Scale Invariant Pareto Optimality

A Meta–Formalism For Characterizing and Modeling Cooperativity in Evolutionary Systems

Mark Fleischer
Johns Hopkins University
Applied Physics Laboratory
Laurel, Maryland
Mark.Fleischer@jhuapl.edu

ABSTRACT

This article describes a mathematical framework for characterizing cooperativity in complex systems subject to evolutionary pressures. This framework uses three foundational components that constitute a *meta-formalism* that can be utilized in a host of research and development settings to improve the management, control, and understanding of large numbers of interacting systems such as in communication, computer, and sensor networks. A new concept, *Scale Invariant Pareto Optimality*, provides a mathematical basis for the *efficient tradeoffs of efficiency* on many scales and the measurement of *cooperativity* in complex systems. A mathematically oriented definition of *self-organized behavior* is also described. Discussion and conjectures are offered.

Categories and Subject Descriptors

I.2.6 [**Artificial Intelligence**]: Learning—*Parameter learning, Connectionism and neural nets* and; I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Coherence and coordination*

General Terms

Algorithms, Design, Economics, Experimentation, Management, Measurement, Performance, Theory.

Keywords

Complex Systems, Self-Organization, Self-Organized Criticality, Multi-objective Optimization, Pareto Optima, Swarm Intelligence

1. INTRODUCTION

The ubiquity of computers and other forms of advanced technology have created a number of difficult problems. Many are of a practical nature relating to the management of large ensembles of interacting systems. To be effective, these systems must work together harmoniously, *i.e., cooperatively*. Unfortunately, their harmonious operation become increasingly difficult to achieve and maintain as their size and numbers increase. Indeed, there is a growing consensus among experts that current approaches for managing these large systems will be insufficient to handle their increased complexity. For example, today's communications networks have become enormously complex systems—for the past three decades, this growth has approximately doubled every 18 months according to Moore's Law [14, p.32]. These new technologies, configurations, protocols, network and computer architectures and so forth are constantly and relentlessly challenging our abilities to effectively manage them.

In addition to these practical issues and difficulties are problems of a more theoretical nature—problems stemming from an incomplete, or incoherent understanding of *fundamental phenomena*. The complete understanding of *complex systems* still lies well beyond our grasp, yet our growing dependence on them impels us to continue to explore ideas and increase our understanding.

Many different approaches for studying the behavior of complex systems have been described. One approach is based on *Swarm Intelligence* (SI) and represents the view that it is possible to control and manage large, complex systems of interacting entities with only "minimal", or *stigmergic* communications channels, where only a relatively small amount of information is communicated [9].

In recent years, new insights have been obtained based on observations of social insects [6]. Ant colonies and beehives, *e.g.,* seem to conduct their affairs in a very *organized* and purposeful way that enhances their *collective* survival. Needless to say, these insects do not have very large brains and their capability to communicate complex information for planning and resource allocation seem very limited. Yet, their collective behavior has often been characterized using the terms "intelligent", "emergent", and "self-organized" [6]. Their behaviors are also reminiscent of those observed in other domains of inquiry such as cellular automata (CA) [17, 18] which have perplexed scientists for many years.

A major problem confronting scientists working in these areas is that no widely agreed upon definition of SI or *self-organized behavior* (SOB) exists. How could or should these terms be *mathematically* defined or characterized? This difficulty is often reflected by the many descriptions of SI and SOB that are couched in terms of "self-organization". Such circular definitions obscure what is really going on and how complex systems can be more simply characterized. Moreover, the absence of precise definitions and therefore, theoretical foundations, itself creates problems caused by many disparate concepts. Yet, progress is continually being made (see *e.g.,* [5, 4, 12]). Nevertheless, it seems useful to provide some new ideas in the hopes of contributing to a greater synthesis.

This article provides some new perspectives for describing fun-

GECCO'05, June 25–29, 2005, Washington, DC, USA.
Copyright 2005 ACM 1-59593-010-8/05/0006 ...$5.00.

damental properties by employing concepts of *efficiency*, and *adaptivity*, in a simple, general, and hopefully compelling way that provides a foundation for further experimentation and study. A set of three *foundational frameworks*, described in [9], are used to develop a mathematically oriented *meta–formalism* that provides guidelines for conducting experiments. Actual experiments are the subject of future research.

One feature of this approach is that it gives us tools to mathematically characterize and model the concept of *cooperativity* using mathematically based descriptions of adaptivity and efficiency. These in turn are described using new results from multi-objective optimization problems (MOPs) and theory. These results show how a *Pareto optimal frontier*, *i.e.,* the set of efficient solutions in an MOP, or indeed any set of non-dominated points in objective function space, can be quantified by a single scalar value (see [10, 24]). This provides an alternative, and we believe more general, method for ranking the desirability of a set of solutions (but see [23] for a discussion on the limitations of unary measures). The measure of this set is a *Lebesgue measure* and entails an *entire regime of efficient solutions* involving various tradeoffs, hence involving adaptation. In short, this Lebesgue measure quantifies efficiency (in some sense) in a more general way that also encompasses adaptivity.

This article further develops this notion of efficiency and adaptivity to mathematically characterize a novel use of this Lebesgue measure: measuring the *efficient tradeoffs of efficiency* or, to put it another way, obtaining a Lebesgue measure of an MOP *where each objective function is itself a Lebesgue measure that measures efficiency in some underlying MOP on other scales*—the *individual and community* scales. These scales define a *high-level* tradeoff curve corresponding to *self-interest* and *community interest* that can also be quantified using a Lebesgue measure. This high-level Lebesgue measure thus quantifies entire regimes of operating characteristics on both the individual (sub-system) and community (system) levels, hence provides a scalar quantity that, in effect, **measures cooperativity** in a homogeneous system of entities. Finite-state machine (FSM) models in a CA are used as a *general model substrate* along with a genetic algorithm paradigm where the fitness function is based on this high-level Lebesgue measure. This approach allows us to mathematically and computationally characterize the *evolution of cooperation*.

