The Global Superorganism:
An Evolutionary-cybernetic Model
of the Emerging Network Society

Francis Heylighen
CLEA, Vrije Universiteit Brussel

ABSTRACT

The organicist view of society is updated by incorporating concepts from cybernetics, evolutionary theory, and complex adaptive systems. Global society can be seen as an autopoietic network of self-producing components, and therefore as a living system or 'superorganism'. Miller's living systems theory suggests a list of functional components for society's metabolism and nervous system. Powers' perceptual control theory suggests a model for a distributed control system implemented through the market mechanism. An analysis of the evolution of complex, networked systems points to the general trends of increasing efficiency, differentiation and integration. In society these trends are realized as increasing productivity, decreasing friction, increasing division of labor and outsourcing, and increasing cooperativity, transnational mergers and global institutions. This is accompanied by increasing functional autonomy of individuals and organisations and the decline of hierarchies. The increasing complexity of interactions and instability of certain processes caused by reduced friction necessitate a strengthening of society's capacity for information processing and control, i.e. its nervous system. This is realized by the creation of an intelligent global computer network, capable of sensing, interpreting, learning, thinking, deciding and initiating actions: the 'global brain'. Individuals are being integrated ever more tightly into this collective intelligence. Although this image may raise worries about a totalitarian system that restricts individual initia-

tive, the superorganism model points in the opposite direction, towards increasing freedom and diversity. The model further suggests some specific futurological predictions for the coming decades, such as the emergence of an automated distribution network, a computer immune system, and a global consensus about values and standards.

INTRODUCTION

It is an old idea that society is in a number of respects similar to an organism, a living system with its cells, metabolic circuits and systems. In this metaphor, different organisations or institutions play the role of organs, each fulfilling its particular function in keeping the system alive. For example, the army functions like an immune system, protecting the organism from invaders, while the government functions like the brain, steering the whole and making decisions. This metaphor can be traced back at least as far as Aristotle (Stock 1993). It was a major inspiration for the founding fathers of sociology, such as Comte, Durkheim and especially Spencer (1969).

The organicist view of society has much less appeal to contemporary theorists. Their models of society are much more interactive, open-ended, and indeterministic than those of earlier sociologists, and they have learned to recognize the intrinsic complexity and unpredictability of society. The static, centralized, hierarchical structure with its rigid division of labor that seems to underlie the older organicist models appears poorly suited for understanding the intricacies of our fast-evolving society. Moreover, a vision of society where individuals are merely little cells subordinated to a collective system has unpleasant connotations to the totalitarian states created by Hitler and Stalin, or to the distopias depicted by Orwell and Huxley. As a result, the organicist model is at present generally discredited in sociology.

In the meantime, however, new scientific developments have done away with rigid, mechanistic views of organisms. When studying living systems, biologists no longer focus on the static

structures of their anatomy, but on the multitude of interacting processes that allow the organism to adapt to an ever changing environment. Most recently, the variety of ideas and methods that is commonly grouped under the header of 'the sciences of complexity' has led to the understanding that organisms are self-organizing, adaptive systems. Most processes in such systems are decentralized, indeterministic and in constant flux. They thrive on 'noise', chaos, and creativity. Their collective intelligence emerges out of the free interactions between individually autonomous components. Models that explain organisation and adaptation through a central, 'Big Brother'-like planning module have been found unrealistic for most systems.

This development again opens up the possibility of modelling both organisms and societies as complex, adaptive systems (CAS). Indeed, the typical examples studied by the CAS approach (Holland 1992, 1996) are either biological (the immune system, the nervous system, the origin of life) or social (stock markets, economies [Anderson, Arrow, and Pines 1988], ancient civilisations). However, this approach is as yet not very well developed, and it proposes a set of useful concepts and methods rather than an integrated theory of either organisms or societies.

The gap may be filled by a slightly older tradition, which is related to the CAS approach: *cybernetics and systems theory*. Although some of the original cybernetic models may be reminiscent of the centralized, hierarchical view, more recent approaches emphasize self-organisation, autonomy, decentralization and the interaction between multiple agents. Within the larger cybernetics and systems tradition, several models were developed that can be applied to both organisms and social systems: Miller's (1978) living systems theory, Maturana's and Varela's (1980, 1992) theory of autopoiesis, Powers' (1973, 1989) perceptual control theory, and Turchin's (1977) theory of metasystem transitions.

These scientific approaches, together with the more mystical vision of Teilhard de Chardin (1955), have inspired a number of authors in recent years to revive the organicist view (de Rosnay 1979, 1986, 2000; Stock 1993; Russell 1995; Turchin 1977, 1981; Chen and Gaines 1997). This gain in interest was triggered in particular by the spectacular development of communication networks, which seem to function like a nervous system for the social organism. However, these descriptions remain mostly on the level

of metaphor, pointing out analogies without analyzing the precise mechanisms that underlie society's organism-like functions.

The present paper sets out to develop a new, more detailed, scientific model of global society which integrates and builds upon these various approaches, thus updating the organicist metaphor. The main contribution I want to make is a focus on the process of evolution, which constantly creates and develops organisation. Because of this focus on on-going development, the proposed model should give us a much better understanding of our present, fast changing society, and the direction in which it is heading. The 'cybernetic' foundation in particular will help us to analyze the increasingly important role of information in this networked society.

The main idea of this model is that global society can be understood as a superorganism, and that it becomes more like a superorganism as technology and globalization advance. A superorganism is a higher-order, 'living' system, whose components (in this case, individual humans) are organisms themselves. Biologists agree that social insect colonies, such as ant nests or bee hives, can best be seen as such superorganisms (Seeley 1989). If individual cells are considered as organisms, then a multicellular organism too is a superorganism. Human society, on the other hand, is probably more similar to 'colonial' organisms, like sponges or slime molds, whose cells can survive individually as well as collectively. Unlike social insects, humans are genetically ambivalent towards social systems, as illustrated by the remaining conflicts and competition between selfish individuals and groups within the larger society (Heylighen and Campbell 1995; Campbell 1982, 1983).

The issue here, however, is not so much whether human society *is* a superorganism in the strict sense, but in how far it is useful to model society *as if it were* an organism. This is what Gaines (1994) has called the 'collective stance': viewing a collective as if it were an individual in its own right. My point is that this stance will help us to make sense of a variety of momentous changes that are taking place in the fabric of society, and this more so than the more traditional stance which views society merely as a complicated collection of interacting individuals (cf. Heylighen and Campbell 1995). More generally, my point is that both societies and biological organisms can be seen as special cases of a more general category of 'living' or 'autopoietic' systems that will be defined further on.

The paper will first try to determine what it exactly means for a system to be an 'organism', and look in more detail at two essential subsystems of any organism: metabolism and nervous system. It will then argue that society's metabolism and nervous system, under the influence of accelerating technological change, are becoming ever more efficient and cohesive. This evolution will in particular give rise to the emergence of a 'global brain' for the superorganism. Finally, the paper will try to look at some of the radical implications of this development for the future.

SOCIETY AS AN AUTOPOIETIC SYSTEM

If we want to characterize society as a living system, we will first need to define what life is, in a manner sufficiently general to be applicable to non-DNA-based systems. Perhaps the best abstract characterization of living organisation was given by Maturana and Varela (1980, 1992): *autopoiesis* (Greek for 'self-production'). An autopoietic system consists of a network of processes that recursively produces its own components, and thus separates itself from its environment. This defines an autopoietic system as an autonomous unit: it is responsible for its own maintenance and growth, and will consider the environment merely as a potential cause of perturbations for its inner functioning. Indeed, a living cell can be characterized as a complex network of chemical processes that constantly produce and recycle the molecules needed for a proper functioning of the cell.

Reproduction, which is often seen as the defining feature of life, in this view is merely a potential application or aspect of autopoiesis: if you can produce your own components, then you can generally also produce an extra copy of those components. Reproduction *without* autopoiesis – which can be designated more precisely as *replication* – does not imply life: certain crystals, molecules and computer viruses can replicate without being alive. Conversely, autopoiesis without reproduction *does* imply life: you would not deny your childless aunt the property of being alive because she is no longer capable of giving birth.

Taking autopoiesis rather than reproduction as a defining characteristic removes one major obstacle to the interpretation of societies as living: although societies generally do not reproduce, they undoubtedly produce their own components. The physical components of society can be defined as all its human members to-

gether with their artefacts (buildings, cars, roads, computers, books, etc.). Each of these components is produced by a combination of other components in the system. People, with the help of artefacts, produce other people, and artefacts, with the help of people, produce other artefacts. Together, they constantly recreate the fabric of society. (To the non-human components of society we may in fact add all domesticated plants and animals, that is to say, that part of the global ecosystem whose reproduction is under human control. As human control expands, this may come to include the complete biosphere of the Earth, so that the social superorganism may eventually encompass Gaia, the 'living Earth' superorganism postulated by some theorists.)

These processes of self-production clearly exhibit the network-like, cyclical organisation that characterizes autopoiesis (see Fig. 1): a component of type a is used to produce a b component, which is used to produce a c, and so, on, until a z is again used to produce an a.

Although societies rarely reproduce, in the sense of engendering another, independent society, their autopoiesis gives them in principle the capacity for reproduction. It could be argued that when Britain created colonies in regions like North America and Australia, these colonies, once they became independent, should be seen as offspring of British society. Like all children, the colonies inherited many characteristics, such as language, customs and technologies, from their parent, but still developed their own personality. This form of reproduction is most similar to the type of vegetative reproduction used by many plants, such as vines and grasses, where a parent plant produces offshoots, spreading ever further from the core. When such a shoot, once it has produced its own roots, gets separated from the mother plant, it will survive independently and define a new plant. Thus, the growth of society is more like that of plants than like that of the higher animals that we are most familiar with: there is no *a priori*, clear separation between parent and offspring. As we will discuss further, in the present globalized world geographical separation is no longer sufficient to create independence. Yet, we could still imagine global society spawning offspring in the form of colonies on other planets.

A society, like all autopoietic systems, is an open system: it needs an input of matter and energy (resources) to build its components, and it will produce an output of matter and energy in the

form of waste products and heat. In spite of being thermodynami-cally open, an autopoietic system is *organisationally closed*: its or-ganisation is determined purely internally. The environment does not tell the system how it should organize itself; it merely provides raw material. The autopoietic system contains its own knowledge on how to organize its network of production processes. *Closure* means that every component of the system is produced by one or more other components of the same system. No component or sub-system of components is produced autonomously. If it were, the subsystem would itself constitute an independent autopoietic sys-tem, instead of being merely a component of the overall system.

This requirement of closure is perhaps what makes the appli-cation of autopoiesis to social systems so controversial. Closure distinguishes what is inside, part of the system, from what is out-side, part of the environment. Maturana and Varela's (1980) origi-nal definition of autopoiesis adds to this that an autopoietic system should produce its own *boundary*, that is, a spatial or topological separation between system and environment. Unlike biological or-ganisms, most social systems do not have a clear spatial boundary. Moreover, for most social systems the closure requirement is only partially fulfilled. For example, a country may produce most of its essential components internally, but it will still import some or-ganized components (people, artefacts) or knowledge from outside. This means that any boundary we could draw around a social sys-tem will be porous or fuzzy. The only way to fulfill the require-ment of organisational closure is to consider *global society as a whole* as an autopoietic system. None of its subsystems, whether they be countries, corporations, institutions, communities or fami-lies is properly autopoietic. All of them are to some extent depend-ent on outside organisation for their maintenance.

This observation may explain why different authors disagree about whether social systems can be autopoietic. Although Maturana and Varela, the originators of the autopoiesis concept, would restrict it to biological organisms, several others (*e.g.* Luh-mann 1995; Robb 1989; Zeleny and Hufford 1991; see Mingers 1994, for a review) have suggested that social systems can be autopoietic, while disagreeing about exactly which systems exhibit autopoiesis. To me, it seems that the controversy can be resolved by only con-sidering global society, the supersystem which encompasses all other social systems, as intrinsically autopoietic.

The problem of the boundary can be resolved by relaxing the requirement that an autopoietic system should produce a physical boundary in space (like the membrane enveloping cells). Although countries, cities or firms sometimes do produce physical boundaries, such as walls or an 'Iron Curtain', planetary society has no need for such a boundary. Indeed, the Earth on which we live offers its own boundary, consisting of the atmosphere which protects the social organism from cosmic rays and meteorite impact, and the lithosphere, which protects its from the heat and magma inside the planet. If an organism, such as a hermit crab, uses a readily available encasing or shell for its protection, rather than invest effort in producing one of its own, then we can hardly blame it for not being sufficiently autopoietic.

If we take the concept of the boundary in a less literal, not purely physical sense, then society clearly does separate its internal components from the environment. The mechanism an organism uses to distinguish and separate insiders from outsiders is the immune system. The immune system is programmed to recognize and expel all alien material, all 'trespassers' that do not obey the rules of the game. These trespassers may in fact include internally produced components, such as cancer cells, that for some reason have stopped obeying the laws that govern the organisation. Society too has an immune system that will try to control both external invaders (*e.g.* wild animals, infectious diseases, hurricanes, foreign enemies) and internal renegades (*e.g.* criminals, terrorists, computer viruses). Basic components of a society's immune system are the police, justice and army.

