Cybernetics And Human Knowing. Vol. 11, no. 1, pp. 79-98

The Holonic Enterprise and Theory Emergence:

On emergent features of self-organization in distributed virtual agents

Mihaela Ulieru[1] and Robert A. Este[2]

In this paper we explore the holonic enterprise as an example of the phenomenon of emergence. To conduct this exploration, we describe some of the development work now being carried out in the holonic enterprise field and examine anticipated near-future outcomes of this work. We then provide a brief overview of significant achievements and innovations such as the Copernican revolution, the development of the calculus, and the explication of conceptual analysis as a methodological consequence of the demise of the logicism project, for example. We briefly discuss the paradigm-shifting consequences of these achievements, suggesting that the holonic enterprise development work discussed in this paper as a phenomenon of emergence may have distinct parallels with such well-known innovation events. Developing a framework of 1st, 2nd and 3rd order conceptual analysis based on the demands and circumstances of unfolding technical, political and conceptual change, we then explore some implications of paradigm shifts in general, and speculate on the nature of the holonic enterprise in light of that overview. We suggest that a new 4th order of conceptual analysis may well be necessary to fully understand and embrace the concept of emergence, and finally conclude that the technical and conceptual work being done to create the holonic enterprise necessarily addresses theory emergence.

Keywords: conceptual analysis, emergence, evolution, fuzzy entropy, holistic, holonic enterprise, Internet, multi-agent system, technology, theory, theory emergence

1.0 What is the Holonic Enterprise?

The main idea of the holonic enterprise model addressed in this paper stems from the work of Arthur Koestler (1968). Koestler postulated a set of underlying principles to explain the self-organizing tendencies of social and biological systems, and proposed the term *holon* to describe the elements of such systems. The term is a combination of the Greek *holos*, meaning *whole*, with the suffix *–on*, meaning *part*, as in prot*on* or neur*on*. This term reflects the tendencies of holons to act as autonomous entities, and also to cooperate to form apparently self-organizing system and subsystem hierarchies. The nested arrangement of cell / tissue / organ / system illustrates this notion of a holonic hierarchy in biology; biological and other network and boundary-

---

[1] Director of the Emergent Information Systems Laboratory and Associate Professor of Software Engineering Department of Electrical and Computer Engineering, University of Calgary, Alberta, Canada
Email: ulieru@acs.ucalgary.ca

[2] Ph.D candidate working with the Graduate Division of Educational Research, the Faculty of Continuing Education, and the Department of Electrical and Computer Engineering at the University of Calgary, Alberta.
Email: raeste@enel.ucalgary.ca

related parallels have been explored elsewhere (Pirolli & Card, 1995; Barabási, 2002; Cadenasso, Pickett, Weathers, Bell, Benning, Carreiro et al., 2003).

The context of the holonic enterprise is the tremendous progress that has taken place in information and communication technologies since the turn of the 20th Century. This progress, frequently illuminated by comments about economic implications, has recently been described and received various emphases by Negroponte (1995), Gleick (1999), Davis and Meyer (1998), Axtell (2001), and Kurzweil (1990, 2001, 2003), among many others. Practical articulation of this progress can be seen today with developments such as Bluetooth (Bluetooth SIG, 2003), the Anthill Project (Montresor, Meling, & Babaoglu, 2002; Babaoglu, Meling, & Montresor, 2001), Project JXTA (Sun Microsystems, 2003), and a variety of military and commercial projects and applications (Oakridge National Laboratory/ Southwest Research Institute, 2000; Microsystems Technology Office, 2003), each of which is based on concepts having to do with distributed networked multi-agent systems (MAS) with a variety of purpose-built capacities (sensing, reporting, acting, collaborating) coupled with varying capabilities for hybrid peer-to-peer (P2P) and networked server-based communications. Implications of such ongoing developments of such broad-scale, networked, independent yet collaborative agents for societal functions such as trade, commerce, military applications, energy, transportation, education and entertainment are arguably extremely significant (Broda, 2003).

Such progress in networked systems has transformed our world to an extent to which it can be argued that very large numbers of entities (mostly modeled as software agents) now exist virtually in a universe of networked information, with distinct parallels to the universe we can normally apprehend through natural and extended human senses (Burke, 1999).

Today, the Internet connects such entities by invisible links through their virtual "clones" forming "societies" in which the virtual entities interact autonomously, generating their own interaction patterns and producing complex adaptive behaviours that are not a part of individual entities' programs (Nicolis & Prigogine, 1989; Lewin, 1992; Casti, 1994; Kauffman, 1993, 1996). Implementing an organizational holarchy (modeled after a real-life, complex adaptive system) into software using the multi-agent system (MAS) framework opens the perspective of regarding the Internet as the equivalent of multiple societies of agents comprising a virtual ecosystem that emulates different contexts of the real world, cloned in software according to the abstractions needed for the specific contextual purposes that in turn determine the nature of the holarchy. In such holarchies, enterprises can thus be cloned as self-organized, self-directed agents which can interact and form global virtual organizations.

The holonic enterprise framework allows information and resource management in global virtual organizations by modeling enterprise entities as self-directing software agents linked through the Internet, engaged in integrative adaptive agency (Christensen & Hooker, in press a, in press b). In this universe of information that, when articulated, will interface with (and in some ways parallel and mirror) our own, enterprises enabled with proposed holonic emergence mechanisms will evolve towards increasingly improved structures, while at the same time self-organizing their resources and processes to optimally accomplish their objectives.

Encapsulating the dynamic evolutionary search strategy into a mediator agent and designing the virtual clustering mechanism by a fuzzy entropy minimization strategy (as proposed) empowers the holonic enterprise with self-directing, self-adapting properties which enable it to evolve in cyberspace like a social organism by mating its components with discovered partners in a continuous incremental improvement search process (Ulieru, 2002). In this manner, exchanges of information among agents allow adaptive exploration of many state spaces to achieve optimal outcomes. The holonic enterprise work we are carrying out at present has many substantial implications for enhancing the capacities and efficiencies of economic systems and has the potential to be generalized to a great many other systems of distributed computation.

2.0 Emergence and the Holonic Enterprise

To this point we have suggested that the holonic enterprise as proposed is capable of emergence. However, for the purposes of this paper, a clear notion of what is meant by this term is required. Emergence has been thought to describe the appearance of patterns, structures or properties of systems at the macro-level, where such features stem from or are generated by the dynamical properties of and interactions among system elements and components at the mid- and micro-level of organizations. The term is commonly used today in this way by complexity theorists (Holland, 1992; Kauffman, 1995, 2000; Lissack, 1999; Waldrop, 1992), but has been a topic of thought in a variety of forms since the time of the ancient Greeks. Goldstein (1999), for example, points out that Aristotle successfully dealt with Zeno's paradox by arguing that the whole was greater than the sum of its parts—essentially that not resolving the paradox was a consequence of not recognizing the emergence of the whole.

We might usefully turn to natural living systems for other illustrations of this phenomenon. For example, we can examine the features, extent, nature and overarching characteristics of an ant hill, wasp nest, termite colony, school of fish, or a flock of birds, and in so doing recognize that our thinking of what constitutes the hill, nest, colony, school, or flock connotes the dynamic emergent properties of the countless interactions and transactions that take place over time among all the individual members of the community of interest, all within their unique ecological contexts. Employing our well-schooled tendencies for scientific reductionism, close examination of such examples has permitted recognition, exploration, highly detailed representation, and simulation of the relatively simple operational programs determining the behaviour of individual agents in these aggregates.

