



**SUPPLY AND OPERATIONS
MANAGEMENT COLLECTION**

M. Johnny Rungtusanatham and Joy M. Field,
Editors

Mapping Workflows and Managing Knowledge

*Using Formal and
Tacit Knowledge to
Improve Organizational
Performance*

Volume I

John L. Kmetz



BUSINESS EXPERT PRESS

Mapping Workflows and Managing Knowledge

Mapping Workflows and Managing Knowledge

*Using Formal and Tacit
Knowledge to Improve
Organizational Performance*

Volume I

John L. Kmetz



BUSINESS EXPERT PRESS

Mapping Workflows and Managing Knowledge: Using Formal and Tacit Knowledge to Improve Organizational Performance, Volume I

Copyright © Business Expert Press, LLC, 2016.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopy, recording, or any other except for brief quotations, not to exceed 400 words, without the prior permission of the publisher.

First published in 2012 by
Business Expert Press, LLC
222 East 46th Street, New York, NY 10017
www.businessexpertpress.com

ISBN-13: 978-1-63157-387-3 (paperback)
ISBN-13: 978-1-63157-388-0 (e-book)

Business Expert Press Supply and Operations Management Collection

Collection ISSN: 2156-8189 (print)
Collection ISSN: 2156-8200 (electronic)

Cover and interior design by Exeter Premedia Services Private Ltd.,
Chennai, India

First edition: 2012
Second edition: 2016

10 9 8 7 6 5 4 3 2 1

Printed in the United States of America.

To my grandsons, Owen, Nicholas, and Carson

Abstract

This is a book that does what the title says, and is different from most business process mapping information in three key ways. First, it lets users capture all the knowledge that goes into a workflow in any kind of organization, including the most difficult kind of all, the tacit knowledge people bring to the job and carry in their heads. Second, it is simple, powerful, flexible, and easy to learn. Third, it does not require installing, learning, and applying a complicated program (sometimes requiring reorganization to support the software rather than the software supporting the organization). It was developed by the author in a 15-year long program of studying, analyzing, and improving avionics maintenance processes for the U.S. Navy and the Royal Canadian Air Force, and then applied to organizations of all kinds ever since, for more than two decades. It has been taught and applied by the author and others in many short courses. It works.

Keywords

Business process, business process mapping, workflow mapping, knowledge management, tacit knowledge

Contents

<i>How to Use This Book</i>	xi
<i>Objectives of This Book</i>	xiii
<i>Introduction</i>	xv
Part I Do Workflow Mapping Now	1
Chapter 1 Workflow Mapping Fundamentals.....	3
Chapter 2 WFMA Data Collection and Analysis.....	49
Chapter 3 Implementing Workflow Mapping and Analysis	91
Part II Why Do Mapping?	113
Chapter 4 Knowing What We Know.....	115
Chapter 5 Systems, Processes, Organizations, and Workflows.....	141
Chapter 6 WFMA and Knowledge Management	161
<i>Appendix: A Brief Summary of the NAVAIR Study</i>	189
<i>Notes</i>	203
<i>References</i>	207
<i>Index</i>	215

How to Use This Book

This book, the first of two volumes, has been written primarily as a how-to guide to Workflow Mapping and Analysis, and is focused on mapping organizational processes. At the same time, it explains *why* we need to do certain things, and it provides some thinking tools to help achieve that.

As we will see in Chapter 4, “processes” are combinations of actions and information. To really understand a process and then map it, you have to understand both of these parts. The actions and the material we work on (often information itself) are the easy parts to capture. Information is trickier—the formal part is the information we can see in rules, procedures, policies, and the like—and that is usually pretty visible. But people do things in different ways, for many reasons, and so the real hard part is capturing this “tacit” knowledge, which gets built into how everyone does his or her part of a process. If a company or organization wants to understand its processes, and perhaps set up a knowledge management system at some point, it will have to include this tacit knowledge.

For a person who needs to get a handle on processes right now, the place to start is Chapter 1—this is the core chapter on process mapping. Chapter 2 then helps figure out how to collect data on a workflow from the maps we create, and analyze and interpret it to do things like process improvement and change. Much depends on intelligent use of our eyes, but a spreadsheet helps. Simplicity is the key to both making and analyzing the maps.

At some point, the user ought to read Chapter 5 to have a fuller understanding of organizations as “systems,” which they all are, and of processes within them. Since information is fundamental to everything we are trying to do, the user should take some time to read Chapter 4. From the perspective of Chapters 4 and 5, doing what Chapters 1 and 2 show will make for better workflow maps and the applications users put them to later. Chapter 3 will also help, in that the organizational issues that frequently come up in process mapping are treated in more detail. In comparison with Chapter 1, Chapters 4 and 5 may seem somewhat

academic, but those chapters provide thinking tools that are as important as the mapping tools, and they work better together.

Finally, Chapter 6 examines an area that has been important but somewhat controversial over the past few decades, and that is knowledge management. Everyone recognizes the significance of this topic, but exactly what it is and what to do with it are not. I think that a well-developed system of process mapping, and more importantly, process understanding, is an integral part of an organization's knowledge management system, and perhaps is the one that creates flexibility to determine who lives or dies in a competitive environment.

Can we actually do workflow mapping without software, as I suggest many times in this book? To draw the maps, we really don't need anything more than the drawing programs that are built into word processors and spreadsheets, but a dedicated program like *SmartDraw* or *Visio* (among others) can be very helpful. For data analysis, you need a standard spreadsheet. What you don't need is a costly, complex, proprietary process mapping program (and often all the support staff that comes with it).

Objectives of This Book

This book is volume I of two volumes. This is designed to achieve the following objectives, without burying the reader in academic jargon:

1. Provide a general framework for understanding knowledge and information, the foundations of both processes and workflow maps.
2. Define Workflow Mapping and Analysis (WFMA) as a descriptive and analytical tool.
3. Illustrate the need for WFMA to be understood and applied as a disciplined approach to the analysis of workflows and processes.
4. Demonstrate the utility of WFMA for analyzing and collecting data on organizational processes.
5. Illustrate the applicability of WFMA to universal organization issues such as efficiency, quality, and control; it also supports many functions in HRM, training, certification, and process improvement.
6. Illustrate how WFMA enables users to capture both the formal knowledge required to do the work of an organization and embedded process-relevant tacit knowledge.

WFMA is a *generalist* tool for managers, supervisors, and professionals needing to understand how work is processed through an organization. WFMA is *not* a specialist tool such as programming, data flow diagramming, and so on, and it *does not require* software training or extensive software mastery. Nevertheless, like nearly any current information-intensive task, software can be a powerful tool that can support WFMA if it is used to proper advantage.

Using WFMA effectively takes a little practice, like most new things in the world, and I have provided some simple exercises in the context of familiar terms and situations in Chapter 1; these exercises also illustrate some of the differences between WFMA and general flowcharting, which are helpful to keep in mind. For those who might be curious about the origins of this approach, I have added a short appendix on the original NAVAIR work.

For those who are interested in going farther and using simulation software to investigate potential future states and configurations of their organizations, Volume II is recommended. I have attempted to convey important content and several stories of my own experiences and interactions with other organizations, and I always learned more than I gave in these episodes. Enjoy.

Introduction

7:00 AM is a miserable time of the morning to have to start work; it is even more miserable time to begin a presentation on a difficult job to an even more difficult audience, only to have it end by nearly being booted off the platform about 15 minutes later. I know this because it happened to me a long time ago, when I first tried to present a complex diagram of a flow of work to a meeting of the Common Automatic Test Equipment Integrated Logistics Support Management Team (CATE/ILSMT, to insiders) in San Diego one beautiful summer morning in 1980. Out of this experience, in some ways, grew the procedure I am going to share in this small volume. This technique, called workflow mapping and analysis (WFMA), is a way of visually capturing a flow of work by using a small set of symbols in a very consistent way; it was born from a need to visually portray a challenging flow of work in the maintenance of aviation electronics (“avionics”) for the U.S. Naval Air Force (“NAVAIR”) during the Cold War. This is a method that was developed from a need to communicate important information about that workflow to military and civilian navy personnel who needed to know it, and not get booted off the platform before that communication could be accomplished.

Within a few years, I had developed a revised approach to WFMA, which met with a much better reception, and hence was much more effective in helping me and several colleagues meet our objectives in our research with the CATE/ILSMT. A large part of this had to do with mapping a hugely diverse range of workflow activities and procedures, along with their supporting flows of information. This required a technique that was flexible and robust enough to capture everything that happened in that flow of work. I refined this technique even further when I began working on an article for the academic journal *Administrative Science Quarterly* that was published in 1984. By that time, I had come to realize that while WFMA used a small set of standard flowcharting symbols that I had learned in my college programming courses, WFMA was not “flow-charting” by any stretch of the imagination; that was one of the things I had to unlearn to really make this technique work. Of much greater

importance was the realization that WFMA was a potent tool to capture all kinds of knowledge embedded in a workflow.

Among the chief lessons I had also learned was that the symbol set necessary to communicate graphically about flows of work and information could also be the greatest obstacle to success. The symbols are necessary, but they must be as simple and unobtrusive as possible if they are not to smother the information they provide. The particular symbols and the disciplined approach to WFMA that I present here may seem overly simple at first and overly rigid at other times; I argue that neither of these is true, but rather that my approach has been developed through many years in the infamous school of hard knocks, and what has emerged is what works. Some of my experience tends to support an old, somewhat cynical definition: “‘Experience’ is that which makes you wonder how it ever got a reputation for being the best teacher.” But it was, and what I learned in those years is one of the principal influences in shaping my approach to WFMA.

My experiences in the NAVAIR environment were given another challenge in the early 1990s when I was asked to examine the avionics maintenance processes used by the Royal Canadian Air Force (RCAF). In 1980, the Canadians purchased a new wing of CF-18 Hornets from McDonnell-Douglas (now Boeing) to replace an aging mix of various aircraft in the RCAF. The CF-18 is an export model of the U.S. Navy’s F/A-18, and is supported by a similar type of automatic avionics tester, which was the subject of my studies in the U.S. Navy; within a few years, the Canadians had begun to experience the same kinds of avionics maintenance workflow problems as had the U.S. Navy, and similarly suspected the automatic tester to be the culprit. They came to the tester manufacturer for help with solving the problem; I was asked to take this on, and I found that analyzing their workflow through WFMA was again the necessary first step, just as it had been for NAVAIR. The end result was a 1991 management manual that helped the RCAF deal with the challenges of this workflow.

By the early 1990s, I had begun to develop this approach into a course that I still offer, and applied it in a number of companies and nonprofit organizations. I found that my approach functioned just as well in service workflows as it did in maintenance or manufacturing environments. This

is not very surprising in retrospect, since the original NAVAIR workflow was a repair process; while it handled physical materials, it was really a knowledge-intensive service activity that required the diagnosis of faults and the replacement of parts, with follow-up testing to be sure that a repair was successful. All of the major issues in that workflow, we came to find, were issues of information and the processing of it, and not issues with the hardware or software; the avionics maintenance process is really an excellent example of “knowledge workers” at work.

A final realization from these years of experience and investigation was that one of the principal benefits of WFMA is that it makes many aspects of the knowledge embedded in a workflow explicit even when these involve information processing and decision making that are the results of many cycles of trial and error, perhaps involving many people and long periods of time. In today’s discussion of “knowledge management,” this has been recognized as a principal form of “tacit knowledge.” WFMA is capable of capturing both formal and tacit knowledge and making it accessible to a company or organization. In the case of NAVAIR, it became an absolute necessity for resolving a very serious problem of Cold-War readiness for the fleet’s carrier aircraft.

Knowledge management (KM) is one of the most challenging tasks facing many companies today, and is often one of the most frustrating. WFMA can not only make both formal and tacit knowledge visible and accessible, but it can also provide a repository for storing both types of information. Thus, it can be a tool for creating and storing job descriptions; it can serve as a training tool; it can be a key tool in personnel accession planning; and, of course, it is a fundamental requirement for Six Sigma programs, quality management, and process improvement.

As the years have passed, one additional capability that I have learned in this environment has come to the foreground, that is, the use of simulation to investigate alternatives to many organizational variables—fundamentals such as sources of revenue, cost comparisons, structural variations, and many others. In fact, the whole of organizational life can be experimented with using this technique, and so I have split this book into two volumes. This is the first; the second volume deals with the next step, which is using simulation to investigate such what-if futures. Interested readers are invited to join me in that inquiry.

PART 1

Do Workflow Mapping Now

CHAPTER 1

Workflow Mapping Fundamentals

Introduction

This chapter is designed to help those who need to do workflow mapping and analysis (WFMA) of their business processes immediately. Chapter 1 covers the fundamentals of workflow mapping, and Chapter 2 covers the analysis of data obtained from workflow maps. Chapter 3 provides some guidelines for implementation of WFMA in different organizational circumstances.

Understanding its processes is fundamental for any organization to achieve its goals, no matter what they are. The highest-level process might be thought of as the “business model,” and being able to understand and control that is of paramount importance. Consider the case of Blockbuster Video, the literal textbook success story of the 1990s—customers came to their stores, rented videotapes and DVDs, and brought them back. Netflix realized that with the growing shift to home DVD players and the growth of the Internet, the “store experience” could be done online and transportation provided by the U.S. Postal Service, all with very low overhead. With no change in the overall service goal but a change in the customer service process, Netflix crushed the Blockbuster success story in the next decade, grew again through early and aggressively priced streaming video, and then stumbled badly by attempting to raise the price for its mail service DVDs by 50 percent at one time. Since then, Netflix has expanded its content to include popular television series in addition to motion pictures and original productions, and has become a threat to the standard cable-television model. The business-model wars in the entertainment industry will be a constant source of pressure to capture relevant information about these processes and adapt them as tastes and technologies change.

This chapter provides a singular approach to process mapping for all types of users, including those who want to analyze their business model. But it also serves users ranging from those needing to study and understand complicated problems bound up in a flow of work, to those needing to capture and codify the knowledge used in their processes, and to those needing only to diagram a well-understood procedure for purposes of training new employees or documenting a process for an ISO 9000 or other quality certification. This system creates the maps that are the basis of these and more extensive analyses for process improvement and other applications. In most cases, these applications will have much in common with each other, and most users will find that this chapter applies to them no matter what their objective. To differentiate between my methods and those of others in the discipline, I have termed my system the “Kmetz method.”

The Kmetz method of WFMA is quite unlike the majority of approaches to process mapping in two significant ways. First, it does not require the installation or application of any specific supporting software. The core objective of the method is understanding and mapping of workflows *as they exist* in companies and organizations, and this is key for any further analysis or process change. Drawing software is certainly helpful to record a map, but even that is not an absolute requirement. Second, tacit knowledge is not an impediment to applying the method, and is actually captured by the method and incorporated into its workflow maps. The Kmetz method reveals and subsumes tacit knowledge rather than considering it an aberration in a workflow or an obstacle to process mapping.

Most WFMA projects encompass some elements of routine recording and some elements of diagnosis and discovery. WFMA meets these objectives by showing how the “stuff” of any organization’s work, whether tangibles like manufactured materials, or intangibles like information or changes in patients’ states of health, moves through that organization. In today’s economy, the “material” that most organizations work on is information itself, and for most of the time I will also refer to this as material, even though it is not. Any of the things that organizations do to accomplish their objectives involves both the processing of their materials and the processing of information, in the form of formal and tacit knowledge.

While WFMA is an increasingly important tool in Six Sigma programs and other aspects of quality management and performance improvement,

there is no standard format for creating workflow maps;¹ not uncommonly, these are referred to as “flowcharts” or “flow diagrams,” “process maps,” “business process maps,” and so on.

I use the term “workflow mapping” to differentiate what is being done here from these other methods, and to reinforce the idea that what we are creating with this method really is a “map,” more often through uncharted territory than we might imagine. In discussing the unknown aspects of processes, always with the unintended outcomes of processes and the like in mind (see Chapter 5), we should expect to find these and other aspects of processes whenever we begin a mapping project. I always encourage anyone beginning such a project to adopt a mindset of “assume nothing.”

Hasn't This All Been Done Before?

Before we go on, we should answer a good question—hasn't this type of book already been done, or isn't this material all out there on the Web? I should point out that to those who have done some Web surfing to examine materials on flow mapping, business process mapping (BPM), business process reengineering (BPR), knowledge management (KM), and related topics, it may seem rather strange to propose that we should start with another clean slate.

There is no question that a lot has been done to bring workflow mapping to the level of management awareness it deserves, and that much effort has been expended to create a body of tools and techniques to enable managers and analysts to do it. In particular, organizations like the Workflow Mapping Consortium (www.wfmc.org) have been a great benefit to the field.² Founded in 1993, the WfMC is a global nonprofit organization dedicated to promoting awareness and uniform standards for WFMA, and publishes an annual handbook. It comprises over 300 member organizations and has developed standard methods for graphing and modeling workflows (Wf-XML and XPDL), which enable companies and organizations to use information technology in standardized ways and with maximum interoperability between systems and locations. Many members and nonmembers of WfMC have made major contributions to this field.

In March 2010, IBM Corporation and Lombardi LLC held a webinar to discuss “lessons learned from the first decade of BPM.”³

This was promoted by another Web-based organization, the Business Process Management Institute (www.BPMInstitute.org), also an excellent resource for process mapping information.⁴ However, one might ask what WfMC believed itself to be doing if it had been founded eight years earlier than the beginning of that “first decade.” While I like much of what IBM, Oracle, WfMC, and BPMInstitute have to offer, and am especially glad for the support they provide to an important function like process mapping, these organizations also illustrate the diversity, and sometimes the confusion, that surrounds it (Harmon and Wolf, 2014).

So why do we need anything more? The short answer to this question is that while there is a great deal of material on the Web or published elsewhere, much of it excellent, it is unfortunately true that confusion and complexity are big problems with this body of work. I will expand on my reasons for this a bit more in a moment, and Chapter 5, where I discuss the broader issue of KM.

Much of what I find on the Web suffers from several problems. The biggest problem is that to use the tools and methods found there, the user is going to have to invest considerable time and energy into learning a fairly complex software-driven mapping tool before they can begin mapping their own processes. These mapping efforts become major projects, with all that such an undertaking entails, to plan a major program, detail what must be done at each step, gain support for doing it, and the like.

In making that investment, those users will have to learn a number of conventions and rules (sometimes a large number), and then map their processes using a predefined structure that is usually associated with a specific software package. Increasingly, the programs associated with these “BPM” approaches are proprietary top-down applications where “process” effectively means “automated process,” despite occasional nods in the direction of nonautomated methods.⁵ Along with the growth in sophistication that has accompanied these methods over the past few years, there is therefore an accompanying dependence on the use of information technology; I am a big fan of IT, but there are cases where the software and hardware can get in the way of things we need to do quickly and simply. These approaches are also often costly. Of course, companies can always hire consultants to do this for them.

A second associated problem is that many of these methods are overprescribed in at least one of two ways. In a word, these all are, or become, *specialist* approaches—they ultimately require the level of mastery only a specialist can provide through careful study if they are to be applied correctly. In many cases they require the user to apply a relatively rigid set of methods to workflow mapping, often involving special software and predefined approaches to mapping and storing data. This requires learning the vocabulary of that approach, its software application, and basically how to fit everything we need to learn about a process into that framework. At its worst, it essentially requires users to redesign processes to accommodate the software and its assumptions, that is, the tail wags the dog.

Specialist methods inherently limit accessibility to process mapping and the interpretation of data for many who might benefit from the information these tools would provide. An individual in a large company in need of help with an inefficient process may find that a specialist approach to deal with the problem must go through a review and approval procedure that may delay needed fixes for far too long; similarly, the small firm with a limited budget may not be able to afford the fix, even if it could be done quickly. Complexity in process-mapping approaches does not serve these potential clients well.

The other type of overprescription is often that the providers of mapping software and techniques have attempted to catalog processes and constrain the user to applying predefined selections from that catalog, regardless of fit. The logic behind this creation of catalogs is ironically derived from the basic idea of workflow mapping—capture the information and steps involved in a process by drawing a picture of it. Having done that, it is easy to extend the logic a bit and conclude that there must be a finite number of small but similar process components, and so creating an exhaustive catalog of them would permit anyone to “assemble” maps of complex processes from the contents of the catalog, expedited with a bit of software support.

This idea is so appealing that in the early 1990s the Massachusetts Institute of Technology (MIT) undertook exactly such a project to create and commercialize a master repository of over 5,000 business processes.⁶ With this repository and its accompanying software, it was thought, any

and every organization could map its processes directly by examining the company, matching or building its processes from the repository, and then assembling them into a master map. The project was commercialized and turned over to the Phios Corporation in 1996.⁷

The project appears to have foundered. On visiting the company website, one finds little reference to the catalog (it is now a subset entry of a larger page), and the emphasis of the company appears to have shifted from process mapping to KM. As of the first printing of this book, a few entries from the late 1990s, 2003, and 2006 could be found, but little evidence of more recent activity. MIT had little to say about it, but in 2006, Phios announced that it had turned its process catalog over to MIT in an open-source format. Neither of the links to the MIT open-source materials worked as of the first printing, and there is no evidence of such a project on the MIT website 10 years later.

What happened? My guess, in two words, is “the inevitable.” For reasons that will be discussed more fully in Chapter 5, workflow processes simply cannot be predefined in the way the MIT project assumed. In addition, the diversity of companies and organizations, the speed and extent of change and evolution in them over time, along with the difficulties in collecting valid information on much of what they do in their processes, conspire to make assembly of a comprehensive catalog nearly impossible. In short, the MIT project was doomed from the outset. It made the fatal assumption that processes are fixed entities, which, once specified, will be found in the same configuration no matter who does them. I think this is most unlikely, and moreover, even if it were true, the variety of organizations out there will find many ways to create processes in their systems that were not foreseen when the repository was formed. People will combine formal and tacit knowledge to create unique and individualized processes that were simply never envisioned in the original organization. Both the MIT catalog and software were certain to have major problems in adapting to these variations, and I would guess they could not. Tacit knowledge is always the major obstacle for such programmed workflow mapping; further, exception handling is always a major challenge, and as we will see is often the most difficult part of a process to map and understand. Exceptions, by definition, cannot be predicted, and their resolution cannot be predetermined.

In contrast to the MIT project, APQC (the former American Productivity and Quality Center) may have had more success with their Process Classification Framework (PCF), begun in 1992, in part because it is a classification scheme, not an exhaustive inventory.⁸ APQC has a number of case studies showing successful applications of its framework in different industries, and, like WfMC and several other organizations, has expanded its coverage of KM in recent years.⁹

In many ways, what seems to have happened to the MIT project is that it ran headlong into information issues that will be discussed more fully in Chapter 4; these are that in working organizations, the “imperfections” of information, the adaptability of systems, and the creativity of people always combine to create truly unique entities to which one-size-fits-all solutions will never apply. In a complex world, we have to come up with our own answers to the challenges of life. Even in a specialized and delimited field such as software development for “knowledge engineering,” characteristics of the “domain” or general environment of a business are so diverse and difficult to capture that creation of a meaningful knowledge base and the collective definition of what its content should be is a major challenge.¹⁰

The issues I have with many existing systems have long been recognized by others, as summarized by Alonso and his colleagues in the abstract to their analysis of automated workflow systems:¹¹

Workflow systems hold the promise of facilitating the everyday operation of many enterprises and work environments. As a result, many commercial workflow management systems have been developed. These systems, although useful, do not scale well, have limited fault tolerance, and are inflexible in terms of interoperating with other workflow systems.

I believe these challenges continue to explain why major vendors of BPM systems all seem to agree that many such software projects fail, wholly or in part,¹² and that the quest for the successful IT-based system that will solve these problems also continues.¹³

For all these reasons, the Kmetz workflow mapping method in this book is deliberately the opposite of these overprescribed approaches—after

learning the simple symbols and rules in this chapter, anyone can quickly begin to model and map workflow processes. My method is better described as “disciplined” rather than “standardized,” and while it is likely there will be much variation in the way different users might describe the same process, this method allows for multiple “right answers.” My approach is *nonspecialist*, and can quickly be learned and applied by nearly anyone. Most importantly, rather than being blocked by tacit knowledge in the workflow, the Kmetz method simply captures and describes it. This method can be understood by anyone who is not a trained user in just a few minutes. Finally, for those who are considering the next step into an IT-intensive system of workflow modeling, they will find that my method is quite consistent with the basic approach used by WfMC and provides invaluable preparation for successful implementation of such systems.

Workflow Mapping and Analysis

What Is WFMA?

As I pointed out earlier, I differentiate my approach to WFMA from IT-dependent mapping methods. WFMA is a graphic method of completely describing the joint material and information flows necessary to accomplish one or more specific objectives of work, in their correct sequence, in a single job, a process, an organizational unit, or an entire organization. As we will see very shortly, WFMA is designed around *simplicity* and the application of a *discipline*, rather than software or more elaborate graphical methods. It certainly may take advantage of drawing software support, but no specific program or package is required.

To begin, a workflow map shows exactly how all materials and information are combined in the correct sequence to accomplish a known end purpose. Unlike a system diagram (shown in Figure 4.4), which shows the overall structure of the system, the workflow map shows the linear sequence of operations on material and information to achieve a system objective. A system diagram and a workflow map for the processes in that system are quite different, and it is often a useful exercise in gaining understanding of a system to diagram both of these. Our present concern, however, is the workflow map.

Workflow maps have a number of important properties. First, they are *graphic*—they show workflow processes visually, diagramming them as a

combined flow of activities and information. Second, the method here is *singular*, meaning that it uses one specific symbol set for all workflows or processes being mapped, and that these are used with a specific set of rules. Third, these rules constitute a *discipline*, and this is as much part of the Kmetz method as the symbols. Fourth, they are *scalable*, in that they can encompass all parts of a process at whatever level of detail is selected by the user. Fifth, they are *robust*—the mapping method we will see below can be applied to any flow of work in virtually any kind of organization. Sixth, they are *verifiable*—maps can describe existing processes as they are, and any map can be audited or evaluated against the actual flow of materials and information and the behavior of jobholders to determine its accuracy. Finally, these properties make WFMA an important precursor for workflow measurement and *process improvement*—all workflow activities, flows, and decisions can be measured in a variety of ways that support change and process improvement.

Workflow mapping creates a “static model” of a process within a business system. When a workflow map is created and verified, the user has a graphic depiction of everything needed to do the work the process includes and identify all information needed to perform the process and handle exceptions to it. If a manufacturing process or an insurance claim process is mapped, people can “walk through” the process as if they were on the plant floor or at the claim site.

There are two major components of the Kmetz method—the symbols used to diagram the workflow maps, and the discipline that governs the methods and conventions used to create the maps. Both are important.

Part 1 The WFMA Symbols

In Figure 1.1, the reader will find the basic symbol set used for all WFMA in my method. The usual reaction to first seeing these is, “Is that all there is”? or “Big deal”! That is entirely by design. One of the hard lessons learned in developing this system in my Naval Air Systems Command (NAVAIR) studies (see the Appendix) is that the symbols have, in fact, very little to do with WFMA; they are a means to identify and order the universal elements comprising any flow of work, but the real objective is to understand the information content of the map, and that has much more to do with the words and expressions in the symbols, along with the

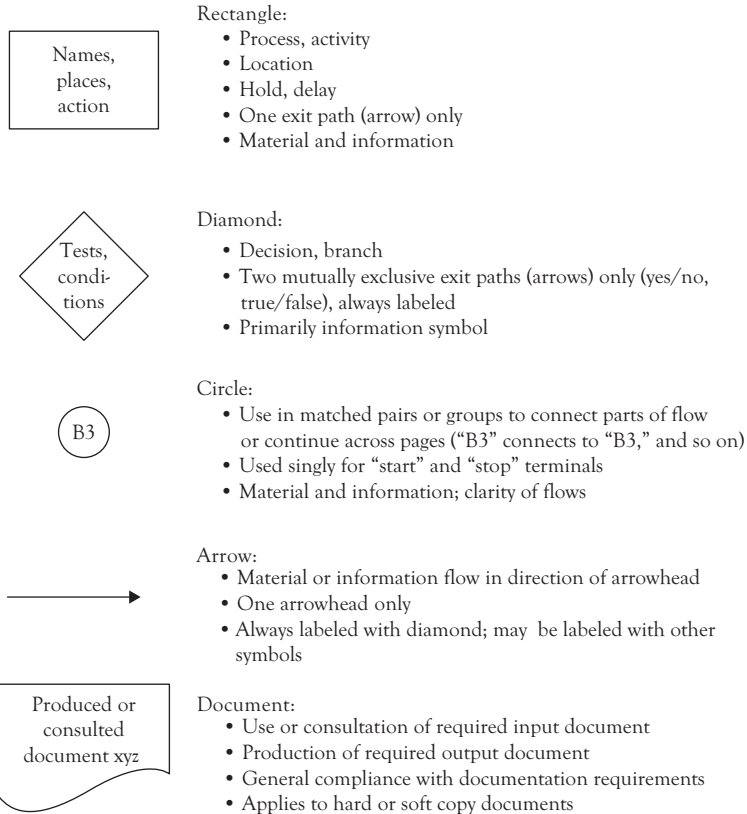


Figure 1.1 *The WFMA symbol set*

Source: Adapted by permission of the publishers, in *The Information Processing Theory of Organization* by John L. Kmetz (Farnham: Gower, 1998), p. 378.

relationships between them, than the symbols themselves. The symbols are important, as we will see, and must be used in a consistent way. But they cannot be, and should not be, the dominant part of the mapping process. This is one key element of the designed simplicity of this method.

The mapping symbols in Figure 1.1 are a very basic set of flowcharting symbols that have been adopted as the entire set for creation of workflow maps. No other symbols are needed, and as we will see, no others are desired. There is nothing unique or original about these, and many different flowcharting programs and systems use them as well. Why these were selected is simply that through multiple trials of mapping workflows, validating these maps with operators on the site, seeking

diagrams that could communicate clearly with other analysts, managers, and specialists in the Navy and elsewhere, these are the ones that did all the things I needed. These five survived a Darwinian selection process. They work. They have been used in every conceivable circumstance and are so robust that any organizational process can be mapped using them.

The first is the simple rectangle (and this must always be a rectangle—a square is not an acceptable substitute). It can show a huge variety of events and operations, including processes, activities, and location; it can be used to designate delays or holds (the absence of activity over time), such as waiting for paint to dry. It can be used for materials and information. It is often referred to as a *process block*. It has only one exit path, although multiple paths may enter a process block.

The second is the diamond. This is simply a small square rotated 45°, and represents a decision or a branching point in a flow. It is because this is a rotated small square that the rectangle must never be replaced with a square—these symbols must be kept visually distinctive. The diamond is the *only* symbol that has *two* exit paths, which must be *mutually exclusive*. It is primarily an information symbol, although it can be used to show how physical items are parsed or divided. It may also have multiple entering paths.

Third, there is the circle. This can be used individually as a terminal to designate “start” and “stop” points, but is most commonly used as a connector for clarity—continuations where matched circles enable a large, complex map to be arranged in sections. A connector circle labeled “A” at the end of an arrow designates that another “A” at the head of an arrow leading into the workflow is where the flow continues. These may also be used to indicate where different parts of a path might lead back to a single reentry point in a flow. (Since many programming purists over the years have harassed me on the use of circles for “start” and “stop” events, I have allowed the small circle to be stretched into an oval, without repercussions for those who do this, but an oval is not really necessary and I always use the circle.) A properly labeled circle will do everything an oval might do.

The fourth symbol is the arrow. This is a single-headed arrow to be used individually between other symbols. Multiple arrows from any source other than a decision diamond are not allowed, nor is a double

head, that is, an arrow with opposing heads on each end. Arrows exiting a diamond are always labeled; those exiting a process block usually are not, but may be for clarity.

Finally, there is the document. This is a bit of an anachronism. The truth is that most of what is shown using a document symbol can be represented just as effectively with a rectangle, but I have found that there are often circumstances where the need to show that a document was produced, or a rule referred to in a process, or a report generated, and the like, were very important. For that reason, it survives. Entry and exit flows are the same as for a process block.

All of these symbols are associated with brief descriptive information. Process blocks (rectangles) usually contain *names* and *statements*. Diamonds usually contain *conditions* in the form of *questions* or *tests*, and these should be unambiguous. Circles tell us what connects to what, or where we start and stop workflows. Arrows are often used without descriptions, but exit paths from diamonds should *always* be labeled. Finally, documents should have a description of either what was done (like a process block), or what the document is, and sometimes both.

The selection of these symbols was rather Darwinian, as I said earlier. What I mean by this is that they are the “survivors” of a number of different candidates that I tried over time. There are several related reasons for their survival and for my insistence that they be used only as shown here. First, they are all *visually distinctive*—if drawn as shown, there is no possibility for any of them to be confused with any other symbol. Second, they can be *learned very quickly*, without requiring a user to first memorize a large number of symbols or the mastery of any particular software. Third, they are *robust*, such that they can be used to map any flow of work, whether regular flows or exceptions. Finally, they are *comprehensive*—they are able to capture all activities and information in a workflow or process, whether material, formal knowledge, or tacit knowledge. Together, they establish a “graphic vocabulary” both for mapping workflows in organizations and for effectively and efficiently communicating what is found.

Part 2 WFMA Discipline—the Rules for Symbol Use

Organizational processes are a combination of actions and information, which we will elaborate on in Chapter 5. But however simple the symbols

are, do not be deceived by that simplicity! After many years of experience with workflow mapping, there are three issues that always confront new users of WFMA: (1) Mapping the material actions in a workflow is by far the easiest task, since these are the things that actually move through the process. This can create a challenge for new mappers, because they tend to focus too exclusively on the materials, and have difficulty with information. (2) Seeing and capturing *information* in a workflow map is one of the most challenging parts of the job. We run into both explicit and tacit knowledge everywhere, and capturing that is always a far bigger challenge than capturing the primary flow of materials. The tacit information is tricky because it is characterized by “four I’s”—it is *individualized*, *internalized*, frequently *invisible*, and *idiosyncratic* (often to the company or organization). (3) A significant part of the mapping process is to create a document that is clear and orderly on the page, so that it is as informative as possible to its users. This will inevitably take more than one draft to get into the final form we want. Like any other skill, *practice* at workflow mapping is necessary, and makes it easier every time we do it.

While we diagram material and information flows separately in a system model (see Figure 4.4), in a workflow map materials and information are usually tightly bound together. We use information to acquire, transform, and deliver the materials in the workflow. Decisions are made, tests conducted, protocols followed, documents completed, and much more, all as part of the combined flow of materials and information. The map should show these parts of the process as they occur, usually together and usually in a specific sequence.

There are a number of rules for symbol use, which are discussed in the next section. I refer to these as a *discipline*—“rules” rather than “suggestions” or some gentler term—because I have found that violating them, even slightly, leads to still more violations, and that these progressively destroy the consistency and simplicity that makes the Kmetz method of WFMA robust and effective.

The rules are the syntax for constructing workflow maps. When applied correctly, they make it possible for the symbol set to quickly retreat into the perceptual background of any WFMA user. *The objective of such a simple set is to make it “disappear” in use—if more than two minutes is needed to completely explain the symbols, they are too complex.* Literally, if

one flips through a five-slide show, with one symbol and the bullet points explaining it as shown in Figure 1.1, an audience with handouts should need no more than 25 seconds to understand each symbol. They will need practice using them, of course, but accurate recognition should be achieved quickly.

Keeping the symbols simple and consistent means the basic WFMA methods here can be learned quickly and efficiently. By keeping both the symbols unchanged and the discipline (syntax) consistent, anyone—jobholder, analyst, consultant, or manager—will be able to create and interpret workflow maps in one morning. These will not be perfect at first, but a little practice will correct that.

The “consistent way” of symbol use discussed above is the most critical part of this method, and thus is what I refer to as the *discipline* that must be used to apply WFMA successfully. This discipline must be learned and ingrained at the outset, and required in all workflow maps created with the symbols. This discipline is necessary to *preserve the simplicity* of the WFMA system, and thereby preserve its power to describe a vast array of workflows. At the same time, users should recognize that this discipline does not result in a single “correct” map for a given workflow or process; in fact, it makes it more likely that it will result in multiple possible “correct” maps, none of which will necessarily be the same.

How can this be so? It is simply that any workflow map is a joint product of the flow of work being mapped or described, and the knowledge and perceptual processes of the mapper, whether this is a jobholder or a specialist. There is no way to isolate these or their effects on each other, and so every map produced will be a function of what work is done interacting with the eye of the beholder. This also underscores the critical need for a disciplined approach to workflow mapping.

With an understanding of the importance of the discipline, we are ready to examine it in detail. There are seven basic components of the WFMA discipline.

Keep It Simple (Stupid)—KISS

The basic WFMA symbol set is based on a sound psychological principle that favors fewer rather than more symbols.¹⁴ Busy people have a hard

time keeping large numbers of information “chunks” (including symbols) straight in their minds, and overload is never far away.

Many people suggest that a method that makes workflow mapping easy is to simply write down, in step-by-step fashion, what is done in a process; some even contend that this is more than enough to provide the information a map would yield. I disagree with both of these arguments. Very few processes are so completely linear as to allow a clear, well-organized written description. What do you do, for example, each time you come to a branch or decision? How do you keep track of where you have been when you're on page 26 and don't quite remember which set of branches you took to get there? When you have to depart from a main process to deal with exceptions, how do you track through frequently complex exception-handling procedures and remember what was happening when you return to the main flow? Diagramming the workflow is preferable to verbal descriptions for all of these reasons, and more.

Use the Symbols Consistently and Correctly

A major benefit of simple symbols is that they have highly specific meanings. Using them as intended prevents confusion; be especially careful not to violate these rules:

- Use the symbols only for the purposes shown. For example, never show decisions in rectangular process blocks (a fairly common error, actually).
- Use only the symbols needed to map the process clearly and correctly. Even though there are only five symbols, not all of them are required for every process map, as we will see, and if they are not needed, they should not be used.
- Use brief, clear text within the symbols (as much as possible) to communicate what is happening at each step. If additional annotation is needed, this can be added in other locations on the page or through notes—this will be discussed in more detail later.
- Use only single arrows, always with the arrowhead shown to designate direction. Two-way arrows (pairs of arrows facing

in opposite directions) or double-headed arrows are never allowed.

- Multiple arrows can converge on an individual symbol (three paths might lead into one process block or diamond), but *with the exception of a diamond, only one arrow exits any symbol*. There are many cases and situations where multiple paths and subroutines in a workflow can converge on a single process or decision; *in no case, however, does a rectangular process block, circle, or document have more than one exit arrow*.
- Use only two mutually exclusive exit paths (arrows) from a decision diamond (unless there is truly no other choice, and that is extremely rare). Mutually exclusive paths mean that there can be no confusion over what each path means or why one would be selected rather than the other. This is an extremely important rule, and in some cases may make a workflow map segment larger than using more than two exit paths would suggest. Nevertheless, follow this rule.

Here are two examples: if we apply a first coat of paint to a wall and want to recoat it, we might show this as a decision (“first coat dry?”) with “Yes” and “No” exit paths. However, which one applies to “touch dry?” We might instead use three paths labeled “fully dry,” “touch dry,” and “wet.” To make sure that we do not mistakenly apply the second coat too quickly, we may need to add one or more additional conditions (decision diamonds) to the map to be unambiguous. Alternatively, we might ask in a first diamond whether the paint is “touch dry,” and in a second diamond whether at least eight hours have passed since the last of this coat was applied (surface-dry paint might still wrinkle if recoated too soon). A “yes” to both of these might actually make “fully dry” redundant. Even a simple matter like this may be represented by more than one “correct” map, but clarity is always served by two mutually-exclusive exit paths.

In a second (historical) case example, credit card customers were being reviewed at random, and for those with some questions about their credit histories, a possible outcome of the review was to limit their account or suspend it entirely; most, however, were at least renewed and frequently offered a higher line of credit. The student with whom I had this argument

took the position that there could be five outcomes from a single review decision. These were (1) raise the limit; (2) renew the account at present limit; (3) renew the account with lower limit; (4) suspend the account for a limited time; and (5) close the account. The student argued that all five exits paths could come from a single “Account review outcome?” diamond, and that there was no need to make this more complicated by the addition of a stepwise series of two-outcome decisions. (I insisted on a progressive series of decision diamonds with yes/no exit paths.)

In both cases, it outwardly sounds logical to use more than two exit paths from a decision, since this summarizes the actual outcome states; however, I disagree with the association of all possible outcomes with a single decision diamond. My contention is that if we do not insist on mutually exclusive paths at all branches in the workflow map, it is inevitable that ambiguous conditions for branching will creep into the map, and what seems absolutely clear to the person(s) creating the map will be completely confusing to other users. My personal experience with this issue bears this out—in another credit card case, one student argued for three decision exit paths, labeled “Yes,” “No,” and “Maybe.” I asked the student to define “Maybe,” and of course this was a nested set of further questions and conditions that had to be detailed in order to understand the full process. This confusion never arises when decisions or branching conditions are mapped in mutually exclusive terms. Moreover, the consistent use of two paths becomes habitual for both mappers and users, so that when a decision diamond is encountered, there will always be two paths to evaluate—no more, no less.

The most important reason for consistent use of two mutually exclusive paths, however, is that these require the creators of the map to fully articulate the logic for the branching that occurs at that point. This is often harder to do than we might think, because much of what gets done within many jobs is a process of evolutionary change and internalized (tacit) knowledge. Having to spell out all that is done, and why, involves more head scratching than is immediately apparent, but the payoff is that a great deal of very important tacit knowledge is uncovered this way.

- Put brief text *labels* on exit paths from a decision diamond. This enables the user to follow the map clearly, and helps

to keep the logic of the map visible and accessible to the user. Arrows between other symbols, and those entering decision diamonds, are simply flow paths and usually are not labeled; there are cases, however, where a label adds valuable information, and if so, then label the arrow.

- Avoid complex backward flows. These are acceptable on small areas of a map, but can be confusing on large areas. The problem is typically that in some part of the diagram, flow arrows will have to cross or “jump” each other. Use connector circles instead to break the map into coordinated sections and pieces, which can be followed more easily.
- Be careful of creating endless loops. This can happen for several reasons, but is often because of the incorrect use of double (or double-headed) arrows in a flow, or ambiguous criteria for a decision or branch. In more complex situations an endless loop may be created by a condition that sends one far upstream in a workflow, and you get caught in a loop without realizing it. Validation, discussed in detail in Chapter 3, will catch this.
- Avoid making a single page too dense, whether with large numbers of symbols and arrows, too much text, small symbols and fonts, or a combination of these. In an earlier lifetime in the printing industry, I learned to be concerned about the amount of “white space” on a page, meaning that there was enough unprinted area to make the printed parts accessible to the reader and the page visually appealing. While one always has to struggle with the tradeoff between detail and content per page versus overall document size, adequate white space is beneficial to workflow maps, as it is to most documents. Connectors help solve this problem.
- When using connectors, there is usually only one start location in a workflow map, but there may be multiple stopping points. These *terminal* points should be identified clearly. It is also best to match *continuation* connectors clearly, so that a point where a loop or another part of a process may rejoin an activity flow from several other locations cannot

be confused—all “C” connectors that are exits from part of a flow should connect to only the single correct “C” reentry point, and so on for all connectors.

Validate the Map

We will hear this term quite often in the following chapters, and it connotes a very important step in WFMA. Basically, validating the map means that we have a designated, nonjudgmental individual who can talk their way through the map with the jobholder, and certify that what has been mapped is an accurate description of the process. Unlike many process mapping programs or systems, the Kmetz method does not assume that anyone really knows what the processes in a complex organization are. This may sound naive, but it is actually quite significant—if we “assume nothing” we have a very good chance of finding out what actually happens in organizational processes. What we do with this information remains to be articulated (and it may be that we will leave things as they are), but any kind of process improvement must begin with an accurate portrayal of what the process is now. We will hear of this again throughout this book.

A “Correct” Process Map Captures an Unbroken Sequence of Processes and Branches from Beginning to End

In short, from the start of any process map, it must be possible to follow the sequence of activities to any end point(s) with no gaps in the sequence and with no stopping point(s) other than completion of the work cycle. In most cases, we want to map the full “work cycle” in the system, showing the input (or how we get it), the transformation process (probably in some detail), and identify the output (or the stopping point).

The stopping point in the workflow being mapped may not always be the end of the full cycle. For example, it is common in many maps to trace the workflow to the customer, and sometimes into the customer’s application of the output. However, if the map you create is only for one department of several on that workflow, the “end point” you select is the point at which the department completes its part of the process; the

remainder of the process is completed elsewhere. The same could be true for the “starting point”—it may be where the work in this department begins.

When the symbol set is properly used, there will be no places where the map stops if there are still parts of the process remaining to be completed; there will be no gaps; and all branches from diamonds will flow to a completion point or will merge back into the flow. It bears repeating that no two analysts will be likely to do the “correct” map the same way.

Basic mapping conventions

Figures 1.2–1.5 illustrate the major conventions for constructing workflow maps. Figure 1.2 illustrates a very simple workflow map, with




Figure 1.2 Single-cycle process flow

symbols of four processes in sequence. Such a simple map requires only four process blocks, a “start” and “stop” point in circles, and the arrows to indicate the flow.

Figure 1.3 illustrates a workflow map where a branch occurs by selection of the process path (main or alternate) to follow. If the main path is followed Processes 1 and 2 are used, but for the alternate path Processes A, B, and 2 are used, and Process 1 is skipped. The results of Process B are also documented. This map uses all five symbols in the set. Annotations of various parts of the process can be added if they add clarity, as shown, and this will be discussed in more detail in the next section of this chapter.