This article is organized as follows: Section 2 provides more detail on the current research environment regarding SI and SOB. Section 3 describes the *meta–formalism* that provides a basis for further research of evolving cooperative and self-organized systems of systems. Section 4 describes how the meta–formalism can be used to model the components of these systems and articulates some system requirements in an abstract and general way. Section 5 describes possible implementation schemes for a CA composed of FSMs. Section 6 presents some discussion and conjectures about how this meta–formalism could address issues involving concepts such as *self-organized criticality* (SOC) and related phenomena. Finally, Section 7 provides concluding remarks.

2. BACKGROUND

2.1 Swarm Intelligence

Observations of social insects such as ants and ant colonies provide a great deal of insight into their behavior and SI in general. These insect colonies have several ways of solving different but related problems. The main mechanism for solving them is through the use of chemical substances known as *pheromones* which have a scent that decays over time through the process of evaporation [6, p. 26]. These pheromones form the basis of what amounts to a clever, and apparently simple, communications and information storage and retrieval system. Because pheromone intensity decays over time, it also provides a very simple information processing mechanism that can implement forms of positive and negative feedback [6, pp. 9-10, 41] and *reinforcement learning* mechanisms [6, p.96].

Many optimization algorithms and heuristics have imaginatively captured aspects of SI. Indeed, many difficult optimization problems have been solved by *ant algorithms* (see [6] for a large number of examples and citations). These algorithms generally use some analogue of pheromone or *stigmergic* signalling mechanism or reinforcement learning mechanism to increase the probabilities of using certain routes in a routing algorithm. In these attempts to implement SI, however, researchers are often forced to creatively sidestep what SOB is.

2.2 What Exactly is "Self-Organization"?

Perhaps the most vexing issue in complex systems is: How do we identify SOB when we see it? Certainly the term seems useful and descriptive in association with social insects that have brains. They do, after all, organize themselves. But the term is also used in many contexts and in association with other concepts that ultimately make its meaning less clear. Indeed, entire books on the subject have been written suggesting this is not a trivial question (see *e.g.,* [12]). Moreover, many phenomena can be described as forms of SOB. Solow [21] describes a mathematical model of function specialization in economic systems that can be considered a form of SOB. SOB is often associated with SI, *stigmergy*, *reinforcement learning*, and in many other descriptions of complex systems behavior [9, 8, 15]. Bonabeau *et al.* [6, p.9-11] highlights four important elements all relating to how large numbers of simple entities interact. In his view, SI involves: 1) forms of positive feedback; 2) forms of negative feedback; 3) the *amplification of fluctuations* that give rise to structures; and 4) multiple interactions of multiple entities. Although this is a significant step towards characterizing *some* mathematical attributes of SOB, it is not sufficient to fully define it. These attributes do not fully explain how the elements of systems exhibiting SOB actually lead to SOB. For instance, why or how did pheromones evolve the way they did? Why do they have certain evaporative properties? How do these properties lead to SOB? See [9]. Wiener [22, p.156-7] noted the significance of this issue:

> How then does the beehive act in unison, and at that in a very variable, adapted, organized unison? Obviously, the secret is in the intercommunication of its members . . . [and] can vary greatly in complexity . . .

As noted earlier, the difficulty of pinning down the concept of SOB seems to have forced even the most logical and insightful researchers to resort to forms of circular reasoning and definitions. Serra *et al.* describes the concept of self-organization as "highly organized behaviour even in the absence of a pre-ordained design." [20, p.1] (what is 'organized'?) and as "unexpected and complex behaviours [ ]." [20, p.2] (what is 'unexpected and complex'?). Even the likes of Ilya Prigogine[1] and Norbert Wiener have been subject to this conundrum. Prigogine *et al.* [18, p.181-186] cites the spontaneous emergence of structures in chemical reactions as examples of self-organization. Weiner [22, Ch.16] used it to describe brain waves. Thus, despite a lot of study and descriptions in many domains of inquiry (*e.g.,* [12]), *self-organization* remains, at best, a rather nebulous concept leaving us with many unanswered

[1]Prigogine won the Nobel Prize in 1977 for his work on the thermodynamics of non-equilibrium systems.

questions: Is SOB that which leads to "structures" or patterns? If so, what distinguishes them from random effects?

Finding answers to these questions and problems is quite daunting and perhaps ultimately elusive. Yet, by attempting to highlight some first principles, we might at least gain additional insight—our ultimate goal. In considering first principles in the context of SOB and SI, therefore, it must be borne in mind that the behavior of systems affected by biological systems (*e.g.,* systems designed using human intelligence) are all heavily influenced by the forces of evolution. The successful performance of such systems must therefore incorporate modes of cooperativity and it is this concept and the related concepts of efficiency and adaptivity that requires mathematical gloss. Therefore, let us attempt to impose some structure on this amalgam of observations, questions, and ideas in the hopes that it leads to further insights and ideas.

3. THE META-FORMALISM TRIAD

Developing this mathematical structure requires a set of basic principles. Fleischer [9] describes three basic principles by drawing analogies to the theoretical development of the simulated annealing (SA) algorithm and puts them in the context of SI. This section provides a brief description of this triad.