What is interesting here is that when we remain in our reductionist mode and only focus on and analyze the features, characteristics and behaviours of the individual agents and limit our view or thinking about the whole, or how the whole comes to be, we cannot predict much, if anything, about the whole at all. This is in keeping with Tasaka's (1999) observation that something is lost when an object is reduced to its component parts. Another way of making this point is: if we dissect and analyze an individual ant down to the smallest detail (for example, we might be able to intimately know the bioelectrochemical neuronal pathways and processes that comprise an individual ant's nervous system, and thereby know the individual "agent ant" program)—and even if we also know that many millions of such individual ants make up what we see as the ant hill—by virtue of relying only on the detailed individual neuronal mapping which we have successfully carried out, we cannot then predict or know the macro-scale behaviours, features, characteristics, and capacities of the ant hill (Kennedy & Eberhart, 2001). In other words, standard analysis and synthesis founded on what Schwarz (2002b) calls the "empirico-analytical paradigm," although helpful, will not permit a full understanding of emergence. The only way we can know *something* about what *may* emerge is to run individual agent programs together in relatively large numbers, so that their aggregate behaviours become holonic, autocatalytic and generative—that is, so that they produce macro-scale behaviours, features, and characteristics that are not in any way written into or known in advance as a part of the micro-scale programs (Theraulaz & Deneubourg, 1992; Holland, 1995). This way of thinking about how the whole comes to be is echoed by Schwarz (2002a) who proposes a metamodel that does not account for pre-existing objects (the things that are known and reasonably well understood; that is, the equivalents of the above-mentioned micro-scale programs) but rather accounts for entire systems that consist of objects, relations, and wholes. These are the complex notions that underpin the concept of emergence.

Emergence has therefore been considered an essential concept in fields of inquiry where a variety of antecedent elements interact and combine in novel ways to generate new consequences or phenomena, things that could not be predicted in advance by virtue of what was known about the constituent parts. Emergence has thus been commonly viewed in terms of the unfolding macro-scale dynamics of interactions among relatively simple system components—the formative interplay between the parts and the whole that generates new features, characteristics, or behaviours that could not be known or determined beforehand, and for which we may not even now have adequate models to permit comprehensive understanding of such phenomena (Cilliers, 2002; Gervasi & Prencipe, 2003).

If we think of how the holonic enterprise is being considered in this paper—a self-organizing hierarchy of systems and subsystem elements or agents—we may usefully attempt to determine if this enterprise, too, can be thought of as an emergent phenomenon, or, as such a thing is sometimes called, "an emergent." How would we know if the holonic enterprise as it is proposed is an emergent? Goldstein (1999; 2000) suggests that all emergent phenomena demonstrate:

- radical novelty—that is, the phenomenon has properties not previously observed, and which could neither be predictable nor deducible from lower, micro-level components

- coherence / correlation—that is, the phenomenon has properties that maintain their identity over time, and correlate lower, micro-level components into a higher, macro-level unity
- global / macro organizational level—that is, the observed behaviour(s) of the phenomenon occurs at the macro, not the micro level
- dynamical—that is, the phenomenon is not predetermined, but arises in terms of new attractors in dynamical systems comprised of dynamically interacting components
- ostensive—that is, the phenomenon shows itself and is recognized

Goldstein further suggests that when viewed through the lenses of complexity theory, emergent phenomena also demonstrate:

- non-linearity—that is, beyond the notion of non-linear positive and negative feedback loops, they also include "small cause, large effect" non-linear events
- self-organization—that is, beyond the notion of simple self-regulation, they also refer to creative, self-generated behaviours that seek adaptation
- beyond equilibrium (multi-, non-, or far from equilibrium)—that is, beyond the notion of homeostasis or "equifinality," to include amplification of random events and dissipative structures in far from equilibrium conditions
- attractors—that is, beyond simple system equilibrium, to include dynamical attractors as features of complex state spaces where concepts such as fitness landscapes successfully account for dynamical system behaviours

As a networked complex system, the holonic enterprise as proposed is chaotically non-linear, self-organizing and adaptive, beyond equilibrium, and features dynamical attractors on a variety of interconnected and evolving fitness landscapes. Individual micro-scale agents may be defined by the extent and nature of their relatively small programs, some of which may be "speciated" and specifically "tuned" for particular micro-scale outcomes, but the generative capacities of their resulting system-wide holonic macro-scale characteristics, features and behaviours are not written into those micro-scale, individual agent programs. The holonic enterprise as proposed can therefore be characterized as possessing the features of elements and relations, positive and negative feedback (generating stability and control as well as transformation), order and self-organization, recursivity and self-reference, and finally, self-production or autopoiesis (Schwarz, 1997). As such, the holonic enterprise can be understood to be an emergent phenomenon, capable of emergence; as such, it has the distinctive characteristics of being able to generate radical novelty, coherence and correlation, to operate in a self-directed manner on the macro-level, to be dynamical, and to be ostensive.

3.0 The Mathematics of Emergence

To be able to coherently make our point on this work we will now summarize the results presented in (Ulieru, 2002) on the mathematics of emergence in the holonic enterprise. The main idea is to minimize the entropy in the information spread across the virtual enterprise (modeled as a multi-agent system)—such that each holon maximizes its knowledge of the task it is assigned, in order to best accomplish it. This naturally leads to the (self-)organization of the virtual enterprise in a holarchy (which defines a holonic enterprise).

We consider the holarchy having its resources predefined and represented as software agents at the logical level of the holonic enterprise; therefore, at this level, the holarchy is regarded as a multi-agent system (MAS). Our purpose is to organize the holonic enterprise such that it can accomplish the goal (say manufacturing of a certain product) with minimal cost. This calls for all the resources to be loaded at optimal capacity through a harmonized flow of information and material across the holonic enterprise. Given that we aim to attain and preserve this perfect order in the holarchy it seems natural to attempt this by minimizing the entropy measuring the degree of order in the information spread across the holarchy's resources.

To enable a mathematical formalism that can support this purpose we regard a MAS as a dynamical system in which agents exchange and organize information through reasoning into knowledge about the assigned goal. Optimal knowledge at the holarchy's highest level of resolution (inter-enterprise level) corresponds to an optimal level of information organization and distribution among the agents within all levels of the holarchy. We consider the entropy as a measure of the degree of order in the information spread across the multi-agent system modeling the holarchy. One can envision the agents in the MAS as being under the influence of an information *field* which drives the agent interactions towards achieving *equilibrium* with other agents with respect to this entropy.[3] Does this really need to be set off with extra line spacing?

