

Introduction to Systems Thinking

Daniel H. Kim

System. We hear and use the word all the time. “There’s no sense in trying to buck the system,” we might say. Or, “Mary, she’s a systems analyst.” Or, “This job’s getting out of control; I’ve got to establish a system.” Whether you are aware of it or not, you are a member of many systems—a family, a community, a church, a company. You yourself are a complex biological system comprising many smaller systems. And every day, you probably interact with dozens of systems, such as automobiles, ATM machines, retail stores, the organization you work for, etc. But what exactly is a system? How would we know one if we saw one, and why is it important to understand systems? Most important, how can we manage our organizations more effectively by understanding systems?

This volume explores these questions and introduces the principles and practice of a quietly growing field: systems thinking. With roots in disciplines as varied as biology, cybernetics, and ecology, systems thinking provides a way of looking at how the world works that differs markedly from the traditional reductionistic, analytic view. But this is not an either-or distinction we are making here. Because some problems are best solved through analytic thinking and others through a systemic perspective, we need both to better understand and manage the world around us.

Why is a systemic perspective an important complement to analytic thinking? One reason is that understanding how systems work—and how we play a role in them—lets us function more effectively and proactively within them. The more we understand systemic behavior, the more we can anticipate that behavior and work with systems (rather than being controlled by them) to shape the quality of our lives.

It’s been said that systems thinking is one of the key management competencies for the 21st century. As our world becomes ever more tightly interwoven globally and as the pace of change continues to increase, we will all need to become increasingly “system-wise.” This volume gives you the language and tools you need to start applying systems thinking principles and practices in your own organization.

Contents

What Is Systems Thinking?	2
What Is a System?	2
<i>Collections Versus Systems</i>	
<i>Defining Characteristics of Systems</i>	
<i>The Importance of Purpose</i>	
Putting Systems in Context: “The Iceberg”	4
What Do Systems Do? A Close Look at Systemic Behavior	5
<i>Fun with Feedback</i>	
<i>The Building Blocks of Systemic Behavior: Reinforcing and Balancing Processes</i>	
<i>Looking for a Sign: Loops and Labels</i>	
<i>The Good, the Bad, and the Ugly: A Closer Look at Balancing Loops</i>	
<i>Delays: The Hidden Troublemakers</i>	
Putting It All Together: Two Examples of How to Manage Systems	12
<i>Managing Product Quality at FitCo</i>	
<i>Fixes That Backfire at DevWare Corp.</i>	
Working <u>on</u> the System, Not <u>in</u> the System	16
Appendix: “Acting” in Different Modes	17
A Glossary of Systems Thinking Terms	19

Introduction to Systems Thinking
©1999 by Pegasus Communications, Inc.

All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or by any information storage or retrieval system, without written permission from the publisher. For additional copies contact:
permissions@pegasuscom.com



What Is Systems Thinking?

What exactly is systems thinking? In simplest terms, systems thinking is a way of seeing and talking about reality that helps us better understand and work with systems to influence the quality of our lives. In this sense, systems thinking can be seen as a perspective. It also involves a unique vocabulary for describing systemic behavior, and so can be thought of as a language as well. And, because it offers a range of techniques and devices for visually capturing and communicating about systems, it is a set of tools.

For anyone who is new to systems thinking, the best way to “get your feet wet” is to first learn about the defining characteristics of systems; in short, what is a system? But to be a true systems thinker, you also need to know how systems fit into the larger context of day-to-day life, how they *behave*, and how to manage them. The final three sections of this volume tackle those issues.



What Is a System?

In the most basic sense, a system is any group of interacting, interrelated, or interdependent parts that form a complex and unified whole that has a specific purpose. The key thing to remember is that all the parts are interrelated and interdependent in some way. Without such interdependencies, we have just a collection of parts, not a system.

Collections Versus Systems

Let's illustrate this point with the following exercise. Take a look at the list of items below and determine for yourself

which ones are systems and which ones are just collections of parts. Ready, set, go!

- Bowl of fruit
- Football team
- Toaster
- Kitchen
- Database of customer names
- Tools in a toolbox
- A marriage

So, which ones are systems and which are merely collections? This question isn't as easy to answer as it might seem at first. Your responses depend on what assumptions you are making about the item in question. Let's walk through each example (starting with the simpler ones first) and make our assumptions as explicit as we can.

Kitchen, database of customer names, and tools in a toolbox. These are all collections, because none of them meets our original criteria of interrelatedness and interdependence. Even though the kitchen itself is full of systems (refrigerator, microwave, dishwasher), it is still just a place that has a collection of systems and other elements in it. None of those things interrelate or interact in an interdependent way. (Note, though, that once humans enter a kitchen, they, together with the other elements, form a system. It's a curious fact, but *whenever you add people to a collection, you almost always transform a collection into a system!*)

Football team and toaster. Both are systems. Notice that in addition to our criteria of interrelatedness and interdependence, a team and a toaster are each put together for a specific purpose. Indeed, purpose acts as the predomi-

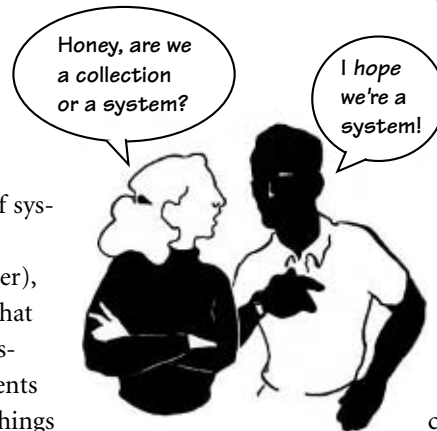
nant organizing force in any system. If you want to understand why a system is organized in a particular way, find out the system's purpose.

Bowl of fruit. Most people would classify this as an obvious collection, because the pieces of fruit are not interrelated in any way and do not interact with each other. In truth, however, they *are* interacting—at a microscopic level. For instance, if you put certain fruits together, they are apt to decay faster because they interact at a molecular level. Someone for whom these interactions are important (a fruitologist?) might even consider this bowl of fruit to be a very interesting system—one whose purpose is to maximize the biodegrading process.

Marriage. For any of you who saw this one as a collection, please seek marriage counseling immediately! All kidding aside, the question of whether one has a healthy marriage has a lot to do with whether the relationship more resembles a collection or a system. Marriage is essentially a voluntarily chosen state of interdependence with another

person (not codependence, which is something altogether different). This state actually characterizes any long-term relationship, including friendships. Is there anybody among us who has not been reminded by someone that our actions have an impact on him or her? Sometimes, that is how we first encounter systems, and how we learn (often painfully) that we are part of a larger system than we may have realized.

Well, that was quite an excursion. I hope this tour has revealed that systems



are indeed all around us and that they take many different forms. In spite of these differences, though, all systems share several defining characteristics. It may be helpful at this point to summarize those characteristics.

Defining Characteristics of Systems

Systems have purpose. As we saw in the examples above, every system has some purpose that defines it as a discrete entity and that provides a kind of integrity that holds it together. The purpose, however, is a property of the system as a whole and not of any of the parts. For example, the purpose of an automobile is to provide a means to take people and things from one place to another. This purpose is a property of the automobile as a whole and cannot be detected in just the wheels, the engine, or any other part.

All parts must be present for a system to carry out its purpose optimally. If you can take pieces away from something without affecting its functioning, then you have a collection of parts, not a system. In the toolbox example, if you remove a wrench, you have fewer tools, but you have not changed the *nature* of what is in the box. Likewise, if you can add pieces to a collection without affecting its functioning, it's still just a collection.

The order in which the parts are arranged affects the performance of a system. If the components of a collection can be combined in any random order, then they do not make up a system. In our toolbox, it doesn't matter whether the screwdrivers are piled on top or buried at the bottom of the box (unless, of course, you really need a screwdriver now!). In a system, however, the arrangement of all the parts matters a great deal. (Imagine trying to

randomly rearrange the parts in your automobile!)

