

# Dependable Information Infrastructures as Complex Adaptive Systems

Nong Ye,\* Ying-Cheng Lai, and Toni Farley

Arizona State University, Tempe, AZ 85287

Received 20 October 2002; Accepted 20 May 2003

DOI 10.1002/sys.10048

## ABSTRACT

Information infrastructures, such as the Internet and Computational Grids, have enabled a networked computing and communication environment. Although many organizations depend on network-centric information operations to support critical missions, existing information infrastructures provide little guarantee of the dependability of network-centric computing and communication. This paper discusses some problems with the dependability of existing information infrastructures, such as stateless or centralized resource management. A Complex Adaptive Systems approach to dependability of futuristic information infrastructures is then presented with emphasis on the detection of emergent states at the regional and global levels of these infrastructures, and the self-synchronized control of such infrastructures in response to emergent states. The key concept underlying our control strategy at the regional or global level is that the two attributes of the Complex Adaptive System in a dynamic environment, accessibility to many states and sensitivity to small perturbation, present us with an opportunity to manipulate the system's dynamics. © 2003 Wiley Periodicals, Inc. *Syst Eng* 6: 225–237, 2003

**Key words:** information infrastructures; Internet; Computational Grids; computer security; complex adaptive systems; complexity theory

---

\*Author to whom all correspondence should be addressed at: Nong Ye, Professor of Industrial Engineering, Affiliated Professor of Computer Science and Engineering, Arizona State University, Information and Systems Assurance Laboratory, P.O. Box 875906, Tempe, AZ 85287-5906. (e-mail: nongye@asu.edu).

Contract grant sponsors: U.S. Department of Defense and U.S. Air Force Office of Scientific Research; contract grant number: F49620-01-1-0317.

*Systems Engineering*, Vol. 6, No. 4, 2003  
© 2003 Wiley Periodicals, Inc.

## 1. INTRODUCTION

An information infrastructure connects information systems belonging to different administrative domains. Information infrastructures, such as the Internet and Computational Grids, have enabled a networked computing and communications environment, which has been used by many organizations to perform network-centric information operations for critical missions.

Network-centric information operations are characterized by information moving between computational nodes on the network, towards information superiority with 100% relevant content and accuracy, and zero time delay for the content, quality, and timeliness of information. Network-centric information operations depend on network-centric operational architectures, such as a high-performance information infrastructure that provides a speedy, dependable backplane sharing of various computational resources (including data resources, computation resources, communication resources, and visualization resources) on the computing and communications network.

This paper addresses the dependability of information infrastructures by addressing performance goals and control strategies at the regional and global levels of an information infrastructure based on Complexity Theory for Complex Adaptive Systems in an effort to overcome existing problems. We first discuss problems with the dependability of existing information infrastructures. We then describe an information infrastructure as a Complex Adaptive System, followed by an outline of control strategies for dependability according to Complexity Theory for Complex Adaptive Systems [Badii and Politi, 1997; Lewin, 1999]. Next, we give some background concepts from nonlinear dynamics and complex systems [Ott, 2002], followed by discussions on detecting emergent states and controlling about desirable and undesirable emergent states. Finally, we provide a brief summary of the paper.

## 2. PROBLEMS WITH EXISTING INFORMATION INFRASTRUCTURES

Existing information infrastructures, such as the Internet and Computational Grids, provide little dependability guarantees for networked computing and communication. The Internet supports the sharing of computational resources on a network based on the “best-effort” model, in which trust among participating parties is all but assumed. That is, participant A will satisfy the need of other participants as long as the current capacity of participant A’s computational resources can still do so. Hence, resources are made available for use by anyone regardless of the state of the resources until those resources are completely depleted. This “best-effort” model has provided an environment for crafty exploits and denial-of-service attacks that have occurred and presented a significant threat to the realm of information superiority. For example, if a malicious hacker sends large amounts of data in a very short period of time to a router covering a certain sector of an information infrastructure, the router keeps filling

up its queue with incoming data streams until its queue is full and no other data streams can be accepted and routed, thereby causing a denial-of-service attack. Moreover, quality of service (QoS) for a user’s application is not guaranteed because other users may emerge at any time to compete for and share computational resources. Thus, little resource management and QoS guarantees exist for dependable information operations on the Internet.

Computational Grids [Foster and Kesselman, 1999], the so-called next-generation information infrastructure, address the resource management problem by using centralized brokers or mating agents to assign networked computational resources to users’ applications. The centralized authority of resource management (top-down, command-directed synchronization of network-centric information operations) works only within a small-scale networking environment of closely-coupled administrative domains, such as a network of several national supercomputing centers, as currently demonstrated [National Science Foundation, 2000].