Both the greatest strength and the greatest weakness of the concept of autopoiesis is its all-or-none character: a system is either organizationally closed, or it is not; it is either alive, or dead. In practice, the distinction between internally and externally produced organisation is not always that clear-cut. Organisms do not just need raw matter and energy as input: these resources must exhibit some form of organisation. For example, an animal, unlike a plant, cannot produce its components on the basis of air, water and minerals. The resources an animal needs must already have gone through some degree of organisation into complex organic molecules, such as lipids, carbohydrates, proteins and vitamins. Similarly, society is to some degree dependent on organisation in the outside world. For example, our present society is dependent

for furniture and firewood on trees, and is dependent for energy on fossil fuels produced by plants millions of years ago.

This observation suggests that we distinguish *degrees* of auto-poiesis: a system will be more autopoietic if it produces more of its organisation internally, and thereby becomes less dependent on its environment. As we will discuss later, the evolution of society will typically lead to more autonomy and a greater capacity to inter-nally produce organisation with a minimum of external input.

To understand how society achieves autopoiesis, we must look in more detail at how the network of production processes can pro-duce a stable organisation, in spite of a variable input of resources and various perturbations in the environment. This mechanism can be functionally decomposed into different tasks to be performed by different subsystems. The most important decomposition is the one distinguishing metabolism, responsible for the processing of matter and energy, and nervous system, responsible for the processing of information. The purpose of both subsystems is to maintain a sta-ble identity by compensating or buffering the effect of perturba-tions. We will now discuss in more detail the different components for each of the subsystems, and the way they are connected.

METABOLISM: PROCESSING OF MATTER-ENERGY

Organisms are *dissipative systems* (Nicolis and Prigogine 1977): because of the second law of thermodynamics, they must export entropy or heat in order to maintain a dynamic steady state. This means that matter and/or energy must enter the system in low en-tropy form (input I in Fig. 1) and leave the system in high entropy form (output O in Fig. 1), after undergoing a number of conver-sions. The entropy that is dissipated or 'wasted' by the system is needed to keep up the cycle of production processes that maintains its organisation.

Although autopoiesis theorists focus on the closed, internal cy-cle of processes inside an organism, the fact that this cycle has an input and an output allows us to make a more of less 'linear' de-composition, which follows matter sequentially from the moment it enters the system, through the processing it undergoes, until the moment it exits. The systems theorist James Grier Miller (1978) has proposed a detailed decomposition scheme which can be used to analyze any 'living system', from a cell to a society. It must be emphasized that such decomposition is *functional*, but not in gen

eral *structural*. This means that the functional subsystems we will distinguish do not necessarily correspond to separate physical components: the same function can be performed by several physical or structural components, while the same component can participate in several functions. Although complex organisms tend to evolve *organs*, *i.e.* localized, structural components specialized in one or a few functions (*e.g.,* the heart for pumping blood), other functions remain distributed throughout the organism (*e.g.*, the immune system).

Since this decomposition does not take into account autopoiesis or organizational closure, Miller applies his living system model also to systems – such as organs or communities – which are organizationally open and which I therefore would not classify as 'organisms'. It seems to me that to fully model organism-like systems, we need to integrate organizational closure with its focus on cycles and thermodynamical openness with its focus on input-output processing (cf. Heylighen 1990). In the following I will discuss the main functional subsystems of an 'organism', using examples both from the animal body and from society. For the societal examples I will focus on artefacts, so as not to repeat the bodily functions that society's human components share with other biological organism.

Table 1

Functional subsystems of the metabolism (processing of matter-energy) in animals and in societies

Function	**Body**	**Society**
Ingestor	eating, drinking, inhaling	mining, harvesting, pumping
Converter	digestive system, lungs	refineries, processing plants
Distributor	circulatory system	transport networks
Producer	stem cells	factories, builders
Extruder	urine excretion, defecation, exhaling	sewers, waste disposal, smokestacks
Storage	fat, bones	warehouses, containers
Support	skeleton	buildings, bridges
Motor	muscles	engines, people, animals

The first function in Miller's model is the *ingestor*, the subsystem responsible for bringing matter and energy from the environment into the system. In Fig. 1, for example, the components *a* and *b*, that directly receive input from the environment, participate in the ingestor function. In animals, this role is typically played by mouth and nose, to swallow food and inhale air. In society, the ingestor is not so clearly localized. Its role is played by diverse systems such as mines and quarries, which extract ores from the soil, water pits, and oil pumping installations. The next processing stage takes place in the *converter*, which transforms the raw input into resources usable by the system. For example, in Fig. 1 insofar that *a* and *b* have not already processed the input they received from the environment, we could situate the converter function in components such as *g* that receive their input from *a* or *b*. In the body, this function is carried out by the digestive system, which reduces diverse morsels of food to simple sugars, fatty acids and amino acids, and by the lungs which ensure that the oxygen fraction of the inhaled air is dissolved into the blood, where it is taken up by the hemoglobin in the red blood cells. In society, the converter function is performed by different refineries and processing plants, which purify water, oil and ores.

A usually subsequent processing stage is transport to those places where the resources are needed. This is the responsibility of the *distributor*. In an autopoietic network such as Fig. 1 all components whose output is similar to their input, but delivered at a different location, can be said to partake in the distribution function. In animals, the distributor function is carried out by the circulatory system: heart and blood vessels. In society, this is the role of the transport system: pipelines, ships, railways, planes, roads. Resources that have arrived at their destination are then processed in order to produce components for the organism. In animals, this *producer* function is carried out by stem cells and glands that produce either other cells or specific chemicals, such as enzymes and hormones. In society, this is done by different plants and factories, producing specialized goods. These products can again be transported by the distributor to wherever they are needed.

One destination where many products end up is *storage*: since the supply of resources from the environment is variable, and internal production cannot always be adjusted to the present need, it

is necessary to have a reserve of resources and products that will help to buffer against fluctuations. In an autopoietic network, components whose output is similar to their input, but delivered at a later time, can be seen as contributing to the storage function. In the body, different organs can fulfill the function of storage for different products. The most general reserve is the one of fat, which can be used as an all-round supply of energy. In society, products are stored in warehouses, silos and containers. Another important destination for products is the *support* function, which physically upholds, protects and separates different parts of the organism. In the body, this function is performed by the skeleton. In society as a whole, which does not have a clear physical structure, the support function is not really needed, but locally it is performed by structures such as buildings, bridges and walls. Another destination is the *motor*, the subsystem that uses energy to generate motion for the organism. In the body, the motor function is performed by muscles, in society by different engines and machines.

Products are typically transformed and recycled into other products. For example, when a cell dies, the lipids that form its membranes will be reused by the body to build other membranes, or stored in fat reserves. In society, the steel of discarded cars will be reprocessed to build cans, steel rods or new cars. Because of the second law of thermodynamics, processes can never be completely reversed: there is always some loss, which is accompanied by the production of entropy. This means that processes will always bring about waste, which cannot be fully recycled. These waste products must be separated from the still usable products and collected. In the body, this is the function of the liver and kidneys, which filter waste products out of the blood. In society, it is carried out by garbage collectors and installations for the treatment of waste. The final matter processing subsystem is the *extruder*, which expulses the waste products out of the system. In the network of Fig. 1, the components *d* and *e* that deliver output straight into the environment, can be seen as part of the extruder. In the body, this function is performed by the urinary tract, the rectum, and the lungs, which get rid respectively of the liquid, solid and gaseous wastes. In society, the respective subsystems are sewers, garbage dumps, and chimneys or exhausts.

NERVOUS SYSTEM: INFORMATION AND CONTROL

Before proceeding with Miller's functional decomposition of information processing, we must discuss the overall function of information in a closed organisation. As Maturana and Varela (1980, 1992) like to emphasize, an autopoietic system is not *informed* by the environment: its form is determined purely by its internal organisation. Autopoietic systems are self-organizing. Data from the environment are only needed to warn the system about perturbations of normal functioning, that may damage or destroy its organisation. By appropriately counteracting or compensating for these perturbations, the system can maintain an invariant organisation in a variable environment (homeostasis).

Thus, organisms are by definition control systems in the cybernetic sense (Ashby 1964): they *regulate* or *control* the values of certain essential variables, so as to minimize deviations from the optimum range. For example, to sustain their intricate organisation, warm-blooded animals must maintain their body temperature within a close range of temperatures (for humans, roughly around 36.5 degrees Celsius). If the temperature of the environment changes, internal processes, such as transpiration or shivering, will be activated to counteract the effect of these perturbations on the internal temperature.

Possibly the clearest overall model of such regulation is proposed by William Powers' (1973, 1989) theory of living control systems. In this model, the behavior or sequence of actions of an organism is explained solely as an on-going attempt to bring the situation perceived by the organism as close as possible to its goal or preferred state ('reference level'). Actions change the state of the environment, and this state is perceived by the organism in order to check in what way it deviates from the goal. The sensed deviation triggers another action, intended to correct the remaining deviation. The effect of this action is again sensed, possibly triggering a further action, and so on, in a continuing negative feedback loop (see Fig. 2). This loop, if it functions well, keeps the system in a remarkably stable state, in spite of the continuous tug of war between environmental perturbations and compensating actions.

Although the resulting state may look largely static, the power exerted to counteract perturbations requires a constant supply of energy. As Powers shows with his mathematical models, an effective control loop is characterized by *amplification*: small deviations must be compensated by relatively large actions. Otherwise, the

result will be merely a give and take between organism and environment, and the result will depend as much on the perturbation as on the action. With large amplification, on the other hand, the result will be much closer to the system's goal than to the external disturbances. In addition to energetic action, such amplification requires very fine-grained, sensitive perception, so that deviations can be detected at the earliest stage where relatively little energy may be sufficient to counteract them.

The different goals or reference levels for the different variables that an organism tries to optimize are typically arranged in a hierarchy, where a combination of perception and a higher level goal determines a goal at the lower level. Thus, goals are not static but adapt constantly to the perceived situation. This perception is not an objective reflection of the state of the environment: it is merely a registration of those aspects of the environment that are relevant to the system's goals, which themselves are subordinated to the overall goal of survival and reproduction of the organisation. Therefore, the epistemology of both autopoiesis theory and perceptual control theory is *constructivist*: an organism's knowledge should not be seen as an objective reflection of outside reality, but as a subjective construction, intended to help find a way to reconcile the system's overall goal of maintaining its organisation with the different outside perturbations that may endanger that goal.

For most non-cyberneticists, the word 'control' connotes the image of a central controller, an autocratic agent that oversees and directs the system being controlled. A cybernetic analysis of the control relation, such as the one of Powers, on the other hand, is purely functional. The 'controller' does not need to be embodied in a separate structural component. In fact, I have argued (Heylighen 1997) that the market can be seen as a *distributed* control system in the sense of Powers. The goal of the market system is to satisfy 'demand', by producing a matching 'supply', in spite of perturbations such as fluctuations in the availability of resources or components. Demand for any particular commodity is itself determined by the overall perception of availability of other commodities and the higher level goals or values (survival, quality of life, ...) of the collective consumer.

This is a negative feedback loop with amplification: small fluctuations in the supply will be sensed and translated into

changes in the commodity's *price*, which is a measure for the difference between supply and demand. Small increases in price (perception) will lead producers to immediately invest more effort in production (action) thus increasing the supply. This will in turn decrease the price, thus reducing the deviation. Similarly, reductions in price will trigger decreased production and therefore decreased supply and increased price. Thus, the market functions to regulate the availability of commodities that the system needs. In spite of this unambiguous control function, no single agent or group of agents is 'in control'. The demand variable, which directs the process, emerges from the collective desire of all consumers, while the supply variable is the aggregate result of all actions by all producers. The control function is not *centralized*, but *distributed* over the entire economic system.

With a few generalizations, this analysis can be developed into a general model for the control mechanism of the social superorganism. Here, Miller's analysis can once more come to our support. Again, we must note that while Miller's functional subsystems are arranged more or less linearly, in the order of processing for information that enters the system, the mechanism as a whole is cyclic: the information that exits the system in the form of actions affects the environment, which in turn determines the information that comes in through perception.

Table 2

Functions of the nervous system (processing of information) in animals and societies

Function	Animal	Society
Sensor	sensory organs	reporters, researchers, etc.
Decoder	perception	experts, politicians, public opinion, etc.
Channel and Net	nerves, neurons	communication media
Associator	synaptic learning	scientific discovery, social learning, etc.
Memory	neural memory	libraries, schools, collective knowledge
Decider	higher brain functions	government, market, voters, etc.
Effector	nerves activating muscles	executives

Miller's first two subsystems fulfill Powers's function of perception: the *input transducer* brings information from the environment into the system, similar to the ingestor bringing matter into the system; the *internal transducer* plays the same role for information originating inside the system. The function of this information is to signal real or potential perturbations away from the goal (dangers, problems), and/or opportunities to achieve the goal (resources, tools). These dangers and opportunities may originate both inside and outside the system, but for simplicity we will discuss them together as if they all come from outside, thus merging 'input transducer' and 'internal transducer' into a single *sensor* function.

This can be motivated by the observation that a truly functional logic implies that we should not consider the actual physical location of a problem or opportunity, but its functional characteristic of being or not being under the control of the system. Remember our discussion of the immune system as the functional 'boundary' of an autopoietic system: internal renegades, such as tumors, are as much perturbations to the system as external invaders, such as pathogenic micro-organisms. Similarly, external extensions of the body, such as clothes, tools or vehicles, are as much under the control of the system as its own components and can therefore be functionally interpreted as parts of the system.

In the body, the sensor function is performed by the sensory organs: eyes, ears, nose, tongue and various cells sensitive to touch, heat and motion in skin, muscles and joints, but also by internal chemoreceptors for hormones, etc. In society, many components take part in sensing: the market, reporters, scientists, polling institutions, voters, and various automatic sensors such as seismographs, thermometers, and satellite sensing installations.

