Chapter 11 – The Governance Model

Abstract

Purpose of this Chapter

In this appendix we give a deeper description of the Hierarchical Cybernetic Governance System (HCGS) introduced in the context of complex adaptive and evolvable systems discussed in Chapter 9. The HCGS is that management subsystem that internally regulates the behaviors of subsystems within the larger complex system.

Agents make decisions and take actions in the context of governing the workings and behaviors of complex systems. In Chapter 9 we elaborated the model of a complex adaptive system (CAS) and a complex, adaptive and evolvable system (CAES) and introduced the role of governance as that process that maintains near optimal operations and coordination of the CAS/CAES with its environment, its sources and sinks, and mitigates the effects of disturbances. We introduced the concept of a hierarchical cybernetic governance system (HCGS, see Figure 9.3) that partitions decision types, as described in the prior appendix, into levels corresponding to time scales of decision-making and action.

In this appendix we will elaborate the nature of the HCGS as a generic governance model archetype that can be used to guide the analysis of CAS/CAESs and/or their design. This is a further elaboration of the governance model introduced in Mobus & Kalton (2015, Chapter 9).

Our thesis is: *The structure and operations of an HCGS when enlightened agents are working at the decision nodes, is what makes a CAS/CAES stable and sustainable* (Mobus, 2017). By enlightened we mean agents that make decisions based on reasonably veridical decision models as described in Appendix C and without biases or hidden agendas. We will show examples from natural systems that have evolved HCGSs in which the decisions made are based on constrained models that correspond with the actual environments in which they operate. Of course, of greatest interest to us is the governance of human activity systems such as municipalities, businesses, non-governmental organizations (NGOs), nation states, and, ultimately, the whole human social system as introduced in Chapter 6.

Meuleman (2008) provides an argument for the notion of generic governance archetypes. He reviews what he takes to be three different generic architectural patterns of what he calls meta-governance structure/dynamics[1]. He categorizes these as: hierarchical (as in top-down command-and-control), networks (depending on trust and cooperation among decision makers), and

[1] Recall from Chapter 9 that we elected to not use the meta- prefix for our generic models. Rather we put these models forth as archetypes found to be isomorphic at all levels of complexity and organization among CAS/CAESs.

markets (using business-centric measures of performance and exchanges of value as discussed in Chapter 8 and again in the next chapter). He refers to these as "styles". Examinations of both public and private governance subsystems reveal that organizations may tend toward one of these architectures being dominant. Militaries, for example, are usually strongly hierarchical whereas small to medium-sized high-tech companies lean more on networks. The economic system is argued to be governed by market mechanisms. More generally, and especially when we look at natural governance in living systems, we can find all three of these superimposed on the organization. Meuleman argues that this is, in fact what should be expected as each style addresses different mechanisms for achieving cooperation and coordination. The optimum in management is achieved when all three mechanisms are used. He argues that it takes a strong form of "metagovernance" to make sure these architectures work in consonance with one another. We assert that the HCGS to be described here (and covered also in Mobus & Kalton, 2015) provides exactly the kind of metagovernance architecture that naturally incorporates all three of Meuleman's styles.

There is a fundamental concept that needs to be addressed before going into the details of an HCGS structure, the core principle of cybernetics, using information feedback and feedforward to make adjustments in processes. And we will do that below.

Complex, dynamic systems, interacting with a complex and generally non-stationary environment need to have means for adapting to ordinary fluctuations in that environment. But even more importantly, systems that interact with non-stationary dynamics need to be able to undergo evolutionary changes (that is heritable modifications to phenotypes) internally in order to sustain their existence into an indefinite future. Systems not possessing mechanisms for internal self-regulation as well as coordination with environmental entities will not long persist in the real world.

The core of living systems, is the metabolism of single cells[2]. Cells are CASs that have mechanisms for adapting to moderate changes in their environments. Multicellular organisms are also CASs up to the evolution of the hominids with larger prefrontal cortices. Some great apes show some capacity to form new concepts when in captivity in novel environments (De Waal, 2016), but as of the present the evidence suggests that only human beings have the capacity to think truly strategically, i.e., for long time scales into the future and broad spatial scales. The human brain is thus an example of a CAES making individual humans evolvable systems. Prior to human existence the natural world's degree and rate of changing was relatively low over long geological periods, thus Darwinian evolution accounted for adaptive changes in species; only species and higher genera were CAESs with evolution providing the strategic component. Occasionally, cataclysmic events occurred that were fatal for some substantial portion of species, but not all. In general, the core of life has been maintained somewhere on the planet such that life

[2] In the following appendix we show that core metabolism is a generic economy model for life.

has been preserved[3]. With the advent of human culture, with artifacts and technologies, the Ecos has become strongly non-stationary. The changes taking place in the Ecos due strictly to human activity are now being considered sufficient to rename the current epoch as the 'Anthropocene.' Those changes are taking place rapidly compared to non-human caused changes which take place on geological time scales.

We have already used the concepts of the governance model in several previous chapters (e.g. Chapter 8) using the model archetype to guide our gaining of understanding of the HSS economy, or at least a subsystem of the economy. Now, all of the methods of systems analysis and modeling discussed up to this point will be used here to gain understanding of how very complex systems manage to maintain their existence in a constantly changing world. We present a general model of management and governance subsystems for larger complex adaptive and evolvable systems (CAES). We then will demonstrate how these concepts work in both natural and manmade systems.

As complex systems are being analyzed, the questions of how a system is self-regulated internally and how it manages to interact successfully with its environment guide the examination of its governance architecture and management subsystem. This appendix will provide the basic structures and functions that are part of every such system. What we propose is a step toward a general theory of governance.

CAESs are found to be composed of subsystems that have *agency* and degrees of *autonomy* as discussed in the previous chapter. An agent, as described in that appendix, is a special information processor that can take decisions and behave in variable ways depending on their information input along with the decision models with which they operate and a computational engine for doing processing. Decision agents are not capricious or should not be. They follow a set of guidelines (the decision model) for arriving at decisions for what to do next. This does not mean, of course, that their model is necessarily veridical. Indeed, as we will argue below, in the case of human agents there are too many times when the models are contaminated by a number of biases and flaws that make them suboptimal – a major source of what goes wrong in societal governance. Their behavior, in turn, however, has causal impact on other components in the system including, and especially, other agents. Every work process in any CAS or CAES will have an agent decision maker with sufficient autonomy to make the decisions that should, in principle, be able to maintain optimal operations of that process, subject to constraints imposed by interactions with other work processes and agents.

Cellular metabolism is fundamentally working at the operations level (recall from Chapter 6 that metabolism is the first form of an economy model). Decisions are made in the sense of homeostatic regulation. The agent model that applies generally is the homeostat (see below,

[3] Indeed, such cataclysmic events, such as the meteor or comet that is thought to have struck the Yucatan peninsula 65 million years ago, causing massive die offs and extinctions create opportunities for surviving species to undergo adaptive radiation subsequently. Life manages to flourish in spite of non-stationary environments.

section [Feedback] for the description of a homeostat) using error feedback to signal a response. It is implemented in biochemical reactions.

However, an argument can be made that some coordination level decisions are also involved. For example, the production of a specific protein is regulated by the relative concentration of messenger RNAs for that protein in the cytosol near ribosomes. That concentration changes more slowly and the mRNAs are manufactured by enzymes reading out the code from the gene's DNA. And that rate of read out was triggered by demand somewhere in the cell for the protein needed. In other words, a deficit of that protein somewhere in the cell sets in motion a chain of communications that result in the increased production of the mRNAs, which, results in increased production of the protein and eventually reestablishing the needed concentration. Similarly, tactical decisions are made when the cell receives messages from its environment, say a particular molecule that is an agonist, activating a membrane channel allowing a current of ions to flow into the cell and set off a cascade of signaling mechanisms as well as work activities. For example, the activation of a postsynaptic membrane by the receipt of neurotransmitter molecule. Penetrating channels through the membrane open to allow the influx of sodium ions and that leads to depolarization of the membrane. If the depolarization event is sufficiently strong it gets propagated outward from the synapse and if reinforced with other similar depolarization events may lead to the neuron firing off an action potential that will propagate along the axon to signal other neurons. Thus, the neuron is a sophisticated information processor. Many more examples of cells interacting through tactical mechanisms are available.

In multicellular organisms we can similarly point to many instances of operational level and coordination level decision mechanisms. In most cases, for example the endocrine system, these are implemented in biochemical interactions, but now at the level of tissues and organs. In some cases that system is triggered by brain processes, i.e. computation that more closely resembles what we usually think of as an agent.

In all of life up to the human species the strategic level of agency is subsumed in the process of Darwinian evolution. Nature makes the de facto strategic decisions through selection of what works best in any particular econiche. Humans are still subject to evolution but not entirely dependent on it.

Humans are the epitome of agency and the need for management and governance of human organizations, from families up through nations, is certainly well recognized. Not only are humans agents, but their degrees of freedom in behavior are nothing short of spectacular. Thus, whatever a human decides and does can produce significant amounts of information in the system, i.e. they can generate a great deal of surprise to other humans. Humans can also suffer information overload, or data processing jambs, that impair their decision-making capabilities. Humans rely heavily on internal tacit and explicit memory in the form of concepts to model what they have experienced in the world. Memories can be slow or faulty and lead to information overload as well.

Finally, humans are subject to sensory distortions and noise – they don't always see the world as it really is. Thus, humans as agents are notoriously bad when it comes to situations involving great complexity and fast dynamics. All of these factors lead to the need for an overall governance architecture for organizations and management practices carried out so that harmful mistakes in judgments can be minimized.

This appendix will provide a general model of governance and management that can be applied to any CAES. The implementation details will obviously vary from one kind of system to another. But the basic model will be found operative in all.

This appendix provides a general overview of the governance model archetype[4] and will provide some examples from various CAS/CAESs found in nature and human activity systems. This model archetype describes a hierarchical cybernetic network of distributed decision makers, agents, as described in Appendix C. The hierarchy is a leveled structure of cybernetic processes and we refer to it as a hierarchical cybernetic governance system (HCGS).

At the base of the hierarchy is the operational level which is comprised of all of the organized work processes that constitute the CAS/CAES. The organization of the work processes is such that the CAS/CAES carries on a basic economic process (Appendix E), supporting processes (e.g. reproduction for all living systems), and auxiliary processes (e.g. philosophy for humans). The governance system is tasked with decision making duties that facilitate all of these processes working in concert, meaning long-term synchrony, and continuing to be stable in the face of disturbances from the environment (e.g. temporary loss of a resource). The proper governance of a CAS/CAES should lead to its sustainability in its environment.

Operational level subsystems – the identified work processes – are able to cooperate to some limited extent and may achieve some degree of local stability and synchrony by inter-process communications. However, this ability does not scale in very complex networks with many work processes that interact directly or indirectly.

Overall stability and synchrony is achieved by the level above, the coordination level. This level operates on a longer time scale than the real-time scale of the operations level. In fact, there may be several sub-levels within it each higher level operating on a longer time scale than the one below. Two kinds of coordination governance are observed. The first is the logistical coordination, which ensures the synchrony and stability of the internal working processes of the system. The second is the tactical coordination needed to synchronize and keep stable all operations interfacing with external entities and forces, for example the importing of resources,

[4] Another term that can be used is 'governance framework', which is an outline of what a governance structure should look like. A major difference between what we call a governance archetype and the way governance framework is used is the former is based on a formal model of governance that is applicable to all CAS/CAESs from living cells to the HSS as a whole, whereas the latter is often used in the context of a specific kind of organization and so may not necessarily be general. See the Wikipedia article: https://en.wikipedia.org/wiki/Governance_framework for background. Accessed 6/25/2018.

exporting of products, services, and wastes, and mitigating effects of disturbance within the range of adaptive response of the system.

Very long-term sustainability of the system calls for its ability to strategically modify itself (or be adventitiously modified) so as to adapt to major changes in the environment in order to remain fit for continuation. In systems that are evolvable they may either have a top level that engages in strategic decisions or they may be modified through Darwinian modification and selection.

This model relies heavily on cybernetic and control theories. We will describe these qualitatively. The details of various principles in controls and quantitative models were reviewed in Mobus & Kalton (2015) and where appropriate we will provide references to rigorous treatments.

The Concept of Governance

Every CAS/CAES relies on a hierarchical cybernetic governance subsystem to regulate its long-term behaviors such that it keeps producing products that keep it fit within its environment[5]. This is as true for biological systems as it is for manmade organizational systems like enterprises. The principles of such a system are covered in Mobus & Kalton (2015, Chapter 9). In this chapter, we will expand and extend the basic model provided there to see how it applies to the process of understanding complex systems. Governance and management systems are a consistent pattern of system archetypes that emerged, particularly, with the emergence of life (Coen, 2012; Morowitz, 2002, chapter 11). They are ubiquitous throughout nature and manmade systems. We have understood the core principle of feedback loops used in operational management for some time (Wiener, 1950; Beer, 1966, chapter 13). Researchers have recognized the need for a general concept of hierarchical cybernetics underlying the nature of more complex systems, especially human organizations that have to operate in a complex environment for very long periods (c.f. Beer, 1972, chapter 5).

In another domain, the author has explored how such a system, a brain, should be designed for a mobile, autonomous robot operating in a fairly complex and non-stationary world (Mobus, 1999; Mobus, 2000). The author simulated the brain of a simple animal (like a gastropod) to control a robot's search for positive reinforcing stimuli and avoid negative reinforcing stimuli. The robot had to learn cue stimuli that were causally associated with those it sought to go to or to avoid. The brain employed a model of an adaptive hierarchical cybernetic system for the management and control of behavior.

[5] A stronger claim is that every system needs an internal control subsystem in order to continue existence against an uncertain environment. Even atoms maintain themselves via feedback mechanisms, both in the nucleus (strong forces) and between nuclei and electrons in orbitals around the nucleus. The mutual attraction of positive protons and negative electrons keep the electrons bound unless a surge of input energy overly excites an electron (forming an ion).

The Purpose of Governance[6]

In the following chapter we will introduce the concept of a generic and universal model archetype of an economy. An economy is the aggregate work activities of a CAS/CAES which produces "assets" or highly organized, low entropy, material structures that are of use in the on-going sustenance of the system. We will see a number of examples of economies at different scales relative to human perception. But we will also see that, for example, the economy of the HSS is actually just an extension of, and wrapper around, the biological economy of population physiology, and that is an extension and wrapper around the individual economy of body physiology. And, in turn, that, it is argued, is just an extension of and wrapper around the metabolic economy of individual cells, which are the fundamental units of living systems. At all of these concentric scales an economy is just the interactions among numerous work processes, all operating to produce that which is needed to maintain the health and well-being of the whole system. The governance subsystem is that which ensures the smooth and efficient working of the whole economy for the benefit of all of its components. The existence of a governance system is to provide the needed constraints on the dynamics of component subsystems. Under the assumption that the supra-system is well designed (or in the biological world, evolved for fitness) the governance processes keep the system from taking excursions that would prove to be dysfunctional. These constraints need not be rigid (e.g. dictatorial). It can be advantageous to the whole system to allow some kinds of excursions in order to explore the adjacent space of possibilities. This will be discussed below.

All material/energy work processes are inherently unstable left to their own devices. They are subject to disturbances from outside as well as degradations of structures from within (entropy). At very minimum a work process will need maintenance and repair. They are also subject to fluctuations in the rate of inputs and the absorption capabilities of the sinks. All of this is just another way of saying there is no such thing as a perpetual motion machine! Even information work processes (computation and communications) suffer from entropic decay (parts wear out) and disruptions of the energy supply (power outages).

No complex system can enjoy a sustained existence, with a life time that is substantially longer than that of any of its components, unless it has the ability to detect problems with structures and functions, and the ability to make changes that will restore nominal function. Sustaining nominal functions in light of disturbances or disruptions in source/sink relations is the job of homeostatic mechanisms. Repair and replacement of components is the job of autopoietic mechanisms.

But on top of these basic regulation and maintenance requirements another problem poses itself. In any truly complex system, there will be many heterogeneous work processes that need

[6] In this chapter we will be primarily concerned with CAS/CAESs and not so much with CS (complex systems) or SS (simple systems). These later are governed by built in mechanisms (e.g. governors for dynamic systems) or constraints imposed by environmental forces. Our primary interest is with adaptive and evolvable systems.

to work in concert with each other to produce final products. The environmental interface processes, the importers and exporters, need to regulate their activities in concert with the entities in the environment, the sources and sinks, as well as considerations for changes in the general milieu. Mere *cooperation*, that is when several work processes must communicate with one another to make adjustments to flow rates to the benefit of the composite work, breaks down when the number and kinds of processes exceeds some limits. When that happens the only way that the composite work can maintain efficiency, effectiveness, and stability is to have a new layer of regulation for *coordination*. This is a set of new kinds of agents using much more complex decision models than is the case for mere homeostasis. The addition of a layer of coordination has associated costs; the overhead in energy consumed is proportional to the complexity being managed. But the benefits of coordination are great. The fitness of any CAS/CAES that has evolved a comprehensive coordination level of management is substantially greater than that for systems that fail to do so or have an inadequate coordination level.

The coordination level just described is responsible for managing the operations level work processes over a longer time scale than the real-time scale of the work processes themselves. This is because it is concerned with trends in performance deviations rather than the real-time error that an operations agent reacts to. Coordination operates to counter deviation trends rather than reacting to more-or-less instantaneous errors. This is a major factor in maintaining overall stability of the whole system.

Some CAESs with very sophisticated computing power are able to achieve a form of second order coordination with the environment over much longer time scales. They can produce models of the environment beyond just the entities with which the system ordinarily interacts. They can use those models to anticipate longer term changes in the environment that could have an impact on the system sometime in the future. This is the capacity to act strategically. The system must be evolvable in order to be able to modify itself to be able to handle whatever changes may ensue.

In this appendix we examine the architecture of a governance archetype model comprised of these three levels of management. We will situate agents within the decision nodes in the hierarchical network and we will outline the basic mechanisms involved in achieving the objective of governance.

Of course, the governance system that probably we are most concerned with is that of the human social system (HSS) introduced in Chapter 6 and further examined in Chapter 8. The history (and prehistory) of human societies is one of an evolution of governance from simple systems (e.g. tribes and groups) to larger communities (city states) to empires. That evolution has been progressively toward the architecture described in this appendix but is greatly complicated by the influence of cultural evolution, i.e. the coevolution of human knowledge, cultural practices, and technologies. The HSS is viewed as a CAES that is still evolving and seeking fitness in the larger context of the evolution of the Ecos itself. Various forms of governments have been ‘tried’ and generally found wanting. The single most important question we could ask is: Is it possible that a form of HSS governance could be instantiated that would make it truly sustainable in the Ecos? Our suspicion is that if the answer is yes, that the form of that governance will embody many of the features presented in this model archetype.

Definitions

There are a number of terms we should explain as they will be used in multiple contexts yet retain their basic meanings. Governance, management, and administration are used to designate different levels and scopes of responsibility and authority. Governance is the overall umbrella term covering all forms of regulation mechanisms. Management is used to describe decision types that need to be made when something isn't functioning properly within the scope of the manager. Administrators are those agents tasked with keeping normal operations going. They monitor operations and activate appropriate responses based on established policies and procedures (or naturally evolved operations in all living systems).

It will be important to first make some clear distinctions about the words that we are using. Governance (and government), management, and administration are too often conflated with the result of there being poor delineation of duties in human organizations. Natural governance systems sorted things out by virtue of the evolution of working models through natural selection. Some examples of the latter will be provided below. But in human organizations the situation is more muddied. An example of the problem might help explain.