This information is usually uncertain, requiring several ways of modeling to cope with different aspects of the uncertainty. Fuzzy set theory offers an adequate framework for dealing with this uncertainty. We will therefore use the *generalized fuzzy entropy* to measure the degree of order in the information spread across the holarchy. The generalized fuzzy entropy is the measure of the "potential" of this information field and *equilibrium* for the agents under this influence corresponds to an optimal organization of the information across the MAS with respect to achievement of the assigned goal. When the circumstances change across the holarchy (due to unexpected events, such as need to change a partner that went out of business, machine break-down, raw materials unavailable, etc.) the equilibrium point changes as

---

[3] The information "field" acts upon the agents much in the same manner as the gravitational and electromagnetic fields act upon physical and electrical entities respectively.

well inducing a new re-distribution of information among the agents with new emerging agent interactions.

We start with the assumption that only the set of resources available for the holarchy formation is given (that is, we know the enterprises that will collaborate to accomplish the pre-set goal) and we aim to organize these enterprises such that their resources are optimally used to accomplish the goal most efficiently (minimal cost and time). In short, we have a virtual enterprise (VE) with several distributed partners linked via the dynamic Web and we want to organize it such that it accomplishes a certain goal optimally. Based on these assumptions, we have proven mathematically that the optimal organizational structure of the distributed organization is a holarchy (Ulieru, Stefanou, & Norrie, 2000; Ulieru, 2000), thus proving emergence in the virtual environment. We model agent interactions through fuzzy relations considering that two agents are in relation if they exchange information. As two agents exchanging information are also in the same cluster, one can describe the clustering configurations using these fuzzy relations. In construction of the fuzzy relation, one starts from the observation that associating agents in clusters is very similar to grouping them into *compatibility* or *equivalence classes*, given a crisp (binary) relation between them. A measure that evaluates "the fuzziness" of a fuzzy set by taking into consideration both the set and its fuzzy (non-binary) complement is the *Shannon measure*, derived from the generalized Shannon's function:

$$\left[\begin{array}{l} S:[0,1]^M \rightarrow \Re_+ \\ (x_1,\ldots,x_M) \mapsto S(x) \overset{def}{=} \\ -\sum_{m=1}^{M}\left[x_m \log_2 x_m + (1-x_m)\log_2(1-x_m)\right] \end{array}\right. \tag{1}$$

If the argument of this function is a probability distribution, it is referred to as Shannon entropy. If the argument is a membership function defining a fuzzy set, it is referred to as (Shannon) fuzzy entropy. Denote the fuzzy entropy by S_μ. Then, according to equation (1), S_μ is expressed for all $k \in 1,K$ by:

$$\begin{array}{l} S_\mu(\mathrm{R}_k) = -\sum_{i=1}^{N}\sum_{j=1}^{N} \mathrm{M}_k[i,j]\log_2 \mathrm{M}_k[i,j] - \\ -\sum^{N}\sum^{N}\left[1-\mathrm{M}_k[i,j]\right]\log_2\left[1-\mathrm{M}_k[i,j]\right]. \end{array} \tag{2}$$

Although a unique maximum of Shannon fuzzy entropy (2) exists, we are searching for one of its minima. When the associated fuzzy relation $\mathcal{R}_{k_0}$ is a *similarity* one, then an interesting property of the MAS is revealed: clusters are associated in order to form new clusters, as in a "clusters within clusters" holonic-like paradigm (see Section 5 in Ulieru, 2002). Moreover, a (unique) similarity relation $\mathcal{R}_{k_0}$ can be constructed starting from the proximity relation $\mathcal{R}_{k_0}$, by computing its *transitive closure*. Thus, the potential holonic structure of MAS can be revealed, even when it seems to evolve in a non-holonic manner. When $\mathcal{R}_{k_0}$ is only a *proximity* relation, tolerance (compatibility) classes can be constructed as collections of eventually overlapping clusters (covers). This time, the fact that clusters could be overlapping (i.e., one or more agents can belong to different clusters simultaneously) reveals the capacity of some agents to play multiple roles by being involved in several tasks at the same time.

We further develop in (Ulieru, 2002) an iterative, incremental search strategy for the agents that best fit the optimal configuration (11), that expands the search domain over time. For this we use the property of global optimizer inherent in genetic algorithms. Our construction is based on the observation that the search process in a set of agents is analogous to the genetic selection of the most relevant ones relative to the goal of the multi-agent system. The genetic operators mutation probability and crossover probability in the probabilistic model that generates the new structures (genotypes) in the evolutionary processes can be defined for a population of distributed agents in a dynamic, open search domain in cyberspace. In this way, the most relevant agents with respect to the goal will be naturally selected as "best" through the evolutionary process. The search problem can be described as follows:

(a): Inside a finite agent domain space search for the most relevant P agents according to a predefined goal. The group of found agents represents the initial *population* consisting of P members. They are the phenotype. Any evaluated agent has a numerical index encoding its relevancy with respect to the search context (that is the goal).

(b): As the search space expands, the members of P are changing continuously, those members with low indexes being eliminated and by this making room for new members found to have higher relevancy indexes.

(c): For each agent its set of indexes constitutes the genotype. They are numeric (e.g., degrees by which the agent characteristics and/or behaviors match the overall goal) and represented as binary strings. Theoretically the index *term frequency* is defined (e.g., on the interval 0% to 100%; that means in binary from 000000 to 1100100, therefore seven bits maximum). *Concatenating* the binary domains for all seven indexes we need 49 bits; therefore, in this case, the chromosome will have 49 bits length.

(d) The initial population evolves by reproduction based on the two major genetic operators: *mutation* and *crossover* which are the probabilistic parameters p_m and p_c. Each chromosome of the population (i.e. relevancy index) will be randomly affected. The isomorphic consideration of genetic operators in the context of information search process interprets mutation and crossover operators as modeling the probability of finding software agents inside the partial domain considered at each iteration.

(e) The population's evolution generated from the previous search is controlled by the *selection* mechanism. This is possible by defining a certain *evaluation function* as a selection criterion. The essence of this evolutionary search process stems from the recursive modification of the chromosomes of the concatenated indexes in each generation while monitoring the evaluation function. In each iteration, all members of the current generation are compared with each other. The best results are placed at the top and the worst are replaced with the new members. The subsequent iteration resumes this process on the partially renewed population. The link between the evaluation function and the relevancy is made by the search query criterion which is defined via the fuzzy entropy minimization.

4.0 Outcomes of Emergence

Now that we have examined the mathematical presentation of the holonic enterprise and also determined that the enterprise as proposed is capable of emergence, we shall at this point generally explore the outcomes of emergence. This will allow us to then move on to consideration of paradigm shifts in order to better understand implications of the holonic enterprise in terms of emergence.

Creating the conditions and initial state spaces within which the holonic enterprise can be developed, implemented and diffused arguably has three primary types of anticipated outcomes—the technical, political, and the conceptual. The organizational policy process framework can be used to view these three interrelated types of outcomes (Downey & Este, 1984; Lindblom & Woodhouse, 1993; Sabatier, 1999).

The first and most obvious type of outcome of the holonic enterprise is technical: this is the emergence of an evolutionary, multi-agent, self-organizing cyberspace supported by and expressed through all extant and future computing and networking technologies as has been described in the previous section. The technical consequences of holonic enterprise deployment will build on this foundation and change the nature, utility and effectiveness of the Internet (and all systems that will interact with it) in profound ways, many of which may not be predictable. However, as with other major technical and scientific innovations that have occurred in the past, the political and conceptual outcomes of the proposed holonic enterprise are also extremely significant. They are integral components of the outcome set. Simply put, technical/scientific innovations do not occur in isolation: they are developed, implemented, and diffused in human organizational and societal contexts, and, as Falconer (2002) suggests, occur in an entirely non-linear manner.