Briefly, the theoretical development of SA required application of a set of *first principles* based on the laws of nature and its relevant implications. These first principles were based on the theory of thermodynamics. The relevant implications of this theory is that systems tend toward a state of maximum entropy. Maximizing the entropy then leads to the definition of the Boltzmann distribution, a central component in SA [2].

The second component of this triad is a *dynamical framework*. In SA, this was based on the theory of Markov chains and related to the Boltzmann distribution. This led to the famous Metropolis Acceptance Criteria. Finally, the third component, the *problem framework*, was based on combinatorial optimization problems (COPs) and permitted the application of the first two components onto test-bed problems for experimentation and analysis. These three components comprise the entire edifice on which SA theory and practice is based (see *e.g.,* [2]).

This foundational triad can be appropriately recast within the context of SI, SOB, and complex systems and imposes a useful and helpful structure for conducting research on evolving and cooperative systems. Each component of this triad, 1) the relevant laws of nature and their significant implications; 2) a dynamical framework, and 3) an associated problem framework are described in greater detail below.

3.1 The Relevant Laws of Nature

The most important part of this foundational triad is the set of *first principles*. For SI, these first principles are based on the laws of evolution and natural selection [9]. Certainly, the laws of evolution and natural selection apply to insect swarms. Because we want to take a more general approach, we do not limit this component to laws of nature *per se*. Stating it in this way however provides a useful way to characterize how systems evolve over time and conveys the notion that the synthesis developed here is based on some overriding or governing principle and helps to frame the issue of what its *relevant and significant implications* are. In effect, stating that the laws of evolution apply implicitly says that over time, systems become more *efficient, adaptable* and *useful* either due to human design and experience, or through the sometimes unkind pressures of natural selection. Indeed, for complex systems, this efficient behavior may exist on *many scales*.

This notion of efficiency seems reasonable on its face, but is also based on observations of social insects and how they *adapt* to changing environmental conditions (see [11]). Certainly, the efficient allocation and use of resources provides a distinct survival value to any species. It also seems reasonable that effective methods for *determining* efficient modes of behavior can be quite valuable for enhancing the survival value of a species (or managing a complex system) because it enables *adaptive* behavior. Efficient operations are then possible when changes occur in the environment.

In general, systems have several operating modalities and goals and the notion of efficiency must therefore often be related to *multiple measures of performance* or goals. These multiple performance measures thus require some tradeoffs based on utility functions. One could, of course, refer to many theories of economics which explicitly entail many measures of performance in an environment with many interacting entities. In such complex systems, tradeoffs must often be made on fairly short time scales as "individual" entities attempt to navigate through a changing environment imposed, in part, by other similar (or possibly different) entities (or systems) each of which may compete for resources. Thus, to be truly "efficient", an entity must make *efficient* tradeoffs in goals or objectives and do so consistently if it is to survive the rigors of natural selection.

So what mathematical concepts can we use here to characterize *efficiency* and *adaptability*? In SA theory, we have the Boltzmann distribution. What is its counterpart in SI, SOB, or complex, evolving systems? It turns out that the notions of efficiency and adaptability do have a mathematical formalism based on the concept of *Pareto optimality* best described in the context of MOPs.

Efficiency and Pareto Optimality: Complex systems, whether man-made or natural, are either designed or evolved to perform in certain ways. The measures of performance in these systems can often be mathematically modeled using *objective functions*—functions of decision variables that produce a scalar to be either minimized or maximized. Such optimization problems are ubiquitous and these objective functions essentially provide a mechanism for ranking the desirability of a decision—the higher (or lower) the value of the objective function, the better. The highest (lowest) value of the objective function corresponds to the *best* decision (or set of decisions if there are multiple optima). Such problems are often very difficult to solve.

Complex systems are often made even more difficult to solve because they may involve several objectives that must all be considered in assessing system performance. In such cases, what is *best* is not necessarily the same as what is *optimal*. *Best* refers to a single solution that maximizes utility from among a set of the efficient solutions. *Optimal* refers to the entire set of efficient solutions that together dominate all other feasible solutions. In single objective optimization problems, these two notions always coincide while in MOPs they can be distinct. See [10].

In systems with multiple objectives, these objectives are often in conflict with one another where one objective function value must be "traded off" for another. In MOPs, therefore, *optimal* solutions are characterized by a set of *Pareto optima* or an *efficient frontier*—a set of points in objective function space often referred to as a *tradeoff curve*. In these problems, each set of decision variables (operational parameters) produces *several* objective function values, *i.e.,* a single point in objective function space corresponding to a single *Pareto optimum*. Such a point has the property that when it is compared to *any other feasible point* in objective function space, at least one objective function value is superior to the corresponding objective function value of this other feasible point. Pareto optima therefore constitute a special subset of points that

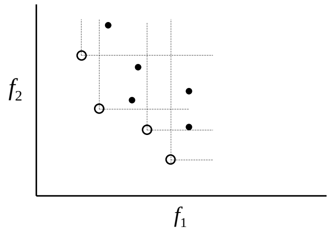

Figure 1: The Pareto Optimal Frontier

collectively dominate all other feasible points in objective function space. The Pareto optima necessarily have the property that improving one objective function value is possible only by incurring some cost in worsening another objective function value—hence the term *trade off* curve. Figure 1 illustrates Pareto optima with open O's in a system with two objective function values f_1 and f_2 that are minimized. Note that these points dominate (are superior to) the other points in this 2-dimensional objective function space. It is among this set of points that operational decisions must be restricted if operational efficiency is to be maintained.