The modern education system, especially in higher education, provides a good example of blurring the distinctions between the kinds of 'work' that needs to be accomplished in a well-functioning governance system. The situation is a result of the rapid evolution of the social demands put on the education system over the last century. Universities and colleges tend to be fairly conservative institutions that found themselves in a radically changing environment without the internal mechanisms for rapid adaptation. Initially the governance of colleges and universities was shared between a board of regents (like a board of directors) composed of prominent private citizens and a few senior professors, and the representative body of the (generally senior) faculty, a senate. The regents' role was merely to provide a degree of oversight in coupling the institution with its constituencies and only in the sense of making sure that the institution was fulfilling its strategic commitments. This task was not very complicated because the strategy of higher education had been worked out centuries before and had not changed substantially until around the end of World War I. Nor were there many management tasks (as defined below) needed by the faculty senate or departments. The pedagogy and curriculum, as well as the divisions of departments, had been well established so that the primary concern of universities was simply maintaining operations of teaching courses. The main form of governance, therefore, called for administration of policies and procedures. That is why those who take on management-sounding titles/roles, e.g. deans and provosts, in the education system today are called "administrators" and not managers.

Unfortunately, in the modern university, especially since the end of World War II, the mission and strategic situation has changed considerably. Universities are now expected to operate more like businesses and compete for students, faculty, and research grants and not just on the basis of providing high-quality education (especially for undergraduates) but on many fronts with many new constituencies/stake holders (Cuban, 1999; Geiger, 2004; Newman et al,, 2004; Rhodes, 2001). Public universities have to do this in the face of diminishing state funding meaning that they are faced with the problem of finding new sources of revenues. Now the strategic situation is incredibly complex and requires a whole new notion of governance and management that goes far beyond mere administration. Yet, sticking to historical methods, most "administrators" in academia come from the ranks of professors who ended up chairing a department and getting promoted from there to a deanship. They are not trained in management decision making but find themselves continuously having to make new sorts of tactical, logistical, and operational decisions in systems that are far from stable in their operations. University administration used to be just that, manned by administrators. Once the nature of higher education started to change from just being institutions of education to encompass research (with the incursion of monetary concerns), professional-like athletics, marketing, finance, and other commercial-like activities, those administrators were called upon to make decisions they were not particularly prepared to make. In other words, they were not effective agents for the kinds of decisions they were called upon to make[7]. The results have been very disturbing for the institution of higher education in the USA[8].

University administrators are very smart people – most have earned PhDs. But they have not studied management science or worked in management/governance environments (e.g. a profit-oriented business such as they are supposed to emulate) where they would have thus learned the difference between administration, management, and governance[9]. Thus, we see middle managers (e.g. deans) too often taking on operational level decisions when they should stick to logistics or tactical decisions. This distinction will be delved into in much greater detail in this appendix.

Figure 11.1 provides a graphic of the relation between governance, management, and administration as a hierarchy of responsibility.

---

[7] We should be cautious in painting with too broad a brush. There have been many very capable and thoughtful university administrators promoted from the ranks of the professoriate over the years and they have grappled nobly with problems of the modern academy. However, if we look at the state of the academy in the US, honestly, we might reasonably conclude that the majority of administrators never really achieve the capabilities of real strategic management.

[8] From various colleagues around the world we hear today that the same trend is diffusing around the globe.

[9] To be fair some education administrators have studied management concepts and know the vocabulary of terms, like "strategic planning." However, few have actually had to perform management functions in any kind of organization other than academia. Understanding of the deeper meaning of terms comes from experience, not from the mere application of logic.

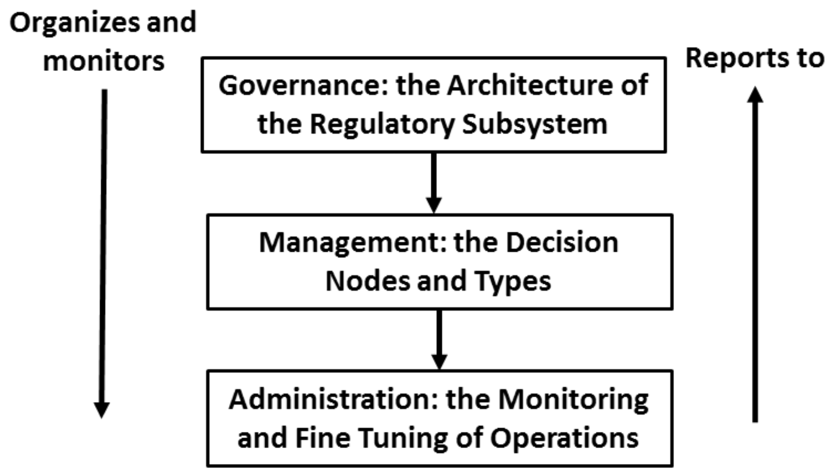

Fig. 11.1. The relations between the three primary functions involving decision agents constitute a hierarchy of responsibility.

Definitions of these should help in understanding the roles of each.

Governance: Establishing an Architecture that will Succeed

Governance is the process (subsystem) by which a CAS or CAES maintains its internal functions such that the whole system sustains existence for extended time, what Stafford Beer called a ‘viable system’ (Beer, 1972)[10]. Governance includes the structural specification of a government (architecture, e.g. the offices and their relations to one another depicted in a classical organization chart) and functional specifications of decision processes that occupy specific nodes in a hierarchical network (management duties). The lowest level of governance is the set of policies and procedures that constitute instructions for the on-going operations of the component work systems, those directly involved in obtaining resources and producing products. These are what need to be administered.

The governance architecture is the ultra-long-term plan for how the whole system is to be controlled, regulated, adjusted when needed, repaired when needed and so on. It designates the structures and regular functions of the system so that it fulfills its long-term mission of remaining viable within the environment in which it operates. It is the design for the ‘brain’ of the system in terms of the major subsystems and their functions. Within those subsystems operate decision processes – agents – that are responsible for managing those functions.

Governance needs to be established early in the evolution of systems. Those that do can expect to have long lives. Those that don’t will be short-lived with a possible tortuous existence. For naturally evolved organic systems this is guided by variation through genetic mutation and natural selection in brains and physiologies. For supra-organic systems, such as organizations

[10] See Stafford Beer’s Viable Systems Theory: the Wikipedia article: https://en.wikipedia.org/wiki/Viable_system_theory.

and states, the design of the architecture of governance is the result of a meta-governance process, such as developing a charter and bylaws for an organization or a constitution for a nation. The process can be seen to be a subsystem of the larger social system in which it is embedded. That governance architecture has to be designed such that the system is responsive to its environment (a CAS – think constitution of a nation), and in general is able to modify its own structure when the environment changes substantially (CAES as described below – think amendments to constitutions).

Once a governance architecture is in place, the on-going operations of internal subsystems and the interactions with the environment require decision making processes for a variety of purposes. That is, the various offices in the architecture require the emplacement of agents able to make veridical decisions and guide the activities of the system.

Additional terms that will be used in the discussions below are:

Authority and Responsibility

The term 'authority' shall be used to designate the agent that has the ability to effect a change in structures or functions lower in the governance hierarchy. It means the management of lower level managers.

Responsibility is a designation of what an agent's purview should be. The basic responsibility of any manager is to monitor the activities of lower-level processes, receiving reports (below) from those activities (measuring qualitative and/or quantitative attributes), exercising the decision model, and sending commands to the lower-level as needed. In other words, it is a reapplication of the basic cybernetic principle at a higher level of organization and scope (see section [Feedback] below).

Command

A command is any message sent from a higher-level manager to a lower-level manager in the governance hierarchy. The lower-level manager acts on the information content of the message to make a change in the process it controls. In natural systems commands are simply communications that arise in the course of a higher-order recognition that a lower-order process is not behaving as needed. In human organizations commands may arise from other than operational needs, motivated by a variety of emotional states. One reason that humans so often do not come close to 'ideal' agents (see discussion below) is that they can succumb to the 'power' relation, an evolutionary hold-over of the status structures in primate societies.

Report and Monitoring

Higher-level managers receive regular reports from lower-level managers summarizing the on-going operation of the lower-level process. The summaries are often in the form of time averaged measures of behaviors of the process. For example, a purchasing manager may receive a report from the receiving department quality control manager regarding the number of defective parts of each kind that were received during the month. The purchasing manager uses this information to monitor the behavior of the various parts suppliers and may need to take actions (such as switching to a different supplier) if the number of defects goes above a threshold.

Management: Decision Making Functions within the Governance Architecture

The word 'management' is used broadly to describe a number of activities involved in controlling the productive operations of an organization and its interfacing with its environment. In this appendix, we will refine the concept of management, and differentiate its activity from governance and administration, so that we can clarify a number of systems issues with respect to the overall success of a CAES. However, we should note that management is not strictly isolated from governance in CAESs that need to change some aspects of their own structures. Officers of a corporation, for example, need to act as managers of their offices, e.g. the Vice President of Finance may also manage the Controller's office, who, in turn, administers the accounting system. They may also be called upon to make what are strategic decisions, as for example when a chief financial officer decides to restructure a firm's investment portfolio for potentially better returns. Such decisions lead to possible changes in some aspects of governance architecture. For example, the CFO's decision may lead to the creation of a new portfolio management function[11].

It is very difficult to distinguish, at times, between a management activity and its relation to the governance of an organization, but such distinctions are important in order to provide clarity to large-scale systems analysis of such organizations.

Management is a sub-activity within the framework of the governance architecture and, itself, a framework for administration (defined below). Management, in the present context, is the process of *making decisions* involving adaptations of operations when the environment changes or when internal disruptions cause dysfunctions (management by exception[12]). Management is performed by decision agents with higher degrees of autonomy to assess local conditions and affect actions as required to restore normal operations as soon as possible. Here, 'normal operations' refers to the operational processes that have been defined *a priori* as needed for the overall accomplishment of the whole system's mission. Their subsystem status is defined within the overall system to be governed. *Management, then, is all of the decision activities directed at keeping a system subsystem operating within pre-established nominal ranges*. Managers are agents that have much more autonomy than administrators (see below). They are

[11] We realize that the most common use of the term 'governance' with respect to corporations or other organizations is used more narrowly to mean something like 'oversight.' But this is just a peculiar usage. Clearly, governance as applied to the running of a nation state or a municipality involves government where active decision making results in changes at the operational level. We use governance in a much more broad sense as described in the text.

[12] See the Wikipedia article: https://en.wikipedia.org/wiki/Management_by_exception for background. Accessed 12/17/2016.

problem recognizers and solvers and, so, require a great deal more intelligence and judgment than administrators (who might be characterized as homeostats – see below section [Classical Cybernetics]).

What is most crucial to the performance of good governance is that managers in every level and in every subsystem completely understand their function (decision types they are supposed to be making) and stick to the agenda. This is not particularly a problem in natural systems like cells and animal brains, except in the event of serious malfunctions or disease. However, it seems to be a problem in human organizations where the human agents all too often forget their function, or as in the case of the university administrator, don't completely understand what it was in the first place. And, with human agents there is the problem of personal agendas that can run counter to the objectives of the organization[13].

Administration: Keeping the Wheels Turning Smoothly

As long as conditions are nominal (as they are supposed to be under normal operations) administration is the process of monitoring and adjusting operations. An administrator's duty is to follow procedures to ensure the fulfillment of policies, not to deviate from those procedures. Should the administrator determine that a disturbance to the process has occurred in the face of following procedures, their responsibility is to refer this fact to the manager of the operational unit for correction (which may require deviation from procedures).

Administration, in natural (living) systems is homeostasis. An administrator is free to respond to error with the normal mechanisms available. The same principles of 'same-staying' apply to human societal systems at all scales. In human governance parlance this implies bureaucracy.

Administration is a more-or-less mechanical process with a minimal amount of discretionary decision processing (e.g. should I use a #2 pencil or ink?) or autonomy. An administrator is essentially a homeostat with a narrow range of reaction. It is the typical dotting of i's and crossing of t's. It is also the bane of bureaucracies when some tasks could use a bit more autonomy but the administrator chooses to fall back on the 'rules' to avoid having to take any

[13] In human organizations and governance processes there is a new emergent process owing its origins to the degree of individual autonomy held by the agents (human) – the political process. This is actually a dual level process that is a recurrent version of the governance/management process discussed in this section. The normal political process that corresponds to on-going management once a governance architecture is established involves the interpretation of policies and procedures – the various political parties, representing and holding different ideological interpretations of, say, the constitution of a nation – and the decision process for selecting which interpretation shall be used in the near term. There is a higher-order or meta-political process that involves decision making regarding the very architecture of the governance system itself. It too may be influenced by ideological positions but it results in the evolution (or sometimes revolution) of the governance process for a given human social system. The subject of political process is, unfortunately, beyond the scope of this appendix since it applies only to human social systems. But, clearly, it is an important subject to be explored. Our hope is that the methodology presented in this book might help in such an effort.

out-of-the-ordinary actions. A major problem in human governance is bureaucratic creep, or over reliance on rules when change is needed.

Confusion of the Roles in Human Governance Systems

Often the same human agent has to operate in all three roles, as governor, manager, and administrator in the course of operations but under varying conditions (e.g. the VP of Finance example above). This is certainly true for human agents in social systems, since human beings can and do present different personas depending on who else they are interacting with. But this same kind of role multiplexing can be found in a few instances in natural systems as well[14]. As a result, it is easy to become confused about what role is to be played in various situations. As described above, many higher education administrators have been unwittingly cast into all three roles without a clear understanding of the requirements and constraints of situations that demand specific actions (decisions or mere monitoring). Thus we see university presidents embroiled in logistical decisions simply because they do not really understand the difference between strategic-level decisions and logistical-level ones. Unfortunately, this phenomenon is not limited to higher education institutions. For human social systems in general, the complexity of modern systems has reached a point where human agents are easily confused about these matters. For another example, take the case of a representative in the national House of Representatives for a district in a particular state getting involved in the details of a pork-barrel issue (an operations level issue) when the intent of governance at this level is determining how the citizens of that district 'fit' into the national narrative of income distribution. The representative succumbs to trying to grab perceived immediate benefits for constituents rather than working on the larger national picture of economic well-being that would benefit those constituents in the long-run.

The cognitive confusion of roles and duties seems to be wide spread. Consider the plight of a department manager in an organization. Most of the time they are just administrating the operations of that department. Occasionally, something goes awry requiring them to make corrective decisions; they might be called upon to make tactical or logistical decisions. And then, very occasionally they are in the position of needing to make a strategic decision, perhaps observing an opportunity for their department to take on additional responsibilities within the organization.

This capacity for a 'manager' occupying a mid-level operation position to make both management and strategic decisions is due to the autonomy of the human brain. Every individual has the necessary 'equipment' in the brain (specifically the neocortex in the frontal and

[14] For example, our emerging understanding of the control networks in the genome give some examples of non-protein coding regions of DNA (what used to be called junk DNA) that code for regulatory strands of RNA (e.g. interference RNAs) in response to signals from the cytoplasm. Thus, the transcription machinery can serve a dual role of making coding RNA and, as needed, non-coding RNA. The complexities of this machinery are just starting to be elucidated but the notion of dual purpose of some proteins is now established. The same kind of phenomenon is found during development of an embryo where various proteins are used for different purposes at different stages of development.

prefrontal lobes) to make all three kinds of decisions for themselves (e.g. a strategic decision about what major to take in college) and that ability can be extended into the social domain so long as the complexity of the social organization is not too great.

However, the ability to switch between decision modes is not well understood. Most managers cross lines between operational, logistical, tactical, and strategic decisions without ever realizing what they have done. The result is considerable confusion in fulfilling roles and mistakes being made.

A major objective of this appendix is to sort out these roles in decision types and levels in a CAES in order to better understand the governance model archetype and how it applies in many kinds of CAESs. We are particularly interested, of course, in how the HSS is to be governed to achieve sustainability in the rapidly changing world.

Decision Agents in Governance Systems

We now situate the agent archetype model within the context of a governance system. In Chapter 6 we introduced the hierarchy of decision types relevant to the governance of a CAS/CAES (see Figure 9.3). Below we provide general explanations for the roles agents (Appendix C) take on in a governance system. Those roles are determined by the kinds or types of decisions that are need to be made in the various levels of the governance hierarchy.

The three levels of decision types are: 1) operational, 2) coordination, and 3) strategic. Here we briefly explain, reiterating what was covered in the last appendix. Below, in the section [Hierarchical Cybernetic Systems] they will be explained in detail.

Operational Level Decisions

Operational decisions are generally made in what we call real-time. They involve making local decisions that affect the operations of the local subsystem, specifically the material, energy, and message work processes (harking back to Chapter 2, figures 2.9, 2.10 and 2.11). The time scale for these decisions is the same as the rate scales for flows of material, energy, and messages into and out of these subsystems. Examples include the time scale for a ribosome to process the stream of amino acids being linked into a polypeptide (microseconds), the time scale for a liver to process glycogen and output glucose into the blood (milliseconds), or the time scale for a manufacturing work cell to produce a product (minutes to hours).

Every work process has at least a minimal decision agent that adjusts the operations to match the real-time flows being processed. This is the basic cybernetic feedback principle used to maintain the quality and quantity of desired outputs feedback (see section [Classical Cybernetics] below). However, in a network of processes forming supply chains (outputs of some processes are inputs to other processes) mere feedback is not enough to maintain the work flow.

Operational level decisions employ real-time feedback but also *cooperative* feedforward messages from other operating units. Multiple operating units often form a network of communications that permit cooperation to play a large role in the decision-making process. These networks operate based on well-defined, formal communications protocols, at least in systems that are well regulated at the operational level. For example, the concentration of various amino acid-tRNA complexes in the neighborhood of a ribosome signal the latter as to the availability of those needed to fulfill the mRNA specification being processed. The liver is signaled through several channels about the need for glucose in the bloodstream. And a work cell foreman sends parts requisitions to the parts inventory kiting function.

Cooperation through feedforward signals and response feedback is an efficient mode to achieve smooth operations, but it has limits based on the time delays inherent in communications channels and the number of 'hops' a message might need to take. Here a hop means a message or some variant on it must pass through a node (a neighboring work process) before being passed on to the next node. Then processing delays in the neighbor further delay the signal before it is sent to the next node in the network. Propagation delays, channel bandwidth, possible noise, and other classical communications 'problems' act to restrict cooperation signals to only a few nearby nodes. Thus, analyzing a complex system for the extent of cooperation and its effectiveness should take these issues into account. In complex human organizations the dysfunctions that occur when there is an over-reliance on cooperation are legion. What typically happens is the formation of informal communications networks that depend on personalities and non-standard messages with protocols essentially absent.

When cooperation becomes impractical due to the scale of the whole operations level, the number of work processes, it is necessary to introduce a new level of management, the duties of which are to make decisions that help coordinate the activities of a cluster of operations level subsystems.

Coordination Level Decisions

Coordination decisions constitute the next higher level in the hierarchy and are needed to keep multiple operational units working in concert with one another when the cooperation mechanisms discussed above cannot provide optimal operations. This is often the case when, for example, a supply chain gets long and response times for downstream processes can be negatively affected by variations in upstream processes. Coordination is achieved by giving commands to operational level decision agents that modify their decision models or change an operational parameter (e.g. the set-point in Figure 11.5 below), at least temporarily in order to, for example, change the rate of processing of a local process. Coordinators seek to balance processing so as to achieve a larger optimum production.

There are actually two kinds of coordination decisions. Internal operational units need to be coordinated since the outputs of some of those units are inputs to other units (see section [Cooperation and Coordination] and Figure 11.7 below). The other kind of coordination decisions involve operations that coordinate the activities of the whole system with the entities in its environment with which it interacts, the interface processes for importing to and exporting from the SOI. The first kind of decision types belong to the category of *logistical* coordination. The second kind are called *tactical* decisions.

