

I INTRODUCTION

If you are an organizational researcher or business educator, imagine that you had a systematic and powerful way of organizing vast numbers of things we know about business: basic principles, key scientific results, and useful case examples. Imagine that you could easily create and share this knowledge electronically with researchers, educators, and students all over the world. And imagine that all this knowledge was structured in a way that helped you quickly find the things you needed and even helped you come up with new organizational ideas that no one had ever thought of before.

If you are a computer scientist, information technologist, or software developer, imagine that different versions of this same kind of knowledge base could help you systematically organize and share many of the basic patterns and components that are used in a wide variety of computer programs. And imagine that computational tools that use this knowledge base could significantly reduce the effort required to develop new software programs from existing components and tailor them for use in specific organizations.

Finally, if you are a manager or consultant, imagine that you could use all this general knowledge about “best practices,” case examples, and software from all over the world. And imagine further that you could also create your own specific versions of these knowledge bases to share detailed information about the key activities in your own company or your clients’ companies: what needs to be done, who is responsible for doing it, and what resources are available to help.

That is the vision that has guided the MIT Process Handbook project since its beginning over a decade ago, and that is the vision that continues to guide our work. There is still much to be done to achieve the full promise of this vision, but we believe that the work we have done so far demonstrates that the vision is both feasible and desirable. This book is the story of what we have done, what we have learned, and what is left to do. It is also an invitation to others to join in the quest to help make this vision a reality.

What Have We Actually Done?

Our goal in the Process Handbook project has been to lay the foundations for the vision we have just described. To do this, we have developed an extensive, publicly available on-line knowledge base,¹ including over 5,000 activities, and a set of software tools to maintain and access this knowledge base.

1. See ccs.mit.edu/ph.

More specifically, the Process Handbook today is a combination of four things:

1. A set of fundamental *concepts* that can help organize and analyze knowledge about any kinds of activities and processes. The two key concepts we use involve the notions of “specialization” and “coordination.”
2. A specific *classification framework* for organizing very large amounts of knowledge using these concepts. Even though parts of this framework can be used to classify activities of any kind, we have put a special emphasis on developing categories for business activities.
3. A representative set of generic *business templates* and specific *case examples* to illustrate how the concepts and framework can be used. This knowledge base includes, for example, generic templates for activities like buying and selling, and case examples of companies doing these things in innovative ways.
4. A set of *software tools* to organize and manipulate large amounts of knowledge (e.g., these templates and examples) using the concepts and framework.

In principle, one could use any subset of these things without the others. But the combination of all four elements provides a uniquely powerful set of capabilities.

As the examples throughout this volume illustrate, this on-line Process Handbook can be used to help people: (1) redesign existing business processes, (2) invent new processes, especially those that take advantage of information technology, (3) organize and share knowledge about organizational practices, and (4) automatically, or semiautomatically, generate software to support or analyze business processes.

What Other Things Are Like the Process Handbook?

One of the best ways to convey an intuitive understanding of the Process Handbook is to describe other, more familiar, things that are like it.

For example, one key element of the Process Handbook is a classification system for business activities. Classification systems are ubiquitous in scientific fields. They provide a way to divide up the world and name the pieces. In this way classifications provide a language for scientific communication and a filing system to organize knowledge about the world. The best go deeper, and provide a conceptual basis for generalization and new discovery.

Periodic Table of the Elements

Perhaps the most widely known and unequivocally successful such system is the Periodic Table of the Elements, whose design is usually credited to Mendeleev in 1869. Though numerous other researchers made proposals to bring order to the ele-

ments, Mendeleev got credit because he used his Periodic Table to predict the existence and even the basic properties of as yet undiscovered elements and to rule out the existence of others.

Of course, the success of the Periodic Table is due, in part, to the nature of the elements themselves. Elements are unarguably distinguishable from each other based on chemical tests and have properties that do not change. The ordering of elements in the Table is based on an essential property, atomic number, and the arrangement of elements into groupings is based on other essential properties, such as the valence electron configuration (though these properties were in fact only fully understood after the discovery of the Periodic Table). In other words, the Periodic Table is a success because its order reflects a deeper order within the elements.

While we doubt that it will ever be possible to describe business processes with the same degree of precision as is possible for chemical elements, we do believe that a classification system like ours can significantly help organizational researchers and others to represent the deeper order within organizational activities.