Centralized authority of resource management does not scale up to an information infrastructure where many independent administrative domains exist, as these domains may not necessarily obey this authority. First of all, administrative domains for different organizations likely act on their own to rapidly respond to local situations and meet local contingencies as they see them. Secondly, information operations through interactions among different administrative domains are typically complex and nonlinear. These two factors make even local influences from closely coupled, friendly administrative domains hard to predict. That is, the information infrastructure itself is a Complex Adaptive System which best organizes from the bottom up through self-synchronization according to Complexity Theory [Badii and Politi, 1997; Lewin, 1999]. The top-down, centralized synchronization of resource management in Computational Grids cannot scale up to a large-scale, complex, adaptive information infrastructure in a dynamically changing environment.

Theoretical research on marketing mechanisms of sharing resources, such as auctioning, bidding, and negotiation, also exists [Foster and Kesselman, 1999]. However, auction and bidding are traditionally used to ensure a *fair* competition for a *static* object, such as an art piece. Auction and bidding may not be appropriate for administration domains in the information infrastructure. It does not make sense for an administrative domain to auction its resources, because service requests for the resources and the state of resources change *dynamically*. It is also improbable that an administrative domain for a private organization or a

military organization would auction its resources for bidding. Moreover, it is less reasonable to assume that negotiations for sharing resources can be carried out in a completely automatic fashion. It is equally unreasonable to assume that administrative domains will be able to negotiate with each other about the sharing of resources when contingencies occur in real time. It is more realistic to have a dynamic coalition of administrative domains. For example, a dynamic coalition of corporate administrative domains form contracting relationships and agree as to how resources are allocated and where the resources are located when the dynamic coalition is created. For an administrative domain *dependability* is of greater concern than *fairness*.

### 3. AN INFORMATION INFRASTRUCTURE AS A COMPLEX ADAPTIVE SYSTEM

An information infrastructure is like a marketplace where both the demand for, and supply of, computational resources exists. System resources, such as those used for computation, data, communication, and visualization, provide services to processes generated by users' applications. These resources include such things as CPUs, files, databases, routers, communication links, and display devices.

An information infrastructure consists of a large number of administrative domains, among which computational resources may be shared. Each domain is an agent that independently manages its own resources and schedules arriving processes, as shown in Figure 1. Some administrative domains may form a dynamic

coalition of agents (see Fig. 1) over a period of time in order to share computational resources for superior information operations, such as distributed information storage/retrieval, parallel computation, and fault tolerance. For example, administrative domain A can use resources (e.g., host machines) or keep redundant resources (e.g., data files) in administrative domain B as back-up resources for when A experiences intrusions or failures of service. Administrative domains belonging to companies with contracting relationships, such as an auto maker, its distribution centers, and its retailers, may form a dynamic coalition to have closely-coupled relationships.

There are a variety of resources in the administrative domain of an agent. A resource (R), and processes (P) requesting services, along with the supply-demand relationships between the resource and processes, form a dynamic system as shown in Figure 1. Queues (Q) may be used to store processes waiting for a resource. For example, an administrative domain may consist of a router, several host machines, and the communication links connecting them. A user's ftp application enters the domain to request some information that resides on one of the domain's host machines. This application yields a series of processes: the process requesting the router's service to send a request to the host machine, the process requesting the host's service to retrieve and send the requested information to the router, and the process requesting the router's service to forward the information to the source of the request. The router may have two queues: an inbound queue for data packets