Coordination decisions involve monitoring various units in the coordinator's purview over a longer time scale relative to operational level cooperation time scales (real-time with some small lag added). Logistical decisions generally seek to optimize overall operations on this longer time scale. Tactical decisions are also made over longer time scales. They involve monitoring the internal subsystems that are engaged in acquiring resources or exporting products and wastes as well as observing the behaviors of the external agents (sources and sinks) with respect to acquisitions and exports. Recall from Chapter 3 that sources and sinks are generally not modelled in the same way that the SOI is. Tactical decisions, therefore, involve considerably more uncertainty than do logistical decisions. Tactical decision agents often use more complex decision models of those entities that are essentially opaque-box. They also generally collect more information on the behavior of those entities.

Strategic Level Decisions

Finally, at the highest level in the hierarchy decisions may need to be made involving the very structure of the whole system itself in view of very long-term changes that might take place in the environment. These are strategic decisions that can and need to be made by evolvable systems. In some sense these are meta-decisions or decisions about what sorts of decisions need to be made in the future. Not all CAES can make explicit strategic decisions. For example, a biological species (or genera) does not make strategic decisions to achieve evolutionary change. Normal neo-Darwinian evolution makes those decisions (about fitness) through mutation/variation and selection. But human beings are, by virtue of their learning neocortices, evolvable systems that are faced with making such decisions (what career do I want? What kind of mate do I want?) These decisions are based on far more complex decision models, memories, and a more powerful computational engine than can now be accomplished with electronic computers. The time scales for strategic decisions are far greater than mere coordination decisions. The spatial scope in terms of external entities is far greater than for tactical decisions. For example, such decisions often consider possible disturbances that the system should take into account when deciding how to reorganize.

As the logistical decision agents influence operations level decision agents when needed and tactical decision agents influence environmental interface operations when needed, strategic decision agents influence tactical and logistical agents through something we could call a "plan" for future activities. Both logistical and tactical coordinators, as part of their decision models, can be programmed, so to speak, to take certain actions (and thus influence operations) at certain points in time, in a longer time scale, given certain external conditions obtain. This will be expanded below in section [HCS.Controls and Controlled Systems.Strategic Level].

The HCGS model to be presented below maps these decision types onto the hierarchical control model, for example of Findeisen, et. al. (1980). Whereas the latter is a theoretical model of temporal domains for decisions – higher levels in the hierarchy operate over increasingly longer time scales – the former maps the kinds of decisions to be made onto the temporal domains. Each higher-level works, essentially, on time-averaged values from the level below it and issues commands downward when needed.

Considering Ideal Agents

In Appendix C we introduced the notion of an ideal agent as one having a sufficiently veridical decision model, sufficient sensory inputs, and sufficient actuator outputs that it will make satisfactory decisions the majority of the time (statistically). This notion of an ideal is not of a perfect agent, obviously. This is because the world is fundamentally stochastic and no agent can have absolute information or knowledge or processing capacity to make perfect decisions. What makes an agent ideal is that it does not inject its own non-relevant agenda into the decision-making process. We would not want our thermostat to 'feel' cold and decide to turn on the furnace for its own comfort.

The governance agents of evolved systems achieve ideal-hood by virtue of natural selection. Here is where we butt into a fundamental problem for human agents. The author suspects that human evolution has not yet produced a mental capability to approach the ideal agent with sufficient autonomy and veridical decision models. Humans are largely guided by beliefs and a background of ignorance about how the world actually works. We can hope that the progress of scientifically-derived knowledge about the world, and the guidance of systems science, will help inform human beliefs such that they will move closer to the role of ideal agents.

While we can hope for improvements in the decision-making capabilities of individual human minds due to the knowledge that science and systems science has to offer (i.e. the evolvability of the human brain), it may yet prove to be the case that human beings, as a species, will need to evolve further. As suggested in Appendix A this might involve a further advancement of the prefrontal cortex. But that is speculation.

Hierarchical Cybernetic Governance System (HCGS)

Energy Requirements per Level

In each section below, on the various levels in the hierarchy, we will include a short discussion regarding the energy requirements devoted to the management in each. Recall from the previous appendix that making decisions requires energy for computation, memory storage, and communications. Actuators require power. There is a proportionality between the amount of computation and other work done by an agent per unit time and the power requirements, that is, the free energy consumed per unit of time, for the various decision types and levels in the hierarchy.

The energy requirements at each level are different because the higher you go into the hierarchy, that is from operations up to strategic, the more energy each agent requires because the decision models get more complex, hence require more power.

As we will discuss in the next appendix, there is a relation between money and energy (the former being a token marker for the latter). Hence there is a simple rule that follows from the energy requirements relation: The more complex a whole system is, the more levels and sub-levels required to manage, the more taxes need to be levied to accomplish the goals of governance. That isn't what many citizens want to hear, perhaps. But it is the way things work.

It should be noted, however, that this subject is pregnant with research opportunities. Not much has been documented about energy costs associated with various agent and decision types. All we do know is that in human organizations executives get paid outrageous salaries while line supervisors barely make a living wage. Whether that is a true reflection of the energy consumed making strategic decisions or an anomaly of human thinking is an open question.

Administration of Work Processes - Classical Cybernetics

Norbert Wiener (1961) noted that regulation of any dynamical process subject to disturbances from the environment required a periodic measurement of error between the current status of a system's path toward a goal state, and the ideal path. The amount of error would then be used to correct the path to reduce the error to zero. This is the famous negative feedback mechanism that constitutes control of a process to achieve its goal state. All CAS/CAESs are based on mechanisms that implement this principle. He termed this cybernetic. It is the basis of all management and governance processes.

Feedback

Cybernetics is the science of control and coordination in systems (von Foerster, 2003; Wiener, 1961). The basic theory of cybernetics is the use of feedback information to cause a system to modify its activities in order to maintain an output function in a viable or nominal value range in the face of disturbances that might otherwise cause the output to deviate from a desired value. The system is goal-maintaining in this sense. Figure 11.2 shows a basic control system and the principle of negative feedback used to counter whatever deviation from the goal state. The 'Product' in this sense could be a physical product (specific chemical, specific object) or a directed motion or force. This system is described more completely in Mobus & Kalton (2015, Chapter 9). The agent in this case is an administrator as described above. The "policies and procedures" are completely embodied in the control model with the "set point" embodying the current management decision.

The principle is that a value (either quality or quantity) of the output (product) is a result of the work process operating normally. If a disturbance to the process causes the value to vary from an ideal, as represented by the 'set point' constant, either higher or lower, an error signal is generated and fed back to the computational engine which uses the control model. This information is used to generate a control signal that activates an actuator (e.g. a motor) that changes the internal operations of the work process in opposition to the error.

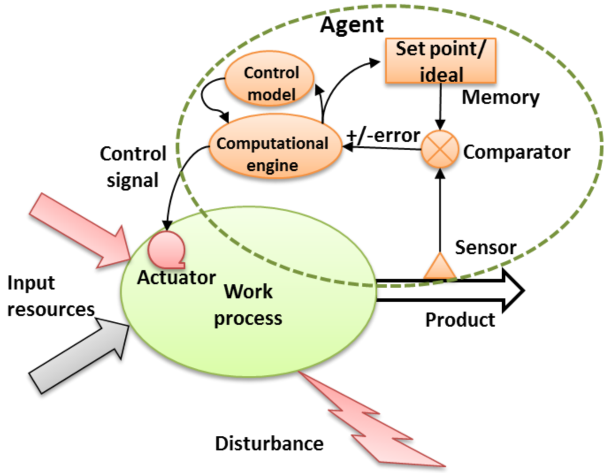

Fig. 11.2. The fundamental cybernetic system operates to restore the proper function of a work process in light of some (manageable) disturbance. The agent in this figure is essentially the same structure as depicted in Figure 11.1. The circuit works by comparing a measure of the product (quality or quantity) with an ideal value (set point) and the consequent error (information if the error is non-zero) is used by the control model to generate a control signal to an actuator inside the work process. (Modified from Mobus & Kalton, 2015, Figure 9.5, with permission.)

This is the basic principle of cybernetics, maintaining nominal operations in spite of disturbances. The actuator must be able to affect the operation in a range sufficient to counter the normal range of variation caused by the normal disturbances encountered by the system[15]. We will borrow the term ‘homeostat’ from Ashby (1958) who designed and built a mechanical/electrical device that demonstrated the ability to counter disturbances giving it that name[16]. We wish to appropriate the name for any mechanism, chemical, electrical, mechanical, or human-based, that fulfills this principle[17], and is acting as a process administrator.

Requisite Variety

Every agent must have a range of actions that can be taken to restore nominal activity to the process being controlled. The “law of requisite variety” proposed by Ashby claimed that the stability of a system’s operations required that the range of actions must be equal to or greater than the deviations from normal operations that disturbances might cause. A homeostat must be

[15] This is often referred to as the “law of requisite variety” formulated by W. Ross Ashby (1958). See the Wikipedia article: https://en.wikipedia.org/wiki/Variety_(cybernetics)#The_law_of_requisite_variety for background. Accessed 5/28/2018.

[16] See the Wikipedia article: https://en.wikipedia.org/wiki/Homeostat for background. Accessed 6/8/2018.

[17] See the Wikipedia article: https://en.wikipedia.org/wiki/Homeostat for background. Accessed 6/8/2018.

able to generate countering actions to the disturbance or the system could not effectively remain stable[18]. Stafford Beer also addressed this in the context of management in human organizations.

A homeostat must be possessed of sufficient range based on the ‘ordinary’ range of disturbances that could affect the system operations. In naturally evolved systems this is determined over the course of a species collective experiences and selection for those representatives that have a capacity that lets them survive the marginal excursions. In human designed systems we try to use data about environmental conditions to estimate what that range should be. In many instances of engineered systems, we design to more than expected variances just to be safe. We have to specify the boundary conditions somehow.

Homeostasis, as described for a wide array of CAS/CAESs, by many authors, is the mode of administration. This is the pre-programmed capacity to handle normal ranges of disturbance.

Management

Autopoiesis and adaptive response mechanisms as mentioned in Chapter 9 [section Adaptivity] go beyond simple homeostasis. They provide a higher level of adaptivity in which a homeostat might be modified or adjusted so that it can respond over a greater range of disturbances over a longer time scale. In our vocabulary this is the role of management over administration. CASs have built-in capabilities to extend or enhance responses to changes as needed. As described in Chapter 9 when a disturbance or critical parameter is pushed out of nominal range or is kept at an extreme for extended time, the system responds with making modifications that help keep it viable even in the face of these higher stresses. Living systems such as cells and multicellular individual organisms have evolved elaborate management of resources such that extended capabilities are called upon only when absolutely needed. This is a cost savings strategy.

Variations on this basic model include the use of feedforward signals (to be explained below), for example measuring the value of inputs as compared with ideal values and using this information in the control model. Some forms of anticipatory control are obtained from such an arrangement, but the control model is made more complex as a result. For engineered systems where the variations in conditions can be calculated in advance it is often possible to use feedback alone, but for natural systems anticipation is generally necessary in order to avoid damage.

Figure 11.3 is a more general model that will be used for agents in an HCGS structure.

[18] This result should not be confused with overreaction which can lead to increasingly wild oscillations that are also a decline of stability.

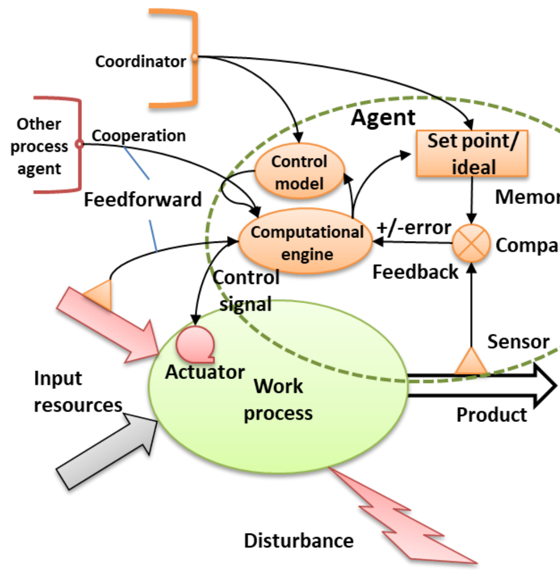

Fig. 11.3. This is the same basic management model as in Figure 11.2, but now with feedforward and 'feed-down' signals. One feedforward signal is the result of sensing the input factor. The other is sent to the agent from another process agent providing information about resource quality and/or quantity for cooperation. The feed-down signals are shown coming from a coordinator. One of these signals might direct changes to be made to the control model. The other is a simpler modification of the set point (as when someone changes the desired temperature in a thermostat).

The basic cybernetic principle discussed above is the basis for the lowest level of management and administration in complex systems. All processes within a CAS/CAES are managed in this way. Recall that all processes do some kind of work, which may be transforming materials or messages. Thus, all complex processes will involve some kind of basic management, including processes that are part of the governance system itself.

Figure 11.3 includes message feedforward and feed-downward signals. One of the feedforward signals comes from measuring an input. Together with the feedback from the output these constitute all of the local process control. The other feedforward signal is shown coming from another entity, another process that is sending information (either as a customer or a supplier) about its situation. The feed-downward signals are shown coming from a 'coordinator'. Both of these additional sources of information will be explained below.

Operations Level Networks – Cooperation

Work processes do not exist independently in complex systems. They form a complex network of processes that generally work in concert to accomplish the goals of the whole system. The next appendix on a model archetype of the economy will drive this point home. The network of processes that do the major transformation work of the system are at the operations level

In such a complex network of processes we will often find local clusters of nodes that are more tightly linked together. That is, they are found to be near neighbors in space and linked by direct flows[19]. For optimal operations these nodes need to coordinate their activities. For example, a supplying process may need to notify the receiving process of fluctuations or interruptions in the flow of what it is supplying. The latter must have a more complex control model that takes this information into account and has the ability to regulate some additional actuators internally to compensate. This is usually accomplished in the receiving process by it having a stock reservoir in which it can buffer the flows, storing excess when the flows are higher than required and drawing down on the supply when the flows are lower than required.

Operational networks with clusters of this sort are directly related to what, in the economy model archetype (Appendix E, next), we call a ‘market’, with the provisos stipulated there that market mechanisms work best in tightly coupled clusters (small worlds). As we will explain, the market is generally also a hierarchically organized small world model with nodes in localized clusters having the strongest coupling with the most frequent interactions and clusters acting as nodes in larger (fractal-like) clusters over longer distances and lower coupling strengths. Local clusters generally operate through a node that acts as a hub and is responsible for the longer distance interactions with other cluster hubs. See that appendix for a more detailed account and Figure 11.9 below for a graphic representation.

Basic Organization

The problem can be characterized thusly: A complex system is comprised of many interacting work processes, each under local basic control, but each also subject to disturbances that might affect downstream receivers of upstream products. If each of the processes is monitored and controlled by a local controller (as in Figure 11.2 above) then the whole system may generally be expected to operate in nominal form. Local disturbances should have minimal impact on overall system operations. Figure 11.4 shows a typical organization of basic work processes at level 1 in a complex system. These processes, must work in consonance with one another, each doing what it is ‘supposed’ to do to satisfy the needs of downstream processes or customer sinks. What is shown in the figure constitutes the ‘operational’ level of the system and the processes are controlled by local real-time agents.

[19] An exception to nodes needing to be near neighbors in space is the case of information processors linked by high speed communications flows.

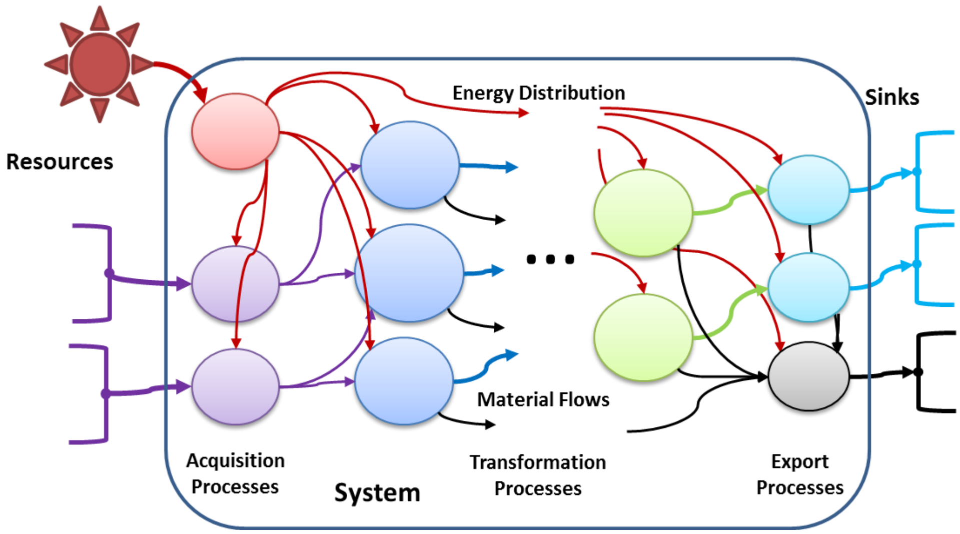

Fig. 11.4. Every CAS or CAES is composed of many internal subsystems that perform work processes, transforming materials and energies. Acquisition or import processes bring resources into the system while export processes push material out to the environment sinks (e.g. customers and garbage dumps). All processes use energy in a high potential form (red arrows) and produce waste heat (not shown). Generally speaking, the internal flows may be very complex depending on the number and kinds of processes.

We start from the assumption that a system has evolved or was designed to have this organization. The central concern is how all of the various processes are to be kept working together in a 'more-or-less' optimal fashion. Moreover, the various acquisition and export processes need to operate in consonance with the sources and sinks.

The problem of maintaining system-wide optimal or near optimal operations becomes increasingly difficult as the system increases in complexity. Both the exposure to and timing of disturbances become increasingly difficult to manage by feedback control alone. For example, should two (or more) disturbances co-occur with sufficient magnitudes, the combined impact could overwhelm a downstream process's ability to compensate.

In all CASs and CAESs that have been examined in detail, we find three fundamental control mechanisms operating at this level. The first is, as mentioned, the fundamental negative feedback mechanism of cybernetics. A local agent (e.g. as in Figure 11.5) monitors the performance of the work process and generates control adjustment signals to the internal actuators as needed to keep the process producing at optimum (given the constraints). An example from the metabolic processes in a living cell would be the way in which the synthesis of ATP from ADP and a free phosphate molecule with energy supplied by the oxidation of glucose (through a complex set of reactions and intermediates) is regulated by the local concentration of the various intermediate molecules within the matrix of the mitochondrion. As more ATP is needed by the cell, more of these intermediates are produced to up-regulate the production in the citric acid cycle. When less is needed, fewer of the intermediates are produced to down-regulate the process[20].

In human organizations, the role of the familiar 'line supervisor' or 'foreman' is to make sure the work being done under their watch is being done correctly. They take action to correct variations before there are significant problems. In a real sense every human worker is a local agent with respect to the work that they perform.

The other two mechanisms are also familiar. They are ways to achieve non-local cooperation between work processes and they employ messaging.

Cooperation Between Near Neighbors

There are several options for establishing cooperation, especially between near neighbors. The simplest mechanism is shown in Figure 11.5, assuming agents as depicted in Figure 11.3 above (e.g. with feedforward signals). Here, Process **A** produces a product that is used by Process **B**, which, in turn, produces a product used by a downstream customer. Both processes contain more complex decision agents that are in direct communications. The problem here is that both agents have to share a common language as well as protocols for interacting based on their own states. Assuming that such protocols have been worked out and encoded into their control/decision models[21], the agents inform one another of variances in their situations that allow them to take some kind of actuation response to help correct the problem.