The next information processing function is the *decoder*, which transforms the incoming stimuli into internally meaningful information. In the control system model, this interpretation process functions basically to relate information about the external situation to the system's goals or values, thus making it easier to use this information as a guideline for action. This implies that information irrelevant to any of the system's goals is ignored or filtered out. The decoded information is then used by the *decider* subsystem to select a particular action or sequence of actions in response to the perceived state of the environment. In higher order control sys

tems, where there is a complex hierarchy of goals and subgoals, the actual actions selected may have little to do with the present situation, but rather anticipate potential situations in an as yet far away and uncertain future (cf. Heylighen 1992a). Exploratory behavior is an example of such action that does not seem to have direct relations to the present situation and goal, although in general it helps the system to find new opportunities to achieve its goals. Only urgent danger signals will require immediate counteraction. In animals, both decoder and decider functions are performed by the brain. In society, they tend to be concentrated in political, scientific, legal and commercial institutions, although basic forms of interpretation and decision are distributed throughout the whole of society, as illustrated by market demand 'deciding' which types of commodities are to be produced, or voters deciding which political values should steer the country.

The next step consists in implementing the decision, that is, translating the information generated by the decider into a concrete plan and executing the corresponding actions. This is the task of the *effector* function. This function is absent in Miller's scheme, who proposes the encoder and output transducer functions instead. The reason is that Miller's decomposition hinges on the linear sequence of information entering the system, being processed and finally leaving again, not on the cyclical control function, where the only function of information is to help select the right control action. Although actions may be informative to other systems, they are not generally intended to transmit information, but to compensate for perturbations. Therefore, there is no *a priori* need for the encoding and output of information.

Of course, certain actions, such as speech, have a communicative intention. But this intended transfer of information is subordinated to the more general purpose of achieving the organism's goal. Typical goals of linguistic expression are to make another person do something (a command or request), to get specific information (a question), to get general feedback about one's own state (free expression), or to provide information that might help the other and thus indirectly – through reciprocation, social ties or kinship – help oneself. In such cases, Miller's encoder and output transducer are present as specialized subfunctions of the more general effector function.

In animals, the effector function is performed by motor neurons that activate muscles. In society, it is performed by 'executives' of various ilk, including government ministers, managers, engineers, drivers, and by automated systems that drive machines.

The nervous system also has its analogues of the metabolism's distributor, storage and producer functions. Miller's *channel and net* function is responsible for the communication of information between the various subsystems, such as sensor, decider and effector. In the body, this function is performed by various nerves, in society by communication channels such as mass media, telephone and post. Another destination for information circulating in the system is *memory*, where information about previous interactions is maintained to support future decisions. Unlike the storage function, memory does not simply accumulate incoming data chunks – the way a computer disk records bytes – but maintains a selective, ever adapting trace of correlations between various perceptions and actions so as to increase the effectiveness of decision-making when similar situations are encountered later. The function responsible for creating this network of associations is Miller's *associator*. In animals, associator and memory are distributed over the neurons in the brain. In society, memory is supported by written documents, libraries and databases. The associator function is performed among others by scientists, scholars and archivists.

EVOLUTIONARY DEVELOPMENT
OF THE SUPERORGANISM

Evolution of cooperation

Although Darwinian theory provides a robust model of the evolution of individual organisms, the evolution of societies of organisms does not fit so obviously into that model. The main issue is the tension between individual selection and group selection. Darwinian theory predicts that if an organism is to choose between behavior that will promote its selfish interests, and behavior that benefits the group or society to which it belongs, then in the long term only the selfish behavior tends to be selected. The reason is that selfish individuals in an altruist group ('free riders') profit more from altruist behavior in others than the altruists themselves do. Therefore, altruist behavior tends to be eliminated. Yet, animal

and human groups provide plenty of examples of altruism, *i.e.* behavior that contributes more to the fitness of others than to the fitness of the altruist individual. Several explanations have been offered for this development (see *e.g.* Campbell 1983; Axelrod 1984; Dawkins 1989; Stewart 1997). Since I have discussed this issue in depth elsewhere (Heylighen 1992b; Heylighen and Campbell 1995), the following paragraphs will merely sketch the main arguments.

The mechanism of group selection (groups of altruists being more likely to survive than groups of selfish individuals) seems rather unsatisfactory because of the free rider problem mentioned earlier. Yet, group selection has recently again become more popular (Wilson and Sober 1994), in part because of the observation that not all behaviors beneficial to the group have a high cost to the altruist individual. The most popular explanation, which is at the base of the sociobiological approach, is kin selection: the principle that it is evolutionarily advantageous to be altruist towards individuals that carry the same genes ('kin'). This mechanism seems sufficient to explain insect societies, where all individuals are closely related to each other via their shared mother (the 'queen' of the nest). Another popular mechanism is reciprocal altruism, or 'tit for tat' (Axelrod 1984), but this seems insufficient to explain cooperation in large societies where there is often no opportunity for reciprocation.

To explain the emergence of human society, for me the most compelling mechanism seems to be cultural conformism or 'meme selfishness' (Campbell 1982; Heylighen 1992b; Heylighen and Campbell 1995): if a cultural norm ('meme') prescribing altruism manages to spread over a group, conformist pressures will make it very difficult for would-be 'free riders' to deviate from that norm. Since different groups in general follow different norms, there will be a cultural group selection promoting the more altruist norms, which have the strongest benefit to the group as a whole. Stewart (1997) has proposed a more general mechanism, where a 'manager' (which may be a dominant individual, a subgroup, or a cultural norm) takes control of a group for selfish purposes, to appropriate part of the group's production, but undergoes selection for

promoting altruistic behavior within the group: groups whose manager does not efficiently suppress cheating and free riding will be less productive, and thus their manager will be less fit.

Whichever its precise origin, once a stable pattern of cooperation had been established as a basis for human society, it quickly led to a division of labor. Division of labor is based on the principle that if individuals specialize in carrying out particular tasks, they can be more efficient. However, if an individual is exclusively busy producing one particular type of commodity or service, then that individual will be dependent on reciprocation by others for providing the other resources (s)he needs. Therefore, division of labor can only evolve on a solid basis of cooperation. But once the process has started, division of labor will spontaneously increase, driven by a positive feedback mechanism, as illustrated by Gaines's (1994) computer simulation: individuals who were successful in providing a particular type of service – because of opportunity, competence or simply accident – will get more requests for that type of service, and thus get the chance to develop a growing expertise in the domain. This in turn will increase the demand for their specific service, stimulating them to further specialize. For example, an individual who happens to live near fruit trees may find it easier to make a living by exchanging fruit for meat and other resources than by participating in the communal hunting and gathering, and therefore will tend to invest increasingly more time, attention and resources in developing fruit harvesting capacities.

The increasing division of labor entails an accompanying increase in mutual dependence and therefore cooperativity. *Cooperativity* could be defined positively as probability or dependability of cooperation, and negatively as lack of cheating or free riding. This property of social systems is related to the concept of 'social capital'. It is implicit in the legal system, the organisation of the economy, and the unwritten rules which individuals follow in their interactions with others. For example, a society in which no one trusts anyone and everybody is constantly trying to take advantage of the others without doing anything in return, has low cooperativity. More concretely, the failure of the market to quickly produce economic growth after the fall of communism in the former Soviet states may well be due to a lack of cooperativity in these societies:

without enforceable contracts or fair-play between economic actors, market transactions will in general not bring mutual benefit.

The cooperativity of a society may be estimated by indicators such as level of corruption or crime (negative), and tolerance or trust in other people (positive). These have all significant correlations with the overall quality of life or development level of a nation (Heylighen and Bernheim 2000a), Third World and Eastern European countries scoring in general much lower than North American and Western European societies. This does not mean that there is no cooperation in primitive societies, but only that it tends to be limited to small 'in-groups' such as an extended family, village, or clan, with a lack of care, distrust or even hostility towards outsiders (Campbell 1982). In conclusion we might hypothesize that the evolutionary development of a social system is generally accompanied by an extension of cooperativity, and thus of 'organic cohesion'.

Network evolution

Once there is division of labor, the main engine of evolution in a society will be neither group selection nor individual selection, but what might be called bootstrapping or *network selection*. This mechanism can perhaps be explained most simply in economic terms. An individual or subgroup specialized in supplying a particular commodity can be seen as a subsystem of the overall social system. In return for its 'product', the subsystem receives payment ('reciprocation'), which it invests in resources (components, raw materials, energy, people, information, infrastructure, etc.) necessary for further production. The product defines the subsystem's output, the resources its input. The subsystem's 'function' within the larger whole is to process input into output. The output of one subsystem is used as input by one or more other subsystems, which in turn pass their output to a third line of systems, and so on. Thus all subsystems are linked to each other via the input and output they exchange, together forming a huge network of processes that feed each other, as illustrated in Fig. 1. If the global system is autopoietic, then this network will be largely closed in on itself, and exchange only a limited amount of raw materials and waste with the outside environment. But before we can analyze the evo

lution of the global network, we must examine the evolution of its individual links.

The input it receives and the output it supplies determine a system's relation with its local environment. Systems will be variably adapted to that environment. For example, a system that produces poor output, for which its 'clients' are not willing to pay much, or that cannot get the input it requires, will be ill-adapted. If different systems compete to perform the same function, then those that are best adapted will survive, while the others will be eliminated. Even if only one system is available to perform a given function (*e.g.* a government agency or a commercial monopoly), there will be external pressure on it to improve its efficacy. This means that if the system undergoes variation, the more productive variants will be preferentially retained, while the less productive ones will be pushed to undergo further variation. Thus, all subsystems or components in the network of production processes are under constant pressure to increase their *productivity*, that is, produce more or better output, while requiring less or more easily available input.

This is the mechanism from the point of view of a single system: variation and selection produce ever better fit to the constraints and opportunities of the given, local environment. From the point of view of the entire network, all the components are constantly adapting to each other's input and output. The network as a whole will adapt to its overall input and output, but this will have only an indirect effect on its subsystems. If we ignore this relatively small effect of the global environment, the network's evolution can be seen as self-organisation (Heylighen 2001b), or 'bootstrapping' (Heylighen 2001a): its co-evolving components are mutually adapting, thus increasing the overall efficiency and coherence of the network, without need for external selection. No component can afford to ignore this drive for mutual improvement, since all components are dependent on the others for their input because of their specialization, as discussed earlier. And no component can afford not to specialize, because otherwise it would lose the competition with the more efficient specialists.

The dynamic we sketched applies to all complex systems that can be analyzed as a network of interacting subsystems: markets, ecosystems, organisms, chemical reaction networks, neural net

works, etc. A number of authors in the complex adaptive systems tradition have proposed formal models and computer simulations of different aspects of such network evolution. Kauffman's (1993) models of autocatalytic chemical cycles, and Holland's (1992) 'bucket brigade algorithm' for rules evolving in a cognitive system are worth mentioning in particular. The network systems that interest us here are those that achieve organizational closure. As Kauffman proposes for chemical networks, it is precisely the emergence of closure that characterizes the origin of 'living systems' or organisms. However, since the whole preceding argument assumes that society already has achieved a basic form of closure, we will now focus on the concrete implications of this network dynamic for our present, globalizing society.

Evolution of complexity

The increasing division of labor leads to a *differentiation* of the system into ever more specialized subunits. The increasing dependency of these units on the rest of the system, to compensate for the capacities they lost through specialization, leads to increasing *integration* and cohesion. Differentiation and integration together produce *complexification* (Heylighen 1999a) of the global system, and an ever greater independence from the environment. The positive feedback relation between integration and differentiation leads to the accelerated development of a complex organisation out of an aggregate of initially similar components. This is a *metasystem transition*: the evolutionary emergence of a higher level of cybernetic organisation (Turchin 1977; Heylighen and Campbell 1995; Heylighen 1995; see also Maynard Smith and Szathmary 1995). This overall dynamic is at the base of both the evolution of multicellular organisms out of similar cells and of societies out of individuals. It is in a number of respects similar to the phase transitions, such as crystallization, magnetization or condensation that characterize self-organizing systems in physics.

The ever-accelerating differentiation and integration of societal components is particularly striking in our present age of globalization. It becomes ever more difficult for individuals, groups or countries not to participate in the global economic and political system. If some manage to escape, such as a few primitive tribes

still living in the rain forest, it is largely because global society artificially tries to maintain their niche, as a kind of relic of the past. At the same time, society becomes ever more complex with ever more businesses, organisations and institutions providing ever more diverse goods and services, interacting through ever more wide-spread networks of exchanges and influences, and subjected to ever more intricate systems of standards and rules.

Increasing complexity is merely a side effect of this dynamic, though (Heylighen 1999a): the underlying drive is increasing efficiency or productivity, which itself results from the selective pressure for increasing control, and – most fundamentally – increasing fitness. A component in a societal network will fit its environment better if it can produce more of what is in demand, while being less dependent on the resources it needs as input. In particular, a fit component should be able to provide a dependable output under conditions of variable input, *e.g.* by using reserves, or shifting methods of production to work with different resources. Thus, the component should have good control over its production. The eventual purveyor of the production does not care how it was produced, or what resources or components were used, as long as quantity and quality are satisfactory. This means that if the client can find equivalent products that, because of a different production method, are sold more cheaply, he or she will switch suppliers, possibly *bypassing* a whole chain of production processes. For example, in Fig. 1, the component h may perform the same function for component l as j and k, and therefore l might decide to 'bypass' the longer process $i \rightarrow j \rightarrow k \rightarrow l$, in favor of the shorter process $i \rightarrow h \rightarrow l$.

