3

Advances in Systems Thinking

Man is only a reed, the weakest thing in nature; but he is a thinking reed.

Blaise Pascal, 1623–1662

Scope, Limits and Values

Systems thinking is a developing discipline. It grew out of the systems approach, in which it was seen as necessary to view systems as open; to view them acting and interacting as wholes within their environment and context. Thus an open system, or whole, is observed, analyzed, understood, or posited, as an interactive part of something greater. The systems approach has brought considerable insight and benefit to those seeking enlightenment in nearly all fields of human endeavor. It owes some of its heritage, too, to operational research, which developed spectacularly during World War II.

Looking at an existing or putative system as an open, interactive part of some greater whole is not necessarily simple or straightforward. The system of interest (SOI) may interact contextually with its environment in various and complex ways, adapting as it does so. System thinking may be viewed as the process of envisioning these actions, interactions and adaptations. For some systems, and for some people, this is an exercise for the mind; systems thinking is cerebral. For others, the complexity may be too great to accommodate mentally, and there may be a need to resort to models. These may be of several types, including, but not limited to:

- Rich pictures — symbolic diagrams highlighting the essential features of some situation or problem space, looking particularly at paths of communication, lack of communication, areas of conflict, and so on. These are associated with so-called soft systems, but may be of much wider application.
- Causal or influence diagrams, usually formed from signed digraphs.
- Dynamic simulations of phenomena using one or other of the many systems dynamic simulation models. These encourage the exploration of the dynamic aspects of problem space, whereas the previous two bullets may describe the dynamics, but cannot simulate it — that must remain in the observer's imagination
- Dynamic models, not of phenomena directly, but of interacting open systems which exhibit properties of their interactions within the simulation

Systems Engineering: A 21st Century Systems Methodology Derek K. Hitchins
 ISBN: 978-0-470-0585-5

Rich pictures

Figure 3.1 shows a simple example of one kind of rich picture (there are various formats, some including stick figure of people and sketches of factories, vehicles, counters, etc.).

Simple though it is, the rich picture has a strong message. We try to predict/predetermine the effectiveness of our artificial systems, i.e., the effect one system has upon another. If the rich picture is correct, however, once our systems start to interact with each other, they will be changed in the process as they adapt to inflows, so that effectiveness, instead of being some fixed measure, is largely an indeterminate variable with time.

You think not? Well, if Red and Blue were World War II naval destroyers, pounding each other with heavy guns, then it seems reasonable to suppose that each will experience damage, each will attempt to repair the damage, each will find its capability to respond impaired, and effectiveness will, indeed, turn out to be time variant. An alternative version of this 'rich picture' might show a sketch with two ships, shells flying from each to the other, and perhaps a stick figure of a man on deck with a speech bubble saying: 'How can I undertake damage repairs when the other side is still shelling us?' 'Meanwhile, another stick figure, perhaps the opposing captain, might be saying: 'I don't like this — perhaps we should withdraw and fight again another day.' Conventional rich pictures of this type often highlight personal reactions as a way of expressing conflict, or lack of essential information.

Yet another way of presenting a rich picture is the ubiquitous causal loop model, as in Figure 3.2. In this example, the rich picture highlights the factors leading to the collapse of the Old Kingdom of ancient Egypt. The example shows an early instance of counterintuitive behavior. Successive kings granted lands to nobles, who set up hereditary 'mini-dynasties' of their own, not only depending less on the favor of the king, but even possibly starting to challenge his divine authority. A succession of poor Nile inundations resulted in famine, internecine war between the mini-dynasties, civil war, and the breakdown of the Old Kingdom. (In later Kingdoms, land would be granted to nobles, but

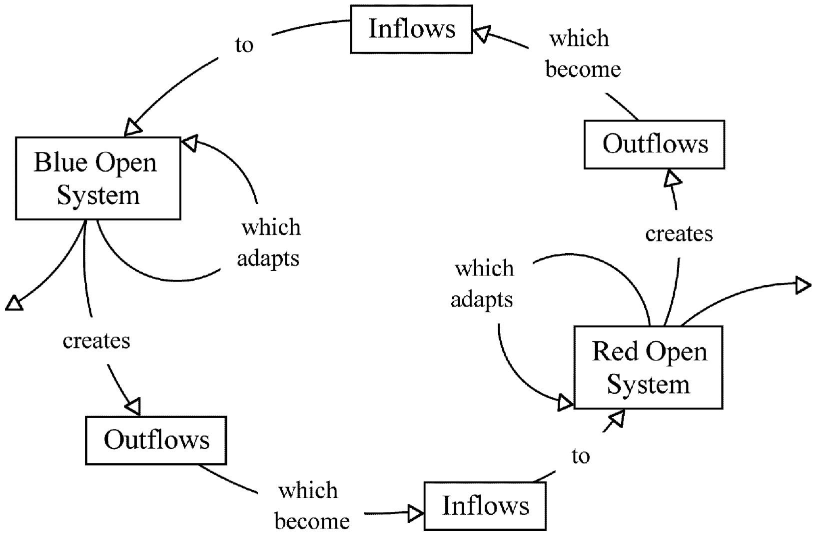

Figure 3.1 Simple rich picture showing how perceived effectiveness may vary with time.

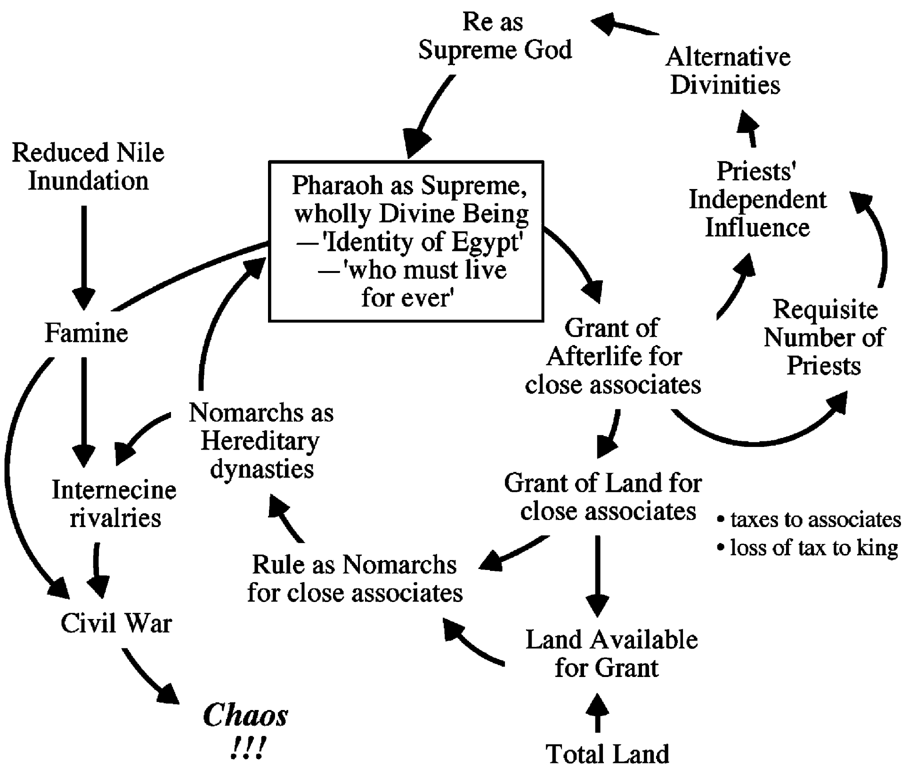

Figure 3.2 Rich picture in the form of a causal loop. The rich picture is a means of depicting the principal features, conflicts, etc., in a situation. In ancient Egypt, the king (pharaoh) reined supreme as a divine being. He granted land to nobles who became rulers of nomes (localities). These nomarchs established hereditary dynasties which fell into civil war when famine the land, leading to the collapse of the Old Kingdom.

only for the period of their life, so that they could not establish independent power bases up and down the country. . .).