Although we may not be able to predict all outcomes, it is the case that we can posit many plausible technical effects of holonic enterprise development. For example, when based on a holarchic web-based system, such things as manufacturing, supply-chain, and distribution network technologies will almost certainly be improved and rendered more efficient. But, can we also know about political and conceptual consequences sufficiently well, clearly, and well enough in advance to be able to meaningfully address questions of what their plausible implications might be, beyond the anticipated economic benefits, to create and implement an emergent, evolutionary, multi-agent, self-organizing cyberspace? Will the political and conceptual outcomes also be characterized as *enhanced*? Do we have the requisite knowledge and skills to know if the unfolding dynamic interactions among the technical, political and conceptual outcome elements will be familiar and which ones will not, which are clear opportunities and which are threats, and especially, what we should expect to do about these things? We can re-state these questions in another way: is the developmental work now being carried out in the holonic enterprise arena the beginning of a deeply profound, multifaceted, and wide-ranging paradigm shift? If so, what can we reasonably anticipate about the extent and nature of this shift, what types of thinking are required to be able to handle such a shift well, and what might the theoretical significance of this shift be?

To speculate meaningfully about these questions, we can usefully examine and briefly reflect upon the consequences of past paradigm shifts, as well as the relationship between the holonic enterprise project and other parallel developments that may have an amplifying role to play in the extent and nature of what will, in fact, emerge.

The next sections of this paper provide a brief overview of innovations that shifted paradigms, examines some present paradigm-shifting innovations, compares the potential impact of the holonic enterprise on current paradigms, and addresses some implications of the future of the holonic enterprise.

5.0 Paradigm Shifts: the past to the present

We have many useful historical examples of the unfolding of deep paradigm shifts (Kuhn, 1996; Conant & Haugeland, 2000). For the purposes of this paper, to begin with we shall delimit our considerations and not address Platonic worldviews beyond recognizing the initial rationalization of philosophy that took place at the time of the early Greeks (Lenzer, 1975). Rather, we shall commence here by noting the renewal of and powerful, extended focus on the scientific method that took place during the European renaissance and enlightenment. The works of Galileo and Copernicus were seminal in this regard, and in essence began the processes of deep paradigmatic change that forever moved the position of humankind from the assumed centre of the universe to that of an increasingly enlightened, questioning, capable (and perhaps occasionally humbled) participant (Kuhn, 1957).

By the late 18th Century, the work of Immanuel Kant created a rigorous philosophical foundation, a worldview based on and generated from a combination of the rationalism of Descartes, the empiricism of Bacon, and the mathematics and physics of Liebniz and Isaac Newton. We recognize here that, with increasingly rigorous methodologies of scientific reductionism and robust calculi in hand, Kant's work at the time still necessarily admitted intuitionism: he recognized, in a manner similar to how Plato had done before him, that some things could not be clearly perceived, analyzed, systematized or understood. It is interesting to consider at this point that Kant may have, in a sense, bumped into a *conceptual ceiling* regarding what we now see as emergence in complex adaptive systems. That is, even though he recognized that burgeoning scientism and analytic investigation could provide increasingly robust and reasonably reliable answers to many important questions, and allow secure and increasingly detailed explication within fields of inquiry that had not been previously illuminated, reductionism (even while creating the foundation for logical positivism) could not explain everything—all could not be known. Some things that could be perceived and thought of were not amenable to the new tools of scientific analysis and reductionism.

In recognizing this, Kant provided a foundation that motivated the later works of Boole, Peano, Frege, and then Russell to expand and enhance the extent and nature of rigorous mathematical logic as the foundation of all human knowing (Dummett, 1991; Russell, 1911) with the goal in mind to create a full logical calculus that could, eventually, explain and completely account for who we are and what we do—starting with language and arithmetic, for example. Here, the inadequacies of Aristotle's syllogistic reasoning that had been hidden for two millennia, revealed through ongoing exploration and discoveries in science, were increasingly overcome and more deeply illuminated through advances in mathematics and logic (Van Heijenoort, 1967; Este, 2003).

Logicians like Russell worked to pursue proof of logical consistency, but even with his (and Whitehead's) pivotal work in *Principia Mathematica*, such proof could not be established. Gödel then demonstrated that such consistency proofs are impossible, namely: (i) that any adequate axiomatizable (meaning the axioms can be computably generated) theory is logically incomplete; and (ii), in any consistent axiomatizable theory which can encode sequences of numbers (with syntactic notions of *formula*, *sentence*, *proof*), the consistency of the system is not provable within that system (Casti & DePauli, 2000). These revelations were devastating to what had become known as the logicism project—but importantly, they did not reduce the significance or importance of the philosophical analysis that Russell, in particular, had spearheaded.

We shall pause here for a moment to catch our breath, and briefly step outside our story. Against this abbreviated backdrop of the emergence of mathematical logic as the foundation for unifying scientific discovery and explanation, we should not forget the surrounding context of technical innovation. Less than three hundred years earlier, the printing press had been invented. The consequent diffusion of literacy and the increasingly broad distribution of print (and thus of logical argumentation) were both central to the spread and sharing of knowledge and development of epistemological frameworks which were becoming increasingly rigorous, systematic and scientifically-based. The point being made here is that the widespread paradigmatic effects of such things as growing scientism, or the logicism project and its inevitable failure, for example, would not have been possible without the consistent application and spread of mass communication technologies. This reminder regarding the diffusion of pivotal communications technologies supporting the global spread of knowledge and paradigmatic shifts will serve us well at a later point in this paper.

Let us return to our overview of the unfolding of events. We should here recall that progress in logic and the philosophy of mathematics did not occur in a conceptual vacuum. Near the turn of the 20th Century in particular, we also see a host of other powerful innovations and discoveries occurring in parallel. To name but three: Einstein clearly demonstrating that Newton's vision of reality—that is, the one we also see ourselves as regularly inhabiting and commonly sharing, with time, gravity, laws of motion and new and useful tools such as the calculus (Berlinski, 2000)—is merely a special case of a relativistic universe; Freud offering explanatory conceptual frameworks to plausibly account at least in part for the dynamic aspects of what we think of as the psyche, and much of human behaviour besides; Darwin, after his long journeys, suggesting that his evolutionary theory could account for the unfolding diversity of life on the planet. No insignificant or isolated illuminations these.

This brief retrospective is one that highlights major conceptual shakeups—a series of profound paradigmatic shifts, alterations in the way people at the time viewed themselves, everything that surrounded them, and especially, their place in the world. It is interesting to note that these shifts began to emerge dramatically in the enlightenment with what we might term a "jump start" convergence of insights, discoveries and innovations, then with increasing in frequency and significance as time progressed, and spreading out in diverse patterns of generative diffusion (Kauffman & Este, 2004).