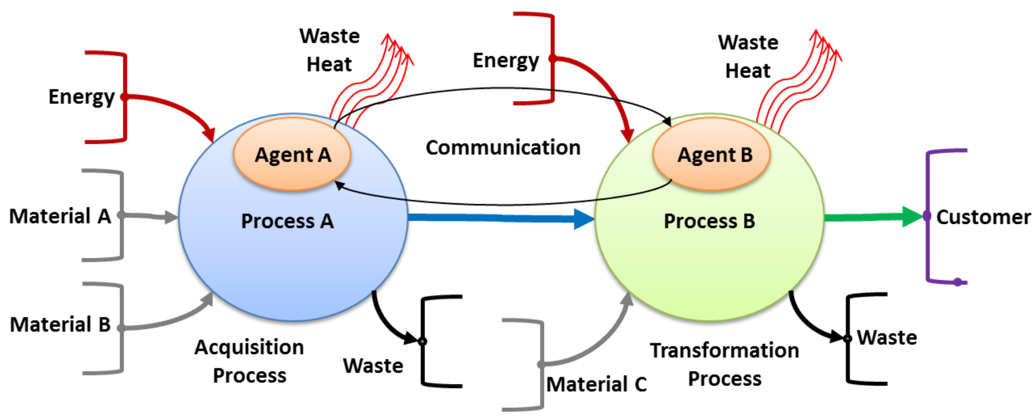

Fig. 11.5. Two systems (processes) cooperate when their management agents exchange messages that help regulate the flows from the producer to the consumer process (blue arrow).

Suppose Process **A** is disrupted by an interruption of flow of Material **A**. Agent **A** can notify Agent **B** of the situation (that is provide feedforward messages to Agent **B**) by virtue of

[20] See the Wikipedia article on the citric acid cycle: https://en.wikipedia.org/wiki/Citric_acid_cycle for some background on this process. Accessed: 01/09/2017.

[21] In real systems this is actually a result of auto-organization and emergence processes as covered in Mobus & Kalton (2015). In designed systems such as organizations the communication protocols have been a priori designed. However, it should be noted that most such designs originated in a evolutionary process of essentially trial and error.

monitoring its own status receiving Material **A**. The agent is projecting a disruption to the flow of its intermediate product to Process **B**. Agent **B** can then use this information to make internal adjustments, for example it might have internal stocks of the intermediate product (stored in anticipation of such a disruption) that it can use in order that it does not experience a diminishment of its production. Of course, these kinds of compensation actions are time and duration sensitive. But this is an example how, if the two relatively independent agents are communicating with the appropriate message protocols they can cooperate to smooth out the overall operation and the ultimate customer is not (immediately) affected.

Figure 11.6 provides a more complicated situation, three of the processes extracted from Figure 11.4, are shown in a more complete cooperative *network*, combining elements of the basic cybernetic model in Figure 11.3 with the communications shown in Figure 11.4 above. This depiction shows that agents in various work processes will need to communicate with multiple other agents. In this case an acquisition process supplies two (or more) internal work processes (the figure only shows the agent/work relation in one process). The flow rates to each of the receiving processes could be very different with different timing. There can even be priorities that must be accounted for in the case of emergency conditions.

Clearly there are going to be complexities involved in providing control models that can adequately provide correct responses in all of the agents involved in these local connections.

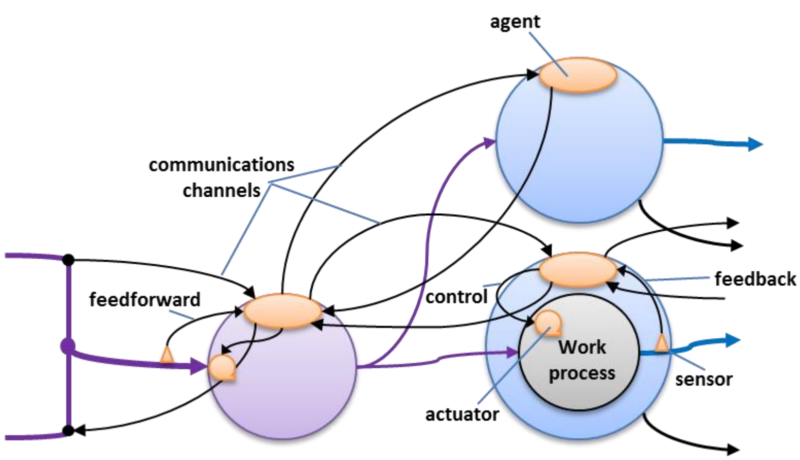

Fig. 11.6. More details of the communications channels needed to provide internal control (lower process on the right) and cooperative messages between processes. The acquisition process (purple on the left) also shows the use of sensing inputs to provide feedforward information regarding the flow of the input. Additionally, the acquisition process may communicate with the environmental source (or just be able to observe what the source is doing).

As should be becoming clear from these diagrams the nature of communications between cooperating agents and the local control models used by each become increasingly complex just to obtain cooperative interactions. After taking a look at two near neighbor forms of cooperation we will revisit the complexities involved in achieving cooperation in the whole network (such as in Figure 11.4).

Sub-processes within systems comprised of many cooperating sub-processes do not attempt to communicate with every other process. The overhead for communications would soon overwhelm any possible advantage. Moreover, the further the distance between sub-processes would mean the greater the time delay distortions that would occur in the messages, further aggravating the problem of cooperation. Instead, successful systems employ a hierarchy of coordination controllers whose job it is to maintain a modicum of 'local' cooperation while also providing a level of more global coordination.

Cooperation by Messages

In figures 11.5 and 11.6 we show messages communicated between the internal agents making operational decisions for the work process in which they are embedded. Not shown, but necessary, are the message interfaces. Recall from Chapter 3 that interfaces involve protocols for sending and receiving, and this is the case whether what is flowing is material, energy, or messages. In general, for processes that have a tight coupling as shown in those figures messages can be exchanged to help coordinate the activities of the work processes so long as there are no major systemic problems.

Cooperation by Exchanges

When systems are in a formative stage of development (e.g. during the origin of life) the first form of cooperation that is established is the exchange of substances, material and/or energy. Such exchanges are the result of fortuitous auto-organization among aggregates of heterogeneous processes (Mobus & Kalton, 2015, Chapter 10, Section 10.3). This is depicted in Figure 11.7 below.

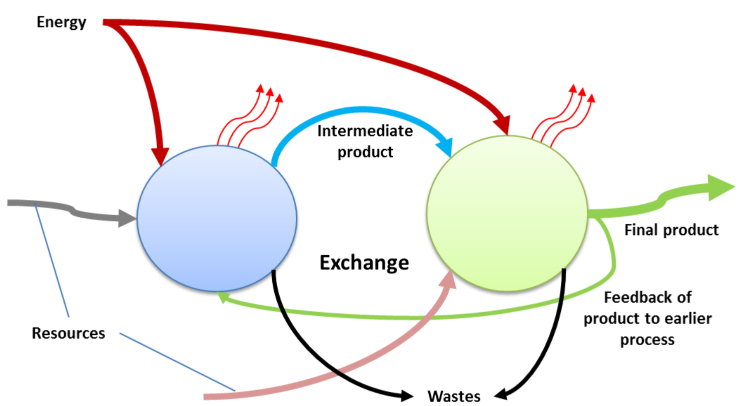

Fig. 11.7. Processes may establish cooperative relations by exchanging substances that are useful to each. In this example the blue process produces an intermediate product that is needed by the green process. The latter produces a 'final' product (so far as these processes are concerned) some of which is sent back to the blue process that might be using that product for consumption purposes. As long as the two processes cooperate they will maintain an optimal overall production of the final product.

If one or both of the processes is evolvable (and from the first primitive auto-catalytic cycles active in the origin of life are thought to be), over time the exchanges might be complemented by message flows to help coordinate the exchanges.

Feedforward

The fundamental cybernetic feedback control in Figure 11.2 can be augmented with the inclusion of feedforward signals that provide anticipatory information. Feedforward can be obtained from simply monitoring the inputs to have advanced warning that some disruption might occur in the output (affecting the feedback error) in the near future, with a time delay based on the transit time through the process. Using an anticipatory model plus, generally, the use of internal buffers to supply the substance of the input in case of diminished flow or absorb excess flow, the operations controller can maintain optimal output in spite of these kinds of disruptions.

In the context of cooperation, however, the feedforward signal comes directly from the supplying process. In Figure 11.5 processes **A** and **B** are communicating with one another directly. **B** is supplying **A** with a form of feedback and **A** is supplying **B** with a form of feedforward[22]. There must exist a protocol for message interpretations in both processes. Both messages convey error values. The information content (value) conveys the significance of the message, while the message is about the flow state (or dynamic). For example **A** may signal **B** that a lower rate of flow is about to be experienced. Or, another example, **B** may signal **A** to increase the flow rate if **B**'s stock of the substance is running low. In either case, neither **B** or **A** are necessarily expecting those messages – they are thus informational. And they relate to the flow rates of the substance going from **A** to **B** at some near future time.

Feedback and feedforward are thus complementary messages that, with the right protocol in place, can act to coordinate behaviors without an explicit command from above to do so.

Decision Models Required

The decision model required by simple feedback mechanisms is relatively straightforward (once time delays are taken into consideration). For example for general analog controls the proportional, integral, derivative (PID) control algorithm suffices nicely (Mobus & Kalton, 2015, Chapter 9, section 9.4.2). However, once you incorporate feedforward information the decision model becomes much less straightforward. The basic PID controller needs to be modified with additional decision functions. The Adaptrode model discussed in Appendix A is one example of such modifications that incorporate adaptive response based on anticipatory signals (feedforward). Other examples in natural systems will be provided below.

---

[22] These conventions are the result of considering the flow of the substance in question going from the left (**A**) to the right (**B**). Thus, **B** is feeding back information to **A** regarding what it has received and **A** is feeding forward information regarding what **B** is about to receive.

In human organizations the use of feedforward from upstream processes is quite well known. For example, a parts inventory management system can ‘tell’ a production kitting manager how many of a certain part will be available in stock. The kitting function can then use that information to schedule the withdrawal from stock to make up the production kit. As will be seen in the next appendix (the Economy) this capability is the essence of a supply chain process.

Anticipatory feedforward is relatively easy to incorporate in the decision models agents use in cooperative clusters. However, the more complex a cluster of cooperating processes grows, at some point a limit in the effectiveness of cooperative signaling is reached. At that point cooperation alone cannot provide solutions to optimal behavior.

Complexity and Breakdown of Cooperation

Near neighbors in a network may be able to achieve stable, coordinated operations using homeostatic and augmented decision models that take the other entities into account (Chapter 1, Principle 9). Cooperation depends on mutual respect for protocols (and no malicious agents sending false messages).

There are several additional concerns in cooperative situations. The message exchanges are reciprocal and there is danger of establishing a positive feedback loop in which one agent, responding to an abnormal condition signaled by the other agent, inadvertently amplifies the abnormality and feeds that back to the first agent with the consequences that cooperation breaks down.

Even if near neighbors can successfully cooperate over a wide range of operating conditions, both may simultaneously be neighbors to other processes, such as the situation depicted in Figure 11.6 that will be affected by their interactions (attempts at cooperation with one of the two processes) in a way that is not readily handled by the recipient. Figure 11.8 hints at the problem. Process **B** receives a message from process A informing it that a disturbance will be felt in the flow of intermediate product from **A** to **B**. **B** must process this message and then inform process **C** that, with an appropriate time delay, its own output to **C** will be disturbed. In theory this allows **C** to prepare (preadapt) in advance.

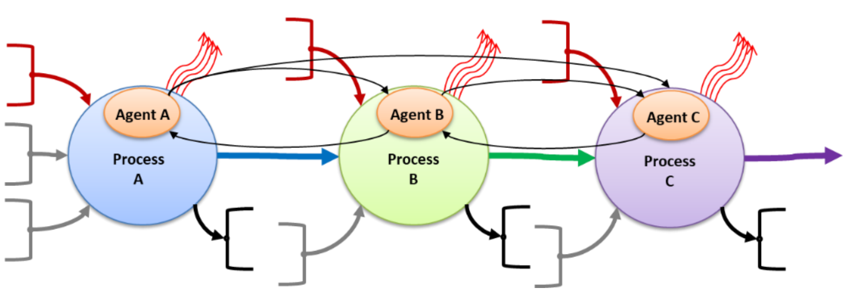

Fig. 11.8. A communication channel can be extended to nodes further distant in the network. This arrangement is meant to mitigate propagation delays in transmitting information to more distant nodes through intermediate ones that have to process their messages before they can transmit to their neighbors further downstream.

However, there is a processing delay in **B** that shortens the time that **C** will have to prepare. Worse yet, upon receiving the message from **B**, **C** may start preempting just as the disturbance is actually reaching it. In a worst-case scenario this could lead to overcompensation with ensuing instabilities or oscillations that could damage the process. What Figure 11.8 suggests is that **A** could communicate directly with **C** so that the propagation delay could be avoided. This is certainly technically feasible, but now imagine complexifying this with the need to communicate with many more downstream recipients. Or imagine every process in Figure 11.4 communicating with every other process, including feedback along long-distance channels (called a peer-to-peer network architecture). In theory this should lead to the maximum in cooperation, but at what cost?

The actual technical problems with designing a cooperation-only architecture, as this might be called, are many and beyond the scope of this appendix. Two important ones will be mentioned to give the reader an idea why this scheme cannot work in practice and even, likely, in theory. First there are still communications delay problems even when the channels are set up in a direct peer-to-peer (P2P) network. Moreover, the longer the length of the channel the higher the probability of interference or noise injection that would possibly distort the signal and lead to false messages. The longer the channel (physically) the more complex any communications protocol would have to be, including error detection and re-send, and security concerns.

Second, the decision models used by the agents (who are doing the communicating) are suddenly much more complex in order to send the right messages to the right other agents. In human designed systems, today, it is possible to mitigate the communications problem through the use of an internet protocol (in fact the TCP/IP and protocol suite could be used directly[23]). This is exactly what is being contemplated in the design of social-technical hybrid systems under the general name “Internet of Things” (IoT) in which human-human, human-machine, and machine-machine communications are facilitated by the Internet (especially including wireless).

However, the admonition against over complexification (even if it can be done cheaply) still stands. If you look carefully at all of the P2P or even client-server applications you will find excessive complexities owing to the “desire” that these applications take all possible decision nodes into account – made more difficult by having so many connections. As a general result of this level of complexity foisted on each agent, there are usually many software bugs lurking in the code waiting for just the right decision condition to occur to wreak havoc with the application. How many times has the reader had to reboot their home router, or a running application on the computer or smart phone? Since software engineering is still somewhat in its infancy (from the standpoint of systems theory) as more complexity is added to communications/computation systems we can expect a multiplicative increase in such bugs. In

[23] See the Wikipedia article: https://en.wikipedia.org/wiki/Internet_Protocol for explanation and background. Accessed 6/17/2018.

all operations-level process networks it is neither feasible nor even desirable to use P2P communications and complex decision models to achieve cooperation.

Natural governance systems have evolved a different approach.

Figure 11.9 provides an abstract version of a set of subsystems within a larger system. This map derives from a second level decomposition of an SOI into its subsystems and then those, into their sub-subsystems. As shown in Chapter 5 this is one of the methods for verifying selections of boundaries when no physical one is found.

Note that there is a detectable difference between any node taken at random, and groups of nodes that cluster. That is there are more (and sometimes stronger) links between small groups of nodes than between those groups and other nodes. Intuitively we realize that the nodes with a higher density of links are more tightly coupled and must be working together more so than with other such groups or independent nodes[24]. In a concrete system context analysis of the dynamics of work and flows into and out of as well as among the members will identify the validity of considering them as a group.

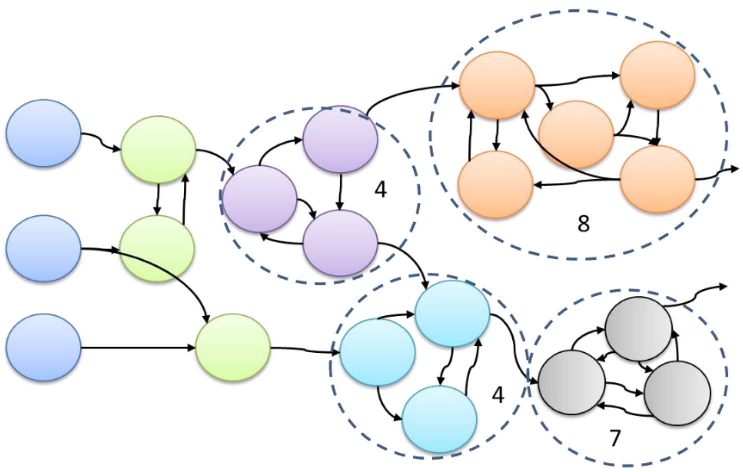

Fig. 11.9. Processes within systems tend to form clusters or groups. In fact, this is just another reflection of the notion of subsystems within the larger system. Each of the blue, dashed ovals contain a group of processes that are tightly linked (counts of the number of links between neighbors can be used to indicate clustering). The implication of this kind of clustering is that members of a group are better able to cooperate since they are close together in process space. These groups generally correspond with subsystems.

It turns out that cooperation between members of any such sub-group is more feasible than between groups or over the entire system, so long as the clusters do not extend too far. Groups of four or five processes may be able to achieve a reasonable level of cooperation within themselves. But then we reach a limit.

[24] In graph theory groups of vertices with more edges between the members of the subgraph than with other vertices in the rest of the graph are called cliques and there are algorithms for detecting their presence. Strictly speaking this applies to undirected graphs, but variations on directed graphs, as in Figure D.9, are known.

However, we need to keep in mind that a map such as in Figure 11.9 only represents one level of decomposition. Each process shown in the figure could, in fact, be further decomposed if it is itself a complex system. Thus, it is reasonable to conclude that P2P schemes for communications in the interest of a 'democratic'-style cooperation-only governance architecture is really not workable. Throughout the world of naturally evolved CAS/CAES (and this includes human societies) cooperative interactions are possible and desirable insofar as they are efficient for getting local groups working together near optimally. But the problems associated with long-distance cooperation presented above will defeat success for any large and very complex system.

Therefore, another approach, suggested by the kind of topology of real systems shown in Figure 11.9, emerges. When, in a group or cluster, one of the nodes takes on the role of an agent whose task it is to coordinate the cluster the problem of long-distance cooperation is elevated to a new level in the cybernetic hierarchy.

Energy Requirements for Operations Level Administration/Management

The energy consumed by operations level agents is relatively small owing to the simplicity of the decision model using error feedback to adjust the operational parameters. Of course the more complex each operational unit is, then the complexity of the decision model increases and the agent power requirements increase accordingly. A surrogate for energy costs in a corporate operations level unit is the number of clerks, data collectors, etc. are needed to support the supervisor. This is analogous to the number of different enzymes and intermediate products being catalyzed and the total amount of ATP consumed in anabolic metabolism. More complex pathways take more energy.

Additional energy costs come from the degree of cooperation being employed at the operations level. Cooperation increases the complexity of each agent's decision model since they now have to take into consideration the direct communications and interpretation of messages that increase the load on computation.

Coordination Level Decision Managers

Ultimately complex systems cannot rely on simple cooperation to achieve optimal performance. Moreover, they cannot always use cooperation when it comes to working with external entities in the environment. By definition a system does not have adequate access to the internal operations of an environmental entity – a source or a sink – and therefore cannot necessarily have the kind of communications channels necessary to work with such entities[25].

Coordination is a form of inter-process control that arises when the communications situation has become sufficiently complex and the channel delays are long such that reliance on

[25] The possibility is not, of course, prohibited. But neither is it a given possibility. A corporation may have a very deep communications capability with its suppliers but will not necessarily know what is going on inside its major competitor.

direct communications are no longer reliable. An attribute of communications between distant process agents is that the time scale over which the dynamics are affected is, itself, long. The change in an output of process **A** may not be impactful on process **Z** for an extended period. Such time delays are disruptive. And it should be clear that if process **A** is supposed to notify **Z** of its situation, the problem of complexity is deeply aggravated. **A** would have to communicate not only with **Z** but every process from **B** to **Z** and every one in between. The overhead of computation and communications is untenable.