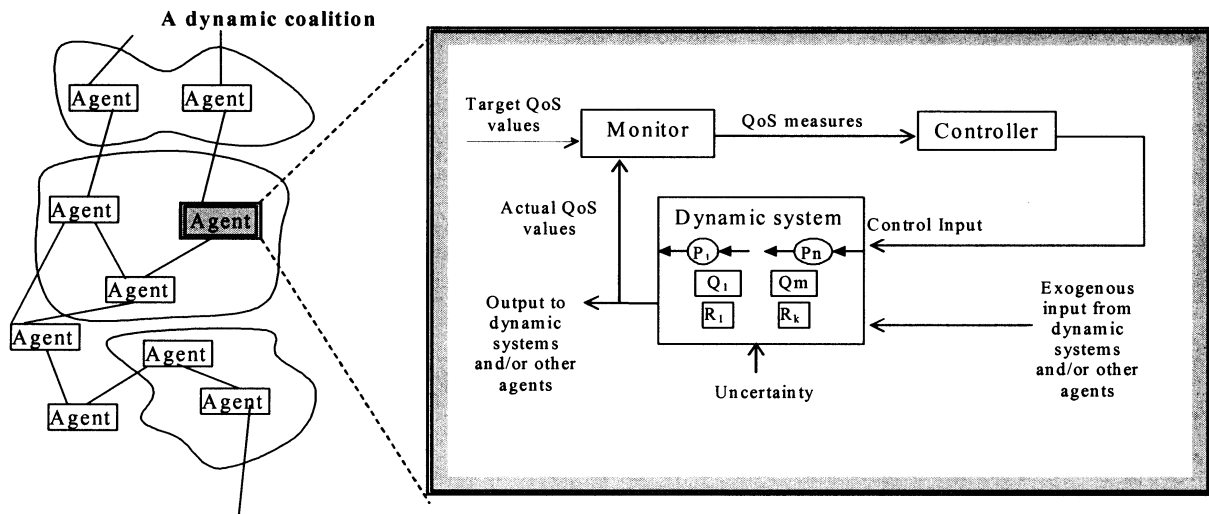


Figure 1. The dependable information infrastructure as a Complex Adaptive System.

entering the domain and an outbound queue for data packets leaving the domain.

An information infrastructure is a Complex Adaptive System with a large collection of agents interacting in a dynamically changing environment. Complexity Theory for Complex Adaptive Systems states that the information infrastructure best organizes from the bottom up through the self-synchronization of agents and their local interactions [Badii and Politi, 1997; Lewin, 1999]. The aggregate system behavior of a set of agents and interconnections emerges from the evolving, local interactions of the agents in a dynamically changing environment. We describe the bottom-up self-synchronized control strategies of QoS-centric stateful resource management for dependable information infrastructure in an effort to overcome problems with existing information infrastructures: the top-down centralized resource management in Computational Grids and the stateless resource management in the Internet. The control strategies operate on resources and processes with respect to performance goals at three levels: the local level within an individual agent, the regional level within a dynamic coalition of agents, and the global level in the information infrastructure.

Considering individual processes that request services from individual resources in the administrative domain of an agent, an agent's local-level performance goal is to provide QoS guarantees for the processes that are admitted. The following QoS-guarantee functions are usually required: admission control, process scheduling, and QoS conformance monitoring. There are many research efforts regarding QoS metrics for the output performance of processes, and QoS-centric management of individual resources within an administrative domain [Chatterjee et al., 1997; Giroux and Ganti, 1999; Lawrence, 1997; Liu, Xiong, and Sun, 2000; Sabata et al., 1997]. This paper focuses on the performance goals and control strategies at the regional and global levels of the information infrastructure.

#### **4. CONTROL STRATEGY AT THE REGIONAL AND GLOBAL LEVELS OF THE INFORMATION INFRASTRUCTURE**

A dynamic coalition at the regional level of the information infrastructure usually includes a large number of administrative domains. When the dynamic coalition is formed, the administrative domains reach an agreement as to how resources are allocated and where the resources are located, including both operational and backup resources. The information infrastructure at the global level also includes a large number of administra-

tive domains, some of which may form dynamic coalitions.

Neither the dynamic coalition nor the information infrastructure as a system has a meta-agent for centralized control. Each system is a Complex Adaptive System with a large number of agents interacting with each other in a nonlinear fashion. Within each system, individual agents determine local control actions, which affect the resources and processes within their administrative domains, as well as interactions with other agents. For example, processes rejected by one agent may be directed to other agents with similar resources. Such information operations are usually non-linear [Coombs and Birx, 2000]. The same control strategy based on bottom-up self-synchronization from Complexity Theory [Badii and Politi, 1997; Lewin, 1999] is appropriate for both the dynamic coalition and the information infrastructure.

The individual resources and processes in an administrative domain are managed by its agent. Therefore, a dynamic coalition and the information infrastructure as a Complex Adaptive System are more concerned with aggregate system performance; specifically, aggregate traffic flows at the regional and global levels respectively. Such flows emerge from the inbound and outbound traffic flows of individual administrative domains. The performance goal at the regional or global level is to optimize aggregate traffic flows at that level. Hence, with respect to control strategy and performance goals, a dynamic coalition and the information infrastructure differ mainly in the size or scope of the system.