For a more concrete example, a company that requires quick and constant news feed in general does not care whether this information reaches it via written reports, computer disks or telecommunication networks. If it can get the same information more quickly and cheaply via electronic mail, it will stop its contract with the organisation that can only supply paper reports, thus bypassing a whole production chain that transforms news into printouts, transports these printouts across continents and delivers them on the company's doorstep. Instead of requiring a chain of three organisations, one that collects news, one that prints documents and one that delivers packages, the company will now rely on

a single organisation that directly enters and transmits the news via its computer terminals.

This is an effective simplification of the organisation, by the elimination of processing stages that have become redundant. The more common development, though, will be complexification by the creation of novel products or services. Every demand that is not perfectly satisfied, or every resource that is not completely consumed, determines a *niche* in which a new type of subsystem can potentially make its living. For example, the huge amount of grape seeds left as waste after the production of wine provides a valuable resource for the extraction of bioflavonoids that can help cure circulatory problems. Without the enormous surplus in grape seeds, these valuable medicines might have to be extracted from a scarcer source (*e.g.* leaves of the Ginkgo tree), at a much higher cost. But this new grape seed processing subsystem will itself create a demand for certain products or services (*e.g.* solvents and reactors to extract the flavonoids from seeds), and supply some products waiting for a purveyor (*e.g.* most people are not yet aware of the benefits of grape seed extracts). Thus, the filling of a niche will itself create a number of new niches, providing opportunities for new subsystems to evolve (Heylighen 1999a; Wilson 1992).

Increasing efficiency in the social metabolism

Now that we have a general qualitative understanding of the evolution of a societal system, we can look in more detail at the quantitative evolution of some of its components. As noted, there is a universal selective pressure for subsystems to become more efficient, that is, produce more or better output while using less or more readily available input. This general tendency is easy to observe in society: employees, tools, technologies and organisations become in general more efficient or productive as time goes by. Perhaps the most spectacular illustration of the underlying technological progress is Moore's Law, the observation that the speed of microprocessors doubles every 18 months, while the price halves. This improvement results mainly from miniaturization of the components, so that more (processing power) is achieved with less (materials).

Buckminster Fuller (1969) called this on-going trend to progressively do more with less '*ephemeralization*' (see also Heylighen 1997). Ephemeralization is at the basis of all evolutionary progress (Heylighen and Bernheim 2000b). The increasing productivity means that less resources and labor are needed to produce the same amount of goods or services. It leads to a steadily decreasing importance of *physical* production factors, such as matter, energy, space and time, and a concomitantly increasing importance of *cybernetic* factors, such as information, communication, intelligence, knowledge and organisation, which are necessary to efficiently regulate the processes. Ephemeralization explains the stable or declining prices (corrected for inflation) of raw materials and energy, in spite of largely non-renewable supplies. The decline is particularly evident if the value of a resource is expressed as a percentage of the average income (Simon 1995). The more society develops, the less its members need to spend on physical resources such as food, energy and materials, and the more they tend to spend on non-material ones, such as information, education and entertainment.

The increasing efficiency of society's metabolism is particularly noticeable in the distributor function: the transport of people, goods and services becomes ever faster and less expensive. As a result, distances become ever less important. A few decades ago, intercontinental travel was still a luxury that could be enjoyed by only a select few. Nowadays, people in developed countries routinely fly to other continents for business or leisure. Moreover, the globalization of trade means that increasing amount of goods are transported across the globe by ships, planes or pipelines, at costs too low to make it worthwhile producing the goods locally.

A general characteristic of ephemeralization is that more and more functions are *automated*, that is, human effort is substituted by more efficient technological systems. At first, only physical work was replaced by machines, but recently technology more and more takes over tasks that require intellectual effort. In this domain, there is still a very wide range of possible improvements. For example, apart from pipelines, transport is still largely controlled by humans: truckers, drivers, navigators or pilots. In spite of the great efficiency of modern shopping, having to drive through dense traffic, find a parking space, enter the shop, collect goods,

pay for them at the cashier, load them in the car, and drive them back to your home is still a quite inefficient way to get goods from the distributor to the consumer. At present, great advances are made in electronic shopping so that goods can be ordered and paid automatically, via the computer network. However, this still requires a truck being loaded with the goods, being driven to your home, and being off-loaded there.

In densely populated urban areas, it would seem much more efficient to build an automated distribution network that would connect all homes and warehouses, *e.g.* via tunnels. Containers filled with goods by robots could then be transported on 'conveyor belts' and automatically switched at crossing points in order to reach their destination in the basement of the house from which the order was made. Packaging, used products, and other waste could similarly leave the house, to be transferred automatically to the appropriate recycling installations. This would strongly reduce human effort, traffic congestion, energy usage, and pollution. Of course, building such a network of tunnels under all streets and buildings would demand a huge investment, but it would not be intrinsically more difficult or costly than developing the roads, railways, sewage systems and communication networks that are already there. In the densely populated Netherlands, prototypes of such a distribution system are already being tested out. Such an automated network of tunnels would be a real equivalent of the body's circulatory system.

Reduced friction and it effects on control

The net result of the drive towards increasing efficiency is that matter, energy and information are processed and transported ever more easily throughout the social organism. This can be seen as the reduction of *friction* (Heylighen 2007). Normally, objects are difficult to move because friction creates an opposing force, which dissipates energy and thereby slows down the movement, until standstill. Noise plays a similar role in information transmission: over noisy channels, parts of the message get lost on the way.

The elimination of noise or friction is beneficial for desired processes. However, it can be dangerous when there is a risk for unwanted processes. For example, ice as a surface produces much less friction than earth or asphalt. That is why you can reach higher speeds and sustain them for a longer time when skating than when

running. However, walking on ice is much more difficult and potentially dangerous than walking on asphalt: once you start slipping there is very little to stop the movement getting out of control.

In a similar way, technology smoothens or lubricates all the mechanisms of society. Movements of matter and information run freely, with very little loss or resistance. But this applies to unwanted movements too. It has become much easier to distribute weapons, drugs or poisonous materials, like plutonium or pesticides. Once such a movement gets started, it can develop very quickly, making it difficult to counteract it in time. For example, an infectious disease can spread much more quickly in a world where people travel frequently. Computer viruses are a more modern variant of the same principle: the easier and faster the exchange of information between computers, the more quickly viruses can spread.

The reduction of friction is particularly dangerous for such self-reinforcing processes. The typical example of such a positive feedback process is speculation on the stock exchange, where buying triggers more buying (causing a 'boom') and selling triggers more selling (causing a 'bust'). For that reason, a number of processes in our present low-friction society have become intrinsically less stable, with a higher risk for catastrophic outcomes. This has led to proposals to artificially increase friction, such as restrictions on computerized trade or the imposition of a 'Tobin' tax on international movements of capital.

On the other hand, reduced friction also improves the efficiency of the negative feedback cycles that characterize regulation. It basically leads to greater amplification in the control loop: smaller error signals lead to appropriate reactions more quickly and more easily, correcting the deviation before it has had the chance to grow. Another buying and selling example may illustrate the principle: the negative feedback between supply and demand in a normal market will become more effective when friction is reduced; prices will come down more quickly when supply increases or demand decreases, and an increase in price will more readily trigger an increase in production to fulfill the unsatisfied demand. Thus, it should take less time for the market to come back to equilibrium after a perturbation. Although the changes in the external situation

are so frequent, large and complex that a market will never actually reach equilibrium, reduced friction should at least help it to remain closer to this theoretical state where all labor and resources are optimally allocated to the different demands. This phenomenon may well be at the base of the low inflation and surprisingly stable growth characterizing the 'new economy' which is thought to have been ushered in by global communication networks.

In the longer term, gains in stability due to low-friction, negative feedback are likely to be maintained or amplified, while the loss of stability due to low-friction, positive feedbacks will increase the selective pressure for evolving new regulatory mechanisms. For example, the appearance of a circulatory system in multicellular organisms strongly reduced friction, making it much easier for nutrients and hormonal signals to travel through the body. This gave the organism better overall control over its dispersed tissues and organs, allowing it to become larger and more differentiated. The disadvantage was that microbes too could travel more efficiently. This was compensated by the evolution of white blood cells, allowing the immune functions to keep up with the flow. Similarly, the spread of computer viruses and other threats on the computer network might be stemmed by an artificial immune system, inspired by the mechanisms (such as self-other distinction) used by white blood cells to recognize intruders (Somayaji, Hofmeyr and Forrest 1998).

More generally, the easier flow of goods, people and money in a globalizing society has led to new dangers, such as international crime syndicates, terrorist networks, money laundering, and competition between countries that erodes wages, governments' capacities to collect taxes, and working standards. These can only be solved by the creation of global control mechanisms, such as transnational crime fighting agencies, rules for international bank transfers, and agreements on minimum social standards (*e.g.* prohibiting the import of goods made by child labor). The need to monitor and control such problems is at the base of the increasing importance of global institutions, such as WTO, World Bank, WHO, and UN. Thus, if short-term catastrophes can be avoided, reduction of friction will further enhance the robustness of societal autopoiesis.

A more subtle effect of reduced friction is the *lengthening of cause-and-effect chains*. Imagine a row of billiard balls, each ball a short distance from the next one. If you hit the first ball with your cue (cause), it will hit the second ball (effect), which will itself hit the third ball (second effect), and so on. Because of friction, energy is lost, and each next ball will move more slowly than the previous one, until the point where the ball stops before it has reached the next one in line: the causal chain has broken. With lower friction, the chain will be longer.

This same mechanism applies to our low friction society. In earlier periods, an event happening in one country would have little or no effect on events happening in another country. The chain of effects would have died down long before it would have reached the national border. Nowadays, the world as a whole has become interdependent. A poor harvest in one country will affect the production in another country, which will affect the stock market in a third country, which will affect the employment in yet another country, and so on. Each of these events will not just have one effect, but many, each of which will again influence many other events. The on-going decrease of friction makes these interactions more numerous and more complex.

Because of this ever-greater interdependency, subsystems of the superorganism need to keep informed about an increasing number of events taking place in other subsystems. Longer causal chains mean that many more potential causes and effects need to be monitored, resulting in the concurrent dangers of unpreparedness – when relevant information is not available – and information overload – when the system is incapable to effectively interpret the available information. This augments the selective pressure to develop a more sophisticated capacity for the gathering and processing of information, and therefore to increase the efficiency of the superorganism's nervous system.

Organizational restructuring

Yet another phenomenon affected by reduced friction is the segmentation of the superorganism into cooperating/competing subsystems. First, friction reduction leads to increased 'liquidity' in the markets: capital is more easily available and can more quickly flow

from one investment into another one. This makes it easier to start up new ventures, providing novel products and services. This will accelerate the overall trend of differentiation and innovation, and the emergence of an ever-greater diversity of specialized suppliers.

Second, it leads to increased *competition*: in earlier periods, competition was largely constrained by geographical proximity. A producer or service provider residing in a far-away region could not compete with a local supplier, because of the extra cost of transport of the good or service. The essence of economic globalization is that distance nowadays contributes very little to cost, and therefore competition becomes global. This means that for a given commodity far fewer providers will be able to survive. For domains where size confers advantage, this leads to mergers and acquisitions among firms that provide the same type of commodities. In some sectors (*e. g.* operating systems for computers), this may even lead to the formation of global monopolies.

The concurrent reduction of diversity will be compensated by another trend, though: *outsourcing*. To explain this, we must understand why organisations arise in the first place. An organisation can be defined as a system of individuals with diverse skills and specializations who cooperate for a common purpose. In a pure market logic, it might seem strange that these individuals collaborate in a rigid organizational structure, instead of flexibly providing their services to whoever is the highest bidder at that particular moment. Williamson (1975) and other economists have developed a theory according to which hierarchical organisations arise in order to minimize *transaction costs*. When two components in the societal network engage in the exchange of goods or services, this costs them effort in addition to the effort needed to produce the goods. They need to explore the market, compare different suppliers, compete with others, exchange information about the goods or services they provide or require, build up a relation of trust, sign a contract, establish a channel of exchange, etc. These costs can be minimized by entering into a fixed arrangement, so that the whole process does not need to be started anew each time another exchange is to be initiated. Thus, organisations reduce transaction costs.

Reduced friction (better communication, information processing, more efficient exchanges, etc.) together with increased coop

erativity produces a marked reduction in transaction costs. This means that there will be less need to group a variety of services into a single organisation. Different subsystems of the larger organisation can become more autonomous, exchanging products and services via different flexible channels rather than through a rigid structure. This produces many benefits. Let us illustrate this by the example of a typical organisation: a hospital (cf. Drucker 1993).

The main function of the hospital is curing patients. However, in order to achieve this purpose, it must perform a number of supporting functions, such as maintaining the infrastructure, providing meals to patients, staff and visitors, cleaning the rooms, doing the administration of the bills, etc. Since these various tasks have little to do with the specialized goal of curing patients, the management of the hospital will have little time or interest in overseeing, improving or developing these supporting functions. Yet, the fact that these tasks are immediately needed for the main function means that they tend to be performed by subsystems of the same organisation. The hospital's cooks, technicians and cleaners are employees of the hospital, just like the doctors and nurses.