Biological Classification

Another classification system with strong analogies to the Process Handbook is the system biologists use to classify living organisms. In fact the search for a way to organize the chemical elements was inspired by the hierarchical classification of living organisms first proposed by Linnaeus in 1758. Biological classification serves many of the functions we envision for the Process Handbook: it provides a standard nomenclature for describing species (so scientists can be sure they are talking about the same animals); it organizes information about different species; and it serves as a basis for generalization in comparative studies (a fact about one species is more likely to apply to other closely related species).

However, classifying living organisms is more problematic than classifying chemical elements for several reasons. First, scientists study individual specimens (a “holotype,” or representative individual), but the basic unit of the classification system is a species, that is, the population of similar individuals. Unfortunately, the definition of a species is not unequivocal, and scientists may disagree about whether two individuals are members of the same or different species. Second, the properties of species can and do change over time. Both of these properties also hold for the processes in the Handbook.

Finally, species (and processes) are much more complex than elements. As a result it is not obvious which properties should be used to organize a collection. A classification will ideally group species that share more than a surface similarity so that the groups serve as a basis for theoretically grounded comparisons. Linnaeus’s original system formed families of species on the basis of common characteristics. More

recently some biologists have proposed classifying species on the basis of their hypothesized common ancestors (e.g., Wiley et al. 1991).

Though the biological classification system is intended to be objective, it also has a strong social component. The classification system is supported by a well-developed social structure, including codified rules for naming, a bureaucracy for registering names, and conferences for vetting and accepting changes to the hierarchy. Development of some kind of similar support structure will be necessary for the full potential of our vision to be fulfilled.

Human Genome Project

Perhaps one of the closest analogies to the Process Handbook project is the Human Genome Project (HGP). The first five goals of the HGP are to:

1. “identify all the approximately 30,000 genes in human DNA,
2. determine the sequences of the three billion chemical base pairs that make up human DNA,
3. store this information in databases,
4. improve tools for data analysis,
5. transfer related technologies to the private sector”

(<http://www.ornl.gov/hgmis/project/about.html>).

The goals of the Process Handbook are broadly similar, though more modest. In our version of goals 1 and 2, we aim to identify a large number of processes and to develop a comprehensive classification for organizing them. Because of the diversity and detail of organizational processes, it would be impossible to completely describe all processes in all organizations, but the HGP will probably not sequence every variation on every gene either. Goals 3, 4, and 5 can be adopted with little change, the most significant difference being that we will organize processes in a hierarchy, implying a different set of tools for storing and analyzing them.

Engineering Handbooks

A final parallel can be drawn to engineering handbooks. Handbooks of various kinds are common in engineering disciplines to present and organize information to support designers. For example, the *Multi-media Handbook for Engineering Design*, created by the Design Information Group of the University of Bristol offers:

... a concise source of ... elementary engineering design principles, design details of machine elements and specific component information. It provides:

- design guides for a variety of design situations including the design, selection and application of components and systems
- catalogue information from component manufacturers to provide standard sizes and dimensions, ratings and capacities
- good practice guides to the proper design of components and systems in terms of increased strength, reduced cost, more efficient manufacture and assembly
- materials data for common engineering materials including properties, standard forms of supply, special treatments and typical applications.

Similar handbooks exist for chemical engineering (Perry, Green, and Maloney 1997), civil engineering (Merritt, Loftin, and Ricketts 1995), electrical engineering (Fink, Beaty, and Beaty 1999), industrial engineering (Maynard and Zandin 2001), mechanical engineering (Avallone and Baumeister 1996), and so on. Most of these handbooks include sections on basic science as well as specific applications. The Process Handbook is intended to provide at least the application-type information to support the design of business processes. Such information is represented as semi-structured information associated with various process descriptions.

The Process Handbook is not quite like any one of these other examples from various branches of science and engineering, but each of these other examples illustrates important aspects of our vision for the Process Handbook.

History of the Project

Even though we had been working on its intellectual precursors for years, the first work specifically on the Process Handbook project began in 1991. Since that time, over forty university researchers, students, and industrial sponsors have worked on developing the software and knowledge bases that today constitute the Process Handbook. For all that time this project has been one of the primary projects in the MIT Center for Coordination Science.

Even though we have refined our ideas over the years, the key conceptual ideas of specialization and coordination were present in the first full proposal we wrote for this project in 1992. For the first few years of the project's life, our main focus was on developing software tools to manipulate knowledge about processes using these theoretical concepts. Over the course of the project there have been at least four complete re-implementations of the software tools and uncounted variations and improvements along the way.