Causality and causal loop models

A useful way to explore the behavior of systems is to trace and predict cause and effect. There are different ways of looking at causality — see Figure 3.3. One view tends to isolate one cause and one effect, so that there may be several causes and several consequent effects, but there is no association between them. This is a view seemingly taken by some politicians and government agents. For example, they might (intentionally?) see no relationship between raising the price of rail fares on commuter trains and the concomitant increase of road traffic congestion to and from the same city during morning and evening rush hours. . . .

A different view is often seen by managers and those seeking to impose control. This viewpoint sees a linear chain of causality, in which one cause creates an effect, which, in its turn, becomes a

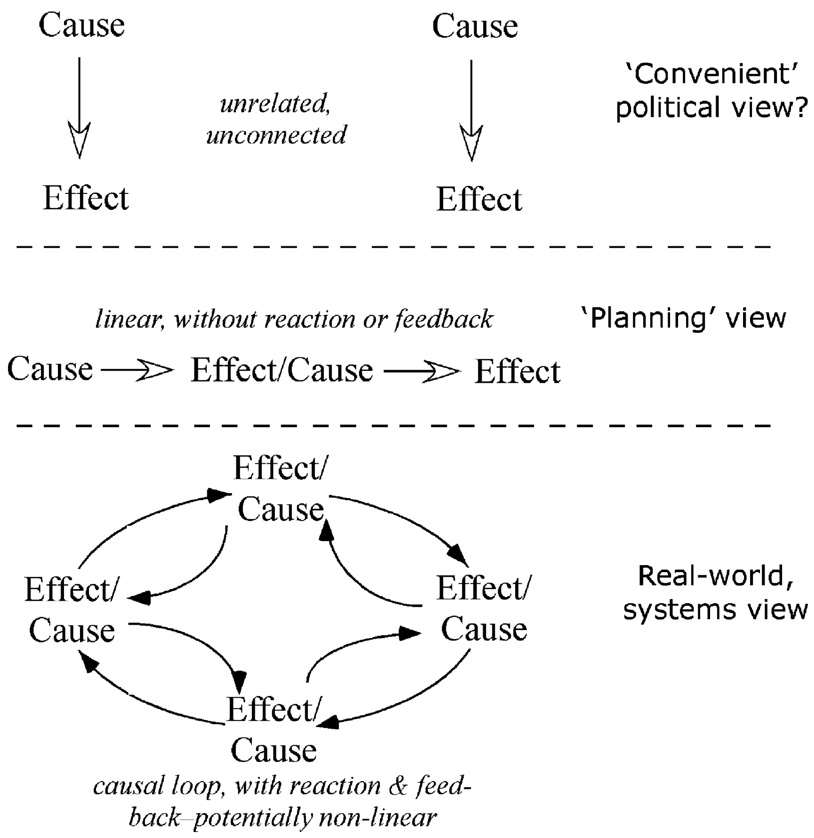

Figure 3.3 Views of causality. Some view cause and effect as singular, one-on-one; this is a reductionist view. Some consider that there are 'chains of causality,' where a cause produces an effect, which is seen as cause, so producing a further effect, and so on. In practice, however, causes not only produce effects, they also provoke reactions, or feedback, giving rise to the third viewpoint — the causal loop, nonlinear feedback viewpoint.

cause producing another effect, and so on. This is the viewpoint taken, for example, by planners, project managers, systems engineering managers, etc., when they formulate plans and when they seek to cost and implement those plans.

The third viewpoint is that every cause creates some reaction, which will either reinforce the cause (positive feedback), or counter the cause (negative feedback). At the most basic level, this is a reiteration of Newton's Third Law: 'action and reaction are equal and opposite.' However, while that can never be wrong, loops of causality may be more complex. Figure 3.4 shows an example of a causal loop model (CLM), which shows a simplistic view of population dynamics. Such a simple model would serve as a seed or nucleus on which to build a more sophisticated model showing, for instance, infant mortality, the effects of famine, nutrition and medical care, etc., as they affected births and deaths.

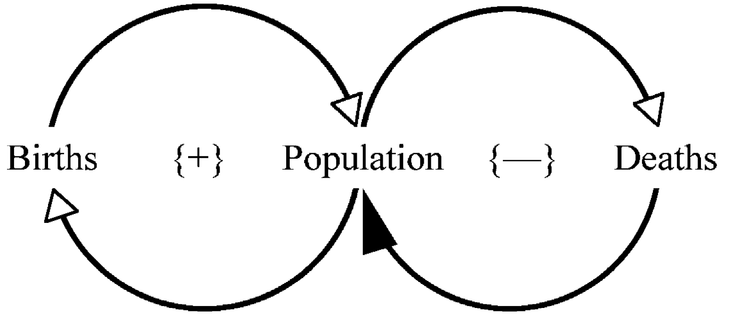

Figure 3.4 Population dynamics. The arrows are termed 'signed digraphs,' and indicate causality. Read open-headed arrows as 'increase(s);' read solid-head arrows as 'reduce(s).'

Note that the CLM in this instance is proposing causality in loop form, rather than illustrating any process. This is a classic open system concept, in that the population numbers are a balance between the opposing influence of births (increasing) and deaths (decreasing).

The CLM is a useful device:

- It presents a viewpoint simply and clearly, such that an observer is in a position either to understand and to agree, or to challenge, disagree and suggest changes, corrections and improvements.
- For this reason there may be 'good' CLMs and 'poor' CLMs. A good CLM would be one that is simple, clear and convincing, in that it appeals to reason and logic, and — by virtue of its emphasis on causality — addresses the core of the subject in hand.
- CLMs are dynamic in nature; rather than presenting a static, structural view, the use of causality results in a dynamic view capturing the essence of change — in the case above, the dynamics of population.
- CLMs provide a useful basis for dynamic simulation modeling; often, a good CLM can be transformed directly into a dynamic simulation, which can then examine those aspects that the CLM cannot address, such as rates of change. In the example, for instance, there might be an interest in assessing the likely change of population over the next fifty years. This cannot be sensibly assessed from the CLM, but a simulation based on the CLM may be more successful.

Influence diagrams also use signed digraphs. Influence diagrams introduce, in addition to chains of causality, 'influences:' an influence is a factor which may affect something, but which may not be considered directly, or uniquely causal. Influence diagrams may be more complex than causal loop models; causality is harder to justify than influence. CLMs tend therefore to be both sparser, and more pointed.