The reduction of transaction costs means that now a number of these functions can be performed by independent organisations. As long as these organisations have an efficient and reliable communication channel with the hospital, and therefore are able to react quickly and adequately to every demand, they do not need to be physically or organizationally subordinated to the hospital. For example, the cleaning of the hospital rooms can be contracted-out to a specialized cleaning firm. Such a firm can perform this same function for a multitude of hospitals and other organisations, and thereby profit from an economy of scale. For example, it can have a large pool of cleaners and specialized products and machines at its disposal, so that it can quickly respond to special demands, without having to hire more personnel or to order additional tools. Because it is specialized in this one function of cleaning buildings, it can also devote more of its resources to research and development of more efficient cleaning methods than the hospital, thus optimizing its resources. Finally, it can better motivate its personnel to perform their cleaning task, since this is the basic

mission of the organisation, rather than a mere supporting activity that is low on the management's list of priorities.

Such transfer of a particular function to an external organisation is called 'outsourcing'. A similar scenario can be developed for other supporting activities. For example, many organisations now use specialized firms to do their administration and accounting. This is facilitated by modern communication technologies. In principle, a firm selling some product would only need to enter the number of items sold and the money received into a computer network, and, using those data, administration of the accounts could be done by a specialized organisation located anywhere in the world.

The reason why outsourcing increases efficiency is the cybernetic principle of *functional autonomy*. In a complex control system, consisting of many interrelated subsystems performing a variety of tasks, the higher order system cannot oversee the activities of all of its subsystems, since, according to Ashby's (1964) *law of requisite variety*, its own variety (complexity) would need to be at least as great as the variety of all the subsystems combined. In order to minimize the complexity of its own decision-making, the control system should as much as possible *delegate* decision-making to the subsystems, that is, make them autonomously responsible for carrying out their function. The only thing that needs to be controlled is *whether* the subsystems carry out their function; *how* they carry it out is up to them. This defines functional autonomy.

The same principle underlies efficient organisation: the hospital management should not be concerned with the precise way the rooms are cleaned; it should only ensure that they are clean. Therefore, there is no reason for the hospital to keep tight control over the cleaning department: this would only burden its regulatory abilities. The only control it needs is being able to tell the cleaners which rooms should be cleaned and up to what standard of cleanliness. Therefore, it can delegate the implementation of cleaning procedures to an outside organisation.

It is only when the subsystems are incapable of making the right decisions that the higher-order system must intervene and tell them what to do. This means that the more autonomous subsystems are, that is, the more control they have over the way they carry out

their function, the less hierarchical supervision they need. This is the essence of Aulin's (1982) *law of requisite hierarchy*: the required number of hierarchical levels decreases with increasing capacity for control in the subsystems. Since friction decreases and overall efficiency and control increase in all components of society, this means that present organisations and society at large can function with a strongly reduced number of hierarchical levels, thus making perception and action much more efficient (Heylighen and Campbell 1995). This explains the present trend towards the flattening of hierarchies. The towering pyramids of hierarchical levels in traditional bureaucracies merely reflected the poor regulatory ability of the individuals and subsystems in such a bureaucracy (and the tendency to institutionalize an intricate pecking order).

Hierarchies not only tend to flatten, but to turn into *heterarchies*, that is, networks of mutual influence without subordination. Let us go back to the hospital example. Initially, the cleaning department is subordinated to the hospital management. After outsourcing, the specialized cleaning firm may now service many different hospitals, without being subordinated to any one of them. Neither are the hospitals subordinated to the cleaning firm: their relation is one of reciprocity, or exchange among equals. The only system to which these various organisations remain subordinated is society as a whole: only the superorganism can exert general control on the transactions between its subsystems (*e.g.* ensuring fairness and honesty, and precluding the exchange of drugs or nuclear weapons).

In conclusion, the improvement in communication, processing and control in all components of global society has a far-reaching impact on the structure of that society: the number of organisations performing the same function tends to decrease because of mergers and competitive exclusion, whereas the number of organisations performing different functions tends to increase, because of outsourcing, innovation, specialization and the discovery of new niches. At the same time, hierarchies are flattened or turned into heterarchies, and organisations become more autonomous in how they perform their functions, while becoming more dependent on society as a whole to determine which functions are in demand.

Also, organisations become less dependent on specific individuals or geographic regions, and more defined by their activity or function.

Thus, society increasingly resembles a complex organism, with its specialized cells, organs and tissues that are functionally autonomous, but tightly integrated in a global, self-organizing network of mutually feeding processes. This is in clear contrast with the more traditional view of society as a bunch of essentially interchangeable individuals, groups and subgroups, separated by geographic distance or historic accident, that are jostling for power, while making temporary alliances. An important remaining difference is that cells in an organism tend to specialize early and irreversibly, whereas individuals and organisations in society remain able to switch from one function to another as demand or opportunities change, keeping the overall system very flexible.

DEVELOPMENT OF THE GLOBAL BRAIN

Automation of nervous system functions

Whereas the previous section mainly discussed the development of the components and metabolic functions of the social superorganism, the present section will focus on the concomitant evolution of its nervous system, that is to say, the specialized subsystem responsible for the processing of information. Like the other functions, this subsystem becomes more efficient through automation, that is, the use of artifacts such as archives, cables, and computers to extend the capabilities of the human nervous system for the storage, transmission and processing of information. Just like the increasing efficiency of the metabolic functions of production and distribution has led to a globalization of the economy, the automation of information-processing is leading to a globalization of humanity's cognitive and decision-making mechanisms. The most direct support for this global nervous system is the Internet, the network that connects most computers on this planet. The following discussion will focus on the present and future development of this network, arguing that it forms an embryonic 'global brain' for the social superorganism.

The issue here is not the specific implementation of the Internet: most of its functions could probably be implemented in other media and communication protocols, such as cellular phone, digital

TV or rival types of computer networks (such as the no longer existing BITNET or CompuServe systems). The Internet's main strength is its overall flexibility and the fact that it has very quickly become a standard. This made it attractive to integrate competing methods of information exchange into the Internet so as to make them all accessible through a single interface. As a result, the historical accidents that created particular standards for particular types of communication are becoming less and less important in shaping the overall organisation of the global nervous system.

In the society as superorganism metaphor, telecommunication channels play the role of nerves, transmitting signals between the different sensors and effectors (Turchin 1977). In more advanced organisms, the nerves develop a complex mesh of interconnections, the brain, where sets of incoming data are integrated and processed. After the advent in the 19th century of one-to-one media, like telegraph and telephone, and in the first half of the 20th century of one-to-many media, like radio and TV, the last decades in particular have been characterized by the explosive development of many-to-many communication networks. This has led to the metaphor of the worldwide computer network as a '*global brain*' (Russell 1995; Mayer-Kress and Barczys 1995; Heylighen and Bollen 1996; Heylighen 2007).

In organisms, the phylogenetic evolution of the nervous system is characterized by a series of *metasystem transitions* producing successive levels of complexity or control (Turchin 1977; Heylighen 1995). The level where sensors are linked one-to-one to effectors by neural pathways or reflex arcs is called the level of *simple reflexes*. It is only on the next level of *complex reflexes*, where neural pathways are interconnected according to a fixed program, that we start recognizing a rudimentary brain. I will now argue that the present global computer network is on the verge of undergoing similar transitions to the subsequent levels of *learning*, characterized by the automatic adaptation of connections, and *thinking*. Such transitions would dramatically increase the network's power, intelligence and overall usefulness.

The present global network already automates Miller's functions of channel and net (distribution of information), memory (storage of data), sensor (collection of data, *e.g.* through web cameras, keyboards, counters, etc.), effector (use of the net to activate

processes from a distance, *e.g.* remotely controlled robot arms, electronic ordering of goods to be shipped, etc.), and decoder (processing of data to make them more meaningful, *e.g.* 'mining' of client ordering data in order to find relevant patterns) (cf. Heylighen 1999b; Heylighen and Bollen 1996). Apart from memory and channel-and-net, most of these functions are still supported only marginally on the network – at least in comparison to their presence outside the electronic medium. However, it should be clear that the large-scale migration of these functions to the Internet is only a question of time, as there are plenty of benefits and no apparent technical obstacles to the implementation of more sensing, decoding and effecting devices.

Learning and thinking

Less obvious is the automation of the functions of associator and decider, which correspond to the higher cognitive functions that we normally associate with intelligence. Yet, recent work by my colleagues and me (*e.g.* Heylighen 1999b; Bollen and Heylighen 1996, 1998; Goertzel 1999) provides evidence that such forms of high-level, creative intelligence can be directly supported by the network, without need for human supervision. This does not even require sophisticated artificial intelligence programs: it suffices to support the self-organisation of the information streams on the network, thus giving rise to a collective intelligence that is much more than the sum of the individual intelligences of the network's users (Heylighen 1999b). The present paper does not intend to discuss the technical details of a possible implementation, as these can be found elsewhere. It will suffice to outline the general principles, thus showing how the increase of efficiency by automation that accompanies the self-organisation of the superorganism extends to its highest cognitive functions.

The main function of the associator is for the network to learn new associations between data or concepts. The World-Wide Web, through its distributed hypermedia architecture (Heylighen and Bollen 1996), already connects associated documents by a mesh of links. Up to now, the creation of the links is done manually, by the authors of the documents who decide which other documents are relevant to their text. Given the hundreds of millions of potentially relevant documents that are available on the web, this process is

very inefficient, and will catch only a fraction of what is really relevant. This makes it very difficult for a user browsing the web to find the most relevant topics on any given subject. Search engines that return documents containing particular keywords only partially alleviate this problem, as typical keywords will return far too many 'hits', submerging the most interesting documents in noise, while many relevant documents remain elusive because they use different keywords.

In the brain, learning is based on the rule of Hebb for neural networks: neurons that are activated subsequently within a short time interval become more strongly connected. The equivalent of neurons in the web are documents or pages, and the equivalent of subsequent activation is being read or used by the same individual within a short time interval. The more people attentively read document B shortly after they have attentively consulted document A, the stronger the link between A and B should become. For linked documents that are rarely used together, the link should weaken and may eventually disappear. The more people 'surf' the web, moving from page to page by following links or by doing subsequent searches, the faster the web would be able to create good associations, creating strong links between documents that most users would consider mutually relevant. Since every page is indirectly linked to every other page in the web, this means that strongly related pages will sooner or later establish direct links, however far apart they initially were in web space. Such methods would transform the web from a huge collection of weakly connected documents into a coherent associative network, similar to the neural network that constitutes our brain (Bollen and Heylighen 1996, 1998).

Given such associations learned from users, the next function of the associator is to use these links in order to solve problems or answer queries. This process may be called 'thinking'. It can be implemented on the Web by the automation of another neural mechanism: *spreading activation*. If in the brain certain concepts (or the corresponding assemblies of neurons) are activated – because of perception or previous thought about the issue – then this activation will spread to neighboring concepts, following the links in proportion to their associative strength. This will activate new concepts, which in turn may activate further related concepts, and

so on, sustaining a continuous train of thought. Spreading activation can be implemented on the web through a software agent, a program that takes an input of concepts defining the problem, locates the pages most relevant to these concepts, and then explores the links from these concepts, activating neighboring pages in proportion to the initial activation and the strength of the intervening links.

The advantage of this approach is that a problem does not need to be defined by precise keywords, since the activation will automatically spread to pages that contain different keywords but are still closely related to the initial problem formulation. For example, in our prototype learning web (Bollen and Heylighen 1998), the activation of the concepts 'building', 'work' and 'paper' would automatically bring forward 'office' as most relevant concept. This is much closer to the way our brain solves problems by intuition and association than to the way traditional artificial intelligence programs solve problems by logical deduction.

The decider function

The last critical function that needs to be automated is the decider. Given the information about the situation produced by the sensor and decoder, and the overall goals or values of the organism, and using the process of thinking, the decider should be able to select the most adequate sequence of actions that would lead from the initial situation to the goal. We have argued that for the social superorganism, the 'goal' or value system emerges from the aggregate demand by the public. The market is a system of transactions that manages to translate this fuzzily defined 'demand' variable together with the supply into a concrete action signal: the price.

Obviously, network technology can support the market mechanism in determining the optimal price for a commodity. Software agents have been developed to automatically compare prices for any given item in different on-line stores and thus find the best deal for the consumer. This forces suppliers to quickly align their prices downward for goods that are plentiful. On the other hand, automatic auctioning systems (such as eBay) have been created where consumers from all over the world can bid for desirable goods. This forces the price up for scarce goods where the demand is

higher than the supply. The two mechanisms together, one surveying supply and one surveying demand, can accelerate the adjustment of prices so as to optimally reflect the balance between supply and demand. The easy accessibility over the web of prices for the most diverse goods and services will stimulate suppliers to invest in those commodities where the difference between demand and supply is highest.

Not all values can be expressed in terms of price, though. Many valuable things (friendship, ideas, beauty spots...) are free, but still it may be difficult to find them, or to decide which out of several attractive options to choose. But here too the web can support the decision process. The mechanism can be illustrated most simply with documents that offer information. Suppose you find a number of pages that suggest different ways to treat a cold. Which should you take most seriously? In society, this problem is normally solved by relying on *authority*: some sources of information (*e. g.* your doctor, or the medical encyclopedia) are considered more trustworthy than others (*e. g.* your neighbor, or a family magazine). On the web, where the supply of information is huge, extremely diverse in origin, and ever changing, traditional ways of establishing authority (academic degrees, reputation, etc.) are not very efficient. Yet, the linking pattern of the web itself can be used to automatically determine authority.