Instead of trying to extend the nature of cooperation to all of the sub-processes in a complex system, the natural tendency is to introduce a higher level cybernetic processor, a coordinator, whose job it is to intervene in the normal feedforward and feedback loops of process controllers to provide coordination. Figure 11.10 introduces the implementation of a coordinator agent whose responsibilities are to monitor the activities of a cluster of closely associated work processes and provide them with longer-term guidance so as to optimize their collective behaviors with respect to the final output of products (blue arrows) going to down-stream processes. The depiction in Figure 11.10 is of a coordinator that is working to balance the flows of materials imported by the purple process to two (or more) work processes that must share that material for their independent processes. As can be seen in this depiction the role of communications and agent decision-making is starting to dominate the internal flows that have to be understood.

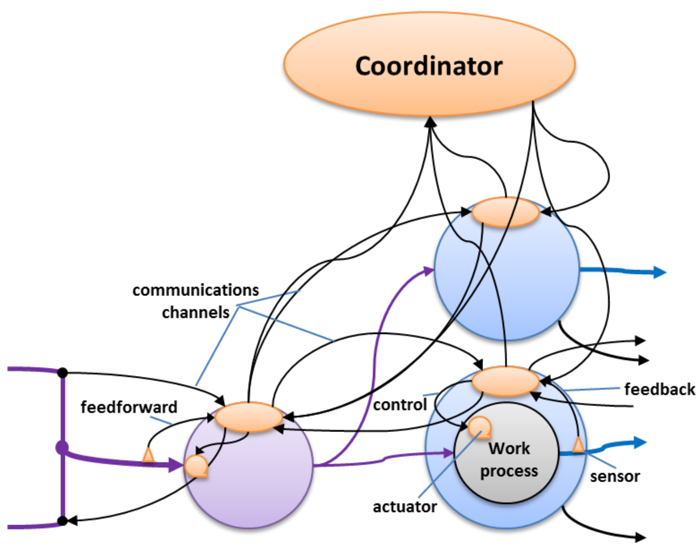

Fig. 11.10. At some level of complexity and dynamics subsystems are no longer able to merely cooperate and need a higher-level perspective coordinator agent to keep them operating at optimal performance.

Figure 11.10 includes some communications channels between the three work processes representing the cooperation channels. To the degree that the messages sent along these channels can assist cooperative decisions on the part of the three agents, then the complexity of the decision models of the coordinator is reduced (the coordinator does not have to tell each process what to do all the time!) In such a system, a hybrid of cooperation and coordination, the computational load on all agents is kept manageable.

Coordination at Larger Scales

The clustering of near neighbors (in figures 11.9 and 11.10) that allow local cooperative behaviors to facilitate the cluster behavior as a subsystem, and only requiring coordination in unusual circumstances, is a prototype for a more general approach to coordination management as shown in Figure 11.11.

Just as a sufficiently large number of neighbors in a cluster warrant a coordination level of management, so too, a large number of subsystems will require a higher order (and longer time scale) coordination. This is achieved by a new level of coordination above the first level of coordinators at the subsystem level.

The figure depicts two (out of many) work subsystems (clusters) that are sufficiently complex as to require some form of coordination level management. The tactical and logistical coordinators in each subsystem have models of the whole subsystems' behavior given inputs and outputs as well as the models of each work processes' time-averaged behavior (see discussion of these models below) by which they provide feed-downward messages.

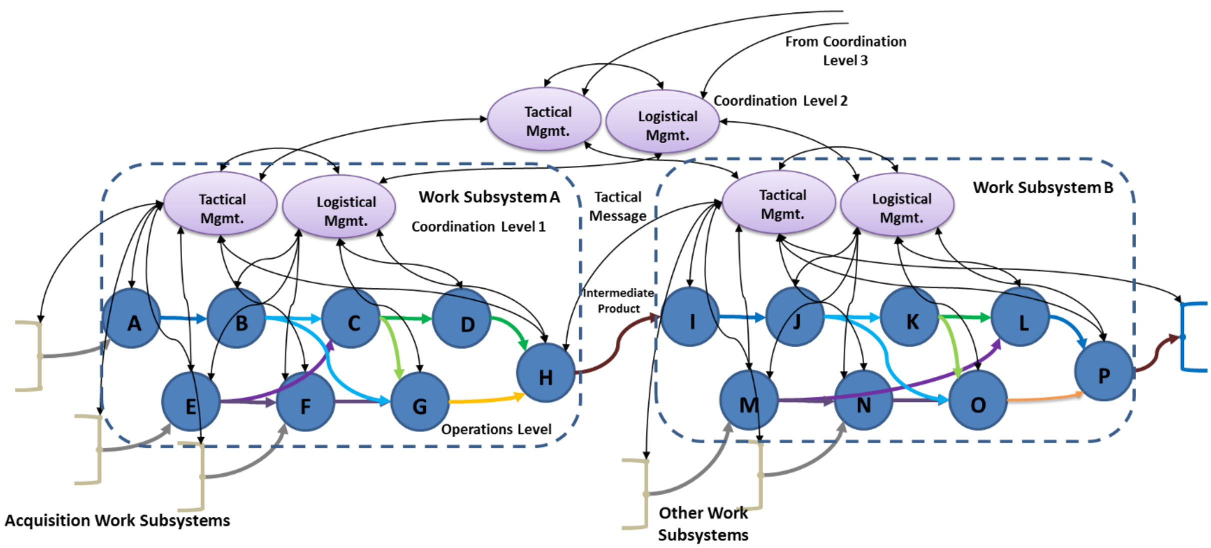

Fig. 11.11. Coordinated clusters of work that correspond with our idea of a subsystem. The output of subsystem **A** is input to subsystem **B** (along with other work subsystems in the network). Each subsystem has its own local coordination level and the two subsystems (and possibly several more) are themselves coordinated by a sub-level of coordination with the larger scope. Blue ovals are operations level work processes. Purple ovals are coordination level processes. Coordination level 1 is local to the subsystems. Coordination level 2 works to coordinate the coordinators in level 1. If the whole system is sufficiently complex a coordination level 3 would be responsible for coordinating the level 2 coordinators, and so on.

However, the introduction of multiple subsystems creates a need for an additional level (super-level) of coordination, that between subsystem coordinators. In the figure we only show two subsystems in order to convey the basic concept. In a real complex system there will be many subsystems and all of their coordinators will need to be in communications with one another. Not unexpectedly, this architecture is the same as cooperation between work processes, but now between coordination processors. We will revisit this situation below in [Coordination between Logistical and Tactical] below. Note for now, that all logistical and tactical coordinators at the same level have communications between them for cooperation purposes. The tactical coordinator at the subsystem level reports to the tactical coordinator at the next level up and the same is the case for the logistical coordinators.

This architecture forms a pyramid (hierarchy) of coordinators from the operations level upward that is as deep as is needed in order to facilitate coordination across the span of the whole system. The architecture is fundamentally that of so-called 'middle-management', with the number of sub-levels dependent on the number of operations subsystem overall – roughly the number of levels is $log_s(m)$, where $s \geq 2$ is the "span of control" and m is the number of subsystems.

Span of control is medium dependent, i.e. the kind of *thing* the system is, and is affected by transmission delays, bandwidth, computing power, and decision model complexity. A generalization can however be made. This architecture can be shown to reduce the overall overhead of management and lead to maximal stability of the system behavior if, and only if, the subsystem sizes and complexity are within the range of capabilities of the coordinators. This is a recursive feature that extends upward to however many sublevels are required. In naturally evolved systems we can see this principle in operation in examining the details of actual near neighbor work processes and their coordination controls (to be discussed below).

Work of the Coordinators

Monitoring Longer-term Behavior

A coordinator is strictly an information processor[26]. In other words, it is basically just a decision agent but with a much more comprehensive decision model and more computational power than operational level agents possess. A coordinator has to monitor the behavior of all of the operations level processes under its ken. This is generally done using time-averaged or integral quantitative methods. The operations level processes (their agents) can 'report' data to the coordinator on intervals longer than the real-time scales on which they operate. In the section [Examples of the HCGS in Natural Systems] we provide examples of this phenomenon in the coordination of synaptic compartments (operations level work processes). The coordinator must collect data from all of its subject work processes using the same time constants so that it can efficiently compute errors in any of them in any particular time frame.

[26] By information processor we still mean that the process does work but now only on messages (data flows) for the purpose of extracting information.

Decision Models

Coordination models are more complex than operations level ones. At base a coordination model has, as its components, the operation level subsystems' models of themselves in a coupled array of models. These are, of course, also time-averaged and/or integral terms summing up the longer-term behavior of each component subsystem and comparing those with the array of ideals for each. The algorithms employed are some form of optimization over the entire array in which deviations from individual ideals are translated into an overall performance metric for the group. From there, the coordinator has a simpler time computing an appropriate response for the overall group and mapping that back to actions that individual units can take.

Another class of decision models that are becoming better understood and 'more popular' in human-built systems is pattern recognition and learned responses. Artificial neural networks with Bayesian learning capabilities are being employed increasingly in all levels of real human governance/management systems and will be further discussed below. Of course, the premier such system is the human brain. We will also reflect on its capabilities for governance of the individual and small social groups below.

Additional modeling approaches found in both natural and man-made systems include finite state machines (e.g. models used to regulate biochemical interactions in metabolism), hidden Markov chains (used to approximate models of sources and sinks), and other statistical methods. Since all of these modeling methods have been thoroughly explored elsewhere in great detail, we will not belabor their uses here.

Feed-downward Messages

Coordinators have two basic ways to influence the operations level agents. They can alter the value of set points (in Figure 11.2) or they can alter the decision models (programs) of the operations level agents (see Figure 11.11). Altering a set point simply causes the operations agent to adjust its actuator commands up or down in accordance with the new set point. This is easily seen in the thermostat model. By setting the thermostat desired temperature the system's 'seeking' behavior is altered and a new 'optimum' is established.

Altering the operations level agent's decision model is more complicated. There are a number of ways in which a PID model, for example, can be altered to produce a new mapping from input sets to output sets. The simplest modifications would be to change one of the formula constants. For example consider the simple output proportionality, $o = af(i)$, where $\boldsymbol{o}$ is the output signal and $\boldsymbol{i}$ is the input signal, $\boldsymbol{a}$ is the constant of proportionality. Simply changing the value of $\boldsymbol{a}$ will change the mapping from $\boldsymbol{i}$ to $\boldsymbol{o}$. Of course, the coordinator must have a model of the behavioral change in order to know by how much to change the value of $\boldsymbol{a}$. The realm of adaptive control theory provides the basis for analyzing and designing such systems.

We can now differentiate between the two major categories of coordination control in the HCGS, logistical and tactical coordination. Logistical coordination involves getting internal subsystems to work together so as to produce a near optimal set of behaviors, thus a near optimal overall system behavior. It involves getting the right stuff to the right place at the right time and at the right cost. Tactical coordination is the tricky business of coordinating the embedding system's behavior with that of the entities in the environment of the system with which it must interact.

Energy Requirements for Coordination Agents

The power requirements vary among agents in the coordination level (and its sub-levels). But in general the wider the scope of management, e.g. a second-order logistics coordinator managing a number of first-order logistics managers, the more complex the decision model and the more computational, communications, and memory capacities will be required. We will discuss the differences between logistical and tactical agent requirements in each discussion below.

Logistical Management

Logistics is a familiar term in human organizations and especially the military. In general it means getting the right things to the right places at the right times and for the right (optimal) prices (costs). It is generally viewed as problem in optimization over a complex operation with requirements and constraints shaping what constitutes an optimal or near-optimal solution. But, as we shall see, even natural systems display elements of logistical coordination management with respect to coordinating internal flows of materials, intermediate products, and getting final products ready for export.

Logistics is one of the more mature management sciences with extensive mathematical models used to solve the problems of complex operations (e.g. operations research was developed in the early and mid-twentieth century coming out of the military needs during World War II to move supplies and troops around efficiently in an ever changing war topography. Since the science of logistics is treated well in many other places, just as with control theory, we will not attempt to rehash that subject here. Rather we will mention just a few aspects of logistical management that are relevant across the spectrum of CAS/CAES models. These will be revisited in the next appendix regarding the management of economies.

Time Scales

Logistical decisions operate on a longer time scale than the real-time scale of the operations level as a rule. If operations level behaviors are measured in, for example, minutes, then logistical decisions will be made using data that is time averaged over hours or even days. There are a number of different kinds of logistical decisions, which we will discuss in the next section. Some types might require quickness, say in emergency situations or to respond to an ad hoc condition, such as in adaptive response. But by-and-large logistical coordination depends on observations of the multiple operations that need to be coordinated over time to compute, for example, trends.

Decision Models

As stated above, logistics is an area of management science that has been studied intensively and there are well understood models that can be adapted to particular instances. We find these models already in place in natural CAS/CAESs. Those systems achieved sustainability through the evolution of implementations. Below we will examine just a few examples of logistical processes in some of these systems.

Strategic requirements

The whole system is supposed to do something. It is supposed to interact with its environment in such a way as to be 'fit' in that environment. It is supposed to produce actions or products that serve the larger embedding supra-system to achieve this. The structure/function organization of an existing system was the product of a strategic process. In the case of systems up to humans and their organizations, strategic 'decisions' were (are) the result of evolution. By definition, evolution of species leads to modifications that either do make the species fit or result in negative selection. That a species (and its populations) exists is *prima facie* evidence that evolution has found a solution (or set of solutions) to that species' strategic problem. We will revisit the concept of *intentional* strategic management, as it applies to human individuals and organizations in section 11.x.x.x [Strategic] below.

The strategic requirements are those that motivate logistic behavior. In the case of CASs they have been incorporated into the structures/functions of the system and so the highest level of strategic management is automatically doing the 'right' thing. In CAESs, however, the situation is different. There can be an explicit strategic level of the HCGS which sends commands to the upper echelon of the logistics agents causing them to alter their decision models and thus causing long-term changes in internal processes. For example, an individual human can, at a young age, decide they want to go to college and they change their devotion to study in order to get good grades. The upper echelon logistics managers affected are those involved with directing the whole human's behavior with respect to extracting useful information from the study materials that the tactical agents have obtained.

The numerous homeostatic mechanisms that respond to changes in conditions have limits within which they can act. Those limits are set by strategic requirements. For example, when a species of animal evolves an ability to survive in a shifted temperature regime from that of the parent species, this is a case of a strategic move that results in setting new limits for temperature tolerance.

Moving stuff – Coordinating sources and sinks

The most general form of logistical decisions involve the transport of material through the system. Once a decision is made that an intermediate product should be sent from subsystem **A** to subsystem **B**, along with volume/weight and timing considerations, the logistical manager must compute the best way to make this happen in the context of the rest of the on-going operations[27]. The logistic manager then needs to monitor the behavior of the flow over time and analyze statistics regarding that flow to determine if it is operating correctly. If not, the manager has to, perhaps, adjust the set points of the two process to bring the flow into the correct regime. Or the manager might need to modify the decision models of either or both of the operations agents.

Budgeting allocations

Logistical managers are responsible for making decisions regarding how much of a resource should be allocated to each process. Energy is a general example of something that needs budgeting based on how much work any given process is expected to accomplish per unit of time.

Optimizing

An on-going problem in most logistical decision models is how to find optimal or balanced solutions subject to numerous constraints and requirements.

Commands

Logistical managers issue messages (generically called “commands”) to the operations level work processes to get them to change their behavior. There are two basic kinds of commands, one changes one or more parameter set point to nudge the process toward slightly different rates of production. The other message involves making changes to the operations level agent’s decision model itself. These are, understandably, more complex sorts of messages with a more complex interface protocol.

Energy Requirements of Logistical Agents

Often, logistical agents are distributed according to the type of decision models as described above. A transportation logistics agent need not be concerned directly with budgeting issues and vice versa. Both may have optimization problems to solve, however. The division of labor across different logistics problems, however, means that any one agent has a relatively simple decision model to compute. Logistics problems are often elaborations on the basic cybernetic model in that they deal with managing errors and so their mappings from inputs to outputs are computationally tractable with modest computing power.

However, coordinating multiple logistics coordinators generally involves an increase in complexity that is at least multiplicative if not exponential. Indeed, the pressure that directs the implementation of higher-order coordination sub-levels, is exactly the same that has driven the creation of sub-process formations throughout hierarchical structures. At some boundary of complexity and scale, it is cheaper to abstract to a set of subsystems and introduce a new higher

[27] This could be done to modify the system given strategic decisions to do so, or it might be from a need to adapt to temporary conditions.

level of control. The energy requirements for this new level (or sub-level) is greater than the sum of energy requirements for all of the subsystems. It is overhead, the cost of increasing complexity. In natural evolution, the cost is offset by the rewards of stability, resilience, and sustainability of the whole system. In human societies it may help sustain profits (i.e. creation of wealth).

As we will discuss below, in general, logistical energy costs are modest and manageable in comparison with tactical management.

Tactical Management

Tactical decisions are somewhat similar to logistical decisions except that one of the two cooperating entities is outside the system and not exactly under the control of the manager the way the logistics manager has control over both (or multiple) operations entities. This means the tactical manager needs more information processing capabilities with respect to building some minimal but sufficient model of the external entity (following Principle 9 in Chapter 1)[28]. They also need some sensory apparatuses set up to observe the behavior of the external entity. In some more advanced forms of entity-entity cooperation, communications channels very similar to those between cooperating work processes might be established to facilitate the interface. A great example of this from the natural world is the communications between flowering plants and pollinators in the form of flowers that signal the presence and location of nectar. The plants know what will attract a bee and the bee knows how to read the message in the flower.

Tactical managers are responsible for situating the whole system in a sustaining relation within its environment. They are responsible for the importing of all resources needed by the internal work processes and for exporting products and wastes to the environment. They manage all of the work processes that are associated with in- and out-flows but do so in cooperation with logistics managers (shown in Figure 11.11 as two-way message arrow between logistic and tactical agents). The tactical manager is responsible for the flows coming into or out of the system through these interface processes, but the logistics managers are responsible for the flows out of the import processes into the internal work processes and the flows into the export processes from the internal work processes.

Higher order tactical managers may be tasked with searching the environment for resources. Search methods depend on the substrate system, i.e. animals forage, humans seek information (informavores), a parts-purchasing manager looks for alternate suppliers with better prices. There is, surprisingly an underlying algorithm to all stochastic searches, that is searches for a resource that might be found in seemingly random distributions in space and time (e.g. food patches like

[28] This is actually also the case when the system must coordinate with the milieu of the environment and not just an identifiable other entity. For example, an animal monitors the weather conditions and seeks shelter when it turns cold or rainy.

flowers for bee foragers). Non-random searches are used in more structured environments. We mention a few of these below.

Time and Space Scales

As shown in Figure 11.11, tactical management agents are generally leveled based on the scales of time and space. Here, the space scale is considerably different since the tactical decisions involve ranges outside of the boundary of the whole system that vary considerably. The time scales relevant to the import or export processes must be in concert with the real-time scales on which the internal work processes operate. But the time scales relevant to the entities in the environment must also be considered. Internally, a system might enjoy a certain amount of synchronicity through logistic management, but externally the world is asynchronous and, quite often, seemingly haphazard. Even quasi-periodic phenomena such as hourly temperatures over 24 hours may be chaotic. Real-world events are, in general, "sporadic (recurring but not necessarily periodic), episodic (lasting for variable amounts of time) and erratic (variance in amplitude within an episode)" (Mobus, 1999).

Thus, tactical management is generally more complex than logistic management. Moreover, tactical managers need significantly more sophisticated sensory and actuation apparatuses. They must be able to recognize the existence of a resource source or product customer against a background of the co-containing environmental background. And, of course, they must be able to do so at distances much further away than are the case for logistics managers that are close to the work processes they coordinate.