The CLM presentation, using loops of positive and negative feedback, is also used to present interactive processes. In such instances, they should, perhaps, be more properly called process loop models. Nevertheless, the CLM is valuable in exploring, understanding and anticipating complex problems and behavior.

Dynamic simulation of phenomena

It is possible to dynamically simulate phenomena, without necessarily representing whole systems. This is still generally referred to as systems thinking.

An example of this approach to systems thinking is shown in Figure 3.5. At the time of the Pyramids, *ca.*2600BC, the population of Egypt was about 1.5 million, and some 4000 men were directly employed on building the Great Pyramid of Khufu, ably supplied and supported by food, beer and other supplies from the whole of Egypt.

The figure, which uses the simple STELLA™ notation, shows the Nile producing the annual inundation or flood, on which the early civilization of ancient Egypt depended. The component representing the Nile is part of the weather model according to Edward Lorenz, shown in Graph 2.3 in Chapter 2. This model of chaotic weather has been used to 'drive' the population model, since variations in the annual flood affected the food supply each year.

The dynamic simulation uses the amplitude of the annual inundation to determine the amount of food grown each year. This in turn feeds the population, which grows exponentially under this 'regime of plenty,' until a point is reached, at which the population's needs match farming's capabilities. Thereafter, a poor inundation that results in little food will either slow population growth, or cause famine and major loss of life. This is known to have happened unpredictably in ancient Egypt. See Graph 3.1.

The simulation of Figure 3.5 is evidently the simulation of the phenomenon of self-organized criticality, in that the population was auto-expanding to a critical state, where it perched 'on the edge of chaos,' to use the popular expression. Graph 3.2, a phase-plane chart drawn from the same figures that generated Graph 3.1, highlights the situation. The phase-plane chart shows the variability in population, year-on-year, and, although complicated, the graph indicates that there is a central 'attractor,' i.e., the population may go above and below some level, but always returns from the extremes of low and high, to this intermediate level, as with Bak and Chen's sandpile. The situation was always self-correcting: famine reduced the population, restoring balance between food supplies and consumers; it was never possible for the situation to cross over to the 'region of chaos.'

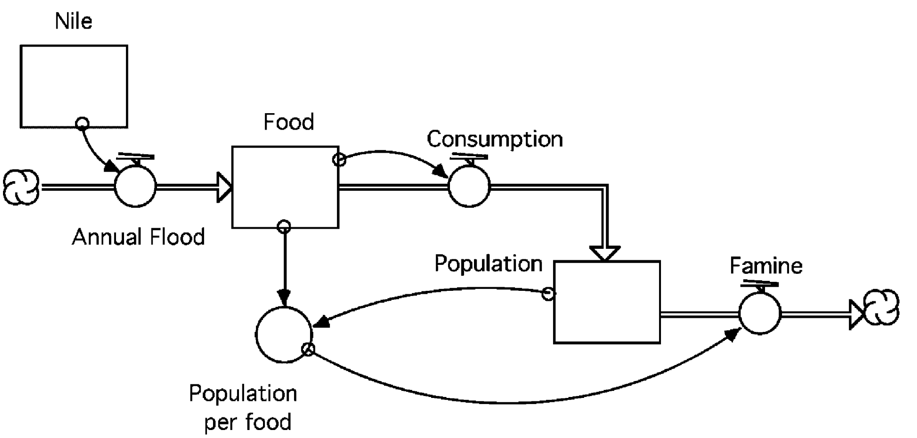

Figure 3.5 Dynamic simulation of the annual inundation of the Nile and its effect on population growth, *ca.* 5000 BPE.

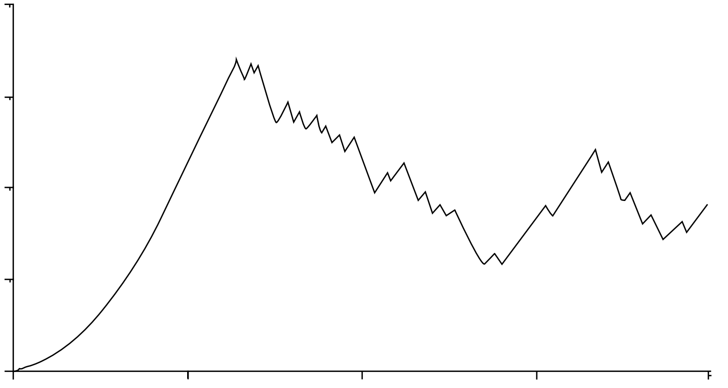

Graph 3.1 Population growth from the simulation of Figure 3.5. This is an example of self-organized criticality.

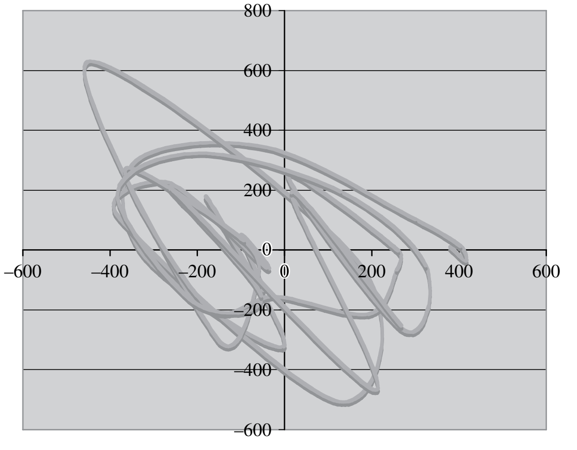

Graph 3.2 Phase-plane chart, drawn from Graph 3.1, showing the variability of the population, year-on-year over a 100-year period — the 'signature' of self-organized criticality?

Dynamic simulation may be conducted for almost any kind of phenomenon. Figure 3.6 shows a dynamic simulation of an archetypal systems engineering process, starting from requirements provided, presumably, by a customer, and ending with the solution system being commissioned. This particular simulation was occasioned to refute the claim by a tool vendor that, no matter how long was spent at the outset of such a process rectifying errors and omissions in customers' requirements, such efforts would not increase the overall project time. On the face of it, this was a nonsense claim. Unexpectedly, the simulation indicated that it was quite correct, however.

Time to clear defects found in requirements varies in practice. Simple, obvious errors are found quickly and easily. As work progresses, further errors prove more difficult to locate. The simulation assumes a square law time distribution: if it takes one unit of time to find 50% of errors, then it takes two units to find 60%, four units to find 70%, and so on up to 32 units of time to find

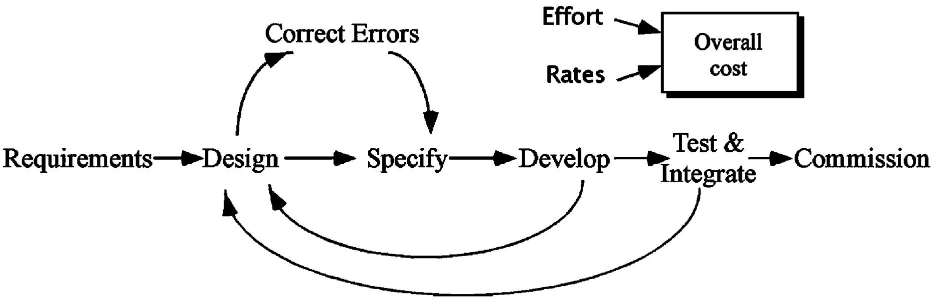

Figure 3.6 Model for dynamic simulation of a simple systems engineering process working from requirements at left to commissioning at right. The requirements contain errors and omissions that take time to correct, with some even going undetected until later in the process where they cause rework to be needed — indicated by the two right-to-left feedback loops.