The Measure of Pareto Optima: Recently, Zitzler [24] and Fleischer [10] described a new way of mapping a set of Pareto optima or set of *non-dominated*[2] points to a scalar. For a given set of bounds[3] the size of this space is its *Lebesgue* measure or hypervolume and is illustrated by the shaded region in Figure 2 depicting two minimizing objective functions f_1 and f_2 with the indicated upper bounds. The black dots constitute a non-dominated set of points.

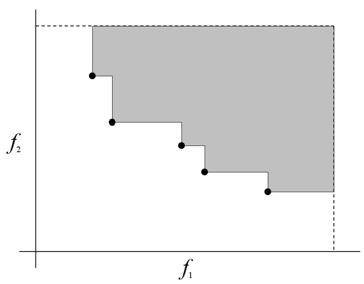

Figure 2: The hypervolume of Pareto optima.

The Lebesgue measure *quantifies efficiency* in a useful way: the Lebesque measure of a set of non-dominated points attains its maximum value if and only if those points are Pareto optima [10]. Consequently, changes in the *overall* efficiency can be reflected in changes in the Lebesgue measure and can be depicted as a receding or expanding Pareto optimal frontier.[4] Thus, the evolution of cooperativity can be effected by using evolutionary pressures to increase some appropriately defined Lebesgue measure.

[2]A set of points is *non-dominated* when a comparison between any two points in the set indicate a tradeoff in objective function values.

[3]The Lebesgue measure of a non-dominated set achieves its maximum value when the non-dominated points are Pareto optima. This measure will be different with different bounds. Nevertheless, once the bounds are set, points that maximize the Lebesgue measure will correspond to Pareto optima. Changing the bounds thus changes *which* Pareto optima maximize the Lebesgue measure.

[4]Note that although two sets of non-dominated points may produce two different Lebesgue measures (we assume both measures use the same set of bounds), it does not necessarily follow that one is *better* than the other. As noted earlier, the notion of *best* (or better) entails aspects of utility which we do not consider here. On the other hand, if one set of points dominates another set, it *necessarily* has a larger Lebesgue measure.

Fleischer [9] utilized the implications of this framework component and the concept of Pareto optimality to articulate a mathematically oriented definition of SOB and is restated (and slightly modified) here:

> DEFINITION: *Self-organized behavior in a complex system involving multiple performance measures is a sequence of system states corresponding to movement along a Pareto optimal frontier.*

Of course this is only one of many possible definitions for SOB or emergent behavior (see [5]) and is not meant to be applied in all situations described as "self-organized". For example, this definition may not be appropriate for simple dissipative systems or various properties of fluid dynamics. This definition does however seem appropriate for

- systems involving several measures of performance that,
- depend on decisions, *i.e.,* some actor is involved, and
- to which natural selection pressures are applied.

Using this definition of SOB, we now consider the issues involved in "moving along" this Pareto optimal frontier, how it is computed, and how it can be used to effect cooperativity.

3.2 The Dynamical Framework

The second component of the triad, the *dynamical framework,* is based on a new concept—*Scale Invariant Pareto Optimality* (SIPO) that pertains to the tendency of evolving systems towards increased efficiency and cooperative behaviors on many scales.[5]

The value of articulating this component of the meta–formalism is that it forces us to ask useful questions: How should the system *change, i.e.,* what is the basis of the evolutionary dynamics? Because the first principle relates to efficiency and adaptability, the dynamics must be associated with Pareto optimality. This in turn implies the need for some *memory and learning* mechanisms. The dynamical framework thus imposes the following four requirements:

1. a mechanism for changing the behavior of a system
2. a form of memory.
3. a learning mechanism
4. a basis for natural selection

Before delving into these details, let us discuss the third component of the meta-formalism—the problem framework.

3.3 The Problem Framework

The problem framework provides a concrete way for implementing the first two frameworks and helps to define specific problems and questions, narrow the issues and focus research and development efforts. Perhaps the most general and flexible (hence useful) of the many possible problem frameworks are CAs where each cell or component is an FSM. CAs provide a basis for modeling system interactions hence captures important aspects of cooperativity. We can also control the neighborhood structure to suit our research needs. FSMs provide flexibility for controlling the levels of complexity of the system as each cell can be defined with any number of states depending on what is necessary and convenient. We also have flexibility in defining various (even arbitrary) measures of performance. Together, CAs and FSMs provide the substrate on which to define and implement the first two elements of the meta-formalism.

[5]This includes the *graceful degradation of performance* in complex systems. See [9].

It is worth mentioning another approach similar to the CA concept. One could instead use a network paradigm or graph where each FSM is connected to an arbitrary set of neighbors that can then interact. On the other hand, CAs have a more regular, or uniform, structure which might prove useful in studying various phenomena. Keeping these considerations in mind, we can now start to fill in the blanks and further develop this meta-formalism and provide guidelines on how it can be implemented for actual experimentation, research and development.

To summarize, the first component of the meta–formalism suggests that evolved systems should behave in a manner consistent with Pareto optimality and do so on many scales. The second component suggests that efficient/adaptable behavior corresponds to movement along Pareto optimal frontiers associated with several *system* scales. The third component provides the basis for implementing these ideas using the general concepts of FSMs and CAs. The next step is to engineer methods which embody these ideas and satisfy their requirements. Before we do this however, some discussion on modeling cooperativity is helpful.

4. MODELING COOPERATIVITY

Efficient Tradeoffs of Efficiency: What is cooperation? The term "give and take" comes to mind: sometimes entities must sacrifice a short-term benefit for a longer term benefit. Or, it may suggest *sacrifice* where an entity gives up, perhaps permanently, any hope for improving its "performance" so that a larger number of other entities can. As noted in [9], soldiers must often make the ultimate sacrifice for the good of their countries. It seems plain that cooperation is a form of SOB, but how can it be modeled mathematically?