In thinking about the patterns of such shifts, we can now rush forward into the 20th Century, leave the broad compass of diffusion of innovations behind, and narrow our focus to the realm of computation. We can recognize that based particularly on the work of Gödel and those who followed, and supported by the fine earlier accomplishments of individuals like Babbage (Lee, 1994), the notions of enhanced mechanical and then electrical computing became not only feasible but were recognized as eminently practical. Even though the logicism project of Frege and Russell became more of an historical artifact than a burgeoning program for the generalized application of logical principles, it became increasingly possible to translate the explorations and determinations of rigorous mathematical logic into what we now call software and hardware (Hillis, 1998). This translation, originating in the realms of wetware and pen and paper, moving through mechanical expressions and then into bits and electrons, has indeed taken place, first starting with very simple devices and expanding and diversifying to the present day to include a huge and increasingly large variety of machines and mechanisms that we have devised to accomplish computational work of every description.

Today, diffusion of and increasing inter-relatedness of computational innovations and applications continues to infuse all aspects of how we define and work within our realities. Now, as we pass through the first decade of the 21st Century, we see that we have recently introduced further very significant elements that are analogous in terms of effect to those we have briefly reviewed earlier in this paper—they have the potential to create major paradigmatic shifts; indeed, many developments such as the transistor and then the microchip already have; in combination with waves of parallel innovations, macro-level developments such as the Internet have pushed and ballooned multiple paradigmatic shifts through many dimensions. As a result, we are today surrounded by and immersed in blended advances in sciences and technologies that allow us to successfully explore, address, work within and expand all things computational. We have the rapidly expanding and diversifying realms of nanotechnology, biotechnology, AI, robotics, and networking, for example. All of these are leading to the inevitability of universally distributed, applied, networked, ubiquitous, autonomous, embedded computing.

To this point in this paper we have reviewed a very brief historical explication of the sharpening and ongoing unfolding of mathematical logic leading to our present conceptions of and advances in computation—essentially, the set of lenses that permits us to see, understand, engage with, manipulate and create a multiplicity of interconnected systems. In the limited scope of this paper, much of this story is necessarily omitted; however, major antecedent landmarks and achievements have been noted, and signposts to the future have been indicated.

If we reflect on the story of this unfolding, we can think about the effects of these achievements on man's understanding of his place in the universe. Let us very quickly review these achievements in terms of three plausible orders of conceptual analysis (Nersessian, 1984, 1998; Andersen & Nersessian, 2000). The purpose of doing so will be to illuminate as much as possible of the implications of the development and emergence of the holonic enterprise.

6.0 Orders of Conceptual Analysis

We could argue that 1st order conceptual analysis took place with the establishment of the earliest foundations of basic mathematical concepts, basic logic, and basic science. In other words, 1st order conceptual analysis took place with the initial move into the realm of basic concepts. We can thank our Greek predecessors for their pioneering work here for recognizing that fundamental knowledge of our world in terms of logic and concepts was humanly possible, and that the journey of exploring and understanding had only just begun. We know there are countless intervening historical variables that determined the course of many events between this time and the beginnings of the enlightenment, but we immediately recognize that with the Copernican revolution, humankind moved from the unique position of being the chosen ones at the center of universe to simply being a very small part in a universe that was much larger and more mysterious than previously imagined: this was no small move in the realm of human knowledge and self-definition. Conceptual pressures were therefore increasingly placed on 1st order analysis. Then, fundamental laws of motion were discovered, systematically articulated and rigorously explicated so that it was possible to consider what was thought to be a true clockwork universe; accurate prediction and calculation about such things as the movement of all bodies became possible; humankind's new position in this revised vision of the universe was again no small expansion of human knowledge and redefinition of self. Rigorous scientific analysis became possible.

Other historical variables continued to steer the overall course of events, not the least of which were powered by such things as the politics and technologies of war (Dyer, 2001), but in the process, as technologies advanced and increasingly deep and broad knowledge of our world was accumulated, axiomatic reduction of logic became seen as distinctly possible and perhaps even inevitable; this was generalized both in intent and through systematization in an attempt to account for language and even more within the burgeoning field of mathematics. From this, the grand notion took shape that an overarching logical calculus might be achieved, and with it, the promise of perhaps explaining the logic of all systems; thence, logical positivism and the Vienna Circle. But this was a short-lived romance with certainty, and in short order, space and time were proven relative, and the goals of the logicism project were demonstrated impossible. All the while, astounding advances in science and technology continued, and orders of magnitude of perception and capacities for calculation were greatly expanded; in this process, we began to reveal fundamental and increasingly detailed keys to understanding chemistry, physics, and biology. We began to see and were increasingly challenged to understand and make sense of extents and natures of our universe that had not only previously been unknown, but even unimaginable.

Again we pause to catch our breath. At this point, we can recognize that what we can call "2nd order conceptual analysis" had emerged through new analytic practices (Nersessian, 1984, 1989, 1998, 2002); Thagard & Shelley, 1997). Such analysis was necessary to systematically and rigorously deal with the new explicated systems of logic emerging in the fields of mathematics and science which could be explored, contemplated and calculated. We must recall that along with this massive 2nd order conceptual shift based on sound empirical evidence, evolution was theorized as being eminently plausible; and, supported through advances in mass communication technologies, these concepts began to diffuse through and be incorporated into humanity's worldview. New conceptual models emerged in all fields (e.g., chemistry; see, Del Re [1998, 2000]) to account for experimental evidence, and new theories were developed to provide foundations for increasing experimentation which, in turn, provided evidence that looped back to build, confirm or modify flexible theories which in turn continued to feed back to create or modify networks of conceptual models. Short leaps forward with great import demonstrated the indeterminacy inherent in many systems, the outcomes of which lead not only to such things as quantum theory (Gribben, 1984; Moore, 1989), but also generated theories of chaos where small and apparently insignificant things can have unpredictable large, system-wide consequences (Hall, 1991); thereafter, self-organization was shown to be the determining feature of all complex adaptive systems, thus complementing and greatly enhancing plausible answers to what Darwin had first posited (Kauffman, 1996, 2000).

From this we have the present unfolding of what we can think of as 3rd order conceptual analysis, where indeterminate emergence is the logical consequence of all self-organizing systems; today, developmental work proceeds to mathematize the design conditions and variables within the state spaces where such emergence will be engineered to take place. This is, of course, the foundation of the holonic enterprise.

The nested consequences of 1st, 2nd and 3rd order paradigmatic shifts in the scientific, technical, and societal realms have forever changed humankind's world views and required (and continue to require) increasingly sophisticated conceptual analytic skills to pose and deal with difficult problems. Regardless of what unfolds or which scenarios continue to be spun (Schwartz, 1991; Schwartz & Leyden, 1997), it would seem that this trend shows no sign of diminishing. The development of the proposed holonic enterprise appears to be a significant specific example of that trend.

7.0 Implications and Future Directions

We are now at the stage of the holonic enterprise—soon to be developed, implemented and diffused. We would not be doing this if the benefits were not clear, but it also seems we are at the point of creating networked computational tools having unprecedented levels of complexity, connectivity, power, and autonomy. Given what we have so far learned about the mathematics of emergence in complex systems, and given the robust virtual space within which we can place purpose-built autocatalytic communities of holonic agents, we will indeed create an emerging, multi-agent, self-organizing, self-directed, cyberspace. Latest applications in this field have shown high suitability for the holonic enterprise evolutionary strategy in the design of Internet-enabled soft computing holarchies for telemedicine, for example (Lacher & Nguyen, 1995). Many other system enterprise applications appear to be eminently possible; it would seem this is logically inevitable (Riegler, 2002).