Decision Models

Tactical agents require decision models that are considerably more complex than those of logistical agents. The models are used to compute a 'plan' of action, that is a set of actuation signals that produce sequenced behaviors that are in concert with the entities and milieu of the environment. Such plans can be as simple as a single actuator command issued directly in response to the model being used, or as complex as a set of actions to be carried out over extended time scales. For example, an earthworm exercises a repetitive sequence of undulations as it crawls through the soil. It responds to sensory inputs that detect food concentration gradients to change course. There is no planning as such. A predator, on the other hand, needs to build a plan of attack that is based on the current conditions of the environment and the prey. We will examine plans below.

Strategic requirements

As with logistic agents, commands that condition tactical agents are strategic in origin. And as with logistic agents, the strategic decisions were made via Darwinian evolution of the species for natural CASs (as individuals) and speciation for natural CAESs. For humans and human societies (organizations) strategic decisions are achieved through intentional modification. That is, the decisions are volitional, a matter of choice by a strategic thinker. Whether that thinker is good at strategic thinking is another matter, of course.

Strategic commands are in the form of goals set for the system to achieve, e.g. find food. The hunger signal is built into the natural system by evolution. The tactical agents involved then go into action to fulfill the need for food. For a human deciding they want to go to college, the goal can be achieved by studying hard and making good grades. The tactical agents include all of the facilities that student has for studying and finding information.

Plans

A *plan* is a sequence of steps that *should* be executed in order to achieve a goal. Plans are not perfect algorithms. Rather they represent a sequence of actions that are the best approximation of actions that, if successful will lead to the goal designated. But as Robert Burns, the 16th century poet noted, “The best laid schemes o' mice an' men / Gang aft agley (often go awry)[29].”

Tactical agents work according to a plan of action. This plan is generated from a set of plan steps known to the tactical agent. It is assembled in response to commands from the strategic management level regarding a sequence of goals to be achieved. In a simple case, a predator feels hungry and its tactical brain modules start formulating a sequence of actions that will lead it to find prey. In a more complex situation a strategic decision to lower costs of the manufacture of a product will lead purchasing managers to look for lower priced parts. A plan of action is assembled (perhaps modeled and simulated to test for reasonableness) as a sequence of actions to follow in order to achieve the goal. See the discussion of strategic goals below.

Minimal models of environmental entities and milieu

Tactical agents need to have some ability to compute expectations regarding the behaviors of entities in the environment with which they interact. This is analogous to logistics managers having models of all of the operational level processes that they manage. The problem is that, by definition, the tactical agent has no way to gain insight into the internal operations of those entities – they are outside the system. The same observation obtains for the milieu which is relevant to some tactical agents (e.g. those tasked with keeping a mobile system in a proper temperature domain). Thus tactical agents are required to develop or otherwise have a minimal model of the entity’s behavior over time in order to make predictions or anticipations regarding that behavior in the future. This is a hard problem and a computationally complex one. Thus tactical agents have need of much more sophisticated decision models and more computational power to run them. Indeed, the decision models must include a simulation of the entity given observations of other environmental conditions.

---

[29] Robert Burns, To a Mouse (1785), see the Wikipedia article: https://en.wikipedia.org/wiki/To_a_Mouse. Accessed 6/25/2018.

Thus, a tactical agent is more challenged with making a decision than is a logistical agent simply because the environment is more complicated than is the internal milieu. Minimal models, that is, models sufficient to make reasonable anticipations, will depend entirely on the complexities of the environment and the entities/milieu with which the system must interact.

A minimal model for a tactical agent would be able to capture enough of the entity's behavior over time such that the agent could make reasonable predictions about what the entity might do in the near-term future (long-term behavior is the subject of strategic management covered below). Tactical agents may, of course, have more than just observational communications established with the environmental entity. A parts purchasing manager may have the ability to query a parts supplier regarding the availability and cost of certain parts needed as inputs. But this is not quite the same as having a causal model of the entity in question. All that a tactical agent can do is recognize that an entity exists, that it has behaved in such-and-such a way in the past, correlated with what it sends in messages, and that is all. Unless we invoke industrial spying, there is no way to gain access to the details of its internal operations, or a systems model of how it works. The tactical agent has to rely on the proposition that its minimal model of the entity is reasonably predictive of that entity's future behavior.

Pattern recognition

Tactical agents are more often required to observe and interpret patterns of state or activities of entities in the environment. Thus, their decision models are more likely to require some capacity for learning patterns to form a mapping from those patterns to appropriate actions by the interface processes they command. Pattern mapping learning is an area of very intense research today. The field of deep learning using artificial neural networks is extremely active and finding applications in many cyberphysical systems such as autonomous vehicles. Animals, of course, rely on biological neural networks that are capable of more than just forming a mapping. In humans, at least, we know that they can formulate exploratory mappings autonomously.

Once a tactical agent has learned a pattern and a mapping from that pattern to an appropriate behavior plan it is competent to manage the system's behavior resources to accomplish the goal.

Actuation capabilities - working through import/export work processes

The tactical agent commands one or more import/export work processes that receive or push out materials, energies, or messages. Since these work processes manage flows (with perhaps some minimal transformation processing) the main focus is on the operation of the interfaces they have with the external environment. These process will tend to depend on flow rate monitoring as the basic operational control, i.e. flows are at their required levels with errors derived from higher or lower than ideal flow rates (and at specified flow times). The tactical manager, in a similar fashion to the logistical manager's influence over internal work processes, can influence the import/export processes by adjusting set points for flows. However, since generally there is also a reservoir or stock buffer involved in managing flows, another variable that might be resettable is the stock level maximum. The tactical agent has to compute what the long-term ideal level of a buffer should be based on past flow activities (and overflows or underflows of the stock) and issue commands to the import/export processes to adjust accordingly (assuming they have adaptive or evolvable capabilities).

The set point adjustments for flow rates assumes that the import/export work processes have some adaptive capacity for ramping up or down the actuators associated with the flows. Pump motors might be sped up or slowed down. More pore channels might be inserted into a cell membrane to allow greater influx or outflux of molecules and ions. A household may have to increase or decrease the number of trips to the grocery store to stock the refrigerator.

In the event that the equipment built into the system for ordinary flow rates is inadequate to handle short-term fluctuations, then if the system is evolvable, the tactical manager may call upon the modification mechanisms to increase the flow actuators. Perhaps a more powerful pump is needed. The tactical manager requests the plant management department (in charge of maintenance and modifications) to install one.

Search

One of the more important decision models for higher level tactical managers is the search for satisfactory situations relative to the environment. This includes the search for satisfactory resources, e.g. a plant having access to good soil, water, and sunlight (accomplished by seed dispersal mechanisms) or a foraging animal's search for food, the search for protection (a fund manager seeking a good hedge position), or the search for a suitable place to dispose of wastes.

There are two general classes of search algorithms that tactical agents can use depending on the behavior of the target of their search and the environment through which the search must proceed. The first, and most general, is called stochastic search. This is where the target resource (or other situation) is distributed in space and time in a quasi-random fashion (i.e. the sporadic, episodic, and erratic characterization given above). In other words, the agent does not have a priori knowledge of where to find the resource and must conduct a search that will produce results in a reasonable time, i.e. before the system is depleted of its resource stores. A foraging animal needs to find food before it starves. Mobus (1999) developed a model of the kind of search that is not a random walk but a semi-directed quasi-randomized walk he termed "the drunken sailor walk." Implemented in a wandering robot it looks very similar to the path taken by searching scout ants as they walk across a relatively smooth surface. It weaves back and forth in the same sporadic, episodic, and erratic fashion, yet remains going in a general direction.

The second class of searches are ordered (i.e. algorithmic). When a search is being conducted in an organized structure, such as a list or a street (looking for a name or an address respectively) there are rules to follow that are guaranteed to produce a result if the item is in the list or the address is on that street. Search algorithms for networks (graphs) and many other topologically ordered structures are well known. However, it should be noted that the structured environment had to be constructed in the first place, before the search can be done efficiently.

For example, a list might first be sorted by lexicographical order. Then a binary search (recursively splitting the list in the middle and determining whether the searched-for item comes before or after the middle item and taking the upper or lower sub-list accordingly until the search ends in finding the item or failure) can be done in roughly $\log_2 N$ time, where N is the length of the list.

Natural environments tend to be unstructured in the way a human-made world is. Thus the first class of search algorithms dominate the world of tactical governance.

Commands

Tactical agents coordinate the work processes that acquire resources, export products, and dispose of wastes. They determine timing of activation of the interfaces with the external environment, flow rates of the inputs and outputs, and, in cooperation with the logistics agents, the flow rates from the interfaces (or local buffers) into the work processes internal to the system.

Tactical agents in mobile systems (e.g. animals or military units) order motor units (muscles or attack formations) to position the system in accordance with its objectives (eat or attack). These are coordinated through plans as previously described.

There are many more aspects to tactical management than is the case for logistical management and so this description must necessarily be abbreviated. We cover only those that are common across all CAS/CAESs.

The agents managing an interface process operation are responsible for activating an action that is coupled with actions of environmental entities such as sources or sinks, but also with milieu conditions. For example, a mobile system may be under the control of a foraging or hunting tactical agent, following a plan of attack to fell a prey. It has to find the prey, chase it, and bring it down. Once that act has taken place it hands off control (cooperatively) to a tactical agent responsible for assimilating the prey (eating it). Afterward, the food is managed by operations under the control of a logistics coordinator that issues commands for digestion and absorption.

There is a great deal of commonality between this description of a hunter chasing and eating its prey and a purchasing department looking for the best price and quality requirements for parts to be used in its manufacturing company. And this is not unlike an army attacking and vanquishing an enemy and then turning the prisoners of war over to prisons to be prosecuted.

The actuators that are managed by tactical agents are, as a rule, more complex and more powerful than those managed by logistics agents. They are often manipulated in temporal sequences (like the running gait of a cheetah) and also need to have a fair amount of flexibility built into the execution of those sequences[30].

Energy Requirements of Tactical Agents

Not surprisingly, given the huge increase in the computational complexity of tactical decisions (over logistical ones), they require much more energy to execute. Put simply, those systems that obtain benefits from a tighter coordination with their environmental entities must pay a higher cost of processing interactions. Animals that move about, both hunters and hunted, must have larger and more energy consumptive brains in order to survive.

Tactical decision models, recall, require not only the basic internal cybernetic mechanisms as logistic models, but also need to have constructed and compute a partial model of the entities they are coordinating with. If those entities are cooperative, the models need only be minimal and correlative rather than complex and causal. But if the entities (or for that matter the milieu) are competitive or threatening, then the models need to be much more elaborate and hence computationally much more expensive.

Coordination between Logistical and Tactical

Tactical decisions are often driven by factors that arise from entities in the environment. These, recall, are not necessarily modeled (at the tactical level, see below for strategic level models of environmental entities). At the same time, the decision process of a tactical agent is constrained by logistics from the insides of the system. The tactical agents are responsible for acquisition of resources and expulsion of wastes or export of products (in a timely fashion). But those processes, in turn, are conditioned by the needs of the internal work processes that are, basically, under logistic control. This means that tactical agents and logistical agents need to be able to cooperate at any given sublevel in the coordination level.

Strategic

To achieve long-term viability in a forever changing world, CAESs must be able to generate internal changes to its own structure and functions that better match up with the forces in the environment. They need to be able to adapt what they do and how they do it so as to continue to be fit in a different set of circumstances.

Fortunately, the world we are in is relatively stable for long enough periods to allow CAESs generally to take advantage of their capacity for evolvability, whether serendipitous or

[30] A fascinating example of this flexible programming may be exhibited by linear activation circuits in the cerebellum of mammal brains. These circuits are strongly implicated in muscle activation and coordination where, under the control of the motor cerebral cortex a sequence of activations can be modified in mid execution. The linear circuits are like computer programs (that have been learned) that have "if" statements embedded and the cerebral cortex, responding to real-time input, is able to change the parameters. Hence, a cheetah can run after a haphazardly escaping gazelle.

intentional to become adapted to changes that do occur. In contrast, we know that this wasn't the case for the world of 65 million years ago when an asteroid crashed into the Yucatan peninsula and led to the mass extinctions that included the dinosaurs. The amount of change involved in an incident like that was too great for the vast majority of species then extant. Some, fortunately, survived to provide the seeds of a new efflorescence of life, the adaptive radiation of new species.

Darwinian evolution serves the purpose of making strategic 'decisions' for natural CASs. Human engineering provides the same service for CAS artefacts. But human beings as individuals, and societies of human beings, undertake intentional evolution. The human brain[31] is a CAES and by virtue of that human organizations and societies are also intentionally strategic. That is the sorts of decisions to which we now turn.

Observing Other Entities and Forces in the Environment

A prime responsibility of the strategic agent is to observe more of the environment than just the entities with which the tactical agent is concerned. The reason is that those many other entities may have direct or indirect causal relations with the resource sources or product sinks that will change their behaviors in the future. The tactical agent may supply information regarding the long-term behaviors of the sources and sinks (as longer time averaged behaviors) but the strategic agent must directly observe the rest of the environment. The strategic agent also receives long-time averaged performance data from the logistical agent. Figure 11.12 depicts the situation.

[31] In reality it may be that many mammals and birds are also capable of possessing concepts that are modifiable to some extent. We know, for example, that chimpanzees organize hunting parties and war parties, indicating that they have an ability to think into the future and lay plans. Crows and parrots are also known to behave as if they were thinking ahead and assessing their environment to change their behaviors somewhat. De Waal (2016) asks if we are smart enough to know how smart animals are (from the title of his book)? And he suggests that many species demonstrate great cleverness and adaptability when confronted by changes in their circumstances. We suspect that the capacity to think strategically (which is considerably more than just being clever) is something achieved only by animals with brains with prefrontal cortices (as discussed in Appendix A) or their equivalent in birds.

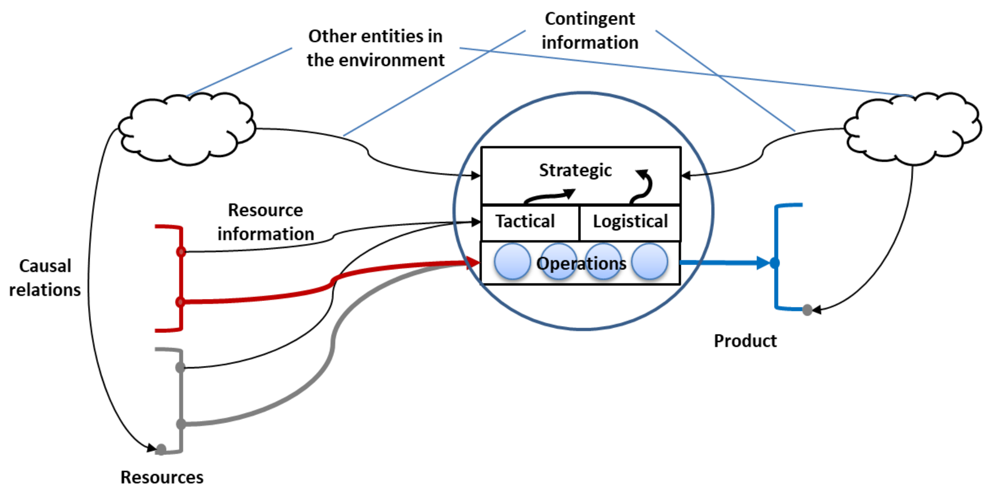

Fig. 11.12. The strategic manager observes other entities and the milieu in the environment. There are many possible 'other' entities that have direct or indirect causal relations with the resource entities (sources) and/or sink entities that affect the future behavior of those entities.

The clouds represent other entities (milieu observations not shown) that may or may not have causal relations with the sources and sinks. The strategic agent observes the behaviors of such entities over a longer time scale and is responsible for finding the causal correlations if they exist. The reason this is necessary will be explained below. It has to do with the need to anticipate those changes and take preemptive action to avoid damage.

Anticipatory Models of the Environment - What to Expect in the Future

The purpose of monitoring the long-term behavior of other entities in the environment in addition to the ones that have a direct relation with the whole system is to detect and exploit any causal relations that exist between those entities and the ones upon which the system depends. The strategic decision model of the environment can then be used to form expectations for what may happen in the future that would require evolutionary change in the system. For example, an automobile manufacturer that depends on a sole-source supplier of a critical part may take note that that supplier has lost another major customer that could put them in jeopardy and thus threaten the companies' future supply line. Or perhaps the supplier is known to be carrying a heavy debt load and the central bank raised interest rates that would squeeze the profits of the supplier. This is an example of a milieu condition change.

An anticipatory model is one that *learns*[32] about causal correlations between entities and milieu conditions in the environment and uses that model to anticipate scenarios that may play

[32] Learning causal correlations is covered in Appendix A with respect to how biological neural networks appear to do it. In general, any learning mechanism that can deal with nonstationary relations should be able to be used for the purpose. In this era of Big Data and Deep Learning, we expect to see these applied increasingly to more structured strategic model building.

out in the future (out to some time horizon that is within the computational competence of the agent). Mobus (2000) has provided a working model of this learning in a robot.

Initially, the strategic agent may not have any information about which entities are relevant to its model. Clearly the agent cannot observe everything at once even if their sampling frequency is low. If the agent has background knowledge (either explicit or implicit) it can "make educated guesses." That is, it can select entities that seem most likely to be related to its sources and sinks. For example, if the sensory arrays and modalities available to the agent are sophisticated (e.g. vision, auditory, etc.) the agent may be able to discern systemic patterns in the manner discussed in Appendix A.

Causal Correlations

A causal correlation imposes a time ordering on events that must consistently not be violated as discussed in Appendix A. The brain's ability to encode causal relations and to modify the model in nonstationary environments is an archetype for modeling in general. Machine representations of networks of causal relations are the basis for computer-based models. These models are both structural and dynamical, expressed in Systemese from Chapter 3 and Appendix B.

Anticipating New Sources or Sinks

Another aspect of building a causal model of the environment, even if some entities included are not found to be immediately related to current sources of sinks, is that the model can be used to identify new opportunities, new sources of current resources or sinks (customers) for current products. If the agent is sufficiently sophisticated and capable of recognizing affordances it might be able to determine completely new resources that might be exploited or new products that it should produce for new customers. Human beings are quite good at this. Unfortunately, not much is known about how the brain computes affordances so we have little understanding of how this might be accomplished in a machine form. For the present, then, and probably into the near future, the ability for an organization or society to recognize affordances will rely on human insight.

Anticipating New Threats

A causal model of the environment can be used to identify new threat conditions that may require evolving the system to counter (be pre-adapted to the change). For example, human societies are currently learning about the threat of global warming and climate change (e.g. sea level rises that will inundate coastal cities and island nations). The leaders of these societies, the putative strategic agents, are beginning to consider what adaptations may be needed to counter these threats. The CEOs of corporations are constantly concerned with product and financial market conditions and spot potential problems before they negatively impact the firm.

Observing the System Itself

Strategic agents actually have two models to work with. The model of the external environment provides anticipatory scenarios about what might happen in the environment to provide new opportunities or pose new threats. But they also have a model (usually very abstract) of the whole system itself (Principle 10 in Chapter 1). This model can arise from the same kind of learning procedure as described above and in Appendix A. Or it can be "given," that is provided by the structure of the system itself. For evolvable systems the model is adaptable to changing conditions.

The purpose of the self-model is to assess capabilities relative to external opportunities and threats. For example, in the case of a new opportunity, e.g. a new product to serve a new market, the strategic agent has to assess the system's ability to produce it. It may determine that significant investment in new structures/functions will be required, in which case it issues commands to the tactical and logistical agents to make this happen. The system then evolves to take advantage of (be fit in) this new environment.