100% of all defects. 100% may be impossible in practice, but it serves as a benchmark of 'best conceivable performance.'

As the simulation model shows, many errors are found and corrected during the system design phase, from which will emerge systems specifications. Some may slip thorough, however, to be found during development — at which stage, they are likely to invoke rework, shown in a feedback loop in the model. Some even more obscure errors may emerge only during test and integration, where they, too, are likely to invoke re-work.

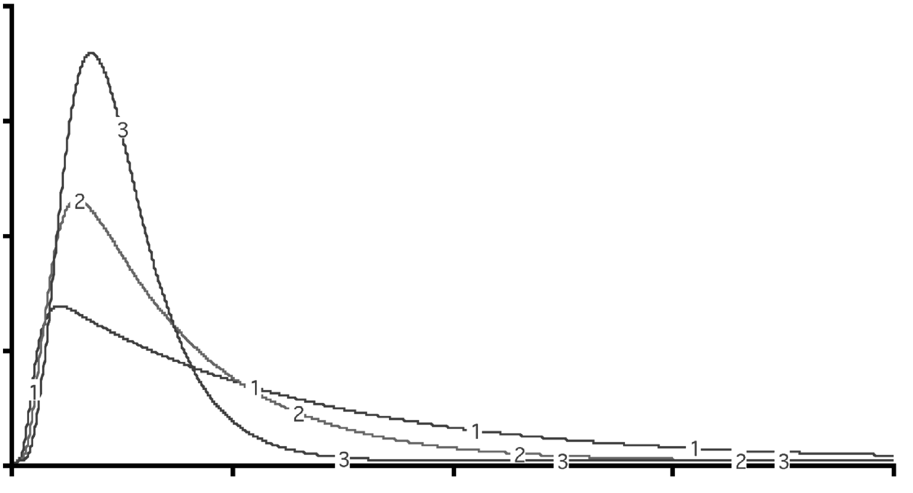

Graph 3.3 The graph shows the rate at which requirements are effectively commissioned. The x-axis is time, while the y-axis is effectively effort. Lines 1, 2 and 3 correspond to increasing design error detection rates Line 3, the 100% line, takes longer to climb due to the increased error detection time, but it peaks higher (implying a greater concentration of effort) and finishes earlier — i.e., completes the commissioning phase of the project. STELLA™ simulation of Figure 3.6.

The results from running the simulation are shown at Graph 3.3. These show that the overall time from 'requirements' to 'commissioning' is dominated by feedback — anything that can reduce, or eliminate feedback will have a beneficial effect on both time and cost. In this case, any time spent at the outset eradicating errors in requirements is well spent, reducing the overall time and cost out of all proportion to the time taken to eradicate errors.

Dynamic interactive systems simulation

There is a difference between dynamic simulation of phenomena, such as the requirements error correction simulation above, and a systems simulation. The latter requires that the essential characteristics of the systems approach be in evidence, i.e., the simulation must present the system of interest in its context of being open to, and adapting to, its environment, including other interacting systems. This accords with everyday common sense: we would not model the performance, say, of car driving through London without due consideration of the road, road surface, cross-wind, other traffic on the road, etc; were we to look at the car in isolation, or in a wind tunnel, we might achieve rather pointless results.

In the general case, there is an implication that we will need a model, both of a specific system of interest (SOI), and of its environment. This view of dynamic systems modeling is organismic: it works for flora, fauna, the body, and for any complex system. Figure 3.7 shows a notional system model in which the system of interest might be, say, the one on the left. The system on the right might represent the whole environment, of which the left-hand system is effectively part. So, in the case of the human body, the cardiovascular system can be seen as interacting with all the other systems within the body, all of which together constitute its environment. Similarly, the pulmonary system can be seen as interacting with all the other systems within the body, which together constitute its environment, including the cardiovascular system. So, perhaps, the SOI should be seen, not so much as interacting with its environment, but as being an intimate part of the environment along with all the other interacting systems. . . .

It is the essence of complex systems that there is a high degree of interaction, and of mutual adaptation, between the parts. Not to represent this in any systems simulation is to risk misconstruing systems behavior.

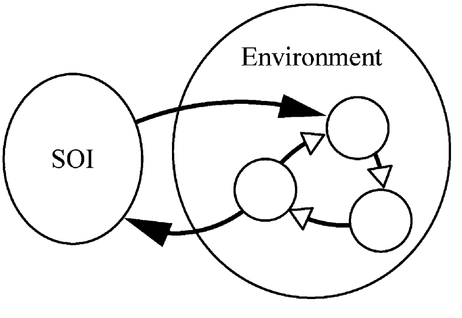

Figure 3.7 Systems modeling. Any system of interest to be examined, explored, understood, or modeled, can be sensibly addressed only when active, interactive with, and adapting to, its environment. For systems modeling, this indicates that any system model should appear only as a part of some larger, wider, containing, whole.

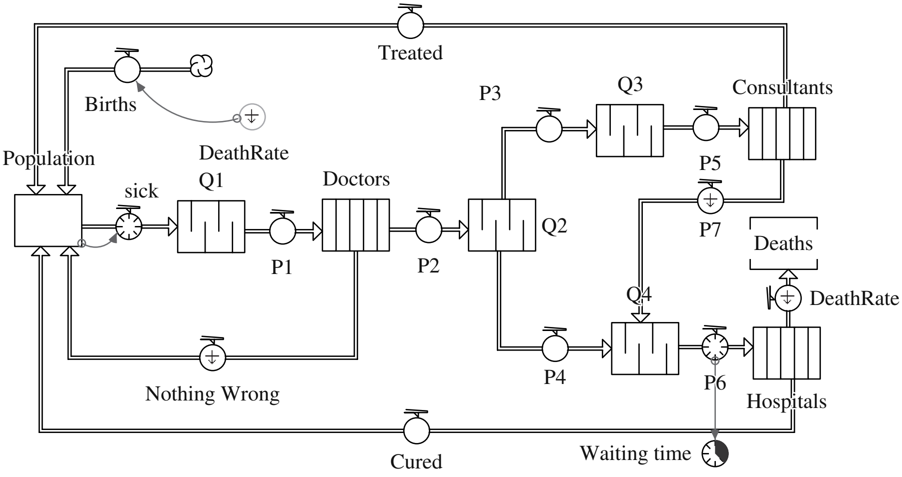

Figure 3.8 Closed loop system model with systems free to adapt to their environment.