To measure aggregate traffic flows at each level, we can set up a number of observation points within the scope of the respective system to obtain collective measures of traffic flows at those points. For example, at the regional level of a dynamic coalition, we can set up a number of observation points at key routers that control the coalition's main traffic flows. Queue lengths and traffic flow rates [Giroux and Ganti, 1999; Liu, Xiong, and Sun, 2000] are typically used to measure traffic flows through routers.

In either the regional or the global system, we can designate an agent to collect time-series measurements of traffic flows at the observation points. The agent then reports situations occurring at the level of interest to individual agents within that scope. Such a designated agent is analogous to the National Weather Forecasting Center that collects weather data and disseminates weather reports.

The environment associated with a regional (dynamic coalition) or global (information infrastructure) system is generally uncertain and extremely dynamic. It is highly unlikely that in such an environment, aggregate traffic flows in a Complex Adaptive System can

settle into a stable equilibrium so that system performance is stable. Given that the system is persistently unstable, our tasks are to:

1. Detect an unstable aggregate or emergent state embedded in the phase space
2. Develop a control strategy to stabilize the system about an emergent state that gives desirable performance of aggregate traffic flows
3. Develop a control strategy to drive the system out of an emergent state that gives undesirable performance of aggregate traffic flows.

Emergent states are typically ordered states, each with its own basin of attraction. That is, in a noiseless situation, each emergent state has a predefined set of initial conditions in the phase space that asymptote to it. Chaotic or random behavior is usually due to chaotic sets on the basin boundary. A chaotic set typically has embedded within itself an infinite number of unstable periodic orbits. In an unstable or noisy environment, the system can move from one basin to another, generating various temporal emergent behaviors in the course of time evolution.

At the regional or global level, the performance goal of optimizing aggregate traffic flows is to minimize the number and time of crises, in that level's aggregate traffic flows, which correspond to undesirable emergent states of aggregate system behavior. To illustrate this point, let us consider a transportation network (a simple analogy of a dynamic coalition or information infrastructure). In a transportation network, vehicle drivers and transportation departments are the agents. Vehicle drivers decide on the movement of vehicles that request services from road resources. Transportation departments manage the use of the road resources by setting and adjusting traffic lights, opening and closing streets, designating detours, and so on. The movement of a vehicle, the on-off switch of a traffic light, or the open-close of a street produces only local impact. However, a crisis, such as traffic congestion arising from a major traffic accident and its chain effects of stopping or slowing down traffic, can produce much larger regional or possibly global impact, which in turn becomes the concern of many agents in the scope of the impact. This is why many vehicle drivers not only watch local traffic situations but also listen to radios for regional or global traffic reports. When no crises occur in a transportation network, its agents will not try to interfere with the aggregate behavior of the network. Therefore, aggregate traffic flows are left alone to follow a natural course, except when signs of crises arise and call for action.

Similarly, in the dynamic coalition or information infrastructure, an individual agent of an administrative domain must monitor not only QoS of processes in the administrative domain with respect to the local performance goals, but also watch for crises in regional or global traffic flows with respect to the regional and global performance goals.

The control strategy to minimize the number and time of crises in aggregate traffic flows at each level is based on the concept that the dynamic coalition or information infrastructure as a Complex Adaptive System best organizes from the bottom up through self-synchronization. Consider the handling of crises, such as traffic congestions, in a transportation network. If all vehicle drivers listen to radios reporting regional or global traffic flows and adjust the movement of their vehicles, such as taking a detour to avoid the area of traffic congestion, the traffic congestion can be greatly alleviated before corrective actions are taken to clear the accident scene and restore the normal road condition.

Hence, the control strategy at the regional or global level includes the following steps:

1. Each agent closely monitors the state space of the regional or global system in which it belongs to detect embedded emergent states.
2. If the emergent state is desirable, each agent in the system acts locally to stabilize the system about that state; if the emergent state is undesirable, each agent in the system acts locally to drive the system out of that state.

Since stabilizing the system about desirable emergent states prolongs the time of desirable system performance, and driving the system out of undesirable emergent states shortens the time of undesirable system performance, the control strategy will minimize the number and time of crises in aggregate traffic flows at the regional or global level. The following sections discuss detecting emergent states and controlling the system about those states. Since several concepts in complex systems and nonlinear dynamics will be used in the discussion, we first present some background material.