Structures and SOB: This type of tradeoff suggests that an entity, or in this case an FSM, must have some notion of how to modify its behavior for the good of its neighboring FSMs. It further implies the possibility that certain states of neighboring FSMs may be 'incompatible'—a certain Pareto optimum in one FSM may preclude a particular Pareto optimum in its neighbor. This incompatibility is analogous to the concept of *frustration* in Ising spin glass models (see *e.g.*, [4]). Such a phenomenon suggests the possibility that only certain combinations of Pareto optima among neighboring entities may coexist at any given time. If this is true, it may explain the notion of *fluctuations* as Prigogine calls them [18] and the formation of structures as an attribute of SOB as the effects of frustration propagate through a CA. It seems reasonable therefore that the patterns of system states spatially and temporally are governed by initial conditions, various state incompatibilities due to frustration, and the nature of efficient tradeoffs in the midst of these state incompatibilities.

The 'Graceful' Degradation of Performance: The foregoing discussion also suggests a general and mathematical way of describing the "graceful degradation" of system performance. Such graceful degradation has been seen as an important component in the management of large scale systems such as the Internet [1, Ch.5]. Rather than suffering catastrophic changes, it may be possible for system efficiency to be degraded gradually. This is because system efficiency can be measured using the Lebesgue measure described earlier, hence the possibility that the *measures of efficiency for different scales can be traded off efficiently*. This idea of the efficient tradeoffs of efficiency measures provides a *quantifiable guide* of how best to achieve this sacrifice of efficiency on one scale for improved efficiency on another scale.

These types of issues can arise in a variety of contexts: if state incompatibilities exist among neighboring entities, then it is possible that some of the entities cannot operate at their highest levels of efficiency, hence there are tradeoffs of efficiency between individual entities and its neighboring entities. Since the Lebesgue measure is a quantification of efficiency *for entire regimes of behavior*, a mathematical way to characterize this type of tradeoff is to use the Lebesgue measures associated with individual entities and with entire neighborhoods of entities to define a Pareto optimal frontier based on them. In effect, *the Lebesgue measures associated with Pareto optima on different scales themselves become objective functions for a higher-level Pareto optimal frontier.* To the author's knowledge, this type of tradeoff has never been characterized or described in this way, hence provides a novel way to characterize system interactions, behavior and efficiency on different scales and constitutes the basis of the SIPO concept. Figure 3 depicts such a tradeoff curve where the Pareto optimal frontier is the bolded part of the curve. Note that improvement on one scale comes at the cost of worsening the Lebesgue measure on the other scale. See [9].

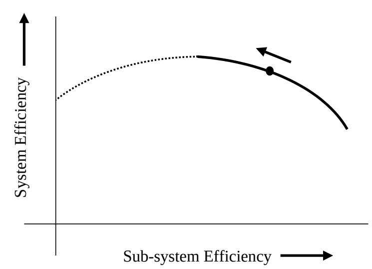

Figure 3: A tradeoff curve of different measures of efficiency.

In Figure 3, the $x-$axis is an estimate of a particular entity's efficiency measure and the $y-$axis is an estimate of a more global efficiency measure. The black dot corresponds to the entity's operating point under current environmental conditions and suggests a possible tradeoff of efficiency where the system's efficiency can be improved at the cost of decreasing the entity's efficiency.

Obviously, if enough subsystems are degraded, it becomes increasingly difficult to adequately compensate for their degradation and the entire system's efficiency measure will be reduced. The point is that this degradation can be influenced or even controlled by actions that minimize the decrease in the Lebesgue measure. Thus, because the system degrades *efficiently*, the system can be said to degrade *gracefully*. How this occurs poses an interesting and, needless to say, difficult problem. The next section describes ways to study this and related problems.

5. IMPLEMENTING THE FORMALISM

5.1 Defining Measures of Performance

Within the FSM/CA framework, virtually any suitable and convenient objective functions can be defined as well as the output functions used to "communicate" with an entity's neighbors. For each entity i in a CA, let $\mathbf{x}_i$ be a vector of state variables, $y_i = g(\mathbf{x}_i, \mathbf{y}_{N_i})$ be some output function for entity i where $\mathbf{y}_{N_i}$ is the vector of output values of i's neighbors, and $f_{ij}(\mathbf{x}_i, \mathbf{y}_{N_i})$ the j^{th} ($j = 1, 2, \ldots, n$) objective function of entity i. These output functions effectively allow entities to communicate and connect them to their neighbors thereby capturing the "community" phenomenon and the attendant interdependence. The function g can also be defined to incorporate aspects of *game theory* [19]. Thus, each objective and output function is dependent on the entity's state variables and the output function values of its neighbors as befits a CA.

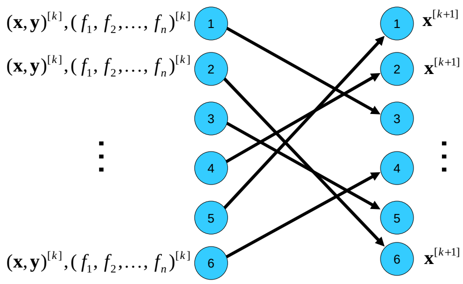

Figure 4: A transformation function that corresponds to a permutation of a list of Pareto optima.