While contemplating such diverse holonic enterprise applications, we must not forget the parallel development of so many other computer-mediated communications technologies in almost every field of human endeavour. The full emergence of the holonic enterprise across firms, organizations, governments, and cultures—and eventually networked and expressed through all technologies across the planet—will create a very deep and significant multifaceted paradigm shift, perhaps more significant than any we have previously seen. There will be many interrelated consequences—scientific, technical, political, and conceptual.

In this paper, based on an overview of the proposed holonic web-based system and a brief overview of major technical, scientific, political and conceptual changes and consequent paradigmatic shifts, we have created a foundation from which to address such plausible consequences by suggesting an introductory framework that describes the development of 1st, 2nd, and 3rd order conceptual analysis. This introductory framework has allowed us to view the unfolding of conceptual skills necessary to discover, explore, understand, and anticipate the workings of highly complex systems that, as we continue to reveal and understand them more fully over time, provide evidence for what we today consider to be emergence.

Against the backdrop of the ongoing development of the holonic enterprise, three deeply interrelated implications having to do with the future of conceptual analysis and its relation to our ongoing work can be drawn from our examination of this introductory framework.

The first is methodological purely from the perspective of what is necessary in terms of rigorous conceptual work (Brown, 1998). This has to do with the need to continue ongoing exploration, development and refinement of plausible conceptual models to account for and further explore stubborn problems that surround the concept of emergence, for example (Emmeche, Køppe, & Stjernfelt, 1997), as well as to examine prospects for and plausible outcomes of what we now think of as self-organization and the emergence of novelty in virtual complex systems (Kauffman & Sabelli, 1998; Castellano, Fanelli, & Mencar, 2002; Conte, 2001; Quartz & Sejnowski, 1997). Following the thinking of Neuss and Kent (1994) as recently underscored by Turkle (cited by Coutu, 2003), we suggest that such essential, challenging, conceptual work can be thought of as a move to a 4th order of conceptual analysis (O'Hara, 1995); this appears to be the direction in which we are headed.

The second implication has more to do with computational and system advances now being developed to mathematize and launch the proposed holonic enterprise described herein, as well as those that will likely emerge in the future from this type of work. Clearly, the development of the proposed holonic enterprise provides an excellent opportunity to explore and study in detail the logical, mathematical, and technical aspects of complex, evolving, virtual, multi-agent systems, both in terms of simulation and real-world articulation. Such exploration and study, in turn, will permit ongoing reformulation of the conceptual models employed to further explicate emergence and explore novelty in such virtual complex systems. Indeed, Andersen (2000, 2001) echoes and explores these very points in his examination of the necessity for what he terms "elastic systems," computer-mediated control interfaces that augment and support human decision-making in very demanding, complex, dynamic task environments such as maritime navigation.

The third implication stands firmly on the first two, which in their primary forms emphasize the conceptual and the technical, respectively. If we are dealing with the enhancement of our conceptual models to explore emergence and novelty in virtual complex systems on the one hand, and at the same time dealing with advances in the technical and mathematical aspects of the proposed holonic enterprise and other computational spaces on the other, we are by default dealing with a larger theoretical system related to self-organizing complex systems in general that encompasses both theoretical models and varieties of their practical articulation. In exploring this terrain, we are dealing with theory emergence. That is, consideration of the conceptual work in the mathematical and technical realms encompassing the proposed holonic enterprise, when coupled with the parallel conceptual work having to do with the exploration of emergence, novelty, and enhanced problem-solving through virtual systems, suggests we are exploring significant elements of theory emergence. This conclusion has direct isomorphism with the thinking that underpins Schwarz's (2002a) holistic metamodel advanced to account for the shifting from the useful but inadequate dualist rationalist and reductionist paradigm (reflected in what has been described here as the unfolding of 1st, 2nd, and 3rd order conceptual analysis) to a holistic paradigm, echoed in part by the work of Bostrum (2003), capable of accounting for spontaneous self-organization, complexity and autonomy (suggested in this paper as a necessary and inevitable 4th order of conceptual analysis).

The proposed holonic enterprise as explored in this paper may therefore have important implications for our understandings of many elements of complexity science (Cooksey, 2001), how we develop and employ models and analogies in our conceptual and technical work (Nersessian, 2002; Andersen, 2000), and how we explore, develop and apply essential conceptual skills focused not only in scientific and technical realms, but in the essential overarching realm of theory emergence (Schwarz, 2002b).

References

Andersen, H., & Nersessian, N. J. (2000). Nomic concepts, frames, and conceptual change. *Philosophy of Science, 67,* S224-S241. (Proceedings) Retrieved Nov. 2, 2003 from http://www.cc.gatech.edu/aimosaic/faculty/nersessian/papers/nomic-concepts-frames-and-conceptual-change.pdf

Andersen, P. B. (2000). *Elastic systems*. Retrieved Dec. 21, 2003 from http://www.cs.auc.dk/~pba/Homepagematerial/publicationfolder/ ElasticSystems.pdf

Andersen, P. B. (2001). Pervasive computing and space. In R. Stamper, K. Liu & E.-S. About-Zeid (Eds.), *Organizational semiotics*: *Evolving a science of information systems (pp* 106-125). Montreal, Quebec: Concordia University. (Proceedings of IFIP WG8.1 Working Conference) Retrieved Dec. 21, 2003 from http://www.cs.auc.dk/~pba/Homepagematerial/ publicationfolder/PervasiveComp.pdf

Axtell, R. L. (2001*). Economics as distributed computation.* Unpublished paper. The Brookings Institution, Center on Social and Economic Dynamics.

Babaoglu, O., Meling, H., & Montresor, A. (2002).*Anthill: A framework for the development of agent-based peer-to-peer systems (*Tech. Rep. No. UBLCS-2002-09). Bologna, Italy: University of Bologna. Retrieved Dec. 2, 2003 from http://www.cs.unibo.it/projects/anthill/papers/2001-09.pdf

Barabási, A.-L. (2002). *Linked: The new science of networks*. Cambridge, MA: Perseus Publishing.

Berlinski, D. (2000). *Newton's gift: How Sir Isaac Newton unlocked the system of the world.* New York: Touchstone.

Bluetooth SIG. (2003). *Bluetooth*: *The official Bluetooth website*. Retrieved Nov. 2, 2003 from http://www.bluetooth.com

Bostrom, N. (2003).*Transhumanist values*. Discussion paper. Oxford University, Faculty of Philosophy. Retrieved Nov. 3, 2003 from http://www.nickbostrom.com/ethics/values.html

Broda, H. (2003).The Six Faces of the Web. *Sun Journal, 5*(2) Retrieved Dec. 19, 2003 from http://www.sun.com/executives/sunjournal/v5n2/feature1.html?redirect=false&refurl=http://wwws. sun.com /software/jxta/

Brown, H. I. (1998). *Conceptual comparison and conceptual innovation.* Paper presented at International Congress on Discovery and Creativity, University of Ghent, Belgium, May 14-16, 1998. Retrieved Oct. 29, 2003 from http://csmaclab-www.cs.uchicago.edu/philosophyProject/sellars/brown/ccci/html

Burke, J. (1999).*The knowledge web: From electronic agents to Stonehenge and back -- and other journeys through knowledge.* New York: Simon and Schuster.