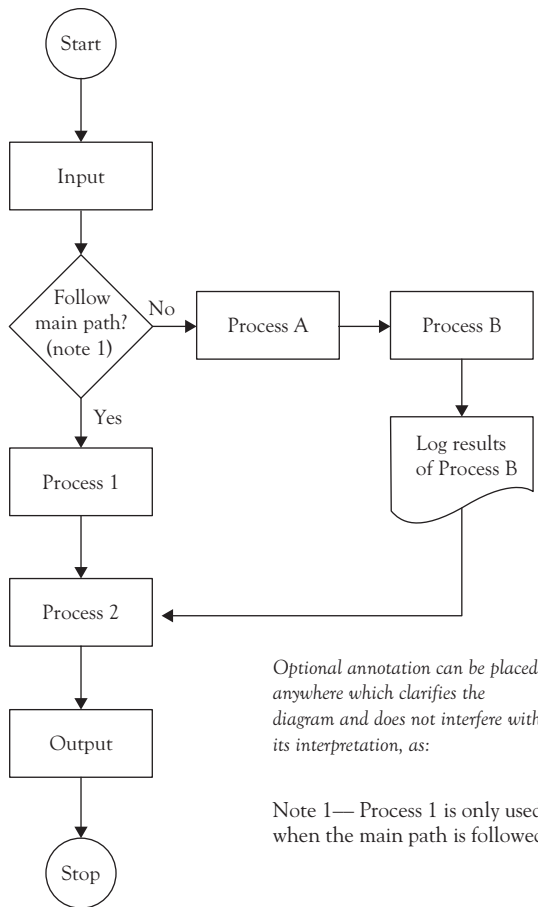


Figure 1.3 Branching process flow

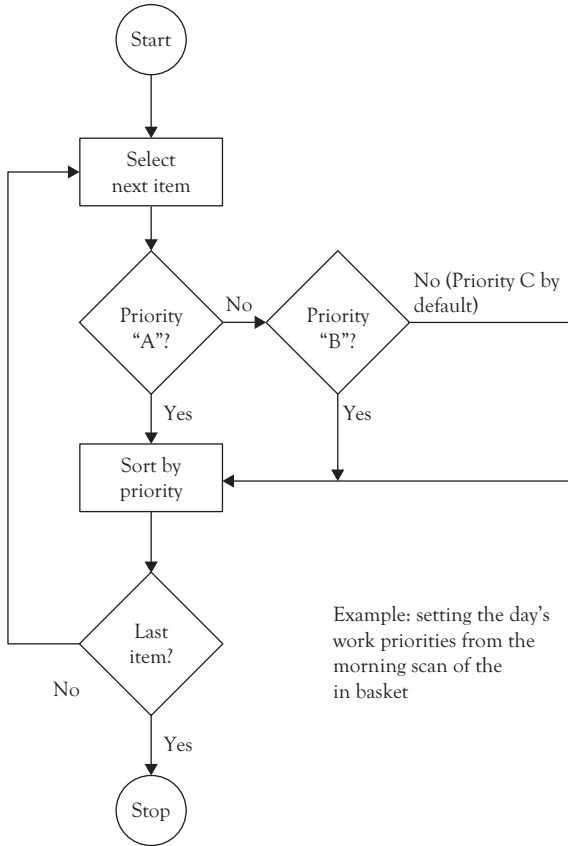


Figure 1.4 Multiple-cycle (“looping”) process flow

Figure 1.4 illustrates a repetitive loop in a workflow map. While I generally discourage backward flow loops in maps, this is a case where it clearly makes more sense to keep the map on one page rather than force some kind of stylistic compromise to avoid such a loop. Proliferation of such backward loops, however, can reduce the clarity of a map, and one excellent application of connector circles is to use them to reduce the extent to which backward (and especially crossing) flows are needed.

Figure 1.5 illustrates how processes can be viewed at differing levels of detail. On the left side of the figure the four high-level steps necessary for completion of a transaction in this financial institution are shown. The right side shows the first of these four steps, the taking of the customer order, in more detail. Here we can see the three different ways that orders can be received (telephone, fax, or via dedicated computer link), and the

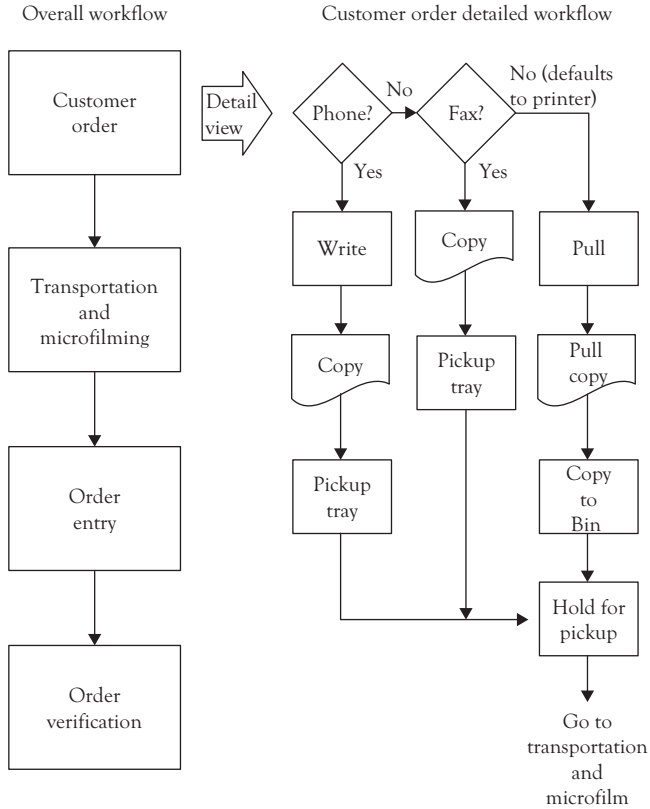


Figure 1.5 Two levels of detail in WFMA (“drill-down”)

additional activities needed to prepare them for the next major part of the process.

Using Annotation in WFMA Maps

The general rule for the use of text in WFMA is that we want to keep as much of the necessary text and information about a process flow inside the symbols as possible, or closely associated with the symbols. This is part of the WFMA discipline and a payoff of it—when the user follows the map, there will be an automatic association between the symbols and the information in them. Together, these fully define the workflow we are mapping.

Occasions arise where annotation apart from the symbols is necessary, and there is a wide variety of such circumstances. In making a workflow map for training new personnel, we might note to the trainee that things

done in one step may be referred to another office under some conditions; that there may be supplier or other external relationships that could affect a part of the flow; that random samples of work may be taken for quality assurance; and the like.

Whatever the situation requiring annotation, there are several rules of thumb that should be followed to the greatest extent possible:

1. Try to keep annotation within the area of the relevant symbols. This simply associates the annotation with the process steps it refers to, and prevents the user from having to go somewhere else in the map or another document to get the desired information.
2. Do not mix annotation directly with process information in the mapping symbols. Annotation implies supplemental information, and it is best to keep it that way. When annotation is used, it should be apparent that it is annotation.
3. Avoid using straight arrows to associate annotation with symbols, since these are used as part of the maps themselves; connective symbols for annotation should generally be avoided. If there is a need to associate comments and annotation with specific symbols in a map, the famous engineer's "curving arrow" (see the "system boundary" in Figure 4.4) may be appropriate in some cases, but too many of them clutter the map and make misinterpretation of it more likely rather than less. My preference is to use relatively standard typographic tools such as asterisks, parenthetical note references, footnotes, and endnotes to associate supplemental information with the map and its symbols. Footnotes are often a very effective way to make this association, and are automatically fit onto a page by word processors. This is a convenient association for making printed copies of a map, and is also useful for electronic documents, since relatively large displays with high resolution have become common.
4. Use annotation appropriate to the level of detail in the workflow map. Annotation is sometimes used to address matters at levels of detail much higher or lower than the level of the map, and this can be confusing and distracting. If one is looking at the high-level (left side) view of Figure 1.5, for example, it might be appropriate to note why the sequence is what it was at that time. (This is an historical

case where time is money, and these are high-value transactions; thus, nothing was done until the order was logged and officially in the system. Needless to say, this was automated a long time ago.) It would not be appropriate to explain at the high level that pulling a copy from the computer printer meant separating five-part forms with carbon sheets between them (as was the case then); this would be appropriate to the detail view on the right side, or perhaps even a more detailed drill-down map.

Cases like these are where the linking capabilities of many drawing and workflow-mapping software packages can be used to excellent advantage. Figure 1.4 could be separated into two parts (the high level and detail views), and each could be annotated as needed. A user wanting to have easy access to information at different levels would only need to select the desired link to see the alternate view. Some packages refer to these links as “layers,” but that term has multiple meanings.

5. Hyperlinks, the familiar blue-colored underlined connections we use to browse the Internet and link documents, are easily created with nearly every major software package, and can be a very useful form of annotation. As an example, suppose that a particular step in the administration of an experimental drug to a patient requires the consultation of an attending physician; specific attention to several important patient health details as well as other matters; and that a form must be signed by the appropriate physician upon giving a dose. Trying to annotate this much supporting information at a single point in a workflow map might be dysfunctional given the length of the procedures and forms to be completed. However, this requirement could be met with annotation such as, “Only attending physicians are allowed to give this drug and they must document their actions every time.” The hyperlink associated with “document” may connect to a lengthy and complicated research protocol; the workflow map, on the other hand, can immediately and clearly go on its way.

Somewhat more time is needed to learn the rules, but these should also quickly become part of the disciplined approach to WFMA. Learning

the rules is best done with hands-on practice, beginning with something relatively simple and then progressing to a more complex process. *Nearly anyone should become an “expert” in WFMA workflow mapping in no more than a few sessions.* Validating the map and doing the rest of what WFMA entails takes a bit more time, but the basics should be successfully acquired very quickly.

Again, what I learned in my early experiences with WFMA is that the information conveyed by the map, through the brief descriptors used in conjunction with the arrangement of symbols that graphically describes how something gets done, is the only thing that matters. If the mapping graphics and rules get in the way, they obscure what they should reveal, and they distract from the analysis rather than support it. In no case is this good, and in situations where there may be political or other opponents of what is being done with WFMA, these distractions provide ammunition for them to discredit the entire program and its objectives. The power of the Kmetz method is its simplicity and consistency.

Making Bread

Let me use an example of a common activity to illustrate these basic conventions. The example is making bread, one of the most fundamental forms of food preparation in human society. I have included my own bread recipe at the end of this chapter for those who might want to know more about it, and I can attest that I have never had a failure with this recipe. Many people are intimidated by baking (the expression, “cooking is art, but baking is chemistry” is true in my experience), but this chemistry always works and tastes wonderful. Figure 1.6a illustrates a high-level view of the process, where everything done to combine water, flour, and yeast to make bread is compressed into a single process. While technically correct, it provides very little information beyond identification of the process itself. In cases where we either do not need details about the process or can assume an expert audience who already knows the details, we may need nothing more than this level of information.

Most users of a workflow map will want more detail than this map provides. Figure 1.6b is a “drill-down” of the one process in Figure 1.6a, revealing a series of four principal processes—mixing the dough from its




Figure 1.6a High-level view of the process of making bread




Figure 1.6b Four principal processes in making bread

basic ingredients, allowing the dough to rise from the action of the yeast, forming the dough into sections appropriate to the type of final product being baked (large loaves, small loaves, individual rolls, etc.), and then baking the formed loaves.

The next level of detail is seen in Figure 1.6c, where we now see that the baker has a choice regarding whether the bread dough will go through more than one rising, with the reason for this decision explained in the annotation to the “Yes” exit path from the diamond. We also see more detail in other parts of the process, where mixing ingredients is now seen to be two steps; this is a drill-down of the single process in Figure 1.6b, and the entire figure is a further drill-down of Figure 1.6a.

The highest level of detail shown is in Figure 1.6d, which is a further drill-down of all process steps and conditions shown in Figure 1.6c.

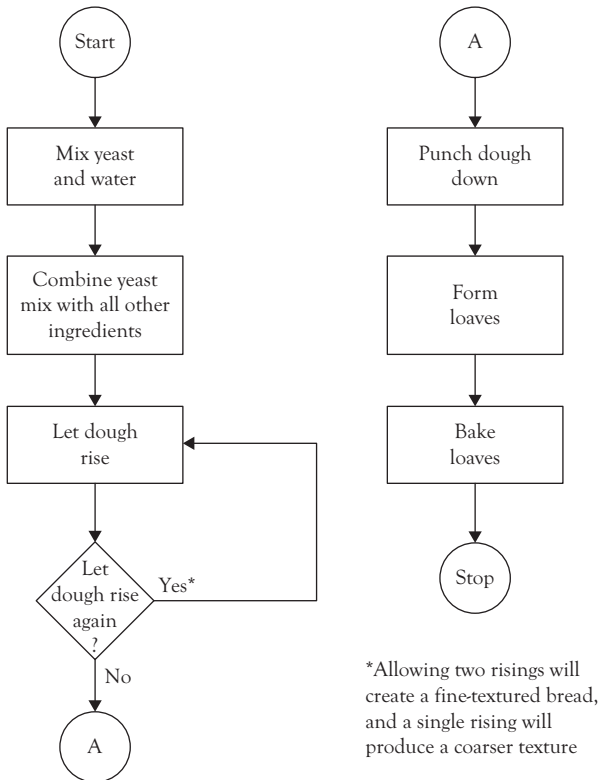


Figure 1.6c Making bread showing election of multiple risings if desired

Figure 1.6d is a sufficiently detailed map of the process of making bread that it could be followed as a recipe, and is annotated to show where individual process steps correspond to the written recipe shown at the end of the chapter. But even more detail could be added to this map, showing things such as the tacit knowledge of what “mix well” means (if using an electric mixer, the dough should release from the sides of the bowl with no residue sticking to the bowl, and the dough should have a

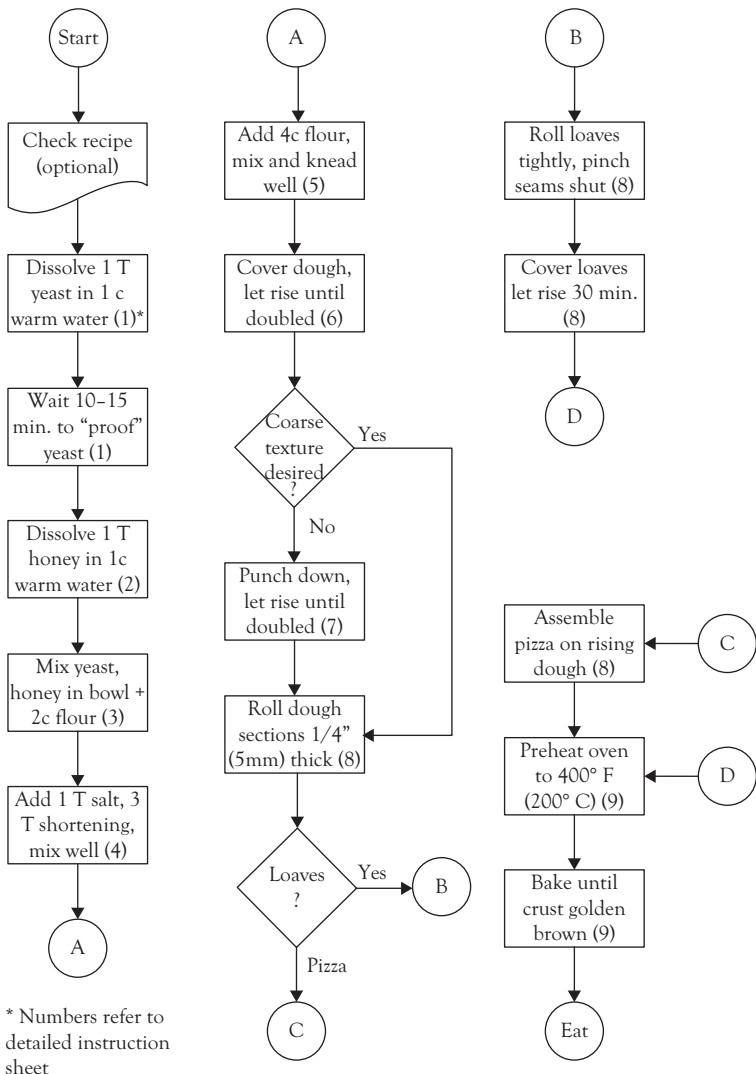


Figure 1.6d Making bread detailed workflow map

high degree of elasticity; other tests apply if the dough is mixed by hand, but it should always be elastic when it is well mixed and kneaded), or how to assemble a pizza if that is how the dough is being used. Other forms of tacit knowledge affect how we form loaves—some bakers prefer the “country” style breads made by simply dividing the dough and forming loaves from the sections rather than rolling it with a rolling pin, as I do.

Since “baking is chemistry,” there will be less variation in a process map for making bread than there would be for making soup. The action of yeast creates carbon dioxide that makes the dough rise and expand, and this requires time and temperature control; it also dictates upper and lower limits on how long a baker can wait to perform steps in the process. But even in this case tacit knowledge is extremely important and has a great impact on the recipe. I spent many years trying bread recipes and experimenting with variations until I felt confident I understood the chemistry and could change the type and proportion of ingredients to my liking. In many ways, it is accurate to regard the entire recipe as tacit knowledge, even though the chemistry of baking imposes boundaries on what we use to make bread. We should note that there is only one document symbol in these maps, but one could be added to any of the last three figures for a baker to make notes on their own recipe or modifications to mine, as I did. My own recipe documents a type of filled Italian bread called “pinolata,” where the dough is rolled like a pizza, then covered evenly with the filling, and rolled up (pinch the seams to seal them) to form a loaf. This is not shown in Figure 1.6d, and would be part of a more elaborate subsequent map.

In closing this culinary example, we should note that what has been illustrated by it applies to nearly every process we can imagine. We can start at a very high level and progressively drill down into greater levels of detail. As we do so we will likely use more of the symbols, but note that only at the fourth level of detail were all of the symbols used. It is equally worth noting that all of the details we want to show can be mapped using the basic symbol set and the discipline—this is a highly robust mapping technique.

Mapping Parallel Processes

A question that often arises is how to map “parallel” processes, where a number of discrete processes are carried out and all seem to have begun

from a single trigger event. The mapping issue this raises is that it would appear to violate the rule that a process block never has more than one exit path, requiring a single process block to immediately trigger a number of other process blocks. This incorrect practice (which is not that unusual to actually see in many maps) is shown in Figure 1.7a, along with the correct mapping method in Figure 1.7b. I have encountered this apparent dilemma in a number of situations, and it is only an “apparent” dilemma. Let me explain.

First, these apparent parallel processes occur in two forms, external and internal to the organization. Externally, there are processes that are simultaneous and not controlled by the organization, as in the case of a home-security company. The security firm monitors clients’ home systems for evidence of (1) break-ins, (2) fire, (3) flood, (4) medical emergencies, (5) earthquake, and (6) loss of electrical power, and has a predefined response for each of these. In addition to the home systems, other sources of information are scanned, including (7) firm personnel in satellite locations whose observations may initiate responses, and (8) local emergency responders such as fire companies, which are monitored for information that may affect customers. Finally, (9) news media reports of

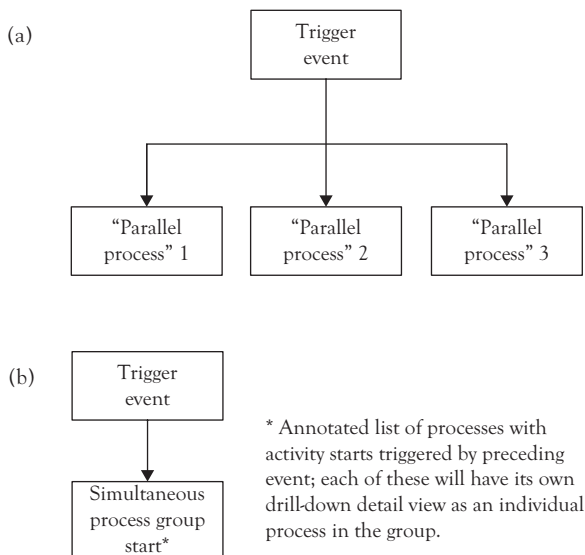


Figure 1.7 *Incorrect (a) and correct (b) mapping of “parallel” processes*

emergencies that could affect customer residences are also monitored for potential response.

Any of these nine information sources could provide a cue that would trigger a response by the security service. This cue is received through rapid sequential surveillance of these sources (like an emergency radio scanner), and the cue triggers one or more predefined response processes. This is in fact rapid execution of a serial process, and is not truly parallel at all. If the process is automated, the program flowchart would reflect this sequence of steps; if not, and human intervention is required, execution is still a sequence, albeit probably slower. For example, for a worker to be able to stop an assembly line in an automobile plant when a quality problem is detected requires a great deal of prior work, involving the creation of company policies that allow such a step, installation and testing of the “kill switch” that would stop the line, training the worker to detect the problem and understand the correct response, and installing corrective processes to solve the problem and resume production. There is a long series of steps required to establish what outwardly appears to be a process that is “parallel” to the normal production process.

Internally, a trigger event sometimes sets a number of processes in motion, almost at once, and can again seem to be a parallel process. A type of production I am personally quite familiar with is the writing of a proposal for a major weapon system, when the official release of a Request for Proposal (RFP) from the Department of Defense triggers rapid mobilization of a team to write the proposal, as shown in Figure 1.8. Teams will be engaged to write a document covering basic system engineering, manufacturing engineering, logistics support, maintenance, personnel training, technical documentation, and a myriad of other requirements. All of these seem to start at once on release of the RFP, but in fact, there are many serial processes at work well in advance of that trigger event. Potential contractors devote much energy to gathering intelligence on what the customer might require, what competitors might offer, and above all, to assembling the teams of people to write the various sections of the proposal. It is literally the case that for most of these programs, years of preparation have gone into being ready to launch the actual proposal effort, and the latter is only apparently a set of “parallel” processes. Many companies in many industries live and die by the same kind of bid-and-proposal business, and the seemingly automatic, parallel processes in

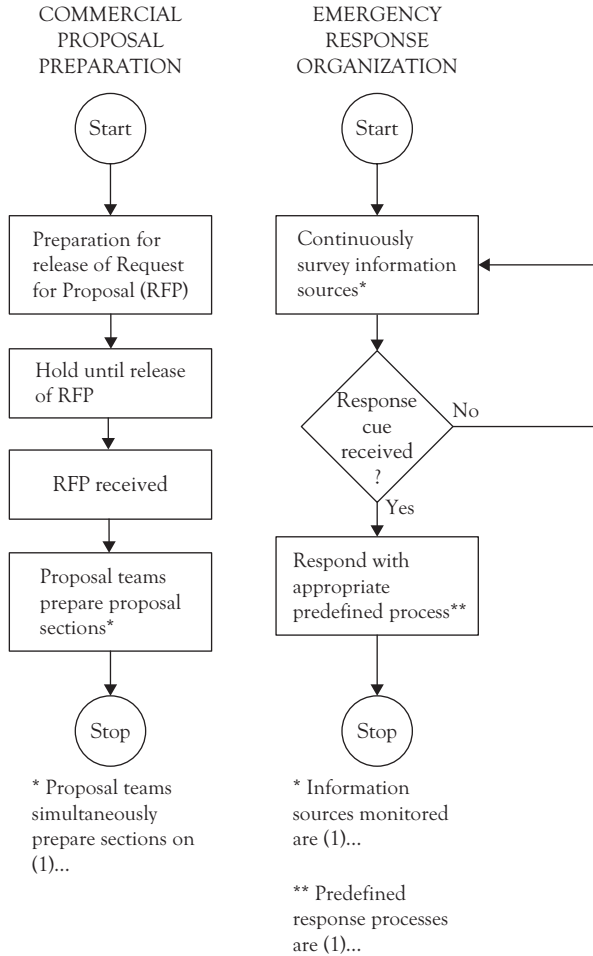


Figure 1.8 *Two illustrations of organizational response to external events with simultaneous internal processes*

their responses are actually the culmination of years of serial preparation, and embody a huge amount of tacit knowledge as well.

Rather than referring to these as “parallel” processes, I prefer the term *simultaneous* processes. I am prepared to argue that “parallel” processes only exist with respect to a point in time, when the typically lengthy and voluminous series of preparatory steps that enables them have been completed. In the case of internal simultaneous processes, most of these have been predefined through design, experimentation, and experience, and often embody tacit knowledge that has often been formalized because experience has showed it worked.

Mapping these is often less of a challenge than may seem to be the case. Consider the emergency services firm—only two process blocks with a decision diamond are required to show how it works. The first process block is “Continuously survey information sources,” connected to a diamond querying, “Response cue received?” If no, that branch loops back to the survey process block; if yes, it connects to the second process block, “Respond with appropriate predefined process.” The first process block would be annotated to indicate the kinds of information sources surveyed (home monitors, scanners, media), and the second annotated with a list of the different processes for each of the six emergencies to which the company responds. A similar technique can be used for the launch of proposal-writing teams. The first process block would identify the kinds of pre-launch preparation needed to be ready for the release of the RFP, which is the query in the diamond. The second block is “Hold for RFP release,” to indicate that preparations were complete and a delay in launching the effort was for reasons beyond company control. The third shows the release of the RFP, and the final block is “Proposal teams prepare proposal sections,” annotated with a list of those teams and sections, as shown in Figure 1.8.

This is a case where hyperlinking and drill-down support from drawing software can be very helpful. For each of the first process blocks, details on individual predefined information source monitoring for emergency services, and details on specific steps needed to be ready to launch a proposal effort, can be provided; these can go into significant further detail as needed to document complex processes. The second process block drill-downs can detail the predefined emergency responses (an excellent way to assemble training and review information on these processes), and the detail needed for each of the proposal sections, all of which is highly specific to the proposing firm or organization. It is not at all uncommon for many of these kinds of predefined processes to depend on legal, procedural, or other documentation, all of which can also be hyperlinked to the appropriate processes at the appropriate level of detail.

Types of Workflow Maps

What types of workflow maps can we produce with the WFMA symbol set and rules? In one sense, there is only one kind of map, because

the disciplined use of the symbol set can be applied completely under the user's control. Different kinds of maps can be produced, however, depending on the focus of the user. There are basically two questions to consider—do we want to focus on a process or do we have a different objective, and do we want to be comprehensive or map a segment of a larger process?

First, we can obviously map specific types of workflows, without respect to whether the map captures the entire content of any particular job on the flow. This is the most common type of map and implicitly has been the type that has been the focus of our discussion thus far. To treat a patient in an emergency room requires the services of a receptionist (if the patient is self-transported), an orderly, one or more nurses, and an ER physician, at the minimum. Some or all of these people may have things to do that go beyond the emergency patient-care workflow, but our concern here is with the handling of the ER patient. The workflow map for a cardiac emergency will not be the same as for a splinter in the left big toe. Using maps to document processes for quality certification is still another example. A map prepared to “document what you do” for ISO 9000 certification will not be the same as a map prepared to train a new worker to do that process—the latter will be much more detailed. ISO 9000 maps are typically not exhaustive descriptions of the individual jobs involved in them, either.

Second, maps can be constructed with a focus on *tasks* or *jobs*, as well as processes. That is, we can create workflow maps of individual *tasks* within a larger job; or we can map the entire *job*, which will always have at least one task, but nearly always has more. We might want to map the work done in or through a *group*, whether the group is a team, part of a unit, or a whole unit, and could also include work done by contractors and others who contribute to the output but are not members of the organization. At the upper extreme of group size, we can even map an entire organization. No matter what the scale, the objective of this type of map is to capture the “technical” (material) content and knowledge in the relevant system, where that system is not a “workflow” as we have been using that term. I have seen this kind of application used by several clients to create graphic job descriptions and supporting training documents, among other things.

We could also map the flow of *materials* or specific items of *information*, without initial regard to how these are used. The latter types of flows are often useful for investigatory purposes such as process improvement, security, waste reduction, cost control, and the like; in some cases, we just want to see what happens to a piece of paper or a bit of information for one reason or another.

These categories are not necessarily mutually exclusive. It is entirely possible for a process to be contained within a single job. The main point to be aware of in considering types of workflow maps is that whether they are comprehensive of either jobs or workflows is a matter of what the user wants. Many workflows only require a segment of the individual jobs they pass through, and conversely, many jobs contribute partially to several workflows. WFMA is flexible enough to accommodate all of these, but it is necessary for users to be clear about the content they want in a workflow map.

Flow Relationships

In Chapter 5, we will discuss two properties of flow relationships: convergence and synchronicity. For present purposes, it is sufficient to note that these distinctions in the nature and timing of workflow elements do not necessarily require separate types of maps, but that these may be disaggregated into distinctive parts. If we are mapping the *process* of “cash management,” meaning depositing and disbursing cash along with periodic reconciliation of the cash account, this will comprise several flows of activities with related exception-handling processes, much like any other. These would indicate that one person prepares and deposits cash receipts, and another will be the one to write checks and disburse cash. At certain dates or on certain random intervals, reviewers who are not part of this regular workflow will verify the overall transactions and the resulting balances (i.e., an audit). The overall process is continuous and cyclical, even though it is done by a group of individuals with responsibility for discrete activities in the process. Divergence is a matter of where the activities that make up the workflow are done, and timing is a matter of when these activities get done; but all of them are part of one process—cash management.

The relationship between the overall process and its detailed parts would be similar to Figure 1.5.

Much of the purpose of WFMA is to describe, and in some cases to discover, exactly this kind of deeper understanding. The understanding begins with a *valid* map of what is actually being done, and this is sometimes a very challenging map to create, for reasons mentioned earlier and discussed more fully in Chapter 3. Having this map, however, enables properties like the functional or dysfunctional nature of the workflow to be seen as well as matters of synchronization or convergence of actions and decisions.

Additional Examples of Workflow Maps

Several additional examples of workflow maps are shown in this section to illustrate several points. First, Figures 1.9 and 1.10 are both relatively complex transactions from financial services companies. Names and details have been changed to protect their identities, but these figures illustrate processes involving typical client services from such institutions. Both of these were prepared using Microsoft *Visio* software, and show how all symbols are used in a map and how they can show considerable detail and yet be arranged to fit on a single page. In truth, Figure 1.10 is about as crowded as I would allow a page to become, but there is always a benefit from being able to see a process map in one piece, as opposed to having to change pages.

Figures 1.11–1.13 are two-part figures, where the first part shows a workflow map as prepared by a client organization, and the second part shows that map constructed using the Kmetz method. Figure 1.11 is a prepared-foods firm, and this figure illustrates the packaging process for a Laptime TV Dinner, as done by the client and as redone. Figure 1.12 shows a proposed new book selection system at the University of Delaware, again in client form and as redone by the Kmetz method.

Figure 1.13 also shows the conversion of a map done in Data Flow Diagram (DFD) format, which the client used. This was an interesting case in that the client wanted to capture both the process and the interplay of parts of their information system, and believed that the best way to do this was to focus on the data. As can be seen, the actual process flow

is much simpler and clearer in part (b) rather than in the DFD format in part (a); this process is really a straightforward linear sequence of steps, as shown on the left in part (b), with the subprocesses needed to retrieve and export data shown in the two detail views to the right of the main flow. This is the type of map where software like *Visio* can be very helpful—the detail views could be part of the file where the main flow is stored, and by hyperlinking the words “See Detail *n*” we can jump to those details when needed, and otherwise focus on the main process flow.

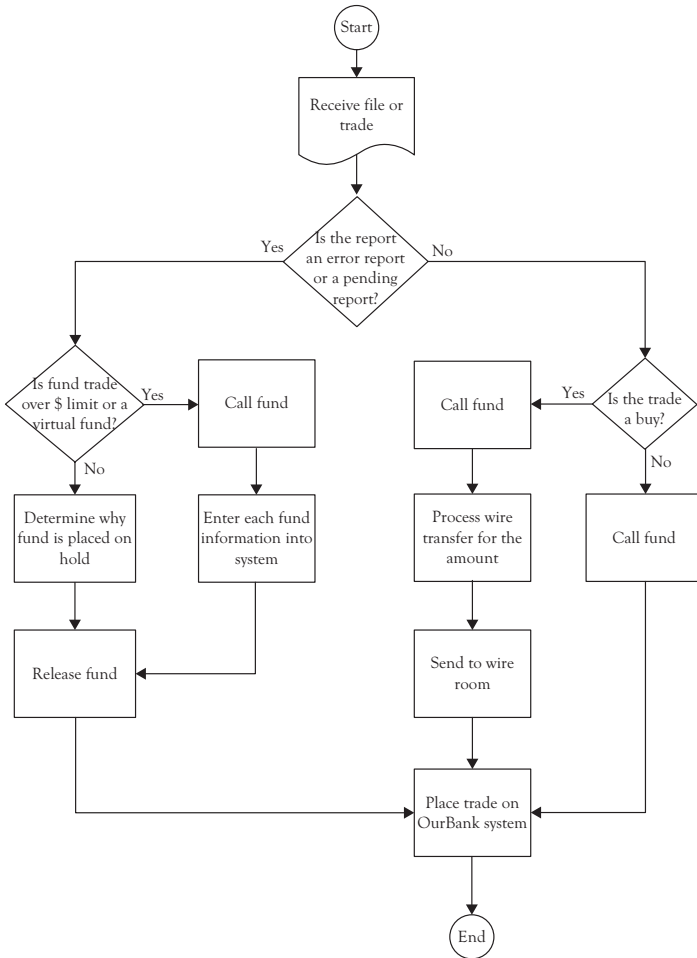


Figure 1.9 Processing a fund trade

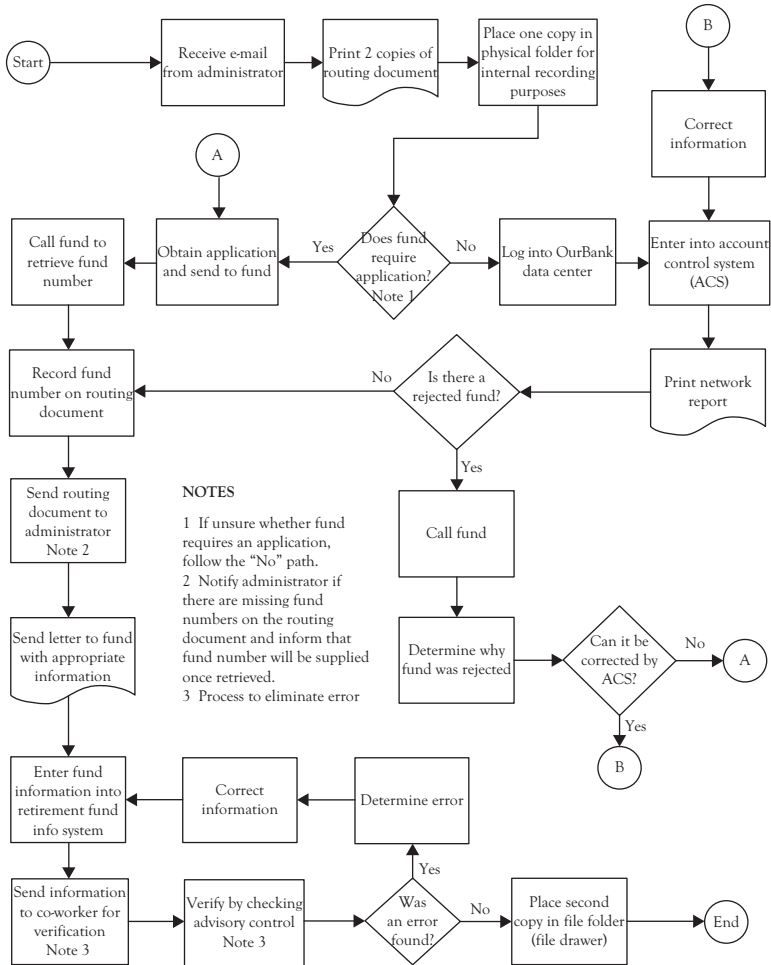


Figure 1.10 Opening a corporate retirement account

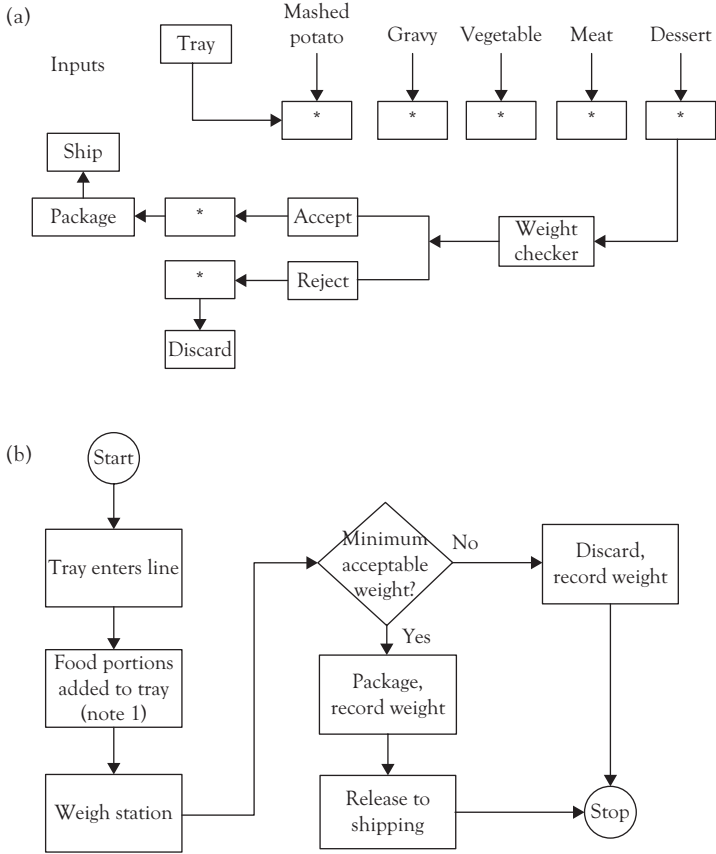
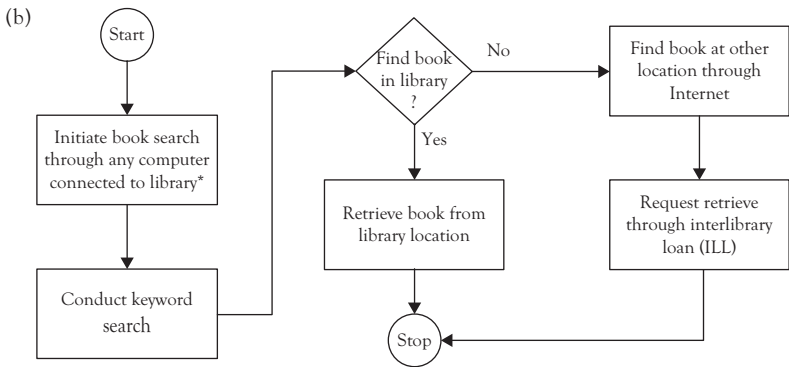
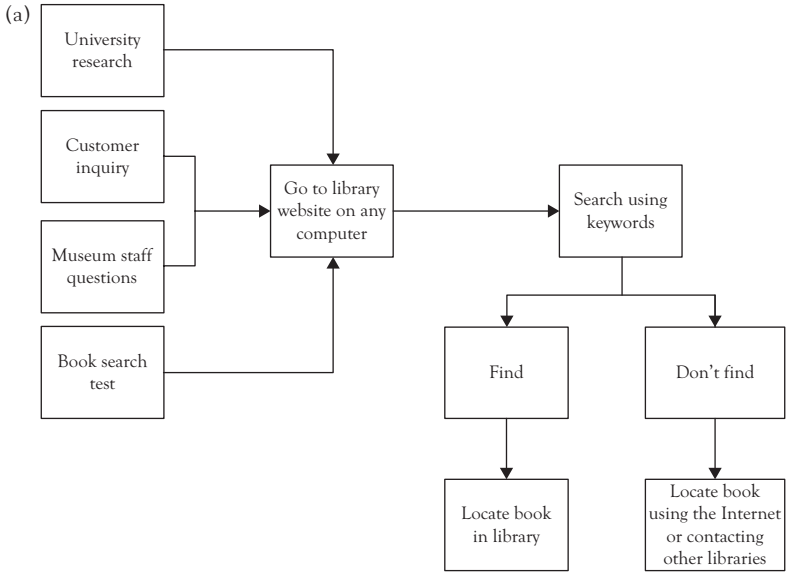


Figure 1.11 (a) Client process map for Laptime TV Dinners
 (b) Kmetz method process map for Laptime TV Dinners



*Book search may be initiated for (1) university research, (2) customer inquiry, (3) museum staff questions, or (4) test purposes

**Figure 1.12 (a) Client process map for new book selection
(b) Kmetz method process map for new book selection**

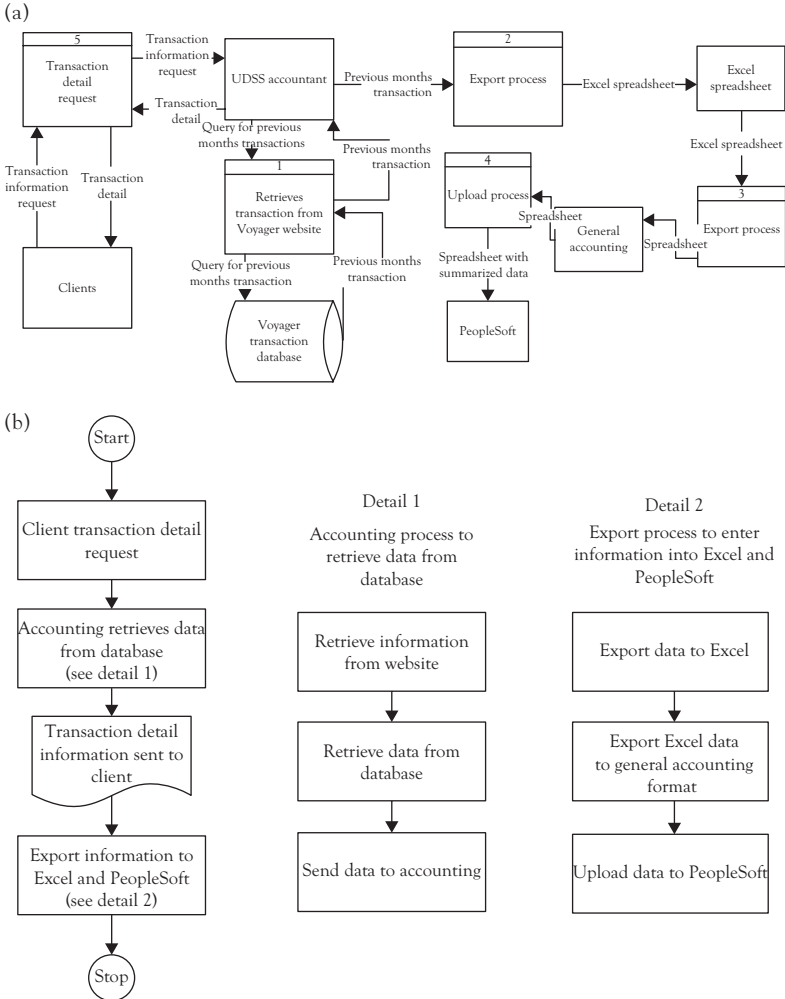


Figure 1.13 (a) Client Data Flow Diagram map for travel accounting (b) Kmetz method process map for travel accounting

Summary

This chapter, in many ways, is the major “nuts and bolts” of the WFMA method I have developed, in that it shows how to develop the workflow maps that are the core of the system. By design, the method is simple

and disciplined; however, “simple” is often misperceived as “simplistic,” and nothing could be further from the truth—“simple” is “powerful,” but only so long as the simplicity is preserved.

The idea behind this method is straightforward—both mappers and users of map information can quickly learn the symbols and rules for their use, so that either can create, validate, modify, or question any map produced with this system. The maps can be consolidated and aggregated easily, whether in soft (electronic) or hard copy. Equally important, the maps can be changed quickly when processes change. As we will see in the next chapter, valid maps create a basis for measurement of process outcomes, and can be linked to important company and organizational metrics. Finally, the method is highly robust, and can be applied to any organization, no matter what it does.

Exercises

Unlike other chapters in this book, it seemed that several practice exercises would be helpful, since one of the principal objectives is to learn how to use the WFMA symbol set and its related discipline. I have included two groups of short exercises here to allow readers to apply what they have read in this chapter to different kinds of processes. The first two are exercises in doing personal things we all have to do, and which most of us probably do “on autopilot.” But try to be as complete in your mapping of both processes as you can.

The second set are situations more typical of what we might find in many workplace and office settings. Both of these have objectives relevant to the success of the hypothetical company in which they are done, and are not as likely to be personalized as the first group.

A hint—in both exercise sets, try to account for the parts of the processes that are purely information as well as those which may be movement of physical materials, and don’t be surprised if the information-intensive parts of the map are bigger and more involved than those tracking the material.

Individual Exercise Group I

On a separate piece of paper, draw a workflow diagram of:

1. How you mow the lawn, or (for those who don't have one), how you clean your residence; and
2. How you do your laundry.

(Additional exercises can come from any common household activity—load the dishwasher; wash the dog/car/boat/windows, etc.)

Individual Exercise Group II

On a separate piece of paper, diagram the workflows for the following two problems.

1. Your job is to periodically review customer credit card accounts in the bank where you work. Periodically, you retrieve accounts, review them, and must decide to either renew them, suspend them, or terminate them. If you suspend an account, you must change its status so that it should not be used, and then notify the customer of that change. If you terminate the account, you must notify merchants that the account is terminated, and then notify the customer. Naturally, if the account is renewed, you want to notify the customer of that, too.
2. You work in a mail-order catalog house where you receive orders from your customers by telephone. If a customer calls, you need to know if this is an existing customer with whom you have done business in the past; previous customers have a PIN. If so, you retrieve a customer number from your terminal, take the order and complete billing information, transmit the order to shipping, and then mail the documentation to the customer. If this is a new customer, and someone who is using a credit card, you must obtain the credit card data before completing the billing information. As with existing customers, you then transmit the order to shipping,

and then mail the documentation to the customer. If the customer doesn't have a credit card (and this sometimes happens), he or she will have to go through a credit review process that you pass on to another department; if approved for credit, billing information will be completed there and the order returned to you to finish as with any other customer.

(Just for fun, have several people do these and compare results—remember that no map is “wrong” so long as it accurately describes what that person does.)

John's Bread
2 cups water 105°F to 115°F
1 pkg. or 1 Tbsp dry yeast
6 to 6½ cups flour (“bread” flour is best, not “all-purpose” flour)
1 Tbsp honey or sugar (I prefer honey)
1 Tbsp salt
3 Tbsp shortening (any light oil is good—I use safflower oil)
Directions:
1. Dissolve yeast in one cup of the warm water, and let stand for 10 to 15 minutes. This is “proofing” the yeast, and it will develop a foamy head—use all of this in the dough.
2. Put yeast solution in large bowl or mixer (a good mixer can do the whole job, including the kneading). Dissolve honey in second cup of warm water and add to yeast solution in bowl.
3. Add 2 cups of flour to water and mix thoroughly. (Hint: I use 2 cups of “stone-ground” (from Newark Natural Foods, Newark, DE, bin 123) or coarse grind whole wheat flour to give more flavor, and 4 cups white flour for the rest; I use King Arthur unbleached white bread flour, and recommend it. You can use all purpose flour if you prefer.)
4. Add salt and shortening and mix thoroughly.
5. Add rest of flour. Mix in until moistened, then knead by hand or work in mixer at low speed until dough is smooth and elastic. I find that 10 minutes of mixer kneading (or 15 minutes by hand) is usually enough.
6. Cover dough bowl with plastic wrap and let dough double in bulk. (Hint: wrap a large towel around the bowl to keep it warm; rising will take about 45 minutes).
7. Punch down. At this point the dough can be used, or you can let it rise again to get a smoother, more chewy consistency. Punch down again after second rising.