5.2 Transformation Functions

Implementing the dynamical framework requires a mechanism for transforming the current state and environmental information into a decision, *i.e.*, the state at the next time step. A *transformation function* for each entity must therefore be defined. The system is designed so that this function which controls an entity's behavior evolves over time to increase the value of the high-level Lebesgue measure (see below). As such it is the component that ultimately leads to cooperativity and other forms of SOB. Because selection pressures are applied to the transformation functions of each entity, there needs to be a way of relating them among all these entities, otherwise the evolutionary process would be meaningless. Some common ground for comparisons and selection must be identified. This issue is addressed in Section 5.3. For now, we examine what a transformation function looks like.

For an entity to maintain efficient system states, it must be capable of moving from one Pareto optima to another using a transformation function where for a given entity i, this function in generation h can be denoted by

$$\mathbf{x}_i^{[k+1]} = \mathbf{T}^{[h]}(\mathbf{x}_i^{[k]} \mid \mathbf{y}_{N_i}^{[k]}) \qquad (1)$$

where the $\mathbf{x}_i$ are associated with a Pareto optimum at time index k. This function outputs state variables associated with a Pareto optimum at the next time step.

The function in (1) requires an *archive* of the Pareto optima $(f_{i1}, f_{i2}, \ldots, f_{in})$ and their associated values $\mathbf{x}_i$, and $\mathbf{y}_{N_i}$. It is essentially a database of these values where each entry can be designated with a *position number* in the archive. The transformation function thus constitutes a mapping of one position number to another position number, hence is a *permutation*. For a system with a fixed number of operating points p, this mapping can be based on a simple table lookup method and depicted as a bi-partite graph. Figure 4 depicts the permutation

$$\begin{pmatrix} 1\,2\,3\,4\,5\,6 \\ 3\,6\,5\,2\,1\,4 \end{pmatrix}.$$

Changes in (1) thus correspond to a new permutation.

The simplicity of the transformation function (1) leads to several problems and possibilities worth mentioning. One obvious problem is that the inputs depend on the $\mathbf{y}_{N_i}$, but as we shall see, entries in the archive are based only on the objective function values. Thus, an input in (1) may not actually exist in the archive. Given an entity's dependence on $\mathbf{y}_{N_i}$, simply moving from one Pareto optimum to another is also unlikely. How can possible dependence on $\mathbf{y}_{N_i}$ and maintenance on the Pareto optimal frontier be reconciled? This depends on how the meta-formalism is implemented and many possibilities are conceivable. It is worth mentioning a few of them.

One approach would be to use the archive as a training set for a *feed-forward neural network* where all of the inputs values in (1) are matched to the corresponding outputs. Another scheme would be to determine the entry in the archive that is closest to the input values using some metric. That entry's position number would then be used in conjunction with the permutation to yield the output values. Other schemes are no doubt possible.

The next section describes more details of this archive and address issues related to learning, how these transformation functions will be used, and the necessity mentioned earlier of establishing some way to relate the transformation functions among the entities.

5.3 Memory and Learning

For the transformation function to work properly some method for learning and remembering the efficient solutions is needed. This requirement can be satisfied by a search algorithm that stores the efficient solutions as they are discovered. Knowles *et al.* [13] recently developed an archiving method based on the Lebesgue measure that meets these requirements perfectly. Such an archive stores the necessary information in the following way: During the search process, non-dominated solutions are added to the archive up to its maximum storage capacity. Any new points dominating ones already in the archive replace those they dominate thereby maintaining an upper bound on its size. In this way, the archive continually filters out undesirable (dominated) solutions and so *learns* while the Lebesgue measure of the archive monotonically increases.

To function properly, the archive must work *with* the transformation function in some meaningful way. This means some method for classifying them is needed, which in turn, requires that some *order* among the transformation functions be imposed. One approach is to list the points in the archive *lexicographically* based either on the vector $(f_{i1}, f_{i2}, \ldots, f_{in})$ or the vector of state variables $\mathbf{x}_i$. *Ceteris paribus*, every entity's archive will then have a *similar* ordering of their elements yet preserve some essential diversity. For example, the first entries of two distinct archives would have a greater likelihood of being close to one another in terms of some metric than the first and last entries of these archives. Other ordering schemes are no doubt possible. Computational experiments will be required to assess whether this scheme works in terms of evolving cooperativity. The next section describes how such experiments could be designed.

5.4 Evolving Cooperation

In this section, several approaches for *engineering* the evolution of cooperativity are described. These approaches all attempt to increase the Lebesgue measures, *i.e.*, the efficiency, associated with self-interest and community interest *and* the Lebesgue measure associated with the efficient tradeoffs of these interests. This can be done in a straight-forward manner using the Lebesgue measures associated with the appropriate scales.

Self-Interest: Note that a single point in objective function space dominates a certain region of that space which has an associated Lebesgue measure. Thus, the m^{th} vector of objective function values in the archive of entity i maps to some scalar:

$$(f_{i1}, f_{i2}, \ldots, f_{in})_m = \mathbf{F}_{im} \mapsto l_{im} \qquad (2)$$

Also, the entire archive with p entries for entity i maps to a scalar L_i, the Lebesgue measure of the union of the dominated objective function space of each vector in the archive:

$$\left.\begin{array}{c} \mathbf{F}_{i1} \\ \vdots \\ \mathbf{F}_{ip} \end{array}\right\} \mapsto L_i$$

(see [10] for an algorithm that computes this Lebesgue measure).