Cadenasso, M. L., Pickett, S. T. A., Weathers, K. C., Bell, S. S., Benning, T. L., Carreiro, M. M., & Dawson, T. E. (2003). An interdisciplinary and synthetic approach to ecological boundaries. *Bioscience, 53*(8), 717.

Castellano, G., Fanelli, A. M., & Mencar, C. (2002). A neuro-fuzzy network to generate human-understandable knowledge from data. *Cognitive Systems Research, 3*(2), 125-144.

Casti, J. L. (1994).*Complexification*. New York: Harper-Collins.

Casti, J., & DePauli, W. (2000).*Gödel. A life of logic*. Cambridge, MA: Perseus Publishing.

Christensen, W. D. & Hooker, C. A. (in press a). Self-directed agents. *Canadian Journal of Philosophy.*

Christensen, W. D. & Hooker, C. A. (in press b). Representation and the meaning of life. In H. Clapin, P. Slezak & P. Staines (Eds.), *Representation in mind: New approaches to mental representation.* Westport, CT: Praeger.

Cilliers, P. (2002).Why we cannot know complex things completely. *Emergence, 4*(1/2),77-84.

Conant, J., & Haugeland, J. (Eds.) (2000).*The road since structure: Thomas S. Kuhn: Philosophical Essays 1970-1993*. Chicago: The University of Chicago Press.

Conte, R. (2001). Emergent (info) institutions. *Journal of Cognitive Systems Research, 2*, 97-110. Retrieved Dec. 11, 2003 from http://www.elsevier.nl/gej-ng/10/15/16/58/31/25/abstract.html (for access see: Science Direct)

Cooksey, R. W. (2001). What is complexity science? A contextually-grounded tapestry of systemic dynamism, paradigm diversity, theoretical eclecticism, and organizational learning. *Emergence, 3*(1),77-103.

Coutu, D. L. (2003) *Technology and Human Vulnerability. A conversation with MIT's Sherry Turkle. Harvard Business Review, 81*(9), 43-50

Davis, S., & Meyer, C. (1998). *Blur: The speed of change in the connected economy.* New York: Warner Books.

Del Re, G. (1998). Ontological Status of Molecular Structure. *HYLE – International Journal for Philosophy of Chemistry, 4*(2), 81-103.

Del Re, G. (2000). Models and analogies in science. *HYLE – International Journal for Philosophy of Chemistry, 6*(1), 5-15.

Downey, L., & Este, R. A. (1984). *Curriculum outline, Education 540 (Educational Policy).* Unpublished document., Vancouver, BC: University of British Columbia, Faculty of Education, Division of Educational Administration.

Dummett, M. (1991). *Frege and other philosophers*. Oxford: Clarendon Press

Dyer, G. (2001).Commentary at private consultative meeting of Foresight Canada, May 4, 2001. Kananaskis Village, Alberta

Emmeche, C., Køppe, S. & Stjernfelt, F. (1997). Explaining emergence: Towards an ontology of levels. *Journal for General Philosophy of Science, 28*, 83-119. (also at http://www.nbi.di/~emmeche/coPubl/97e.EKS/emerg.html)

Este, R. A. (2003). *Conceptual analysis as a useful methodological approach for the study of innovation: A discussion*. Unpublished paper. Calgary: University of Calgary, Graduate Division of Educational Research, and the Faculty of Continuing Education.

Falconer, J. (2002). Emergence happens! Misguided paradigms regarding organizational change and the role of complexity and patterns in the change landscape. *Emergence, 4*(1-2), 117-30.

Gervasi, V., & Prencipe, G. (2003). Coordination without communication: The case of the flocking problem. Unpublished paper. Pisa: Università di Pisa, Dipartimento di Informatica. Retrieved Nov. 11, 2003 from http://circe.di.unipi.it/~gervasi/Papers/dam02.pdf

Gleick, J. (1999). *Faster. The acceleration of just about everything*. New York: Pantheon

Goldstein, J. (1999). Emergence as a construct: History and issues. *Emergence, 1*(1), 49-72.

Goldstein, J. (2000). Emergence: a construct amid a thicket of conceptual snares. *Emergence, 2*(1), 5-22.

Gribben, J. (1984). *In search of Schrödinger's cat: Quantum physics and reality.* Toronto: Bantam Books.

Hall, N. (1991). *Exploring chaos: A guide to the new science of disorder.* New York: W. W. Norton

Hillis, D. (1998). *The pattern on the stone*. New York: Basic Books

Holland, J. (1992). Echoing emergence: Objectives, rough definitions, and speculations for echo-class models (# 93-04-023). Santa Fe, NM: Santa Fe Institute Working Papers, Preprints, and Reprints Series.

Holland, J. (1995). *Hidden order. How adaptation builds complexity.* Reading, MA: Helix / Addison-Wesley.

Kauffman, L., & Sabelli, H. (1998). Bios: Creative organization beyond chaos. Paper presented at the International Society for Systems Sciences, Atlanta, July 1998. Retrieved Dec. 9, 2003 from http://www2.math.uic.edu/~kauffman/BIOS.pdf

Kauffman, S. (1993). *The origins of order.* New York: Oxford University Press

Kauffman, S. (1995). *At home in the universe*. New York: Oxford University Press.

Kauffman, S. (1996). Investigations: The Nature of Autonomous Agents and the Worlds they Mutually Create. Santa Fe, NM: Santa Fe Institute Preprint. Retrieved Sept. 20, 2003 from http://www.santafe.edu/sfi/People/kauffman/Investigations.html

Kauffman, S. (2000). *Investigations*. New York: Oxford University Press.

Kauffman, S., & Este, R. A. (2004)."Conceptual analytic thinking about innovation: Some economic implications," Unpublished paper.

Kennedy, J., & Eberhart, R. C. (2001). *Swarm intelligence*. San Francisco: Morgan Kaufmann.

Koestler, A. (1968). *The ghost in the machin*e. Basingstoke, Hampshire (UK): Macmillan.

Kuhn, T. S. (1957). *The Copernican revolution.* Cambridge, MA: Harvard University Press

Kuhn, T. S. (1996). *The structure of scientific revolution*s (3rd ed.). Chicago: The University of Chicago Press.

Kurzweil, R. (1990).The Age of Intelligent Machines. Cambridge, Mass.: The MIT Press

Kurzweil, R. (2001). *The law of accelerating returns*. Retrieved Apr. 17, 2003 from http://www.kurzweilai.net/articles/art0134.html?printable=1

Kurzweil, R. (2003). Promise and Peril in the 21st Century. *CIO Magazine*, (Fall/Winter) (Special Issue: Technology's impact on everything). Retrieved Dec. 2, 2003 from http://www.cio.com/archive/092203/kurzweil.html

Lacher, R. C., & Nguyen, K. D. (1995). Hierarchical architectures for reasoning. In R. Sun, & L. A. Bookman (Eds.), *Computational architectures integrating neural and symbolic processes: A perspective on the state of the ar*t. Boston: Kluwer Academic Publishers

Lee, J. A. N. (1994).*Charles Babbage*. Retrieved Mar. 21, 2003 from http://ei.cs.vt.edu/~history/Babbage.html

Lenzer, G. (Ed.).(1975). *Auguste Comte and Positivism: The Essential Writings*. New York: Harper.