## 5. BACKGROUND: NONLINEAR DYNAMICS AND COMPLEX SYSTEMS

A dynamical system can be mathematically described by either a set of first-order differential equations that do not depend explicitly on time, or by a set of discrete-time maps. The dynamical variables in these equations

constitute the *phase space*, and the number of them is the phase-space dimension.

Given a set of initial conditions for the system, its evolution corresponds to a *trajectory* in the phase space. It can occur that trajectories from many different initial conditions approach in time to some asymptotic sets in the phase space. These are the *limit sets*. *Attractors* are one class of limit sets. In order for a limit set to be attracting, the phase-space region from which trajectories approaching the set are initiated must have a finite volume. The finite phase-space region, containing all initial conditions leading to trajectories approaching an attractor, is the *basin of attraction* of that attractor.

Attractors can be nonchaotic or chaotic. The former include limit cycles and periodic or quasiperiodic attractors. Chaotic attractors are characterized by a sensitive dependence on initial conditions. That is, two trajectories starting out with a small difference in their initial conditions are separated exponentially in time. The rates of the exponential separation in different eigendirections along a trajectory define the spectrum of the *Lyapunov exponents*.

Chaotic limiting sets are not necessarily attracting. It is common for dynamical systems arising from many natural phenomena to have *nonattracting chaotic sets*, for which the set of approaching initial conditions has zero phase-space volume. Thus, for a random initial condition, the trajectory can stay near a nonattracting chaotic set for only a finite time before leaving the set and approaching an attractor. Chaos in this sense is merely *transient*. For a system with multiple coexisting attractors, each with its own basin of attraction, the presence of a nonattracting chaotic set on the basin boundary can lead to serious difficulties in predicting to which attractor a specific initial condition approaches. When the system has many coexisting attractors, nonattracting chaotic sets on the basin boundary can give rise to a high level of complexity in the dynamics, rendering the system complex.

A complex system is characterized by the following: (1) it consists of many components that are interconnected in a complicated manner; (2) the components can be either regular or irregular; and (3) the components exist on different length and/or time scales, i.e., a complex system exhibits a hierarchy of structures. These are also called the *three traits* of a complex system [Badii and Politi, 1997].

According to this characterization of complexity, a nonlinear dynamical system, even if its equations are simple, can exhibit complex behaviors. Examples include situations where there are many coexisting attractors and nonattracting chaotic sets on the basin boundaries, or where there are a few coexisting attractors but a nonattracting chaotic set permeates the entire

basin of attraction [Poon and Grebogi, 1995; Lai and Grebogi, 1996; Lai, 1999].

## 6. DETECTION OF EMERGENT STATES

To detect emergent states, the nature of emergent states must be fully understood. An emergent state is associated with recognizable changes in the invariant properties of the underlying state. In particular, an emergent state is accompanied by a significant reduction in the number of degrees of freedom, or dimensionality, of the dynamics and, possibly, by a high level of functional correlation (or generalized synchronization) among the interacting agents. There exist well-established methodologies in nonlinear dynamics that can be adapted to computing system dimension and detecting possible synchronization of Complex Adaptive Systems. In this regard, two situations need to be distinguished: (1) An accurate mathematical model of the system, represented either by a set of continuous-time differential equations or discrete-time maps, exists; or (2) such a model does not exist and the only information available about the system is the measurements of a limited number of intrinsic system variables. In the first case, possible emergent states can be detected by direct numerical simulation of the system equations to yield, for instance, (i) the set of unstable periodic orbits embedded in the dynamics and their eigenvalues, and (ii) the stability properties of all coexisting emergent states. In the second case, the system dynamics has to be reconstructed from the time-series measurements to yield information about the various emergent states. From the reconstructed phase space, one can detect unstable periodic orbits, emergent states, and their stability properties.


We emphasize that, as we described in the background materials, simple but nonlinear systems can exhibit complex behaviors [Poon and Grebogi, 1995; Lai and Grebogi, 1996; Lai 1999]. Thus, for complex systems, both situations can be expected where the system equations can or cannot be put on paper. In the former case, the availability of a mathematical model can facilitate searching for emergent states. In the latter case, identification of emergent behaviors can be done through measurements, or time series. If only a single time series is available, the requirement is that it must be sufficiently long so that the corresponding trajectory in the phase space visits many emergent states. The availability of multiple time series from a set of independent measurements can alleviate the otherwise stringent length requirement for time series. Emergent states from a complex system are commonly unpredictable, due to the three traits characterizing complexity. This

unpredictability is caused, however, by the intrinsic property of the system itself. It does not depend on whether an explicit mathematical model for the system is available or not. From model or time series, one can learn a great deal but only about the statistical properties of the possible emergent states. This is so because simulations using different sets of initial conditions or independent measurements generally result in different trajectories or time series, which are useful only in the statistical sense. The unpredictability has little to do with whether the system equations are simple or the measured time series are sufficiently long.