The main idea is that a document or website is considered authoritative if it is referred to by other pages that are authoritative. Although this definition may seem circular, it can be implemented by a recursive algorithm, which uses a number of iterations to determine the overall authority of a page. Two variations of this approach have been proposed: PageRank (Brin and Page 1998) calculates overall authority, whereas HITS (Kleinberg 1998) determines authority within a specific problem context (*e. g.* all documents about colds and respiratory diseases). Both seem to work surprisingly well in practice, and are likely to work even better if the linking pattern in the web would automatically learn from its usage, as discussed earlier.

In principle, similar algorithms could determine the 'authority' not just of pages or sites but more generally of ideas, people, services or organisations to whom reference is made via the web. As such, they potentially propose an automated means of deter

mining the value of something to its users, without requiring people to offer money for it. However, the disadvantage is that some kind of average value is established for the group, which may be very different from the specific value for a particular user. For example, your taste in music may be very different from the one of the average person, and therefore you would find little value in the list of most popular records. However, you would be interested to hear the recommendations from people whose taste is similar to yours. Such personalized recommendations too can be automated, by using the set of techniques known as 'collaborative filtering' (Shardanand and Maes 1995; Heylighen 1999b). The basic principle is that the system records the personal preferences of a great variety of people, and then uses the preference profiles that are similar to yours to determine options (*e. g.* music records, web pages, movies) that you are likely to appreciate as well. Such a system can even be used to help create personal relations, under the simple assumption that people who have tastes and friends similar to you are likely to get on with you as well.

Merely determining which products or services are valuable is a rather trivial aspect of the decider function. The core of that function is to use data about the perceived situation and about goals and values in order to infer an adequate sequence of actions. This is the most difficult part to implement on the network, although some examples can already be found. For example, a search engine like Google (www.google.com) could be used to enter a number of keywords describing symptoms of a problem. It would find documents that not only discuss these symptoms, but that have a high value according to the PageRank algorithm. Thus, the returned documents would have a high probability to contain a reliable solution to the problem – if such a solution is known, and if the problem is accurately described, using the right keywords. The requirement of accurate keyword description could be relaxed if an algorithm based on spreading activation, as described earlier, would be included into the system.

The problem-solving system would become even more intelligent if the web would be organized in the form of a semantic network, where pages and sections of pages are classified and linked according to an ontology of concept types and link types, as conceived in the new XML standard for the web. This would allow the system to make logical inferences, deducing aspects of the problem

situation that were not entered by the user (Heylighen and Bollen 1996; Heylighen 2001a). For example, if a user would describe the disease symptoms of his poodle, the system would automatically infer that since poodles are dogs, it should look for the same symptoms in documents that describe dog diseases, even if they do not mention poodles.

Ideally, the decider function on the web would connect directly to the effector function, so that actions are not only chosen automatically, but executed as well. For example, a shopping agent might not only gather data about available products and available prices in order to select the 'best buy' option, but might actually order the chosen product.

Integrating individuals into the global brain

Metabolically, most individuals are already strongly integrated into the superorganism: they are wholly dependent on society for shelter, energy, food, water, health and waste disposal. Even the birth of a new human being nowadays is difficult to imagine without a complicated socio-technical infrastructure of hospitals, doctors, nurses and machinery. Intellectually too, individuals get most of their information, knowledge and values from the surrounding social system. However, the latter information exchanges between individual and superorganism are relatively slow and inefficient, at least compared to the speed of the individual's own nervous system. In contrast, the time needed by an individual to get food from the superorganism (*e. g.* by visiting a fast-food restaurant) is not longer than the time needed for that food to be digested by the individual's own metabolism.

This relative inefficiency of information transfer is likely to vanish in the near future. In order to use the cognitive power offered by the global brain effectively, the barrier between internal and external brain should be minimized. The explosive spread of wireless communication, portable devices and, soon, 'ubiquitous computing' (Weiser 1993; Gellersen 1999) heralds the constant availability of network connections, whatever an individual's location. An emerging research domain is that of 'wearable computers' (Starner *et al.* 1997): small but powerful processors which remain available continuously, for example integrated into clothing. Users could wear special glasses that allow them to see the information

from the computer superimposed on a normal view of the sur-roundings. Thus, the computer can constantly provide them with information about the environment, and warn them *e.g.* when an important message arrives.

Such computers would use sophisticated multimedia inter-faces. This would allow them to harness the full bandwidth of 3-dimensional audio, visual and tactile perception in order to communicate information to the user's brain. The complementary technologies of recognition of speech, gestures, facial expres-sions, or even emotional states, make the input of information by the user much easier. For example, the wearable computers would be connected to a small microphone, in which the user can speak, and a glove or sophisticated trackball kept in a pocket, with which the user can steer a cursor or manipulate virtual objects.

Even more direct communication between the human brain and the Web can be conceived. First, there have already been experi-ments (Wolpaw *et al.* 1991) in which people steer a cursor on a computer screen simply by *thinking* about it: their brain waves associated with particular thoughts (such as 'up', 'down', 'left' or 'right') are registered by sensors and interpreted by neural network software, which passes its interpretation on to the computer inter-face in the form of a command, which is then executed. Research is also being done on *neural interfaces*, providing a direct connection between nerves and computer (Lusted and Knapp 1996).

If such direct brain-computer interfaces would become more sophisticated, it really would suffice for an individual to think about a problem in order to see recommended solutions pop-up on the screen (or spoken into an earplug), and the corresponding ac-tions executed in reality. For example, it might suffice for you to think 'It's time to go home' to have a cab automatically directed to the place where you are, pick you up, and bring you to your home address, while being paid from your electronic account, your soft-ware agent having made sure to select the cab that would provide the service most quickly and inexpensively. Thus, the boundary between individual cognitive processes and processes inside the global brain would be minimized, integrating the individual into the superorganism not only physically, but mentally (Heylighen and Bollen 1996).

ISSUES RAISED BY THE SUPERORGANISM MODEL

As noted in the introduction, the view of society as an organism has elicited many worries, objections and general questions that need to be addressed. A discussion of these recurring issues will allow us to argue that global integration is not only likely, but moreover desirable, and, in fact, inescapable.

Totalitarian control, collectivism or freedom?

The most common objection to a superorganism model is that people tend to interpret it as a thinly disguised way of promoting a totalitarian, collectivist system. Especially the use of words such as 'control' and 'collective' evokes immediate associations with Stalinism and the brutal oppression of individual liberties. These negative connotations may be understandable, but they are wholly misdirected. The societal evolution I have sketched is mostly an extrapolation of existing trends, and these show an on-going increase in freedom, individualism, democracy and decentralization rather than a decrease (cf. Heylighen and Bernheim 2000a). These trends can be explained straightforwardly by the postulated mechanisms of differentiation, which opens ever more possibilities for an individual to choose a role, education, or occupation, of reduced friction, which increases the general freedom of movement, of expression, and of consumption, and of increasing autonomy, which reduces the need to tightly control or monitor an individual's activities.

The complementary mechanism of integration could be seen as a source of new constraints or limitations, but these are likely to restrict the freedom of powerful individuals – such as a Stalin-like dictator or a robber baron – and organisations to abuse the system for their own ends, rather than the freedom of ordinary people to realize their individual ambitions. Global integration means an increasing mutual dependency of various organisations, and therefore an increasing difficulty for any one organisation to dominate the others. This is understandably resented by those who have most power to lose, but should be welcomed by the less powerful. For example, this anticipated loss of power may explain the common distrust of global institutions, such as the United Nations, in the presently most powerful nation, the USA.

Historically, totalitarian regimes, such as Hitler's Germany, Stalin's Soviet Union, or Saddam Hussein's Iraq, were the result of

an individual or select group's desire to gain and maintain power and privileges at the expense of the larger population, by suppressing their freedom to question those privileges. The underlying mechanism is simply individual selfishness augmented by social power structures (cf. Heylighen and Campbell 1995). There is nothing particularly modern about such social systems: apart from more sophisticated methods for propaganda and control, the same type of ruthless, centralized organisation can be found in the kingdoms and empires of Antiquity and the period before the French and American revolutions.

Insofar that totalitarian societies were based on an ideological or political system, such as Soviet communism, this system was very different from the self-organizing, cybernetic, 'organism-like' system that this paper proposes. As discussed by the cybernetician and Soviet dissident Valentin Turchin (1981), the Soviet system lacked the most crucial component of cybernetic control: feedback. Instead of a distributed feedback loop, constantly adapting to the changing circumstances, the Soviet economy was based on a rigid, mechanical, top-down command structure, with little regard for the effect of those commands in the real world. This led to the well-known 'calculation problem', where the central planning agency would find it impossible to determine exactly how many shoes would need to be produced to satisfy the needs of a given population. The resulting economic inefficiency contributed to the eventual collapse of the Soviet system.

The emphasis of the present paper on distributed control is not meant to imply that centralization is necessarily bad: concentrating control knowledge in a separate subsystem has a number of benefits (Heylighen 1995). The main advantage is that by giving the control system an explicit, physical form it becomes more open to scrutiny and improvement. For example, the control of a cell is centralized in the DNA in its nucleus. This makes it easy for evolution to try out new forms of organisation by making small changes in the DNA. A cell where control would be distributed over the whole of the participating molecules, as assumed by autocatalytic cycle models for the origin of life (Kauffman 1993), might seem more flexible, but appears less likely to evolve a complex organisation.

Similarly, there is a role in society for some form of government: although the market mechanism can solve many problems, it has certain intrinsic shortcomings (*e.g.* speculative bubbles or disregard for 'externalities', such as pollution) which cannot be corrected by replacing its components (*e.g.* producers or consumers). Since the market reacts as a whole, it can only be steered by an outside system (the government) that imposes constraints (such as regulations, taxes or subsidies) on all participating actors (cf. Heylighen 1997). The advantage of such a separate system is that if it does not function well, it can be replaced by a different one (as when an unsuccessful government is voted out), unlike the global market.

The absence of centralization is at the base of another nightmare vision associated with the superorganism model: the true 'collective', where everybody thinks the same and does the same, and where there is no room for individual initiative or decision-making. This vision is more inspired by insect societies, such as beehives or ant nests, than by existing political systems. Its most popular recent instantiation is the 'Borg', the race of cyborgs imagined by the creators of the science fiction series 'Star Trek'. Again, from a cybernetic point of view a Borg-like organisation would be most inefficient. As noted earlier, Ashby's and Aulin's laws imply that the global organism, in order to maximize its own control over its environment and its chances for survival, should maximize the capacity for autonomous decision-making among its components. Moreover, it should maximize the diversity or variety of the strategies used by its components. This can only be achieved by stimulating individuals to develop themselves freely (cf. Heylighen 1992a), and as much as possible choose their own path, rather than merely conform to the collective point of view.

Even for ants, it can be shown that the colony will be most efficient in finding food if individual ants do not merely follow the paths laid down by their fellow ants, but regularly deviate and create a path of their own (Heylighen 1999b). If people understandably dislike the analogy between human societies and insect societies, it is not so much because insect societies are organized in an intrinsically more totalitarian or collectivist manner, but because insects are simply very dumb, characterless creatures com

pared to humans. An isolated insect, whose behavior is governed by a few simple and rigid rules, is not intrinsically freer or more creative than an insect living in a colony. If you would have to choose, would you rather be a (social) termite, or its individually living cousin, a cockroach? Would you rather be a 'collectivist' bee, or an 'individualist' fly? Neither of these alternatives seems particularly attractive.

Dropouts, conflicts and shared values

Another recurring issue brought forth by the superorganism model is whether all individuals and groups will agree to become part of such a global system. In principle, an individual, nation or group of nations could refuse to be 'integrated' into the transnational social system.

On the individual level, the phenomenon has always existed of tramps, hermits or adventurers who were in practice living 'outside' of society. This phenomenon has always been marginal and is likely to remain so. In principle, there is no reason why the social organism should not tolerate the existence of such individuals or small groups (*e. g.* communes or isolated monasteries) that do not really contribute to society and do not follow its rules. The only condition will be that such outsiders should not harm or endanger those inside, as may be the case for criminals or people with mental disturbances. In practice, though, it seems unlikely that many people would choose that option. The benefits of belonging to society, such as security, comfort, companionship, knowledge, medical support, etc. are so great that it will be very difficult to resist their lure. These benefits are likely only to increase as the superorganism further develops.

On the other hand, the common idea that what you lose in comfort by dropping out, you gain in freedom, is based on a misunderstanding of what 'freedom' means. Without technology and social support systems, life is basically a struggle for survival, where most energy and time must be directed towards finding the necessary food and shelter. By removing these requirements, society has given us the real freedom of doing what we want, where we want it, and (most of the time) when we want it, without having to worry whether we will be able to survive. Especially technology, such as the transport and communication systems, has enormously

expanded our freedom of movement and of communication. The more the superorganism increases its differentiation and integration, the more options we will have to choose our occupation, or go wherever we want whenever we want.

Of course, belonging to an encompassing system does impose certain constraints, aimed at maximizing the synergy between interacting components and minimizing mutual obstruction (Heylighen 2007). However, such constraints do not generally reduce overall freedom. This may be illustrated by the traffic code. Being able to travel with your car wherever you want is a great freedom, which people from previous ages could hardly even imagine. However, for many people to drive safely and with a minimum of interference on the same roads, traffic rules must be followed. Though some of these rules, such as speed limits, are self-evident ways to avoid danger to self and others, others may seem largely arbitrary. For example, there is no intrinsic reason why cars should stop at a red light rather than at a green light, or that they should drive on the right side of the road rather than on the left. Yet, these arbitrary conventions become useful ways of regulating traffic if everybody follows them. Without these seeming restrictions on your freedom to drive on the side of the road that you prefer, driving would become much more difficult and dangerous, effectively limiting your freedom of movement. The freedom lost by following the rules is more than compensated by the freedom gained because of a fluid and safe flow of vehicles.