Starting in about 1995, we also began to devote significant efforts to developing business content for this framework. At first we had very ad hoc classification structures and a few more-or-less randomly chosen business examples. Over time we

added many more examples and developed much more comprehensive and systematic classification structures.

In part because of our belief that the potential for this vision would never be realized without commercial-scale efforts, several members of our project team helped start an MIT spin-off company, called Phios Corporation (www.phios.com), in 1996. Under a license from MIT, Phios developed commercial versions of the Process Handbook software tools and extended the knowledge base. For example, one of the two main versions of the Process Handbook we use at MIT today uses the commercial version of the software tools.

Over all these years, we have also used the basic knowledge base and software tools in classes, presentations to business audiences, and other research projects. In the last few years, our primary focus has shifted to demonstrating the utility of the tools and knowledge base in different applications. Today, for example, we are working on projects that integrate the Process Handbook with other tools for visualizing supply chain processes (Goncalves et al. 2002) analyzing organizational change (Brynjolfsson, Renshaw, and van Alstyne 1997), and classifying company's business models (Herman, Malone, and Weill 2003).

Structure of the Book

This book includes a number of articles previously published in a variety of different publications, as well as several chapters published here for the first time. Together, this collection of readings presents a comprehensive view of the work we have done in our first decade of work on this project.

Introduction

This initial section of the book gives an overview of the whole project. It contains a chapter by Malone and colleagues that gives a comprehensive summary of all the key concepts and major results of the project as of 1999. This chapter is both a summary of, and a foundation for, the rest of the book.

The main body of the book contains three more detailed subsections on theoretical foundations, current contents, and uses of the Process Handbook.

Theoretical Foundations of the Process Handbook

The first main section (section II) focuses on the theoretical foundations of the Process Handbook. Subsection IIA presents in three chapters the basic ideas of *coordination theory*, the source of some of the key concepts embodied in the Process Handbook. The basic premise behind coordination theory is that many activities in a

process can be viewed as coordination activities whose purpose is to manage the relationships among other activities. A key insight of the theory is that many of these coordination activities are very similar across many different kinds of processes. Furthermore, for any given coordination activity (e.g., assigning resources to a task), there are several plausible alternative approaches (e.g., first come–first served, managerial decision, auction). This means that one coordination mechanism can often be substituted for another to generate many different possibilities for how the same basic process can be performed.

Subsection IIB is comprised of a single chapter that discusses the concept of *specialization of processes* in detail. Processes in the Handbook are organized in an extensive hierarchical network, somewhat similar to the organizing principle used in biological classification. In the Process Handbook, however, we also take advantage of the concept of *inheritance* from computer science. We apply that concept here in such a way that the specialized versions of a process automatically “inherit” characteristics from more general processes.

Subsection IIC presents two discussions of what is meant by a process in the first place. One chapter uses concepts from linguistics to describe processes as grammars; the other shows how process descriptions themselves can constitute an important kind of theory for organizations.

Current Contents of the Process Handbook

Section III describes the current contents of the Handbook. Subsection IIIA begins with a summary of all the knowledge currently represented in the Handbook. This chapter shows how the basic concepts described in section II lead to a comprehensive, intuitive, and theoretically based classification framework for a wide range of business knowledge, and how this framework can be used to classify a number of specific business templates and case examples.

Subsection IIIB provides in two chapters examples of two very different kinds of knowledge included in the Handbook: organizational methodologies for business process redesign and coordination methods used in computer programs.

Subsection IIIC shows how more content can be added to the Process Handbook. It describes an approach to using the basic concepts of the Process Handbook to analyze business processes from real organizations in order to include them in the on-line Handbook.

Uses of the Process Handbook

Section IV gives examples of how the Handbook has been used in research and in practice. Subsection IVA includes three examples that demonstrate the Process

Handbook's usefulness in redesigning business processes. For some of these cases the Process Handbook serves as a well-organized but essentially passive knowledge base; for others, it is used to actively generate new organizational possibilities for people to consider.

Subsection IVB contains three chapters that show how the Process Handbook can be used for knowledge management. The first discusses managing knowledge about operational business processes, the second potential problems in product design, and the third communication genres used in organizations.