Systems attempt to maintain stability through feedback. In simplest terms, feedback is the transmission and return of information. The most important feature of feedback is that it provides information to the system that lets it know how it is doing relative to some desired state. For example, the normal human body temperature is 98.6 degrees Fahrenheit. If you go for a run, the exertion warms your body beyond that desired temperature. This change activates your sweat glands until the cooling effect of the perspiration readjusts your temperature back to the norm. Or, in our car example, imagine that you are steering your car into a curve. If you turn too sharply, you receive feedback in the form of visual cues and internal sensations that you are turning too much for the speed at which you're traveling. You then make adjustments to correct the degree of your turn or alter your speed, or some combination of both. If you are a passenger in a car driven by someone who is not paying attention to such feedback, you might be better off getting a ride with someone else!

The Importance of Purpose

We talked about systemic purpose a bit, but let's take a closer look at it. A key to understanding any system is knowing its purpose, either as a separate entity or in relation to a larger system of which it is a part. In human-made (or

mechanical) systems, the *intended* purpose is usually explicit and reasonably clear, at least at the outset. The purpose of a washing machine, for example, is to wash clothes. The washing system is designed so that all the components work together to accomplish that purpose as effectively as possible.¹ In mechanical systems, the purpose is usually "hard-wired" into the design and therefore does not evolve over time. Your car, for example, was designed to take you places and will continue to operate with that purpose (provided you do your part in taking regular care of it). You'll never encounter a situation where you wake up one morning and your car has changed its purpose to be a lawnmower (though it may turn into a big, heavy, unmovable paperweight!).

Living (or natural) systems, on the other hand, are continually evolving and have the capacity to change their purpose, temporarily or permanently. For example, one of the most basic assumptions people make about animals is that they are driven only by survival instincts and the need to pass on their genes. As we deepen our understanding of nature, however, scientists are discovering that many animals seem to have much more complex set of purposes—some of them quite social—that govern their behavior. (Of course, we humans take it for granted that we have higher purposes beyond survival.)

Natural and social systems can be far more difficult to understand than nonliving systems, because we can never know for sure what their purpose

¹ Beware: Customers who buy these systems may use them for other purposes that fit their own needs. In such situations, where a system is used for a purpose different from the one for which it was originally designed, the system is likely to degrade or fail. An unexpected use of washing machines actually occurred in Japan, where farmers employed the machines to wash their potatoes—and then complained to the manufacturer about the frequent breakdowns! The company had the option of trying to redesign the machine to accomplish both purposes effectively or to persuade the farmers not to wash their potatoes in them. In this case, the company chose to change the design and tout the machine's ruggedness as an extra feature.

or design is. As a result of this inability to truly know their purpose and design, we tend to take actions in these systems without really understanding the impact of our actions on the system. Whenever we do this, we risk causing a breakdown of the system. For example, people smoked tobacco for years before it was discovered that one of smoking's long-term consequences is lung cancer. Even though we had a fairly good understanding of the purpose of our lungs, we did not have a sufficient understanding of how the lungs worked and what impact smoking would have on them—and us—over a long period of time. Since we aren't the designers of the human body, we have to learn about how it works as a system largely by trial and error. Similarly, farmers have had to learn about ecological systems in much the same way, and managers struggle with organizational behavior for the same reasons. Like the human body, nature and human social systems don't come with an owner's manual.

Despite our ignorance about natural and social systems, we still can't seem to resist attributing some purpose to them. We even tend to impose a purpose on natural systems and then behave toward them in a way that is consistent with that purpose. For example, in some countries, people view dogs as pets for families to enjoy. In such regions, people might treat dogs almost as members of the family. In other parts of the world, dogs are seen as a source of food, and people treat them accordingly. In both situations, the practices toward dogs are consistent with each different, perceived purpose. Neither viewpoint is intrinsically right or wrong, although each may seem wrong when viewed through the "lens" of the other.

Clearly, there are lots of systems to choose from if you want to study sys-

temic behavior. But as we will see, social systems make up the most complex class of systems—which you probably already know from direct experience in trying to manage some of them!



Putting Systems in Context: "The Iceberg"

Before we dive more deeply into the world of systems, it's helpful to see how systems fit into a broader context. We can actually view reality from the following multiple levels of perspective: *events*, *patterns*, and *systemic structures* (see "The Iceberg"). As we'll see below, systems occupy a key position in this framework. But what do these levels mean? Some basic definitions and a few examples might help:

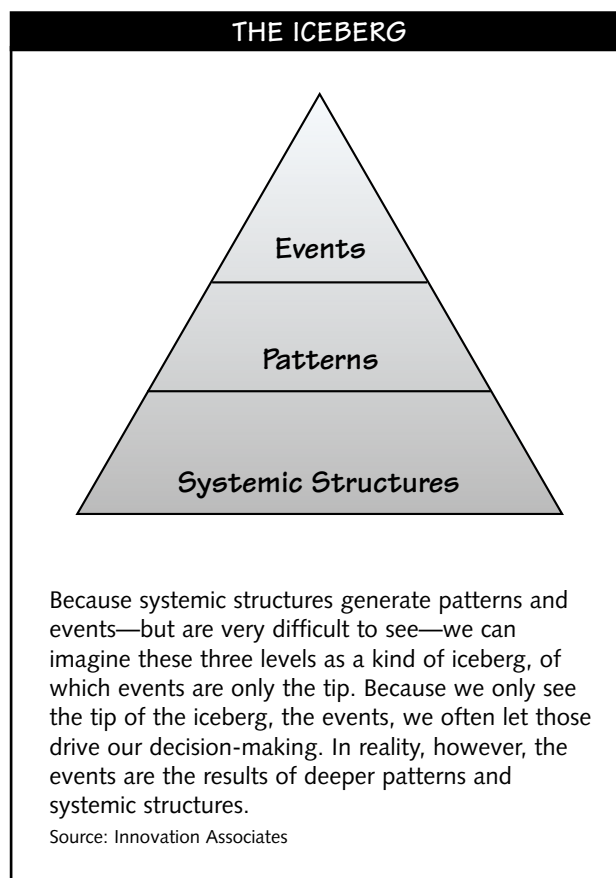
Events are the occurrences we encounter on a day-to-day basis. For example, we catch a cold, a fire breaks out, or a defective product comes off the assembly line at our company.

Patterns are the accumulated "memories" of events. When strung together as a series over time, they can reveal recurring trends. For example, we catch colds more often when we're tired, fires break out more frequently in certain neighborhoods, or we notice a higher number of product defects during shift changes.

Systemic structures are the ways in which the parts of a system are organized. These structures actually generate

the patterns and events we observe. In the example above about defective products, perhaps shifts are scheduled such that there is no overlap between the outgoing and incoming work crews—hence, there's a greater likelihood of defects during these times. Note that systemic structures can be physical (such as the way a workspace is organized, or the way a machine is built) as well as intangible (such as ways employees are rewarded, or the way shift changes are timed).

A key thing to notice about the three different levels of perspective is that we live in an event-oriented world, and our language is rooted at the level of events. Indeed, we usually notice events much more easily than we notice patterns and systemic structures even though it is systems that are actually



driving the events we do see. This tendency to only see events is consistent with our evolutionary history, which was geared toward responding to anything that posed an immediate danger to our well-being. As we'll see later in this volume, it's redesigning things at the *systemic* level that offers us far more leverage to shape our future than simply reacting to events does.



What Do Systems Do? A Close Look at Systemic Behavior

We've explored what defines systems and how systems generate the patterns and events we see around us. But how do we actually start looking at reality from this intriguing viewpoint? We need to do two things: deepen our understanding of how systems behave, and gain familiarity with some terms and tools of systems


thinking in order to communicate our understanding of that behavior. This section “walks” you through some basic systems thinking tools—causal loop diagrams and behavior over time graphs—to illustrate the concepts.