<p>8. After rising, cut and shape into loaves (makes two generous “French” loaves or four “baguettes”) in whatever form you like. To make “French” loaves, I roll each half of the dough out into a rectangle $\frac{3}{8}$” to $\frac{1}{4}$” thick, and then roll this into a long loaf. Let loaves rise for 30 to 45 minutes covered with plastic wrap to prevent drying out. Pinch the seams in rolls or loaves shut. The dough can be used for nearly any purpose, from bread to rolls to pizza dough (pizza can be assembled on the dough while it rises if ingredients are near room temperature—not too hot or cold).</p>
<p>9. Preheat oven to 400°F. Cut $\frac{1}{4}$” deep diagonal slashes into top of loaves (or cut down middle) to allow steam to escape. Bake 35 to 45 minutes. If top and bottom shelves are used, switch loaves from top to bottom halfway through baking.</p>
<p>Bread is done when tapping the bottom of the loaf with your finger makes a hollow “thump” sound. Crust should be nicely browned, not burned, not pale.</p>
<p><i>Tip:</i> Put a pan of water on the bottom of the oven to steam the loaves while they bake. This gives an extra-crunchy crust.</p>
<p>You can also use the dough for pizza with just a partial first rising (30 minutes or so). It will have a coarser texture than with multiple risings, but the taste is fine. You can also use the dough for filled breads—roll it as if making regular loaves, but then spread a room-temperature filling on it (try browned and drained ground beef, onions, and chopped Greek calamata olives, for one), roll it up and pinch the roll shut, and bake as usual.</p>
<p><i>ENJOY! I’ve never had this recipe fail.</i></p>

CHAPTER 2

WFMA Data Collection and Analysis

Introduction

This chapter discusses application of workflow mapping and analysis (WFMA) to the measurement of workflow properties after the completion and validation of a workflow map, and is the second of two “nuts and bolts” chapters on WFMA. Obviously, the most common measures of interest are process characteristics like cost, time consumption, error rates and locations, and similar things, usually with the objective of performance improvement. The general approach to data collection and the use of business metrics in this chapter will be familiar to those who have previous exposure to approaches such as DMAIC (Define, Measure, Analyze, Improve, Control), including Six Sigma programs and other quality and performance-improvement procedures.¹ While these methods all have advocates and critics, nearly all of them agree on the use of metrics, measurement, and analysis of data to select improvement targets and determine results, and as we will see, that approach is often fundamental to what we do with the results of workflow mapping.

There are two broad issues that emerge in the application of workflow maps for analysis and change. The first is the environment in which WFMA is done, which considers the procedure for WFMA projects and some of the issues involved in using a workflow map as a data-collection tool. The collection of data from a workflow map may serve various objectives, one of which is often some type of quantitative analysis as a precursor to process change, and is our first concern here. The second is how to analyze and interpret data provided by measurement on the workflow map. Since most of the data analysis methods used in WFMA are well known and easily supported by spreadsheet functions, and also

because simple analytical procedures are best, actual analysis will be given very brief coverage—the greater part of this chapter will discuss the questions involved in collecting data from a workflow map and using that to experiment and collect data for potential changes and improvements.

This chapter will proceed incrementally, by first looking at business metrics in more detail, and then extending this information into a three-phase approach to process improvement. Owing to the close relationship between measurement and process improvement, many of the issues of implementation of findings will also be discussed in this chapter. Implementation issues which are more general are covered in Chapter 3.

Business Metrics and Process Workflows

Woven throughout this chapter will be the subject of business metrics. Business metrics are, of course, the measures companies and organizations use to evaluate their operations. We are all familiar with fundamental measures such as financial accounts and ratios, scrap rates, recruiting yields, and so on, and these are valuable for many purposes. As useful as these are, they alone may not always be the most appropriate measures for the evaluation of processes, and here we may want to be more expansive.

Most of the WFMA examples we have considered up to this point have been concerned with operations of one kind or another, implying that “bottom-line” performance is the only kind we need to evaluate. However, as Peter Drucker pointed out years ago,² there are at least eight different kinds of performance standards companies should judge themselves against: profitability, of course, but also market standing; innovation; productivity; use of physical and financial resources; managerial performance and development; employee performance and attitudes; and public responsibility. His fundamentally sound ideas have been “rediscovered” many times over the years. In a few words, Drucker is saying that companies have to pay attention to a balance of both effectiveness and efficiency.

Of these two major performance indices, efficiency is the easier one to evaluate. Time and cost, the two universal measures that any organization can apply to its work, are primarily concerned with efficiency, as are many other input-to-output measures. Efficiency is necessary for effectiveness,

but it is not both necessary and sufficient—to be effective, a company has to perform well on nearly all of the eight criteria Drucker outlined, and this is a long-term process that requires managers to have a vision that extends beyond quarterly “bang for the buck.” Moreover, a company has to establish the right criteria for its purposes, since not all are equally important to every firm.

Two recent large-scale studies illustrate this claim. Both of these were done by consultancies (who excluded client companies from their samples) working with academics to investigate the contribution of management to company performance. In the first one, Joyce et al.³ selected a sample of 160 companies in 40 industries and used extensive public-document archival data to identify the management practices that explained differences in performance among these over the previous 10 years; these practices were extracted from a set of over 200 suggested by a reported survey of management academics.⁴ Of these, four mandatory practices were found to be necessary to the success of these firms: clear and focused strategy; flawless operational execution; a performance-oriented culture; and a fast, flexible, flat structure. Two factors from a second group of four also had to be present: keeping high talent; committed leadership; industry-transforming innovation; or growth through mergers and partnerships. Which two of these second four were selected depended on the individual firms. Firms that were successful in the mandatory management practices and two of the secondary ones in effect had an 80 percent chance of being the best performers over the decade of performance studied. As Joyce et al. point out, however, keeping six balls in the air for that long is no easy feat.

In a McKinsey and Co. study of 700 manufacturers in the United States, the United Kingdom, France, and Germany, Dowdy⁵ and his colleagues at London Business School concluded that there is a strong relationship between management practices and the return on capital employed (ROCE). They examined three areas of management practice: (1) shop floor operations, particularly the extent to which the companies had adopted “lean management” methods; (2) target setting and performance management, that is, whether companies set and track the right performance measures and take action to correct processes when necessary; and (3) talent management, whether the firms are attracting,

hiring, developing, and retaining the right people. Companies were rated on 1-to-5 scales on 18 items to measure these practices, and scores were accumulated from these ratings for the three areas of practice. The bottom line on this study was the title used by McKinsey: “Management Matters.” In all four countries, the quality of management practices was directly and strongly related to bottom-line performance and to the factors that drove that performance. Not only did better-managed companies perform better, they also showed evidence that job satisfaction and general “quality of life” measures for employees were higher for the best performers.

In both of these studies, business metrics that both “measured the right things and measured things right” were fundamental to the success of the best-performing firms. None of the best-performing firms were able to achieve that status by paying attention to only one thing. At the same time, both studies show that there is considerable room for companies to find their own ways to succeed so long as they track appropriate performance indices.

Nonfinancial measures and factors which require indirect or approximation measures were also used in these and other studies of firm performance. Many of the things that are important to long-term firm success, such as customer satisfaction, do not lend themselves either to simple or easy measurement or to direct translation into financial impact. Nevertheless, they are often among the most important factors that affect long-term company health and success, and measures of these are as important to include in business metrics, as are financial ratios. After many years of study of job satisfaction and its relationship to job performance, academic management researchers have come to the conclusion that the two appear to be related, although often moderated by other variables. Job satisfaction is never easy to measure, but it matters.

It is also important to realize that the relationship between many of the metrics we might use is not direct, and that the time between a measured change and an observed effect varies tremendously between factors. Likert⁶ proposed a very useful model of the time lag between factors contributing to performance and intermediate factors that moderate the relationship. Recognition of the fact that the time lag between many of these variables is a matter of three to five years is an important aspect of the use and interpretation of business metrics in many companies today—it

is simply unrealistic to expect immediate change in many situations. For a hospital procedure to respond to changes in one case required nearly 2½ years.⁷ Quality is also one of the outcomes for which this kind of lagged relationship is evident.⁸

Many processes have properties similar to “projects” as we use that term in the context of project management. Milestone measures, used to monitor significant progress and intermediate outcomes in projects, are also an important business metric, even though in many cases they are not “measured” as such—we either hit the intermediate target or we miss. But as indicators of progress and potential problems, these milestones are often of major importance, and should always be used where possible. Performance targets of all kinds are among the most important metrics.

The General WFMA Approach to Data Collection and Analysis

The majority of the discussion that follows assumes that we are working with validated workflow maps, and using these as analytical tools to measure the properties of the workflow. WFMA is therefore usually associated with business metrics at one level or another, meaning that WFMA might be used to assess how well a jobholder is performing on a metric (how long does it take for a customer service call to be handled?), or it might be used to create new data that allow that goal to be assessed. No matter which end of the “analysis” part of WFMA a user starts with, there is typically a rather predictable sequence of steps that must be done. These can be grouped into three phases, as follows. In each of these, we will progressively introduce methods that can be used to analyze and interpret the data our metrics provide.

Phase I Develop a Valid Baseline of Current Operations

This phase revolves around four basic steps that enable data collection, with the preparation of the workflow maps as discussed in Chapter 1 as a critical step. For organizations doing WFMA for the first time, it is important that each step be done and that each be given enough time to be properly completed. How much time that is needs to be answered on

a case-by-case basis—for a company under financial pressure, employees may be threatened by the expectation of layoffs, or they may be motivated by the need to regain financial health as quickly as possible. The only thing that can be predicted with accuracy is that a rush through the process will waste a lot of time and accomplish very little.

1. Develop an objective-based approach to WFMA. Do you want to analyze the entire business or begin with a few key processes and then expand? Is the purpose to start a process improvement project, to develop a training document for a job or process, to analyze the effectiveness of alternative approaches to a process, or a myriad of other possibilities?

Many aspects of what you do in the design of the analysis will depend on the objective. For example, if a plant manager is trying to cut down error rates on the part of new trainees in a particular job, the objective of a WFMA project may be to understand how the newcomers do the job and how mistakes happen. This might support development of a better method of training than now used. In another case, the objective might be to determine how the fastest operators in a variable process do what they do, as opposed to others who get the job done correctly, but more slowly. A third example might be to map the documentation of a part of a manager's job so that others can do it while the manager is on travel. In all three cases, who is selected to be the "subjects" of the project, what kind of map is desired, and what the application of the analysis will be, are determined by the objective of the study. In short, be clear in your own mind about what you are doing and why you are doing it. Large organizations may need to create a team, task force, or committee to agree on an approach, and in some cases, to first agree on the objective.

2. Inform everyone about the plan. Mapping and analysis without explanation is almost always threatening to people, and fear motivates self-protection. That self-protection will result in distorted data and general resistance to the whole WFMA project. Saying nothing is the worst thing to do—silence will be interpreted as a threat.

3. Have everyone in the selected process map what they actually do using the five basic symbols and discipline discussed in Chapter 1. The maps should show where materials come from (input) and where they go when finished (output). The map should show what people do, and the information they need to do it. More detail on what to mapping actual versus “should be” workflows is provided in Chapter 3.
4. Validate the process maps. Have co-workers, committee or team members, or supervisors actually follow the workflow map and check it for accuracy. Be sure no linkages in the workflow are missing. More detail on validation is provided in Chapter 3.

Starting with a description of the process as it exists is necessary to get a baseline for accurate measurement or to plan changes and improvements. While it seems logical to many to omit this step when we are trying to improve things (and it is a very American bias to want to “cut to the chase” and move on), it is important to get an accurate initial bearing on where we are before planning the next steps. Tradeoffs will usually be necessary in designing the new workflow, but full knowledge of what those tradeoffs mean can only be gained from an accurate baseline.

Four Common Phase I Questions

Validation is the major contributor to getting a map of what we actually do. It is needed for many reasons, but the major one is simply that it takes thought and work to accurately describe what most people do. Again, material flows are the easily described parts of a job or process and are not a challenge to map, as a rule, but when people get into the details of these, there is much that becomes challenging. Four types of questions typically arise at this point: (1) how much detail should be provided; (2) to what extent should the map concentrate on the “normal” process as opposed to exceptions; (3) should the entire job be mapped, or just the part relevant to process X; and (4) how does one deal with the parts of these on either extreme—the routine parts like scheduled staff meetings and the one-time events like crises and emergencies, which affect all jobs and processes.

Level of Detail

How much detail is needed is largely dependent on user objectives for the map, and that should be the criterion used to make decisions. What do we plan to do with the map? If we think about the variations in simply checking the oil in one's car, there are many possibilities. If we need to train someone who has no knowledge of anything that is done on a car, we may need to include steps such as warming the engine, making sure the car is parked on a level surface, locating and pulling the correct dipstick (many cars have several), wiping and reinserting the dipstick to get an accurate measurement, and so on. To someone who knows about cars and engines, most of this would be "intuitively obvious"; to someone who doesn't, it is the only way to prevent disaster. (As an example, a young neighbor who was hired to mow my lawn while we were on vacation decided to check the oil in my riding mower; it was low, so he added a quart. Two weeks later, as I was mowing the lawn, the overfilled sump caused a crankcase gasket to blow out, spewing hot oil all over that part of the lawn and costing \$300 to repair. That is why most dipsticks have "Do not overfill" stamped into them, and it is necessary to know what that means.)

Graphic software can be very helpful in this capacity, and a number of these programs are available both commercially and as freeware. Most programs allow for multiple "layers" of a drawing to be created as we saw in Figure 1.4, or hyperlinking to more detailed views. As one goes to progressively deeper layers or follows the hyperlinks, more detail can be seen. The user who needs the "intuitively obvious" level of detail can easily coexist with the raw trainee in using the same program, and it may not be necessary to map all the details until the need arises.

An additional consideration here might be regarded as more "strategic." In developing workflow maps for quality certifications such as ISO 9000, it may be necessary to refrain from including all details in a workflow map, or to be very clear that training or other maps not used as overall guidance in performance of a task are identified as such. Quality certification audits generally live by the rule of "document what you do, and do what you document." If the fine detail and correction and safeguard procedures used to train someone to check oil (or a far

more complex task) are also used as the certification document, it will be necessary to demonstrate that this detailed map is followed in all future audits, and that is probably not what we want to do. Certification documentation needs to be clearly identified, and protected.

Normal Versus Exceptional Tasks

In most cases, the default assumption should be that the “normal” and “exceptional” tasks are all part of the same process, and the most likely case is that both are necessary to have a complete picture of the process. This is especially true when the objective is to create a map that is comprehensive of everything done in a position, as in creating a job description. Many jobholders compartmentalize their daily activities into those that are the “regular” or “real” job, and the tasks that are “problems,” “issues,” “goof-ups,” and many other colorful names. As discussed earlier, these exceptions are often the most important components of the job for many reasons (keeping customers happy, for example), and it is not uncommon that significant time and energy go into them, often more than the “normal” work.

The tendency to consider exception-handling tasks as distractions, “junk work,” or generally things that should not need to be done, can lead to serious distortions in a workflow map, and thus in our ability to accurately measure and understand what goes on in a process. If a job consists entirely of being available to assist customers no one considers time taken to answer inquiries and solve problems to be a distraction; if one’s “primary” job is being a software engineer, however, that same time may seem to be the biggest single “time-waster” an engineer faces in a day. The Disney organization trains its maintenance workers, especially those responsible for routine cleaning while guests are in their parks, to be prepared to answer questions and provide help to them at all times—they are typically one of the first and most visible points of contact, and while talking to guests is not picking up trash, it is very important that responding to guests is done, and done well.

Both job design and reward systems factor into this matter of how the “real” job is perceived. Software engineers who are paid and promoted for being productive, when “productive” really means “average lines of code

written each day,” will logically be more likely to consider customer help requests to be distractions. A job design that rewards both coding and customer relations will be less likely to have this perceived bifurcation in task value for the same work. That means both have to be explicitly measured and evaluated in performance appraisals.

This is also a driver in the argument that we want to avoid jumping to the ideal job design when we create a workflow map. Depending on the present design and reward system, that “ideal” may leave customers twisting in the wind when they need help, and over the long run would be a huge mistake. The real issue is to determine what we do and why, and if part of that work is perceived as junk, what do we do with that?

Whole Job or Process-Relevant Tasks

There are cases where the general advice to map the whole job may not be the best path to follow. If a mapping objective is to show only those parts of several jobs that contribute to one process, many positions may be involved in doing the process in its entirety—this is somewhat analogous to an assembly line, except that the workers on the “line” also have things to do other than just put one part into a final product. The tasks of immediate interest are those relevant to the process being mapped.

The same points made earlier regarding exception-handling activities apply to process-specific maps just as they do to job-specific maps. Exceptions occur in tasks done for specific processes just as they do in off-process work, and both regular and exception-driven process activities should be accounted for in a comprehensive map of either a job or a process.

Ennui and Extremes

“Ennui” may be an overstatement, but at one end of the spectrum of the daily routine are the meetings, mail, and recurring things that organizations require, but which are hardly exciting or interesting. The other end of the spectrum is the extreme and often scary events of life—fires, robberies, acquisitions, and the like. How are these to be handled in a workflow map?

To the extent that these are involved in a job, they may be put into two separate categories, if they are handled at all. The extreme and exceptional cases are often not part of either “the job” or “the process” because they are not a recurring business-driven event. As such, while they happen, they are not something for which ordinary use of resources is an issue, nor are they part of a process ever likely to happen again. They are quite simply not on the map.

The routine stuff of staff meetings and taking care of predictable recurring duties can be shown as one or more processes, usually needing little or no explanation. These are typically shown as a minor part of a job description. If they are part of a process, they may be subsumed by a description of a larger task; for example, they might be part of a coordination procedure used to keep different units in an organization informed of progress and aware of problems. Here, again, it is possible staff meetings may not show up in a process workflow at all, although they are likely to show up as one item in a job description. In no case should we overlook the significance of information from scheduled meetings and the like for coordination; these “ennui” are often absolutely irreplaceable connective tissue in the workflow.

To a considerable extent, these questions are a matter of existing job design. Obviously, an “enlarged” or “enriched” job design is more likely to comprise a number of different tasks and behaviors, and thus likely to be involved in more processes, than a job with a more restrictive design. Partly for this reason, no two organizations ever come to exactly the same answer on many of the questions regarding meetings and the like, and the final answers on how to do it will have to be devised by the user.

It also bears repeating that the primary new content produced in validating workflow maps is information. In the large majority of cases, the new validation content will be maps of how information is acquired, evaluated, and consumed for resolution of the exceptions and issues that validation exposes. Much of the change in the validated map will consist of information flows related to decision making and problem solving necessitated by exceptions rather than the “normal” flow. It should come as no surprise to find that the new information content often creates a larger map, in terms of the number of symbols, the physical size of the

map, and the amount of annotation, than the main workflow in the original map.

The answers to these four questions about map content will always be somewhat idiosyncratic to the organization creating the map. Understanding a job or process, and describing it sufficiently for the purposes intended, will guide the decisions to answer them. When the objective is to create a map that supports measurement for process improvement and change, the most important issue in mapping is to enable the collection of data that will support the change.

Before moving on to Phase II, we should note that validation is necessary because people map what they think they should do, what they think the boss wants to see, or what they think will save their jobs. Validation means both accuracy and full disclosure; it may be threatening—see the discussion of that point in more detail in Chapter 3. “Minor” discrepancies may create major costs and problems in the downstream flow of work. Nonvalidated workflow maps prevent meaningful process improvement, effective training, or whatever their purpose might have been—“garbage in, garbage out.”

Phase II Develop Opportunities for Improvement

We can begin by using the workflow maps themselves to look for process improvement opportunities. We will consider the analysis of data more fully in a moment, but if the workflow map has been carefully developed, it provides a basis for at least three categories of opportunities:

- Examine the workflow maps to look for four nearly universal types of process improvement: (1) physical flow improvements; (2) information processing improvements; (3) cost or resource-consumption improvements; and (4) quality improvements. These are typically not separate or discrete—if you find one, you almost always find another accompanying it.
- Look for points where “hard” data can be extracted or collected. Measures with hard data (actual time used, actual costs, etc.) are preferred to estimates or subjective measures,

although in some cases estimates are the only option. When possible, measure key processes and then use distribution diagrams (histograms, bar charts, etc.) to show where time and costs are consumed, as will be illustrated later.

- There will be some things that appear to be process improvement opportunities for which no clear data are available. The first step may be to measure these (e.g., time or cost). In many cases, using a spreadsheet to play “what if” with differing data estimates scenarios may be helpful.

A valid process map enables a large variety of analytical tools to be used. A valid map supports the use of many business metrics that can be associated with the tasks and activities performed on the map, either individually or in related groups. Phase II assumes that a valid map has been developed and that our principal objective is now to use business metrics already in place or to find or create data that enable measurement of the workflow, wholly or in part. In addition, some types of analysis require nothing more than careful examination of the steps and actions in a process, and a willingness to ask why things are done as they are. As suggested previously, simple inspection of the process and the data that describe it can be a powerful tool.

Opportunities for change and improvement are based on thorough understanding of what we are doing now. The point of departure for improvement in an ongoing organization is to know the situation we are really in (sometimes the “swamp” we’re in), and from this, plan an orderly form of change that will accomplish our objectives. This is a psychological necessity as much as, and sometimes more than, a basis of physical or organizational need. Organizations change incrementally over time, and when we reach a point where it is apparent that things need to change, there will inevitably be one group of stakeholders who see the current process as the best possible, all things considered, and others who see it as the dumbest thing there ever could be. Both will require persuasion, and trade-offs in the redesign of processes will have to be made. Doing these on the best factual basis we can is a good place to start.

One of the characteristics of many processes and jobs today is that they are directly or indirectly connected to various forms of information

and communication technology (ICT), which enables the collection of data either as an automatic by-product of doing the job, or with a bit of software “tweaking,” expedites such data collection. Other cases require some kind of direct effort or intervention to collect data, and much of the discussion related to this phase assumes that this kind of new data collection is needed. All too often, that is the case.

The validated maps are a framework to measure costs and resource consumption on the workflow. We should acquire data by measuring against the actual workflow; in all cases, time is a universal resource which can always be measured. In the early stages of a process-improvement study, measuring against the workflow may seem difficult because the business metrics used do not match the characteristics of the workflow. That is, there may a major (or complete) mismatch between what has been measured and what we do, to such an extent that connecting existing measures other than time to the workflow map, in a meaningful way, is extremely difficult.

This may seem an unlikely state of affairs, but it is not—over the years I have encountered quite a number of situations that illustrate this reality. In one case involving preparation for the notorious Y2K bug at the turn of this century, a regional utility company sent several participants to my WFMA course, and I learned that their concern was that they did not fully comprehend how sale of their commodity finally resulted in invoices to customers! (This was confirmed by examination of a large color print-out they brought to the last class, where politically correct worker symbols along the top of the page performed tasks through organizational units shown along the left side.) This “map” connected tasks to workers to show processes that produced the company’s outputs, and the linkages were portrayed by hair-fine colored lines running through this huge sheet. One of these outputs was billings, and several cost-generating inputs literally “trailed out” and simply disappeared, or else merged into others in this “map,” before they got to the invoice! One consequence was that for the first few months of the year 2000, numerous customers received wildly erroneous billings. Some got credits they later learned were mistaken, and they then had multiple-month bills to pay; others got inflated billings, the most notorious case being a bill exaggerated by a factor of 10,000. It took months for the company to clean up this mess.

In another case, a successful financial services company was growing so fast that very little current or accurate documentation of jobs, functions, or procedures could be created or maintained. By the time someone wrote up these kinds of descriptions, the job content had changed, the jobholder had moved on, a new product had been introduced, or a reorganization had occurred. The problems being created by these issues were not evident for a long time because the rapid growth, and the revenue stream it brought, made it difficult to see (or sometimes to care) about the fact that what was being measured and rewarded often had little discernible relationship to the everyday duties of many people, or to the success of the business.

About halfway through the lifespan of this firm, I was retained to design a WFMA program and train a small group of trainers, so that this company could get internal control over what it was doing. After designing and piloting the company-approved program, the three people I trained were reassigned to different divisions of the company, two in other states. Five years later, a participant in an open-enrollment WFMA course said he was taking it “to learn how to use the program you designed years ago.”

Eventually the stresses from this state of affairs took their toll. What had begun as a highly motivated, highly cohesive team-based firm grew into a large company of different fiefdoms. In these, the reason to work changed from “do it for the team” to “Just [Bleeping] Do It,” known as JFDI; the reason to stay with the company went from being an elite member of a unique team to “golden handcuffs”—they paid so well for so long that it quickly became hard to find comparable compensation elsewhere if you got tired of the pressure, and there was always pressure. This company no longer exists—it became necessary to curb much of the exuberant spending as the firm aged, and eventually the company was forced to sell itself to a competitor.

In Chapter 6 we will see how radically intended versus actual (validated) workflow processes differed. This was the aircraft electronics repair process for in the U.S. Naval air force during the Cold War that I have mentioned several times. The designed workflow is shown in Figure 6.1, and a composite of the actual workflows is shown in Figure 6.2 for those who would like a peek in advance.

I have never seen an organization that is exempt from issues of “what do we do versus what do we measure.” It is therefore not uncommon that workflow metrics may need to be created or modified to accurately track what gets done in a process. Workflow maps can be used to acquire data on many individual steps and subprocesses, as well as to measure aggregate data for processes as a whole. Clearly, different kinds of business metrics can be generated using such information, and the data these metrics provide can be used to test and model different alternatives for processes. Examples of the kinds of data that can be collected include:

- Who performs each step or process
- How much time is taken for each step or process
- Labor costs at each step
- Volume of work at each step
- Delays and time lags
- Quality measurement and other control times and costs
- Normal versus exception handling activities and costs
- Rework
- Turnover rates
- Job satisfaction (periodic)

As suggested earlier, information technology is an inherent part of the workflow, and in most organizations today it may be possible to obtain much data directly from the ICT system. In many cases the technology itself records data on individual tasks and activities and allows data to be extracted from existing files (although not always at a level of detail consistent with the workflow map). When this is available, significant additional work to record data may not be necessary. However, a special log is sometimes needed even when information technology is there, to obtain the process-specific data we require.

Activity Logging—Using Workflow Maps to Capture Data

The primary method for setting up a recording device or procedure is known as “activity logging.” There are three basic steps in logging activity data:

Explain the Process

If you are going to do a workflow analysis and log, the first step is always to inform your people of what you are doing and why you are doing it. Workers must know that they are not being watched by “Big Brother,” that we are not trying to get data that can be used against them, and that complete honesty is necessary for accurate analysis. For the great majority of your people this will not be an issue, but as suggested before, for those who are concerned the reassurance is very important.

Create a “Form” to Log Observations

The second step is to create a format to record data in a consistent and organized manner. This needs to be nothing more than a single sheet of paper divided into columns (or a spreadsheet), one row or column for each task, action, or decision, and labeled to record data on these events. Each sheet should be dated and the period of observation (a shift, a number of hours, or whatever time period is chosen) should be recorded. The person working during the time should record his or her name or be noted by the observer. Each person should have a watch or clock that can allow easy measurement of time to the second (computer clocks and smartphones can do this, of course).

There are two basic types of forms—a sequential data sheet and a tally sheet. In a sequential sheet, the observer records the data on events in the order in which they happen, usually using the actual time (number of seconds, for example) for each activity. A tally sheet usually records data as a “tick mark” in predetermined intervals. For example, a fairly short work cycle might be broken down into 10-second intervals, with activities being recorded as from 0 to 10 seconds, from 10.01 to 20 seconds, and so on.

Information technology can also be applied in many cases to record live actions for analysis at a different time or place, with the added benefits of not disrupting actual work to collect data, and being able to review data for accuracy.

Record Data on Observations

Recording the data is relatively simple. The observer needs only to record the time each activity in the workflow diagram starts during the period of observation. Event duration is easily calculated by subtracting the beginning time for an event from the beginning time for the next event in the sequence. Spreadsheet macros can do this easily and automatically.

Who should do the recording? Data collection is best done by the worker who actually does the job, although it can be done by a supervisor observing the worker, or a specialist or outside observer; but in most cases, it is best to have the worker record the data. Experience with such analyses shows that this method is usually accurate. When it is not, it is usually because the worker has not been trained to record data correctly, or else feels threatened and distorts the data.

In workflows where individual processes and decisions normally happen very quickly, it is a good idea for the person collecting the data to measure the “cycle time”—the time to complete an entire series of events in a segment of the workflow from beginning to end. This avoids adding a significant amount of time (to do the data recording itself) to the entire cycle. The only way to detect this “measurement inflation” is to time the whole cycle without stopping, that is, as it is normally done end-to-end. An outside observer is often the best person to do this.

Differences between full process times and the sum of times for individual activities and decisions may or may not be important. In many cases, the greatest interest is in the individual activities and decisions within the workflow. If there is need to be concerned with full process times, the correct way to evaluate the difference between the full process time and the sum of its parts is to take a sample each way, compute the average for the full processes and the average for the sums of parts, and subtract the difference between the two averages from the sum of individual times. This correction removes additional time to record data and can only be safely used for the full process, not the individual parts.

Practice

A final part of activity logging is very important—it needs to be practiced. Measuring is a skill in its own right, whether recording the time needed for tasks in a process, taking physical measures of products, collecting survey data, or whatever form is used.

When I conduct seminars in WFMA, I use an in-class exercise consisting of playing several hands of poker, and I will have one group of observers time the major steps in playing several games of poker: shuffling, dealing, exchanging cards, and determining the winner. After everyone has both played and logged the process times, we evaluate the data we get. Inevitably, a major source of error in the data comes from the observers themselves—they realize they did not really know when one task ended and the next began; they get caught up in the hand and forget to record times; they record only the part of the task that they believe is “significant,” not the end-to-end time; and many other sources of error.

A conclusion that is universally applicable from this experience is that for activity logging to produce meaningful data, observers need to practice taking measurements. I always recommend at least one trial session, a detailed assessment of all the things observers can think of that can reduce error in the measures, and if possible a trial of the “new and improved” observation technique before any “real” activity logging is done.

An Example

While this may sound as if it is becoming a complicated process, it seldom actually is. A real example of how simple activity logging of a cash transfer method yielded valuable insight into a process some years ago, and soon thereafter suggested an effective process improvement, is shown in Figures 2.1 and 2.2.

Figure 2.1 shows the workflow map for this process, which is the one I had worked with in 1989. This client was a financial services company, and part of its business was to handle large cash transfers for global customers who were generally managing cash flow, engaged in multilateral netting of payables and receivables, and also investing surplus cash in overnight and other investments through open-market operations with the

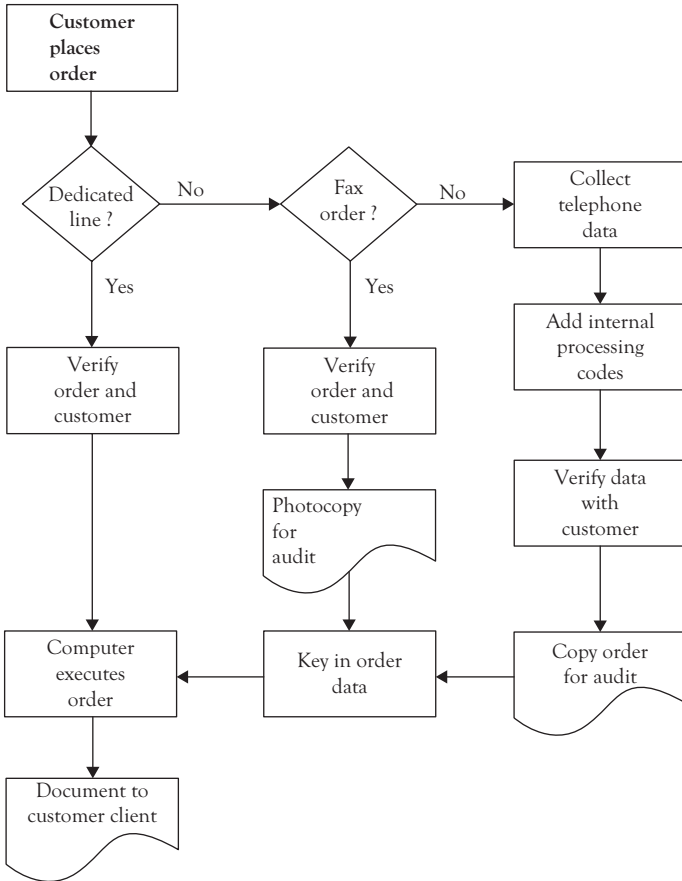


Figure 2.1 Workflow for customer cash-transfer order by method of receipt, ca. 1989

government. Many of the firm’s customers moved large cash deposits in and out of various instruments frequently, so the fees generated from even small charges for this service should have been a very profitable business. Nevertheless, internal audits had revealed that over the years, profitability had eroded to the point where this was barely a break-even operation.

The initial workflow map showed what was generally already known—there were three methods for orders to be taken. One was through a dedicated line on a secure mainframe computer, the second was through fax orders, and the third was through telephone orders. The company prided itself on both speedy and personal service, and over the

years its volume of business for many customers had grown because of this service-oriented business model. That was the good news.

When it was agreed that the workflow map in Figure 2.1 was valid, the volume of transactions coming in through each mode of entry was counted, and the steps involved in taking and executing a transfer order were timed. It was found that nearly 50 percent of orders were being taken by telephone, 32 percent by fax, and the remaining 28 percent through dedicated-line connections.

The next step was to measure the average time required for each method of ordering. This was done by observation and logging the activity of a sample of orders and order-takers; the time required for each step for each of the three methods of ordering were summed to generate an overall time for each type of order. The cumulative average time in minutes required for each type of order are shown by the bar graphs in Figure 2.2.

Not surprisingly, the dedicated-line orders shown by the left bar in Figure 2.2 took only three minutes on average since verification of a customer code was the only operator intervention needed (and this, of course, has since been automated). Fax orders (the middle bar) took nine minutes on average, partly because customers had to be called and have orders verified, and then keying time was needed in addition to making copies. Telephone orders (the right bar) were the longest, by

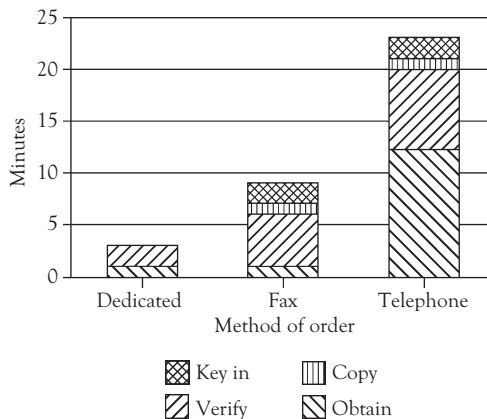


Figure 2.2 Minutes required for cash transfer by method of order

far—averaging nearly 23 minutes for each order. This seemed an anomaly, and triggered further investigation. The company soon learned that this was an accurate measurement, and was an unintended consequence of its service orientation. Customers would call in, often toward the end of the business day in their time zone, and over the years many first-name personal relationships had developed. Order-takers would chat about family, weather, upcoming holidays, and all manner of interesting topics as they took the orders, and customers who were surveyed genuinely liked the warm, personal service they received. Customers were known as people and valued clients, not just impersonal sources of revenue.

The bad news was that when the average revenue generated by telephone transactions was converted into equivalent minutes of operator time, it was discovered that the average transfer generated a fee that covered 18 minutes of contact, while the average contact was going on for five minutes longer than that. Hence, the telephone orders were costing the company money on average, not making it, and these were the largest volume of orders!

What to do? The company did not want to risk losing its service reputation, but had to reduce the amount of time spent on the average telephone order. The solution to the problem was to analyze the volume of business arriving by telephone, and send the larger customers a free fax machine with a preprogrammed dedicated number. Some customers took to the fax immediately, and over time those who still called orders in on occasion became accustomed to the convenience of the one-button call, so volume was gradually shifted from telephone to fax. As fax machines became cheaper, and more microcomputer support came into the market, this operation returned to profitability; the Internet finished the job in the late 1990s, about 10 years after this study was done. Interestingly, the time needed for the die-hard phone orders never fell very much, and it seems that the social interaction was as important to these clients as completing the transaction; however, that volume is now very small.

As a final observation, it is worth noting that in the near term the processes themselves did not really need to be changed, but rather

the volume of orders flowing through each of them. There was active discussion of the possibility of taking immediate steps such as providing incentive awards, price discounting, and the like, to encourage customers to stop placing telephone orders and use other methods. In the end, the company decided to do nothing that would potentially have a negative effect on its reputation as a service-oriented firm, and so a “thanking you by helping you” approach was used—the fax machines were the first wave, followed by more of them as prices fell, and later there was dedicated software and other forms of assistance to encourage movement to the Web, where nearly all business is now conducted. Over time, the process was nearly completely automated through ICT and is now quite different from that shown in Figure 2.1. But by choosing this approach, the company lost neither friends nor a profitable business.

A General Workflow Metric

With this background on what numbers might tell us, I want to suggest a general measure that can be used with any workflow, capitalizing on the fact that every process requires time to complete. This general measure is Mean Process Execution Time (MPET). MPET is the sum of two parts: Mean Primary Process Time (MPPT) and Mean Exception Handling Time (MEHT), or

$$\text{MPET} = \text{MPPT} + \text{MEHT}.$$

The mean (average) time to complete a process depends on how long it takes to complete the primary process flow, where nothing goes wrong, and how long it takes to handle exceptions when things do go wrong. Both of these are important in their own right, but neither should be ignored when considering overall performance in a process. A company may be very good at handling normal demands that follow the primary workflow process, while being terrible at handling things that do not go smoothly. Whether this is a “good company to do business with” may well depend on whether the customer we ask was the recipient of the primary process or had to have extraordinary help to deal with an exception; both are important, and adding them together keeps both visible.

This metric can be used in the form suggested here for any workflow; at the same time, it can be modified in many ways. The sum is the simplest version of this metric—we can also examine the two components of MPET individually, of course, or weight the two components by the proportion of production that goes through normal and exception-handling processes; by converting them into a standardized metric on a single scale (zero to one, for example); by normalizing them (expressing them in terms of standard deviations); and others. Each modification yields additional information, but the basic sum of the two reflects how well we do on both the normal and exceptional outputs that make up our total production.

Phase III Experiment and Implement

The third phase of process improvement is to experiment with potential changes and modifications to the workflow, and to collect data on the results for future decisions. An excellent, well-regarded and well-established basis for this kind of experimentation and data collection is the Shewhart cycle in Figure 2.3, named after Walter A. Shewhart,⁹ the “father of statistical quality control” (and often referred to as the “Deming cycle”). The fundamental idea is that of scientific experimentation, trying things to find out if they work better than what we have in place now.

The logic of the Shewhart cycle is the basis of its long-term success, as has been amply demonstrated by the Japanese automobile industry, and




Figure 2.3 *The Shewhart experimentation cycle*

by many in many other industries as well. “Experimentation” is usually better than decreed change, and solves many problems in the change process. One positive psychological aspect of experimentation is that it encourages getting input from the people doing the work and who are affected by the experimental change. After the experiment, we keep what works and abandon what does not, and by having those affected involved in the experiment, it is much more likely that they will regard this as “their” experiment as much as “management’s,” to the benefit of personnel “buy-in.”

1. *Plan a change or test.* Define the problem, suggest possible causes.
2. *Do the test.* Small scale at first; collect data.
3. *Study the effects of the test.* Let the data speak.
4. *Act on what was learned.* Improve, test, monitor, improve; recommend and implement.

Working in an experimental mode eliminates the need to deal with “loss of face” or a need to protect one’s status—that alone can often stop good decisions or perpetuate bad ones. Staying with a bad idea is not a good decision, and neither is failure to accept someone else’s idea when it might work.

Experimentation encourages communication. Keeping people informed about when we are experimenting, what the experiment is, letting them know the results and what we plan to do next, all contribute to successful process improvement. Knowledge drives out fear.

Experimentation allows control of the scale of investigation. Process improvement change can be a major restructuring or a small change in one part of a workflow. Most experiment-based changes will be small and incremental, with the cumulative results being large over time.

Experimentation, finally, creates a specific context in which many business metrics can be used, but more importantly, such metrics must be used if we are to evaluate experimental results on a rigorous basis. The best expression of this is in the statistical design of experiments for quality improvement, but the spirit of experimentation does not require that level of measurement control every time. What is more important is that we follow the four steps in Figure 2.3, that we measure as objectively as we

can, and that we look at all the results of the experiment before making final decisions about what we plan to change or implement. No method is perfect, and even the “exact” sciences make mistakes, but this approach has been shown to be an effective way to evaluate process improvements in many organizations.

Process improvement is continuous improvement—it is never “done.” Experimentation establishes a culture of trying new ideas and always looking for a better, cheaper, faster way to do things. It is hard to find any business that is truly the same as it was 10 years ago, or even five years ago, and many die long before reaching that age. Complacency is not an alternative in a competitive world.

Analyzing Data from Workflow Maps

Whatever business metrics we use to collect data, the primary methods of data analysis are familiar to most readers, and so this section will be very brief. Simple data analysis methods found in all spreadsheets are emphasized. Statistical tools (most of which are now embedded in spreadsheets) may be used to analyze data, but our minds and eyes are the most important tools in a great many cases.

The most basic idea to keep in mind when analyzing data from any source is that very few things in the universe come in one form only—there is enormous variability in things, and data from workflow processes and experiments will always be characterized by this variability. Consider a simple “operation” like shuffling a deck of cards. We might take a sample of people from our office or workplace, give each of them a deck of cards, tell them to “thoroughly” shuffle the deck, and time them as they do.

What will we find? In a word, variability, for all kinds of reasons. One person is a regular poker player and can do a thorough shuffle in 10 seconds; another has never played cards, has no idea either how to shuffle or what “thorough” means, and takes two minutes; and so on. If we have 10 people in our sample, we will get 10 different times, varying from 10 to 120 seconds if these are our two extreme cases. Most will take between 15 and 30 seconds to shuffle the deck, with 10 being very fast

as a reflection of the practice that person has had, and 120 seconds as an “outlier” value produced by a complete novice to playing cards.

We might decide there is so much variability that we need more data, and so we get another 10 people to do the same thing. In general, the larger the sample, the better; but we have to trade that off against the cost or feasibility of obtaining a large sample. Whatever the decision about sample size, we will eventually have a data set, and from that we can obtain a number of useful items of information about the process.

The first thing to do with any set of data is to describe it—get an overall look at what you have. The best way to do this is to use a few simple lists, tables, and graphic tools to help you see the data. We call these basic methods descriptive statistics, because their primary function is to describe our data. For most analyses we will ever want to do, simple description is all that is needed. The only tools needed to do the analysis of data will be your eyes and your spreadsheet.

At this point, unless the sample of data is really large or many people have somehow shuffled the deck in exactly the same time, no two values will be the same, so the data set will simply be a list of times. We will next use the spreadsheet to sort them from low to high and form a frequency distribution. A frequency distribution shows two things: (1) the values of the things we have measured, and (2) the number of times each value occurs, or the frequency of each value. In this example, “values” refer to the times taken for each person to shuffle the deck of cards, and the “frequency” refers to the number of times each value occurred in our measurements. When arrayed from lowest to highest values (or the other way if it suits our purposes) you literally see a distribution of how often each time occurred.

There are a number of types of frequency distributions, but some of them are especially useful in analysis of WFMA data. One of the most useful is a histogram, which is the best form of distribution for variables that only occur in whole numbers (a family can have two or three children, but not 2.6; in the case of our card shuffles, we might have discrete values for each person shuffling, so we would want to group the data into categories). A histogram of shuffling data shows the frequency of each category on the vertical axis, and the individual times to complete

a shuffle falling into each category on the horizontal axis, as a series of groups. There are usually from three to nine groups, depending on the number of observations and the range. A very simple and useful chart of this kind of data, and one we have seen in Figure 2.2, is a bar chart. It is simple to make, and very informative.

Spreadsheet versus Manual Construction

The best way to build a histogram is to let a computer do it for you. Spreadsheet software like MS Excel, OpenOffice, and Quattro Pro are all able to do the job, and give the user many choices of types of graphs and different views of the data. In Excel, if one types “histogram” in the Help dialog box, a response for Statistical Analysis Tools will be generated. Following that, guidance will give the user access to a histogram tool, as well as many other statistical tools available in this program.

To build the histogram manually is not difficult, but takes a little bit of time and isn't as flexible as the computer. To do it manually, one only needs to record the number of occurrences of each value, and arrange these in order from lowest to highest. Observations having the same value are all recorded in the same category, so that one sees a number of “stacks” of values. The height of the stack is the frequency of that value, that is, the number of times that value occurred. A histogram summarizes data which have been grouped on some basis—for example, we might show time for card shuffling in five-second intervals, so that events taking 36, 38, and 39 seconds all fall into the interval of “35.01 to 40.00 seconds,” and the frequency of that interval is three events. But in all cases, the spreadsheet is the fastest and easiest tool to do this.

How wide each interval should be is a matter of judgment, and there are no fixed rules. You want to have enough groups to see the variation in the data—neither too many nor too few should be used, and you may need to experiment to determine the range of intervals that give you the most useful information. The only mistake to be careful of is that all intervals must be defined so that events fall into one or the other with no ambiguity. For example, an interval of 1 to 5 seconds cannot be followed by an interval of 5 to 8 seconds; into which one does an event of exactly 5 seconds fall? The solution is to either make the first interval 1 to

4.99 seconds and the second 5 to 8, or else 1 to 5 and 5.01 to 8—these are mutually exclusive.

Bar charts are just a special form of grouped data presentation. Many values of things in the world are mutually exclusive by nature: male versus female; day shift versus night shift, and so on. These data are usually shown in discontinuous bar charts (i.e., the intervals do not touch each other), with the height of each bar representing the number of each of the values. Again, our friendly local computer is the best way to make these.

Extracting Information from the Data

Frequency distributions can be built for each of the activities logged in a WFMA study, for a series of them within a job, or for entire workflows. Different frequency distributions can be constructed for individual workers, for different shifts, and for as many different ways of doing the job as there are. An overall frequency distribution, combining data for individuals and shifts for each process, can also be made. A huge benefit of spreadsheets is that they allow us to enter not only the raw data, but also additional data to identify sources and other characteristics that enable further search, through sorting, regrouping, and combination. Whenever possible, entering data into a spreadsheet rather than keeping it on paper allows software support for many additional analyses.

The most basic information to get from a frequency distribution of your data is from visually examining it, and seeing several of its features. This kind of simple observation is often one of the best ways to use data. First, you can observe what the minimum and maximum values are—the shortest and longest times to complete a shuffle if playing cards, or to complete a series of steps in a workflow. The difference between them is called the *range* of the data. The range gives some idea of how much variability there is in the data, and we will be able to measure this in numerical terms.