Lewin, R. (1992). *Complexity. Life at the edge of chaos*. New York: Macmillan.

Lindblom, C., & Woodhouse, E. (1993). *The policy-making process*. New York: Prentice-Hall.

Lissack, M. R. (1999). Complexity: The science, its vocabulary and its relation to organizations. Emergence, 1(1), 110-126.

Microsystems Technology Office. (2003).*DARPA/Microsystems Technology Office Programs*. Retrieved Nov. 2, 2003 from http://www.darpa.mil/MTO/RADPrograms.html

Montresor, A., Meling, H. & Babaoglu, O. (2002). Towards self-organizing, self-repairing and resilient peer-to-peer systems (Tech. Rep. UBLCS-2002-09). Bologna: University of Bologna, Department of Computer Science, Retrieved Oct. 19, 2003 from http://www.cs.unibo.it/projects/anthill/papers/2002-09.pdf

Moore, W. (1989). *Schrödinger. Life and thought*. Cambridge: University of Cambridge Press.

Negroponte, N. (1995). *Being digital*. New York: Alfred A. Knopf.

Nersessian, N. J. (1984). *Faraday to Einstein: Constructing meaning in scientific theories*. Dordrecht: Kluwer Academic Publishers.

Nersessian, N. J. (1989). Conceptual change in science and in science education. *Synthese, 80*, 163-184.

Nersessian, N. J. (1998). Conceptual change. In W. Bechtel, & G. Graham (Eds.), *A companion to cognitive science* (pp155-166). Malden, MA: Blackwell. (also available at http://www.cc.gatech.edu/aimosaic/faculty/nersessian/papers/conceptual-change.pdf)

Nersessian, N. J. (2002). Kuhn, conceptual change, and cognitive science. In T. Nickles (Ed.), *Thomas Kuhn, (Contemporary Philosophers In Focus Series)*. Cambridge: Cambridge University Press. Electronic version Retrieved Nov. 21, 2003 from http://www.cc.gatech.edu/aimosaic/faculty/nersessian/papers/kuhn-conceptual-change-science.pdf

Neuss, C., & Kent, R. E. (1994). Conceptual Analysis of Resource Meta-information. Discussion paper. Fraunhofer Institute for Computer Graphics. Retrieved Aug. 30, 2003 from http://www.igd.fhg.de/archive/1995_www95/papers/94/www3.html

Nicolis, G., & Prigogine, I. (1989). *Exploring complexity*. New York: W. H. Freeman.

Oakridge National Laboratory/Southwest Research Institute. (2000). I*ntelligent Wireless Sensors and Systems for Insect Tracking*. Retrieved Nov. 3, 2003 from http://tent.appliedphysics.swri.edu/insecttrack/smith_abs_ornl.pdf

O'Hara, M. (1995). *New minds for new times*. [VHS video of conference presentation, ~50 minutes, colour]. Calgary: The Alliance for Capitalizing on Change

Pirolli, P., & Card, S. (1995). Information foraging in information access environments. Proceedings, 1995 Conference on Human Factors in Computing Systems (Chicago, May 1995). Retrieved Nov. 25, 2003 from http://www.acm.org/sigchi/chi95/proceedings/papers/ppp_bdy.htm

Quartz, S. R., & Sejnowski, T. J. (1997). The neural basis of cognitive development: A constructivist manifesto. *Behavioural and Brain Sciences, 20*(4), 537-596.

Riegler, A. (2002). When is a cognitive system embodied? *Cognitive Systems Research, 3*, 339-348. (Special issue on Situated and Embodied Cognition). Retrieved Dec. 18, 2003 from http://www.univie.ac.at/constructivism/people/riegler/papers/riegler02embodiment.pdf

Russell, B. (1911). *The philosophical importance of mathematical logic*. Retrieved Apr. 20, 2003 from
http://www.marxists/org/reference/subject/philosophy/works/russell.htm.
Currently available at http://www.marxists.org/reference/subject/philosophy/index.htm

Sabatier, P. (1999). *Theories of the policy process*. New York: Westview Press.

Schwartz, P. (1991). *The art of the long view*. New York: Doubleday.

Schwartz, P., & Leyden, P. (1997). The long boom: A history of the future, 1980-2020. *WIRED Magazine*, Issue 5.07 (July). Retrieved May 27, 2003 from http://www.wired.com/wired/archive/5.07/longboom_pr.html

Schwarz, E. (1997). Lesson 1 – General System Theory. Selection 2. Excerpted from: Toward a holistic cybernetics. from science through epistemology to being, *Cybernetics and Human Knowing, 4*(1), 19-23. Retrieved Dec. 21, 2003 from http://www.unikk.ch/course/PDF/lesson1-text-2.PDF

Schwarz, E. (2002a). A systems holistic interpretation of the present state of contemporary society and its possible futures. paper presented at the fifth European Systems Science Congress, Heraklion, Crete, Oct. 16-19 2002 Retrieved Dec. 21, 2003 from http://www.afscet.asso.fr/resSystemica/Crete02/SchwarzB.pdf

Schwarz, E. (2002b). can real life complex systems be interpreted with the usual dualist physicalist epistemology – Or is a holistic approach necessary? Paper presented at the fifth European Systems Science Congress, Heraklion, Crete, Oct. 16-19 2002. Retrieved Dec. 21, 2003 from http://www.afscet.asso.fr/resSystemica/Crete02/SchwarzA.pdf

Sun Microsystems. (2003).*Project JXTA*. Retrieved Nov. 2, 2003 from http://www.jxta.org

Tasaka, H. (1999). Twenty-first-century management and the complexity paradigm. *Emergence, 1* (4), 115-123.

Thagard, P., & Shelley, C. (1997). *Abductive reasoning: Logic, visual thinking, and coherence*. Retrieved May 5, 2003 from http://cogsci.uwaterloo.ca/Articles/Pages/%7Fabductive.html

Theraulaz, G., & Deneubourg, J. –L. (1992). Swarm intelligence in social insects and the emergence of cultural swarm patterns (# 92-09-046). Santa Fe, NM: Santa Fe Institute Working Papers, Preprints, and Reprints Series.

Ulieru, M. (2002). Emergence of holonic enterprises from multi-agent systems: A fuzzy evolutionary approach. In Frontiers in AI and Applications Series and V. Loia (Ed.), *Soft computing agents* (pp. 187-215). Amsterdam: IOS Press.

Ulieru, M., Stefanou, D., & Norrie, D. (2000). Holonic metamorphic architectures for manufacturing: Identifying holonic structures in multi-agent systems by fuzzy modeling. In J. Wang & A. Kussiak (Eds.), Handbook of computational intelligence in design and manufacturing (pp. 3-1—3-36). Boca Raton, FL: CRC Press.

Van Heijenoort, J. (1967). *From Frege to Gödel. A source book in mathematical logic, 1879 – 1931*. Cambridge, MA: Harvard University Press

Waldrop, M. M. (1992). *Complexity.* New York: Simon and Schuster.