Fun with Feedback

To hone our systems thinking perspective, let's look again at feedback. As we saw earlier, feedback is the transmission and return of information. The key word here is *return*—it is this very characteristic that makes the feedback perspective different from the more common perspective: the linear cause-and-effect way of viewing the world. The linear view sees the world as a series of unidirectional cause-and-effect relationships: A causes B causes C causes D, etc.

A → B → C → D →

The feedback loop perspective, on the other hand, sees the world as an interconnected set of circular relationships, where something affects something else and is in turn affected by it: A causes B causes C causes A, etc.



As trivial as this distinction between these two views may seem, it has profound implications for the way we see the world and for how we manage our daily lives. When we take the linear view, we tend to see the world as a series of events that flow one after the other. For example, if sales go down (event A), I take action by launching a promotions campaign (event B). I then see orders increase (event C), sales rise (event D), and backlogs increase (event E). Then I notice sales decreasing again (event F), to which I respond with another promotional campaign (event G) . . . and so on. Through the “lens” of this linear perspective, I see the world as a series of events that trigger other events. Even though events B and G are repeating events, I see them as separate and unrelated.

From a feedback *loop* perspective (see “Thinking in Loops” on p. 6), I would be continually asking myself “How do the *consequences* of my actions feed back to affect the system?” So, when I see sales go down (event A), I launch a promotions campaign (event B). I see orders increase (event C) and sales rise (change in event A). But I also notice that backlogs increase (event D) (another eventual effect of event B), which affects orders and sales (change in

MENTAL MODELS AND VISION: MORE LEVELS OF PERSPECTIVE

We can gain even richer insights into systems by adding two more levels of perspective to the events/patterns/structure model. The two additional levels are *mental models* and *vision*.

Mental models are the beliefs and assumptions we hold about how the world works. We can view these assumptions as “systemic structure generators,” because they provide the “blueprints” for those structures. In our example about defective parts, maybe the production-line folks believe that they are responsible only for what they produce, not what the shift *after* them produces. This mental model may have led the company to create a *structure* whereby there is no overlap of staff during shift changes.

Vision is our picture of what we want for our future. It is the guiding force that determines the mental models we hold as important as we pursue our goals. For example, perhaps the people on each assembly-line shift hold a vision of competition—of striving to produce higher-quality products than any other shift. This vision would drive the mental model that says that each line is responsible only for what it produces.

See the “Acting’ in Different Modes” appendix on p. 17 for how to incorporate mental models and vision into the events/patterns/structure framework and take high-leverage actions to address a problem.

events C and A), which leads me to repeat my original action (event B).

After looking at both the linear and feedback representations, you might be saying to yourself, “So what? I’m too busy to draw pretty pictures about my actions. My job is to produce results—so I have to take actions *now*. Describing what has happened in two different ways still doesn’t

change what *actually* happens, so why do the two perspectives matter?” But here’s a key insight in systems thinking: *How we describe our actions in the world affects the kinds of actions we take in the world.* So, let’s reexamine the linear and feedback perspectives. Notice how the feedback view draws your attention to the interrelationships among *all* the events, whereas in the

linear view, you are probably drawn to each cause-and-effect event *pair*. By becoming aware of all the interrelationships involved in a problem, you’re in a much better position to address the problem than if you only saw separate cause-and-effect pairs.

The point here isn’t to “wax philosophical” about the intrinsic merits of two perspectives, but to identify one that will help us understand the behavior of complex systems so that we can better

manage those systems. The main problem with the linear view is that although it may be a technically accurate way of describing *what* happened *when*, it provides very little insight into *how* things happened and *why*. The primary purpose of the feedback view, on the other hand, is to gain a better understanding of all the forces that are producing the behaviors we are experiencing.

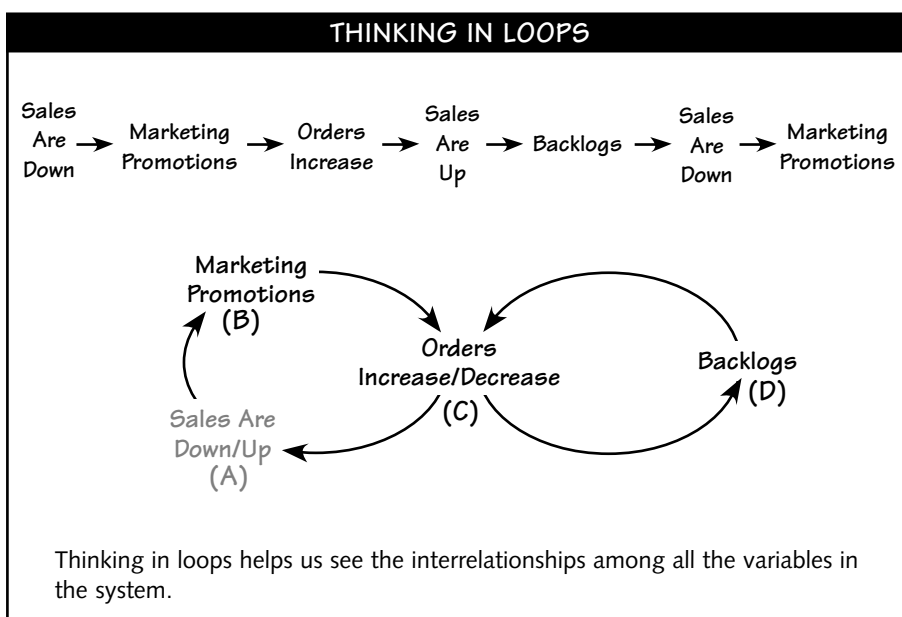


The Building Blocks of Systemic Behavior: Reinforcing and Balancing Processes


Feedback is just one piece of the picture when we’re thinking about how systems behave. To fill out the picture, let’s consider some examples of systemic behavior that we’ve all experienced. For instance, maybe you’ve worked in a company that was initially growing exponentially in sales, only to collapse a few years later. Or, maybe you’ve engaged in one of America’s favorite pastimes—dieting—where you kept losing the same 15 pounds over and over again. Or, you may recall that, when you were first learning to ride a bicycle, you wobbled down the street trying to stabilize yourself and eventually fell down (wondering what was wrong with *three* wheels anyhow).

All of these examples might seem completely unrelated on the surface; however, if we take a closer look at them, we can identify some very basic things that they have in common. In fact, all systemic behavior can be described through just two basic processes—called reinforcing and balancing processes. Both of these “building blocks” of systemic behavior involve distinctly different feedback. And, it’s the *combinations* of these processes that give rise to the vast variety of dynamic behavior in the systems we see all around us.


Reinforcing Processes: The Engines of Growth and Collapse. Reinforcing processes arise from what’s known as positive feedback. No, this isn’t praise for a job well done. In systems terminology, it means information that compounds change in one direction with even more change in that direction. In other words, successive changes *add* to the previous changes and keep the change going in the same direction.



Let's take a simple example of a savings account. If you have a positive balance, each time there is an interest payment calculation, the amount will be slightly bigger than the preceding payment period. This is because the balance has grown since the previous calculation. The time period after that, the interest amount will be bigger still, because the balance has grown a little more since the time before. Of course, all this is assuming that you are not making withdrawals during this time (which may be a big assumption for many of us!).



Another example is the wonderful growth engine that every marketer knows about: the word-of-mouth effect. As you increase the number of customers using your products, there are more “mouths” to tell others about your products. The resulting awareness leads



BEHAVIOR OVER TIME GRAPHS

Throughout the rest of this volume, you'll notice a few diagrams that look like this:

These are called behavior over time graphs. They're valuable because they show how certain variables that may be of interest to us—such as our savings balance, the number of customers we have, or our weight—are changing over time. They also provide clues to the kind of systemic processes that may be at work. A rapidly rising or falling graph, for example, indicates a reinforcing process, whereas an oscillating graph suggests what's called a balancing process.

to more sales, which leads to even more happy customers telling others. (Of course, this scenario is based on the assumption that your customers have nice things to say about your product!)