The shape may be very informative—for example, are the large majority of the events either very short or very long (i.e., is the histogram “piled up” on either end or very symmetrical on both sides)? If piled up in one direction, one can ask what accounts for the extremes.

This could be very interesting for many reasons—are the short times low because people are highly skilled, or because they are cutting corners; are the longer times an indication of rework to correct errors, lack of timely delivery of components, or driven by some other cause? Depending on our objective, either or both of these extreme values may be of considerable interest.

Measures of the “middle,” or measures of central tendency are also very useful. Three of these measures are most often used for different purposes:

- The *mean* or average of a group of data is the most common—it reflects both the number of times something occurs and the range of values. Computing a mean by hand requires only adding all the values and dividing by the number of observations. The mean is often symbolized as \bar{X} (“x-bar”) or M . This is calculated by selecting a spreadsheet function; in Excel and Open Office this function is = AVERAGE(range).
- The *median* is the value dividing the number of observations in half—half fall above that value, and half fall below. In Excel and Open Office this function is = MEDIAN(range).
- The *mode* is the most frequently observed value in the distribution, the highest bar in the chart (there can be multiple modes). In Excel and Open Office this function is = MODE(range).

Many of the variables we use to measure processes along workflow maps can only take on whole-number values, like the number of loaves of bread out of 100,000 that fail to rise correctly before baking (they are not what statisticians call “continuous” variables). The histogram discussed previously is the appropriate way to portray these data, and the mode and median are the most appropriate measures of the middle. A cumulative table can also be constructed, and these tables are also informative.

Over the years, I collected data on the time taken to shuffle cards as part of the in-class poker exercise in my WFMA course. Figure 2.4 shows the histogram for 163 card shuffles. Each bar in the histogram is the frequency of shuffles that occurred in a 10-second interval, going

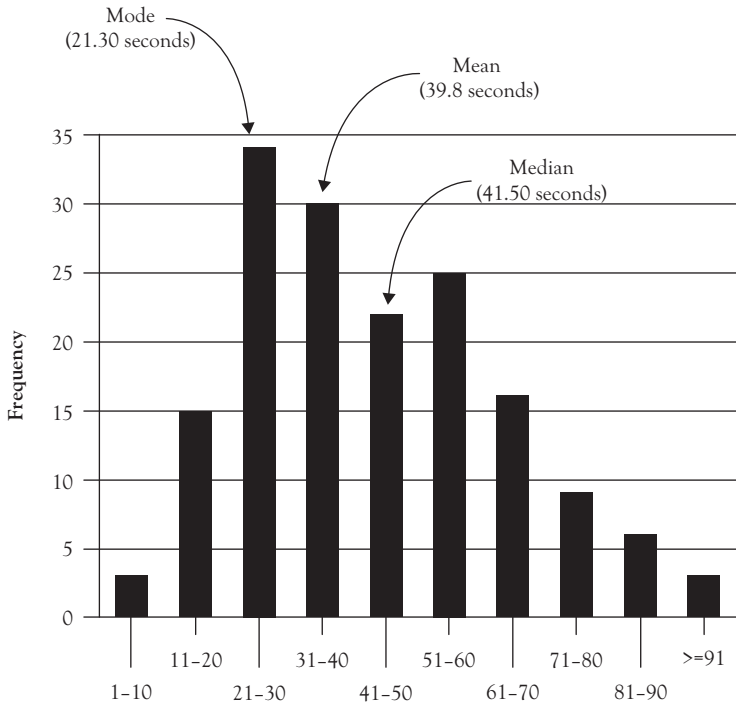


Figure 2.4 Histogram of 163 card shuffles in 10-second intervals

from left to right across the horizontal axis. For example, we can see that we had three shuffles done in the 1–10 second interval and three in the 91-second-or-greater interval on the far right; we had 15 shuffles in the 11–20 second interval, 34 in the 21–30 second interval, and so on. As we can see, this histogram is not completely symmetrical, in that there are more high bars in the left half of the histogram as compared to the right half.

We can easily find the measures of central tendency discussed above, and these are shown in Figure 2.4. The mode is the 21–30 second category, the highest bar in the histogram. Since there are 163 cases, the median is the value for shuffle 82, which falls in the 41–50 second category. The mean is simply the average, where you add all the times and divide by 163, and that is 39.8 seconds, in the 31–40 second category. (Not having the actual times, the reader cannot check this—the best you could do is multiply the number of shuffles by the midpoint value of each interval, which is 5.5, 15.5, etc. If you do that, you will get a value of 43.7 seconds

for the mean. The reason for the difference is that many of the actual scores in the 21–30 second category happened to be on the low side, around 23 seconds.)

Variability can be measured and expressed in numeric terms, and usually is. The measure of variability can be the *variance*, measured in squared units, or the *standard deviation*, measured in original units. The reason the variance is expressed as a squared value is that it is measured in terms of differences above and below the mean, and because the mean is the exact weighted middle of the distribution, values above (positive numbers) will cancel values below (negative numbers), and variability would be zero. We square both types of values, add them and average the sum, and this solves the problem, except that variance is expressed as the square of original values. If we want to convert this back to the original units of measurement, we simply take the square root. Better yet, we let the computer do both by using = VAR(range) and = STDEVA(range) in either Excel or Open Office.

Larger values for either variance or standard deviation mean the data are more spread out. Thus, if we collected data for a hospital laboratory process on two shifts and found that while the mean (average) was the same, the standard deviation was considerably larger for one shift, we might want to know what accounts for the greater variability on one shift as opposed to the other. A situation like this would mean that while the averages were the same for the two shifts, one was both completing some processes in less time than the other shift, while some processes were taking more time.

Frequency distributions for different workers, shifts, or processes may vary on any of these characteristics. Some activities may be done either very quickly or after a long time, with few times between the extremes. In all cases, the question to be asked is “Why?” What does the shape, the difference between workers, or other information tell us? Why is the average for one event on the day shift 30 (or 3) percent higher than the average for the night shift? These are questions to be investigated in more detail, and in most cases, all of the questions can be addressed in part by using the data from activity logging and frequency distribution analysis, whether our interest is improving efficiency, quality, or control.

Clues for Performance Improvement

To conclude this chapter, I want to consider how the results of the kinds of analyses we have been discussing can directly contribute clues or leads to performance improvement. The first way to obtain such clues is to examine the measures themselves in detail; the second is to examine the measures and the workflow properties they describe against our organization design and performance objectives at a higher level. This discussion will necessarily be kept somewhat general because the nature of what is measured in workflow analysis is typically very specific to the organization.

What's in a Number? "True" and "Error" Components

Whenever we use a metric, it is important to think about the composition of what has been measured. We noted above that one of the things necessary for data collection was practice, because we want to minimize observer error. This bears repeating here—what this means is that we want the variability we see in our data to tell us what is going on in the workflow, rather than being a measure of relative observer skill.

Any measure of something can be thought of as composed of two parts—a "true" value and an "error" component:

$$[\text{measure or "value" or "score"}] = [\text{true component}] + [\text{error component}].$$

In the symbolic language of mathematics, we can call these x , t , and e , so that our expression above simplifies to:

$$x = t + e.$$

In the overwhelming majority of working situations, each of these two parts of x , in turn, can be broken down into two smaller parts. First the "true" component t can be thought of as a Least Practical Value (LPV) and some degree of Individual Variation (INDVAR) in the way a task or process is done; thus:

$$t = \text{LPV} + \text{INDVAR}.$$

What do these mean, and why do we care? Think about our hypothetical poker game. If we are examining the times taken to shuffle,

where each x is a time someone takes to shuffle a hand, then the LPV component would be the least time practical for the person shuffling the deck to complete that task. Obviously, this will vary depending on skill and manual dexterity, but for most people who have at least some experience with cards, there will be a time beneath which adequate shuffling cannot be done, and an upper limit beyond which additional shuffling is simply unnecessary. For most people, those two values express upper and lower limits for the LPV time to shuffle cards.

If we observe someone who really wants to be sure that shuffling mixes the deck, that person may go through multiple cuts of the deck followed by several riffle shuffles, several more cuts, and so on. Those additional steps are seen as necessary to that person, but may be more than we really want or need for shuffling to produce an acceptable randomization of the cards in the deck; that additional time and effort are that person's INDVAR. INDVAR can also be negative, as in a case where someone is careless and does a minimal job of mixing the deck between hands. A person who simply takes the top half of the deck and puts it on the bottom and says "Done!" is demonstrating a lack of concern for randomizing the deck (INDVAR) as well as violating the minimum time needed to do the job (LPV). Tacit knowledge plays a huge role here.

Thus, in looking at our data, one of the things we might want to think about is what constitutes the LPV for a particular task or process. That may require additional evaluation of data and some work and thought to resolve, but once an LPV is determined (usually in terms of its upper and lower limits rather than a single value), it provides a valuable benchmark to contrast to other scores. In an activity like shuffling cards, values outside the limits for LPV are probably due to INDVAR, and may be reason for a change or intervention to bring the actual times for jobholders within the limits.

It may also be the case that variation outside the limits of LPV occurs because of error, the e term in our equation above. The e term also has two parts: Observer Error (OBSERR) and Individual Error (INDERR). In many cases where we must collect raw data, the observer is a significant source of error for many reasons. Anyone who has tried to observe and time task performance realizes this very quickly, as it is very hard not to get involved in whatever is going on; it is also hard not to be distracted

by outside events. In many tasks there are ambiguous transitions between one task and another—those transitions may be preconceived to be smooth and easy, but there is very little in the world of work where we find the clear separations between movements as in a symphony. If there is a remote or automated “sensor” used to collect data, there may be errors in these that are very difficult to see or correct.

INDERR is similarly subtle, and in some cases is hard to differentiate from INDVAR (and in some cases there may be no real basis to differentiate it at all). The difference may be of greater psychological than substantive importance; for example, a worker may have learned to shuffle cards in the unnecessarily thorough manner described earlier by observing another worker who considered this the right way to do it (tacit knowledge, again). Informing that person of the need to change because they are “making a mistake” is probably less acceptable than being informed that they are very good at shuffling, but need not go through so many steps to do the job equally well. In some cases, it may be easier to detect and evaluate, such as when a person simply mishandles a deck and unintentionally shuffles by playing “52 pick-up.”

Both types of error suggest important possibilities for improvement. To the extent that OBSERR exists, it is necessary to find ways to improve the ability of observers to get accurate data, because with respect to the actual performance of the work being observed, this is pure “noise” in the measurements. Elimination of this noise gives a clearer picture of job performance. INDERR can occur for many reasons, and to the extent that we can identify and eliminate it, performance is improved, benefiting the company and the worker alike.

The main point to be made here is that metrics and measurement are invaluable to our search for ways to improve performance, but that neither should be taken for granted. Selecting meaningful metrics is important, and so is application of them in collecting data. Both take thought, and the latter takes practice, as well.

What’s in a Number? Accumulated Time in a Workflow

The idea of a distribution sometimes seems a bit abstract, although we have illustrated what is in a distribution in our discussion of Figures 2.1

and 2.2 earlier—we just didn't use that term. Here is a way of combining both ideas, which may be helpful (although we have to skip ahead a bit to do this). I have often referred to the study of avionics maintenance in the U.S. Navy that was the beginning point for the WFMA methods explained in this book. A high-level diagram of that process is shown in Figure A.1 in the Appendix.

The focal point of the early parts of that study was the Versatile Avionics Shop Test (VAST) shop in the lower-right corner of Figure A.1, where avionics components (ACs) were actually repaired. There were several predictable and unpredictable steps in the VAST-shop workflow, which contributed to the total time it took to process an AC through the shop. This was measured in hours as Elapsed Maintenance Time (EMT), and a typical distribution curve for an AC is shown in Figure 2.5.

Processing an AC as described in the Appendix consisted of several required steps, just like taking orders in our example of Figure 2.2. The first step is referred to as “buildup”— the technician got the necessary test program tape and connecting cables specific to the AC being tested from

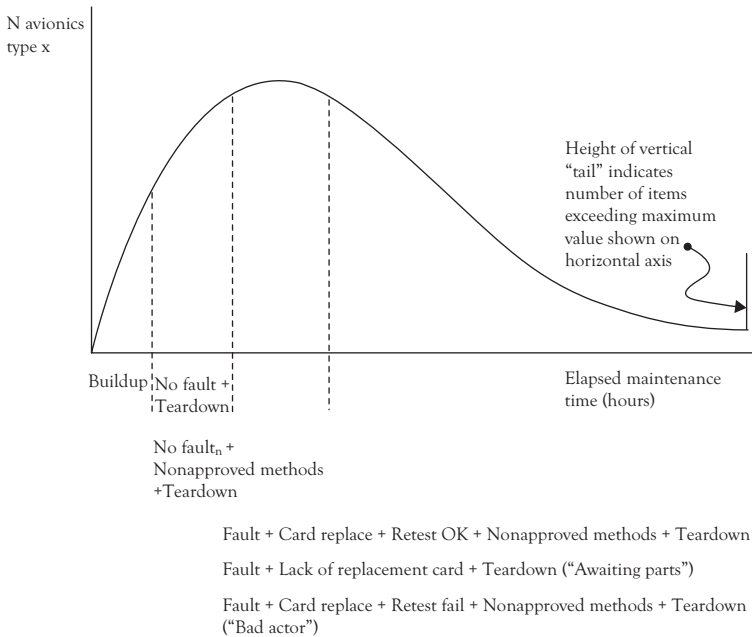


Figure 2.5 *What’s in a number? The composition of EMT in avionics maintenance*

an assigned storage area, and then mounted the tape and connected the cables to the AC and the tester. Next, the technician started the computer test program, which ran until a fault was detected, parts replaced, and the AC passed a retest, or until either all tests were completed and no fault was found (this happened from false error readings on the aircraft, and could be as much as 30 percent of all items tested), in which case the AC could simply be placed back in the Supply pool. However, a “tear-down” was needed to remove the test equipment installed in buildup and return it to its proper storage locations. Buildup, Test, and Teardown are always needed to test an AC. If we were to test a group of the same ACs, Figure 2.5 shows that for that group, the first part of the time curve will always be Buildup, and in the event of no fault, the test and teardown will make up the next part of time under the curve (this will vary for each AC). This figure simply takes the steps in the process and arranges them from left to right, instead of stacking them as in Figure 2.2.

In most cases, the tester indicated a fault at some point, and this required the requisition of replacement electronic “cards” (very similar to computer parts) from the Supply facility, which replaced the ones the tester had indicated to be the cause of the fault. The AC was then retested, and if the card replacement solved the problem, the AC was now OK to return to Supply to be available for an aircraft when it was needed. These two steps, card replace and retest, were next elements added to the cumulative time curve (with Teardown at the end) in Figure 2.5.

However, life is often not this simple or predictable, and this is where complications could set in. These could add a great deal of time to a test run that ultimately might not result in a repair. There were often test and retest procedures that technicians had learned from experience, but which were not part of the approved Navy process (all of which had been extensively tested and evaluated prior to approval as the correct way to do things); these might be used as “workarounds” to a problem test, and added time that was not expected to be needed for the repair. In some cases, the cards requisitioned by the test program would not be in stock (many factors could influence this, but such stockouts were not intended to occur), and the test would be aborted, with the AC being torn down and sent to an Awaiting Parts holding area until the cards arrived. The AC then had to be built up again and the process resumed until repair,

or possibly the next Awaiting Parts status for some new stockout, was done. Finally, some ACs were just extremely difficult to test and troubleshoot effectively, and after long periods of time these simply did not get repaired; these were usually sent to advanced maintenance facilities on shore, which was contrary to the maintenance plan for these ACs, but there was no choice. These were typically categorized as “bad actors,” because one of their effects was to consume a lot of tester time without anything being repaired.

In Figure 2.5, we see the accumulated times for a group of identical ACs (i.e., one type only, not a mix), and how time accumulates from the origin as these parts of the workflow are completed; this accumulation of time can be directly related to steps in the workflow map. Buildup is always the first part and must be done for every AC; at least one test cycle must be completed; every AC has to have one teardown for every buildup, and so on. By having a clear workflow map, our analysis of accumulated times, or any other workflow metric, can be understood more easily and provide more ideas for improvement of that workflow. If we look at different types of ACs we will see different distribution shapes, just as we saw different “stacks” of tasks in Figure 2.2. Each of these ways of portraying data gives us something more to see, and helps us to understand, our workflow.

Workflow Design Trade-offs

The final issue to consider is what we want a changed or improved workflow to do—the performance criteria we measure, and the trade-offs we must make to get the multiple but often conflicting outcomes we need. In many ways, this becomes a matter of a trade-off between efficiency and risk, and the questions this trade-off raises are best understood in terms of two workflow properties—convergence and divergence between flows of material and information, and synchronous versus asynchronous flows. These characteristics are discussed fully in Chapter 5 (and the relationships between them are shown in Figure 5.5).

For the time being, it is worth noting that the optimal design for many process improvements is to align the flows of information and

materials so that they are both convergent, that is, occur in the same place, and synchronous, occurring at the same time. This design usually yields the highest efficiency, but also comes with the highest risk—it usually increases the amount of control a jobholder has over the work, and in some cases this increases the potential for things to go wrong. In one case at the Harley-Davidson motorcycle company, a drive to “enrich” jobs allowed assemblers to decide whether a defect in a chrome-plated part was sufficient reason to not install it on a finished cycle. One purist on the assembly line chose to regard the smallest flaws, even those on the inside of an exhaust pipe that would never be seen by the customer, to be a “defect” and thus be scrapped. Technically, the worker was right, but his approach was far too costly, and he had to be retrained to avoid applying such stringent criteria to his decisions.¹⁰

We might separate material and information flows in time, as is done in much scientific research, where through replication studies, separate examination of results and outcomes are performed, thus giving a double-check on the initial result. This approach to risk reduction, however, comes at the cost of reduced efficiency and speed of response, one of the recurring issues in bringing new drugs to market, for example. An alternative is to manage risk by using divergent paths for materials and information, as done with “double blind” evaluation procedures in the pharmaceutical industry. This may reduce time delays, but requires multiple modes of information processing, at higher cost. “Double blind” means that neither those who administer a drug in development know whether they are giving the drug being tested or a placebo; neither do the patients who receive the dose. Only a select group of principal investigators knows who got what, and they must both sort through the reports of efficacy (or absence of it) and evaluate the outcomes with knowledge of who received the drug; as with any kind of measurement, they face the same problems of individual and observer error we discussed earlier. Obviously, double blinding is slower than allowing those who administer a drug to report the results, but we avoid some serious risks. (Some readers may remember the birth defects that resulted from rushing Thalidomide to market; while useful for treating some cancers and leprosy, the drug was sold as an over-the-counter preventive for morning sickness,

and caused significant deformation of limbs, urinary, and reproductive systems in over 10,000 infants. This outcome was one of the reasons that drug testing was tightened in later years.)

Managing risk requires active consideration of workflow design issues, and must always be done with awareness of the inevitability of the trade-offs as well as awareness of the options. Knowing what we know, and knowing what we don't know, become very important in making these trade-off decisions. These matters are discussed in more detail in Chapter 4.

Summary

WFMA not only produces graphic portrayals of workflows, but also enables the application of business metrics and measures to the processes being mapped, in part or in their entirety. The use of business metrics is a much better way to evaluate process maps than reliance on “gut feel” or intuition, and there is always at least one measurable variable in a workflow, which is time. With a bit of thought about what happens in a workflow or a part of it, in terms of both actions and information (as discussed in Chapter 5), many metrics may be available for analysis.

The simple, stepwise approach one can use to analyze the data and experiment with improvements to the process has been in use for many decades, and has proven its worth in quality and process improvements around the world. It is basically a model of applied science, and is applicable to any kind of process in any organization.

The basic descriptive statistics used to collect and analyze data derived from workflow maps are powerful because of their simplicity, and from the insights we can gain by careful observation of our data with a spirit of asking questions and experimenting with potentially constructive changes. Using one's eyes to really see what the data look like rather than going into more complex statistical procedures to analyze the data, along with applying one's knowledge of the work being done to understand what the numbers we have collected actually tell us, can go a very long way toward potential changes and improvements.

More complex statistical analysis has its place in workflow improvement, and what I am saying here should not be interpreted as an

argument against such analyses. But sophisticated analysis is never a substitute for “knowledge of local waters,” and anyone trying to design process improvements without knowledge of the uniqueness of each company and its culture and workforce is likely to incur the same unhappy status as the Massachusetts Institute of Technology project. The typical review of a first-time WFMA project includes many “aha” moments, and many of these are based on observation of process steps alone, with no measurements being taken at all. As these easy targets are identified and acted upon, other process improvements may be found by doing more intensive quantitative analysis.

It is worth noting the famous “learning curve” that came from the aircraft industry years ago. In brief, this curve shows that the cost of manufacturing an aircraft tends to fall at a relatively constant rate (20 percent is commonly used from early papers on the subject) for each doubling in production of a particular aircraft. This happens because so many parts go into an aircraft that given whatever sequences of assembly are being followed at any given time, there are always alternatives. Some of these alternatives are more efficient than others, and as experience is accumulated we learn through experimentation how to do things more cheaply with no loss of quality or other necessary outcomes. The “discovery” of the learning curve was largely an outcome of industrial experience that later led to the expression of a formal relationship between volume of production and unit cost; but careful observation and a willingness to experiment came first. This still holds true, and is a huge opportunity for many of our service industries, an opportunity I personally suspect is not being exploited to nearly the extent it could be.¹¹

CHAPTER 3

Implementing Workflow Mapping and Analysis

Introduction

In the fall of 2013, I was retained by a major global company to help provide a process map for what I thought would be a simple function: travel for company personnel on company business. The client company had several different geographic divisions but a nominally centralized office to plan and procure tickets and accommodations for its people when they had to be in other places, referred to as its “travel voucher” system; like most firms, it had negotiated special rates with railroads, airlines, hotels, and car-rental firms, so it wanted its personnel to use these to minimize costs. In addition, the company received some funding from governmental groups for basic research, and these funds required special accounting and reporting to the funding bodies. The company knew (and had told me) that there were some idiosyncratic differences between the processes used at different locations to initiate a travel voucher, which often made life difficult for the people in the central office. I was asked to help a select group of these people who knew the process, and find a way to reconcile them after we found out exactly what they were.

While I expected to be there some complications, there was nothing to suggest a major problem. I planned to be on site for a day, which would be spent giving a short introduction to the material in Chapter 1 in the morning, and then in the afternoon the group would work out the application to the company travel system. I would nudge and tut-tut as needed to keep people true to the discipline, and then wrap up with a short presentation on the need to use the new system, but remember to protect it, and call me if anything problematic came up. I’ve done this many times.

On one hand, this program went as planned, and as many others before had gone. I always start these workflow mapping and analysis (WFMA) implementations with a conversation, partly to find out how much differentiation in methods there is, and partly just to engage people, about which more later. What I had not been told before taking the job on was how extremely different the travel-voucher process was between geographic locations (country blocs); in truth, I do not believe the U.S. group that retained me actually knew themselves. What was needed, in short, was not a reconciliation of different parts of the process, but a nearly total redesign—beyond the most basic objectives of getting people out, keeping them reasonably comfortable, and getting them back home, there was little agreement on anything else, except that “our system is best because” I extended my stay for another two days, and at the end of the third day, we had worked out a new, uniform travel voucher system.

This story illustrates a number of things any process mapper should bear in mind. First (and by now I hope this is obvious, but if not, read Chapters 4 and 5), the “same” process will vary between sites and jobholders, for many reasons. Second, once that variation settles in, it becomes the “right” way to do the job, and any change, no matter how well-intentioned, is somewhere between an annoyance and a threat, and will be resisted. Third, the psychological importance of these variations become very personal to the jobholders, and any attempt to change may well be seen as a threat to the person as well as the job; if we don’t devote some time to finding out how important these properties of the process are, we will inevitably encounter a form of “irrational” resistance to change that is puzzling, frustrating, and ultimately may anger us to the point where we become the threat we were first perceived to be. Fourth, if the mapper is an outsider (as I always am), that person will be construed to have motivations that can only be imagined, and until those affected by the project have the sense that you are trustworthy, or at least honest, the mapper will be the last to know what those motivations are. I have found that being completely open and transparent is the only approach that works, and in some cases have turned down retainers when top management would not agree to those conditions. This is an important

point to keep in mind when working for a large firm as I did for the travel-voucher client, but is just as relevant (and often more so) for any internal mapper in any large organization.

More often than not, WFMA implementation is welcomed and seen as a useful (frequently overdue) change. It is an opportunity to improve efficiency, become more competitive, eliminate “because” as a reason for particular procedures and workflow steps, and importantly, to include the tacit knowledge that many workers have been using all along and recognize them for that. But it may also be a precursor to reductions in force, realignment of functions, or the sale of an entire business line. In truth, I have seen jobs eliminated in about half of my engagements in WFMA, and in many cases, this was somewhat of a surprise to everyone; the rationalization of workflows and improvements in performance meant that more could be done with fewer people, and for many small-to-medium businesses, the cost reduction was too important to ignore. Larger firms and organizations can often find ways to absorb displaced workers when such circumstances arise, but small organizations often cannot. So there can be a perceived downside to process improvement, and any mapper should be aware of that and acknowledge it if asked.

Whatever the circumstances, there is a growing body of findings to show that harnessing tacit knowledge improves organizational performance. Whether in well-established, customary applications such as quality¹ applications in health care and health care management,² or in better integration of related services (which might be interpreted as the proverbial “breaking out of silos”),³ tacit knowledge has been shown to be valuable in everyday activities. In many cases, the best way to discover what tacit knowledge exists in an organization is to undertake WFMA.

What follows in this chapter is a general set of guidelines to implement WFMA. Many of these topics have been introduced briefly in previous chapters and are expanded here. Given the extent to which process improvement is dependent on valid workflow maps and the extraction and analysis of data from these, many implementation issues were covered in Chapter 2 in relation to measurement. But there are also several which are new, and more related to the environment of WFMA than to mapping itself. These matters are covered here.

Implementing WFMA

Plan and Prepare

No matter what the organization, or the size of the project, it is necessary to plan and prepare if WFMA is to be successfully implemented. In many cases, this means the mapper will need to get rid of preconceptions and superficial ideas of what will be undertaken; it is better to begin with a Zen-like expectation of “finding what is” rather than “installing what should be.”

An important part of this initial step is to have a conversation with those whose jobs or processes are to be mapped. In a one-on-one case, this should obviously include introductions, an explanation of what the mapper plans to do and why, and an opportunity to answer any questions the subject has. In that connection, the mapper should generally expect that the subject will not ask all the questions at that time, and emphasize the willingness to answer questions at any stage of the process. The same applies if the mapping is done with a group of subjects or a group of mappers. A meeting and conversation is vital.

The other part of the conversation explores what the subjects do in the process, and if more than one individual is involved in it, it is helpful to start at the beginning, with the first person who initiated the process. Active and careful listening are required for this part of mapping, and the mapper should be as willing to ask questions as the subjects are presumed to be.

Somewhere in the conversation, I use the sentence, “I know you believe you understand what you think I said, but I’m not sure you’re aware that what you heard is not what I meant.” This may come up in response to something I said as well as something said by a participant, but it doesn’t really matter—what is important is that if we parse that sentence into its constituent phrases, we find that (a) all of the parts are true, (b) all of the parts are mutually reinforcing, and (c) all must be addressed if we want to change the perceptions of the subject. This may be useful at a point where stress, fatigue, or other difficulties have emerged; it may also be part of a banter that comes after a breakthrough. Not only is it funny, it is a profound truth.

Process Owners

Part of the planning and preparation may require designation of process owners in larger organizations, and possibly the establishment of a “configuration control” designate (or office) to maintain the discipline of the Kmetz WFMA approach.

As WFMA is applied more widely in an organization, it is strongly recommended that a “process owner” be designated (if that has not already been done). Process ownership can be very simple or very complicated. When there is a small organization or a process that is contained within only one department or unit, “ownership” is almost always the responsibility of the head of the unit (although in some cases they may actually not know it!). In larger and more complex organizations, this becomes more complicated—the nature of the organization structure, its culture, and the attitudes of managers, in addition to the complexity of the process itself, all contribute to the properties of process ownership. At one extreme, these matters may all have been resolved and different managers or owners know they are responsible for success in a process, whether all actions in it reside in their department or not; at the other extreme, one can find processes perceived as so screwed up that no one wants to be associated with them, let alone “own” them, or where no one has really thought about the fact that they have “processes” in the first place.

No matter what the case, process ownership will be a necessity, and only makes sense given the importance of processes in organizations. For one thing, designating a process owner gives recognition to the individual who cares about it, and also will most likely result in that owner being given the influence necessary to make decisions and changes to the process; if the owner is now responsible for process success, that status is likely to be accompanied by the authority and resources needed to make it happen. If that is not made clear, the owner needs to pursue it and negotiate working relationships with others who have control over assets and personnel who affect the process. This is often not a simple arrangement to work out, since such negotiations affect the structure of the organization and the relative power of affected individuals.

The Role of the Process Owner

The role of the process owner is often a key to success or failure of WFMA and performance improvement projects. WFMA can frequently be seen as controversial simply because it may bring change; organizational politics are inevitably implicit in such efforts, and it is both foolish and risky to assume that either information or access to it is politically neutral. This is especially true where processes cross unit boundaries in an organization, so that work done within the process requires contributions from multiple units, each of which is headed by a different individual. The “process owner” is likely to be a higher-level manager than these unit heads, but because of their status, they may not take direct control of the process—that will be accomplished indirectly through the unit heads. So who is the (ideal) “process owner”? To summarize, someone with:

- Experience as a senior manager (and often a team of senior managers may be preferred as a “process owner”)
- A predisposition to oversee and work with the teams in the process group
- Major equity in the process (and process group)
- A clear understanding of activities and challenges in the core process
- Knowledge of upstream and downstream activities and processes
- Ability to influence people, to coach, and to support the team.

The process owner plays a critical role in supporting and shepherding WFMA projects. What are the process owner’s responsibilities? To:

- Define performance objectives for the process group;
- Monitor team members’ performance and resolve conflicts;
- Promote and drive continuous improvement and collaboration through the group;
- Develop process plans and budgets;
- Serve as a process “champion”;

- Build a sense of shared objectives and support within the process group;
- Help process team members build complementary skills and work methods;
- Evaluate and recognize progress toward the objective;
- Identify and remove performance impediments in the analysis;
- Keep creative tensions alive in the process group (not all conflict and not all resistance are necessarily bad);
- Represent all areas of activity in the process to others and outsiders;
- Recognize and reward the good work of process team members; coach those needing help; and discipline or remove those unable to contribute.

This is a brief summary of the role and responsibilities of process owners, but it is clear that these are key individuals in the coordination of work across unit boundaries in complex organizations. Without their active support for a WFMA project, it is very difficult to accomplish one—for the individual unit managers, the work involved in workflow mapping is a distraction from the larger job the organization has to do, and is a burden without immediate rewards. It may be true that the unit managers see the need for the WFMA project, but so long as that entails doing “someone else’s” work, it is not likely to happen.

A process owner who meets the specification here is a key to successful knowledge management (KM) in two ways. First, the process owner pilots the WFMA project through politically tricky waters, a necessity because the project is a vehicle for change that may be resisted by many. Second, as the WFMA performance improvement or other change effects emerge, the process owner can maintain the spirit of experimentation discussed in Chapter 2, so that the final changes and improvements are retained in the knowledge base and visible to the KM system (KMS).

The major benefit of process ownership on a day-to-day basis is quickly recognized by most firms. They can track process performance, something often lacking in many companies; be included in any decisions that affect their process; suggest changes when they are needed; test new

ideas; and generally become the “go to” resource to answer questions about the process or its performance. The process owner does not necessarily have to be a formal head in the organization structure; instead, the owner can be recognized as equivalent to a project manager, who is often in charge of a project that spans internal organizational boundaries. I personally recommend that status if the actual process owner is not a supervisor or manager.

Configuration Control

A configuration control system should actually be put in place for most organizations, and like process ownership, the owner of configuration control need not be a member of the formal hierarchy, but should be given authority to control the system. Such control is a good idea no matter what the scale of process mapping, and given the simplicity of the Kmetz method, it is also very simple—it consists of review and final approval of all WFMA maps, and saying “no” to those that do not conform to the discipline.

This may not seem necessary for most organizations, especially small ones, but my experience has shown a perverse relationship between the acceptance of WFMA and the desire to change it. Many of my clients over the years have proudly sent me first- or second-draft maps of their processes after completing my training program, and in more than half of them, there will be instances of minor changes that are not consistent with the symbol set, or other cases of not following the discipline.

In one of these, done by two librarians who worked in a last-generation IT environment, I was sent a rather well-done workflow map except for the addition of two symbols by their supervisor—one was storage on a magnetic tape (a circle with a line extending from the right bottom), and the other was several off-page connectors (a small upside-down Monopoly house). I showed these to my wife, who has no flowcharting training, asked if she knew what these meant, and of course she did not. I relayed that information back to my students, and in doing that I suspect that I filled the role of *de facto* configuration manager, since I didn’t hear from them again. Apparently, their boss got into the mapping and “wanted to make just a few small, constructive changes.”

Repeat this a few times in any organization, and you quickly invite the full flowcharting symbol set to be used, resulting in confusion for everyone. Don't go there, ever.

Who Prepares the Map?

Whenever possible, I recommend that the best approach is to have the person who does the job prepare the map. Prior training in the symbol set and its use is strongly advised, but this is not a time-consuming requirement—having trained hundreds of people to create workflow maps, my experience is that 75 percent of trainees will learn to use the symbols correctly on the first or second try, and nearly everyone else gets it on the third try. This is especially likely if several nonthreatening exercises are done as part of the training, and groups of trainees serve as “consultants” to each other. The “consulting” usually consists of simple tasks like helping people remember to label exit paths from decisions, not mix diamonds and rectangles (a common problem at first), and putting arrowheads on arrows. Clarifying what is meant by the content of the diagram is also important—putting statements in a decision diamond when a question is actually needed may be confusing to the intended user, for example.

While there may be some circumstances where the initial (first draft) workflow map should be prepared by a supervisor or a specialist, my experience is that these are rare, and generally not beneficial, regardless of the size of the organization. First, it takes a great deal of time and interaction between two people for the second one to understand and correctly describe the other's work. Second, there are many perceived reasons why the jobholder may not want to be fully forthcoming to someone else, particularly at first—lack of trust in that person, a feeling of being examined or watched closely (perhaps to be caught in mistakes), or a sense of pressure just to get done, from either or both people. Most importantly, the internalized tacit knowledge on the part of the jobholder may be harder to discover and understand if interpreted by another. These factors interact to make validation more difficult (discussed below), because the onus of correct explanation of a flow of work has been shifted from the person who does it to the person mapping it.

Having said that, the mapper should be aware of the likely need to go through several iterations before getting a final map, and that can result in some relationship problems with jobholders. In many instances the jobholder-mapper will have never done this before, and having to change a version of a map naturally implies that the map was initially not correct, for whatever reason, and this can elicit defensiveness on the part of many individuals. The need to stay with the discipline can also lead to irritation; this is particularly true when digging into exception-handling processes, but can also develop in the course of explaining the tacit knowledge and steps involved in a complex part of a nonexception process.

There are many situations where the objective of WFMA is to develop a map for a process that is performed by a group. An alternative to building the process map from individual ones may be for the group to do the map collectively. For WFMA projects like this, the common “sticky note” can be an engaging way to get everyone in the group involved. The common rectangular (process block) and square notes (decision diamond, when rotated 45°) provide the two workhorse symbols; the other three can easily be drawn on either of these two. These can be put on a wall or whiteboard in a public area (the coffee pool) with an invitation to everyone to change or contribute to the diagram as one sees fit. Not only does this create a great deal of interest in the mapping process, it also creates “buy-in” for the final product, a vital concern for change and process improvement.

As with any job, larger organizations may want to have one or two skilled specialists in WFMA diagramming around; these people will likely be able to help prepare and validate maps more quickly and easily through working with individual jobholders, and can serve as a mapping resource to others. Specialists selected and trained for their interpersonal skills, for example, are often very effective at getting through the validation of a map with less stress than the supervisor or an outsider (and excellent interpersonal skills are recommended). They can take a first “rough cut” jobholder map and work with the jobholder to easily do the editing and revision these typically need. As they learn the nature of work done in a group or unit, they can become very effective at asking good questions

about exception handling and parts of the job where tacit knowledge is important. They can play a valuable role, and are well worth the time and practice taken to develop them.

Create Actual Maps of the Workflow

A sure way to waste lots of time with WFMA is to let people diagram the way (they think) work *should be* done when the mapper needs to know how it is *actually done now*. Most WFMA projects start with the intention of description and diagnosis, and often are intended to support process improvement. Validation, the next step, will help ensure that the map really reflects what is being done now, and that is the point of departure for any application of the map. Mapping “what should be done” is almost always an exercise in fiction writing, to some extent. Fiction writing is *not* the right idea for WFMA, no matter how good the final story.

Mapping the process “as it should be” automatically introduces individual biases and perspectives into the map; no matter how well-intentioned these are, they inevitably are suboptimal for others. What is more, creating a “should be” map shifts everyone’s focus from diagnosing what is actually happening (and the important underlying reasons for that) to trying to install the new order of things, making it unlikely we will stay focused on the question of what existing workflow issues are in the first place.

This should not be interpreted as meaning that a normative, should be map can never be drawn—a normative map might serve as a valuable straw man for discussion of how one or more jobs might be redesigned, for example. The problem is that when this approach is used for many existing jobs and workflows, the version of reality that is reflected in the map will be an idealized view of it, and often a view held by one or two passionate individuals alone. In mapping the workflows on seven aircraft carriers and seven shore sites in my NAVAIR studies, I found no two alike; nevertheless, I had three experiences where officers or chief petty officers provided me with “ideal” workflow and organizational maps for different sites. None of those matched, either.

Validate the Map

When a map has been done, trace it through to the final details and verify the map against the workflow. If the map doesn't completely match the process, it is not a valid map, and you don't know what is actually being done. Revise the map and try again.

Inevitably, that which is crystal clear to me when I describe it in familiar terms is absolute gibberish to someone else; steps that are assumed to be known from birth will be absolute unknowns to someone else; common language and acronyms that "everyone" knows will be misinterpreted, and so on. All of these issues must be resolved before a map is considered valid. While jobholders usually prepare the first draft of a map, having others' eyes and ears on the map for validation is equally valuable, and almost always results in change and clarification.

This is a task that often falls on the manager or originator of the mapping process, and sometimes requires the support of higher management (such as the process owner, discussed previously). In simple terms, the validator often has to be a pest and "walk through" the map with the person who created it. At every step, that "pest" needs to be prepared to ask questions, clarify terms, correct some of the mapping discipline, add or subtract steps, and the like. Typically, a revised map will be needed, and this may happen several times until the map is fully validated. This is a high value-added part of the discipline, and if it is not done there are many things that can go wrong with workflow mapping, producing maps that are ambiguous, incorrect, and misleading.

Getting a fully validated map means that there will be "little things" in the first draft that are omitted or do not match the actual flow, that is, exceptions, and many of these are not "little things." It is not at all uncommon for the major part of the work done in a process to be the smaller part of a map, while exceptions and deviations from the norm take lots of time and energy to diagram, often mimicking the reality of actually doing them. Getting this information will take persistence, but this is where process variation lives, and it is often where the richest opportunities for process improvement are found as well.

Mapping and validating actual workflows are parts of the mapping process where information flows become paramount. This is true for

several reasons. First, most people are familiar with the routine process that works correctly, so it is easy to map it and describe its details. This also seems to be the most important part of the process to map, since it is where most companies add value that customers are willing to buy; it indeed seems logical to pay the most attention to this part of the process.

However, these routine material flows are not the part of the process that necessarily requires the most attention from workflow staff; that distinction is usually associated with the exceptions we referred to above and these are often complicated with respect to both the sources and causes of them, as well as from the work needed to correct them and get the process or product back on track.

In validating a workflow map, the majority of what will be added to the map will be examples of information processing activity undertaken solely to deal with exceptions, and much of this may be for specific types of exceptions. These exception-handling maps may well be larger than that for the entire routine workflow. Generating them is often more time-consuming than considered justifiable when it is first done, given that exceptions are not where value is perceived to be added—people often resist doing this part of the process, for that reason. The tacit nature of much of this information also makes it difficult to describe and to map. In many cases, companies may find it helpful to start with mapping the normal, routine, successful workflow and get that correct (this is good practice for mappers, too), before going into mapping the handling of exceptions.

My experience has been that creating both a complete workflow map—one that attends to all the details as they now exist—and validating that it is in fact the actual process that is followed both when things go right and when they do not is the most challenging part of my method. This is true largely because the mapper must dig out all of the detail that has to be acquired, processed, understood, and put into the map in such a way that both the nature of the exception and the process to handle it are clear and fully communicated. For experienced personnel in many organizations, this is an area where deep internalization of information occurs over the course of handling many exceptions, and this is vital tacit knowledge in the workflow. This can be discovered and mapped, but it is typically much easier to carry out this process than it is

to fully describe the activity, the information required, and the outcome it should generate.

What Happened Before, and What Happens Next?

If a map cannot account for where something came from before it got to its current location in a process, or where it goes next, is it really connected to the rest of the organization? Is it really connected to critical suppliers? Is it really connected to customers? As we will see in Chapter 5, organizations are *systems*. Valid workflow maps must show what the system is actually doing, and that includes predecessor and successor activities and information flows.

This issue may come up in organizations with highly differentiated structures, which are typically made up of “silos” where work is “thrown over the wall” by a unit when they are done with their part of the process. In order to fully understand the process, this obviously cannot be allowed, and the entire process must be mapped from its beginning to its end. But there are circumstances where the mapper has no alternative but to deal with work that is “thrown over the wall.” Military organizations and their suppliers are one example, and there are cases where what is done in a process is absolutely never discussed outside the office where it is done, to safeguard security. Company secrets must be protected (think of the formula for Coca-Cola); some procedures require people in some parts of a process to be kept unaware of what others are doing (think double-blind tests in pharmaceuticals); many small firms have processes that exist as they are because the owners want them that way, and for no other reason, and so on. It is not necessarily the case that such processes can be fully mapped, and we may have to deal with that reality.

Taking only an internal view of our processes can be risky for many reasons. In one study, McElheran found that in the world of electronic business, the leaders that were studied invested in process innovation only when their customers had to invest little to use their new processes, or were already well-aligned with them.⁴ Moving too quickly in these circumstances, it was implied, could result in alienation of customers and have a counterproductive effect. Suriadi found that gaps in risk analysis were common in many forms of business process mapping (BPM) which

were based on information technology.⁵ In many cases, their extensive study concluded that integration of risk-aware BPM models into mainstream management decision making was still in a future realm.

Transparency and openness are usually best. I once worked on a military proposal effort where part of the aircraft being designed had top-secret equipment on it and this was designed in a “black room.” The manager of the black room was asked to attend one of our organization meetings where we explained the need for transparency and understanding of the process. The manager understood, mapped his process and validated it, and asked us to take his word that this had been done; but we would have to accept that when our work extended into his area, our maps would have to say something like “...and then a miracle happens ...” and continue on the other side. This worked rather well, and we won the program based on our proposal. In another case, the owner of a small chemicals firm listened to my recommendations for process change, and then affably, clearly, and firmly stated that things would stay as they are because he liked them that way. And so they did.

Make WFMA Your Own, but Limit “Customization”

Some circumstances may necessitate minor changes to the basic approach here to account for things unique to a process or organization. Modify WFMA when there is truly no choice, but only as much as really necessary. Using “swimlanes” or color, to designate where parts of the process are done, are usually fine, for example, but *do not* add more symbols (or “modifications”) to the symbol set.

Even color can be a change that becomes undesirable. I once had a group of MBA students prepare process maps for a firm going through a changeover to Radio Frequency Identification (RFID) order processing, and they proudly proclaimed an improvement to my system, in that they had mapped all material movement in black, and all information movement in red. I asked our sponsor to bring two personnel who would be involved in the new system into the room, and when they arrived I asked them if they could interpret these maps; one presumed it was the new RFID method, and the other couldn't. I considered that to be a less-than-50-percent success rate, and rejected the color modification.

If color or any other modification is added, be sure to add that as a legend to every page where it occurs. Being very conservative about changes, I find that anything that can be done with color or swimlanes can also be achieved through notation in the map or annotation of it. Even a fairly complex workflow, such as that shown in Figure 6.2, which involves three separate organizational units, was done with symbol notation, using only four of the five symbols in the set, and all in black and white.

This rule is especially likely to be violated by those with flowcharting, data flow diagramming, or similar kinds of experience. The temptation to add a shadow to a rectangle after doing it for years is hard to resist; so is using a round-shouldered as opposed to a right-angled rectangle, or using double vertical walls. But doing this now requires the user to learn, and to carefully observe, these subtle differences while also attending to the logic of the flow and the information contained in the symbols. This introduces noise and inevitable confusion to the system, and these are never good things—they will simply confuse users of the map.

Will Maps of the Same Process All Be the Same?

I have pointed out several times in this chapter and its predecessors that there is a great deal of variation in the way individuals and organizations do things, and that part of what will be in a map depends on the perceptions of the mapper. While it might seem logical to expect that using a restricted symbol set and a consistent discipline would produce the same maps for any specific process, that is not the case. It bears repeating that there will be more than one “right” answer to a process map, and in most cases there will be many right answers.

Consider the simple act of making a list of food and groceries before going shopping for them, an exercise I regularly run in WFMA training programs I conduct. Many people do not actually make a list, but rely on memory (pure tacit knowledge) to do their shopping, and making a list is a somewhat artificial behavior for them in the first place. For those who do make lists, the nature of that list depends on a huge number of variables. Some people rely on preprinted checklists (also available as smartphone and other device apps) and record items needed on them; in some cases everyone in the family is expected to contribute to the

hard-copy list, but in other cases it is done by one primary shopper. Some use a freeform “running” list to accumulate items between shopping trips.