Community-Interest: The Lebesgue measure reflecting the efficiency of the community is based on the Lebesgue measures of the neighbors of entity i. Let L_{i_N} represent the Lebesgue measure of the N^{th} neighbor of i. Each of these measures is now treated as an objective function value as in (2) which also dominates a region of this new objective function space and maps to another scalar reflecting the efficiency of the neighborhood of i:

$$\{L_{i_1}, L_{i_2}, \ldots, L_{i_N}\} \mapsto \mathbf{L}_i$$

Cooperativity Measure: Finally, the Lebesgue measure of the space dominated by both self-interest and community interest is defined:

$$\{L_i, \mathbf{L}_i\} \mapsto \mathcal{L}_i \tag{3}$$

and provides a way to *measure cooperativity*. This Lebesgue measure thus captures the *economic* vitality of a subset of FSMs, hence can be used as a fitness function—entities with higher values of $\mathcal{L}_i$ will tend to cooperate more effectively with their neighbors.

Evolving Transformation Functions: Once entities with relatively high fitness values $\mathcal{L}_i$ are identified, a genetic algorithmic approach is needed to *evolve* these transformation functions. Because transformation functions are essentially permutations, crossover and mutation operators for permutations must be defined. Although permutations are a somewhat strange entity on which to apply these genetic operators, such operators on permutations do exist and have been studied over the years. Many of these techniques are designed to preserve partial orderings in successive generations. See [16] for an extensive survey. The next section describes some general approaches for putting all of these elements together to evolve cooperativity and SOB.

5.5 The Meta-Formalism Process

Several alternative approaches for implementing these ideas are no doubt possible. Here, three distinct phases to the evolutionary process are described although many variations are conceivable. The following is offered to stimulate further research.

Phase 0—Initialization: For each FSM, system states evolve randomly while non-dominated points in objective function space and the associated values of $\mathbf{x}_i$ and $\mathbf{y}_i$ are archived. This phase probabilistically explores the state space and remembers/learns the set of efficient operating points discovered in the process.

Phase 1—Learning & Training: This phase either creates a random transformation function if one does not already exist, or uses a transformation function developed in Phase 2. In this phase, the system goes through the fixed number of iterations under the influence of the transformation function. The archive is updated with new non-dominated points while lexicographic ordering is maintained. This phase can incorporate *supervised learning* using neural networks or some metric to determine the input-output mapping as described in Section 5.2. Lebesgue measures are computed either at each iteration or at the end of this phase.

Phase 2—Genetic Modification: This phase utilizes various evolutionary paradigms to induce improvement in the transformation functions. A fitness value for each entity is computed using (3). Alternatively, the self-interest measure l_i in (2) or its time average $\bar{l}_i^{[k]}$ over the number of iterations can be used in place of L_i in (3) to compute the fitness value $\mathcal{L}_i^h$ for generation h. Entities with relatively high fitness values are then selected out and crossover and mutation operations applied to their respective transformation functions. A new generation of transformation functions then replaces the existing ones. The process then reverts to Phase 1 hence oscillates between these phases. The test of whether this scheme works can be measured by the sequence of the average values of (3) among the entities. This value ought to increase over time although it may not do so monotonically.

6. DISCUSSION

What can we expect to observe from an implementation of this scheme and what sorts of reasonable conjectures, if any, can be made? Clearly, an implementation would likely have all the complexity and unpredictability of typical CAs—yet, its very structure, might lead to some interesting phenomena. One of the more fascinating phenomena to investigate is *self-organized criticality* (SOC) [3]. SOC describes the sudden, often dramatic, changes that occur in complex systems. Many examples involve swarm intelligence, but this phenomena is also associated with the statistical mechanics of *phase transitions*. How might SOC be observed and studied within this meta-formalism? How does it relate to SOB?

Consideration of other paradigms and phenomena such as game theory, economic theory, and the evolutionary paradigm itself might help in answering these questions. For instance, one could design objective function relationships that model some form of competition: a cell's objective in maximizing some function could be inversely related to the outputs of the same function in one of its neighbors. Add to this dependence of a cell's "success", as measured by objective functions or (2), to the success of its neighbors and it is very likely that some very complex behaviors might be observed. Conway's Game of Life CA (see *e.g.,* [17]) displays some aspects of this in a much simpler system. In more complex systems with richer sets of possible behaviors, could such competitive and cooperative features lead to SOB and would it also display SOC? See *e.g.,* [19, 15].

Consider the possibility of the existence of an "ideal" transformation function that leads to very advantageous states for a given cell, but that its long-term success is too impaired by the less than ideal transformation functions of its neighbors. This ideal transformation function may then not be able to replicate sufficiently and ultimately dies off or is mutated. Perhaps this ideal transformation function can only succeed *and propagate* if a sufficient number of its neighbors have a similar transformation function. One could then imagine that although this *proximity of ideal transformations* may have a much smaller probability of occurring, once it does occur there could then be a *cascade of evolution* where viable transformation functions propagate in sufficient numbers and cause the system to behave in a maximally efficient manner. A threshold number of these viable transformation functions may thus manifest itself in a "domino effect" and could be considered a type of SOC. See also [7, 12] for a discussion relating to 'punctuated' evolution. Perhaps the concept of SOC is fundamentally related to some form of chain reaction and requires some form of proximity characteristics for it to occur. This idea might be explored by restricting crossover operations to involve only neighboring entities.

All of the flexibility in experimenting with this meta-formalism can be used to explore this and other possibilities. On the other hand, true understanding of phenomena often requires the simplest of systems to manifest the phenomena where the essential features of the phenomena are more observable (again, the domino example is compelling). The meta-formalism described here could easily become overly complicated in any implementation and so caution is advised. It should be borne in mind, that the crux of this meta-formalism is the use of a metric, as in (3), that directly measures cooperativity and this value will intimately depend on how the objective, output and transformation functions are defined and genetically modified.