The problem with such rules is that everybody should agree to follow them. Since the content of these rules is in part arbitrary, different cultures or traditions tend to evolve different rules (cf. Heylighen and Campbell 1995). For example, cars in Great Britain and Japan drive on the left-hand side of the road, unlike cars in most other parts of the world. Changing an established rule is difficult, costly and stressful, and will be resisted by the groups that traditionally follow them – especially if they have (real or imagined) reasons to believe their rules are superior. Yet, global integration entails an eventual harmonization of rules, so as to make the free exchange of goods, services, people and information as fluid as possible. This also implies a reduction of the freedom of certain groups (*e. g*. governments) to set rules that differ from the rules used by the rest. This provides a strong motive for such

groups to resist integration. For example, the European Union, until now the most successful attempt at transnational integration, experiences constant pressures to block harmonization of laws and standards, as illustrated by some countries still refusing to join the common Euro currency.

The possibility seems real that some groups or countries will effectively want to remain outside the emerging global society. It is also conceivable that different federations of countries will be formed, each following their own set of rules, while minimizing exchanges with each other. This happened to some degree during the Cold War when capitalist countries were politically and ideologically separated from the communist block. At the moment, the more important divisions perhaps oppose developed countries and developing nations, or countries with a Christian tradition and countries with an Islamic tradition. A deepening of such divisions could in principle lead to the creation of separate, competing superorganisms.

Yet, there are several reasons why this scenario appears unlikely. The first reason is similar to the reason why individual dropouts are rare. A country that would decide to leave the international community with its systems and rules would immediately lose a great many benefits: resources, products, services, information, new technologies, solidarity, etc. This would significantly slow down or even reverse its development compared to other countries. This can be illustrated by the fate of 'pariah' states, like Iraq under Saddam Hussein, North Korea, or Albania before the fall of communism. Such growing backwardness would provide an increasingly strong incentive for the regime to change its policies.

The negative effects of disconnecting from the rest of the world could be mitigated if a large number of countries would 'drop out' together, forming a rival block. However, the Cold War has shown that two competing blocks, even if they seem roughly matched in size, resources or military power, are unlikely to remain at the same level of development. Because economic and technological progress is an ever accelerating process (Heylighen 2007), small differences in initial conditions or speed of development will lead to increasing gaps, until it becomes clear for everybody that the one block is more successful than the other one. This will put

increasing pressure on the less successful block to open up towards the more successful one, in order to assimilate its successes.

A second reason why a splitting up of the superorganism seems unlikely is the homogenizing effect of global communication on preferences and standards. If people have to choose between competing, but similarly valuable, alternatives, they tend to choose the one they encounter most often. This reinforces the lead of the most common one, in a positive feedback loop (the 'law of increasing returns') that quickly drives out all alternatives but one ('lock-in', see Arthur 1989). In a situation where communication between groups is limited, this may lead to different standards in different groups (cf. Campbell 1982; Heylighen and Campbell 1995), but in an era of fast, global communication, all groups will tend to converge on a single standard in a rather short time.

The third reason is that the basic values underlying different political and ethical systems are effectively universal. The above scenario assumes that the competing options are about equally valuable (*e. g*. VHS *vs*. Betamax standards for video). But what if different cultures or groups disagree about fundamental values? For example, some nations consider capital punishment to be intrinsically barbaric, while others believe that certain crimes must be punished by death. Such differences have led postmodernist thinkers to argue that values are relative or culture-dependent, and therefore there cannot be a rational mechanism for reaching a consensus. Even if we forget about the 'irrational' mechanism of increasing returns discussed above, there are good grounds for consensus. Although different religions and ideologies may disagree about certain concrete dos and don'ts (such as taboos against eating pork, respectively against eating beef), most of their basic values are shared. All ethical systems condemn murder, theft, rape, lying, incest, etc. On the positive side, people from all cultures basically agree about the value of health, wealth, friendship, knowledge, honesty, safety, equality, freedom, etc. In fact, such universal values can be derived empirically, by examining which socio-economic factors correlate with people's happiness or life-satisfaction across different groups (Heylighen and Bernheim 2000a), and supported theoretically, by examining which conditions are conducive to evolutionary fitness on the individual and the social level. The resulting list is remarkably similar to the Universal

Declaration of Human Rights, showing that universal standards can be rationally agreed upon, even though the practical implementation in most cases remains open to discussion. The emerging global network can only intensify and accelerate that on-going discussion.

Another recurrent worry is that the kind of socio-technological developments we have sketched may increase the gap between haves and have-nots, and more particularly between those that have access to information and those that have not, thus creating an 'underclass' of people excluded from the benefits of the superorganism. Although the 'global brain' technologies that we sketched will be adopted most quickly by the wealthiest and best educated populations, this will not stop the poorer regions from joining a little later. Internet technologies are relatively inexpensive to install, compared to *e. g.* roads, electricity or running water, and are becoming ever less expensive. Moreover, as the interface becomes more intelligent, it will become ever easier to use, requiring an ever lower education level for entry. Speech technologies will soon make the web available even for illiterates, and may teach them to read and write in the process. Thus, the emerging global brain is an inexpensive and efficient medium to increase the education level, access to information, and economic competitiveness in all regions of the world, helping Third World countries to bridge the gap with the wealthiest countries.

As Stock (1993) suggests, if regions such as Central Africa still suffer from the wars, famines, epidemics and other atrocities that have become inconceivable in the developed world, it is because the superorganism's nervous and circulatory systems have not yet really implanted in those regions, making them vulnerable to lack of resources, diseases and other perturbations that would otherwise be under tight control. However, as it is in the superorganism's interest to suppress perturbations not only in its core but also in its periphery, there is an on-going pressure to extend these systems even to the most remote regions.

CONCLUSION

This paper has proposed a first sketch of an evolutionary-cybernetic model of society and its development, seen as the emer

gence of a global superorganism. The reasoning underlying the model can be summarized as follows.

Complex systems composed of a variety of interacting subsystems, such as chemical networks, ecosystems, or societies, tend to evolve towards more coherence and interdependence, as the subsystems mutually adapt. This makes the system as a whole less dependent on its environment, and thus increasingly 'closed'. Once there is a sufficient degree of organizational closure, the system can be seen as *autopoietic*, and therefore 'living' in the abstract sense. All such 'living' or 'organism-like' systems combine organizational closure, realized through a network of internal feedback cycles, with thermodynamic openness, entailing the input of low entropy resources and the output of high entropy waste. This allows us to conceptually divide the system into functional components responsible for the different stages of the processing of incoming matter and energy (metabolism), and for the processing of information needed to maintain cybernetic control over this mechanism (nervous system). As the system continues to evolve, on-going adaptation and division of labor lead to an increasingly diverse, complex, and efficient organisation, consisting of ever more specialized components.

This general model of complex, self-organizing systems can be directly applied to the present development of society. Since society is an organism-like system consisting of organisms (individual people), it can be viewed as a 'superorganism'. Conspicuous trends such as globalization, automation, and the rise of computer networks can be understood as aspects of the general evolution towards increasing efficiency and interconnectedness which makes the superorganism ever more robust. In particular, increasing efficiency explains the growing economic productivity and the decrease of friction, which facilitates all material and informational exchanges. The accompanying differentiation and integration explain the seemingly opposite trends towards outsourcing and mergers, and the growing importance of supranational rules, standards and institutions. Increasing efficiency of communication and control moreover explains the increasing functional autonomy of components (individuals or organisations), and the concurrent flattening of hierarchies and rise of heterarchies.

Although the effects of these trends are mostly positive, for both individuals and society as a whole (Heylighen and Bernheim 2000a), some of the side effects can be detrimental (Heylighen and Bernheim 2000b). Reduced friction in particular increases the risk that positive feedback processes would get out of control. It also leads to increasingly complex causal chains of interconnected events, augmenting the need for information gathering and processing. Controlling these dangers requires a strengthening of the superorganism's nervous system. This control system has both centralized and distributed components. Centralization is exemplified by the growing importance of global institutions, responsible for the formulation and implementation of international standards, rules and laws.

Distributed control can be exemplified by the 'invisible hand' that mutually adjusts supply and demand. Its effectiveness is boosted by the emerging global computer network. The increasing reach, capacity, and intelligence of this network allow it to automate more and more functions of the superorganism's nervous system. This will transform the World-Wide Web into a 'global brain', capable of sensing, interpreting, learning, thinking, deciding, and initiating actions. Individuals will become more and more intimately connected to this intelligent network, through ubiquitous, intuitive interfaces, and eventually a direct brain-to-web connection.

The traditional view of society as an organism is controversial, as it seems to imply a restriction of freedom and diversity, and a subordination of individuals to a faceless collective. The present model, on the other hand, sees the emerging superorganism as a further step in the emancipation of humanity, increasing individual autonomy, diversity, and various freedoms of choice, movement, education, career, expression, etc., while decreasing the power of governments, corporations, or dictators to control society for their own purposes. The integration of individuals and organisations into an efficient, coherent supersystem, though, will require the agreement about a number of universal standards and rules for the exchange of goods, services and information. However, because of the greater flexibility and efficiency of a self-organizing, 'global brain'-like system, these rules are likely to be less constraining than existing national laws and regulations, generally in

creasing the diversity of options and freedom of initiative available to individuals.

In conclusion, the picture of an emerging global organism that I have sketched, like the one of Stock (1993), is an optimistic one: although the increasing complexity and accelerating changes that accompany this social evolution may temporarily add to existing stress, conflicts and confusion, overall developments are for the better, increasing people's wealth, freedom, sense of belonging, level of knowledge, equality of opportunity, and overall quality of life (cf. Heylighen and Bernheim 2000a, b), while creating a more flexible, efficient and sustainable society. Moreover, because of the underlying selective pressures and feedback cycles, this development appears quite robust, and can probably be arrested only by a major catastrophe such as a nuclear war or an asteroid impact.

The model throws new light on several contemporary issues such as globalization of markets, computer networks, and the information economy, and thus may help us to understand better what is going on in our complex and rapidly changing society. Moreover, it makes a number of general, qualitative predictions, such as further reduction of friction, restructuring of organisations, long-term improvement of control over the economy, increasing efficiency in production, information processing and services, greater integration and differentiation in the global socio-economic system, and the emergence of a sophisticated collective intelligence for decision-making and problem-solving supported by the computer network.

The question can be raised in how far a true organicist model is really necessary to explain these developments. Most of them could probably be derived from a weaker evolutionary or developmental theory of society or of globalization (cf. Heylighen 2007). The strength of the superorganism model is that it allows a very detailed analysis, zooming in on specialized functional components, such as immune system, distributor, or associator, that have no obvious counterpart in non-living systems. Applying the general logic of network evolution to each of these functions allows us to produce specific predictions, such as the creation of a computer immune system, a fully automatic distribution network, or a world-wide web that autonomously learns new associations.

There is no obvious way to infer such predictions from a more general model, except by including a number of *ad hoc* hypotheses.

Of course, proposing falsifiable predictions is not yet sufficient to make this into a good model: the predictions must also be tested and verified. The problem is that we cannot do experiments with an encompassing system such as global society. We can only wait and observe. It will take many years before any of these predictions can be convincingly confirmed or refuted. In the meantime, the model itself will undoubtedly have evolved, taking into account factors that have been ignored until now. The refutation of any specific prediction should therefore not be interpreted as a falsification of the model as a whole, but rather as an admonition to reflect more deeply about the exceedingly complex interactions within global society. The refutation of several predictions, on the other hand, would be sufficient ground to abandon the model, and look for a better one. Although the time scale is usually the most error-prone aspect of any futurological prediction, I would venture that most of these developments will have taken place within the next 10 to 20 year, whereas the global superorganism itself should have taken a shape clear enough for everybody to recognize it by the next half century.

ACKNOWLEDGMENTS

REFERENCES

Anderson, W., Arrow, K. J., and Pines, D. (eds.)
1988. *The Economy as an Evolving Complex System*. Redwood City CA: Addison-Wesley.

Arthur, W. B.
1989. Competing Technologies, Increasing Returns, and Lock-in by Historical Events. *The Economic Journal* 99: 106–131.

Ashby, W. R.
1964. *Introduction to Cybernetics*. London: Methuen.

Aulin, A.
1982. *The Cybernetic Laws of Social Progress*. Oxford: Pergamon.

Axelrod, R. M.
1984. *The Evolution of Cooperation*. New York: Basic Books.

Bollen, J., and Heylighen, F.

1996. Algorithms for the Self-organisation of Distributed, Multi-user Networks. In Trappl, R. (ed.), *Cybernetics and Systems '96* (pp. 911–916). Vienna: Austrian Society for Cybernetics.

1998. A System to Restructure Hypertext Networks into Valid User Models. *New Review of HyperMedia and Multimedia:* 189–213.

Brin, S., and Page, L.

1998. The Anatomy of a Large-Scale Hypertextual Web Search Engine. Proceedings of the 7[th] International World Wide Web Conference. *Com. Networks* 30: 107–117.

Buckminster Fuller, R.

1969. *Utopia or Oblivion*. New York: Bantam.

Campbell, D. T.

1982. Legal and Primary-group Social Controls. *Journal of Social and Biological Structures* 5(4): 431–438.

1983. The Two Distinct Routes beyond Kin Selection to Ultrasociality: Implications for the Humanities and Social Sciences. In Bridgeman, D. L. (ed.), *The Nature of Prosocial Development: Theories and Strategies* (pp. 11–41). New York: Academic Press.