One memorable former student had a family of nine children, and his wife had given up a lucrative career to become a full-time mother, requiring them to provide for the entire family on one income. Their “shopping list” began with a month-long detailed menu plan (updated weekly) so that exact items could be purchased in sufficient quantities, and at the same time sales and price reductions could be fully exploited to minimize food costs. An unexpected sale on seafood might cause a change in the menu to take advantage of the cost savings, meaning that the menu plan sometimes had to be updated more often than weekly. All items selected for purchase were then matched with coupons if possible, and sometimes lack of a coupon meant that an item was removed from the list and the menu modified to reflect that. Regular monitoring of newspapers, mail offerings, Internet coupon sites, and other media were a regular part of menu and shopping list preparation, and this task was done daily. The mother had become a self-trained nutritionist, which meant that overall menu planning and shopping had to meet nutritional standards, and that food purchases were subject to those tests. This was by far the most elaborate and rigorous shopping list preparation I have ever encountered, and could not be in greater contrast to the unmarried man whose “preparation” for food shopping consisted of quickly looking at his pantry cupboard, followed by a scan of his refrigerator (sometimes smelling or examining items to see if they had spoiled and had to be discarded), and then going to the local grocery store to get whatever he thought he needed; the trigger event for doing this, however, was usually that he had run out of beer.

One student had an aversion to frozen foods, and shopped every two to three days to purchase only fresh foods. She had been born in Italy, and as with a great many European city-dwellers, shopping was a daily event, and the shops visited were very specific for individual food items. That habit carried over when she emigrated to the United States and no longer found the assortment of small shops she had known in Italy, but the idea of thawing frozen chicken or frozen beans before cooking them was simply incomprehensible to her, and her shopping lists reflected the tacit knowledge that so strongly influenced her shopping behavior.

For the average class group of 20 students, there is typically not one true duplicate among the workflow maps for shopping list preparation. The same is true for every other exercise I have ever run as well as for every organization where I have mapped working processes. Any workflow map that is a complete and valid portrayal of a process is “right,” without regard to whether it matches another person’s map or some presumed “standard” that does not really exist.

Three Things WFMA Is Not

Having discussed guidelines for WFMA and examined some of the issues that may arise in doing it, it is appropriate to end this chapter by examining several things that WFMA is not. While there are some commonalities and shared techniques with these methods, there are also important differences, and we should be aware of these when performing WFMA.

Flowcharting

Most methods of charting or mapping workflows are derived from common flowcharting symbols, as mine has been, but there are no truly standardized methods for mapping.⁶ While well-known flowcharting symbols are used to create workflow maps, workflow mapping is *not* flowcharting, although the outward similarity at first and the frequent use of that term might make it seem so. There are a number of important differences between WFMA and flowcharting as the latter is applied in programming and information technology environments:

1. Only a limited set of flowcharting symbols and rules are used for workflow mapping, and these are distinctive and easy to learn and apply. This is quite important for workflow maps to be useful to those not trained in flowcharting, and makes the symbols a means to the desired end, which is understanding the workflow and communicating about it. Most flowcharting symbols, in fact, are far too specialized for workflow mapping.
2. Workflow mapping works by capturing and describing an actual process used to accomplish a purpose, whatever steps may be

involved. Flowcharting works from a logical program structure and follows rules similar to computer language. In that sense, workflow maps are *descriptive*, where flowcharts are *prescriptive*.

3. Workflow maps can vary enormously; it is unlikely that two different workflow mappers would create the same map for any purpose or process. Therefore, many different versions of a workflow map can be “correct.” Flowcharting will produce a “right” answer that can vary only to the extent the programming language allows. Unlike a workflow map, it is likely that two skilled programmers would solve a programming problem in relatively similar ways.
4. Levels of detail and presentation format can vary enormously in workflow maps, whereas flowcharts must provide sufficient detail to write code, and must be consistent with the rules and syntax the computer language requires. Workflow maps can and should be tailored to the level of detail most useful and informative to the user.
5. Workflow mapping is an iterative process of discovery and investigation, rather than application of externally determined coding rules; the objective of workflow mapping is to develop a valid map of a process. The objective of flowcharting is to develop an efficient, correct answer to a programming problem, and by definition that means it has to run.

Business Process Reengineering

For better or worse, WFMA data collection and analysis is often associated with Business Process Reengineering (BPR), which became very popular in the 1990s. BPR, in many ways, is a view of organizations in which process automation replaces the information processing capacity of humans, and in which their tacit knowledge is implicitly of little or no value. The slogan, Hammer used, during the early days of BPR was “don’t automate, obliterate,” the title of his 1990 article in *Harvard Business Review*. During their presentations on BPR, it was not uncommon for Hammer and Champy⁷ to simulate pointing a gun at a nonbeliever’s head and say “Bang! You’re gone!”

BPR was seen as useful for some companies, but certainly not all.⁸ It quickly became seen as a code word for “downsizing.” The immediate perception of everyone working in a company where BPR was begun was

that many of them were going to lose their jobs, and quite often that is what happened. Not only were individuals “laid off” (and it was then that “laid off” became a substitute word for “let go”), but entire levels of organizations were removed. One acquaintance who went through this process twice likened it to the Grim Reaper being given a broadsword, and entire levels of management and staff were “laid off” *en masse*, as if the sword were slicing the middle layers out of a cake in two quick blows.

The popularity of BPR was enhanced by the effects these changes had on the bottom line, and of course, not having to pay a significant number of people cuts costs immediately, so the savings quickly showed up on the income statement. Nevertheless, BPR and the frequently thoughtless removal of entire layers of organizations was shown, over time, to cause organizations to “lose their minds”—companies found that they could not cope with exceptions or problems, that overworked survivors could not handle all the work suddenly heaped on them, and that when mistakes were made the emphasis was on hiding them rather than fixing them, lest the offender become the next to be “laid off.”⁹ The long-term costs were huge, despite the immediate benefits to the bottom line, and BPR had long-term benefits for companies far less often than desired. Hammer and Champy have since acknowledged that they underestimated the effects of the human factor in BPR, and soon changed their approach.¹⁰ But the identification of process study with layoffs and downsizing is now entrenched, and anyone who starts WFMA and does not address this issue is making a mistake. Even the most positive stories of successful process change and business turnarounds are saddled with the “chain saw” image.¹¹

BPR still has a large following, and undoubtedly brings value to many of its users. The marriage between BPR and information technology (IT) has continued unabated, for better or worse, and with that comes a set of challenges unique to the functioning of IT. For example, Branco et al. found a large proliferation of BPM models in one bank he and his colleagues studied, and the challenge was maintaining consistency between them as they changed and evolved.¹² A similar finding was reported by Wang et al., who have found it necessary to develop a model for querying business process repositories.¹³ vom Brocke et al. report that the enterprise resource management firm SAP has been able to greatly reduce the time required

for “big data” analysis by adopting in-memory technology, and used this to good effect with Hilti Corporation.¹⁴ As these studies suggest, the major payoffs of BPR are achieved in large organizations that are able to support the specialist staffing and expense involved, but given their complexity, this is exactly where a marriage of process mapping and IT may be needed. A recent study by the Forrester organization¹⁵ strongly implied that meeting the collaborative and competitive-advantage challenges created by the Internet of things required unprecedented understanding and integration of their processes, often across national as well as organizational boundaries; again, this is fertile ground for the integration of process mapping and IT.

As discussed in this and previous chapters, there are many applications for WFMA. Just because the initiator of a particular WFMA project knows that its objective has nothing to do with staff reductions doesn't mean that everyone else in the workplace will. The reality is that WFMA typically identifies a number of opportunities for improved efficiency, and unless there is a clear need for additional staff, the result may be some level of staff reduction. Such possibilities are often an unintended consequence of WFMA, but they cannot be ignored, and should never be treated as if they do not exist.

Time Study

One additional comparison that should be added to our discussion of WFMA is that when people map the flow of work through their workplaces, they sometimes confuse workflow mapping with the study of how daily time is spent on the job, or a “time study.” These are *not* the same. It is obvious that time is a universal resource, and is consumed in every aspect of work; it is therefore logical that how time is consumed in a workflow is a nearly universal metric of interest.

However, WFMA and time study are quite different. The objective of workflow mapping is to identify a valid sequence of activities and decisions used to perform activities or functions in a *flow of work*. That flow may be contained in one office or workplace, of course. However, most people work on several types of flows in their jobs (not just one), and what is needed in the end is the ability to see how a process integrates

activities and decisions to achieve some end result. How long that takes is only part of the question and may or may not be relevant to present WFMA objectives.

A process might be entirely contained in one person's workplace, as we have noted, but it might cross several workplaces for all activities in it to be completed. A well-known University of Cardiff study (reported by Carson¹⁶) investigated how long it took to get a can of Coca-Cola to a supermarket shelf. The entire process took 319 days (to make the can and everything in the soda), going through 14 batch and hold steps until completion. The time to actually make an individual can and its contents was three hours. WFMA looks at the whole process, and can also identify and differentiate between the 319 days and the three hours.

Summary

This short chapter has expanded on the procedures recommended for WFMA as well as some of the difficulties one might encounter when doing it. In short, this can be reduced to one overarching rule: stick to the symbols and the discipline, and only deviate from them when there is truly no other choice. Adherence to this rule will always result in process maps that are complete, correct, and accurate. They may not be the processes that the organization ultimately wants, but that is a matter for attention once the maps of what is actually being done have been produced. The Kmetz process-mapping method is powerful because it is simple, and that simplicity needs to be protected.

Like the symbol set, these rules are simple; however, following them consistently is often not. The discipline needed to make WFMA effective is to apply these rules strictly, and not to allow variations to creep into either the maps or the procedures that produce them. If this is not done, any moderately large firm, or firm with geographically separated offices, will inevitably begin to develop divergent and incompatible WFMA techniques, and the power of WFMA will be lost.

PART II

Why Do Mapping?

CHAPTER 4

Knowing What We Know

Knowledge and Information: The Framework

The first three chapters of this book give the “nuts and bolts” of process mapping for those who need to get going with a quick start. Chapters 4 and 5 develop matters underlying process mapping, which turns out to be somewhat more complex than first meets the eye. Why do we have so many variations on the way a workflow might be correctly mapped, and why do these variations arise in the first place? These and other questions lie beneath the surface of what we know and do if we want to correctly map a workflow process.

This entire book is about organizations and understanding how they work. A fundamental idea that we have introduced is that of a “process” because everything that organizations do, in one way or another, can be described as a process. In organizations, we design processes to accomplish specific goals, and I find it useful to think of these designed processes as “workflows,” the term in the title of this book.

Organizations and workflows are both critically dependent on information; in fact, I am going to argue that organizations themselves are information processors, in a very fundamental way. The information they process consists of two main types—information about what we are producing, whether tangible goods or intangible products like services, and information about how we do that. The latter can be thought of as information necessary for coordination, which is an absolute requirement in organizations because their reason for existing is to do work that is beyond the capability of a single individual. Coordination requires both formal and tacit knowledge, two other key terms we have heard often in this book.

The “workflow” used by an artisan to craft an item of jewelry is entirely up to that artisan; as soon as the artisan hires help for that flow

of work, however, it becomes necessary to think about who does what, in what order, what happens when things do not go as planned, and much more. That is what we mean by coordination, and there is no escaping the need for it in organizations. Some of this information may take the form of formal rules and policies, but a great deal of it is individual and worked out on the basis of day-to-day interaction, the way that masters and their apprentices did for centuries. For every person added to the organization, the coordination requirements go up geometrically—coordinating four people takes much more than twice the information processing needed for two.

Organizations do not “just happen,” and therefore neither do workflows—they are designed. Neither one is static or unchanging over time, so what was designed at one time will need to be modified in the future; we are constantly changing organization structures for one reason or another. One consequence is that the workflow designs that made good sense at one time no longer do, but they persist and often become seriously out of whack with the goals of the organization.

So, if there isn't a “Second Law of Organizational Thermodynamics,” there should be. In physics, the Second Law of Thermodynamics says that everything eventually winds down until energy is evenly distributed throughout the universe and everything comes to a stop. Entropy rules! My experience with organizations suggests they follow this law, and this chapter is going to present some underlying reasons why this is the case. We build on this in Chapter 5 to explain the basis for the tools and techniques we have seen in Chapters 1 and 2; combined with this information, these tools can not only help us to manage some of the chaos but actually change and improve processes and performance.

Understanding workflows also requires some fundamental understanding of information and knowledge, and in this way knowledge management (KM) is related to workflow mapping. Given this relationship, one payoff is that mapping by the Kmetz method becomes a valuable way of capturing both formal and tacit knowledge in the workflow. We will discuss KM in more detail in Chapter 6—our immediate concern is to know more about “information” and “knowledge,” two words that we use all the time but seldom appreciate for their richness and complexity.

Knowledge Is Information Is Knowledge

I want to begin with an idea that in some ways is the entire point of this chapter. In the perspective of this book, information is “knowledge,” in the sense that it is a product of human intellect; it is structured, rather than random; and it is communicable to others.¹ This is a utilitarian perspective on the definition of knowledge, in the sense that if you do not know you have information, then you do not have it. Two simple models help explain this perspective.

The first model posits that all information (and thus all knowledge) can be represented by a simple 2×2 framework, shown in Figure 4.1.² This simple model categorizes all information into one of four cells. Known knowns (KK) are those items of information we consider to be “facts,” or to which we attach so little uncertainty as to make them effectively factual; known unknowns (KU) are essentially questions we know to be unanswered. Unknown knowns (UK) are information which we may have but cannot unambiguously interpret—a classic illustration is the problem faced by intelligence analysts, who are confronted with myriad facts that cannot be easily evaluated for truth or accuracy, or what they collectively mean. The final cell comprises unknown unknowns (UU), effectively an undefined area of information, the existence of which might be surmised but cannot be forced to yield to analysis—for example, what is the likelihood that a specific person will break his or her left leg in


		State of the universe	
		Knowns	Unknowns
State of our information	Known	Known knowns	Known unknowns
	Unknown	Unknown knowns	Unknown unknowns

Figure 4.1 *An exhaustive model of states of information*

exactly 27 days; the probability that the Yellowstone volcano will erupt with the same force as its last eruption (and on a historical basis, it is due) and potentially end advanced civilization; the odds that we are actually on a surface in 11-dimension space-time, and that none of the universe we see can even begin to be understood in the four dimensions of space and time as we know them? All of these are serious questions, but with the exception of theoretical physics we have no way to frame a serious question in terms that we can comprehend, let alone get a meaningful answer.

Figure 4.1 provides a way of characterizing the overall state of the information we have in terms of both knowns and unknowns. The contents of these cells are not the same for different observers, however, because information is a product of human intellect and dependent on the observer. One aspect of this content is that for each observer, any item of information may be described as a “vector,” which is our second model. In the terminology of linear algebra a vector may be thought of as an expression of a single path through a multidimensional matrix. In terms of human experience, at least seven properties of any item of information might define a vector, as shown in Figure 4.2.

The “vector” in Figure 4.2 is the dotted line connecting each scale or continuum for seven properties of information. Each of the seven properties is an opposite pair (true–false, consistent–inconsistent, etc.), where the extreme end of each scale might be defined by the associated word—for example, only statements at the extreme left of the first continuum are really “true.” Where the dotted line intersects each of the



“Information” is a vector representing the location of an information element on all of these properties, each of which is a continuum.

Figure 4.2 Properties of information

Source: Adapted by permission of the Publishers, in *The Information Processing Theory of Organization* by John L. Kmetz (Farnham: Gower, 1998), p. 16. Copyright © 1998.

scales defines the value of the vector for a specific item of information as seen by one observer.

How could information have a vector like that shown in Figure 4.2—in particular, how can information be partly true and partly false, as shown on the first continuum? Consider the following statement: “I don’t know whether to believe them entirely or not, but the numbers coming out of the rare-earth explorations we’ve been doing at Site X, even though they don’t agree with a number of other prospectors who’ve looked around the same area, seem to make a strong case for spending some serious development money there.” Going from top to bottom of Figure 4.2, all seven vector properties are embodied in this statement. This is the way information usually comes to us—it is a bundle of qualities that are not necessarily reconcilable with each other, let alone the basis for a firm conclusion or immediate action. We literally need time and thought to figure out what we consider to be a KK.

Moreover, it is highly unlikely that any two individuals will perceive an item of information in identical terms for each of these vector properties; that is, the *meaning* of information (“knowledge”) to one person will inevitably not be the same as for another. Depending on where one person considers an item of information to fall on each vector, a bit of information may be considered highly credible and be placed in cell KK in Figure 4.1; another observer who evaluates the vector properties for that item differently places that item in cell KU. For example, “source credibility” will strongly affect where one places information on these continua, as any follower of marketing or political science can easily attest.

The idea of “known knowns” may ultimately be an oversimplification. Very few things are truly “known” in the sense of being fixed and final—courts review verdicts, analysts recalculate the books for businesses, research outcomes are reviewed, and so on. Because information is a function of both inherent content and human perception and processing, everything is subject to reinterpretation. Much of the tacit knowledge in organizations is derived from these kinds of highly individual processes.

What both of the models in Figures 4.1 and 4.2 emphasize is the importance of thinking about what we know, and also about how much confidence we have in that knowledge. In his highly recommended book, *The Black Swan*, Taleb points out a number of very important

characteristics of human information processing which may lead to error in our conclusions about things.³ We have a tendency to “tunnel,” as he terms it, to look at one or more sources of information and disregard others. An immediate implication is that we need to be as receptive to information as we can, perhaps especially to that we do not really want to hear. The absence of information itself may have value—absence of evidence on a subject is not the same as evidence of absence. We are also strongly persuaded by stories, or “narratives,” which often have the property of making rough knowledge appear to be more smooth and complete than it really is if we look at it closely. These two figures give us some powerful ways to think about what we think we know. In terms of Figure 4.1, categorizing important information on a performance problem into three of the four cells of that model can be a high-payoff application of it, as may evaluating arguments on the basis of relevant vectors in Figure 4.2.

So does this mean we never really know anything? Perhaps in the philosophical sense it does, but in the world of working organizations we deal with the complexities of information differently. Much of what we “know” is a social reality, meaning that through usage, experimentation, and learning, we come to agree on what something “means” to the extent that we can use it as if it were a KK. Working knowledge evolves. Most formal policies and procedures develop in response to a perceived need, to fill a vacuum when it becomes evident; they are changed and replaced in the same way.

Tacit knowledge does the same thing, only on the part of individuals and small groups. Tacit workflow knowledge develops in the environment of formal organizational knowledge, which has many implications (one of them being the old bromide that “we get things done around here not because of the rules, but in spite of them”). So we may not have final answers to anything, but we agree on the information we need to make progress, and that information always includes tacit knowledge.

All Information Is Imperfect

What constitutes KU or UK in Figure 4.1 depends considerably on the individual making the judgment about the contents of these cells. As a commitment to faith, one observer may reject the entire construct of

Figure 4.1, since it rejects the potential for all unknowns to rest in the hands of a higher power. Over time, each of the cells with known elements is a “fuzzy set,” in that the content and classification system may change. A humble example of this is the definition of “dishwasher safe” kitchen equipment and cutlery. As a wooden-handled knife (located in the KK cell as not “dishwasher safe” when acquired) becomes older, duller, and less prized, it is less likely to be hand-washed and more likely to be put in the dishwasher; “dishwasher safe” is partly a matter of who makes the determination as well as the physical properties of the item. Therefore, “information imperfection” can be summarized as either a problem of *incompleteness*, where at a minimum the UU cell in Figure 4.1 can never be eliminated, or as a function of the fuzzy set problems induced by *vector properties* shown in Figure 4.2.

Either of these two forms of information imperfection may be the product of active or passive sources, as summarized in Table 4.1. These may result in simply incomplete information or differing vector properties.

Active processes are shown in cells 1 and 3. In cell 1, active distortion of information or misleading information may be provided by a competitor as a deliberate method for concealment of strategy or intentions; in cell 3, various kinds of analytical error may result in imperfect information—these could include incorrect weighting of information content, misinterpretation of vector properties, and simple mathematical error. Passive forms of imperfection are shown in the other two cells, and are relatively straightforward—the lack or loss of information in cell 2, rendering what we think we know to be incomplete, and the unconscious

Table 4.1 *Forms of information imperfection*

Form	Source of imperfection	
	Active	Passive
Incompleteness	1. Misinformation, disinformation, “jamming”	2. “Uncertainty,” lack of information, signal loss, or noise
Vector properties	3. Analytical error	4. Values, feelings and emotions, source-specific responses, culture

Source: Reprinted by permission of the Publishers, in *The Information Processing Theory of Organization* by John L. Kmetz (Farnham: Gower, 1998), p. 17. Copyright © 1998.

filtering of vector properties or addition of unintended vector properties to information in cell 4. These four archetypal processes are interdependent for any observer; for example, jamming information about a source (person) may create emotional filters that affect the vector properties of all information from that source. Examples of such interactions can easily be imagined for all four sources of imperfection, and these interact over time.

We hear all the time that knowledge is often the most critical asset any organization possesses. The lengths taken to protect the formula of Coca-Cola, to protect innumerable trademarks, and the global concern over protection of intellectual property are abundant testimony to that fact. Thus, the cells in Table 4.1 where information is incomplete may be the product of active processes on the part of external agents who do not want knowledge to be full or complete, in addition to imperfections from our internal thought processes.

It is also important to recognize that actively derived imperfections do not necessarily imply bad intent. Businesses keep at least three sets of books—one to report to shareholders, one to use for internal decision making, and one for tax collectors. While we might view this cynically and suggest that each is intended to keep information away from people, it is equally true that compliance with a hugely complex tax code may not always tell the most accurate story of how the business is doing for the shareholders, and that neither of these is what a manager needs for day-to-day operations. Changing the way we keep accounts changes the properties of the knowledge we have to work with, and we need to actively create different versions of a single “truth.” In the wrong hands, of course, this same need opens the door for the kinds of abuses we have seen with the Enrons and WorldComs of the business community.

Information imperfection is a major issue in the mapping of workflows, as we have seen in Chapters 1 and 2. This imperfection of information illuminates the importance of “slow thinking,” the kind Kahneman refers to as System 2 thinking.⁴ System 2 is slow, deliberative, and logical, and results in very different outcomes when compared to System 1, which is more intuitive, faster, and more emotional. In many ways, the WFMA techniques in those chapters might be considered as “forced” System 2

thinking in a specific framework, in which deeply held tacit knowledge is made visible.

Much of the tacit knowledge in a workflow becomes so deeply embedded in individualized behavior that it becomes difficult to extract. Everyone has had the startling experience of having driven a long distance without really being aware of it until some point near the end. We overlearn a familiar route to the extent that conscious attention to driving it is not necessary, and we navigate by using waypoints and landmarks; if we are asked how we travel, we suddenly realize we no longer know route numbers or street names, but these landmarks. The same thing happens in our work, and a type of “uncertainty” is the inevitable result. But we should be aware that in cases where people feel threatened by a new workflow-mapping project in their company, they may respond by engaging in “jamming” and providing disinformation.

Organizations Are Information Processors

When we talk about “organizations,” the type of social creation we focus on in this book, that word conjures up many images. The one that I personally prefer is that an organization is an information processor. Organizations, small or large, are groups of people using various technologies to accomplish something that cannot be done through individual effort alone. Because we have multiple players, different materials, different objectives, different constituencies, and all the myriad things that come with an organization, it is necessary to process information to coordinate everything that has to be done.

This need is easy to understand. A small team of people can coordinate their actions for a small project relatively easily (especially since they are likely to be self-selected members for the job at hand). They simply ask questions and make suggestions to each other as circumstances require, and with everyone in contact with each other, processing information to coordinate the team is easily managed. But when a job gets bigger, takes more people and more specialized skills, extends over a long time, and so on, the capacity for information processing activity sufficient to coordinate a small team will simply not be adequate.

The solution to this problem is also easy to understand—we break the big organization down into smaller groups (typically by the type of skills people have or the type of output they produce), and have a specialist in charge of each group, so that the amount of information that has to be processed within each group will be dramatically reduced relative to the whole organization. Each small team will only have to coordinate its actions within the group, and between-group coordination can be done by the team leaders. They may need a higher-level team leader, and if so, we have just created a three-level hierarchy.

This spontaneous hierarchy is hardly new, and the discovery of the information-processing efficiency and effectiveness of the hierarchy is as old as organization itself. The Romans are often credited with invention of the hierarchy (a “centurion” was the leader of 10 groups of 10), but hierarchical military organization was used by the early Assyrians, Genghis Khan, and the Mongols, among others. This efficiency is also why the hierarchy is durable, despite the efforts of many thinkers and advocates of alternative forms of organization to discredit it—it persists because it works.

When we design an organization, we have a significant impact on the way the organization will be able to process information, and how much and what kind of processing it will have to do. There are significant trade-offs. If we organize our basic units by skill or type of work, as opposed to grouping people and skills around production of a particular type of output, we create specialist units that tend to pay most attention to their specialization, and often lose touch with the customer; organizing by product may keep us closer to the customer, but at the cost of losing our skill (and innovative) edge. If we make the hierarchy tall and keep all the decision power at the top, it makes it easier for the whole organization to adapt its overall goals over time, but at the cost of “buy-in” and much valuable knowledge that stays at lower levels; if we reverse that and keep decision making at lower levels, we risk having the overall goal lost in the cracks between goals of the individual business units. These are never-ending problems, and they are a constant challenge to large organizations because the trade-offs between them are important.

As an example, Gore Associates, the maker of Gore-Tex and many other nonconsumer products, decided to commit to an organizational

form that reflected William Gore's experience and preferences from his early career in a large company. His decision was to form production units of 200 or fewer people, and when a site grew beyond that size, he opened a new physical unit—a new plant at a new site. The reason was that in his early experience, large organizations always lost touch with individuals and were simply not much fun; he wanted plants small enough to let everyone get to know everyone else. He also abolished hierarchy and status differentiation, so that everyone who works with Gore is an "Associate." As a result, this global company now has small units in roughly similarly sized buildings scattered around the globe. They have their own unique problems in trying to coordinate this kind of operation, but have learned how to do it successfully through several long-term business cycles and the end of patent protection for a major product line.

The scattering of task-related information through an organization thus induces a new type of information imperfection. Anyone who has been in an organization knows that keeping both the left and right hands informed of what the other is doing is an endless job; moreover, what is "important" at any given time depends considerably on the point of view of both the individuals and the unit they represent. It requires time and money to process information to achieve functional consistency, where goals can be met with enough success to keep the wolf from the door over the long term.

Figure 4.3 illustrates the problem in general terms. Knowing what we know is not free. From the vertical axis on the left, two organizations (A and B) might start from much the same level of internal information consistency, and both might agree that this level of consistency is inadequate for their needs—they need to "get on the same page." Doing that requires time and money, and as they move through time to the right, they improve their consistency, but at increasing cost. (It is also worth noting that the problem they are working on becomes less and less current.)

The two organizations may start in a similar position with respect to their internal degree of functional consistency, which might be thought of as increasing the relative size of the KK cell in Figure 4.1. To increase the size of the KK cell requires effort and expense, as does further resolution of the UK and KU cells. How much processing toward these outcomes is

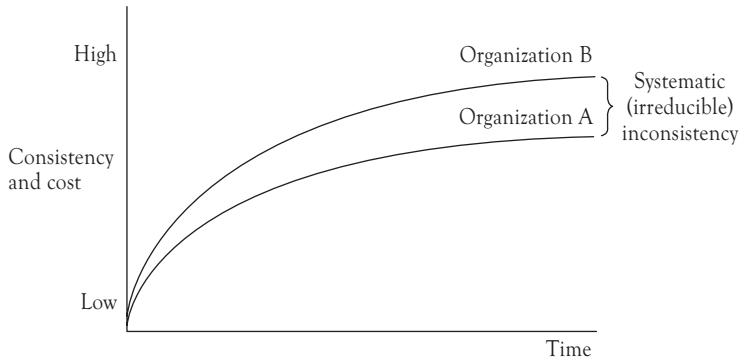


Figure 4.3 Functional consistency lag and cost

Source: Adapted by permission of the Publishers, in *The Information Processing Theory of Organization* by John L. Kmetz (Farnham: Gower, 1998), p. 352. Copyright © 1998.

justified, and how do we know? To what extent is acquiring the n th item of information worthwhile? What is the cost of the time to do this, and what is the time value of money relative to all of these tasks? These are questions that are fundamental to any organization design. The existence of the UU cell means that there is also an irreducible system-level cost, where spending infinite amounts of money will not gain much by way of new information.

Several decades ago, Aaron Wildavsky coined the term “uncertainty absorption” to describe what happens when information in raw or nearly raw form enters the organization, and decision makers have to deal with the unknowns and imperfections in it.⁵ Raw data and information at any level of an organization is partially a mess, and what to do in the face of a problem is frequently not clear. Wildavsky argues that managers acquire much of their information as summaries of it from the level below (an interesting process in its own right), and use this to make decisions that are hopefully consistent with the organization’s goals. From the subordinates’ point of view, once the management has made a decision and passed it down, uncertainty about what to do has been absorbed by the manager for the subordinate—I may or may not agree with management’s decision, for example, but my job is to comply with it. This pattern is repeated through all levels, with all the potential for organizational politicking and infighting one could imagine. Such battles need not originate from an outside problem—having been through several

wars over technical design in the aerospace industry, I can personally attest that the technical battles in the labs and engineering divisions are as bloody as they come.

Against this backdrop, it should come as no surprise that organizations are destined to constantly struggle with the problem of internally getting their act together. Testament to the difficulty of this job is provided by the popularity of the *Dilbert* comic strip, which parodies the role of managers and the problems of running a company (a former MBA student, an engineer by training like Dilbert, once told me in all sincerity that all he needed to know about management in the real world could be learned by regular reading of the strip; the strip author, Scott Adams, has often noted that most ideas for his strips are sent to him by readers in the working world). The challenge to managers is that no matter how hard they try, there will always be some things that slip through the cracks.

Even in the world of international spying and intelligence, the fundamental need for effective information processing cannot be escaped. The “how did we miss that?” or “how could we not have known?” investigations that often follow intelligence failures are as predictable as rain. Even attempts to resolve the problem by restricting information access fail—they only create different types of performance and coordination problems. Harold Wilensky made an observation in his 1967 book that is as true today as ever:⁶

The more secrecy, the smaller the intelligent audience, the less systematic the distribution and indexing of research, the greater the anonymity of authorship, and the more intolerant the attitude toward deviant views.

Organizations must constantly struggle to “get everyone on the same page,” and it is a never-ending battle. Organizations are constantly restructuring, and by some accounts the average time between reorganizations is at a record low. All of this is driven by the need to process the right information in the right place at the right time, and stay competitive in a rapidly changing world. From basic hierarchies with a decision maker at the head of each group, we experiment with pre-made decisions in the form of rules, policies, and procedures; we add staff specialists to

take some of the processing burden off line managers; we add information technology; we split up by region or customer group or on some other basis that makes sense in our industry. Every one of these changes has an impact on our processes, and it is not uncommon to find that a large part of the body of tacit knowledge is directed toward patching the cracks in the workflow left by the last reorganization. Most of these “patches,” actually, are taken care of through voluntary action on the part of employees, who use their tacit knowledge of customers and situations to fix things when they get out of whack.

Organizations Are Systems

The next of the fundamentals to discuss is that organizations are “systems.” On one hand this will seem intuitively obvious when explained, but on the other, it is a powerful and useful way to think about organizational processes and the critical roles of formal knowledge and tacit knowledge in them.

The idea of a “system” that I am using here is based on the everyday observation of complex, organized, often self-regulating entities in the world around us. “Systems” are the subject of a body of knowledge known as General Systems Theory, and they are formally defined in several ways, but they all build on the idea that a system is a whole made of component parts, and which is relatively stable and is both recognized and functions as a whole. There are four properties associated with “systems” as they are defined in General Systems Theory: (1) the whole is greater than the sum of its parts; (2) the whole determines nature of the parts; (3) the parts cannot be understood in isolation from the whole; and (4) the parts of the system are dynamically interrelated—they are interdependent and interact with each other over time.⁷

Much of this sounds theoretical, but the essence of these definitions is captured in my own somewhat tongue-in-cheek definition: “A system is a thing made up of other things, all connected to each other and all other things.” Examples are everywhere. A person is a system; so is a town or city, on a larger scale, or a gut bacterium, on a much smaller scale. What is evident from consideration of these three systems is that any system is on one hand a “subsystem” of a larger entity, while at the same

time a “supersystem” for smaller entities within it. Bacteria in the human gut are independent organisms on their own right, but as subsystems of a human body they are essential, and without them the survival of the human would be impossible. A political entity like a town or city has specific governing bodies that give the entity of “town” the ability to regulate behavior, repair itself, protect itself from hazards, and so on, even as its human subsystems come and go.

In one respect, I tend to prefer my informal definition because it forces one relationship to the forefront—the relationship between the system and its environment. The human body is the environment for the bacterium, as is the town for the citizen. Each system we examine has this relationship to larger and smaller entities, internally and externally. Where “the system” ends and “the environment” begins is a matter of both scientific and philosophical debate. This is something we will not attempt to resolve here, but it has significant implications for the way we define any system we want to examine through workflow mapping.

Figure 4.4 shows the basic relationship a system maintains with its environment, and this will be discussed in more detail in Chapter 5, particularly what goes on inside the system. All systems have a permeable “boundary” with their environments, and take inputs, transform these using internal processes into outputs, and return these outputs to the environment. As I suggested above, the boundary is not fixed or impermeable, and how we define that boundary may have significant implications for workflow characteristics.

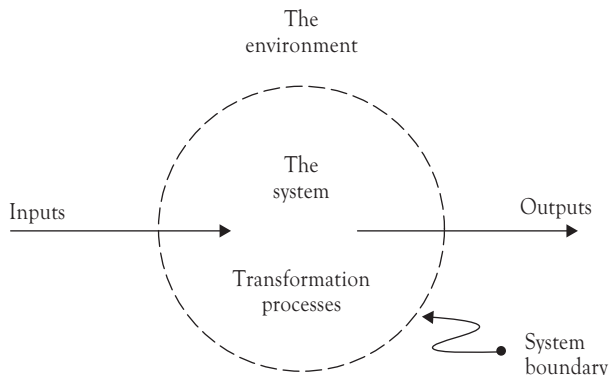


Figure 4.4 *A system and its environment*

Examples of the importance of these boundary relationships can be found in a modern supply chain or Just In Time (JIT) management. In order to make these methods work, companies must share information with “outsiders” on a level that a few years ago would have been considered an unacceptable breach of corporate confidence. What was “the organizational environment” a few years ago is now part of “the operating system,” and cannot work any other way.

Understanding organizations as systems has a number of important implications for understanding what goes on inside them. First, the inputs to one system are the outputs from one or more other systems. We sometimes think of knowledge as an economic “stock” of information, to be categorized and accounted for as a bakery would with different flours being prepared to make bread and pastries, or in other cases, knowledge is treated more as an input variable, without regard to its source(s). My approach to capturing knowledge does not treat it as only a stock or an input flow—rather, knowledge takes on both roles in different times and circumstances in a workflow.

What differentiates the system and its environment is often worth careful consideration. Most companies (and perhaps most organizations in general) want to be as selectively open to the outside world as they can, while at the same time protecting the intellectual property (“knowledge base”) that makes them successful at what they do. This creates interesting problems and interesting opportunities. A few years ago many companies that used telephone back-office customer support felt that it was a no-brainer to take those functions offshore; what was once considered a necessary “internal” part of the business had become redefined as a routine function that could be done by contract employees on another continent. Since then, many of those firms have had to rethink that decision insofar as critical functions and customers are concerned. What constitutes a core body of knowledge, how porous the boundary should be, and how a company manages the relationships between them is quite important to workflow design and performance.

The environment is not static. Companies and organizations must adjust to shocks and environmental disturbances all the time, and many of them show remarkable resilience. In addition, there are things that

go wrong internally for all manner of reasons (“exceptions,” since they were not what we planned), and we have to adjust to these as well. Both types of adjustments require improvisation, jerry-rigging, and the like; they are heavily dependent on the expertise of people at the scene, at the time. These adjustments often become institutionalized because they worked, and they are both an important form and important source of tacit knowledge. From the perspective of the system, however, these tacit-knowledge adjustments are often nearly invisible, simply because they worked.

Another property to recognize is that complex systems, like companies and organizations, exhibit a high degree of self-regulation and adaptability. These properties are critically dependent on the knowledge base within the organization, which itself has to change and adapt as both internal and environmental forces require. Every individual and every group or unit within an organization possesses bodies of formal and tacit knowledge, the latter often a large body. These not only enable the organization to meet its immediate objectives, but to regulate its processes to do that and to react to problems in its workflow and change as necessary. Indeed, Senge and his colleagues argue that mastery of these knowledge bases and the ability to learn over time is a major competitive advantage and a requirement for long-term survival.⁸

Both the self-regulation and adaptability of complex systems depends on what may be thought of as the economic “stocks” of information mentioned earlier, and also “flows” of information. Every part of an organization depends on a knowledge base of formal knowledge, which is principally focused on the technical aspects of work; this knowledge base consists of many components, each of which is closely associated with the differentiated units that make up the organization. Each of these units applies its knowledge to the material in the flow of work, transforming raw inputs into final outputs. But much of this is heavily dependent on the tacit knowledge base, which is partly brought to the organization by its members, and partly created within it as the members interact with each other. It is primarily in this tacit knowledge base where we find “flows” of information, in the broadest sense meaning any information mobilized or used in a way that the formal knowledge base could not

anticipate. Much tacit knowledge is also associated with units of the organization, but much is not, and it is free to “move” and be applied when and where it is needed to make the organization flexible and adaptable.

There is a good bit more to say about organizations as systems, but for the present time we should appreciate that “internal” processes are part of the “connected things” that make up a system. The inputs we bring into the organization from its environment not only include information about suppliers, markets, and so on, but the people who process it; they bring with them many other information inputs, along with a body of skills and interests. Some of these are unknown when we hire them and have unanticipated impact on the organization’s information processing—they are both an input to carry out formal processes, and a stock of their own knowledge which will influence how they do these processes. To fully understand an organization requires recognition of the openness of the system to its external environment as well as the full extent to which formal and tacit knowledge are necessary to meeting its goals. We will expand on this idea in Chapter 5.

Information, Processes, and Performance

To pull the previous four points about organizational processes and information together, we need to consider the relationship of these to performance. Performance or goal attainment is not a foregone conclusion in a world of imperfect information, and this is one of the reasons that we often discuss performance in terms of the *degree* of goal attainment. The linkage between what an organization plans and projects on the one hand, and what actually happens on the other, is neither a sure thing nor a straight path. Thus, any discussion of performance must take factors that cause performance variations into account; these variations and bends in the path to the future necessitate information processing, just as elements of other organizational processes do.

At the same time, the variability of process outcomes and the fact that we are always dealing with imperfect information makes it difficult to rigorously link performance to information. For example, the ability to demonstrate the payoff of investments in information technology has been a major challenge for decades. Strassmann argued that much of

the early investment in information technology failed because it simply automated obsolete methods of doing work.⁹ Since then, information technology has been argued by some to be a key to the rapid increases in productivity of the U.S. economy during the late 1980s and early 1990s (Farrell et al.).^{10,11} But the time lags and lack of one-for-one correspondence between variables in a complex system always make such relationships difficult to identify or measure.

The desired or planned level of performance for a company or organization might be thought of as the outcome that would be attained under conditions of perfect information—but we know that is impossible because we have only imperfect information. Imperfect information causes deviations from the outcomes that we would obtain with perfect information, in the form of both gains and benefits on one hand, or as costs and losses, on the other. Considering both positive and negative outcomes caused by imperfect information, the performance of an organization may be described in terms of the following relationships:

$$\text{Performance} = \begin{array}{l} \text{outcomes as planned or} \\ \text{projected (assuming perfect} \\ \text{information)} \end{array} + \begin{array}{l} \text{net payoff of outcomes} \\ \text{resulting from imperfect} \\ \text{information} \end{array}$$

What is that last term on the right—the “net payoff of outcomes resulting from imperfect information”? As shown in Table 4.2, this payoff is the net value of all benefits or gains, and costs or losses, resulting from both proactive and reactive organizational responses to imperfect information. Companies cannot simply sit and wait for everything to be known, so we take both proactive and reactive steps to deal with risks and unknowns. Costs or losses may be incurred whether the organization attempts to deal with imperfect information through proactive steps, such as planning, market research, and forecasting; or they may be incurred through reactive steps, such as missing market share or having to correct or compensate for the costs of delay. In either case, there are planned costs for coordinating organizational activities in the face of this imperfect information, and there are unforeseen costs and losses. Similarly, benefits may be gained from both proactive and reactive approaches to dealing with imperfect information, either through gains from anticipation and exploitation of new opportunities and competitive advantages,

Table 4.2 Positive and negative outcomes as a function of imperfect information

Response mode	Payoff value	
	Cost or loss (–)	Benefit or gain (+)
Proactive steps	1. Planning and forecasting of future outcomes; coordination costs	2. Avoidance (errors), anticipation and exploitation (payoffs)
Reactive steps	4. Opportunity costs of foregone outcomes and payoffs; coordination costs	3. Passive opportunism (a.k.a. IIABDFI—If It Ain't Broke, Don't Fix It)

Source: Adapted by permission of the Publishers, in *The Information Processing Theory of Organization* by John L. Kmetz (Farnham: Gower, 1998), p. 43. Copyright © 1998.

or through the avoidance of costs for unnecessary information and information-processing activities. An organization of any size usually does most of these things, and obtains many individual payoffs. The sum of all outcomes in cells 1–4 makes up the net “payoff” of imperfect information.

Consider the payoff of what I refer to as “coordination costs.” These costs may be to acquire information for decision making, or may be the costs of tightly coordinating activities within and between organizations. Again, an excellent example of the latter is Just In Time (JIT) vendor–customer relationships, where considerable initial cost is borne by both parties to tightly coordinate their production and logistics flows across company boundaries. The net payoff of that investment in JIT, however, is so great that for many manufacturers any other approach to doing business is inconceivable.

But imperfect information often pays off in terms of benefits. For those companies able to find a competitive advantage in their technology or market niche, returns far above those obtained by competitors may be earned. For those who adopt a wait-and-see approach to dealing with unknowns, problems often go away and the unnecessary costs of coordination and attempted mastery of new technologies and new markets are avoided—“if it ain't broke, don't fix it.” Of course, many firms using either approach guess wrong, and fail—neither proacting nor reacting are totally free of risk.

An interesting implication of all of this is that companies can adjust to the challenges of imperfect information through *lowered* performance,

that is, if an organization lacks the information processing capacity to cope with all its knowns and unknowns, then an adjustive reaction is to reduce the level of output relative to what it might have been with adequate capacity. The hard question in this is “what might have been,” either in terms of opportunity costs or forgone benefits. Most organizations would not choose to lower performance levels, but many do so by not knowing how formal and tacit knowledge interact in their workflows.

The obverse, of course, also holds—if information processing capacity is *increased* in a system, then at a later time there should be a measurable increase in performance, which has clearly been the argument of both the information technology and business process consulting industries over the years.

The bottom line to this is that organizations are systems that function through information processing and what we know about the formal and tacit aspects of this in our workflows has both direct and indirect effects on performance. If this seems obvious at this point, that is excellent; if not, we need to be clear about this fundamental point, which we will expand on in Chapter 5. For now it is necessary to recognize that information is both the stuff of much organizational work, and the “glue” that holds the organization together so that it can work.

Summary and Implications

In some ways, it might be appropriate to return to Figure 4.1 and use that as the summary of this entire chapter, since the real issue is, as the chapter title says, “knowing what we know.” By now it should be clear that this is a more complicated question than it might first seem, and that realization is a good thing.

It is a good thing for two major reasons. First, a fundamental assumption of this entire book is that as organizations change and evolve over time, their internal processes need to do the same. Much experience has shown that this evolutionary change affects not only the overall structure of the organization, but has many subtle and frequently unknown effects on the workflows within it. Indeed, in later chapters we will hear about a number of these effects from many different kinds of companies and

organizations. The information we do have about processes is seldom complete since part of that evolution happens because people bring outside knowledge into the organization with them, and use it in creative, but often unexpected and unknown ways, to get their work done. Unless we understand the role of this tacit knowledge in our processes, we never really know what those processes are. So in short, it is quite reasonable to find that in many organizations, we really do not know how we do things, even though we may think we do before we take a careful look.

Second, much of the information we use for making decisions and controlling the day-to-day activities of a productive enterprise, the kind we consider to be in the KK cell of Figure 4.1, is seldom really examined or questioned as to whether that designation is accurate. Who has not left a meeting wondering what the whole thing was about? Who has not had the experience of being told to manage a financial decision on the basis of a policy that, with little analysis, can be shown to be less cost-effective than an easy alternative? One does not have to look very hard to find examples of companies that spent millions of dollars on an Enterprise Resource Planning (ERP) system, entirely on faith that it will work, only to find that in some respects it never really did. In reality, we do a lot of things in organizations on the basis of “because.” As long as our cash flow enables us to absorb the costs of “because,” we can get away with it, but that may not work over the long term, and we will hear some stories in this book about that, too.

Knowledge is information, and information is always partly incomplete and in some ways imperfect, so it can only be rendered useful through processing. One of the major functions of organizations is to process information and knowledge, so that coordinated progress toward goals is enabled despite the limits to the information we face. Organizations are also systems, and are therefore open to all manner of inside and outside shocks and internal changes, all of which require them to be adaptable. Many organizations do this rather well over the long term, while many others have short, if interesting, lives. How well an organization performs depends on all the outcomes of its actions, whether proactive “offensive” behaviors or reactive “defensive” behaviors. Both of these may result in costs or benefits, and it is the net payoff of these that determines how we do in the long term. I generally dislike sports

analogies for their oversimplification of complex issues, but the idea of batting averages in baseball applies here. A batter can strike out whether he swings or not, and for the batter who produces a respectable average of hits in his at-bats, along with the occasional home run, there is a realistic chance of making the World Series.

This chapter has focused on some basic propositions about knowledge, information, and organizations. In some ways I have stressed the limits to our knowledge and our ability to cope with them. This does not mean that useful management of knowledge is beyond our reach, however—quite the opposite. I have focused on limitations and boundaries because it is important to know what we know as well as what we do not know. We may have to give up on the idea of a full and comprehensive database or boundless wellspring of innovation based on an open organization structure, but there are tools and methods that can be very helpful in increasing the extent to which we know what we know. We will always have to deal with the reality of UU, and the conundrum that we cannot know what these are; there will always be questions about the value of information and the value of obtaining more of it, without fully knowing what the payoff of additional information might be. Nevertheless, there is also the potential to capture more of what we have discovered and learned, and to use what is frequently an “unknown known” to much greater advantage. One of the key functions of workflow mapping is to help the organization know what it knows.