In the bank-account and word-of-mouth scenarios, a reinforcing dynamic drives change in one direction with even more change in the same direction. You can detect this kind of loop at work simply by sensing exponential growth or collapse (such as the rapid spread of an exciting new idea, or a company that suddenly goes out of business).

You can also think of reinforcing processes as “virtuous circles” when they produce desirable behavior. You may have encountered virtuous circles when you heard people talking about coming down the learning curve (the compounded increase in rate of learning as we learn more) or increasing economies of scale (the higher the production volume gets, the lower our unit costs become).

When reinforcing processes produce behavior we do *not* want, they are called “vicious cycles.” Oftentimes, a virtuous loop can become a vicious loop when something kicks it in the opposite

direction. In our word-of-mouth (WOM) example, the loop can turn “mean” if what people have to say about our product is negative. The negative WOM effect leads to lower sales, fewer customers, less WOM effect, even lower sales, etc.


These reinforcing processes are already embedded in our everyday language, which speaks to their pervasive presence. You've probably heard or used expressions such as “we were caught in a death spiral” or “things just kept snowballing.” Mapping such processes explicitly onto feedback loop diagrams (or causal loop diagrams, as they are called in the systems thinking field) lets us see and talk about them collectively so that we can respond more effectively to them.

Balancing Processes: The Great Stabilizers. We know there must be more to systems than just reinforcing loops, because our experience tells us that nothing grows forever (well, okay, except for taxes). We need something else to describe other kinds of behavior that do not look like continual exponential growth or decline. When we look around us, we see a great deal of stability, despite all the talk about the

era of rapid change we are in. For example, despite the rising or falling fortunes of individual companies or industries, the world of commerce continues to thrive around the globe. The world does change, but it does so on a platform of great stability. What accounts for all this constancy? It is balancing loops, the other “building block” of systemic behavior.

Balancing loops are continually trying to keep a system at some desired level of performance, much as a thermostat regulates the temperature in your house. Whereas the snowballing effect of reinforcing loops destabilizes systems (that is, puts them out of equilibrium), balancing loops are generally stabilizing or goal seeking. They resist change in one direction by producing change in the opposite direction, which negates the previous effects. (This is why they are also called negative feedback loops.) For example, when the thermostat in your home detects that the room temperature is higher than the thermostat setting, it shuts down the heat.

There is always an inherent goal in a balancing process, and what “drives” a balancing loop is the gap between the goal (the desired level) and the actual level. As the discrepancy between the two levels widens, the system takes corrective actions to adjust the actual level until the discrepancy decreases. In the thermostat example, gaps between the actual room temperature and the temperature setting of the thermostat (the goal) prompt the thermostat to adjust the heating or cooling mechanisms in the house to bring the actual tempera-




ture closer to the desired temperature. In this sense, balancing processes always try to bring conditions into some state of equilibrium.

It would not be a gross exaggeration to say that balancing processes are everywhere. They are far more ubiquitous than reinforcing loops. However, they’re a lot less visible, because they quietly function to *keep things as they are*. We tend to notice things that have changed much more than things that remain the same. For example, think about the times when you are aware of your body temperature. Most likely, you notice it only when it has “grown” beyond your normal level in the form of a fever, or when it has fallen below normal owing to hypothermia. Similarly, when do you notice how your car engine is running? Most likely, only when it *quits* running. In both cases, there is a massive number of balancing processes at work to keep the system running smoothly. (Quick, which system—you or your car—has more loops? Hint: One is made by humans; the other by nature.)

Balancing loops show up in organizations most often in the form of control loops. The balancing “language” is everywhere you look: “damage control,” “inventory control,” etc. We could say that all managerial responsibilities can be viewed, in one way or another, as balancing processes. Just think: All you really need to do to be a great manager is understand how to manage balancing loops! Sound far-fetched? That is actually the great secret to becoming a good general manager—having the ability to rise above the distraction of the details and see the underlying systemic structures that are producing the results. Seeing the world through the lenses of reinforcing and balancing loops will help you develop these skills.

YOU TRY IT: REINFORCING PROCESSES


Now that you have a feeling for what reinforcing loops are like, try your hand at drawing a few of them. They could be from your personal life (falling in love, making an investment) or professional setting (launching a new product, learning a new skill). The main point is to depict a clear and compelling story of how things mutually reinforce change in one direction in a complete circle.



Looking for a Sign: Loops and Labels


Before we go on to take a closer look at how balancing loops play a key role in systems, let's take a moment to explore two helpful features of causal loop diagrams in general: arrow labels and loop labels. Throughout the rest of this volume, you'll notice that the arrows in loop diagrams are labeled with an "s" or "o." These labels show how one variable influences another: An "s" indicates that as one variable changes, the next variable changes in the *same* direction. (Or, the first variable adds to the second variable.) An "o" indicates that as one variable changes, the other changes in the *opposite* direction. (Or, the first variable subtracts from the second variable.)

For example, each arrow in our savings-account/interest-payment loop would be labeled with an "s," because as savings go up (or down), so do interest



payments. And as interest payments go up (or down), so do savings. Notice that the savings account diagram has an "R" in the middle, too. This means that the loop represents a reinforcing process.

Here's another example, this time of a balancing process. Let's say that whenever you get stressed out, you do some relaxation exercises, which brings your stress level down. In a diagram of this system, the arrow going from stress level to gap is labeled with an "s." (As your stress increases, so does the gap between your actual and acceptable level.) The arrow going from gap to use of relaxation exercises is also labeled



with an "s." (The bigger the gap, the more you try to relax.) But the arrow going from use of relaxation exercises back to stress level is labeled with an "o." (As your use of relaxation exercises increases, your stress decreases.) This diagram would have a "B" in the middle, to indicate that it represents a balancing process.²

One easy way to tell if you have a reinforcing or balancing loop is to count the number of "o's." If there are no "o's" or an *even number* of "o's," the loop is reinforcing. If there is an *odd number* of "o's," the loop is balancing.

2 In classic system dynamics, a plus sign (+) is used instead of an "s," and a minus sign (-) instead of an "o." A plus sign indicates positive feedback; a minus sign, negative feedback.

However, even though this method is convenient, you should still double-check your reasoning by "walking" around the loop and telling the story of what it is depicting.


With these handy labels in mind, let's take a closer look at balancing loops.

The Good, the Bad, and the Ugly: A Closer Look at Balancing Loops

So far, the concept of a balancing loop might sound simple: These processes generally work to keep things stable. But beware: Balancing processes are actually pretty complex in real life. In many cases, we can think of them as a complicated blend of the "good" (the desired goal), the "bad" (the actual situation that we don't want), and the "ugly" (our perception of the situation, which we hate to examine). This makes managing these loops a little tricky, because people often have many different perceptions of a situation—and these perceptions can strongly affect the situation itself.

YOU TRY IT: BALANCING LOOPS

Here's an opportunity to flex your general management skills by seeing your responsibilities in terms of balancing loops. Think of a business goal that is especially important to you. It may relate to employee development, sales, or quality. Try to identify these four critical variables: Actual Level, Desired Level (or Goal), Gap, and the Corrective Action you need to close the gap. You might find the template below helpful.



Let's take quality of a product or service as an example (see "On a Quest for Quality"). In our standard balancing loop structure, we have our desired quality level and the actual quality level. When our desired quality level increases, our internal quality gap also increases (note the "s" on the arrow).

Whenever the gap itself increases, we increase our improvement efforts (again, note the "s"). When improvement efforts increase, we expect actual quality to increase (another "s"). Finally, when actual quality increases, our quality gap decreases (note the "o"). Once the gap decreases, we spin around the loop again: Improvement efforts also diminish, which in turn brings down actual quality. Once more, the gap increases.

Still with me? Good! But hold on: Even in this relatively basic examination of quality, there are many other important variables at work. For example, we often do not operate on the basis of what the actual quality is, but on our *perceptions* of what that quality is. In

addition, our *customers'* desired quality may not necessarily be the same as our own desired quality. And, in turn, customers don't always act on what our *actual* quality

is, but rather on their *perceptions* of what that quality is. Each of these variables introduces a new gap to worry about.