Figure 3.8 illustrates the impact of such considerations in more practical terms. The simulation shows a simplified healthcare scenario: doctors treat patients from the general population; they hand some patients off to consultants, who assign some to hospital treatment. Patients either return to the general population in good health, or expire. The birth rate is coupled to the death rate, although the death rate is lower due to improved medical care; as a result, the population is increasing. Notice that there are no inflows and no outflows — this is a closed model of open systems: population, doctors, consultants, and hospitals. The doctors serve as a choke, or control system between the population, the consultants and the hospitals as systems. The inflow of 'sick' persons to the doctors, then, is from the population. But the outflow from the doctors also goes into the population, either directly or through the other systems/agencies. This is true for all four systems — each finds that the others constitute its environment.

This is an important feature of a system model — the system of interest appears as part of its whole, not in isolation. Moreover, the 'behavior' of the system model depends in part on its interaction with the other systems forming its 'environment;' so, those interactions must be in evidence, too, else neither the whole nor its parts will exhibit sensible behavior.

Graph 3.4 shows some results from running the simulation of the healthcare system of Figure 3.8. The x-axis is time: the y-axis is the number of patients seeing doctors at any time. The graph is limited in the y-direction by the number of doctors available, which is presumed constant in the model. Excursions below this upper limit indicate when the doctors are not fully occupied with patients. Note that this amount of 'less than full occupancy' is reducing with time; the doctors are becoming busier as time progresses. This is an inevitable result of improved healthcare, which leads to reduced death rate, to increasing population, to more patients to see the doctors. The price of success in healthcare, then, is an ever-increasing workload. . . .

Behavior modeling

Systems thinking is often coupled with behavior modeling. Behavior is the way something, a system, acts or reacts. As the simulations above have shown, it is reasonable and practicable to

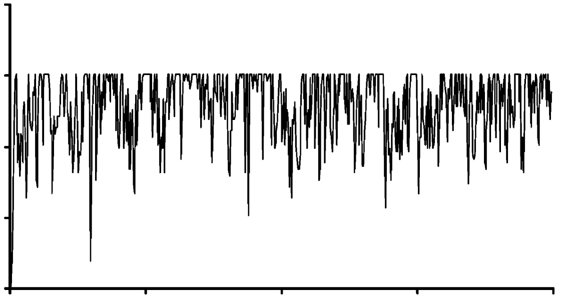

Graph 3.4 Graph from the dynamic simulation of Figure 3.8, showing the numbers of patients consulting doctors. Notice that there is a gradual, but perceptible, increase in the mean number of patients with time. Improved medical care prolongs life and so increases the population and the number of potential patients . . . .

represent and simulate the actions and reactions of systems without explicitly presenting their form or internal functions. This is behavior modeling. It relies for validity upon sensible understanding of the actions and reactions of the various parts within some whole. So, in the healthcare simulation example, how the various systems within the whole function, act, or behave, is well known; there is no need to represent the doctors' surgeries, the hospital wards and the nursing staff, etc. How the whole healthcare system behaves as a result of the interacting population, doctors, consultants and hospital behaviors is less well understood, and is the subject of the simulation.

As this example shows, then, it is possible to bring together the behaviors of a variety of systems making up some whole, to cause these behaviors to mutually interact in closed cause–effect loops, and so to 'observe' the evolving behavior of the whole. In the example, we may even see (simple) emergence in the population increase, brought about by improved healthcare. . . this is reasonable within the definition of emergence: properties, capabilities and behaviors of the whole that are not exclusively attributable to any of the rationally separable parts.

Systems Thinking and the Scientific Method

The scientific method experiments, formulates a theory about how something works, operates or behaves; uses the theory to make predictions about further working, operating or behavior; and, compares the predictions with reality. If the predictions do not match reality, then the theory is wrong. The scientific method can only disprove a theory; it cannot prove it, since it is generally impossible to experience all possible conditions and situations.

Systems engineering seeks to anticipate the future, in that it conceives and furnishes a solution (system) to some problem. The solution system will adapt to its environment (and it is in this *continually adapting condition* that it will be expected to meet its goals and objectives): systems engineering is expected to anticipate such adaptation and create solution systems accordingly. This usually invokes the use of some kind of dynamic simulation to accommodate the complexity, and to address future situations about which there may be only restricted information.

It is not possible to predict the future sensibly: certainly not beyond a near-term time horizon. Beyond that, it may be possible to predict trends, but little more. As we come to understand chaotic systems with their strange attractors better, we may be able to improve on longer-term projections. For the time being, however, systems engineering faces problems in prediction.

No matter how accurate the detail of simulation models may be, they offer no guarantee that the behavior they exhibit will match that of some future real world. This has led some systems engineers to design and build solution systems on the assumption that they will not adapt once in use and operation; and, it has led some systems engineers to abandon modeling and simulation in favor of a simple linear building-block approach to building rigid solutions, in which the linearity encourages them to believe that the solution will behave in its environment just as they designed it to behave in isolation. This is, in effect, abandoning science — and experience — for a more pragmatic approach of meeting the requirements of some customer by creating precisely what the customer asked for. . . which might be considered ideal in business terms — provided the customer does not seek guarantees of systems utility in operation.

A more reasoned approach is to make best endeavors to predict the behavior of the future solution system in its future environment, and to adjust the various (simulated) parts of the solution system so that the whole solution system exhibits the required effects, achieves the required goals and objectives, while interacting with, and adapting to, its future environment and interacting systems. Knowing that accurate prediction is not possible, the kind of simulation model that should be employed is one in which experiments could be carried out — a so-called learning laboratory (Richmond, 1992).

Using a learning laboratory enables the seeker after enlightenment to conduct experiments and to observe the outcome in the simulation. Seekers may observe counterintuitive effects (Forrester, 1971), which could be potentially problematic or advantageous. Since they are experimenting, such findings should serve as warnings or as windows of opportunities rather than precise results. But, not to experiment might be deemed as an oversight at least, and possibly irresponsible.

The simulation model of Figure 3.8 is a simple example of a learning laboratory. Using the model, it is possible to alter the numbers of doctors, consultants and hospitals; to consider what would happen if birth rates and death rates changed, and what effects that might have within the healthcare system; to estimate waiting times throughout for seeing a doctor, seeing a consultant, getting a hospital place, etc. By altering, e.g., the number of doctors working in parallel, it is possible to cut queues and speed up throughput — at a cost, which can also be estimated. In the model, the icon ‘waiting time’ calculates the mean time from an individual becoming sick, queuing for and being seen by a doctor, waiting for and receiving a consultant’s appointment, and waiting for a hospital bed to become available. Waiting time is evidently an emergent feature of the whole.

System thinking in this manner can be reasonably deemed scientific, in so far as the simulation model used exhibits appropriate behaviors, and the experimentation is conducted rationally and logically. It also has the added bonus of falsifiability (Popper, 1972), i.e., a knowledgeable observer can see at once if the model is incorrect, or if the experiments being conducted are unsound, unfair, improper, give unfavorable results, or are insupportable in any real-life context. In systems thinking, it is possible to pursue the scientific method with the only proviso that experiments are conducted in simulation of the real world, since they cannot be conducted in a future, as yet nonexistent, world.