Tacit knowledge is always a key to how organizations cope with their limits to knowledge. Consider three types of organizations—a glass-products company, a software developer, and a hospital. At the beginning of this chapter I pointed out that organizations have to process information to achieve both technical and coordinative functions, and Figure 4.5 shows how formal and tacit knowledge both contribute to these objectives. First, formal knowledge is the basis of technical performance. The properties of materials that make various glasses, the programming rules and syntax for computer code, and sources of infection, are all among the many elements of the formal knowledge base that technical performance depends on; at the same time, coordination depends on related formal knowledge of how glass behaves in its molten state, so that a successful production line can be designed; how (and to whom) to

		Knowledge	
		Formal	Tacit
Performance function	Technical	Properties of materials; programming syntax; sources of infection	“Opening” blown glass; transportable “chunks” of computer code; modes of infection
	Coordination	Design of glass production line; assignment of code modules; sterile-zone rules and procedures	Everything else – all other modes of adaptation and adjustment not predicted in other cells

Figure 4.5 *Formal knowledge, tacit knowledge, and organizational functioning*

assign code modules for new programs; and on the steps and procedures medical staff follow to keep sterile zones sterile, since hospital-derived infections are a major medical problem.

Tacit knowledge, shown in the right column of Figure 4.5, is equally important to organizational capabilities. Much of the technical success of organizations is entirely dependent on what people learn in what might be thought of as “apprenticeships.” One learns to “open” blown glass through trial and error; knowing how to apply transportable “chunks” of computer code is often a matter of deeply knowing how a piece of code works, by the programmer; and how and where infections get started is often as important as the bug that causes it, and sometimes more so.

What is most important to realize about tacit knowledge, however, is the bottom-right coordination cell—this is literally “everything else” we know how to do. It is where individual and group learning and knowledge give the organization response capabilities it never could have anticipated needing, let alone designed. In a universe where we can never have complete and perfect information, an absolute necessity is the ability to compensate and adjust when the UU’s and other unknowns in Figure 4.1 reveal themselves. In many situations this cell defines how organizations survive.

The next chapter provides an expanded framework for understanding how organizations function, and that understanding is the basis for the simple, robust, and widely applicable method of graphically describing workflow processes that we have seen as WFMA in Chapters 1 and 2. The combination of conceptual tools in this chapter and Chapter 5, and applied tools in Chapters 1 and 2, enable managers and analysts to comprehensively describe all that is done with material and information in a process. The ability to accurately capture both formal and tacit knowledge in our workflows has a big payoff. While it will never solve the fundamental limitations to full and complete information, workflow mapping and analysis will certainly go a long way toward letting us know what we know, and experience clearly shows that improvements in the quality of information from that increase our ability to improve performance.

CHAPTER 5

Systems, Processes, Organizations, and Workflows

Introduction

There are two primary objectives of this chapter. One is to establish as much clarity as possible in our thinking and use of terms in workflow mapping and analysis—we all want to be on the proverbial “same page.” This takes a bit of work, since all of the words in the title of this chapter are used very widely but often with many different interpretations. The second objective is to understand why and how organizations and their processes vary as much as they do, and why we need a robust method like the one in this book to capture all that variety. The good news is that when we understand why organizations and their processes can be very complex, the simple workflow mapping methods we have seen in Chapters 1 and 2 enable us to capture those workflows and the formal and tacit knowledge that is part of them.

To achieve these objectives, we will first start by defining a system, the most fundamental type of structure. All systems contain processes, so we will examine the variety of them and how processes operate in systems. Workflows are processes designed by people to get things done, so we will examine these last. Throughout this discussion we will consider the impact on the organization workflows present, so that by the end of this chapter we will have a very thorough understanding of the words and concepts we have been using and will also be prepared to map and analyze workflows as we find them in the real world.

Systems

The concept of a system is part of a long-established body of thinking known as General Systems Theory (GST), often traced back to the scientist-philosopher Gottfried Leibniz (1646–1716), and to more modern development by the American Association for the Advancement of Science in 1954, principally through the work of Ludwig von Bertalanffy.¹ The general system model, that is, the concept of complex wholes made of many interacting parts, can apply to an enormous variety of entities. There are four properties associated with “systems” as they are defined in GST: (1) the whole is greater than the sum of its parts; (2) the whole determines the nature of the parts; (3) the parts cannot be understood in isolation from the whole; and (4) the parts of the system are dynamically interrelated—they are interdependent and interact with each other over time. As we saw in Chapter 4, information is the “glue” that holds all of the components of systems together.

In a classic paper, Boulding proposed a hierarchy of complexity of systems, which in brief is: (1) static structures or “frameworks;” (2) simple dynamic structures, or “clockworks;” (3) cybernetic systems—self-regulating structures like thermostats; (4) open systems, or self-maintaining structures like cells, at that time considered to be the level where life began; (5) genetic-societal level, like plants; (6) the animal level, featuring mobility and use of information; (7) the human level; (8) the social organization level; and (9) the “transcendental” level, where questions may not have answers.² My only minor disagreement with Boulding’s scheme is that even frameworks have long-term dynamics, and are not really “static.” While indiscernible to humans, radioactive isotopes in rocks decay and tectonic plates move, meaning that rocks melt and are formed anew, water transforms everything over time, and the universe is in constant motion. I would argue that all systems known to mankind are dynamic, and change over time. In addition, I regard all systems as existing within a larger “environment,” a point Boulding did not really address, but seems to assume given his discussion of the solar system, the category of open systems, and the like. (His principal concern was with the utility of GST as a scientific model and so this was not a major issue.)

Because of the hierarchy or surrounding environment, whenever we select a “system” for observation or analysis, we must specify our frame of reference. In biology we can apply this to a single cell, or a smaller part of a cell (e.g., mitochondria); it might apply to much larger aggregations of cells (an organ) that in turn are part of a larger organism, and so on. The model also applies to aggregations of people—we can think of a company or a rugby team as a system just as easily as a living organism. We can also think of parts of larger systems as *subsystems*, or a larger system than the current focus as a *supersystem*. This is a robust model that can be scaled up or down as needed. When I refer to a “system,” that always means that we must define what elements are within the system of interest and which are part of the environment. That is sometimes obvious, but at other times is not a trivial matter. So our system model is a simple representation of reality, but it permits the description of very complex wholes.

Because all systems are dynamic and interact with a larger environment, any “system” thus contains “processes,” even if they are nothing more than predetermined natural phenomena. At the most fundamental level, there is at least a minimal *action* or activity between the system and its environment, or within or between elements within the system—even a decaying radioactive element in a rock sends a particle or energy into the rock’s environment, and internally the isotope of that element is changed. In my view, increasing system complexity in the Boulding hierarchy necessarily means increasing process complexity—there is a parallel structure between systems and their processes, in terms of both the number and structure of processes. In all cases, *information* is the link between parts of the processes. (For those who are unconvinced that physical elements like atoms might have informational linkages, I will leave it to them to research *quantum entanglement* on their own.)

A “process” is therefore the product of actions and information. We will expand on this idea in the next section.

Processes

“Process” is one of those words that is very general and used in a lot of different contexts in the world of work—there are also numerous definitions in most dictionaries. Over time, everyone uses it and understands

it, at least within the context where they apply it; whether that meaning is the same as another user's, however, is often another matter.

If we expand on the idea that processes are the product of actions and information, we can assemble some very basic process building blocks, as shown in Figure 5.1. First, when we speak of “actions,” we are really concerned with two basic dimensions of a system—*where* it happens (the *locus*) and *how* it happens (the *mode*). Action modes can either be *active*, where a system or part of it initiates the action, or *passive*, where inherent properties of the system cause the action. A bird's feeding behavior is obviously active; radioactive decay is passive. Second, the locus of the action may be *internal* to the system or a *boundary-spanning* action, where the environment is involved along with the internal system elements. Interactions between the locus and mode of actions enable a wide variety of specific actions; these can also be aggregated into much more complex actions in more complex systems.

Information, which, as we have seen in Chapter 4, is an elemental part of everything we do, can be conceptualized in a way similar to that of actions. One mode (or function) of information can be thought of as *inherent* to an action, meaning that it is part of it and perhaps the whole substance of it. A change in the state of magnetism in a piece of

		Locus of action	
		Internal	Boundary-Spanning
Mode of action	Active	Move; modify	Acquire; dispose
	Passive	Observe; monitor; decay	Absorb, accept; deflect;

Figure 5.1 Examples of actions resulting from interactions between locus and mode

		Locus of information	
		Stock	Flow
Mode of information	Inherent	“Library,” data base	Path, procedure
	Regulatory	Rules; policy	Conditions, comparisons

Figure 5.2 Examples of information resulting from interactions between locus and mode

iron is an example—the magnetic property of the iron is changed, but not the iron itself. In addition to this inherent function, there is also a *regulatory* information mode. This is typically found in more complex actions, such as a human being who perspires automatically to cool down if the temperature of the body gets too high, but even simple “framework” systems rely on this to maintain stability, as we see with the balance of motion and gravity in the solar system.

Information can also be thought of, using the terminology of economics, as located in a *stock* or a *flow*. A *stock*, as the name implies, is an accumulation of information, and a *flow* is the movement of information from one state or location to another. These interactions enable the types of information shown in the cells of Figure 5.2.

Organizations as Systems

Let us now combine these basic ideas and apply them to working organizations. We will immediately recognize many parallels and transferable, applicable ideas. First, since systems interact with their environments, that includes organizations—they take in inputs, transform them into outputs, and return outputs to the environment, as we see in Figure 5.3.

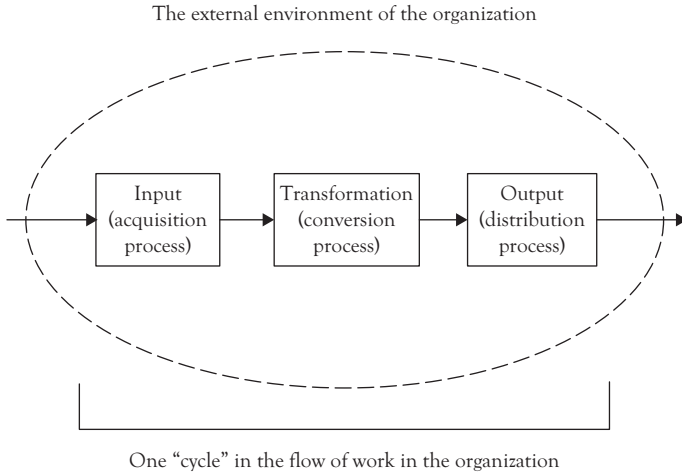


Figure 5.3 *The basic system model*

Source: Adapted by permission of the publishers, in *The Information Processing Theory of Organization* by John L. Kmetz (Farnham: Gower, 1998, p. 38).

In this context, a “process” in the organization may be considered as the transformation of the input into an output, or it could be considered to be the combination of input acquisition, transformation, and output disposal.

The organizational processes in Figure 5.3 are of two general types. The first type is the exchange of inputs and outputs with the environment. These are described as “boundary-spanning” activities because that is what they do; these processes cross the permeable system boundary to interact with the environment. The second is the internal transformation of inputs. For much of what we do in organizations, transformation is the most important type of process, since it is how we add value to the inputs we acquire to create our outputs. Transformation is literally central to our existence, and is therefore one of our primary concerns when we examine organizational processes.

Because they enable the transformation process to go on, boundary-spanning processes are just as essential to the system as the transformation processes themselves. Both of these processes are *mutually interdependent* from the perspective of the organization, since all of them are required for the system to function as a whole (and for this reason I resist labels like “primary” or “secondary” for any of them). As with any system process, these are bound together with information.

Examples of both internal transformation processes and boundary-spanning interactions are everywhere: a human being takes in air, food, water, and so on, from the environment, processes these internally, and returns a variety of outputs, including work (typically with added value) and “waste” products (which are usually only “waste” to that system). Companies take in raw materials of many kinds (capital, raw or unfinished materials, employee talents and skills) and combine these into new products, services, office trash, and frequent-flier miles. The products are sent to the “market,” an open environment where potential customers can take or leave what companies have to offer. If customers accept an offering, they give back money, the common medium of exchange, and the process repeats; if not, the firm loses its capital and fails (and as Williamson argues, the essential alternatives to how these transactions are regulated are either through markets or hierarchies—name your poison³). Transformation processes and boundary-spanning interactions with the environment in both of these examples can be seen to take many forms, and in the case of companies and organizations, the relationships become very complex.

Nearly every organization tries to influence what the “environment” does, through advertising, public relations, gaining favorable regulation, and the like—that is why organizational processes include altering relationships with the environment. If we get cold because we are losing too much body heat to the environment, we try to regulate the exchange by putting on warmer clothes or through other steps. Banks hold some money from depositors and limit access to it without changing it, while at the same time using it as a basis for loans and creation of other assets.

Like all others, organizational processes are *dynamic*—they occur over time, and the transactions in them may be very closely linked to predecessors or there may be material delays between them. Time, of course, is a universal resource and a universal cost, and the temporal relationships between actions and their components have important consequences for overall system functioning. In most complex systems (like companies) there will be multiple processes under way at the same time, so that several subprocesses may go through a significant amount of activity before they (or their outputs) converge in the larger system. When convergence between these processes occurs, we may find that some of the pieces do not fit, and will need rework. This is inevitable, and some

degree of loss or dysfunction is characteristic of every system. Even the human body cannot process 100 percent of its energy into useful work—there is always significant waste heat. But if delays or long linkages between subprocesses and their outcomes occur, it may be difficult to detect that losses are occurring. As we will see, not everything in a complex process is always easily measured or predictable.

Organizational processes are *dependent on information*, and in organizations this is both clearly visible and of major importance to the ability of the organization to meet its goals. One type of information dependency is that which is *inherent to the process*, meaning that some information is used specifically for that process. We collect raw data on customer satisfaction and transform it into new product designs, advertising campaigns, and the like. We monitor patients in recovery rooms very closely to guard against adverse and unanticipated reactions to medical procedures, and move them through a progressive series of steps until we are reasonably sure of full recovery.

As with other systems, there is also *regulatory* information, and in organizational processes this is extremely important. Most regulatory information takes two generic forms: (1) “feedback,” where some information or output is consumed to maintain control and evaluate performance against expectations—some of the work done by the system is consumed by the system for self-regulation, and (2) “feedforward,” which sets targets and establishes criteria used to guide future actions and decisions.

Most of us are very familiar with the idea of feedback, and we use that word all the time in many aspects of life—feedback from customers, significant others, suppliers, the government, and virtually everywhere else in the world. In companies and organizations, feedback includes financial control data, performance appraisals, assessment of responses to and from customers, quality measurements, and the like. We realize how important it is and collect and think about feedback all the time. Feedback about processes is often obtained from internal sources, but many types are also from external sources.

“Feedforward” is a less common term for a familiar idea, but just as important as feedback. Feedforward consists of information like project plans, performance projections, consumer confidence surveys, budgets, benchmarks, and similar types of information. Feedforward establishes

the plans and criteria that guide our actions, and gives us a roadmap to follow, if we use it well. Many feedforward elements are from external sources as well as internal ones.

These are typically used in combination with each other over time, so these are also dynamically related. We formulate a careful project plan and budget before bidding on a job (feedforward), and then if we win we both use the criteria established in the initial plan to stay within performance parameters (feedback) as well as review the process we went through to learn how to win business in the future (feedback on the feedforward). This latter case can be thought of as a “nested” relationship between the two types of regulatory information. We are surrounded by this kind of information: we hold brainstorming sessions to come up with novel product designs; we measure the dimensions of a manufactured part before it goes into the next higher assembly, and if the part is not within specification, it is not used; we edit and “polish” a report before sending it up the line to the boss, and so on.

Figure 5.4 is a more complete system model for an organization, since the coordination and control of work is impossible without information exchange. It is hard to conceive of examples of any process where regulatory information is not involved (even “automatic” ones like shivering

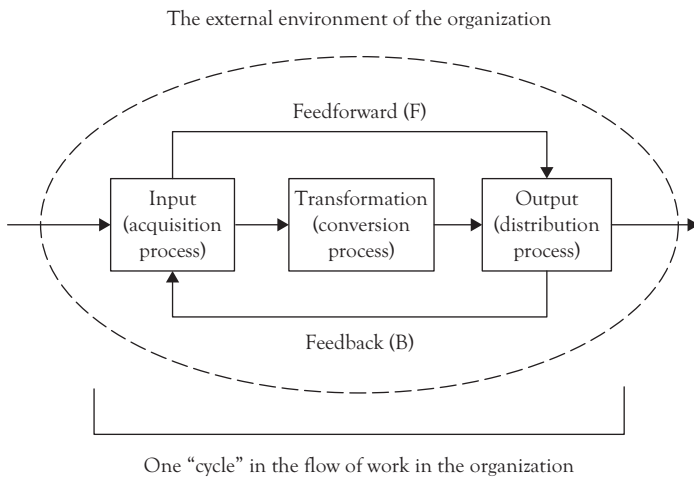


Figure 5.4 *The basic system model with regulatory information flows*

Source: Adapted by permission of the publishers, in *The Information Processing Theory of Organization* by John L. Kmetz (Farnham: Gower, 1998) p. 39.

if we are cold—this involves an internal flow of information about the body’s heat level and its changing status over time, and if we are losing too much heat we shiver to increase our warmth, even though we never think about doing it). The energy required for such internal regulation is thus another reason why no system can ever function at 100 percent efficiency. As simple as they are, these regulatory activities constitute one of the great ongoing challenges of organization, and they are never “done” because both the system and its environment undergo continuous change that creates new regulatory feedbacks and feedforwards.

Knight and McDaniel pointed out that while most of what we seek in inputs can be specified, there are often unknown other properties that accompany the part we want; in addition, for those qualities we want, there may be some parts that are not controllable.⁴ When we hire someone, we get the whole bundle of attributes that make up the person, intended, anticipated, or not. Even when we buy a car we buy a bundle of attributes, some of which we may not really know of when we make the purchase decision, and which may not be controllable later. It is easy to give examples of automotive recalls as examples, but even when there is no recall the same reality applies. A French friend bought a powerful Mercedes-Benz AMG, and found that in a snowy winter, it was best to garage the car and drive his old Subaru—the wide tires and engine power in the Mercedes made it very difficult to control on slippery road surfaces, where the old Subaru, with less power and four-wheel drive, was ideal!

Processes produce both known and unknown consequences, and not all of the known consequences are necessarily intended. In the Soviet era it was common to hear stories about the unintended effects of production-quota systems—one classic story is about the Czech fastener factory that was hundreds of tons short of production at the end of one month, so they solved the problem on the last day of the month by producing one 600-ton bolt! There are, however, so many stories about the unintended consequences of using performance measurement systems like this in the industrial world, that the title of a nearly 40-year-old article has become a cliché in management: “On the folly of rewarding A, while hoping for B.”⁵

Outputs also frequently have unknown properties that accompany the known ones, and among the known ones, some are unanticipated.

Promotions in the automobile industry may often serve their main purpose, to stimulate sales, but they may do it by “stealing” sales from subsequent periods to benefit the time when the promotion was held. As a result, whether overall sales volume or market share benefit from such promotions is often an open question in the industry.

Despite the vagaries of knowledge and imperfections of information, most organizations function very well as self-regulating systems, and once set on a path to achieve specific objectives they are remarkably adaptable. They are so adaptable that many systems theorists consider complex systems like organizations to be “goal seeking,” almost as a living organism. That property is very much dependent on tacit knowledge and the willingness of the people in the organization to use it to help achieve objectives. That tacit knowledge is embedded in the workflows that make the organization the productive entity it usually is.

Workflows

All systems have processes, but only organizations have workflows. Organizations are human inventions—they are social entities designed to achieve a goal. A *workflow* consists of all necessary activities and information to accomplish a specific purpose in an organization—it is a *designed process*, which may consist of a number of subprocesses linked together. Workflows can vary enormously in terms of the number of steps involved, the length of time required, the resources required, and so on. In organizations, many workflows pass through more than one department or organizational unit, and are not contained within only one of them. Most of the people who discuss “organizational processes” are really concerned with organizational workflows; for most purposes, this does not matter, so long as we are clear about our unit of analysis. While there are many parallels to natural processes, there are many ways in which workflows are quite different from natural processes.

Workflows change over time, for many reasons, but most of these reasons are not the sorts of things that change natural processes. Workflows change because technology changes, markets (the environment) change, or because they are forced to change by competitive actions that leave no alternative. Because they are inventions, the change process is

also by design, and this is a major reason why some variations in them may be unknown, unanticipated, and sometimes significantly dysfunctional. This means that in some cases workflow may vary in terms of clarity or definition—*why* are we doing this? This may be one of the most important attributes to understand in the mapping of a process.

Understanding workflows, and the ability to map them, requires understanding of the “raw materials” the company brings in from the environment and the “technology” used to turn them into outputs (and everything, no matter how everyday it may seem, involves a technology of sorts). Input and conversion processes usually consist of many steps and activities, many of which are specific, and often unique, to the organization. While knowing about these is necessary to map a workflow, the technical part of the process alone is not everything—we also need to know about the feedback and feedforward used to control the processes, and these can become very intricate.

As our discussion of processes suggests, all organizations require two closely related workflow elements to accomplish their objectives. First, there are flows of *materials* in manufacturing, or “quasi-materials” such as documents and information in many manufacturing and service companies. Since the majority of the U.S. economy is no longer manufacturing or agriculture, intangibles are the largest element of national production, and most of those intangibles have large information content. For that reason, we might use the word “informaterial” to designate information as production; most of the time, I will use “material” to designate both physical material and informaterial.

Second, there are related flows of *information*, both the inherent information about the activities in the flow of work and regulatory information flows. A marketing research firm provides information about markets to its clients as its product; as we use the terms here, the product is the informaterial flow, and its creation is guided by information flows for coordination and control. Both types of flows must be understood to fully describe a workflow and as we will see, these are usually closely intertwined. To attempt to map the flow of physical material or informaterial through a process without understanding the roles of feedback and feedforward is basically a waste of time.

We focus most of our energies on what goes right in a material workflow, but not everything does—those things that do not constitute *exceptions* to the normal flow of material, meaning that the output does not satisfy all criteria set for it. The flow of materials in all working organizations consists of a combination of normal outputs and exceptions, and as we have seen earlier, the latter are often a major challenge to resolve and to map. Collectively, we refer to these dedicated flows of information (and sometimes materials) as *exception handling*—what we do to correct or compensate for exceptions in our day-to-day work. What do we do if there is a fire drill, or a power failure? Even for highly structured tasks, exceptions will arise. Unloading and reloading a container ship in a modern port is a complex process. Containers must be unloaded in a manner that minimizes delays in moving from the ship to local transportation or storage, and reloaded so that the container order is correct for the receiving port. However, this can never be done perfectly for a number of reasons, a major one being that if the ship's load becomes too unbalanced at any time in unloading or reloading, she might simply roll over and sink! So despite the best planning and the use of intensive automation, exceptions and problems will occur, schedules will slip, and information must be processed to deal with these events. This is why most exception-handling flows are highly information-intensive.

Other exceptions are “bolts from the blue.” An unloading crane breaks down, completely upsetting the schedule for a ship; a package gets lost or delayed; a check isn't cut because of a clerical error by a trainee; a critical meeting is missed because the presenter (a) got caught in traffic, (b) got sick, (c) went into labor, (d) had to deal with a family emergency, and so on. While exceptions such as these are caused by events beyond our control, they must be dealt with nevertheless. It is not uncommon to find that significant amounts of managers' time goes into this one activity, and that exceptions may take more of any given day than the so-called “daily routine.”

The dynamic properties of workflows inevitably involve associated information flows, and it is impossible to conceive of a dynamic system that did not have information at its core. In the physical world atoms and molecules are bits of matter held together by forces like magnetism and

gravity; in the working world, materials are assembled within processes held together by information. This may seem like making too much of a straightforward observation, but as we will see in the next chapter, this point is very significant to the understanding and validation of workflow maps.

Material and Information Flow Relationships

Obviously, material and information flows are connected, but the relationships can vary in two important ways—the tightness of their *coupling*, and the extent to which they are *synchronized*. First, with respect to coupling, they may be fully *convergent* through the workflow or *divergent* in one or more ways and places. In many cases of process improvement and work simplification, we want to make the flows of materials and information converge as much as possible. Separation of work from decision making about the work, for example, is often a principal contributor to delay, error, and the perception of red tape (often the reality), as well as a demotivator for the people who do the job.

This is a point where the “yarn ball theory of organization” should be introduced. In many consulting and training situations, I have used a ball of yarn to illustrate a few simple truths about the way that organizations grow and how their structures change. The exercise is simple—I tie a loop in the loose end of a ball of yarn, put that over my wrist, and toss the ball to someone anywhere in the room. The instruction to everyone from that point on is the same—put a loop around one’s wrist, and toss the ball to someone else. This process continues with random tosses until everyone has at least one yarn loop around their wrist, or happens to be the last to get the yarn ball.

This exercise illustrates two useful things. First, if “a system is a thing made up of other things connected to each other and all other things,” there is no better way to see this than to be webbed into the “Yarn Co.” Second, no one anywhere can move without affecting at least several other people. The yarn linkages between people are the paths that information and materials have to follow for any work to get done. Rationalizing these relationships, meaning that we want to reexamine them and improve how they work after they have been put in place, is often the expressed

or implied reason for any kind of workflow mapping project. Following the paths that materials and information follow in reality is at the heart of workflow mapping, and if we were to “rationalize” the Yarn Co., we would have to make many of the convergent relationships diverge instead, and vice versa.

The second dimension of workflow relationships is *synchronicity*. Material flows and information flows may be *synchronous*, meaning that they are simultaneous or linked in a very tight serial sequence, with the absolute minimum of delay in accessing or using information to support material flow. In synchronous flows, materials and information follow the same paths and are processed in parallel and usually in the same location. However, flows may also be *asynchronous*, in that they are nonparallel, follow different paths, cut through unit boundaries, follow different sequences, or a combination of these. Timing is less likely to be tightly linked between the flows, and tight linkages may be undesirable.

Typically, synchronous flows are preferred in workflow designs, as they keep materials being processed and information about the process in close proximity, so that decisions happen quickly, exceptions are handled on the spot, and the most relevant expertise on the workflow issue is quickly brought to bear where needed. “If you want to learn how to dig a better ditch, ask a ditch-digger” applies here. Asynchronous material and information flows are typically considered undesirable, in that the lack of synchronization both creates delays and necessitates additional information processing.

As summarized in Figure 5.5, the consequences of these two aspects of workflow design may be either *functional* or *dysfunctional* for the achievement of organizational goals. Functional relationships are those in which organizational objectives are served and value is added by the material and information flows as designed. Dysfunctional flows are those that cause unnecessary delays, costs, increased probability of error, or a combination of these, and for which little or no value is gained in return. Because of these differences, the general bias toward convergent flow coupling and synchronization in process improvement is usually the best choice, but may not always be.

Notice that functional or dysfunctional outcomes can occur in any of the four cells of Figure 5.5, and that most outcomes, in fact, are the

		Flow synchronization	
		Synchronous	Asynchronous
Flow coupling	Convergent	High efficiency, high risk	Time differentiation— “scientific method”
	Divergent	Flow differentiation— “double blind”	Low efficiency, low risk

Figure 5.5 Path and synchronization effects on workflow outcomes

proverbial double-edged sword. Convergent and synchronous outcomes in the upper left cell of the figure bring both high efficiency and high risk; the opposite is true on the lower right cell diagonally opposite. By differentiating either flow paths or flow synchronization on the other diagonal (lower left to upper right), one can avoid certain kinds of problems by being very cautious with information, but at the cost of slowing processes down greatly. “Trade-off” is the operative word in deciding how to proceed. As we have seen in Chapter 3, we need to think about all effects that come from a change.

Examples of divergent or asynchronous flow relationships that are functional are not at all difficult to find—think of cases such as those where people who handle money are not allowed to reconcile accounts; where people with vested or financial interests in companies are not allowed to make public policy decisions about these firms; or where those who claim to have made scientific breakthroughs are subject to independent replication and verification of their findings. Even though there is designed separation of information and activities in these workflows, they are ultimately functional and add value precisely because the flows of material and information are divergent.

Scientifically controversial work such as “cold fusion⁶” or scandals like the stem cell research fraud in South Korea⁷ were revealed as

functional benefits of asynchronous information flows. In general, much of the scientific method may be thought of as an asynchronous but convergent workflow, from which bad questions or ideas are rejected in favor of better questions and ideas. Ideally, we think of science as leading to “truth,” but final answers in any field are hard to come by. Much benefit, nevertheless, is gained from asking strong questions rather than weak ones, and a functional combination of convergence and synchronicity enables strong questions, both in science and in business.

Divergence can also be functional. As we have seen in Chapters 1 and 2, no drug is ever cleared for general public prescription today unless it has gone through a double-blind review process during its testing phase. Much experience, some of it fatal, has shown the functional effects of such double-blinding, where neither the doctor administering doses nor the patient receiving them knows whether a particular dose was an experimental medicine or a placebo. Divergent, synchronous double-blind procedures protect the objectivity of the observers and the patients, ensuring the most accurate possible data on drug effects.

A business example may help clarify this point. In the early 1990s I traveled regularly to Bulgaria to work on a University of Delaware USAID-sponsored program to assist the transition of Bulgaria from a centrally planned economy toward a market-based one. Like most of Central and Eastern Europe, Bulgaria had a cash economy, and credit cards were very rare and seldom used. However, in 1993, the Bulgarian central bank contracted with MasterCard to begin offering card services, and largely out of curiosity I went into the main bank building in Sofia on a cold winter day to get a cash advance of USD 200.

I first had to inquire where I could begin such a transaction. After being directed toward the correct window, I submitted my credit card and passport to a teller who completed a cash advance form by hand, and had me sign it in the presence of her supervisor, who verified my identity and that my passport signature matched the request form signature; she also initialed the request. The teller then telephoned the request into the central MasterCard office, which took a few minutes. Upon approval from the central office, she filled out a currency request form, which requested the correct amount in Bulgarian *leva* and designated the correct number of leva bills of different designations to equal USD 200. I signed

the credit card receipt (of course inspected against my passport again), and got my copy.

The cash request form was in four parts, and she initialed and took one. We then went to another office on the same floor, where two different young women counted out the correct bill designations and numbers, initialed the form, and kept one copy. I observed that one teller took the leva from its drawer as directed by the other, and they completed an internal form to record this transaction and note the source of the request. The second teller kept a copy of the initialed request, and took my passport, the last two copies of the request, and the leva and me to the foreign currency exchange teller, handed her the entire package of paper and left with a pretty smile and a “good day.” The foreign currency teller carefully counted out the two USD 100 bills I wanted, twice counted the leva, which she placed in another safe drawer, initialed a copy of the request, and prepared to hand me the U.S. currency, my passport, and my (fourth) copy of the cash request. The transaction was finally completed when I paid a 30 leva processing charge that had not been included in the transaction; I was given a separate receipt for that. As I gathered up my papers, I observed the foreign currency teller take the paperwork she had just completed to a supervisor who examined and initialed the foreign currency disbursement form. The entire process took about 45 minutes and six people working in three separate offices.

This is an example of a divergent workflow in the extreme, and while it seems overboard to us, in that particular cash economy at that time such controls made sense. The novelty of the credit card as a means of obtaining cash when a user did not have an account at the bank probably was perceived as reason for extra caution, and it was clear that nobody in the bank was really comfortable with this kind of transaction. Nevertheless, despite the extreme divergence of this process and the redundant and asynchronous checking of identities and documents, it was still one long (very long!), continuous process. (For those planning a visit to Bulgaria, I should point out that the country is now part of the European Union, and automatic teller machines are everywhere.)

In short, do not make the mistake of assuming that because material and information flows converge or are tightly synchronized with each other they are necessarily optimal, nor that the absence of these properties

is necessarily dysfunctional. The bias toward convergent and synchronous flows of materials and information is quite understandable in many business situations. But examples of science gone awry, frauds, internal thefts of money and information, and other similar examples in the news every day, should make it clear that accurately mapping a workflow, and then using that knowledge as a way to improve it, requires deeper understanding of how we use information to control and coordinate the flow of work to add or create new value.

Summary

Depending on how the reader approached them, this chapter completes the amplification of, or setting of the stage for, the applied tools in Chapters 1 and 2 to be used. We have seen that organizations are systems that conduct transactions with their environments and apply processes within to convert inputs into outputs. What happens on the “inside” is thus always interdependent with what happens on the “outside.” Second, we have seen how processes are combinations of actions and information, and by examining where and how these are formed we find that there are enormous numbers of processes with subtle variations between them. Third, we have seen how organizations, the kind of system of primary interest to us, create and use processes in the form of goal-related workflows. Depending on the nature of the coupling and synchronization designed into these workflows there are many possible consequences for workflow outcomes, and in a world of imperfect information, the decisions we make to structure workflows and organizations bring significant trade-offs between functional and dysfunctional outcomes with them. Simple one-size-fits-all “right” answers to workflow and organization design are not forthcoming in this world, and we must keep this in mind when we begin to dissect and reconstruct workflows for various purposes. This also places an absolute requirement on us to know how existing workflows function before we tinker with them. Those who believe in a fixed “catalog” of subprocesses from which we can assemble all higher-level processes will forever be frustrated, in my view; rather than a fixed catalog, I think a better analogy is music, and with only a few notes, scales, and keys, we have yet to exhaust the variety we can create.

CHAPTER 6

WFMA and Knowledge Management

Introduction

In one important way, this chapter tells us something we already know, which is how to capture the knowledge embedded in a workflow and understand how it is used, and we do not need to master any additional tools to do this. Workflow mapping and analysis (WFMA) is a form of knowledge management (KM) in its own right, in that by mapping workflows we capture the formal and tacit knowledge an organization incorporates into its processes. If the workflow mapping methods described in Chapters 1–3 are followed, both the formal or explicit knowledge and the tacit knowledge in a workflow will be revealed. Much of the basis for process improvement discussed in Chapter 2 rests on having captured both types of knowledge and their application to workflow processes.

This chapter takes a hands-on approach to KM, consistent with earlier chapters. It is not concerned with theoretical or esoteric matters of KM, nor does it become concerned with the *de facto* questions of database definition and manipulation that are so often associated with “knowledge management.” We have shown that the overwhelming majority of relevant knowledge in a workflow can be identified and mapped. As such, the maps themselves and the results of our analyses of them can be important content in a KM system. Unlike many other approaches to KM or process mapping, however, our approach does not implicitly assume that all knowledge is positive or helpful. Variations in a workflow can sometimes be the result of application (or misapplication) of tacit knowledge on the level of a worker or group, and in other cases may be the result of either ignorance of formal knowledge or “malicious compliance” in

the application of it. Knowledge (or information) is power, whether for good or evil. Implicitly, much performance improvement involves the discovery and modification of actions based on incorrect or misinterpreted “knowledge” about how things are done.

An Overview of Knowledge Management

KM focuses on the capture, creation, storage, and retrieval of knowledge and information. The roots of KM as we now think of it are very diverse. Industrial engineering, economics, decision theory, systems theory, communication, psychology, and sociology all contribute to how we conceptualize knowledge and the effective use of it. Indeed, my doctoral dissertation focused on the influence of “knowledge technology” on organization structure.¹ More recent techniques such as benchmarking, many quality management methods, the use of business metrics, business process reengineering, and advances in database development, data mining, and dynamic modeling are all part of the lineage of KM.

In this chapter, I am interested in tools to manage knowledge, not in addressing the philosophical issues of KM. The subtleties and differences between “data,” “information,” and “knowledge” have been the subject of debate since the dawn of classical Greek civilization, wherein the argument is that we first receive raw data which must be transformed into information and then into knowledge.² This places knowledge in a fundamentally different category than either data or information. Like others, I reject that hierarchy as spurious.³ Suppose this entire book were written in binary. Unless one knew what binary code was, it would be absolutely indecipherable to one only familiar with the standard English alphabet. Binary would probably be written as ones and zeros, but there is no reason that other symbols could not be used—a hyphen for one and a space for a zero, or the reverse, could also be used, as well as any two randomly chosen characters from any alphabet, or the Wingding characters accompanying most word processing software. Without prior knowledge of binary code and alternative ways of portraying it, the “raw data” are inaccessible in the first place. This academic debate still goes on.⁴ William Starbuck relates the story of an academic conference on KM

he convened in the early 1990s, and the only thing the participants could agree on was that they did not agree on the definition of “knowledge.”⁵ We are not going to solve these problems here.

Nevertheless, KM has become a growth industry of sorts during the past two decades, although one that perhaps has reached and passed its zenith;⁶ over 10 years ago, Nickols (2000) reported receiving two million hits on “knowledge management” and seven million on “KM” when he did a Web search.⁷ Several factors played a major role in that growth, including the enormous investment in information technology that happened with the introduction of microcomputers, the Internet, and the emergence of postindustrial economies in much of the global “Triad” of North America, Europe, and Japan. These forces shifted attention to information on a huge scale, but also propelled knowledge and information issues to the forefront of strategy and decision making for many organizations. The era of the “knowledge worker” predicted years earlier by Peter Drucker had arrived.⁸ Knowledge was how value was added and created, and “human assets” became the primary asset for many industries.⁹

As a contributor to business success, KM has a somewhat uneven track record. There are claims for its success and illustrative case studies of apparent success,¹⁰ but there are examples of initial successes that have backfired¹¹ and the Phios experience discussed in Chapter 1 is a case where an IT-based KM system appears to have completely bombed. Reviews on the payoff of KM are often mixed and inconclusive;¹² even supporters have to acknowledge that many of the early claims of success never examined the costs of KM.¹³

As a discipline, KM has split into two rather sharply differentiated branches. One might be thought of as an exercise in the creation of the ultimate database and the ultimately intelligent search engine, typically with heavy dependence on the Internet; this is the primary realm of the KM system, or KMS. Davenport organized these IT applications into a relatively comprehensive matrix of eight different types, based on the complexity of work and the level of interdependence between individuals in the organization.¹⁴ He also found some very interesting and insightful methods individuals used to manage personal knowledge, some of which, like lists, were quite independent of information technology.¹⁵

I am also not attempting to summarize or address the issues relevant to advancing the development of KMS. There are some aspects of KMS that are relevant to the approach to workflow mapping in this book, particularly simulation methods discussed in Volume II of this series. My method for workflow mapping is in fact very compatible with many types of KMS. But I do not want to limit my method for workflow mapping by associating it closely with KMS, for the reasons that I alluded to in Chapters 1 and 2. Most importantly, the Kmetz method is designed to be highly flexible and quickly accessible to nearly anyone who might want to use it, with no support from any kind of software.

A second motivation for differentiating my method of WFMA from KMS is derived from our discussion of companies as systems, in Chapter 5. What happens in a company depends critically on events in the “outside” world. This is no news to anyone, I hope, but in many KMSs, the unstated assumption is that we are working in a steady state, and that once the KMS is debugged and put in place, there ought to be fewer surprises than there were pre-KMS. My view is that a well-designed KMS should help achieve the functional consistency discussed in Figure 4.3 more quickly and cheaply than without it, for most of what we do, but that the universe will still deliver external shocks and surprises for which no KMS can prepare us. In terms of our expectations for KM, that is an important consideration—to some extent, experience will always prepare us to fight the previous war, not the one we face next. Companies and organizations of all kinds tend to wobble from crisis to crisis, and I think that is the norm rather than evidence of things being out of control. The value of the KMS is thus to help us cope, not necessarily to predict the future.

The other branch of KM has treated it as an embedded entity within the individual or group, and this is the primary realm of tacit knowledge as we have been discussing it in this book. The idea of “organizational learning,” that organizations can and should learn from experience and analysis is anything but new. Organizational adaptation as a form of “learning” was fundamental to how firms dealt with uncertainty, and Argyris has long held that organizations need to master “double loop learning,” where managers need to simultaneously deal with issues of strategy and the “elephants in the room” that make communication and

open exchange of information difficult.¹⁶ Senge extended this idea with the addition of systems thinking, and has enjoyed considerable popularity and support for his ideas in the business world; the fundamental role of information in organization design and processes has been an underlying theme in the works of March and Simon, and more recently by Galbraith, March, and Quinn.¹⁷

The information imperfections we discussed in Chapter 4 pose a major problem for KM, in that information can never be completely constrained within any fixed set of boundaries. Neither databases nor artificial-intelligence KMS programs can encompass all information or all rules that govern the contents of a cell in Figure 4.1. Moreover, KM based on capture of the knowledge of individuals and teams will encounter the same problems in attempting to put boundaries around tacit information specific to the person or team. Since much tacit knowledge (and the information processing rules that create and apply it) are associated with the experience and characteristics of the observers, the problems of categorizing and delimiting information are greatly complicated by the varying interpretations of information vectors (Figure 4.2) that these experiences cause.

Thus, the primary reason that tacit knowledge is treated differently in the context of a KMS is that it has always proven to be a major challenge to capture and include in a KMS—much of it is specific to individual ways of doing things and is very difficult to standardize, as is possible with articles, procedure manuals, or research reports. Indeed, the KPMG Knowledge Management Research Report for 2000 reported two interesting findings: one was that 50 percent of the 423 firms surveyed found that capturing tacit knowledge was a major problem with KM; second, KPMG reported that companies “still see KM as a purely technological solution.”¹⁸ Along with the difficulties in capturing tacit knowledge, one has to wonder about the extent to which this contributed to the other two major problems with KM reported in the survey—the lack of time to share knowledge (62 percent) and the failure to use knowledge effectively (57 percent).

As IT-intensive KMSs grew during the 1990s, others also commented on the difficulties of using IT to manage knowledge.¹⁹ Factors like culture and trust were also recognized as important conditional variables

in KM.²⁰ A rare but ingenious method to integrate highly individualized views of the world and the roles that different people play in organizations formed the basis of a robust method of organizing knowledge, by using the Internet to its fullest through a framework that claims to capture personal and organizational idiosyncrasies; this method has been applied in a number of complex organizations and is in the public domain, but it is not a simple method to apply.²¹

For those focused primarily on tacit knowledge in relation to KM, many have treated tacit knowledge as a branch of the personal growth or sensitivity training movements, and reified it as a body of knowledge that cannot, itself, be forced to yield to analysis—it is regarded as a deeply personal, internalized form of knowledge that can only be experienced. KM in this context becomes a matter of developing trust, facilitating group interaction and sharing of knowledge, and generally treating KM as a social process. As Birkinshaw noted, many of these processes are as old as the hills, and are familiar to anyone who knows about 1960s-style sensitivity training.²²

I suggest that two major issues limit what IT can do for KM, no matter how dedicated the effort to automate it. First, it is indeed true that much knowledge is tacit and fluid, and highly idiosyncratic; second, all of the social factors that affect people in every other facet of organizational life equally affect organizational knowledge, and perhaps even more greatly. Knowledge, after all, is power, and a somewhat naive assumption that I see in many KMSs is that information is neutral and its application is always objective. Any familiarity with organizational politics rapidly exposes these assumptions as highly simplistic.

For all these reasons, it is not the purpose of this book to attempt to reconcile these differing conceptions of KM or define the ultimate KMS. My view is that to a great extent, the preconceptions of the two “camps”—those who want the giant searchable knowledge base and thus limit themselves to formalized knowledge, and those who insist that knowledge is rooted in the collective psyche of the organization and can never be fully known or captured, both limit the scope of an organization’s knowledge base. I regard this as a false dichotomy, which is an artifact of way the camps define knowledge. My experience with WFMA suggests that both arguments are too narrow.

Capturing Knowledge in a Complex Workflow

Two illustrations from my NAVAIR studies help make this point, and allow us to see that workflow knowledge is neither neutral nor objective. (Readers may want to examine Figure A.1 in the Appendix before going on.) Figure 6.1 illustrates the avionics workflow as designed “by the

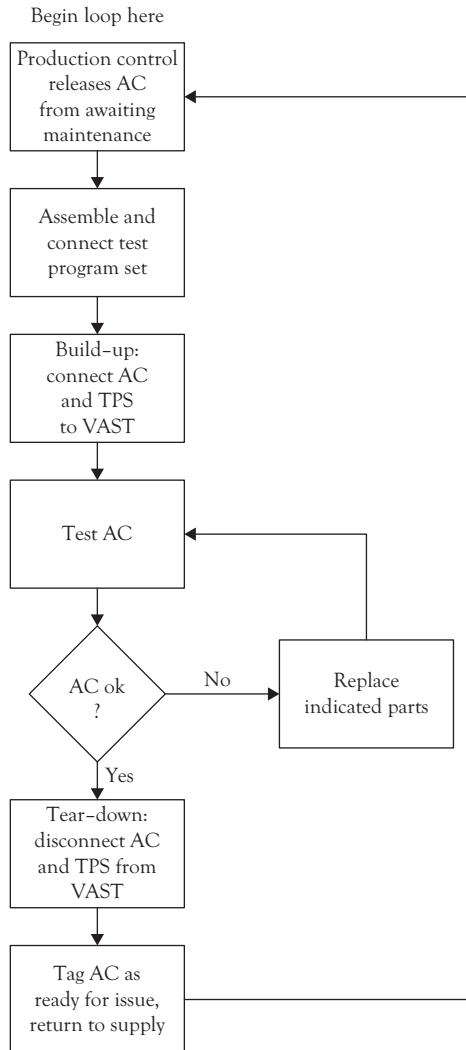


Figure 6.1 Simplified view of intended VAST shop test workflow

Source: Adapted by permission of the publishers, in *The Information Processing Theory of Organization* by John L. Kmetz (Farnham: Gower, 1998), p. 192.

book,” that is, the formal knowledge the system specified at that time to repair avionic devices in the Versatile Avionics Shop Test (VAST) shop in the lower right corner of Figure A.1.²³ The intention was that all of the VAST shops at avionics repair sites would follow this procedure, which was described at length in a military manual, fully approved by the Chief of Naval Operations.

The idea in Figure 6.1 is very simple—avionics components (ACs) enter the system on a First-In, First-Out basis, exactly as that concept is used in accounting for cost of goods sold in merchandising and inventory control. In this case, it simply meant that the sailors in the repair shop (the Versatile Avionics Shop Test, or VAST, system) that serviced these ACs took them in order of arrival, and did not need to be concerned about anything other than the time when the AC entered the shop; even that was determined by the local Production Control officer. As each AC was taken into the shop the sailor got the Test Program Set, the tape with the repair program on it and the cables needed to connect the AC to VAST, and then ran the program. The program provided information through a familiar CRT screen to tell the sailor what parts to remove and replace (usually one or more “cards” like those used in computers), and when that was done the AC was retested, presumably passed, and was returned to Supply to await its next use. We saw how the times to do these tasks in the repair workflow were accumulated in the overall distribution in Figure 5.4.