For example, when Hewlett-Packard first entered the

portable personal computer business, they designed and built their units in line with their usual high standards of quality. Internally, they were proud of the fact that their computers were virtually indestructible, so rugged was their design. But this ruggedness came with a high price tag. As a company, HP had launched improvement efforts that were driven largely by their internal quality gap. Customers, on the other hand, wanted their computers to be "rugged enough"—but they also wanted them to be affordable. Hence, HP's computers did not sell very well. It took some time for HP to shift their attention from their internal gap to their customers' quality gap.

What's the best way of managing these "good," "bad," and "ugly" balancing loops? Well, if you accept the premise that you are better able to manage things that are visible (rather than

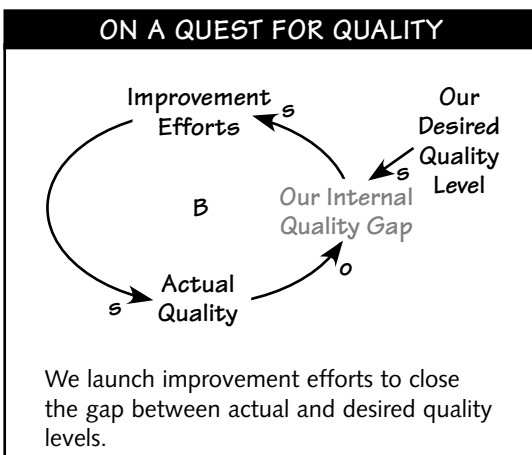

invisible) and relationships that are explicit (rather than implicit), then a good first step is to try mapping your issues onto causal loop diagrams (see "Why Draw a Diagram?" on p. 12). Through this kind of systems thinking approach, you begin to make more visible and explicit the causal structures driving organizational behavior. Creating such a diagram together as a team can be especially powerful, because it leads you to ask questions that you might not have asked before, such as,

- ? Which gaps are driving our system when, and by how much?
- ? How accurately do we know what each of the gaps is?
- ? How are we monitoring the gaps?
- ? What are the different ways in which we can close the gaps?
- ? How long does it take for perceptions to catch up to actual quality?

Addressing all these questions can uncover hidden assumptions and habitual practices that may be contributing to poor results.

Delays: The Hidden Troublemakers

Now that your head is spinning with all these loops, let's add one more layer of complexity. Another thing that makes understanding the behavior of complex systems so challenging is the existence of delays in the system. Every link in a system contains a delay. Sometimes delays are imperceptibly short (like the time between when the traffic signal turns green and when the person behind you honks his horn). At other times, they're interminably long (like waiting for a major marketing campaign to start generating sales).



Delays come in four basic “flavors”: physical, transactional, informational, and perceptual. *Physical* delays represent the amount of time it takes for actual “stuff” to move from one place to the other or to change from one state to another; for example, shipping products from the warehouse to retailers, or converting raw materials into useful products. Every transaction also takes time to complete, whether it’s a phone call or a series of contract negotiations—these can be called *transactional* delays. Then there are the delays associated with communicating information about the physical changes or decisions that have been made. Even with all our modern, high-speed communications systems, *informational* delays can still be quite long, because transmission does not necessarily equal communication. (That is, just because information was sent does not mean it was received and understood accurately.) The fourth

source of delay is often the trickiest—delays in *perception*. The physical changes have taken place (after a delay), decisions have been made, and the information about the change has been communicated. But, our beliefs and assumptions are often so deep that even if the reality on which they are based changes, our perceptions don’t neces-

sarily shift as easily. (It takes a long time to teach an old dog new tricks!)

These four kinds of delays are neither good nor bad; it’s how we handle them that determines whether they’ll cause trouble. In our rush to get things done quickly, we tend to underestimate the true delays in the system or even ignore them. But, delays are important to notice,

YOU TRY IT: DELAYS

Think of a process that you are responsible for managing—landing a new contract, for example. Now think through the whole process and identify the four different kinds of delays that may be involved—physical, transactional, informational, and perceptual. For each delay that you identify, estimate both the current as well as the theoretical minimum delay time. Now assess how your decision-making delay times compare with the other delays in the process. Where are your bottlenecks? You may discover that speeding some delays won’t help if you don’t shorten other delays first.

STOCKS AND FLOWS: ANOTHER SYSTEMS THINKING TOOL

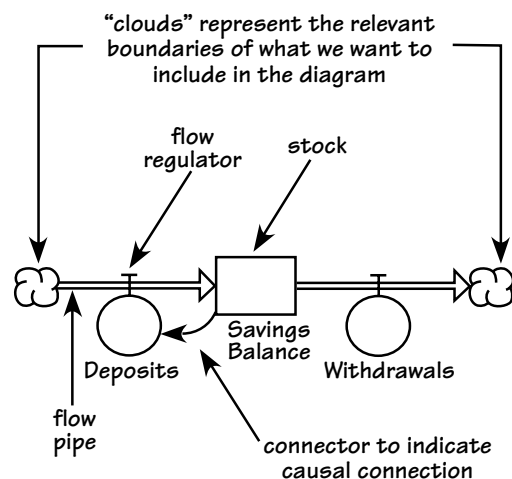
There’s another way besides causal loop diagrams to depict our understanding of systemic structure. It’s called a stock and flow diagram.

To create or read one of these diagrams, you first need to know what stocks and flows are. Stocks (also called accumulators) are anything that accumulate and that can be measured at one point in time, such as savings, population, the amount of water in a bathtub, and so on. Flows (or rates) represent things that change over time, such as deposits into a checking account, the inflation rate, etc.

Unlike causal loops, stock and flow diagrams provide information about *rates* of change. Combined with causal loops, they show how the various stocks and flows in the system influence one another and how the feedback flows through the system.

These diagrams are also used to build computer simulation models; the model builder assigns initial values to the stocks (such as “savings equals \$2,000 at time zero”) and rates for the flows (such as “\$20 savings per month”).

The diagram below identifies the various parts of a stock and flow diagram.



For more about stocks and flows, see *Systems Thinking Tools*, by Daniel H. Kim (Pegasus Communications, 1994).

because they can make a system's behavior unpredictable and confound our efforts to produce the results we want, as we will see in the next section.

Putting It All Together: Two Examples of How to Manage Systems

As we've said before, human systems are complex and challenging to manage. In addition, they tend to behave in counter-intuitive ways. (For example, we do something to fix a problem, but the problem just seems to get worse—and it just isn't clear *why*.) Understanding the different levels of perspective can help us figure out when it's time to design systems that will generate the kinds of events—and the kind of future—that we want. Using tools like causal loop diagrams can also be a powerful way to clarify our understanding of the systems we

want to work on, and to communicate that understanding with each other. Let's consider two focused examples of how to use systems thinking to grasp and manage a complex system.


Managing Product Quality at FitCo

We'll start by taking a peek at the inner workings of FitCo, a company that makes exercise equipment.³ FitCo is struggling with a problem that faces many organizations: managing product quality. We can think of this issue as a simple balancing process that comprises the interrelationships among three common variables: Product Quality, Customer Demand, and Production Pressure.⁴

The Simple Version. For FitCo (as with most other manufacturing firms), the higher the company's product qual-



ity, the more customers want to buy. But FitCo—thinking that the jump in demand might be temporary—doesn't do anything to beef up its production capacity once demand starts to rise. As a result, the folks in the production



department begin feeling enormous pressure to keep FitCo's expanding base of health-conscious customers supplied

WHY DRAW A DIAGRAM?

Systems thinkers work from a central premise: If you don't know how you're producing certain outcomes, you'll have great difficulty determining how to produce better outcomes! Sound obvious? Well, because social systems are so complex, this is not as easy to grasp as it seems. Still, we tend to operate our organizations as if we really knew what implications our actions will have. Worse, we often do so without the benefit of both a diagram that shows us the "wiring" of our system and the proper tools with which to conduct the operation successfully.


