picture. He took a "systems perspective" before I knew what that meant. Our work focused on sonar systems in submarines, but Alec was able to put that in the perspective of the whole submarine, not just the sonar. During my two years at Raytheon, I worked in electrical engineering, mechanical engineering, systems engineering, reliability and maintainability engineering, and even bid and proposal operations, always with Alec as my mentor.

Transdisciplinary perspectives emerge from having to look at the whole problem, because that is the assignment and/or because you are inherently oriented that way. I have always wanted to understand the broader context. I found that operations and maintenance function within the context of design, and that design operates in the context of business. Of course, business, in turn, operates in the context of the economy and society.

These broader contexts are not organized in terms of disciplines. They are holistic phenomena that, occasionally, can be addressed by reducing them to constituent parts. More often, however, key phenomena are not evident among the parts – they only emerge from the whole. Consequently, we have to address the whole problem. Transdisciplinary perspectives provide the basis for transforming your thinking about the whole problem and moving beyond just optimizing the elements of various subproblems.

I truly believe that such thinking is key to addressing national and international issues such as health care, security, and the environment. I intend this book as a contribution to advancing this point of view.

Transdisciplinary perspectives, when combined with an openness to enlightenment via serendipity, can be a powerful means to coping with complexity.

## REFERENCES

Billings, C.E. (1996). Aviation Automation: The Search for a Human-Centered Approach. Mahwah, NJ: Erlbaum.

Booher, H.R. (Ed.).(2003). Handbook of human systems integration. New York: Wiley.

Burke, J. (1996). The pinball effect: How Renaissance water gardens made the carburetor possible and other journeys through knowledge. Boston: Little, Brown.

Card, S.K., Moran, T.P., and Newell, A. (1983). The Psychology of Human-Computer Interaction. Mahwah, NJ: Erlbaum.

Evans, H. (2004). They made America: From the steam engine to the search engine: Two centuries of innovation. New York: Little, Brown.

Klein, G. (1998). Sources of power: How people make decisions. Cambridge, MA: MIT Press.

Klein, G. (2002). Intuition at work: Why developing your gut instincts will make you better at what you do. New York: Currency.

Norman, D.A., and Draper, S.W. (Eds.).(1986). User Centered System Design: New Perspectives on Human-Computer Interaction. Mahwah, NJ: Erlbaum.

Rasmussen, J. (1986). Information Processing and Human-Machine Interaction. New York: Elsevier.

Rasmussen, J., Pejtersen, A.M., & Goodstein, L.P. (1994). Cognitive Systems Engineering. New York: Wiley.

Rouse, W.B. (1985). On better mousetraps and basic research: Getting the applied world to the laboratory door. IEEE Transactions on Systems, Man, and Cybernetics, SMC-15(1), 2-8.

Rouse, W.B. (1991). Design for success: A human-centered approach to designing successful products and systems. New York: Wiley.

Rouse, W.B. (1994). Best laid plans. New York: Prentice-Hall.

Rouse, W.B. (1998). Don't jump to solutions: Thirteen delusions that undermine strategic thinking. San Francisco, CA: Jossey-Bass.

Rouse, W.B. (2001). Human-centered product planning and design. In G. Salvendy (Ed.), Handbook of industrial engineering (3rd Edition, Chapter 49). New York: Wiley.

Sheridan, T.B. (1992). Telerobotics, automation, and supervisory control. Cambridge, MA: MIT Press.

Sheridan, T.B. (2002). Humans and automation: System design and research issues. New York: Wiley.

Simon, H.A. (1957). Models of man: Social and rational. New York: Wiley.

Simon, H.A. (1969). The sciences of the artificial. Cambridge, MA: MIT Press.