Planning

Strategic agents develop long-term plans for their future behavior based on the models of the environment and of themselves. If they identified new opportunities or threats they have to consider how prepared they are to meet the challenges, or how they might recruit the resources needed to effect a change in the structure/function of the organization/society. This is the problem of affordance processing mentioned above, for which there are no current machine-computable models[33]. Some humans are quite good at doing this as the strategic agents of organizations; they are CEOs of various kinds of organizations. Of course, not all organizations will necessarily have a CEO that is good at affordance, so not all organizations do a particularly good job of strategic management. For non-autocratic governments it is very unclear how or even if affordance is accomplished. Other than the conduct of war (in various forms) governments do not seem to be able to think strategically (see discussion below). Autocracy might be successful if the leader is good at strategic thinking and affordance. For most of history, however, such 'wise' leaders have been in the minority based on the fact that most historical civilizations, which were often autocratic or oligarchic, have collapsed, with the reasons understood to be poor strategies (Diamond, 2005). As Diamond (2005) points out, many past societies were caught unawares of changes in climate (a milieu factor) and failed to adapt in time. Others chose to pursue unsustainable practices such as distant conquests or soil-depleting farming practices.

[33] Affordance is more than pattern recognition. Current automated pattern recognition can be trained to recognize an object for its category properties but that isn't enough to suggest uses of the object in other contexts. This would be an interesting area of research to pursue – linking pattern recognition to identifying affordances.

Planning is a process of using the two anticipatory models discussed above, coupled so that the output of the environment model is input to the self-model and vice versa, considering multiple possible states of inputs to the models, and then aligning the system's capabilities against the outcomes to test, in much faster than real-time, whether the system is sustainable (c.f. Rosen, 1985). And if it is not under its current configuration, what would need to change to make it so?

Playing What-if with Model Simulations

A 'what-if' game is played with environment- and system self-models. Factors in the environment, including the other entities and milieu factors described above, are varied in a range of plausible scenarios. The factors' range of changes go from low levels of change to high levels of change in a set of model simulation runs. The outputs of the simulation are what are called best-case to worst-case scenarios. The middle case is generally thought to be the most probable outcome, but this depends entirely on assumptions made about the values chosen for the low end and the high end of the factor values.

Human beings practice this sort of simulation when they rehearse a scene they expect to play out in real life in the future. For example, when an employee anticipates asking their boss for a raises, they will imagine the scenario in their minds and with each run vary some aspect of what they do and imagine how the boss will respond. Since the brain can construct systems models of its own personality and the boss' then it is engaging in exactly this sort of what-if game playing. We can use this simulation approach in much faster than real-time processing so as to project what is likely to happen into the future. We can thus anticipate which things might happen and what our more successful moves would be. Of course, the success of using this kind of planning depends on the efficacy of our models. Not everyone is sufficiently good at constructing veridical models or of eliminating emotional biases in selecting scenarios.

Using Scenarios to Set Up Plans of Action

Above, in section [Tactical Management.Decision Models.Plans] on tactical plans we considered the assembly of action steps needed to accomplish a goal. In most cases the plan is actually a series of plans, each set up to accomplish strategic steps toward a strategic goal. The strategic agent does not determine the details of the actions to be undertaken. Rather its job is to set up a sequence of goals that the tactical or logistical managers should accomplish. The sequence needs to be designed so as to lead ultimately to a strategic goal. For example, a human being may decide that going to college (while, say, a junior in high school) is in her best interest and so considers the tactical goals she will have to achieve to get into college. She might consider her grade point average (needs to be high), her extracurricular activities (need to be diverse), and her community involvement (need to be engaging and sacrificial). She thinks this through strategically but then hands off the goals to the part of her consciousness that deals with tactical details. She will work up tactical plans to realize these goals. The same brain is participating in strategic as well as tactical thinking. This might tend to blur the lines between the two, which is a very general problem for human agents when they try to explicitly participate in either kinds of thinking. We will examine this problem below when considering the human situation in governance explicitly.

For evolvable systems the strategic decisions may involve modifications to the structures and functions of the operational and coordination levels of the organization. For human brains this is limited to the act of changing one's mind about a concept. No changes in biological functions are permitted[34]. However, changes of mind can lead to changes in overt behavior. For example, a human who has had liberal political leanings might reflect on, say, the debt crises, and shift his thinking to become more conservative, after which, voting for conservative candidates.

In organizations and societies there are no real constraints on what evolutions might take place.

Changing the Organization

Ultimately the purpose of strategic management is to evolve the organization so that it remains fit in the ever-changing environment in which it seeks to persist. In naturally evolving systems such as a species or an ecosystem the changes are achieved through semi-random mutations (or invasive species) spread among a large population of individuals followed by natural selection for those few changes that led to phenotypes better equipped to persist, or, as in the case of an ecosystem, greater stability and more efficient energy flux. The changes are chance but based on the fact that they take place in a large population of fecund individuals there is bound to be some change that is beneficial and will be propagated to the next generations.

In intentional, volitional systems such as the human brain or human organizations, the changes are made by design or adoption of and importation into the system of a structure/function observed in other systems. The change is still provisional and risky in that it is the strategic agent's best guess about what will happen in the environment in the future and that the design will be capable of countering the environmental situation or take advantage of a new opportunity. The anticipatory models that produced the guess are only probabilistic and only as good as the model building procedure used to create it. A weak model will produce weak scenarios and coupling those with models of the 'new' organization will likely not produce results in which one can place any confidence. The ability to implement a change that will be successful in boosting the future fitness of the system depends very heavily on the capability of the strategic agent to construct good models (see Chapter 10).

[34] Although neuropsychology is considering forced adaptations of physiology based on held beliefs. The most straightforward such adaptations are seen in the placebo effect, now known to be a real shift in biological function due primarily to the patient's belief about their own condition. This is an area of intense research – the problem of mind over matter in the arena of physiology. Stay tuned.

The strategic agent is responsible for deciding not only what changes should be implemented, but also why they should be implemented. For natural biological evolution that is a given. The genus evolves new species in order to continue into the future. For intentional and volitional systems, the situation is somewhat more complicated.

Considering Motivations, Values, and Beliefs (Ideology)

Humans and human organizations are motivated to continue their existence and are aware of the alternative consequences. At almost any time of day (or night for nocturnal creatures) an individual is aware of some motivating drive such as hunger until they take action to satisfy the need and it is satiated. All such drives are cyclical and may alternate in their presence in the individual's awareness.

Needless to say, human motivations for decisions and actions are much more complex than just the ordinary biological drives. Humans construct additional layers of motivations on top of those. Some of these, for example what we call 'values' are, at least in part, socially constructed or culturally derived and so are owing to historical developments (some of which were volitional change at the time, but most were likely ordinary accidental evolutionary).

The subject of the complexity of human motivations and values is still in the process of explication and is filling volumes of psychology papers/books so must necessarily be far beyond the scope of this appendix. Our purpose in pointing it out is to underscore the need for the strategic agents in human organizations and governments to take this factor into account when contemplating changes to meet a strategic need. A most common example of why this is requisite is the case of an organization implementing some change (ordering it from the top-down) without getting buy-in from the lower echelons in the organization. The cases of subtle passive resistance and outright sabotage are well outlined in, for example, past issues of the Harvard Business Review.

Where the effects of motivations, values, and beliefs becomes obviously potent is in the social subsystem of politics[35]. Political decisions are different from governance decisions. They are about what prevailing ideology is to be used to interpret the decision options given in the governance agency. Ideologies are a grouping of beliefs about how the world works, sometimes called a 'world view'. These beliefs need not be grounded in the reality of how the world actually does work (or should). Neither are they necessarily internally consistent. Ideologies work to shape our concepts of governance

[35] We are using the term, politics, in the very broad sense of human interactions directed at deciding what agents should be situated in the key decision-making roles in a governance structure. We most commonly think of politics in the context of governments but the same processes of argumentation and negotiation leading to decisions about who gets to decide occurs within all human organizations. For example, it is well understood that in a non-government organization official titles do not necessarily point to who is actually providing leadership.

Implementing Change

Figure 9.2 is a model of a complete intentional CAES showing auxiliary processes not directly involved in the on-going economics. Processes such as the adaptation, repair and maintenance, and evolvability handle the exception cases due to disturbances (the first two) and long-term significant changes in the environment (last one). Figure 11.13 shows the relevant processes from Figure 9.2 along with additional details of how a strategic decision to modify an internal work process comes about.

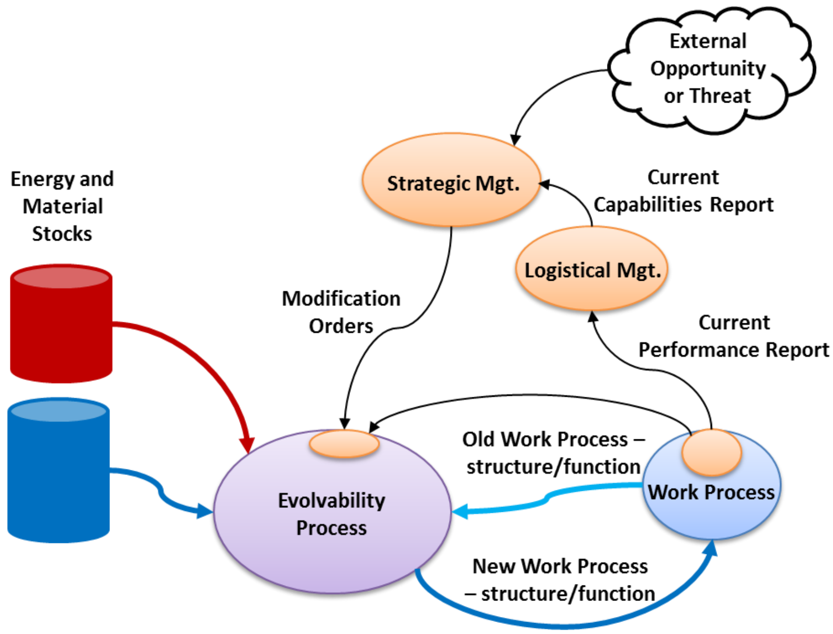

Figure 11.13. The strategic decision to modify a work process is implemented by a special work process, the evolvability process, which is capable of turning a strategic command into the physical modification.

The strategic management process monitors the environment and detects a possible opportunity or a threat by comparing the current capabilities of the lower level work processes to the nature of the threat or opportunity. The capabilities may have been determined by the logistics management (in this example) from a current state of time averaged past performance. The strategic manager, exercising affordance, determines what needs to be done in the form of modifications to a work process that would then meet the future needs of the system to thwart a threat or seize an opportunity. The strategic agent sends a command to the evolvability process directing it to modify the work process. It requires knowledge of the current work process (its structures and functions) and proceeds to 'engineer' the modifications that need to be made. With that specification and material/energy resources from system stores, it takes as input the old process and returns the modified version with new structures and functions.

For example, suppose marketing has identified a currently underserved customer base that would consume a modified product that a company already builds. To serve this base, the work process that manufactures the current product will have to be changed so as to produce the new version. Automobile companies famously bring out new models every few years based on this 'theory.'

The evolvability process should be seen as a complete manufacturing capability for tools and procedures that is a subsystem of the whole CAES. It is like any other work process in the sense of requiring energy and material inputs that it is capable of converting to the changes required by the strategic decision. But it is not a 'routine' process. It is highly flexible and contains a design and engineering component that can translate the strategic command into a usable and efficient design (Part 4 covers this topic).

A very similar process takes place in the human brain as concepts are being 'refined' or corrected (as in Appendix A). Though the details are not yet understood, it appears that the hippocampus is involved in orchestrating memory formations and reformations. Concepts are distinctive patterns of neural activity correlated across many brain regions. Engrams are formed when neural clusters representing percepts and concepts fire in synchrony with signals from the hippocampus that tell these clusters that they are part of a causally correlated assembly. They then strengthen their synaptic connections as described in Appendix A.

Evolvability processes may be as elaborate as needed to achieve changes in structures and functions in the system. Some organizations are able to make whole new processes in situ. Others have evolved means of incorporating smaller autonomous organizations that already do what the larger organization seeks to do. We call those operations mergers. Not too surprisingly, when corporations acquire other corporations to start serving a different market they are following a very old tradition in CAESs that started, perhaps, some two billion years ago when some larger prokaryotes formed symbiotic relations with other, smaller prokaryotes that eventually gave rise to eukaryotic cells[36]. This will be discussed below in the example of governance in cellular metabolism.

The act of making a strategic decision in government is no different from what happens in smaller organizations, except that the detailed architecture of a government may affect where and how such decisions are made. In the case of a representative democracy with a legislative branch of government, the decision to make changes may come from legislation. However, as will be discussed below in the section [The HCGS and Governance of HSS], the layout of the strategic decision-making agency is muddled. The changes affect policies and laws. The latter are determined by the legislature whereas the former fall within the jurisdiction of the executive branch. This division of power can get very confusing as well as contentious. Suffice it for now to say that however the rules for governing the society are decided, there are a set of operational agencies (mostly under the control of the executive branch) that are tasked with carrying out the change order.

[36] The process of symbiogenesis is thought to be the origin of eukaryotic cells and a precursor to multicellularity in animals, plants, and fungi. See the Wikipedia article: https://en.wikipedia.org/wiki/Symbiogenesis for background. Accessed 7/3/2018.

Monitoring On-going Performance

Once a strategic decision is made and the command executed in the actual change to the organization, the duties of the strategic agent are not completed. This is an issue that is sometimes lost on executives who issue commands and then think they have earned their paychecks. In reality the strategic agent has to continue to monitor (through reports supplied by the tactical and logistics managers) the actual performance of the changes to make sure they are doing what the anticipatory models indicated. This is a higher-order, longer time-scale but nevertheless cybernetic activity.

There are many possible criteria for success and no simple set point error signal will tell the strategic manager how the whole system is performing. For Darwinian evolving CAESs such as a species, the monitoring and assessment are performed by natural selection. But in intentional CAESs such as brains and organizations, the monitoring needs to be explicit and verifiable. Moreover, whereas in natural systems where the law of large numbers is at play, in intentional systems there is a limited number of conditions for success. If the marketing manager's projections of sales turns out to be wrong, then the company may take a financial hit that will weaken it.

In the government of a society the need for on-going monitoring and attention to variances is all the more important. One of the aspects of adopting policies or legislating rules for commerce or behavior is the emergence of unintended consequences. In the governing of society these can take generations before they are recognized. And one of the crucial failures of governments is to not continue monitoring the social or economic impact of the changes they mandate and evaluate the effectiveness of their changes, or the consequences they never considered originally. At prime example is the various ways in which the fossil fuel industry has been subsidized financially in order to boost production and profits. Meanwhile, the increasing consumption of fuels has led to major releases of CO_2 into the atmosphere with consequent raising of the global mean temperature due to the greenhouse effect. That has to count as the granddaddy of unintended consequences.

Energy Requirements for Strategic Agents

Strategic management activities are information processes, which in general, require less power than manufacturing processes. However, of the various types and levels of agents in a governance architecture, these require considerably more power. This can be seen in the energy requirements of the human brain, the latest evolved module, the prefrontal cortex, is a power hog compared with other organs, including muscles. The whole brain consumes some 20% of the energy available to the human body at rest. Surprisingly, the prefrontal cortex (where 'thinking' takes place) may only consume about 6% (Marieke et. al, 2008). But this figure is somewhat misleading in that the prefrontal cortex is actually directing the firing of many different neural clusters throughout the brain. So, its power consumption alone is not indicative of the total power consumed in strategic thinking.

Another way to assess the power requirements of strategic agents is to consider the budgets for operations of executive offices in corporations. Here too, the overall budget of the localized office might not be significant with respect to the budgets of other management functions. However, just as with the situation with the brain prefrontal cortex commandeering resources from other parts of the brain, the strategic management of an organization often works by distributing chores to lower down offices that then provide the work required and consume the energy directed at strategic decisions. The corner offices on the top floors are generally responsible for a considerable amount of energy expended on information processing.

Comparing Stafford Beer's VSM with the CAS/CAES Governance Architecture

In Chapter 9 we introduced the generic CAS/CAES model archetype comprised of three subsidiary model archetypes (agent, HCGS, and economy). We also noted how the CAS/CAES model is very similar to Stafford Beer's VSM in terms of describing what a system needs to do to be a sustaining viable one. We now return to the VSM concept and describe here some key differences between the VSM and the CAS/CAES models especially with respect to how the governance subsystems are organized and operate in both. Recall that Beer was mostly concerned with applying principles of cybernetics to the management of organizations like corporations or NGOs.

Table 11.1. Relationships between elements of the HCGS model and Beer's VSM

Stafford Beer's Viable System Model[37] [mapping relation in brackets]
System 1 in a viable system contains several primary activities. Each System 1 primary activity is itself a viable system due to the recursive nature of systems as described above. These are concerned with performing a function that implements at least part of the key transformation of the organization. [Operations level units - work processes] System 2 represents the information channels and bodies that allow the primary activities in System 1 to communicate between each other and which allow System 3 to monitor and co-ordinate the activities within System 1. Represents the scheduling function of shared resources to be used by System 1. [Cooperation communications network] System 3 represents the structures and controls that are put into place to establish the rules, resources, rights and responsibilities of System 1 and to provide an interface with Systems 4/5.

[37] From Wikipedia: https://en.wikipedia.org/wiki/Viable_system_model. Accessed: 11/27/2016.

Represents the big picture view of the processes inside of System 1. [Similar to the coordination level but not differentiated between tactical and logistical] [Systems 1 - 3 concerned with the "here and now"] System 4 is made up of bodies that are responsible for looking outwards to the environment to monitor how the organization needs to adapt to remain viable. [Half of the strategic level] [Concerned with reconciling the "here and now" with the "then and there" (below)] System 5 is responsible for policy decisions within the organization as a whole to balance demands from different parts of the organization and steer the organization as a whole. [Other half of the strategic level - concerned with the "there and then"]

Examples of the HCGS in Nature

Cells and Multicellular Organisms

Evolution selects for those CAS/CAESs that are successful in existing in any given environmental situation. No such system can sustain for long without self-regulation. Recall that the CAS/CAES category of systems essentially starts with the *origin of life* and constitutes the living and supra-living systems we find in our world today.

Lee & Morowitz (2016, especially Chapter 4) have written extensively about the origin of life from the geochemistry that existed on the Earth early in its life and how that origin can be conceived of as a phase transition from mere inorganic chemistry on mineral substrates through autocatalytic processes of carbon reduction cycles to metabolism as we now understand it. In their model the earliest metabolic processes were a natural extension of geochemical reactions, probably at early anoxic geothermal vents in deep oceanic spreading zones, examples of which can be found today[38]. Once this chemistry became associated with encapsulation via lipid vesicles (protocells) it came under the influence of selection for protein-based enzymes that could facilitate the fixation of carbon (into basic carbohydrates, amino acids, phosphoric acids, and lipids) the precursors of biochemical processes. The phosphoric acids combined with ribose carbohydrates were the substrate of RNA and its role in amino acid polymerization – the construction of proteins.

Protein-based enzymes are the work horses of both anabolism (construction of higher molecular weight molecules such as other proteins and nucleic acids) and catabolism (the

[38] In fact, it seems that our most current thinking about how life emerged comes from the examination of the extremophiles found in these environments.

destruction of higher weight molecules to either recover the components or produce waste products). Their theory is that the autocatalytic cycles that constituted the original proto-living systems form an inner core that has basically been conserved from the origin in the last universal common ancestor (LUCA) to the present. That core, the local regulation of which involves the stoichiometry involved in reactions mediated often by cofactors[39] and some metal ions, has been overlain by the networks of enzymes that further regulate it. Additionally, with the advent of an oxygen-rich atmosphere, the reducing cycles were augmented with oxidizing reactions that could provide much more energy to the anabolic processes. For this to happen, the network of enzymes took on the role of regulating the core(s). But the enzymes themselves had to be constructed. A popular current theory suggests that ribonucleic acid (RNA) molecules that have enzymatic capacity may have mediated the early stages, providing a means for primitive protein enzymes to be synthesized. As the protein enzymes evolved they became more specific to reaction types, and more efficient than the ribozymes[40] (the general name given to RNA molecules capable of enzymatic activity) that preceded them.