A diagram of the reinforcing and balancing processes at work in the system we're interested in can be an excellent first step to figuring out how the feedback is generating behavior that we want to change. And, it can help us address problems before breakdowns actually occur. Causal loop mapping is especially powerful when done in a group—because by sharing our understanding of how a system might work, we can get a fuller picture of reality and therefore arrive at much more powerful action plans.

So, whether it is our bodies, our cars, or our organizations, preventive maintenance is a worthwhile investment. There is a great deal of systemic truth in the old adage "an ounce of prevention is worth a pound of cure." (Likewise, "an ounce of systems diagrams are worth a pound of quick fixes"!)


3 The FitCo story, as well as the story about DevWare Corp. that begins on p. 14, are composite stories based on common experiences within many different companies. The company names are fictional.

4 You may have noticed that those variables don't include an explicit gap, unlike the earlier balancing loops you've seen in this volume. However, in any balancing loop, there's always an inherent gap—whether the gap is mentioned or shown in a diagram or not. In a diagram, not showing the gap is a shorthand way of drawing the loop. In the loop diagram on this page, there's an implicit gap between product quality (which represents the actual situation) and desired product quality (which represents the goal, and is not shown in the diagram).

with enough exercise machines. As the frantic production staff make more and more mistakes, and as the company's overused manufacturing machines break more and more frequently, the quality of FitCo's products starts to suffer—and customers begin drifting away. In this story (as we've traced it so far), customer demand and product quality rise and fall in close synchronization. If we were to graph the two variables, the resulting picture would look a lot like something called a steady-state equilibrium (you know, the kind that most economic texts presume is an accurate description of reality!).




Now Add Delay. You may have noticed that this version of FitCo's story is missing a key factor: delay. Because of delays, the situation at FitCo is much more likely to resemble a state of dynamic *disequilibrium*. Customer demand falls fast when FitCo's product quality falls, because people tend to react quickly to visible drops in quality—and because there are lots of other



exercise-machine companies out there for a disgruntled customer to choose from. However, the demand picks up more slowly when (and if) quality improves, because people become skept-

tical about quality improvements and want to wait and see if they're "for real."



The Investment Decision. There's yet another wrinkle to this picture. We know that, like many companies, FitCo doesn't keep its production capacity constant in the face of changing demand. Instead, it tries to adjust capacity so as to produce the right quantity of product at the desired level of quality. So, we have to add "investment in capacity" to our loop diagram (see "To Invest, or Not to Invest?"). If

FitCo is managing all the dynamics well, it should end up with both quality and demand rising ever upward. This is because, as customer demand increases, the company boosts capacity, which takes the heat off the production department and thereby improves product quality, further stimulating customer demand (see the R loop in the diagram).

The Death Spiral. Here's a key thing to realize about this quality-demand-pressure-investment structure: Depending on the impact of delay, this exact same structure can produce the "virtuous" or the "vicious" spiral shown in the "To Invest, or Not to Invest?" graph, in which product quality and customer demand are forced ever higher or ever lower, respectively. (That's the frustrating thing about systemic structures: They don't discrimi-

TO INVEST, OR NOT TO INVEST?

This causal loop diagram shows the bigger picture involved when we consider the impact of capacity investments on the quality-demand-pressure balancing structure we've been discussing. As product quality and customer demand increase, the company decides to invest in capacity. After a delay, the new capacity comes on line, which reduces the production pressure—which once more causes product quality to rise (note the "o" link). The decision to invest creates a reinforcing process. (To see how this works, trace the diagram from link to link; you'll count two "o's.")


As shown in the behavior over time graph, this structure can produce a vicious or a virtuous cycle of quality levels—depending on how skillfully we manage the dynamics.

nate between the two kinds of spirals! It's up to us to anticipate which kind of spiral might be lurking in our future—and manage the system in a way that keeps the “bad” kind at bay.)

To manage that vicious spiral, let's look at what kinds of conditions tend to put it in motion. The vicious spiral is more likely to happen when the delay between rising customer demand and increasing production capacity (the R loop in “To Invest, or Not to Invest?”) is significantly longer than the delay in product-quality changes and shifts in demand levels (see the B loop). Here's how this can happen to a company like FitCo—as well as to any manufacturing company:

1. As demand increases, FitCo holds off investing in additional capacity—perhaps because they've seen temporary blips in demand before, and they don't want to end up saddled with excess capacity.
2. Pressure on the production folks rises, and product quality begins to slip. Yet the drop in quality does not yet affect demand, so demand continues to rise.
3. When FitCo becomes convinced that the rise in demand is “for real,” it authorizes expansions in capacity.
4. New capacity takes a while to come online. If the delay in getting capacity online is significantly longer than the other delays, then the pressure on production will continue to mount, leading to even lower product quality and eventually lower customer demand.
5. When customer demand starts to drop, FitCo now tries to reverse its capacity additions. This prevents the company from getting the additional

capacity it needs. Pressures on production remain high, and product quality drops further. So, demand



- continues to fall. FitCo's managers applaud their supposed good judgment in cutting back on capacity, because (in their view) the customers were being fickle after all.
6. Convinced that they were right about the temporary nature of demand “blips,” FitCo's managers begin cutting capacity *ahead* of falling demand. Now they're thinking they're quite brilliant for saving the company so much money (even though they're totally blind to how their “wise” actions may be driving FitCo out of business).

The lesson here is that we can sometimes make decisions based on a belief about something that can actually *cause* the things we are trying to prevent. In FitCo's case, *beliefs* about falling demand can actually make the demand fall, in a tragic example of a self-fulfilling prophecy. When we're in the middle

of such a situation, however, it can look to us as if the fall is happening to us and that our actions are really a response to customer actions. Such is the nature of complex systems and the world of circular feedback loops: Once a loop gets going, it's hard to tell what is driving what.

As a structure, a reinforcing loop has no directional preference. So, how might FitCo ensure that it gets the loop to go in the desired direction (upward in this case)? Look again at the figure “To Invest, or Not to Invest” on page 13. One way to manage these loops is to realize the importance of the relative delays in the two loops. If the delay in the R loop is longer

than the one in the B loop, FitCo could try to figure out how to shorten the R-loop delay. For example, it could contract with other suppliers or partners who have excess capacity; that way, it could respond more quickly to upswings in demand. If that were not possible, then the company could try to create early-warning indicator systems that would alert it to unexpected jumps in production pressure or drops in quality. Both of these events are important signals that a company needs to expand its production capacity.

Fixes That Backfire at DevWare Corp.

A lot of managers expend energy trying to “fix” things. If sales are too low, we do something to get them higher. If yields are too low, we try to get the team responsible for yields to improve its performance. If profits are down, we cut costs to boost the bottom line. We may be quick to congratulate ourselves when conditions improve in the short term. But, in many cases, the problem

eventually returns to the same level as before—or gets even worse. We end up having the odd sensation that our supposed “fixes” are backfiring on us.⁵

To illustrate, let’s look at DevWare Corp., a hardware-development company. DevWare is facing an all-too-common situation, in which managers’ well-intentioned actions produce the exact opposite of what they wanted. One day, Toby, a manager in the company’s product-development program, notices that the number of parts behind schedule is alarmingly high. If this continues, he decides, the team won’t be able to launch the product on time. His conclusion: that the engineers need tighter supervision and a review of all parts in order to get the message that the number of parts behind schedule has to come down.

Sure enough, once Toby focuses his attention on the parts problem, the late parts start moving briskly through the pipeline. But after a while, the parts problem returns. And when Toby focuses on it once again, things improve again—but not as fast as before. Over time, the more attention Toby places on the problem, the worse the problem becomes. What’s going on?