Subsection IVC focuses, in three chapters, on using the Process Handbook concepts and infrastructure to help generate and customize software systems. The first deals with the fundamental problems in specifying the architecture of any software system; the second more specifically with customizing software-based production processes, and the third with systems to support cooperative work by people in dynamically changing situations.

Conclusion

Section V concludes by a brief survey of what has been accomplished so far in the Process Handbook project. It then discusses the major challenges ahead in fulfilling the vision that has guided the project since its beginning.

A Guide for Readers from Various Disciplines

We believe one of the strengths of this project is the way it draws upon and makes deep connections among different academic disciplines. One consequence of this, however, is that not all parts of the book will be of equal interest to all readers.

To help you find the parts of the book that are likely to be of most interest to you, we therefore wish to offer a small bit of guidance about how to navigate through this book. First, we recommend that all readers start with the overview paper in this introductory section. Most readers might also want to look at chapter 8 which gives an overview of the contents of the Process Handbook.

Most of the other chapters in the book were written with readers from one of two disciplinary backgrounds as the intended audience (see table I). The two primary disciplines are computer science (including related disciplines like information technology, artificial intelligence, and software engineering), and organizational studies (including related disciplines like sociology, political science, and many parts of management).

Table I.1
Primary disciplinary perspectives of different chapters in this volume

		Primary discipline	
		Computer science	Organization theory
I	Introduction		
1	Malone et al.	*	*
II	How can we represent processes?		
<i>IIA</i>	<i>Coordination as management of dependencies</i>		
2	Malone and Crowston	*	*
3	Crowston	*	*
4	Dellarocas	*	
<i>IIB</i>	<i>Specialization of processes</i>		
5	Wyner and Lee	*	
<i>IIC</i>	<i>Different views of processes</i>		
6	Crowston		*
7	Pentland		*
III	Contents of the process repository		
<i>IIIA</i>	<i>Overview of the contents</i>		
8	Herman and Malone		*
<i>IIB</i>	<i>Examples</i>		
9	Wyner		*
10	Dellarocas	*	
<i>IIIC</i>	<i>Creating process descriptions</i>		
11	Crowston and Osborn		*
IV	Process repository uses		
<i>IVA</i>	<i>Business process redesign</i>		
12	Klein et al.		*
13	Bernstein, Klein, and Malone	*	*
14	Klein and Dellarocas		*
<i>IVB</i>	<i>Knowledge management</i>		
15	Carr		*
16	Klein	*	
17	Yoshioka et al.		*
<i>IVC</i>	<i>Software design and generation</i>		
18	Dellarocas	*	
19	Bernstein	*	
20	Bernstein	*	
V	Conclusion		
	<i>Appendix</i>		
	Lee et al.	*	

Here are some suggestions for readers with these (and other) backgrounds:

Computer scientists, software developers, and information technologists may find the theoretical perspectives on coordination (section IIA) and specialization of processes (section IIB) of special interest. They may also be interested in a number of the applications of our framework from the perspective of software engineering (chapters 10, 18, and 19), cooperative work (chapter 20), knowledge management (chapters 15 and 16), and process redesign (chapters 12, 13, and 14). Readers with an interest in artificial intelligence may find it interesting to compare our efforts to develop a comprehensive knowledge base about business intended for use primarily by human readers with Lenat's (1995) even more ambitious efforts to develop a comprehensive knowledge base about "common sense" intended for use by automated reasoning programs.

Researchers in organization studies, management science, and related disciplines may find it interesting to contemplate the possibility of a comprehensive classification system in these disciplines analogous to those in biology and chemistry. The concepts of coordination (subsection IIA), and process as theory (chapter 6) may be of special help in this goal. In addition these readers may be interested in a number of the applications of our approach to research questions in process design (chapters 9 and 12), analytical methodologies (chapter 11), and communication genres (chapter 17).

Business educators may find it interesting to consider the possible uses of approaches like this (especially chapters 8 and 9) in organizing and retrieving business school cases and other course material.

Researchers in cognitive science may find it interesting to think about the theoretical approach to studying organizations described here (especially in section II) as being, in some ways, analogous to the computational approach to studying intelligence in cognitive science.

Researchers in library science and related disciplines may be especially interested in the activity-oriented approach to classification described in chapter 8.

Managers, consultants, and others in business should find the uses of our approach described in section IV to be of special interest.

We hope also that readers from all these different backgrounds will find it interesting to look at some of the chapters outside their immediate field of interest in order to understand better how all these different disciplinary perspectives can contribute to the overall vision.