Representing and Modeling Systems

Systems may be represented in several different ways, according to one’s understanding of what a system is, and what purpose the representation is to serve. Defining a system as a ‘dent in the fabric of entropy,’ suggests that bringing disordered entities into an ordered state would constitute forming a system, and indeed some engineers look at systems this way. There is no right or wrong in such definitions, but there is ‘more useful’ and ‘less useful.’ It is possible, but not useful, to view a brick wall as a system, but it is a static, closed view; notwithstanding, a fundamental aspect of any system is that it represents relative order, or reduced entropy.

Definitions of ‘system’ abound. Dictionaries offer ‘A combination of related elements organized into a complex whole;’ and, ‘complex whole formed from related parts.’ Wikipedia offers: ‘System is an assemblage of elements comprising a whole with each element related to other elements. Any element which has no relationship with any other element of the system, cannot be a part of that system. A subsystem is then a set of elements which is a proper subset of the whole system.’ None of these captures the essence of Gestalt, which is at the heart of systems and systems concepts, although they all recognize the concept of ‘whole,’ which by association with holism, should carry with it the implication of the whole being greater than the sum of the parts.

The Oxford American dictionary offers ‘system:

1. *A system of canals*; STRUCTURE, organization, arrangement, complex, network, informal setup
2. *A system for regulating sales*. METHOD, methodology, technique, process, procedure, approach, practice, means, way, mode, framework, modus operandi, scheme, plan, policy, program, regimen, formula, routine
3. *There was no system in his work*. ORDER, method, orderliness, systematization, planning, logic, routine.
4. *Youngsters have no faith in the system*. THE ESTABLISHMENT the administration, the authorities, the powers that be, bureaucracy, officialdom, status quo.’

Little wonder that ‘system’ is an overworked word. Each of these four ‘definitions’ is more a list of examples of how the word is used and abused, than what a system is, or might be. Ideas of wholeness, synthesis, etc., are conspicuous by their absence.

A definition from INCOSE (International Council on Systems Engineering) offers:

> ‘a system is an interacting combination of elements viewed with regard to function.’

A chemist might apply this definition to common salt, which is an interacting combination of the elements sodium and chlorine. Sodium, a soft, silvery alkali metal, and chlorine, a greenish yellow gas, combine to form the familiar translucent cubic crystals of common salt — a neat example of emergence. Salt occurs naturally, and just *is*, but people have put salt to many uses: preserving and flavoring food; as natron (a naturally occurring mix of sodium and potassium salts) for desiccating and preserving mummies in ancient Egypt; as freezing mixture, making ice cream and preserving food; and many, many more. This form of definition says less about ‘system’ *per se*, more about what a system is perceived as doing: no Gestalt, no whole, no synthesis. . . indeed, it is rather hard to determine just what this form of definition either means, or offers to those seeking enlightenment. In short, it does not appear to be particularly helpful.

INCOSE has another definition of system (INCOSE, 2006):

> A system is a construct or collection of different elements that together produce results not obtainable by the elements alone. The elements, or parts, can include people, hardware, software, facilities, policies, and documents; that is, all things required to produce systems-level results. The results include system level qualities, properties, characteristics, functions, behavior and performance. The value added by the system as a whole, beyond that contributed independently by the parts, is primarily created by the relationship among the parts; that is, how they are interconnected (Rechtin, 2000).

All of which seeks to define by giving examples, and to describe emergence and hierarchy without mentioning either term. The core definition is given in the first sentence, while the remainder is descriptive narrative of parts, emergence and hierarchy. The final sentence just misses the point that the emergent properties, capabilities and behaviors are created, not by relationships and how parts are interconnected, but by the dynamic interactions between parts.

Overall, however, this definition contains useful notions of 'Gestalt,' 'wholeness,' synthesis,' 'organicism,' and 'value.' This last implies that a system can be measured, and that some value can be put upon it. Evidently, too, this definition of system is concerned with human activity in making, forming, or encouraging the formation of purposeful systems. It excludes, for instance, the solar system, but it includes social systems, human activity systems, sociotechnical systems, socio-economic systems, etc. The inclusion of 'hardware and software,' in place of, say, technology, suggests that this definition was brought about through some form of uncomfortable consensus, with some of the consensual members being engineers concerned with processing. . . the definition seems to be aimed at tangible solutions, and to exclude, for instance, a process as a system.

Previously, a system has been defined as:

> An open set of complementary, interacting parts with properties, capabilities and behaviors emerging both from the parts and from their interactions to synthesize a unified whole.

Such a definition encompasses most of the above definitions, and allows a process to be a system, and the solar system to be a system, too. Defining a system in this way indicates that a system could be represented as, or by:

1. An open set of complementary interacting parts, in the same way as the human body could be represented as a set of complementary organs, or organic subsystems.
2. A set of emergent of properties, capabilities and behaviors which together describe the system 'from the outside,' as it were, and with no indication of the systems internal form, functions or behaviors.
3. A unified whole, i.e., without describing either internals or behaviors, etc. This is more normal practice; we call an airplane 'an airplane,' without going into any detail about how it works, what size it is, how many subsystem it has, or whatever.

The three representations offer decreasing levels of detail. A key attribute of systems thinking, science and theory is the ability to represent a whole system using a simple label, such as its emergent properties. This is the key to managing complexity, by concealing it.

From the various definitions given above, a system may be a process, a procedure, a method, a arrangement, etc. What is less clear, is how one might represent a system diagrammatically, or in a model. Figure 3.9 illustrates the dilemma.

In the figure, the upper diagram might be called a structural viewpoint. It shows a bounded containing system, within which are three complementary, interacting subsystems within a bounded environment. One of the subsystems is shown as containing a further three complementary, interacting sub-subsystems. (In reality, it is likely that there would need to be more than three subsystems to furnish a full complement — three are shown for convenience of presentation.)

The top diagram is known, somewhat irreverently, as the 'poached egg' diagram, for obvious reasons. It indicates that an open, interacting system is comprised of open, interacting parts which are themselves systems, and that these subsystem parts are, in their turn, comprised of open interacting parts which are themselves systems, or sub-subsystems. This is a nesting model of systems within systems within systems, *ad infinitum*; like the proverbial fleas. In a real example, it may show

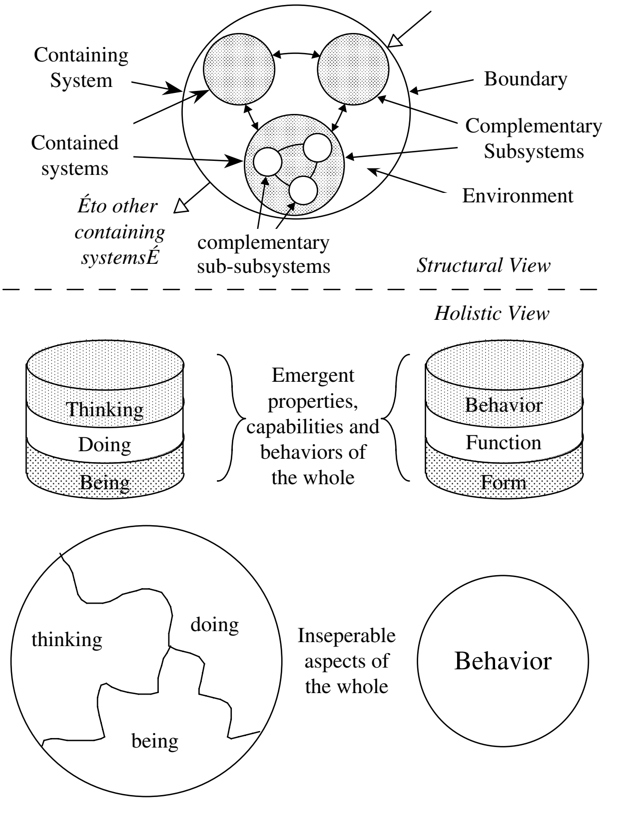

Figure 3.9 Structural and holistic views of a system.

considerable detail: it has the advantage of presenting order, structure and hierarchy to the viewer. It has nothing overt to offer with respect to emergence, unless it is understood that each system, and each system within a system is defined as a system by virtue of its emergent properties, capabilities and behaviors.