Chen, L. L., and Gaines, B. R.

1997. A CyberOrganism Model for Awareness in Collaborative Communities on the Internet. *International Journal of Intelligent Systems* 12 (1): 31–56.

Dawkins, R.

1989. *The Selfish Gene.* 2[nd] edition. Oxford: Oxford University Press.

de Rosnay, J.

1979. *The Macroscope*. New York: Harper & Row.

1986. *Le Cerveau Planétaire*. Paris: Olivier Orban.

2000. *The Symbiotic Man*. McGraw-Hill.

Drucker, P. F.

1993. *The Post-Capitalist Society*. HarperCollins.

Fayyad, U. M., and Uthurusamy, R. (eds.)

1995. *Proceedings 1[st] International Conference on Knowledge Discovery and Data Mining*. Menlo Park, CA: AAAI Press.

Gaines, B. R.

1994. The Collective Stance in Modeling Expertise in Individuals and Organisations. *International Journal of Expert Systems* 71: 22–51.

Gellersen, H-W. (ed.)

1999. *Handheld and Ubiquitous Computing.* Proceedings First International Symposium. Springer Verlag, Berlin.

Goertzel, B.

1999. *Wild Computing. Steps toward a Philosophy of Internet Intelligence.* Electronic book, available at http://goertzel.org/ben/wild/Contents.html

Heylighen, F.

1990. A New Transdisciplinary Paradigm for the Study of Complex Systems? In Heylighen, F., Rosseel, E., and Demeyere, F. (eds.), *Self-Steering and Cognition in Complex Systems* (pp. 1–16). New York: Gordon and Breach.

1992a. A Cognitive-Systemic Reconstruction of Maslow's Theory of Self-Actualization. *Behavioral Science* 37: 39–58.

1992b. 'Selfish' Memes and the Evolution of Cooperation. *Journal of Ideas* 2 (4): 77–84.

1993. Selection Criteria for the Evolution of Knowledge. *Proc. 13th Int. Cong. on Cybernetics* Int. Ass. of Cybernetics, Namur, (pp. 524–528).

1995. (Meta)systems as Constraints on Variation. *World Futures: the Journal of General Evolution* 45: 59–85.

1997. The Economy as a Distributed, Learning Control System. *Communication and Cognition- AI* 13: 2–3, 207–224.

1999a. The Growth of Structural and Functional Complexity during Evolution. In Heylighen, F., Bollen, J., and Riegler, A. (eds.), *The Evolution of Complexity* (pp. 17–44). Dordrecht: Kluwer Academic.

1999b. Collective Intelligence and its Implementation on the Web: Algorithms to Develop a Collective Mental Map. *Computational and Mathematical Theory of Organisations* 5(3): 253–280.

2001a. Bootstrapping Knowledge Representations: from Entailment Meshes via Semantic Nets to Learning Webs. *Kybernetes* 30 (5/6), p. 691–722.

2001b. The Science of Self-organisation and Adaptivity. In: Kiel, L. D., (ed.) Knowledge Management, Organizational Intelligence and Learning, and Complexity, in: *The Encyclopedia of Life Support Systems* ((EOLSS), (Eolss Publishers, Oxford). [http://www.eolss.net]

2007. Accelerating Socio-Technological Evolution: from Ephemeralization and Stigmergy to the Global Brain. In Modelski, G., Devezas, T., and Thompson, W. (eds.), *Globalization as an Evolutionary Process: Modeling Global Change* (pp. 1–26)*.* London: Routledge.

Heylighen, F., and Bernheim, J.

2000a. Global Progress I: Empirical Evidence for Increasing Quality of Life. *Journal of Happiness Studies* 1.

2000b. Global Progress II: Evolutionary Mechanisms and their Side-effects. *Journal of Happiness Studies* 1.

Heylighen, F., and Bollen, J.

1996. The World-Wide Web as a Super-Brain. In Trappl, R. (ed.), *Cybernetics and Systems '96* (Austrian Society for Cybernetics) (pp. 917–922). Vienna: Austrian Society for Cybernetic Studies.

Heylighen, F., and Campbell, D. T.

1995. Selection of Organisation at the Social Level. *World Futures* 45: 181–212.

Holland, J. H.

1992. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial Intelligence.* Cambridge MA: MIT Press.

1996. *Hidden Order: How Adaptation Builds Complexity.* Addison-Wesley.

Kauffman, S. A.

1993. *The Origins of Order: Self-Organisation and Selection in Evolution.* New York: Oxford University Press.

Kleinberg, J.

1998. Authoritative Sources in a Hyperlinked Environment. In *Proceedings of the 9ᵗʰ Annual ACM-SIAM Symposium on Discrete Algorithms* (pp. 668–677).

Luhmann, N.

1995. *Social Systems.* Stanford, CA: Stanford University Press.

Lusted , H. S. , and Knapp, R. B.

1996. Controlling Computers with Neural Signals. *Scientific American* 275(4): 82–87.

Maturana, H., and Varela, F.

1980. *Autopoiesis and Cognition: the Realization of the Living*. Dordrecht: Reidel.

1992. *The Tree of Knowledge: The Biological Roots of Understanding.* Rev. ed. Boston: Shambhala.

Mayer-Kress, G., and Barczys, C.

1995. The Global Brain as an Emergent Structure from the Worldwide Computing Network, and its Implications for Modelling. *The Information Society* 11(1): 1–28.

Maynard Smith, J., and Szathmary, E.

1995. *The Major Transitions in Evolution*. Oxford: W. H. Freeman.

Miller, J. G.

1978. *Living Systems*. New York: McGraw Hill.

Mingers, J.

1994. *Self-Producing Systems: Implications and Applications of Autopoiesis*. New York: Plenum Publishing.

Nicolis, G, and Prigogine, I.

1977. *Self-Organisation in Non-Equilibrium Systems.* New York: Wiley.

Powers, W. T.

1973. *Behavior: the Control of Perception*. Chicago: Aldine.

1989. *Living Control systems.* New Canaan, CT.: Benchmark Publications.

Robb, F.

1989. Cybernetics and Suprahuman Autopoietic Systems. *Systems Practice* 2 (1): 47–74.

Russell, P.
1995. *The Global Brain Awakens: Our Next Evolutionary Leap.* Miles River Press.

Seeley, T. D.
1989. The Honeybee Colony as a Superorganism. *American Scientist* 77: 546–553.

Shardanand, U., and Maes, P.
1995. Social Information Filtering: Algorithms for Automating 'word of mouth'. Proceedings of CHI'95. *Human Factors in Computing Systems*: 210–217.

Simon, J. L. (ed.)
1995. *The State of Humanity*. Oxford: Blackwell.

Somayaji, A., Hofmeyr, S., and Forrest, S.
1998. Principles of a Computer Immune System. In *Proceedings of the Second New Security Paradigms Workshop Association for Computing Machinery* (pp. 75–82). New York: Association for Computing Machinery.

Spencer, H.
1969. *Principles of Sociology*. London: MacMillan.

Starner, T., Mann, S., Rhodes, B., Levine, J., Healey, J., Kirsch, D., Picard, R. W., and Pentland, A.
1997. Augmented Reality Through Wearable Computing. *Presence: Teleoperators and Virtual Environments* 6 (4): 386–398.

Stewart, J.
1997. Evolutionary Progress. *Journal of Social and Evolutionary Systems* 20(4): 335–362.

Stock, G.
1993. *Metaman: the Merging of Humans and Machines into a Global Superorganism*. New York: Simon & Schuster.

Teilhard de Chardin, P.
1955. *Le Phénomaine Humain.* Paris: Seuil. Translated as: *The Phenomenon of Man* (1959). New York: Harper & Row.

Turchin, V.
1977. *The Phenomenon of Science. A Cybernetic Approach to Human Evolution*. New York: Columbia University Press.
1981. *The Inertia of Fear and the Scientific Worldview*. New York: Columbia University Press.

Varela, F.
1979. *Principles of Biological Autonomy*. New York: North Holland.

von Bertalanffy, L.

1973. *General System Theory*. New York: George Braziller.

Waldrop, M. M.

1992. *Complexity: The Emerging Science at the Edge of Order and Chaos*. New York: Simon & Schuster.

Weiser, M.

1993. Some Computer Science Issues in Ubiquitous Computing. *Communications of the ACM* 36(7): 75–84.

Williamson, O.

1975. *Markets and Hierarchies*. New York: Free Press.

Wilson, E. O.

1992. *The Diversity of Life*. Cambridge MA.: Belknap Press.

Wilson, D. S., and Sober, E.

1994. Reintroducing Group Selection to the Human Behavioral Sciences. *Behavioral and Brain Sciences* 17 (4): 585–654.

Wolpaw, J. R., McFarland, D. J., Neat, G. W., and Forneris, C. A.

1991. An EEG-based Brain-Computer Interface for Cursor Control. *Electroencephalography and Clinical Neurophysiology* 78 (3): 252–259.

Zeleny, M., and Hufford, K. D.

1991. All Autopoietic Systems Must Be Social Systems. *Journal of Social and Biological Structures* 14(3): 311–332.
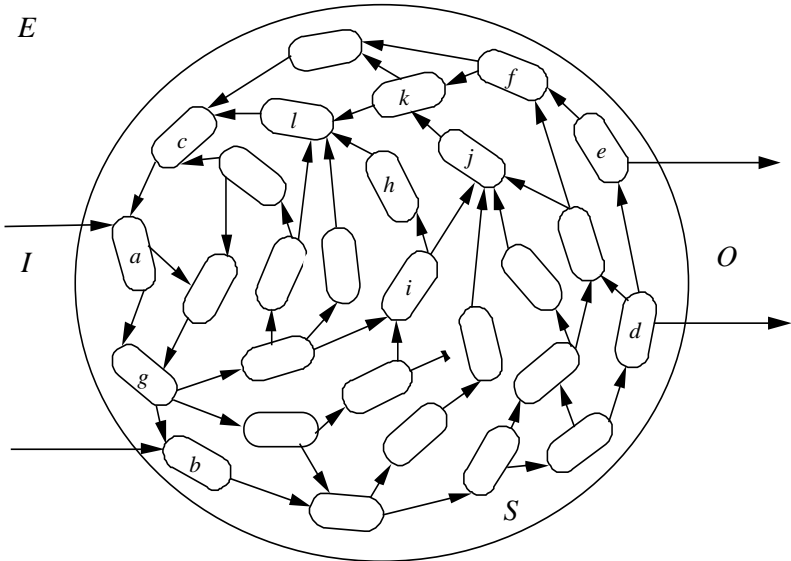
Fig. 1: An autopoietic network. The system *S* consists of a network of components or subsystems {*a, b, c, d*, ...} that are connected to each other via their inputs and outputs, recursively producing its own organisation. For example, component *l* receives input (goods, services, information, ...) from *k* and *h*, processes this input and passes on the resulting product to *c*. The network is mostly closed (the paths connecting components are cyclical) but it still receives some input *I* from the environment *E*, and passes on some output *O* to this same environment. There are in general many redundant or 'parallel' paths that start with the same component (e. g. *i*) and end in the same component (e. g. *l*). In this particular case, the component *h* may be performing the same function for *l* as *j* and *k*, and therefore *l* might decide to 'bypass' the longer process $i \rightarrow j \rightarrow k \rightarrow l$, in favor of the shorter process $i \rightarrow h \rightarrow l$.
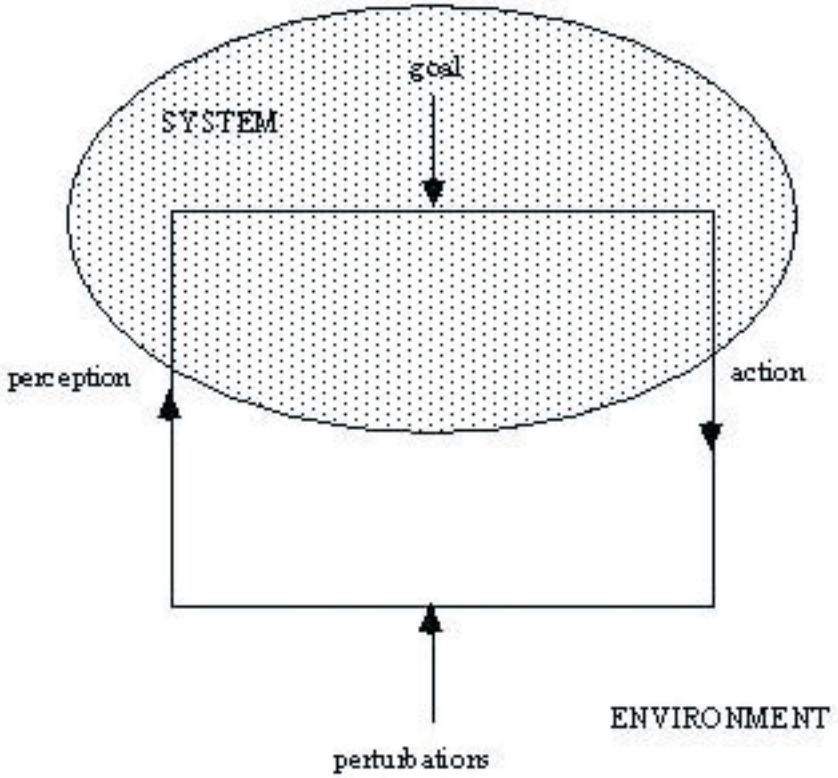
Fig. 2. A control system according to Powers