Well, Toby’s attention to the late-parts problem came in the form of requiring more review meetings to check the status of parts—especially


parts that were running late (see “The Problem with Review Meetings”). All those meetings took time away from actual engineering work. So, rather than reporting problems with their parts as they arose, the engineers began waiting until they already had solutions to the problems. This meant that other engineers would find out about changes affecting their parts much later than they used to (see the R loop). As more and more engineers withheld information, more parts fell behind schedule—a situation that reinforced Toby’s belief that he needed to continue “helping” the engineers. The end result—a steadily worsening problem of late parts—was something nobody in the system wanted. Yet both Toby and the engineers were

unintentionally colluding to create this very situation.

A higher-leverage solution in this situation would be for Toby to take a very different kind of action than the review meetings he had been imposing. For example, if he had encouraged the timely reporting of problems—and promised not to “penalize” the engineers with more reviews or brow beatings—the engineers would have gladly reported problems sooner. Eventually, the number of late parts would have fallen dramatically. (However, this would have happened only after the problem got worse first. This “worse before better” outcome is a classic example of how complex systems behave. Once again, delays are the culprits in this dynamic.)

As you may have begun sensing in the FitCo and DevWare examples, everything really is connected to everything else. Yet no matter how narrowly we choose to define a system, that system ignores our arbitrary definition

Mapping the possible unintended as well as intended consequences of our actions in causal loop diagrams can help us anticipate and address problems before they arise.



5 This example depicts a systems archetype often referred to as “Fixes That Fail.” Systems archetypes are a set of eight classic “stories” of problems or behaviors that occur in many situations and across a broad range of organizations. To learn more about the archetypes, see the Suggested Further Reading list at the end of this volume.

and responds to all the relevant interconnections. As a result, there are many unintended consequences of our actions on a system, in addition to the intended consequences. Indeed, the issue is never whether our actions will have unintended consequences, but rather *to what degree* and *what kind of consequences* they will have. Mapping the possible unintended as well as intended consequences of our actions in causal loop diagrams can help us anticipate and address problems *before* they arise.



Working on the System, Not in the System

If I were to ask you who has the greatest impact on the safety and comfort of your flight on a commercial airline, what would you say? You might answer

that it is the pilots; after all, they're the ones who handle the takeoffs and landings and directly control how the plane operates under various circumstances. But then again, you might answer that it's the flight attendants, given that they have more contact with you during the flight. But if you really think about it, you may want to credit the designers of the aircraft, since they put the systems and structures of the airplane in place. (Now you know who to complain to about the carry-on luggage restrictions!) Whereas pilots and flight attendants work *in* the system, the aircraft designers work *on* the system—and therefore have the most influence on your experience of flying.

This idea—of working *on* the system as opposed to *in* the system—is a key lesson about systems thinking to take with you after reading this volume. How can we become better *designers* of

systems rather than mere *operators*? The concepts and tools introduced in this volume are a good start. We talked about what systems are, how they generate the patterns and events we see around us, and how they behave. We also talked about seeing the world in terms of interconnected reinforcing and balancing loops with delays. Finally, we saw how causal loop diagramming can be a powerful tool for depicting our understanding of systemic behavior, and for gaining insights into avenues for change.

All of these things can help us take the first steps toward becoming true systems thinkers. We hope that this introductory volume has given you a valuable foundation—one that will encourage you to try using these powerful perspectives and tools to begin shaping your future in a new way.

APPENDIX: "ACTING" IN DIFFERENT MODES

As we saw earlier in this volume, events are very compelling because they often require an instant response. For example, if a house is burning, we react by immediately trying to put out the fire. Putting out the fire is an appropriate action, but if that's all we ever did, it would be inadequate from a systemic perspective. Why? Because it has solved the immediate problem (the burning house), but it has done nothing to alter the fundamental structures that caused that event (inadequate building codes, lack of fire detectors, lack of fire-prevention education, etc.). Nor has it addressed the mental models and vision that have generated the problematic systemic structures.

The "Levels of Perspective" framework can help us go beyond responding only to events and begin looking for actions with higher leverage (see Level of Perspective and Action Mode figure). That is, we can begin to move from working *in* the system to working *on* it.

How does the "Levels of Perspective" framework help us take higher-leverage actions? Each level offers a distinctive mode of action. To illustrate, let's revisit our example of a manufacturing plant that is producing defective parts, and take a deeper look at how we would address the problem from each of the different perspective levels.

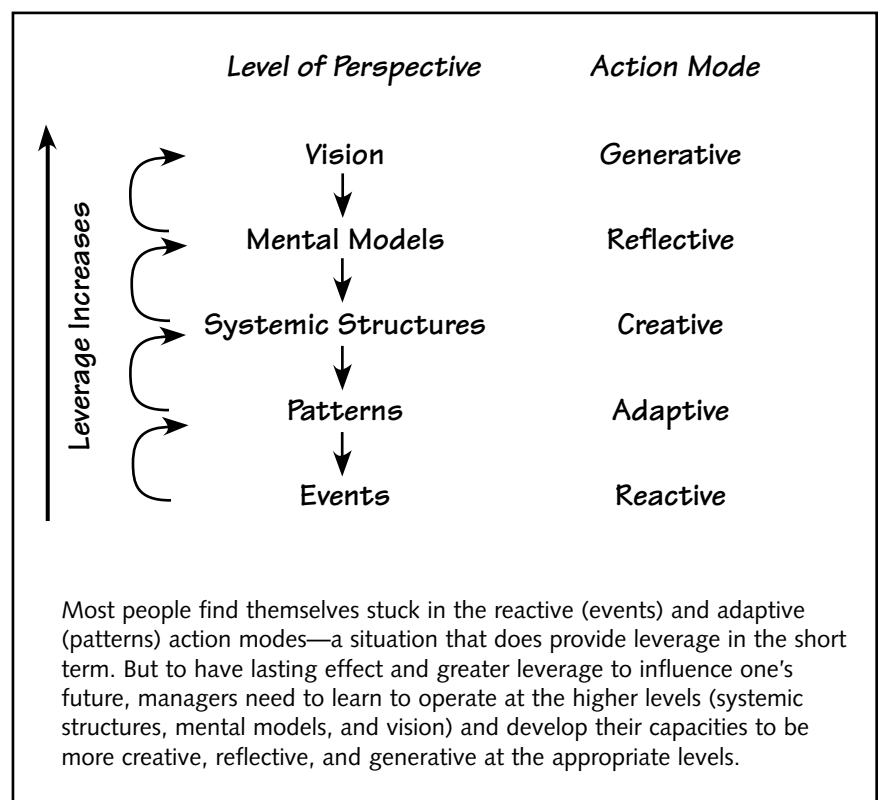
Events—Reactive. Whenever we encounter a defective part, we sort it out and either rework it or put it in the scrap pile. We may try to correct the situation by adjusting the machinery or by inspecting more closely, but our primary mode of action is *reactive*. Although we tend to view reactive actions in a negative light, they can still

be vital to our individual and organizational survival. However, they are not sufficient for sustaining long-term health.

Patterns—Adaptive. If we look at the problem over a period of time (for example, the rate at which we're scrapping parts), we may notice a pattern, such as higher scrap rates at certain times of the day. Specifically, we may notice higher scrap rates during shift changes. We can then *adapt* our processes to make the best use of the current system, perhaps in this case by simply accepting the fact that there's going to be higher scrap rates during shift changes. Notice that we are not trying to change the pattern; instead,

we're simply adapting to it. We can be intentional about these adaptive actions. However, they can also occur invisibly and unconsciously as we struggle to cope with the continued pressure of the situation. In other words, nobody proclaims that they are going to let quality erode by 20 percent, but that's exactly what may happen if we don't address the problem from a higher-leverage perspective.

Systemic Structures—Creative. As we saw earlier, systemic structures produce the patterns and events that make up our day-to-day reality. They are also the mechanisms through which mental models and vision get translated into action (look again at the sidebar on p. 5).



By *creating* new systemic structures (either through redesigning existing ones or making new ones), we can change the events and patterns we get. We alter the system, rather than just adapting or reacting to it. This is the level at which many change efforts operate (reorganizations, process redesign, reengineering, compensation schemes, etc.). In our defective-parts example, we might alter the system by creating an overlap of outgoing and incoming assembly-line crews.