The lower diagrams of Figure 3.9 show a quite different view — they offer the viewer *only* the emergent properties, capabilities and behaviors of the whole, without internal detail, structure, etc. From one perspective, the emergent properties of a system may be seen as aspects of being, or existing, doing, and thinking, suggesting intelligence. An alternative view sees the emergent properties, capabilities and behaviors as comprised of form, or morphology, function and behavior; (some analysts propose technology, process and people for appropriate sociotechnical systems)

In each of the center diagrams of Figure 3.9, there is an implication, not so much of structure, but of substrates. For example, being or form is considered as a foundation: it is difficult to envisage process and function floating in space without supporting mechanism, structure, effectors, etc. Doing, or function, then exists on the substrate of physical entities; processes, as it were, travel over the surface of the substrate which offers pathways and energy to direct, enable and support the processes. Not all system have a third layer: thinking or behaving. Where they do, this may be thought of as sitting above the level of doing or function, in the context of responding to external stimulus.

Thinking and behaving systems may respond differently to the same, repeated stimulus according to experience, history and context. This implies some evaluative and decision-making process . . . .

Finally, in the lower figures, a purely holistic view is presented: on the left, the three aspects of being, doing and thinking are presented as irretrievable intermixed, but without any sense of structure. At the right, only behavior is shown. This is reasonable only where the dynamic interactive properties of some system are to be represented in open interaction with other systems, and where the behavior of the system to be represented can be described, and relied upon to remain dynamically stable for the duration of simulation time.

This is not quite as radical as it might seem. If the systems in question were, for example, trained and experienced soldiers either as a platoon or individually, then it might be quite reasonable to predict and rely upon their behavior in a given set of circumstances when they come into contact with some hostile opponent. We would not, generally, consider it necessary to know, in detail, about each soldier, his particular role in the platoon, his physical dimensions, what make of rifle he is carrying, etc. We would hope and expect that the platoon would act as a whole, would engage the hostile force — if appropriate — and would attempt to minimize own casualties. . . .

Similarly, if the system was a collision avoidance radar, part of an avionics system in a transport aircraft, we might have little interest in its construction, the many functions of which it was capable, etc., provided we could be confident of its behavior in detecting, and alerting aircrews to, relevant danger. In simulating air traffic management, therefore, the details of particular radar on one of many aircraft in the 'system' at any time are largely irrelevant; the behavior of that radar would be sufficient for purpose.

Nonlinear Systems Thinking

Many real world systems behave in a nonlinear fashion. Figure 3.10 presents a classic example, of interlinked rabbit and fox populations. Rabbit population increases with a plentiful food supply. A plentiful supply of rabbits 'feeds' the fox birth rate and fox population. Foxes hunt rabbits, increasing the rabbit death rate and reducing the rabbit population, allowing the grass and vegetation to recover. With less rabbits, there will be less food for foxes, the fox birth rate falls, predation rates fall, and the rabbit population will start to recover. And so the cycle continues.

This cycle may be formulated as a set of nonlinear, simultaneous difference equations, to which there is no fixed solution; instead, there are an infinite number of solutions as the two populations

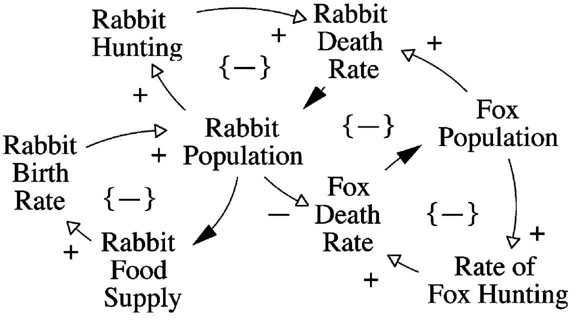

Figure 3.10 Rabbit and fox populations — nonlinear system behavior (Hitchins, 1992).

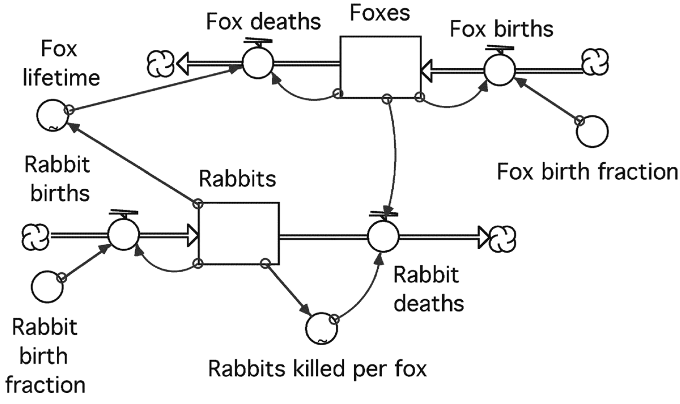

Figure 3.11 Fox and rabbit population simulation.

continuously interact. As Figure 3.10 indicates, it is possible to introduce foxhunting as yet another loop in the model.

Figure 3.11 shows a STELLA™ version of the causal loop model of Figure 3.10; the model has been simplified, and excludes both grass and foxhunting, concentrating instead on the interaction between the fox and rabbit populations. Each population is shown with its births and deaths respectively. Cross-linking between the two populations is shown in both directions: increasing rabbit population also increases fox lifetime, so reducing the fox death rate; and the more foxes, the greater the rabbit death rate.

The results of running this simple simulation are shown in Graph 3.5. As the graph shows, the two linked populations oscillate, some 90° out of phase, with the fox population (line 1 in the graph)

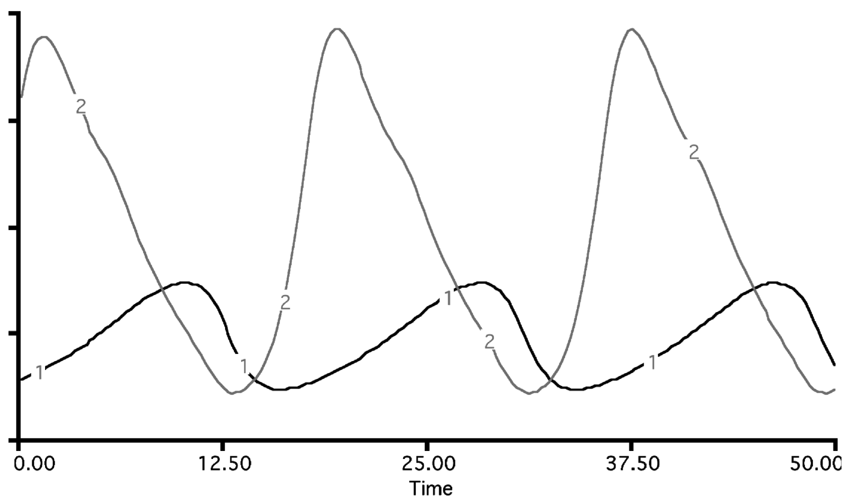

Graph 3.5 Graph showing oscillating fox–rabbit populations from the simulation of Figure 3.11.

lagging behind the rabbit population, line 2. The simultaneous, nonlinear difference equations governing the two oscillating populations are developed in STELLA™ within the dynamic simulation model.