7. CONCLUSION

This article described a meta-formalism that can be used to define formulae and iterative schemes capable of demonstrating the evolution of cooperativity among a set of interacting systems. Using notions of Pareto optima and its measure, a mathematically oriented definition of SOB was described. The meta-formalism itself was based on three foundational components: a set of first principles based on the laws of evolution and natural selection. Its implications naturally led to concepts such as Pareto optimality, efficiency and adaptability as the basis for the other foundational components. The second component was the dynamical framework and described measures of performance and efficiency measures associated with Pareto optima. Transformation functions, an archiving process to effect memory and learning, and an evolutionary process to select out transformation functions associated with a high-level efficiency measure—one that provides a direct measure of cooperativity—were also described. The third component was a problem framework that provided the modeling clay with which to mold a system of systems together. This employed the concept of FSMs linked in a CA to provide both simplicity and generality.

This meta-formalism and the ideas in this article will hopefully provide a basis for organizing ideas, stimulate further research and development, and help in evolving a greater understanding of complex and self-organizing systems.

Acknowledgments

I would like to thank four referees for reminding me that I do not practice law anymore and for their very helpful suggestions which hopefully have led to an improved article. I would also like to thank Donna Gregg, Susan Lee and the Johns Hopkins University Applied Physics Laboratory for its support under the auspices of its Independent Research and Development program.

8. REFERENCES

[1] *Making the Nation Safer: The Role of Science and Technology in Countering Terrorism*, chapter 5: Information Technology, pages 135–176. National Academy Press, Washington, D.C., 2002.

[2] E. Aarts and J. Korst. *Simulated Annealing and Boltzmann Machines: A Stochastic Approach to Combinatorial Optimization and Neural Computing*. J. Wiley & Sons, 1989.

[3] P. Bak, C. Tang, and K. Wiesenfeld. Self-organized criticality. *Physical Review A*, 38(1):364–375, 1988.

[4] Y. Bar-Yam. *Dynamics of Complex Systems*. Perseus Books, Reading, MA., 1997.

[5] Y. Bar-Yam. A mathematical theory of strong emergence using multiscale variety. *Complexity*, 9(6):15–24, July/August 2004.

[6] E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Santa Fe Institute Studies in the Sciences of Complexity. Oxford University Press, New York, N.Y., 1999.

[7] J. Crutchfield and E. van Nimwegen. The evolutionary unfolding of complexity. In L. Landweber, E. Winfree, R. Lipton, and S. Freeland, editors, *Evolution as Computation, Proc. DIMACS Workshop*, LNCS, Princeton University, January 1999. Springer-Verlag. Santa Fe Institute Working Paper 99-02-015.

[8] G. Di Caro and M. Dorigo. Antnet: Distributed stigmergetic control for communications networks. *Journal of Artificial Intelligence Research*, 9:317–365, 1998.

[9] M. Fleischer. Foundations of swarm intelligence: From principles to practice. Technical Research Report TR-2003-10, Institute for Systems Research, University of Maryland, College Park, Maryland, 2003.

[10] M. Fleischer. The measure of pareto optima: Applications to multiobjective metaheuristics. In C. M. Fonseca, P. J. Fleming, E. Zitzler, K. Deb, and L. Thiele, editors, *2nd International Conference on Evolutionary Multi-Criterion Optimization, EMO 2003*, volume 2632 of *LNCS*. Springer-Verlag, 2003.

[11] J. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, Michigan, 1975.

[12] S. Kauffman. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, Inc., New York, NY, 1993.

[13] J. Knowles, D. Corne, and M. Fleischer. Bounded archiving using the lebesgue measure. In *Proc. of the IEEE Congress on Evolutionary Computation*, volume 4, pages pp.2490–2497. IEEE Press, 2003.

[14] A. Leon-Garcia and I. Widjaja. *Communication Networks: Fundamental Concepts and Key Architectures*. Computer Science and Electrical and Computer Engineering. The McGraw-Hill Companies, Inc., New York, NY, 2000.

[15] M. Millonas. Artificial life iii. volume XVII of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 417–445, New York, NY, 1994. Addison-Wesley Pub. Co.

[16] A. Moraglio and R. Poli. Topological crossover for the permutation representation. Technical Report CSM-408, Department of Computer Science, University of Essex, Colchester, UK, 2004.

[17] W. Poundstone. *The Recursive Universe: Cosmic Complexity and the Limits of Scientific Knowledge*. Contemporary Books, Chicago, Ill., 1985.

[18] I. Prigogine and I. Stengers. *Order Out of Chaos: Man's New Dialog with Nature*. Bantam Books, Inc., New York, NY, 1984.

[19] T. C. Schelling. *The Strategy of Conflict*. Oxford University Press, New York, NY, 1960,77.

[20] R. Serra and G. Zanarini. *Complex Systems and Cognitive Processes*. Springer-Verlag, New York, NY, 1990.

[21] D. Solow and J. Szmerekovsky. Mathematical models for explaining the emergence of specialization in performing tasks. *Complexity*, 10(1):37–48, 2004.

[22] N. Wiener. *Cybernetics: Or Control and Communication in the Animal and the Machine*. The M.I.T. Press, Inc., Cambridge, MA., 2 edition, 1961.

[23] E. Zitzler, M. Laumanns, L. Thiele, C. Fonseca, and V. da Fonseca. Performance assessment of multiobjective optimizers: An analysis and review. *IEEE Transactions on Evolutionary Computation*, 7(2):117–132, 2003.

[24] E. Zitzler and L. Thiele. Multiobjective optimization using evolutionary algorithms– a comparative case study. In M. S. H. S. A.E. Eiben, T. Bäck, editor, *Proc. of the 5th International Conf. on Parallel Problem Solving from Nature—PPSN V*, Berlin, Germany, 1998.