In my investigations of this workflow, I found that not a single carrier or shore site followed the formal process in Figure 6.1. The VAST shop was perceived as being a boondoggle and a bottleneck in the maintenance workflow, resulting in many aircraft deployed on the carriers being “not mission capable,” a serious issue at any time, but especially during the Cold War when this study took place. The issue was simply that when the carrier went to sea and air operations began, the available supply of replacement ACs for aircraft was quickly depleted, and it seemed that VAST was unable to repair them fast enough to replenish the Supply shelves and keep enough ACs available to support air activities. That, in turn, led to many disruptions of the larger-scale logistics support mission for the carriers, partly by disrupting avionics support operations at the shore sites, which further aggravated the problem. However, every time

the VAST testers were investigated to find out why they were “failing,” that is, not meeting performance and production expectations, the studies showed that VAST was not only meeting but exceeding its designed performance specifications. It was a very complicated problem, and a quite predictable behavior on the part of VAST shop supervisors was for them to do whatever they thought was right to get production flowing through their shops. What those local decisions and actions were, however, were never agreed upon by everyone in the system, whether shop supervisors or high-level Naval Material officers. Different theories regarding the cause of the problem abounded, and every manager in the different parts of the avionics repair cycle and broader logistics support “supply chain” tended to perceive the problem in terms of his or her function, and not infrequently one that could only be fixed with more money. By the time I became involved in this study, several advocates of a KMS as a “solution” to the problem had emerged, and three prototypes were actually in use at three locations (one on the *USS Forrestal*).

Figures 6.2a and 6.2b show a composite view of the actual workflow that was used in the avionics repair centers, on both the carriers and shore sites. Without explaining all the technical terms in the figures, what is immediately evident from comparing Figures 6.1 and 6.2 are the radical differences between them. The reason for these differences is that the tacit knowledge necessary to successfully repair ACs is significantly different from the information anticipated in the formal design of the system; even more important than the tacit knowledge itself was the information processing capacity required to manage the workflow. It is important to be aware that Figure 6.2 is a composite—in reality, not one of the carriers or shore bases surveyed did this process the same way. The only commonality was that the one shore site that attempted to follow the official procedure in Figure 6.1 simply “ground down” to a deficient level of production, which severely reduced aircraft operational readiness. It took a long time and a great deal of work, but in later years the lessons learned from the model in Figure 6.2 and considerable study of the entire maintenance cycle in Figure A.1 were integrated into an automated management system that became the standard for not only avionics maintenance but nearly all maintenance at this level in the entire carrier Navy.

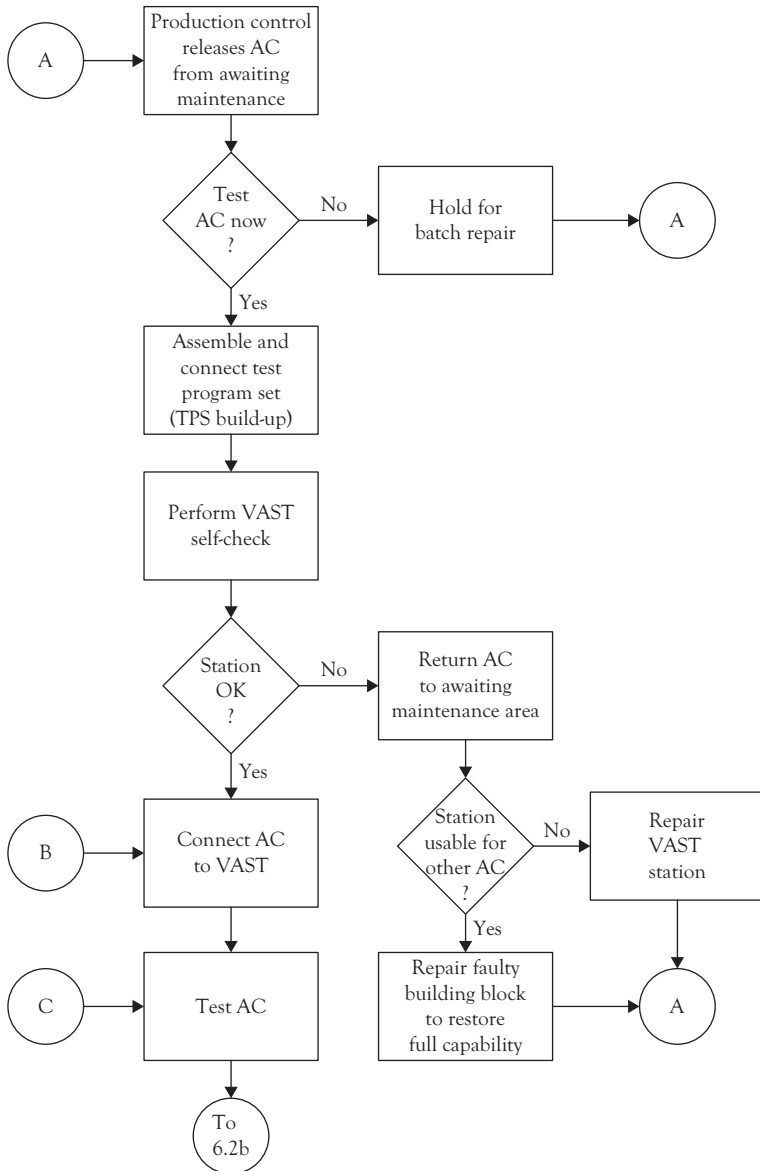


Figure 6.2a Actual VAST shop workflow

Source: Adapted by permission of the publishers, in *The Information Processing Theory of Organization* by John L. Kmetz (Farnham: Gower, 1998), pp. 204–205.

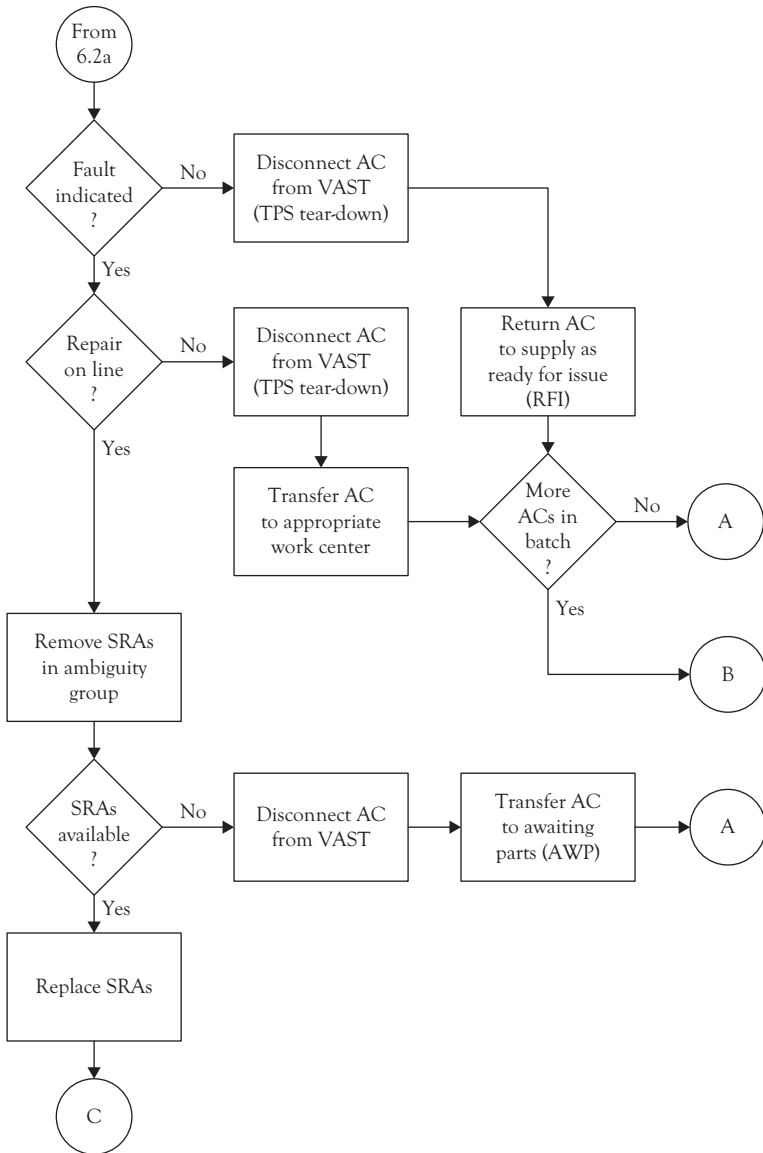


Figure 6.2b Actual VAST shop workflow

Source: Adapted by permission of the publishers, in *The Information Processing Theory of Organization* by John L. Kmetz (Farnham: Gower, 1998), pp. 204–205.

One illustration (of many possible) shows the relationship between tacit knowledge in this workflow and the success of the KMS that became part of the maintenance program. In Figure 6.1, an AC is released to the VAST shop for repair, and the next step is a “build-up” where the correct computer tape, cables, connecting devices, and so on are collected and attached so that the AC can be tested. After testing, a “teardown” was needed to remove all that gear and return it to its proper storage locations. In Figure 6.2, however, the second step is for the VAST shop supervisor to decide whether to do the build-up or hold the AC as part of a batch and go to the next AC in the queue. Why this radical departure from the formal procedure? Simply because the shop supervisor was under pressure to produce, and by holding ACs until batches were accumulated, the number of build-up–teardown cycles could be significantly reduced, allowing more VAST time for production. This tacit local adjustment made excellent sense from the perspective of increasing production, but at the cost of creating Supply stockouts of ACs being accumulated in batches, and thereby grounding aircraft. Unless one understood the complexities of the entire repair cycle, and not just the localized shop problem, it was possible for advocates of either the formal or tacit approach to argue that the other was the cause of the problem! Building a KMS around either approach alone would have been a mistake, while understanding the reasons for numerous tacit adjustments in the repair cycle was an absolute necessity for a successful KMS.

Complexity as Large-Scale Simplicity

These two figures illustrate not only how the Kmetz method of workflow mapping using a simple set of symbols and rules can accurately describe what a process is, but how such a simple tool may be used to discover and dissect processes that are in fact rather complex. Having found the underlying causes of inadequate repair-system production in the carrier Navy revealed that the tester presumed to be the culprit in the problem was not the problem at all—the real problem was that the change in system technology that VAST brought with it required increased information processing capacity to work. VAST itself was a great success, but the repair system could not unleash the potential of VAST until adequate

information-processing capacity to capture and utilize all the formal and tacit knowledge contained in the workflow was also designed into the workflow.

This is not unusual. In Chapters 4 and 5 we discussed the idea of systems. Organizations are typically rather complex systems, although the complexity of them may be difficult to see for many reasons. Those who know the history of a company understand it from that perspective, in ways that others newer to the firm do not; those who know certain technologies understand what those do in ways that nontechnies do not; and so on. One of the earliest analyses of an organization as a workflow illustrated how the information processing capacity inherent in one organization design was not only deficient for the work being done, but further created symptoms interpreted as “personality clashes” that were incorrectly believed to be the root cause of the performance problem in the workflow.²⁴ Very few people would predict workflow change to be an approach to resolve a “personality problem.”

That a system is complex, however, does not mean that it cannot be analyzed or that its processes cannot be understood. Two key ideas from systems theory are very helpful in the analysis of complexity. First, any “stable” organization is a collection of things that are all in motion at once. The analogy to a human body is appropriate—we look at a person sitting in a chair, and describe that individual as “at rest.” But in reality, there are subsystems in action throughout that person, and these never stop. Systems theorists use the term “dynamic homeostasis” to describe this condition, and it applies to most firms and organizations as well as to the people in them. One key to understanding a complex system is discovery and description of these internal subsystem processes. Second, much of the internal activity in a system is the kind of regulatory information flow needed to maintain that outward tranquil stability, in the form of the feedbacks and feedforwards discussed in Chapter 5; understanding self-regulation is thus also a key to unlocking complexity. Together, these are two keys to unlock complexity in any company or organizational system, because they reveal how that system monitors and regulates itself. This is also a much sounder basis for creating a successful KMS.

Figure 6.3 illustrates how a simple process can become very complex through the interactions of different feedbacks and feedforwards.

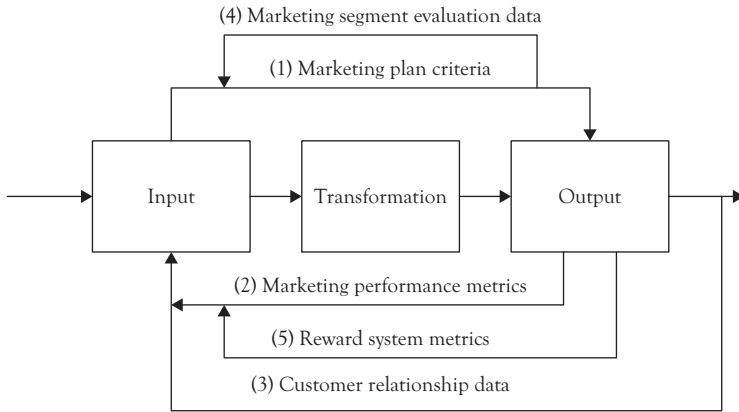


Figure 6.3 *Complex feedback and feedforward relationships*

A marketing plan (1) is produced to guide Company X in its quest to increase market share and performance of its distribution channels. Since market share and channel performance reflect the desired impact of the plan, they are the metrics (criteria) used to evaluate plan performance (2), acquired by standard measures of market share, time to market, and so on. Additional feedback is obtained through our Customer Relationship Management data (3), and these unfortunately indicate that what is being done to increase share and channel performance is alienating some of our major customers. The feedback data indicate that the plan is awry in that it has targeted the wrong market segment (4), and needs to be materially changed if it is to be anything more than a short-term success. However, since the reward system metrics (5) are based on the plan performance metrics, that is, market share and channel performance (2), there is strong pressure from those committed to the plan and doing well in terms of these metrics (and in the short term) not to change it. The die for internal conflict and politicking is cast. While not a real company, stories like this are common and easy to imagine; similar conflicts between R&D and marketing, accounting and manufacturing, and so on, are literally textbook cases of how complexity challenges organizational functioning. Most KMSs would be successful at tracking these metrics, but most cannot cope with the conflicts they might create, partly because the metrics track only specific individual objectives, and not the entire system goal; the tacit political issues that arise are completely beyond their ken.

This kind of complexity underscores the need for validation of workflow maps, and equally reinforces the need to start a WFMA process by mapping the process as it is now, not the one that “should be.” The map that “should be” will invariably incorporate the worldview and perception of the system that is held by those who propose it, and that view is inevitably biased in many ways.

The example in Figure 6.3 also illustrates another reality of organizational life that comes to the fore in any conflict situation, whether latent or overt. In the absence of valid data to defend a contrary case, political alliances will become the substitute for rational analysis and insight. To some extent, organizational politics is inevitable since there is no such thing as perfect information, and thus uncertainty can never be eliminated in decision making or any form of feedforward; it is also a problem in many forms of feedback as well. As Roger Hall stated in his study of management decisions in the failure of the old *Saturday Evening Post* magazine, not having good information on which to base an argument assures working in circumstances such that “power will follow uncertainty and that resources will follow power.”²⁵

“Knowledge is power.” Given that, information is seldom neutral. In the case of Company X’s marketing plan, one of the most potent forces at work to resist changing the plan, even in the face of customer data indicating the need, is the internal reward system of the company. We discussed several aspects of metrics in Chapter 2, when we were interested in measures that would help us find ways to improve processes. But another extremely important aspect of measurement is that it drives behavior. The psychologist E. L. Thorndike expressed it well in his 1911 “Law of Effect”; to paraphrase and shorten, it is that “behavior that is seen to be rewarded tends to be repeated, and that which is punished is likely to be extinguished.”²⁶ There are many implications of this law, but one of the clear implications for any WFMA project of significance is that it must have an organizational sponsor or owner who will shepherd it through company politics. A poorly designed KMS may not only blindly record whatever is measured, but may exacerbate the problems these measures cause. The process owner, whose role was discussed in Chapter 3, must be empowered to deal with these issues.

WFMA and Useful Knowledge Management

This discussion raises a point that I have never seen discussed in the KM literature, and that is the potential for knowledge to have negative impacts. These can occur in several ways. Obvious examples are those of the “skeletons in the closet,” where undesirable behavior or outcomes may threaten some individuals. More common is the case where employee X learns how to do something by asking employee Y how to do it, and then learns unnecessary or incorrect steps by copying that behavior. One of the more maddening things about employee training is that survey after survey shows that the most common form of learning on the job is through exactly this kind of mimicry. In one recent incident involving an old but successful financial institution, student-analysts working under my supervision in an MBA course found employees taking two days at the end of each month to use visual examination of spreadsheet printouts to identify and separately record events of interest, when a common macro could have done this in seconds. When shown this, the workers resisted changing their behaviors. Every organization has these kinds of stories to tell, and all of them are cases where some embedded tacit workflow knowledge has a negative impact, not positive.

Figure 6.4 summarizes both positive and negative impacts that can result from application of formal and tacit knowledge. Examples of all of these can be found in nearly every organization (in fact, I cannot think of any that do not exhibit all of these). Our obvious objective with a KMS is to aid in the generation and storage of knowledge with positive impact and to make that knowledge as available as possible. There are many success stories of this kind, and I have had the good fortune to contribute to some of these. In one case, a global chemical company found that each of the six major regions in their geographical organization structure had developed its own approach to tracking the results of research and development (R&D) work—while the good news was that at least some attempt was being made to capture the basic research that was the long-term lifeblood of this firm, there was neither consistency nor compatibility between the methods from any of these regions. I was able to help them develop a project-management based version of the Stage

		Impact of knowledge	
		Positive	Negative
Type of workflow knowledge	Explicit	Approved methods and procedures; “same page”	Work to rule: “red tape”
	Tacit	Effective local problem solutions; “wheelbooks”	Suboptimal local solutions; “batching”

Figure 6.4 *Positive and negative outcomes from workflow knowledge*

Gate system that was approved and formalized as the global model for managing R&D knowledge.²⁷

One might expect this to be an effort greeted with enthusiasm and a sense of positive expectations. Far from it—the almost universal reaction was “we’ve done OK so far without this, so why do we need it now?” In not a few cases, there was nearly palpable resentment of the project, based on feelings that management had ignored many good ideas sent up the R&D pipeline in the past, often missing good markets, which then got the “responsible managers” beaten up by their superiors. One excellent silicone scientist who had taken on management responsibilities spoke for a number of his colleagues in several facilities: “I know that what we have now is a mess, and that we throw a lot of money away for no good reason, but if we do what you’re suggesting we’re only going to expose a lot of people to criticism for failures that were not of their making—I’m afraid the net result of getting a good handle on our global R&D will be negative for the people we need to do the R&D.” I disagreed, of course, but I completely understood why he felt that way.

Nevertheless, the company forged ahead with implementation of the project, and eventually it became an internal success story for them, resulting in an R&D KMS that was indeed standardized, searchable,

linked to a consistent, stepwise review process for research as it moved from the laboratory toward the market, and was truly “global.” From my closing contacts with the company on this project, I learned that there was a period of disgruntlement with the new system, apparently centered on a group from the home country of this firm, who felt that they had rights to stay with their particular system given their corporate “ancestry.” Even they eventually came around. To the great credit of this firm, one of the things that helped was realignment of the reward system to recognize excellent basic research as well as the ability to shepherd research into product launches. The perception was that the reward system had paid too much attention to the latter, and the new KMS made it much easier to recognize good fundamental science and make it available to other global divisions of the company that might be better positioned to use it than the region that invented it. The new reward system was perceived as resulting in far more winners than losers, and with more ways to win, than in the past. The transition, however, was not easy.

The essential issue in this case is one that I believe confronts every organization, that it is impossible to set up a KM from a “neutral” position. The motivation for many KMS efforts is to resolve problems like the chemical company’s, and prevent them in the future. “Gaining a competitive advantage” is nearly a cliché in the world of KM, but it is difficult to do that without consideration of the current situation the organization is in, and that includes thorough understanding of its current processes.

The payoff of such process understanding can also be illustrated quite easily from another case. In visiting a DuPont company facility in Stevenage, England, in the 1980s, when ISO 9000 requirements were becoming a global standard in many industries, the plant manager gave our visiting student group a nearly evangelistic briefing on the effects of doing the process analysis needed to gain ISO certification. Among other findings, it was discovered that in the interest of upholding the DuPont reputation for quality, nearly 3,000 in-process tests were being run on product throughput batches, when fewer than 1,100 would have yielded the same quality control. At a cost of about GBP 3 per test, this was nearly \$6,000 of cost per batch that added no value. In another case, delivery times from batch start to completion were cut from 15 days to

1.5 days! First-pass yields improved from 72 to 92 percent at one plant, and on-time delivery increased from 70 to 90 percent. Even with results like these, it took about two years for the enthusiasm of the Stevenage plant to migrate back to the United States and finally become part of the formal operating rules there.

Tacit knowledge is even more challenging to get into a KMS. In the Navy studies, an excellent example of the positive impact of tacit knowledge was found in the “wheelbooks” that many VAST technicians made up from their own experience working with ACs in the shop (the term comes from ship pilots who would know local waters and make notes on hazards to avoid as they piloted a ship through a particular body of water). One case at a time, they noted false or ambiguous test results at various stages in a test program that required manual retesting; they noted difficulties in putting a large cable on an AC and the likelihood that this would result in false problem identification; and the like. In a number of cases, these wheelbooks became so well identified with individual sailors that you could tell from production numbers in the VAST shop whether they were at sea or in port.

At the same time, tacit knowledge led to outcomes that were negative for the organization as a whole, although they solved local problems, a condition described as “suboptimal” in the decision science literature. The logic behind holding individual ACs until a batch had accumulated (Figure 6.2) made excellent sense from the shop manager’s point of view, since it reduced the number of times individual build-ups and teardowns had to be done for different ACs. However, this meant that as the batch built up, stocks of ready-for-issue ACs of that type in the supply bins were being seriously depleted, often to the level of a complete stockout. Stockouts meant planes that could not fly missions, and that was serious.

There was a formal system for making tacit knowledge explicit, one part of which was identified by the name “Beneficial Suggestions” (or “benisuggs” to the sailors). Much of what was incorporated into the wheelbooks was put into the benisuggs system; the problem was that a complicated issue in, say, test software, would take years to be put into effect because it had to undergo a formal engineering investigation, followed by a report recommending specific changes that had to be reviewed and approved, and did not become part of the everyday operating procedure

until the next formal revision of the Naval Aviation Maintenance Plan (NAMP). The NAMP was continuously upgraded, but nothing happened quickly, for many good reasons, not the least of which is that “every line of the NAMP is written in blood”—many accidents and deaths over the years had forcibly pointed out the risks of taking shortcuts, no matter how sensible they seemed in other terms.

On one of my survey visits to (then) Naval Air Station Miramar, outside San Diego, I strongly requested that I be allowed to make some copies of one of these “wheelbooks” to use as support for the argument that we needed to create a better system for capturing, verifying, and distributing this kind of information throughout the carrier Navy, to the benefit of Fleet readiness. I was summoned to the base maintenance officer’s quarters, and told in no uncertain terms that if I pressed that issue again I would have my security clearance revoked and would be barred from the base. While wheelbooks were one of the most open “secrets” in the entire Fleet, none of the contents had been formally evaluated and approved—they were therefore technically violations of the NAMP, and did not exist.

For working in a dangerous and technologically complex world, the Navy probably has one of the best KMSs one could find, and it illustrates an underlying problem underlying many IT-intensive KMS designs in the corporate world. Every company wants to avoid being the next Xerox PARC (where the mouse, the graphical user interface, and networking were all invented but never developed or marketed), but the reality is that a corporate “knowledge base,” when considering the full organization, is much more difficult to systematize than simply finding and categorizing everything for the ultimate database, two jobs which are very difficult in their own right.

To be clear, this is not merely a matter of formal versus tacit knowledge. In the world of IT it is common to find many companies and organizations wanting to move from “homegrown” system maintenance and support processes to those more formally organized around project management. (Those “homegrown” processes are often grist for the mill for those who like to point out the failure of project management to meet its cost and schedule goals, and are a great source of material for the “Dilbert” comic strip.) I had a personal experience in this environment when I was

retained to train a large group of IT managers from a local division of a large global insurance company. The managers had struggled for years to work out their own version of a project management methodology, and could see little reason to abandon their system to learn another, especially when neither one could absolutely predict the outcome of an IT project nor guarantee its success; worse still, either approach required active, hands-on management and didn't allow the IT managers to do the "real work" of coding, testing, installing, and integrating software and hardware. Historically, much of contemporary IT matured in parallel with project management, and for many IT managers it was necessary to find a way to maintain and change IT through their own resources, without a formal management system. Many of them did this very well, and their feelings of resentment at having years of hard learning set by the wayside were understandable. I have encountered this in many IT applications of project management. In many respects, I think the growth and popularity of Agile project management in the IT world is a direct result of these continuing issues.

The unspoken and largely untested assumption in most KMSs is that finding and storing local "know-how," as the British like to call it, will always result in discovery of knowledge that is in the positive column of Figure 6.4, or that transforming tacit knowledge into explicit knowledge will have positive results. What I have attempted to illustrate with the examples here is that either of these outcomes is far from given, and that any stable system will first want to remain in equilibrium—resistance to change occurs for many reasons, some of them good from at least a local perspective.

Learning To Be Rigid

One of the close relatives to WFMA and KM is "organizational learning," which, as the name suggests, is the collective set of processes organizations can implement to acquire, upgrade, and make accessible knowledge about events that will help take better advantage of opportunities and avoid mistakes. I have argued that WFMA is a valuable resource to help us "know what we know" about our organization and its processes, and thus is a major source of knowledge content for a KMS as well as permitting

the organization to be modified over time for many reasons. George Patton was a scholar of history and is noted for his reference to Roman triumphs, where a slave rode in the triumphal chariot behind the emperor or general and repeatedly whispered to him that “all glory is fleeting.” Our knowledge base must somehow strive to capture that wisdom as well.

What I am alluding to is the possibility that as we change and adapt our organization, we tend to strive for efficiency and control; we want to adopt best practices that we have learned from others; we want to be “lean” in the best sense of that expression; and we want to understand all of our processes as fully as we can. These are all appropriate objectives, but beware of creating an organization that loses flexibility. While we achieve all of these laudable objectives, we must guard against driving out that which gives us the ability to innovate, to be creative and “think outside the box,” or to respond to the “bolt from the blue.”

More than once, I have heard scientists, engineers, and more than one nontechnical manager say, “this organization is fine—it’s the people who screw it up.” Apart from the obvious non sequitur in that statement, there is a long-recognized model of the “organization as a machine” (it is still worth reading Burns and Stalker’s classic book, where the terms “mechanistic” and “organic” were used to characterize two polar types of organization designs and their characteristics).²⁸ I think it is safe to conclude from much experience and a good deal of structured research that organizations are managed to achieve long-term stability and predictability, and the “machine model” captures that idea and its consequences very well.

The problem is that we have little choice but to manage to the “mechanistic” model to some extent. The Red Cross/Red Crescent and similar organizations are famous for being quick to respond to natural and manmade catastrophes, and are frequently first on the scene after such an event. These organizations are best known for their effectiveness rather than their efficiency; they may not use resources in a way that delivers the most “bang for the buck,” but delivers it first. Companies cannot run that way—they must keep costs under control and extract more from their resources on a continuing basis, while constantly trying to outdo their competitors. The mechanistic model of organization exists for good reason.

Mechanistic organizations need not be a product of conscious design. To a considerable extent, the negative and unintended consequences of innovation are among the major forces at work that led to the inability of the former MBNA America Bank to sustain itself as an independent company. As Christensen pointed out, the innovator's dilemma is that the very processes that make a company a successful innovator may also be the ones that set them up for failure.²⁹ One suggested answer to this dilemma was to set up other processes that help the right questions to be asked to make your company the new-technology "disrupter" in the marketplace, not the one disrupted.³⁰ While they may not explicitly label their approach part of a KMS, I would argue that the only distinction is the name, not what is done. An excellent study of the commercialization of flat-panel displays is much more explicit about the role of information and KM, and shows how these capabilities paid off in a truly global and highly competitive industry.³¹

As in any other kind of decision making, the operative word in organization design is "trade-off." While constantly moving toward a leaner, more efficient organization, companies must be sure to create processes and cultures where present success is questioned for sowing the seeds of failure, where markets are moving away from the business model that has made us successful, where mistakes and missed opportunities are evaluated to learn what not to do (not whom to punish, unless there is compelling reason), and where we are always looking for what the future holds. This is not trivial—as what Joyce and his colleagues found in their study of company success over a 10-year span, managers must learn how to keep at least six balls in the air at all times, and many are simply not that good as jugglers. One of those six balls is a strong performance-oriented culture.³²

In an excellent article for practitioners, Birkinshaw summarized six challenges in making KM work, and these are worth considering in light of the discussion above:³³

1. "KM is never a zero-based activity. It is already happening to some extent in every company." Approaching KM with the assumption that everything is being written on a "clean slate" is a gross oversimplification, one that may simply waste existing knowledge (and its

value) at least, but may be the basis for damaging resistance, at worst. Organization structures and professional networks play a significant role in KM, and these are the framework for a KMS.

2. KM is difficult to do well. It often requires fundamental changes in the way employees behave (and also reward systems that have been designed not to reward A while hoping for B); a good KM must be integrated with the culture, and that takes time. Changes in behavior are much more difficult to engineer (or “reengineer”) than entries in a database, and any company with a history of a few years has already institutionalized a great many processes based on explicit and tacit knowledge.
3. Most KM techniques end up looking just like the traditional techniques we have been using for years. The Balanced Scorecard is a common example of a KMS, whether it is called that or not; better communication practices (as with “centers of excellence”) are often the prescription for implementation of a KMS, but better communication is a general aspiration in nearly any firm.
4. KM is as much about generating new knowledge as it is about recycling existing knowledge. What must be recognized in construction of any KMS is that there are multiple organizational processes at work, each of which generates needs for new knowledge as well as requiring a knowledge base for effective deployment. The R&D process in the chemical company I described earlier is quite different from most of the batch- or continuous-process manufacturing of many chemicals and pharmaceuticals, and an effective KMS for one set of processes will not necessarily be the same as an effective KMS for others.
5. Information technology is never a substitute for social interaction. Whether we like it or not, the highest “bandwidth” of any communication is that found in face-to-face interaction. Even with the availability of the Internet, videoconferencing, and more private personal contact through social networking software, richness of interaction is reduced relative to personal contact.
6. Most firms, to some extent, do not know what they know. This issue was discussed extensively in Chapter 4. WFMA, properly done, can go a very long way toward closing this gap, and thus establish a

basis for next steps in setting up and operating a KM. What those steps are, or should be, will almost inevitably be changed by the findings of the WFMA process itself, but that is even more reason for continuing the effort—valid workflow maps provide a fact-based starting position to both set up the KMS you want and to avoid inadvertently creating a struggle between “old” and “new.” One would think, with all the advances in information technology in recent years, that firms and organizations would know how to use business intelligence systems to build organizational knowledge, but that appears to remain elusive.³⁴

An example of this tension and its potential costs can be found in the former MBNA America Bank mentioned earlier, which began business in 1982, grew explosively for 23 years, and was absorbed into the Bank of America in 2005. MBNA was originally founded in a former food store on the basis of a strong culture expressed as “Think of yourself as a customer.” Everyone was a customer, internally and externally, and as such one treated people with respect and expected it in return; there was a strong expectation of application of the “golden rule”; one never took one’s paycheck for granted—it was imprinted with “brought to you by the customer.” The hours were long, but the pay was excellent for the industry, and being a good corporate citizen meant that everyone pitched in and participated in community and charitable events, and for many early years employees were feasted and feted at an annual summer “corn boil.” Doing the right things the right way, it was believed, would not only make for a great place to work but a company that would grow.

And grow it did. MBNA ultimately became one of the biggest employers in the state of Delaware, surpassing the once-mighty DuPont, expanded operations into other states and other countries beginning with its initial foreign facility in the English Midlands, and was actively pursuing operations in Continental Europe at the time of its acquisition. Internally, the growth created a great need for a more codified and structured system of management to meet the challenges of growth. One of the carryover aspects of the original “food store” culture was that people often had to be moved upward and into different functional responsibilities very quickly, sometimes having only a few months to learn a job before

they were needed elsewhere. Under the old culture, one smiled and did one's best, and no one pointed out that unlike the former boss, you had just screwed up mightily in one of your first decisions—"customers" didn't treat each other that way.

As I learned in numerous episodes of working with MBNA, the bank was ultimately unable to get a handle on what it knew, and perhaps most importantly, what it did not know. One of its justifiable claims to ascendancy in the credit card industry was that it processed customer applications and information through personal processes, not relying on automated decision tools. This was very costly relative to competitors and could only be sustained through superior long-term operating margins. But the cost of doing business in a much larger and more complex organization that had grown well beyond being only a one-product company could not be controlled without knowing what you were doing and knowing it well. Ultimately, for all of its financial accomplishments and excellent corporate citizenship, those margins could not be sustained and MBNA did not survive.

Summary

WFMA necessarily will make much explicit and tacit knowledge visible in ways that did not exist before mapping and validation of workflows was done. As such, workflow mapping is inevitably linked to the "knowledge base" in any company. We should expect that two major types of results will obtain to successful WFMA projects. One will be the creation of a large database where validated workflow maps are an important part of the content of a larger KMS. The other is that many workflows and parts of workflows will be subjected to question and evaluation, not just for purposes of formal process improvement projects, but to basically understand why they are done as they are, and what the implications of change may be. In some cases, it may be that a suboptimal process is better left alone than changed, as with separating the handling of deposits and bank reconciliations in cash management, or things we do in the lab where outwardly "nobody seems to understand that we're running a business here."

The majority of cases where WFMA is applied result in immediate benefits, partly because this is a simple method and partly because it avoids some of the difficulties associated with IT-based process mapping systems we discussed in Chapter 1. It is easy to assume that this is the normal outcome. But while there is every reason to expect benefits and positive payoffs from WFMA, one should never assume that every such undertaking will be the harvesting of the proverbial low-hanging fruit.

The content of any knowledge base is relevant, to some extent, so long as the context of that knowledge remains relatively unchanged. A knowledge base containing technical data on the properties and assembly methods for incandescent light bulbs may be technically perfect, but of little value in a world where energy conservation is a universal need, and where new technologies have not only rendered incandescent bulbs obsolete, sometimes illegal (i.e., where laws have been passed forbidding their use), and where simple economies of use make them increasingly impractical. Once WFMA is begun, it will need to be continued so that workflows and processes are periodically reexamined and updated. Process improvement is indeed continuous improvement.

APPENDIX

A Brief Summary of the NAVAIR Study

Figure A.1 shows a high-level descriptive diagram of a complex repair process, used to capture the flow of materials in the repair cycle for aviation electronics (“avionics”) in the U.S. Naval air force (technically the Naval Air Systems Command, or NAVAIR), which has much to do with the development of Workflow Mapping and Analysis (WFMA) as covered in this book. Each of the blocks in this diagram is a high-level view of a great deal of detailed activity in different physical units on an aircraft carrier or at a shore base.

Note that unlike figures of workflows in this book, this is *not* a WFMA workflow map as discussed in Chapter 1. This diagram was created at the request of a Navy client for legibility and conservation of page space, so that the entire repair cycle could be shown on one page. It does not conform to the discipline of workflow mapping as covered in Chapter 1—it uses nearly square “rectangles” for the process blocks; there are multiple blocks with more than one exit arrow; there are double-headed arrows made of dotted lines instead of solid lines, and so on. The arrows have been numbered and labeled to guide readers through the flow of processes and locations represented in the diagram; however, without extensive documentation this diagram was much more difficult for many users to follow than a workflow map. It was accompanied by 10 pages of text that explained every process and arrow in detail, and took at least 30 minutes to present if there were no questions.

This actually portrays a very effective design for a maintenance process that is entirely contained on the carrier or shore site. It begins with a simple one-for-one exchange of an “avionics component” (AC) from an aircraft, where a technician-sailor pulls a failed unit from an aircraft (arrow 1) and turns it in to a screening unit, at which point it enters the

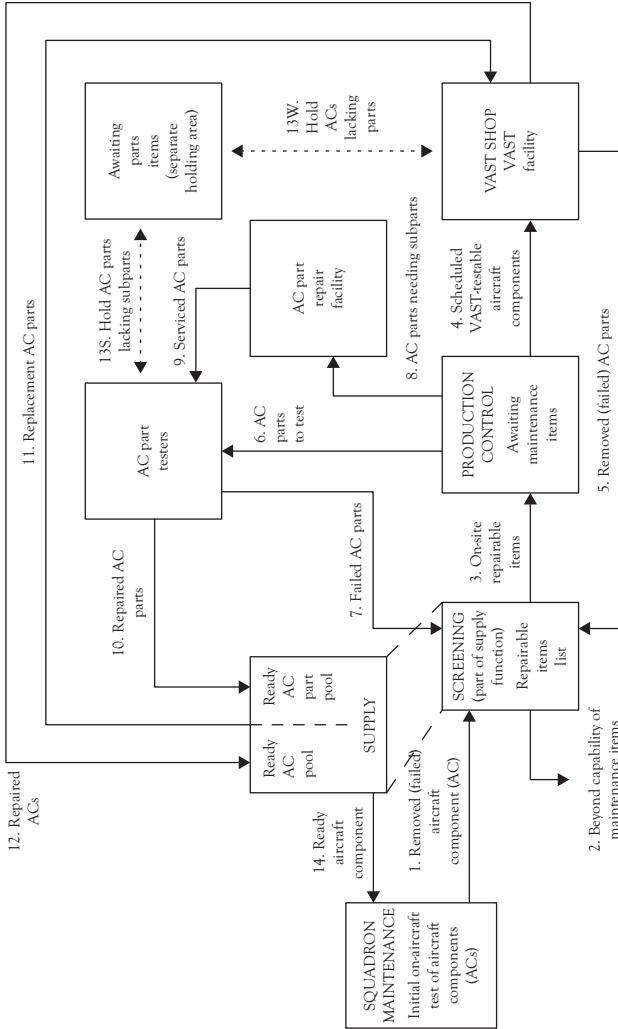


Figure A.1 The avionics repair cycle in the U.S. Navy Aircraft Intermediate Maintenance Department (AIMD)

repair cycle shown in arrows 2 through 13. However, the sailor should be able to obtain a Ready AC from Supply and install it in the aircraft (arrow 14) without delay, and thereby immediately return the aircraft to operational status. In short, this is nothing but a remove-and-replace process where the failed unit is removed, turned in, and exchanged for a unit of the same item ready for installation. All of the repair cycle in the rest of the map is invisible to the sailor on the deck fixing the aircraft.

The problem that the Navy encountered during the Cold War was that this workflow depended entirely on one tester, the Versatile Avionics Shop Test (VAST) system shown in the lower-right corner of Figure A.1. It seemed that this tester was unable to produce repaired ACs in sufficient numbers to support the requirements of the air wing once flight operations began in earnest on a deployment. That meant that the supply of Ready ACs in the Supply pool was depleted, and when a sailor brought a failed AC from an aircraft, there might not be (and often was not) a replacement AC available. This was a major readiness issue for the Navy for many years.

I became involved in the study of this workflow as a member of the faculty of the University of Southern California's Systems Management Center (long since reorganized and merged into other parts of the university), where we attempted to screen and compile a number of engineering and logistics support studies of the VAST tester and its performance, and from this determine why this tester appeared to be failing to meet its production expectations. We found, quite simply, that it was not failing, and that all we could do was propose a more comprehensive study of the entire avionics maintenance process in NAVAIR. This modest start led to my 15-year-long program of work in the avionics maintenance field.¹ We figured out that the problem with VAST was not a problem with VAST at all, but instead with the information processing capacity inherent in the Navy's maintenance system. Part of this required information technology to resolve, but the most important part was understanding the relationship between VAST and its effects on the maintenance workflow. Armed with that understanding, one of the units I worked with was able to design software and process changes to cope with the "VAST problem" by early 1983, and within the next two years

problems of production in the VAST avionics workflow had largely been eliminated.

(In doing this I became far more knowledgeable of the NAVAIR maintenance workflow than I had ever thought would happen, and so worked on the next generation of automatic testers, helped write technical proposals for major prime- and subcontractors, and had opportunity to see the results of earlier work on Fleet readiness. Without realizing it, I had become somewhat of a “national expert” on the subject, at least according to the Royal Canadian Air Force—I was asked in 1989 to do a study of their avionics maintenance process, and was so labeled by the principal logistics officer who had requested me.² The Canadians bought F/A-18 Hornets in 1980 to replace a group of older aircraft in their air force, and since the F/A-18 is a U.S. Navy design that used a newer version of VAST for avionics maintenance, they came to NAVAIR and the VAST contractor for help when they began to encounter the same problems the U.S. Navy did. Their problem was the same as the NAVAIR’s, and it was not the tester.)

Part of the problem was that the “repair cycle” was complicated, and even more so than Figure A.1 suggests. In brief, this figure reduces the complexity to the essential 14 flows in the repair cycle shown, which are:

- Arrow 1. An AC indicated as not working (from aircraft pretest or in-flight failure) is removed by a maintenance technician (sailor). The sailor turns it in to a Screening unit that is part of the larger Supply organization, and should immediately be given a replacement to install on the aircraft (Arrow 14).
- Arrow 2. A few things cannot be repaired on the carrier and are removed from the repair cycle as Beyond Capability of Maintenance.
- Arrow 3. On-site repairables are turned over to Production Control, and are now on Awaiting Maintenance status.
- Arrow 4. VAST-repairable ACs are released to the VAST shop for testing.
- Arrow 5. Parts identified by VAST testing (“cards” similar to computer cards and adapters) are removed and turned

- into Screening, and should immediately be replaced by new ones (Arrow 11) to repair the AC.
- Arrow 6. Cards repairable on site are sent to other specialized shops to be repaired, and some of these generate an additional failed-parts flow (Arrow 7), which is handled like ACs.
- Arrow 7. Small parts are screened and either put into the repair cycle or junked because they are consumables.
- Arrow 8. Repairable cards go to a specialized repair shop.
- Arrow 9. Serviced cards are sent for retesting.
- Arrow 10. Repaired cards (having passed testing) are returned to Supply for reuse.
- Arrow 11. Cards are sent from Supply to the VAST shop as noted in Arrow 5.
- Arrow 12. Repaired (passed testing) ACs are returned to Supply for reuse.
- Arrow 13. Either the AC or a removed card may need parts, and they may not be available (although they should be). When this happens the items are put on Awaiting Parts status, and held in a separate area. When parts arrive, Production Control releases the original items back into the repair cycle.
- Arrow 14. A Ready for Issue AC is given to the aircraft technician, as noted with Arrow 1.

What may not be immediately evident in this is that we have a very simple set of one-for-one exchanges nested in this workflow, which looks like Figure A.2 if there are no instances where a need for a part cannot be satisfied. Some nonrepairable items fall out along the way, but these are very few. By design, this closed-loop system permits carriers to operate with minimal shore support while on deployment, so there is a continual program in place to reduce parts removal from the repair cycle and maximize at-sea repair capability.

Of course, a perfectly uninterrupted supply of parts is not going to happen in the real world, and so the simplicity of Figure A.2 is never found in the operating Navy. “Lack of spares” is one problem that




Figure A.2 A simplified diagram of the avionics repair cycle

Source: Adapted by permission of the publishers, in *The Information Processing Theory of Organization* by John L. Kmetz (Farnham: Gower, 1998), p. 101. Copyright © 1998

confounds this model, and so it is simple and logical to conclude that if the model is not working in reality, the solution is to increase the number of spare parts in the system. At one point this logic was followed and NAVAIR bought \$46 million of spares (in the early 1980s); the result was to saturate the repair cycle with a larger number of items awaiting maintenance and awaiting parts (!), but no permanent improvement in the effectiveness of the repair cycle or operational readiness for the Fleet air wing.

In part, this failed experiment in increased sparing levels illustrated the importance of political factors in WFMA. Organizational politics is inevitable (Chapter 6) because information is inherently imperfect (Chapter 4), and this applies to the Navy as much as any other organization. In the military, politics are often manifested in the logistics support functions because “an army marches on its stomach.” For those familiar with the military, it is a given that there are 99 support personnel in place for every frontline soldier or pilot, and those support personnel are responsible for flows of a huge amount of specialized material. Empires are built on these flows, and these empires become the basis of enormous political and budgetary battles. The number and type of empires is suggested by Figure A.3, which suggests the size of the support structure in aircraft maintenance. This figure shows the support structure for the now-retired F-14 “Tomcat,” of “Top Gun” fame, but it applies to all aircraft and weapon systems (and in the case of the F-14, its retirement was partly because it required an average of 110 hours of collective maintenance for every hour it spent in the air).

All six of the aircraft logistic support elements immediately under the “deck” are individual empires, and there is a secondary universe of these

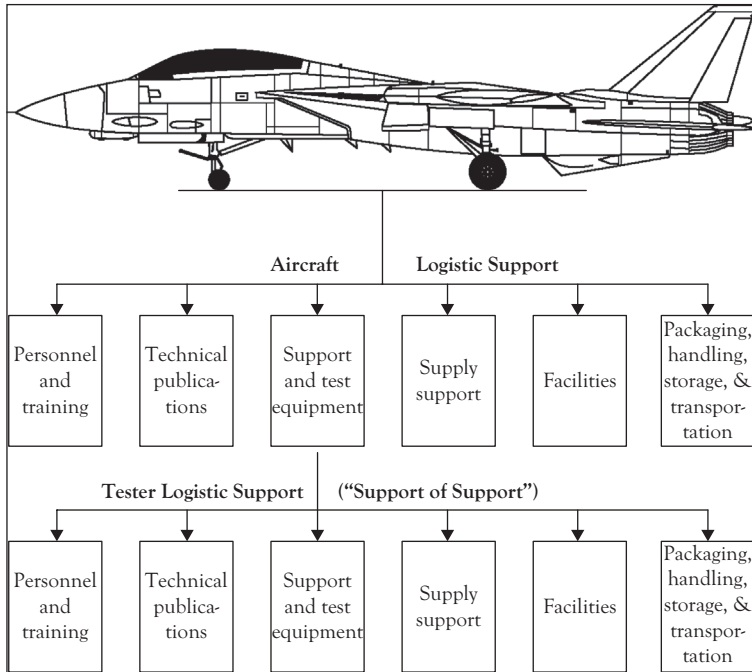


Figure A.3 *The logistics “tail” of aircraft maintenance*

Source: Adapted by permission of the publishers, in *The Information Processing Theory of Organization* by John L. Kmetz (Farnham: Gower, 1998), p. 94. Copyright © 1998

in the support and test equipment arena. I intend nothing pejorative or demeaning when I use the word “empire,” because an “emperor” is needed to organize, operate, and not occasionally defend, these logistic elements. All of them are absolute necessities to keep an effective fighting force deployed, and I have enormous respect for the managers and officers I have met in these functions. But they have to deal with imperfect information, just like the rest of us, and so politics is inevitable.