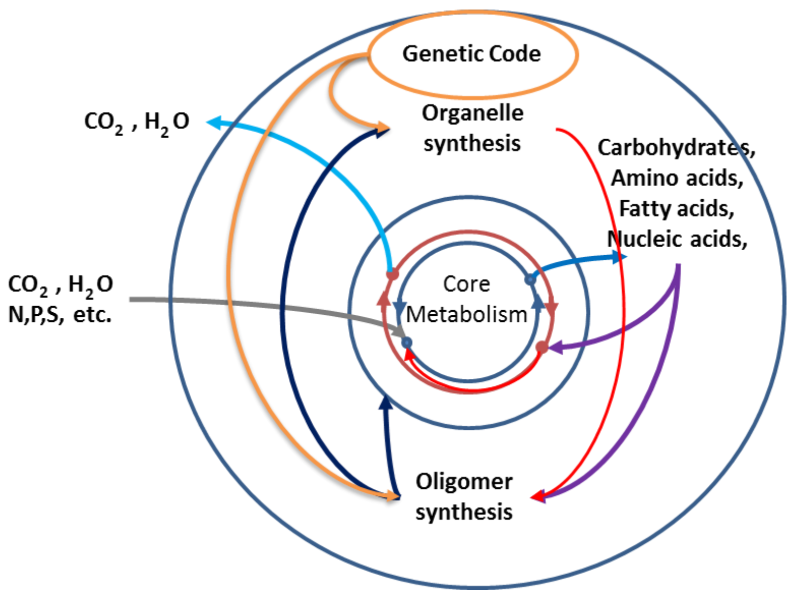

[39] The origin of life is thought to have involved 'naked' cofactors that acted as catalysts directing the sequences in autocatalytic cycles that would become cellular metabolism with the advent of protein enzymes. The latter associated with the cofactors to further increase the specificity of reaction control and its rate. See the Wikipedia article: https://en.wikipedia.org/wiki/Cofactor_(biochemistry) for background. Accessed 7/3/2018.

[40] See the Wikipedia article: https://en.wikipedia.org/wiki/Ribozyme for background. Accessed 7/3/2018/

Fig. 11.14. A rough outline of the layered control structure of a living cell shows the distribution of governance. Core metabolism is the basic production work processes building fundamental metabolites, carbohydrates, amino acids, fatty acids, and nucleic acids. Not shown is the production of cofactors, e.g. biotin, that are used in the autocatalytic cycles. The inner blue cycle is the reductive cycle where carbon is reduced and other atomic components are added in to produce the products. The outer red cycle is the oxidative cycle that uses some of the products (e.g. carbohydrates) to extract energy. The outer blue oval represents the main logistical and tactical management. General coordination is processed by the genetic code through the production of structural (organelle) elements and especially protein synthesis (oligomers).

Figure 11.14 is a maximally simple representation of the nature of management in a cellular context. The concepts of anabolism (synthesis of higher weight molecules) and catabolism (breaking higher weight molecules down for components and energy) are grossly oversimplified but the organization of cell physiology into managed layers recapitulates the HCGS insofar as operational control (embodied in autocatalytic cycles mediated and amplified by critical cofactors), logistical control (embodied firstly in the catalytic activities of higher-weight protein enzymes and secondly in the genetic code transcription process that helps regulate the concentrations of those proteins. Tactical management is embodied in the particulars of how specific kinds of cells obtain their inputs and export their products. Usually this is managed by specialized pore structures penetrating the cell membrane (interfaces) but some substances, like gas molecules, may simply diffuse through the membrane according to differential pressures between the inside and outside of the cell membrane.

At this point an outer core of regulation, essentially homeostatic mechanisms, takes on the task of sensing the concentrations of various metabolites and their enzyme actuators in the inner core (now augmented with some innovations and elaborations that appear to provide greater stability to the balance of anabolism and catabolism. A major transition in cellular evolution occurred

https://en.wikipedia.org/wiki/Symbiogenesis

Monitoring – CA+2 accumulation (integral) during synaptic excitations activates a signal pathway that may ultimately result in gene activation.

Societies of Organisms

Human Brain as an HCGS

The HCGS and the Governance of Human Social Systems

Social systems are aggregates of human participants who have a sense of common purpose. Whether it is a family unit, a local community, a social organization (like a club), a city, or a nation state, the architecture of the HCGS, or something bearing some resemblance to it, can be detected in the organization of management and administration. All such social systems make strategic, tactical, logistical, and operational decisions. Sometimes, in small social systems, one or a very few individuals are responsible for these decisions. In larger systems the roles may be distributed in an "official" hierarchy as in a corporation or the military, with specifically identified individuals responsible for making them and carrying out the implementations.

What is uniquely different about a human social system as compared with naturally evolved CAS/CAESs, is that the agents are individuals with human brains. And as we saw above, the human brain is, itself, an HCGS complete with a strategic level. This sets up an interesting problem and also a great opportunity. Figure 11.y provides a sense of the situation. In any organization we find all three levels and agents within each level capable of all three levels themselves.

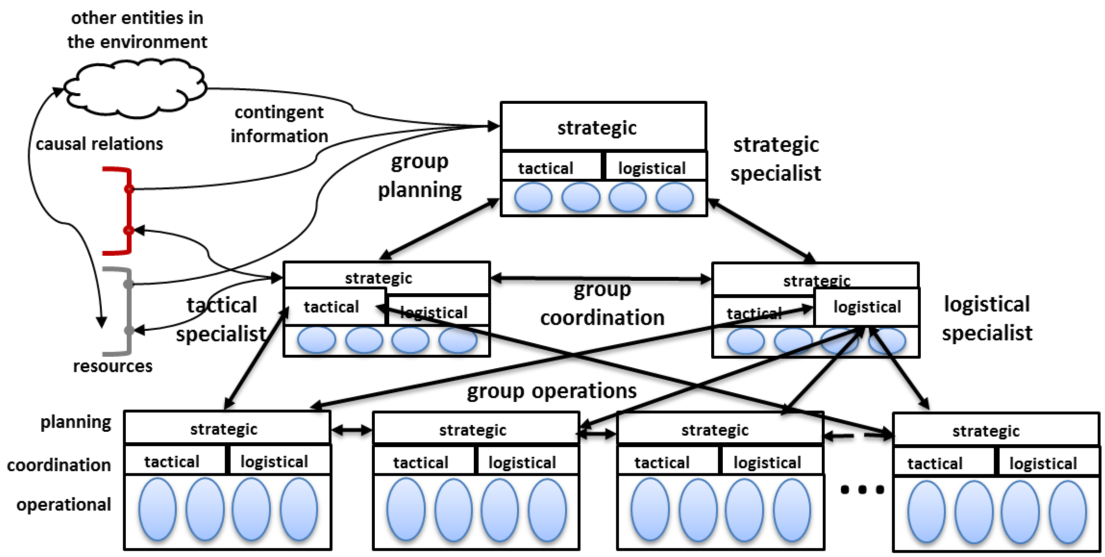

Fig. 11.15. An HCGS for human organizations is a set of nested hierarchies. Blue ovals represent work processes. At the operations level these subsystems may be processing materials and energies as well as information. At the coordination level these represent specific information processes that require management. The same is the case at the strategic level.

The figure depicts the basic architecture. This architecture can be recursively applied as societies grow in scale (basically the number of people) to manage complexity. As noted in section [Coordination at Larger Scales] above, as a system grows in size and heterogeneity it must divisionalize at all levels, adding sub-levels where necessary in the coordination level, due to the span of control problem. But the relations shown in Figure 11.14 remain.

One might ask why operations level agents need to have strategic or tactical, if not logistic, decision-making capabilities. There is an advantage. A human agent can manage far more activities, greater complexity of work processes, meaning that having coordination-level and strategic-level decision-making abilities allows them to interact with the other agents (i.e. other *departments*) strategically. This results in the potential for evolvability anywhere in the social organization. For example, an operations manager given notice that improvements in the product that they are responsible for producing can plan and execute the necessary changes without being given detailed instructions from above. In other words, the hierarchy of human agents permits the efficient distribution of all three kinds of decisions. Top-down micromanagement is, in theory, unnecessary.

This architecture provides a tremendous amount of flexibility and 'modularity' with respect to how any agent in any position in this organization can be replaced (at least in theory) with any other agent (or newly recruited agents). Notwithstanding individual talents (note in the figure that the rectangles associated with the level and decision types are shown a bit larger, representing a specialization relation) thanks to the ability to learn and adapt, many human individuals are equipped with the basic capabilities to be placed in any of the positions. This is exactly the basis for things like promotions in corporations and politicians running for higher-level offices. It is also the basis for being able to cover a critical position if an employee or member is unable to perform their duties. It is thought that one of the main reasons that humans have been so successful as a species is that working in social groups with such flexibility gave them a distinct advantage in the game of survival (see Chapter 8, section [Societies and Economies], Figure 8.7 and section [Tribes and Beyond], Figure 8.8).

The problem, however, is that the agents as depicted in the figure have more to them than ideal decision-makers. We humans have drives, needs, wants, and a panoply of emotions to underscore them. We have ambitions and expectations that can too often supersede the rational decision process and our individual strategic thinking can lead us to less-than-optimal tactics. In other words, humans are too often not motivated to make the best decisions for the well-being of the society. This is a fundamental conundrum.

We will not belabor the human foibles that sometimes keep the governance of social systems from performing the function of maintaining the system as a whole. Psychologically, we are what we are. Below we will compare typical governments of nation states with the HCGS architecture and functioning to get an idea of where the governance of human social systems stand in relation. We will consider some plausible mechanisms that could be implemented to compensate for human foibles along the way. But primarily we seek to demonstrate some major flaws in the current state of thinking about the structure and function of governments as we find them today. We need to remind ourselves that governance structures evolve over time. Those that survive do so because they prove capable of keeping their substrate system viable in the long run. Insofar as human governments are concerned, very few, at the nation state or empire level have proven viable for very long. In light of the rate of state failures and state disruptions currently taking place, we need to ask if this is because the governments we have seen to date deviate from the model of governance that nature has found so useful.

Typical Government Architectures and the HCGS Model

Recursive, nested structures

Governmental Roles

Regardless of whether the political system is based on democracy (in some form) or authoritarian power, three basic functions have to be performed. In a form of democracy such as the representative form of the USA, these functions are usually performed by separate branches of the government. In dictatorships all functions may be under the control of the dictator even if performed by separate agencies.

The typical governmental architecture found around the globe is based on a tripartite division of responsibilities into three distinct, and generally co-equal branches. These are the executive, the legislative, and the judicial branches. This architecture is repeated at various scales from nation states to counties to cities of varying sizes.

Executive branch duties are not dissimilar from chief executive and chief operating officer duties in organizations. The executive branch sees to the administration of the laws of the land and the management of public infrastructure. As a rule, the executive is also responsible for public health and safety, as well as defense against foreign powers (in the case of nations).

The legislative branch is tasked with making the laws of the land. Laws cover a variety of social interaction arenas such as commerce and trade, citizen behavior, and maintenance of individual and group rights.

The judiciary's primary responsibility is to administrate justice, to make findings of legality, to sanction those who are found to be breaking laws or harming others, and to interpret legislated laws in terms of higher-order laws such as constitutions, which are considered the ultimate set of laws.

The Basic Problem – Mapping the HCGS onto the Government Architecture

As a challenge to the use of systems science, thinking, and analysis to the really big problem domains (as in Chapter 8 when we tackled the biophysical economy) we will examine a "typical" government architecture.

Architecture – Constitution

The HSS as a Whole

Conclusions

Chapter References

Ashby, WR (1958) Requisite Variety and its implications for the control of complex systems, *Cybernetica*, Vo1 1, No 2. Available on-line: http://pespmc1.vub.ac.be/Books/AshbyReqVar.pdf. Accessed 11/28/2016.

Beer, S (1959), *Cybernetics and Management*, English University Press, London UK.

Beer, S (1966) *Decision and Control: The Meaning of Operational Research and Management Cybernetics*, John Wiley & Sons, New York.

Beer, S (1972), *Brain of the Firm: A Development in Management Cybernetics*, Herder and Herder, New York.

Burton, RA (2009) *On Being Certain: Believing You Are Right Even When You're Not*, St. Martin's Griffin, New York.

Coen, E (2012) *Cells to Civilizations: The Principles of Change that Shapes Life*, Princeton University Press, Princeton, NJ.

Cuban, L (1999) *How Scholarship Trumped Teachers: Change Without Reform in University Curriculum, Teaching, and Research, 189-1990,* Teachers College Press, New York

Damasio, AR (1994). *Descartes' Error: Emotion, Reason, and the Human Brain*, G.P. Putnum's Sons, New York.

De Waal, F (2016) *Are We Smart Enough to Know How Smart Animals Are?*, W.W. Norton & Co., New York

Diamond, J (2005) *Collapse: How Societies Choose to Fail or Succeed*, Penguin Group, New York

Findeisen, W, Bailey, FN, Brdys, M., Malinowski, K, Tatjewski, P & Wozniak, A (1980). *Control and Coordination in Hierarchical Systems*, John Wiley & Sons, New York.

Geiger, RL (2004) *Knowledge & Money: Research Universities and the Paradox of the Marketplace*, Stanford University Press, Stanford, CA

Hansson, SO (1994) *Decision Theory: A Brief Introduction*, Department of Philosophy and History of Technology, Royal Institute of Technology (KTH), Stockholm. Available: https://web.archive.org/web/20060705052730/http://www.infra.kth.se/~soh/decisiontheory.pdf. Accessed: 11/27/2016.

Heylighen, F, Rosseel, E, & Demeyere, F (Eds.) (1990), *Self-steering and Cognition in Complex Systems: Toward a New Cybernetics*, Gordon and Breach Science Publishers, New York.

Ivanenko, IV (2010) *Decision System and Nonstochastic Randomness*, Springer, New York, available on-line: ftp://nozdr.ru/biblioteka/kolxo3/M/MV/MVsa/Ivanenko%20V.I.%20Decision%20systems%20and%20nonstochastic%20randomness%20(Springer,%202010)(ISBN%20144195547X)(O)(285s)_MVsa_.pdf . Accessed: 1/29/2018.

Kahneman, (2011). *Thinking Fast and Slow*, Farrar, Straus, and Giroux, New York.

Losos, JB (2017) *Improbable Destinies: Fate, Chance, and the Future of Evolution*, Riverhead Books, New York

Marieke LS, Howarth C, & Attwell,D (2008), "The cortical energy needed for conscious perception", *Neuroimage*. 40(4): 1460–1468

Minsky, M (1986) *The Society of Mind*, Simon & Schuster, New York.

Mobus, GE (1999) "Foraging Search: Prototypical Intelligence", *Third International Conference on Computing Anticipatory Systems*, Center for Hyperincursion and Anticipation in Ordered Systems, August 9-14, 1999, Liege, Belgium, Daniel Dubois, Editor. On-line: http://faculty.washington.edu/gmobus/ForagingSearch/Foraging.html

Mobus, GE (2000) "Adapting Robot Behavior to Nonstationary Environments: A Deeper Biologically Inspired Model of Neural Processing" *Proceedings, SPIE - International Conference: Sensor Fusion and Decentralized Control in Robotic Systems III, Biologically Inspired Robotics*, Boston, MA, Nov. 6-8, 2000, pp. 98-112

Mobus, GE (2017) "A Framework for Understanding and Achieving Sustainability of Complex Systems", *Systems Research and Behavioral Science*, 34:5, John Wiley & Sons, Ltd, doi: 10.1002/sres.2482

Mobus, GE & Kalton, MC (2015) *Principles of Systems Science*, Springer, New York

Morowitz, HJ (2002) *The Emergence of Everything: How the World Became Complex*, Oxford University Press, Oxford UK.

Meuleman, L (2008) *Public Management and the Metagovernance of Hierarchies, Networks and Markets: The Feasibility of Designing and Managing Governance Style Combinations*, Springer, New York. ["how do successful public managers design and manage combinations of hierarchical, network and market governance? In other words: what is their rationale to apply metagovernance?"]

Newman, F, Couturier, L, & Scurry, J (2004) *The Future of Higher Education: Rhetoric, Reality, and the Risks of the Market*, Jossey-Bass, San Francisco CA

Rhodes, FHT (2001) The Creation of the Future: The Role of the American University, Cornell University Press, Ithaca NY

Rosen R (1985) *Anticipatory Systems: Philosophical, Mathematical and Methodological Foundations*, Pergamon Press, Oxford UK

Schwarz, E (1997). "Toward a holistic cybernetics: from science through epistemology to being", *Cybernetics and Human Knowing*, 4, 17-50

Schwarz, E. (1992) "A Generic Model for the Emergence and Evolution of Natural Systems toward Complexity and Autonomy", *Proceedings of the 36th Annual Meeting of the International Society of Systems Science,* Denver, Vol. II.

Simon, HA (1957) "A Behavioral Model of Rational Choice", in *Models of Man, Social and Rational: Mathematical Essays on Rational Human Behavior in a Social Setting*. Wiley, New York.

Simon, HA (1991) "Bounded Rationality and Organizational Learning", *Organization Science*, 2 (1): 125–134.

Simon, HA (1998). *The Sciences of the Artificial, Third Edition*, The MIT Press, Cambridge MA.

Smith, E and Morowitz, HJ (2016) The Origin and Nature of Life on Earth: The Emergence of the Fourth Geosphere, Cambridge University Press, Cambridge UK

Türke, R-E (2008), *Governance - Systemic Foundation and Framework*, Springer, New York. [new organizations and technologies for more effective governance (eGovernance) but requiring new theoretical concepts]

Varela, F (1991), *The Embodied Mind*, MIT Press, Cambridge, MA.

von Foerster, H (1974), *Cybernetics of Cybernetics*, Urbana Illinois: University of Illinois. OCLC 245683481

von Foerster, H (1981), *Observing Systems*, Intersystems Publications, Seaside, CA.

von Foerster, H (2003), *Understanding Understanding: Essays on Cybernetics and Cognition*, Springer-Verlag, New York.

von Neumann, J and Morgenstern, O (1944) *Theory of Games and Economic Behavior,* Princeton University Press, Princeton NJ.

Wiener, N (1950) *The Human Use of Human Beings*, The Riverside Press (Houghton Mifflin Co.), New York

Wiener, N (1961) *Cybernetics: Or the Control and Communication in the Animal and the Machine* 2nd ed., Martino Fine Books, Eastford, CT

Ivenenko, VI (2009) *Decision Systems and Nonstochastic Randomness*, Springer, New York
Decision system and control system are the same (chapter 2 page 9)

von Foerster, H. (2003). Cybernetics of cybernetics. In *Understanding understanding* (pp. 283-286). Springer New York

Ostrom, Elinor. 2010. "Beyond Markets and States: Polycentric Governance of Complex Economic Systems." *American Economic Review*, 100(3): 641-72.

https://en.wikipedia.org/wiki/Elinor_Ostrom

1969. The Sciences of the Artificial. MIT Press, Cambridge, Mass, 1st edition. Made the idea easy to grasp: "objects (real or symbolic) in the environment of the decision-maker influence choice as much as the intrinsic information-processing capabilities of the decision-maker"; Explained "the principles of modeling complex systems, particularly the human information-processing system that we call the mind"

- 3rd ed. in 1996, MIT Press.

https://en.wikipedia.org/wiki/Second-order_cybernetics

https://en.wikipedia.org/wiki/Management_cybernetics

Neuroimage. 2008 May 1; 40(4): 1460–1468.
doi: 10.1016/j.neuroimage.2008.01.032
PMCID: PMC2330065
PMID: 18321731
Marieke LS, Howarth C, & Attwell,D (2008), "The cortical energy needed for conscious perception", Neuroimage. 40(4): 1460–1468.