Adding the effects of grass shortage and of foxhunting would be straightforward, but would, of course, result in a more complex graphical result. With such additions, the simulation model might form the basis for a learning laboratory to explore the effects off hunting on, say, vegetation and grassland, when hunting should best be undertaken to benefit the farmers who usually own the land being hunted upon, etc.

There is, then, no single solution to the question of population size; instead, there is an infinite set of solutions. This oscillating behavior is characteristic of much social, biological and physical behavior; such behavior might be viewed as a form of homeostasis, or dynamic equilibrium of, in this instance, the overall system of fox–rabbit interactions. The emergence of such behavior can be counterintuitive, however, which underscores the advantages of dynamic simulation as part of exploring and understanding the problem.

A further consideration stems from the observation that the nonlinear, indeterminate behavior shown in this instance and more widely occurs despite many of the elements in the overall system behaving linearly. This is something to be noted by engineers and other designers who seek to make 'their part of ship,' i.e., their part of some overall design and development strictly linear, thereby expecting the whole to be linear; the whole may well not behave linearly, even where some — many — of the parts do. A simple, biological example might consider a heart-transplant patient who has received a mechanical heart with a fixed rate and volume of pumping blood around the cardiovascular system. Would this steady rate of pumping, enable the patient to act and behave as a normal person with a normal heart? Or, would the linear heart pump impose limitations on the capabilities and behaviors of the patient? That bears some (systems) thinking about. . . .

Summary

Systems thinking is still developing as a discipline: it is a potentially broad subject since it may be possible to think in systems terms about almost anything. The term, systems thinking, is starting to imply some rigorous application of the systems approach, in which a part is considered, not in isolation, but in the context of its containing whole, such that it is open to, and adaptive to, inflows and interchanges with other parts in that containing whole. Systems thinking, the term, is also applied, perhaps with less rigor, to the dynamic simulation of phenomena.

Systems thinking often faces abstract, complex and even obscure issues and problems; methods have been developed to help cope with such vagaries, including so-called rich pictures associated with soft systems methodology (see page 192), causal loop modeling, N^2 charts (see page 370), and many more. Causal loop modeling seeks to identify causal loops, both reinforcing and reducing or opposing, which describe the active and reactive features in any situation, complex system, etc. N^2 charts show interacting entities and their interfaces, which can also be seen as active and reactive pathways between entities. All of these methods seek to capture the essential dynamics, conflicts, synergies, etc., as a basis for understanding the behavior both of the parts and of the whole.

Causal loop modeling (CLM) is notably useful as a precursor to dynamic simulation, as the CLM elements can be transposed almost directly into the model elements. Dynamic simulation similarly can be conducted by examining the dynamics of phenomena, but can be more rigorous when examining representations of whole system-parts interacting dynamically with other system-parts within a containing whole. In such full models, the system part of interest is free to act, react and adapt to the other parts in its 'environment,' and can be enabled to find its natural level, in balance and harmony with the other parts.

One lesson from chaos is that predictions are of limited value, particularly in complex problems and situations. Simulation models that purport to predict accurately may be unable to justify their predictions. This is true of systems thinking and systems thinking dynamic simulations, too. Such models are often about the future, attempting to understand what will happen as a result of some change, the introduction of a new organization, a new process or some new technology. So, are such models useless?

One way to address the problem is to employ the scientific method. This involves conducting experiments, learning, formulating a theory about the problem based on the results of those experiments, and making predictions about the outcome of the problem based on the theory. If the outcome fails to match the theory, then the theory is wrong. If the outcome matches the predictions, then the theory may not be wrong – but can never be proven right.

Dynamic systems simulation models can be employed sensibly as an environment within which experiments can be undertaken, alternative theories tried out, outcomes tested, etc. Such simulations can never prove that the outcome in the real world will match the outcome in the simulation. However, they can identify some theories as flawed, some as likely to be better than others, and so on. With time, experience and further research, systems thinking has the potential to make better predictions; that is, predictions that are wrong less often, and predictions that are less wide of the mark. That potential suggests that it is far better to think, in systems terms, rigorously and to our best ability, than not to think at all. Moreover, the tools and methods for systems thinking are now so freely available that not to think systems might be deemed irresponsible. . . .

Assignments

1. You are advising on the rules to be applied for foxhunting in a rural area of England. The area suffers from rabbits damaging farmers' crops. Farmers also complain that foxes, who would hopefully keep the rabbit numbers down, are often more interested in raiding hen-houses, catching young game birds, etc. Farmers also report periodic increases in the numbers of rabbits, in particular, and that the problem has got progressively worse since foxhunting with hounds was made illegal. You are required to judge and justify:

- Whether foxhunting should be resumed, with or without hounds.
- If fox hunting is to be resumed, at what point in the cycle of varying fox populations should the hunting be undertaken to offer maximum protection to the farmers' crops from rabbits.

2. A typical assembly plant takes in parts from suppliers, assembles them into subassemblies, assembles the subassemblies into assemblies, and ships these out to be sold. The money received from the sales goes to pay the workers, train the workers, equip the plant with effective machinery, purchase more parts from suppliers, etc., and pay profits to shareholders, if any. Present this overall process as a causal loop model (CLM), with view to developing a dynamic simulation of the whole system at some later time. Identify what value and insights the CLM provides, and also identify what aspects of the model are uncertain and would require the full dynamic simulation to clarify and expose.

3. A factory follows the practice of producing items, placing them in inventory store and allowing the consumption rate to reduce the inventory; this is called 'production push' since the rate of production bears no direct relationship to the rate of consumption. At the same time, the factory

has been attempting to increase the productivity of the workforce, introducing new machinery and improved training and practices. An unexpected side effect of this increased productivity is that the company has entered a cycle of alternatively hiring and firing workers, as the inventory builds up and reduces. So, both the number of employed workers and the inventory that they manufacture vary cyclically. Not surprisingly, the workers are unhappy with this ‘sword of Damocles’ hanging over them, and the factory management is unhappy at spending money on training employee only to fire them within weeks or months. Develop a CLM of the situation, explaining how the cyclic variation occurs. Outline a manufacturing and selling scheme that does not result in significant amounts of inventory and which, at the same time, doe not result in cyclic hiring and firing: develop an alternative CLM for this alternative scheme.