An unintended consequence of this kind of organization and the imperfection of information is that “defense of the empire” may lead to real or perceived information misuse. This had occurred in the Navy—the Chief of Naval Operations (CNO) was very aware of the problems with operational readiness afflicting the carrier Navy, and wanted to know how to fix them. The information he received, from different logistics managers inside NAVAIR, from contractors without, from experienced air wing commanders and many others, led to nowhere but confusion. NAVAIR

is part of the broad Navy matrix organization, as was the Bureau of Naval Material, the organization controlling all supply functions as well as test equipment. The CNO finally became so frustrated with his inability to get clear answers and effective solutions to his problems that he created a separate organization that reported directly to him, and kept information flows about this problem out of the regular chain of command. This organization was the Aircraft Intermediate Maintenance Support Office, or AIMSOS. AIMSOS was charged with full investigation of the avionics repair cycle, and as an independent consultant and researcher engaged in the same activity, I was assigned to coordinate my work with them.

This turned out to be a marriage made in heaven. The principal investigator for AIMSOS was a graduate of the USC program I had taught in, and from the first we had an excellent working relationship, from our intellectual orientations to our practical ideas. We were able to really test and evaluate many of our ideas using some simulation methods (some of which are described in Volume II of this series) as well as develop other tools for information technology support of the VAST workflow. Most importantly, AIMSOS provided us with the forum to get our recommendations to the right sponsor (Chapter 3), and thus into application. Of course, after resolving the problem AIMSOS had been intended to investigate, it was merged into the Bureau of Naval Material.

Before that happened, however, the story had a very happy ending—actually, two. One was that what we learned through this study became the basis for an automated decision process that later grew into the overall program used for all Naval maintenance, not just avionics—this was the Automated Production Management Module (APMM). APMM was both a simulation of the avionics maintenance system and a tool to manage it; it was part of the knowledge base for NAVAIR and a tool to manage that knowledge. It is now considered “legacy” software, but the system relationships that were embodied in it and were key to its effectiveness remain in place in the new generation of aircraft maintenance management to this day.

Second, when President Reagan decided to build five new Nimitz-class nuclear carriers as part of the 600-ship Navy he wanted, one of the implications was that the VAST production line would have to be restarted, at enormous cost. By the time this study was completed, all of the VAST

stations ordered by the Navy had been completed, the company that built them had merged into another larger firm, the plant site closed and sold off, and many of the engineers and production people who knew VAST had retired or were scattered around the country.


But by this time the VAST studies had brought us to fully understand what information needed to be acquired and processed to properly manage the material flow in avionics maintenance, and after creating and validating new programs to process that information, the productivity of the VAST shops exploded. Shore sites that had been provisioned with six stations for support and training were now doing the same job (actually, an increased workload) with two stations, and kept a third one for maintenance training. Carriers that had needed four stations were now meeting their requirements with two, and keeping a third for rotating, periodic preventive maintenance, and to be available as a backup. In addition, technical bugs had been investigated and worked out, as with any new system. This combination of active ingredients meant that enough VAST stations were freed from other bases and carriers to outfit the five new carriers, at a savings to the taxpayer of between \$200 and \$400 million, nearly a billion dollars today, depending on whose estimates you choose. That is a very happy ending, indeed.

Very early in the slightly more than 15 years I spent in this program of study and analysis, it became clear that the ability to capture what was happening in the enormous detail that makes up the “everyday” things that organizations do, and to convert that into information that could be communicated to those who could make decisions and affect the future course of events, was the single most important WFMA outcome of all. From that realization grew much of what we see in this book, and I hope it will be as useful and helpful to my readers as it was to me.


A Short Pictorial Gallery

I thought it might be helpful to show at least a line drawing or two, primarily of a VAST station, which very few people have ever seen, and of the primary aircraft and type of carrier this system supported during the height of the Cold War. These follow on the next four pages.


The first illustration is a line drawing of a VAST station. Test equipment is just not as sexy as the aircraft it supports or the carriers on which it was installed, so it seems there are very few pictures or drawings of VAST anywhere, including on the part of the manufacturer and their engineers and others who built it. The line drawing gives an idea of what electronic units were in the station, and some idea of how it appeared physically. It is a big machine—78 inches tall, 34 inches deep, 30 feet, 4 inches long, and weighing nearly 19,000 pounds.




F-14A “Tomcat” air superiority and defense fighter aircraft




E-2C “Hawkeye” command, control, communication, and intelligence aircraft



S-3A “Viking” Antisubmarine warfare aircraft



A-7E “Corsair II” strike (surface attack) aircraft



F/A-18 “Hornet” air superiority and strike aircraft

Notes

Chapter 1

1. Harris (1999).
2. www.wfmc.org
3. <http://www.BPMInstitute.org>; Diaz and Rudden (2010).
4. www.BPMInstitute.org
5. van der Aalst, van Hee, and Houben (2002).
6. Carr (1999a, 1999b); Malone et al. (1999); Malone, Crowston, and Herman (2003).
7. Phios.com
8. APQC (2010).
9. Cowan (2010, 2012).
10. Montero and Scott (1996).
11. Alonso, Agrawal, El Abbadi, and Mohan (1997).
12. Diaz and Rudden (2010); Haugh (2010).
13. Diaz and Rudden (2010); Haugh (2010); Ketabchi (2010).
14. Miller (1956).

Chapter 2

1. Brassard, Finn, Ginn, and Ritter (2004).
2. Drucker (1954).
3. Joyce, Nohria, and Roberson (2003).
4. Joyce et al.
5. Dowdy, Dorgan, and Rippin (2005).
6. Likert (1967).
7. Wysocki (2004).
8. Foster (2007); Hendricks and Singhal (2001).
9. Shewhart (1931).
10. Personal communication from Mr. Randy Yeager, Quality Control Manager at Harley-Davidson, 2003.
11. Palmer (2007).

Chapter 3

1. Desai et al. (2014).
2. Kothari et al. (2012); Song (2015).
3. Arnett and Wittmann (2014).
4. McElheran (2015).
5. Suriadi et al. (2014).
6. Harris (1999).
7. Hammer (1990); Hammer and Champy (1993).
8. Earl and Khan (1994); Strassmann (1994).
9. Fire and forget? (1996).
10. Maull et al. (1995).
11. Greeley (2012).
12. Branco et al. (2014).
13. Wang et al. (2014).
14. vom Brocke et al. (2014).
15. Forrester Consulting (2013).
16. Carson (1998).

Chapter 4

1. Kmetz (1998).
2. Logan (2009). Donald Rumsfeld, the former Secretary of Defense, is frequently credited with this idea, but as Logan points out, this concept considerably predates him. I was first exposed to this in the 1960s in a training-the-trainers program by Dr. David Brown of George Washington University, and when I asked him who had originated it, he did not know.
3. Taleb (2010).
4. Kahneman (2011).
5. Wildavsky (1983).
6. Wilensky (1967).
7. von Bertalanffy (1968).
8. Senge (1990, 1994).
9. Strassmann (1985).
10. Farrell, Mandel, and Weber (1995).
11. The right direction? (1995).

Chapter 5

1. von Bertalanffy (1968).
2. Boulding (1956).
3. Williamson (1975).
4. Knight and McDaniel (1979).
5. Kerr (1975).
6. NBCNews (2006).
7. Pelley (2009); Hambling (2014).

Chapter 6

1. Kmetz (1975, 1977).
2. Alavi and Leidner (2001).
3. Tuomi (1999).
4. Alavi and Leidner (2001); Mertins, Heising, and Vorbeck (2003).
5. Starbuck (2006).
6. Harmon (2007); King, Chung, and Haney (2008).
7. Nickols (2000).
8. Drucker (1972).
9. Hermanson (1964).
10. Rajan (2004); Gooderham and Ulset (2007); Kazi and Wolf (2006).
11. Chua (2007).
12. Choi, Poon, and Davis (2008).
13. Mårtensson (2000).
14. Davenport (2005).
15. Davenport (2007).
16. Cyert and March (1993); Argyris (1977).
17. Senge (1990); March and Simon (1993); Galbraith (1977);
March (1999); Quinn (1992).
18. Parlby (1999).
19. Chua (2007); Walsham (2001).
20. Alavi, Kayworth, and Leidner (2005–2006); DeLong and Fahey
(2000); Renzl (2008).
21. Jacobs (2009).
22. Birkinshaw (2001).
23. NAVAIRINST 4790.2B (1976).

24. Chapple and Sayles (1961).
25. Hall (1976).
26. Thorndike (1911).
27. Cooper, Edgett, and Kleinschmidt (2002a, 2002b).
28. Burns and Stalker (1958).
29. Christensen (1997).
30. Christensen and Raynor (2003).
31. Murtha, Lenway, and Hart (2001).
32. Joyce, Nohria, and Roberson (2003).
33. Birkinshaw (2001).
34. Kmetz and Naik (1999); Shollo and Galliers (2015).

Appendix

1. Kmetz (1984, 1998).
2. Kmetz (1990).

References

- Alavi, M., T.R. Kayworth, and D.E. Leidner. 2005–2006. “An Empirical Examination of the Influence of Organizational Culture on Knowledge Management Practices.” *Journal of Management Information Systems* 22, no. 3, pp. 191–224.
- Alavi, M., and D.E. Leidner. 2001. “Review: Knowledge Management and Knowledge management Systems: Conceptual Foundations and Research issues.” *MIS Quarterly* 25, no. 1, pp. 107–36.
- Alonso, G., D. Agrawal, A. El Abbadi, and C. Mohan. 1997. “Functionality and Limitations of Current Workflow Management Systems.” *IEEE Expert, Special Issue on Cooperative Information Systems* 12, no. 5, pp. 105–11.
- APQC. 2010. “Process Classification Framework.” Ver 6.1.1. <http://www.apqc.org/process-classification-framework> (accessed May 12, 2014).
- Argyris, C. 1977. “Double Loop Learning in Organizations.” *Harvard Business Review* 55, no. 5, pp. 115–25.
- Arnett, D.B., and C.M. Wittmann. 2014. “Improving Marketing Success: The Role of Tacit Knowledge Exchange Between Sales and Marketing.” *Journal of Business Research* 67, pp. 324–31.
- Birkinshaw, J. 2001. “Making Sense of Knowledge Management.” *Ivey Business Journal* 65, no. 4, pp. 32–38.
- Boulding, K.E. 1956. “General Systems Theory—The Skeleton of Science.” *Management Science* 2, no. 3, pp. 107–208.
- Branco, M.C., Y. Xiong, K. Czarnecki, J. Küster, and H. Völzer. 2014. “A Case Study on Consistency Management of Business and IT Process Models in Banking.” *Software & Systems Modeling* 13, no. 3, pp. 913–40.
- Brassard, M., L. Finn, D. Ginn, and D. Ritter. 2004. *The Six Sigma Memory Jogger*. 2nd ed. 12B Manor Parkway, Salem, NH: Goal/QPC.
- Burns, T., and G.M. Stalker. 1958. *Management and Innovation*. London: Tavistock Institute.
- Carr, N.G. 1999a. “A New Way to Manage Process Knowledge.” *Harvard Business Review* 77, no. 5, p. 24.
- Carr, N.G. 1999b. “Redesigning Business.” *Harvard Business Review* 77, no. 6, p. 19.
- Carson, I. 1998. “Survey: Manufacturing.” *The Economist* 347, no. 8073.
- Chapple, E.D., and L. Sayles. 1961. “Work Flow as the Basis for Organization Design.” In *The Measure of Management: Designing Organizations for Human Effectiveness*, ed. E.D. Chapple, 373–94. New York: Macmillan.

- Choi, B., S.K. Poon, and J.G. Davis. 2008. "Effects of Knowledge Management Strategy on Organizational Performance: A Complementarity Theory-Based Approach." *Omega* 36, no. 2, pp. 235–52.
- Christensen, C.M. 1997. *The Innovator's Dilemma*. Boston, MA: Harvard Business School Publishing.
- Christensen, C.M., and M.E. Raynor. 2003. *The Innovator's Solution: Creating and Sustaining Successful Growth*. Boston, MA: Harvard Business School Publishing.
- Chua, A.Y.K. April 28, 2007. "The Curse of Success." *The Wall Street Journal*, sec. R, pp. 8–9.
- Cooper, R.G., S.J. Edgett, and E.J. Kleinschmidt. 2002a. "Optimizing the Stage-Gate Process: What Best-Practice Companies Do I." *Research-Technology Management* 45, no. 5, pp. 21–27.
- Cooper, R.G., S.J. Edgett, and E.J. Kleinschmidt. 2002b. "Optimizing the Stage-Gate Process: What Best-Practice Companies Do II." *Research-Technology Management* 45, no. 6, pp. 43–49.
- Cowan, M. 2010. "Organizing Enterprise-wide IT Efforts Using APQC's Process Classification Framework," <http://www.apqc.org/OrganizingEnterpriseWideITEffortsUsingAPQCsPCF> (accessed November 6, 2010).
- Cowan, M. 2012. "Marathon Oil Corporation." Marathon Oil Corporation: APQC's Process Classification Framework in Action. http://www.apqc.org/MarathonOil_PCFCaseStudy/ (accessed 27 August, 2012).
- Cyert, R.M., and J.G. March. 1993. *A Behavioral Theory of the Firm*. 2nd ed. Cambridge, MA: Blackwell.
- Davenport, T.H. 2005. *Thinking for a Living: How to Get Better Performance and Results from Knowledge Workers*. Boston, MA: Harvard Business School Press.
- Davenport, T.H. 2007. "Information Technologies for Knowledge Management." In *Knowledge Creation and Management: New Challenges for Managers*, eds. K. Ichijo and I. Nonaka, 97–117. Oxford: Oxford University Press.
- DeLong, D., and L. Fahey. 2000. "Diagnosing Cultural Barriers to Knowledge Management." *Academy of Management Executive* 14, no. 4, pp. 113–22.
- Desai, D., P. Kotadiya, N. Makwana, and S. Patel. 2015. "Curbing Variations in Packaging Process Through Six Sigma Way in a Large-Scale Food-Processing Industry." *Journal of Industrial Engineering International* 11, no.1, pp. 119–29.
- Diaz, A., and J. Rudden. 2010. "Lessons Learned from the First Decade of BPM." <http://www.ibm.com/bpm> (accessed January 11, 2011).
- Dowdy, J., S. Dorgan, and T. Rippin. 2005. *Management Matters*. London: McKinsey & Co.
- Drucker, P.F. 1954. *The Practice of Management*. New York: Harper.
- Drucker, P.F. 1972. *America's Next Twenty Years*. Freeport, NY: Books for Libraries Press.

- Earl, M., and B. Khan. 1994. "How New is Business Process Redesign?" *European Management Journal* 12, no. 1, pp. 20–30.
- Farrell, C., M.J. Mandel, and J. Weber. 1995. "Riding High." *Business Week*, 7958, no. 332, pp. 134–46.
- Fire and Forget? 1996. *The Economist*, 7958, no. 332, 51–52.
- Forrester Consulting. 2013. "The State of Modern Product Delivery: Challenges and Trends Making Product Delivery Responsive, Iterative, and Collaborative in the Age of the Customer." http://www.jamasoftware.com/wp-content/uploads/documents/forrester-the-state-of-modern-product-delivery_forrester-thought-leadership-paper (accessed January 5, 2014).
- Foster, S.T., Jr. 2007. "Does Six Sigma Improve Performance?" *Quality Management Journal* 14, no. 4, pp. 7–20.
- Galbraith, J.R. 1977. *Organization Design*. Reading, MA: Addison-Wesley.
- Gooderham, P.N., and S. Ulset. 2007. "Telenor's Third Way." *European Business Forum (Online)* 31.
- Greeley, B. April 30, 2012. "We're from Private Equity, and We're Here to Help." *Bloomberg Business Week*, 55–59.
- Hall, R.I. 1976. "A System Pathology of an Organization: The Rise and Fall of the Old Saturday Evening Post." *Administrative Science Quarterly* 21, pp. 185–211.
- Hambling, D. 2014. "Cold Fusion Continues to Progress Stealthily into the Mainstream." <http://www.wired.co.uk/news/archive/2014-01/15/cold-fusion-moves-into-mainstream> (accessed March 3, 2014).
- Hammer, M. 1990. "Reengineering Work: Don't Automate, Obliterate." *Harvard Business Review* 68, no. 4, pp. 104–12.
- Hammer, M., and J. Champy. 1993. *Reengineering the Corporation: A Manifesto for Business*. New York: HarperCollins.
- Harmon, P. 2007. *Business Process Change: A Guide for Business Managers and BPM and Six Sigma Professionals*. 2nd ed. Burlington, MA: Morgan Kaufman.
- Harmon, P., and C. Wolf. 2014. "The State of Business Process Management—2014." Ch. 1, p. 6. <http://www.bptrends.com/bpt/wp-content/uploads/BPTrends-State-of-BPM-Survey-Report.pdf> (accessed April 17, 2014).
- Harris, R.L. 1999. *Information Graphics: A Comprehensive Illustrated Reference*. New York: Oxford University Press.
- Haugh, K. 2010. *BPM and beyond* (Slide show). <http://www.slideshare.net/deanwhit/bpm-andbeyondslide> (accessed December 17, 2010).
- Hendricks, K.B., and V.R. Singhal. 2001. "The Long-Run Stock Price Performance of Firms with Effective TQM Programs." *Management Science* 47, no. 3, pp. 359–68.
- Hermanson, R.H. 1964. *Accounting for Human Assets*. East Lansing, MI: Bureau of Business and Economic Research, Graduate School of Business Administration, Michigan State University.

- Jacobs, B.E. 2009. "Complex Distributed Process Management Using Process Libraries and Electronic Handbooks (Where Shakespeare Meets Freud)." <http://www.ehbs.org/> (accessed March 11, 2009).
- Joyce, W., N. Nohria, and B. Roberson. 2003. *What (Really) Works: The 4+2 Formula for Sustained Business Success*. New York: Harper Business.
- Kahneman, D. 2011. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.
- Kazi, A.S., and P. Wolf. 2006. *Real-Life Knowledge Management: Lessons from the Field*. Helsinki, Finland: Knowledgeboard.
- Kerr, S. 1975. "On the Folly of Rewarding A, While Hoping for B." *Academy of Management Journal* 18, no. 4, pp. 783–94.
- Ketabchi, M.A. 2010. "Responsive Process Management." <https://www.progress.com/> (accessed May 16, 2012).
- King, W.R., T.R. Chung, and M.H. Haney. 2008. "Knowledge Management and Organizational Learning (Editorial)." *Omega* 36, no. 2, pp. 167–73.
- Kmetz J.L. 1975. *Technology and Organization Structure: The Relationship Between Contextual Variables and Structure Variables in Manufacturing and Service Organizations*. Ann Arbor, MI: University of Maryland.
- Kmetz, J.L. 1977–1978. "A Critique of the Aston Studies and Results with a New Measure of Technology." *Organization and Administrative Sciences* 8, no. 4, pp. 123–44.
- Kmetz, J.L. 1984. "An Information Processing Study of a Complex Workflow in Aircraft Electronics Repair." *Administrative Science Quarterly* 29, no. 2, pp. 255–80.
- Kmetz, J.L. 1990. *Management of Canadian Forces Air Command ATE Facilities*. Manual for Canadian Forces (Royal Canadian Air Force); prepared under Harris Corp. contract ATS0001-02C2-1989.
- Kmetz, J. L. 1998. *The Information Processing Theory of Organization: Managing Technology Accession in Complex Systems*. Aldershot, England: Ashgate.
- Kmetz, J.L., and Chandrasekhar Naik. 1999. "Competitive Intelligence Versus Competitive Knowledge in the Multinational Enterprise." *Proceedings of the 8th World Business Congress*. Monterrey, CA.
- Knight, K.E., and R.R. McDaniel, Jr. 1979. *Organizations: An Information Systems Perspective*. Belmont, CA: Wadsworth.
- Kothari, A., D. Rudman, M. Dobbins, M. Rouse, S. Sibbald, and N. Edwards. 2012. "The Use of Tacit and Explicit Knowledge in Public Health: A Qualitative Study." *Implementation Science* 20, no. 7, p. 12.
- Likert, R. 1967. *The Human Organization: Its Management and Value*. New York: McGraw-Hill.
- Logan, D.C. 2009. "Known Knowns, Known Unknowns, Unknown Unknowns and the Propagation of Scientific Enquiry." *Journal of Experimental Botany* 60, no. 3, pp. 712–14.

- Malone, T.W., K. Crowston, and G.A. Herman, eds. 2003. *Organizing Business Knowledge: The MIT Process Handbook*. Cambridge, MA: MIT Press.
- Malone, T.W., K. Crowston, J. Lee, B. Pentland, C. Dellarocas, G. Wyner, J. Quimby, C.S. Osborn, A. Bernstein, G. Herman, M. Klein, and E. O'Donnell. 1999. "Tools for Inventing Organizations: Toward a Handbook of Organizational Processes." *Management Science* 45, no. 3, pp. 425–43.
- March, J.G. 1999. *The Pursuit of Organizational Intelligence*. Oxford: Blackwell Publishers.
- March, J.G., and H.A. Simon. 1993. *Organizations*. 2nd ed. Oxford: Blackwell.
- Martensson, M. 2000. "A Critical Review of Knowledge Management as a Management Tool." *Journal of Knowledge Management* 4, no. 3, pp. 204–16.
- Maull, R.S., A.M. Weaver, S.J. Childe, P.A. Smar, and J. Bennett. 1995. "Current issues in Business Process Re-Engineering." *International Journal of Operations & Production Management* 15, no. 11, pp. 37–52.
- McElheran, K. 2015. "Do Market Leaders Lead in Business Process Innovation? The Case(s) of E-Business Adoption." *Management Science* 61, no. 6, pp. 1197–216.
- Mertins, K., P. Heising, and J. Vorbeck, eds. 2003. *Knowledge Management: Concepts and Best Practices*. 2nd ed. New York: Springer-Verlag.
- Miller, G.A. 1956. "The Magical Number Seven, Plus or Minus Two: Some Limits on Our Ability for Processing Information." *Psychological Review* 63, no. 2, pp. 343–55.
- Montero, L., and C.T. Scott. 1996. "Improving the Quality of Component Business Systems with Knowledge Engineering." *Proceedings of Eighth International Conference on Software Engineering and Knowledge Engineering*.
- Murtha, T.P., S.A. Lenway, and J.A. Hart. 2001. *Managing New Industry Creation: Global Knowledge Formation and Entrepreneurship in High Technology—The Race to Commercialize Flat Panel Displays*. Stanford, CA: Stanford University Press.
- NAVAIRINST 4790.2B. 1976. *Naval Aviation Maintenance Program*. NAVAIRINST 4790.2, Washington, DC: Department of the Navy.
- NBCNews. 2006. "South Korean Cloning Scientist Indicted: Scandal over Faked Stem Cell Research Rattled the Scientific Community." http://www.nbcnews.com/id/12749497/ns/health-cloning_and_stem_cells/t/south-korean-cloning-scientist-hwang-indicted/#.VbEyXPkVliY (accessed May 13, 2006).
- Nickols, F. 2000. "The Knowledge in Knowledge Management." In *The Knowledge Management Yearbook*, 12–21. Boston, MA: Butterworth-Heinemann.
- Palmer, N. 2007. "A Survey of Business Process Initiatives." [http://www.bptrends.com/bpt/wp-content/surveys/FINAL PDF 1-23-07.pdf](http://www.bptrends.com/bpt/wp-content/surveys/FINAL%20PDF%201-23-07.pdf) (accessed April 17, 2007).
- Parlby, D. 1999. Knowledge Management Research Report 2000. London: KPMG Consulting.

- Pelley, S. 2009. "Cold Fusion is Hot Again." CBSNews.com. <http://www.cbsnews.com/news/cold-fusion-is-hot-again/> (accessed December 11, 2009).
- Quinn, J.B. 1992. "The Intelligent Enterprise: A New Paradigm." *Academy of Management Executive* 6, no. 4, pp. 48–63.
- Rajan, A. 2004. "Getting Tacit Knowledge to Work." *The Financial Times*, March 29, p. 5.
- Renzl, B. 2008. "Trust in Management and Knowledge Sharing: The Mediating Effects of Fear and Knowledge Documentation." *Omega* 36, no. 2, pp. 206–21.
- The Right Direction? 1995. *The Economist*, 7940, no. 329, pp. 23–27.
- Senge, P. 1990. *The Fifth Discipline: The Art & Practice of the Learning Organization*. New York: Currency Doubleday.
- Senge, P.M., A. Kleiner, C. Roberts, R.B. Ross, and B.J. Smith. 1994. *The Fifth Discipline Fieldbook: Strategies and Tools for Building a Learning Organization*. New York: Currency Doubleday.
- Shewhart, W.A. 1931. *Economic Control of Quality of Manufactured Product*. New York: van Nostrand.
- Shollo, A., and R.D. Galliers. 2015. "Towards an Understanding of the Role of Business Intelligence Systems in Organizational Knowing." *Information Systems Journal*. doi: 10.1111/isj.12071.
- Song, L.L., X.Q. Guo, and C. Wang. 2015. "Research on Model-Driven Simulation Approach for Healthcare Information System." *International Conference on Artificial Intelligence and Industrial Engineering (AIIE 2015)* Atlantis Press.
- Starbuck, W.H. 2006. *The Production of Knowledge: The Challenge of Social Science Research*. New York: Oxford.
- Strassmann, P.A. 1985. *Information Payoff: The Transformation of Work in the Electronic Age*. New York: Free Press.
- Strassmann, P.A. 1994. "The Hocus-Pocus of Reengineering." *Across the Board*, 35–38.
- Suriadi, S., B. Weiss, A. Winklemann, A.H.M. ter Hofstede, M. Adams, R. Conforti, C. Fidge, M. La Rosa, C. Ouyang, M. Rosemann, A. Pika, and M. Wynn. 2014. "Current Research in Risk-Aware Business Process Management—Overview, Comparison, and Gap analysis." *Communications of the Association for Information Systems* 34, no. 1, pp. 933–84.
- Taleb, N.N. 2010. *The Black Swan: The Impact of the Highly Improbable*. 2nd ed. New York: Random House.
- Thorndike, E.L. 1911. *Animal Intelligence: Experimental Studies*. New York: Macmillan.

- Tuomi, I. 1999. "Data is More Than Knowledge: Implications of the Reversed Knowledge Hierarchy for Knowledge Management and Organizational Memory." *System Sciences 1999: Proceedings of the 32nd Annual Hawaii International Conference*.
- van der Aalst, W., K.M. van Hee, and G.J. Houben. 2002. *Workflow Management: Models, Methods, and Systems*. Cambridge, MA: MIT Press.
- vom Brocke, J., S. Debortoli, O. Muller, and N. Reuter. 2014. "How in-Memory Technology Can Create Business Value: Insights from the Hilti Case." *Communications of the Association for Information Systems* 34, no. 7, pp. 151–68.
- von Bertalanffy, L. 1968. *General Systems Theory*. New York: George Brazillier.
- Walsham, G. 2001. "Knowledge Management: The Benefits and Limitations of Computer Systems." *European Management Journal* 19, no. 6, pp. 599–608.
- Wang, J., T. Jin, R.K. Wong, and L. Wen. 2014. "Querying Business Process Model Repositories: A Survey of Current Approaches and Issues." *World Wide Web (Springer Science+Business Media)* 17, no. 3, pp. 427–454.
- Wildavsky, A. 1983. "Information as an Organizational Problem." *Journal of Management Studies* 20, no. 1, pp. 29–40.
- Wilensky, H.L. 1967. *Organizational Intelligence: Knowledge and Policy in Government and Industry*. New York: Basic Books.
- Williamson, O.E. 1975. *Markets and Hierarchies: Analysis and Antitrust Implications*. New York: Free Press.
- Wysocki, B., Jr. 9 April 2004. "Adopting Toyota Techniques Can Cut Costs, Wait Times." *Wall Street Journal*.
- Yeager, R. 2003. Quality Control Manager at Harley-Davidson, Inc. Personal communication.

Index

- Action modes, 144
- Active processes, 121
- Activity logging
 - data recording, 66
 - example, 67–71
 - format creation, 65
 - practice, 67
 - process explanation, 65
- American Productivity and Quality Center (APQC), 9
- Animal level system, 142
- Annotation, 25–28
- Asynchronous flows, 155
- Avionics components (ACs), 84–85, 167–168
- Avionics repair cycle, 190, 194
- Bar chart, 77
- Basic mapping conventions, 22–28
- Basic units, 124
- Bottom-line performance, 50
- Boulding hierarchy, 142
- Boundary-spanning action, 144
- Branching process flow, 23
- Bread making, 28–32
- Business metrics
 - best-performing firms, 52
 - efficiency, 50–51
 - management practices, 51–52
 - milestone measures, 53
 - nonfinancial measures and factors, 52
 - performance indices, 50
 - time lag, 52–53
- Business model, 3
- Business Process Management Institute, 6
- Business process mapping (BPM), 104–105
- Business process maps, 5
- Business process reengineering (BPR)
 - benefits, 110
 - downsizing, 109–110
 - information technology, 110–111
- Cash management, 38
- Catalog processes, 7
- Central tendency, 78, 79
- Certification document, 57
- Client process map, 42, 43
- Clockworks, 142
- Complex systems, 147
- Complex workflow, knowledge management
 - avionics workflow
 - avionics components, 167–168
 - build-up-teardown cycles, 171–172
 - larger-scale logistics support mission, 168–169
 - maintenance cycle, 169–170
 - large-scale simplicity
 - dynamic homeostasis, 173
 - feedbacks and feedforwards, 173–174
 - information-processing capacity, 172–173
 - internal activity, 173
 - stable organization, 173
 - workflow map validation, 175
- Components, 81–83
- Configuration control system, 98–99
- Continuous variables, 78
- Convergent workflow, 154
- Conversion processes, 152
- Coordination, 59, 138, 149
- Corporate retirement account, 41
- Cost/resource consumption, 62
- Coupling, 154
- Customer cash-transfer order, 68

- Cybernetic systems, 142
- Cycle time, 66
- Data collection and analysis
 - accumulated time, 83–86
 - business metrics and process workflows
 - best-performing firms, 52
 - efficiency, 50–51
 - management practices, 51–52
 - milestone measures, 53
 - nonfinancial measures and factors, 52
 - performance indices, 50
 - time lag, 52–53
 - central tendency measures, 78
 - experimentation and implementation, 72–74
 - frequency distribution, 77
 - histogram, 78–79
 - measure of variability, 80
 - performance improvement, 81–83
 - spreadsheet *versus* manual construction, 76–77
 - WFMA approach
 - activity logging, 64–71
 - certification documentation, 56–57
 - ennui and extremes, 58–60
 - map validation, 55
 - normal *versus* exceptional tasks, 57–58
 - objective-based approach, 54
 - planning, 54
 - process improvement
 - opportunities, 60–64
 - whole job/process-relevant tasks, 58
 - workflow design tradeoffs, 86–88
 - workflow maps, 74–76
 - workflow metric, 71–72
- Data flow diagram (DFD) format, 39–40
- Decision diamond, 18–20
- Deming cycle, 72
- Descriptive statistics, 75, 88
- Designed process, 151
- Divergent workflow, 154, 156–158
- DMAIC (Define, Measure, Analyze, Improve, Control), 49
- Drawing software, 4
- Drill-down, 25, 28, 30
- Dynamic homeostasis, 173
- Elapsed maintenance time (EMT), 84
- Enterprise Resource Planning (ERP) system, 136
- Environment, 129, 142, 143
- Exception handling process, 38, 100
- Extremes, 58–60
- Feedback, 148, 173–174
- Feedforward, 148–149, 173–174
- First-In, First-Out (FIFO), 168
- Flowcharting, 108–109
- Flow diagrams, 5
- Flow of information, 149, 150
- Flow relationships, 38–39, 154–159
- Formal knowledge, 137–138
- Framework
 - imperfect information, 121–123
 - information, processes, and performance, 132–135
 - information processors, 123–127
 - knowledge, 117–120
 - systems, organizations are, 128–131
- Frequency distribution, 75, 77, 80
- Functional workflow, 155
- Fund trade, processing, 40
- Fuzzy set, 121
- Garbage In, Garbage Out, 60
- General Systems Theory (GST), 128, 142
- Genetic-societal level system, 142
- Goal attainment, 132
- Golden handcuff, 63
- Graphic software, 56
- Graphing and modeling workflows (XPDL), 5
- Hierarchy, 124
- Histogram, 75–79
- Homegrown process, 180
- Human level system, 142
- Hyperlinks, 27
- Imperfect information, 121–123
- Individual error (INDERR), 83

- Individual variation (INDVAR), 82
- Information and Communication Technology (ICT), 62
- Information and knowledge
 - exhaustive model of states of, 117
 - human information processing, 120
 - imperfection
 - active process, 121
 - passive forms of, 121–122
 - tacit knowledge, 123
 - known knowns, 117, 119
 - known unknowns, 117
 - performance/goal attainment, 132–135
 - processors
 - functional consistency lag and cost, 125–126
 - information processing activity, 123
 - international spying and intelligence, 127
 - role of managers, 127
 - tacit knowledge, 128
 - task-related information, 125
 - three-level hierarchy, 124
 - trade-offs, 124
 - uncertainty absorption, 126–127
 - William Gore's experience and preference, 124–125
 - properties of, 118, 119
 - tacit knowledge, 120, 137–138
 - unknown knowns, 117
 - unknown unknown, 117–118
 - vector properties, 118–119
- Information dependency, 148
- Information imperfection, 151
- Information processing
 - improvements, 60
- Inherent information, 152
- Inputs, 14644
- Intellectual property, 122, 130
- Internal processes, 131
- Internal transformation, 146, 147
- IT-dependent mapping methods, 10

- Just In Time (JIT) management, 130, 134

- Keep It Simple, Stupid (KISS), 16–17
- Kmetz method, 4

- Knowledge management. *See also* Information and knowledge
 - academic conference, 162–163
 - binary code, 162
 - complex workflow (*see* Complex workflow, knowledge management)
 - information imperfections, 165
 - information technology, 163
 - IT-intensive KMS, 165–166
 - Knowledge Management Research Report, 165
 - limitations, 166
 - MBNA America Bank, 185–186
 - organizational learning, 164–165
 - Balanced Scorecard, 184
 - fundamental changes, 184
 - mechanistic model, 182–183
 - new knowledge generation, 184
 - social interaction, 184
 - trade-off, 183
 - zero-based activity, 183–184
 - positive and negative impacts
 - IT-intensive KMS designs, 180–181
 - NAMP, 180
 - payoff, 178–179
 - research and development (R&D) work, 176–178
 - tacit knowledge, 179
- Known knowns (KK), 117, 119
- Known unknowns (KU), 117

- Lean management methods, 51
- Least Practical Value (LPV), 81

- Management practice, 51
- Manual construction, 76–77
- Mapping parallel processes, 32–36
- Mapping symbols
 - Darwinian selection proce, 13
 - flowcharting, 12
 - NAVAIR studies, 11
- Market, 147
- Mean, 78
- Mean Exception Handling Time (MEHT), 71
- Mean Primary Process Time (MPPT), 71

- Mean Process Execution Time (MPET), 71
- Measurement inflation, 66
- Mechanistic organizations, 183
- Median, 78
- Mode, 78
- Multiple-cycle (looping) process flow, 24

- Naval Air Systems Command (NAVAIR) studies, 11
- Naval Aviation Maintenance Plan (NAMP), 180
- Nested relationship, 149
- Nonfinancial measures, 52
- Nonspecialist methods, 10
- Nuts and bolts, 44, 49

- Observer error (OBSERR), 82, 83
- Open systems, 142
- Operating system, 130
- Organizational environment, 130
- Organizational learning, 138, 164–165
 - Balanced Scorecard, 184
 - fundamental changes, 184
 - mechanistic model, 182–183
 - new knowledge generation, 184
 - social interaction, 184
 - trade-off, 183
 - zero-based activity, 183–184
- Organizational process
 - acquisition process, 146
 - conversion process, 146
 - dynamic, 147–148
 - environment, 147
 - feedforward, 148–149
 - information dependency, 148
 - information processors
 - functional consistency lag and cost, 125–126
 - information processing activity, 123
 - international spying and intelligence, 127
 - role of managers, 127
 - tacit knowledge, 128
 - task-related information, 125
 - three-level hierarchy, 124
 - trade-offs, 124
 - uncertainty absorption, 126–127
 - William Gore's experience and preference, 124–125
 - known and unknown
 - consequences, 150
 - regulatory information, 149–150
 - regulatory information flow, 149
 - systems
 - environment, 129–130
 - General Systems Theory, 128
 - internal processes, 131
 - self-regulation and adaptability, 130–131
 - transformation process, 146–147
- Outputs, 150

- Payoff value, 25, 134
- Perfect information, 133
- Performance, 132–135
- Performance management, 51
- Physical flow improvements, 60
- Proactive offensive behaviors, 136
- Proactive steps, 134
- Process
 - block, 13, 14
 - building blocks, 144
 - classification framework (PCF), 9
 - improvement, 11
 - locus of action, 145
 - locus of information, 145
 - organization (*see* Organizational process)
 - owners
 - benefits, 97–98
 - designating, 95
 - knowledge management, 97
 - larger organization, 95
 - role and responsibilities, 96–97
 - small organization, 95
- Process Classification Framework (PCF), 9

- Quality certification audits, 56
- Quality improvements, 60
- Quasi-materials flow, 152

- Radio Frequency Identification (RFID) order processing, 105

- Range, 77
- Raw materials, 152
- Reactive defensive behaviors, 136
- Reactive steps, 134
- Regulatory information
 - feedback, 148
 - feedforward, 148–149
- Regulatory information mode, 145, 148
- Request for proposal (RFP), 34
- Return on capital employed (ROCE), 51
- Rough cut jobholder map, 100
- Second law of organizational thermodynamics, 116
- Self-maintaining structures, 142
- Sequential data sheet, 65
- Shewhart cycle, 72–73
- Simple data analysis methods, 74
- Simple dynamic structures, 142
- Simultaneous process, 35
- Single-cycle process flow, 22–23
- Six Sigma programs, 4
- Social organization level system, 142
- Software tweaking, 62
- Specialist methods, 7
- Specialist units, 124
- Spreadsheet, 76–77
- Stable organization, 173
- Standard deviation, 80
- Static model process, 11
- Static structures, 142
- Statistical tools, 74
- Stock, 145
- Stocks of information, 131
- Subsystem, 128–129, 143
- Supersystem, 143
- Supply-chain management, 130
- Swimlanes, 105
- Symbol usage, 14–21
- Synchronous flows, 154, 155
- System boundary, 26
- System complexity
 - animal level, 142
 - cybernetic systems, 142
 - environment, 143
 - General Systems Theory, 142
 - genetic-societal level, 142
 - human level, 142
 - open systems/self-maintaining structures, 142
 - simple dynamic structures/clockworks, 142
 - social organization level, 142
 - static structures/frameworks, 142
 - subsystem, 143
 - supersystem, 143
 - transcendental level, 142
- System diagram, 10
- Tacit knowledge, 4, 93, 120, 137–138
- Talent management, 51–52
- Tally sheet, 65
- Time study, 83–86, 111–112
- Transcendental level system, 142
- Tunnel, 120
- Uncertainty absorption, 126–127
- Unforeseen costs, 133
- Unknown knowns (UK), 117
- Unknown unknowns (UU), 117–118
- Validate, 55
- Variability, 80
- Variance, 80
- Vector, 118
- Vector properties, 121
- Versatile avionics shop test (VAST), 84
 - avionics components, 168
 - build-up-teardown cycles, 171–172
 - maintenance cycle, 169–170
- Work cycle, 21
- Workflow mapping
 - annotation, 25–28
 - basic mapping conventions, 22–25
 - discipline, 14–21
 - examples, 18–19, 39–44
 - flow-charting symbols
 - arrow, 13–14
 - circle, 13
 - diamond, 13
 - document, 14
 - rectangle, 13
 - rules, 17–18

- selection of, 14
 - flow relationships, 38–39
 - making bread, 32–36
 - mapping parallel processes, 28–32
 - map validation, 21–22
 - static model, 11
 - workflow map
 - properties, 10–11
 - system diagram, 10
 - types of, 36–38
- Workflow mapping and analysis (WFMA)
 - data collection and analysis (*see* Data collection and analysis)
 - fundamentals
 - annotation, 25–28
 - basic mapping conventions, 22–25
 - discipline, 14–21
 - examples, 18–19, 39–44
 - flow-charting symbols, 13–14
 - flow relationships, 38–39
 - making bread, 32–36
 - mapping parallel processes, 28–32
 - map validation, 21–22
 - static model, 11
 - implementations
 - business process reengineering, 109–111
 - configuration control system, 98–99
 - map duplication, 106–108
 - map preparation, 99–105
 - modifications, 105–106
 - planning and preparation, 94
 - process mapper, 92–93
 - process owners, 95–98
 - tacit knowledge, 93
 - time study, 111–112
 - knowledge management (*see* Knowledge management)
 - organizational process (*see* Organizational process)
 - system complexity, 142–143
 - workflows (*see* Workflows)
- Workflow mapping consortium (WfMC), 5–6
- Workflow map preparation
 - actual map creation, 101
 - exception-handling processes, 100
 - jobholder, 99–100
 - map validation, 102–104
 - predecessor and successor activities, 104–105
 - skilled specialists, 100–101
 - symbol usage, 99
 - transparency and openness, 105
- Workflows
 - change process, 151–152
 - designed process, 151
 - dynamic properties, 153–154
 - flows of information, 152–153
 - flows of materials, 152
 - input and conversion processes, 152
 - material and information flows
 - convergent, 154
 - coupling, 154
 - divergent/asynchronous flow, 156–158
 - functional/dysfunctional outcomes, 155–156
 - synchronicity, 155
 - yarn ball theory of organization, 154–155
- Y2K bug, 62

OTHER TITLES IN OUR SUPPLY AND OPERATIONS MANAGEMENT COLLECTION

Johnny Rungtusanatham, The Ohio State University and
Joy M. Field, Boston College, Editors

- *Understanding the Dynamics of the Value Chain* by William Presutti and John Mawhinney
- *An Introduction to Supply Chain Management: A Global Supply Chain Support Perspective* by Edmund Prater and Kim Whitehead
- *Production Line Efficiency: A Comprehensive Guide for Managers, Second Edition* by Sabry Shaaban and Sarah Hudson
- *Sourcing to Support the Green Initiative* by Lisa Ellram and Wendy Tate
- *Designing Supply Chains for New Product Development* by Antonio Arreola-Risa and Barry Keys
- *Metric Dashboards for Operations and Supply Chain Excellence* by Jaideep Motwani and Rob Ptacek
- *Statistical Process Control for Managers* by Victor E. Sower
- *Supply Chain Information Technology, Second Edition* by David L. Olson
- *Supply Chain Risk Management: Tools for Analysis, Second Edition* by David L. Olson
- *International Operations: How Multiple International Environments Impact Productivity and Location Decisions* by Harm-Jan Steenhuis
- *Introduction to the Four Cs of Supply Chain Management: Chain Structure, Competition, Capacity and Coordination* by Ananth Iyer
- *Supply Chain Logistics and Applications* by Ananth Iyer
- *Business Models for E-Commerce* by Zinovy Radovitsky
- *Application Models for E-Commerce* by Zinovy Radovitsky
- *Designing and Implementing an E-Commerce System* by Zinovy Radovitsky
- *Better Business Decisions Using Cost Modeling, Second Edition* by Victor Sower and Christopher Sower
- *Improving Business Performance with Lean, Second Edition* by James R. Bradley
- *Lean Communication: Applications for Continuous Process Improvement* by Sam Yankelevitch and Claire F. Kuhl
- *Leading and Managing Lean* by Gene Fliedner

Announcing the Business Expert Press Digital Library

Concise e-books business students need for classroom and research

This book can also be purchased in an e-book collection by your library as

- a one-time purchase,
- that is owned forever,
- allows for simultaneous readers,
- has no restrictions on printing, and
- can be downloaded as PDFs from within the library community.

Our digital library collections are a great solution to beat the rising cost of textbooks. E-books can be loaded into their course management systems or onto students' e-book readers. The **Business Expert Press** digital libraries are very affordable, with no obligation to buy in future years. For more information, please visit www.businessexpertpress.com/librarians. To set up a trial in the United States, please email sales@businessexpertpress.com.

THE BUSINESS EXPERT PRESS DIGITAL LIBRARIES

EBOOKS FOR BUSINESS STUDENTS

Curriculum-oriented, born-digital books for advanced business students, written by academic thought leaders who translate real-world business experience into course readings and reference materials for students expecting to tackle management and leadership challenges during their professional careers.

POLICIES BUILT BY LIBRARIANS

- *Unlimited simultaneous usage*
- *Unrestricted downloading and printing*
- *Perpetual access for a one-time fee*
- *No platform or maintenance fees*
- *Free MARC records*
- *No license to execute*

The Digital Libraries are a comprehensive, cost-effective way to deliver practical treatments of important business issues to every student and faculty member.

**For further information, a
free trial, or to order, contact:**

sales@businessexpertpress.com

www.businessexpertpress.com/librarians



BUSINESS EXPERT PRESS

Mapping Workflows and Managing Knowledge

*Using Formal and Tacit Knowledge to
Improve Organizational Performance,
Volume I*

John L. Kmetz

This is a book that does what the title says, and is different from most business process mapping information in three key ways. First, it lets users capture all the knowledge that goes into a workflow in any kind of organization, including the most difficult kind of all; the tacit knowledge people bring to the job and carry in their heads. Second, it is simple, powerful, flexible, and easy to learn. Third, it does not require installing, learning, and applying a complicated program (sometimes requiring reorganization to support the software rather than the software supporting the organization).

It was developed by the author in a 15-year long program of studying, analyzing, and improving avionics maintenance processes for the U.S. Navy and the Royal Canadian Air Force, and then applied to organizations of all kinds ever since, for more than two decades. It has been taught and applied by the author and others in many short courses. It works.

John L. Kmetz is president of Transition Assistance Associates, a selective firm of professionals who assist companies and organizations in diagnosing challenges and formulating transition strategies to meet them. He is retired from the Department of Business Administration at the University of Delaware, where he taught at both graduate and undergraduate levels for 34 years, in addition to taking a very active role in development and delivery of professional courses and seminars. He combined his academic interests with a wide range of experience. He spent two years as a proofreader and copyeditor in a printing company while working his way through a BS at Penn State. He took his MBA and Doctor of Business Administration at the University of Maryland.

SUPPLY AND OPERATIONS MANAGEMENT COLLECTION

M. Johnny Rungtusanatham and Joy M. Field,
Editors