Mental Models—Reflective.

Altering systemic structures often requires a change in our mental images of what those structures can or ought to be. In the example we've been following, if we truly believe that each assembly-line shift is responsible only for the quality of their products, then we'll never be able to imagine a different structure, such as overlapping crews who are each responsible for more than just their own lines. Taking actions at the level of mental models is *reflective*, because it requires that we develop the ability to surface, suspend, and question our own assumptions about how the

world works and what's most important. This skill also involves inviting others to do the same reflection with their mental models. (Note, though, that reflective actions do not include trying to *change* someone else's mental models—that would simply be another reactive action. Changes in our own and others' mental models come from genuine reflection and clarity of vision, not force.)

Vision—Generative. Surfacing, reflecting on, and changing our mental models is often a difficult and painful process, because those mental models are firmly embedded through years of experience. Why would we choose to put ourselves through the discomfort of changing them? Because we have a compelling vision of a new and different world that we are committed to creating. At the level of vision, our actions can be *generative*, bringing something into being that did not exist before. For example, a vision of providing the highest-quality products at all times through cooperation among assembly-line crews may generate the impetus to reexamine our old mental models that

say that each crew is responsible only for their own work.

Here's another important thing to notice about the levels of perspective: Our ability to influence the future increases as we move from the level of events to that of vision. As we move up these levels, our focus shifts from the present to the future. Consequently, the actions we take at the higher levels have more impact on future outcomes, not present events.

Does this mean that high-leverage actions can be found only at the higher levels? No—because leverage is a relative concept, not an absolute. For instance, if you find yourself in front of a runaway bus, that is probably not the best time to become very reflective about how you got yourself into that situation (because you won't be reflecting for very long!). In this case, the high-leverage action is to react fast and get out of the way; any other action would be inappropriate. There is leverage at every level, and the challenge lies in learning when and how to take the appropriate action for each level.

A GLOSSARY OF SYSTEMS THINKING TERMS

Systems thinking can serve as a language for communicating about complexity and interdependencies. To be fully conversant in any language, you must gain some mastery of the vocabulary, especially the phrases and idioms unique to that language. This glossary lists many terms that may come in handy when you're faced with a systems problem.

Accumulator Anything that builds up or dwindles; for example, water in a bathtub, savings in a bank account, inventory in a warehouse. In modeling software, a stock is often used as a generic symbol for accumulators. Also known as **Stock** or **Level**.

Balancing Process/Loop Combined with reinforcing loops, balancing processes form the building blocks of dynamic systems. Balancing processes seek equilibrium: They try to bring things to a desired state and keep them there. They also limit and constrain change generated by reinforcing processes. A balancing loop in a causal loop diagram depicts a balancing process.

Behavior Over Time (BOT) Graph

One of the 10 tools of systems thinking. BOT graphs capture the history or trend of one or more variables over time. By sketching several variables on one graph, you can gain an explicit understanding of how they interact over time. Also called **Reference Mode**.

Causal Loop Diagram (CLD) One of the 10 tools of systems thinking. Causal loop diagrams capture how variables in a system are interrelated. A CLD takes the form of one or more closed loops that depict cause-and-effect linkages.

Feedback The return of information about the status of a process. Example: annual performance reviews return information to an employee about the quality of his or her work.

Flow The amount of change something undergoes during a particular length of time. Example: the amount of water that flows out of a bathtub each minute, or the amount of interest earned in a savings account each month. Also called a **Rate**.

Level See **Accumulator**.

Leverage Point An area where small change can yield large improvements in a system.

Rate See **Flow**.

Reference Mode See **Behavior Over Time Graph**.

Reinforcing Process/Loop Along with balancing loops, reinforcing loops form the building blocks of dynamic systems. Reinforcing processes compound change in one direction with even more change in that same direction. As such, they generate both growth and collapse. A reinforcing loop in a causal loop diagram depicts a reinforcing process. Also known as vicious cycles or virtuous cycles.

Stock See **Accumulator**.

Structural Diagram Depicts the accumulators and flows in a system, giving an overview of the major structural elements that produce the system's behavior. Also called flow diagram or accumulator/flow diagram.

Structure The manner in which a system's elements are organized or interrelated. The structure of an organization, for example, could include not only the organizational chart but also incentive systems, information flows, and interpersonal interactions.

System A group of interacting, interrelated, or interdependent elements forming a complex whole. Almost always defined with respect to a specific purpose within a larger system. Example: An R&D department is a system that has a purpose in the context of the larger organization.

Systems Archetypes One of the 10 tools of systems thinking. Systems archetypes are the "classic stories" in systems thinking—common patterns and structures that occur repeatedly in different settings.

Systems Thinking A school of thought that focuses on recognizing the interconnections between the parts of a system and synthesizing them into a unified view of the whole.

The Innovations in Management Series

Pegasus Communications' *Innovations in Management* Series features in-depth analyses of both leading-edge and foundational topics in systems thinking, organizational learning, and management innovation. Concise and comprehensive, these volumes are ideal for customizing your learning—whether your interest is in the tools of systems thinking, the disciplines of organizational learning, or the latest management ideas as expressed by the most prominent thinkers in the business world. Watch for new titles, and build your *Innovations in Management* Series library!

For a complete listing of Pegasus resources, visit www.pegasuscom.com.

Daniel H. Kim is co-founder of Pegasus Communications, Inc., and co-founder of the MIT Center for Organizational Learning. He is also a trustee of the Society for Organizational Learning. Daniel is a leader in helping managers apply the power of systems thinking to tough organizational issues. As an international public speaker, teacher, and facilitator, he has worked with dozens of companies in developing their capabilities to become a learning organization. Daniel has a Ph.D. from the MIT Sloan School of Management and a B.S. in electrical engineering from the Massachusetts Institute of Technology.

Editorial support for this volume was provided by Kellie Wardman O'Reilly and Lauren Johnson. Many thanks also go to Janice Molloy, Kara Peters, Paul O'Reilly, and Scott Keller for providing helpful suggestions.

Suggested Further Reading

Learning Fables (available in soft cover or as e-books)

Outlearning the Wolves: Surviving and Thriving in a Learning Organization

Shadows of the Neanderthal: Illuminating the Beliefs That Limit Our Organizations

The Lemming Dilemma: Living with Purpose, Leading with Vision

The Tip of the Iceberg: Managing the Hidden Forces That Can Make or Break Your Organization

Systems Thinking for Kids

When a Butterfly Sneezes: A Guide for Helping Kids Explore Interconnections in Our World Through Favorite Stories

Billibonk & the Thorn Patch

Billibonk & the Big Itch

The Pegasus Workbook Series

Systems Thinking Basics: From Concepts to Causal Loops

Systems Archetype Basics: From Story to Structure

Volumes in the Innovations in Management Series

Introduction to Systems Thinking

Designing a Systems Thinking Intervention

From Mechanistic to Social Systemic Thinking: A Digest of a Talk by Russell L. Ackoff

Pocket Guides

Guidelines for Daily Systems Thinking Practice

The Do's and Don't's of Systems Thinking on the Job

Palette of Systems Thinking Tools

Guidelines for Drawing Causal Loop Diagrams

Other Titles by Pegasus Communications

Pegasus Anthologies

Organizational Learning at Work: Embracing the Challenges of the New Workplace

Making It Happen: Stories from Inside the New Workplace

Organizing for Learning: Strategies for Knowledge Creation and Enduring Change

The Innovations in Management Series

Concise, practical volumes on systems thinking and organizational learning tools, principles, and applications

Newsletter

The Systems Thinker®

Leverage Points™ for a New Workplace, New World is a free e-newsletter

spotlighting innovations in leadership, management, and organizational development. To subscribe, go to www.pegasuscom.com.

For a complete listing of Pegasus resources, visit www.pegasuscom.com.