A technique that has been utilized commonly in chaotic phase-space reconstruction, and that we will apply to a regional or global system as a Complex Adaptive System, is the delay-coordinate embedding method to reconstruct the phase space [Takens, 1981]. Suppose that we have collected the time-series measurements, e.g., queue lengths or traffic flow rates, at a number of observation points in the system. Such time series typically assume discrete values, but they are approximately continuous if the amount of data is large. Let  $u_i(t)$  ( $i = 1, \dots, l$ ) be a set of  $l$  such time series. Since the system is complex, we conceive that in order to describe it, a large number, say  $M$ , of independent variables is necessary. The number of independent variables necessary to describe the system is the dimension of the system. Typically we expect  $l$ , the number of available time series, to be less than  $M$ . In order to be able to estimate  $M$  and assess other dynamical properties of the system, it is necessary to reconstruct a phase space with a sufficient number of independent variables.

In the delay-coordinate embedding technique [Takens, 1981], from each measured time series  $u_i(t)$  ( $i = 1, \dots, l$ ), we build up a vector with  $q$  components, as follows:  $\mathbf{u}_i(t) = \{u_i(t), u_i(t + \tau), \dots, u_i[t + (q - 1)\tau]\}$ , where  $\tau$  is the delay time, as shown in Figure 2. In order for the time-delayed components  $u_i(t + j\tau)$  ( $j = 1, \dots, q - 1$ ) to serve as independent variables, the delay time  $\tau$  needs to be chosen carefully. Roughly, if it is too small, then adjacent components, say  $u_i(t)$  and  $u_i(t + \tau)$ , will be too correlated for them to serve as independent coordinates. If, on the other hand,  $\tau$  is too large, then neighboring components are too uncorrelated for the purpose. Empirically, one can examine the autocorrelation function of  $u_i(t)$  and decide a proper delay time [Theiler, 1986]. Since we have  $l$  time series available, we can then construct a vector with  $m \equiv ql$  components, as follows:

$$x(t) = \{u_1(t), u_1(t + \tau), \dots, u_1[t + (q - 1)\tau],$$



**Figure 2.** A schematic illustration of delay coordinates from a single time series  $u_i(t)$ .

$$u_2(t), u_2(t + \tau), \dots, u_2[t + (q - 1)\tau], \dots,$$

$$u_l(t), u_l(t + \tau), \dots, u_l[t + (q - 1)\tau]\}.$$

In principle, with a given  $l$ , we can choose  $q$  so that the total number components  $m$  in the vector  $x(t)$  is much larger than  $M$ , the number of independent variables needed to describe the system. It is necessary to stipulate that  $m$  be larger than  $M$  because of the mathematical theorem by Takens [1981], which guarantees that the reconstructed phase space in which the vector  $x(t)$  lives is a faithful representation of the original dynamical system. In our case,  $m \geq 2M + 1$ . One then obtains a one-to-one correspondence between the reconstructed phase space and the original dynamical system (Sauer, Yorke, and Casdagli, 1991). We shall refer to  $m$  as the embedding dimension. From the reconstructed phase space,  $M$  can be estimated as the dimension of the system—the number of independent variables necessary to describe the system.

As the system evolves in time, various emergent states appear. The dimensionality of these states, or the number of independent variables needed to describe them, is typically different from  $M$ . For instance, one can conceive a particular type of emergent state where a substantial number of agents becomes synchronized in their dynamical behavior. Due to synchronization, the number of independent variables (also called the number of degrees of freedom) that are necessary to describe the state can apparently be smaller than  $M$ , roughly the number of independent variables required to describe the system when all agents behave independently (no synchronization). Thus, sensing charac-

teristic changes in the dimensionality of the system state is an effective way to detect emergent states. In what follows we shall use  $D$  to denote the dimensionality of emergent states (to distinguish from  $M$ ).

A commonly computed dimension in nonlinear time series analysis is the correlation dimension  $D_2$ , which we take to be a good approximation of  $D$ . Grassberger and Procaccia [1983] showed that  $D_2$  can be evaluated using the correlation integral  $C(\varepsilon)$ , which is defined to be the probability that a pair of points, chosen randomly in the reconstructed phase space, is separated by a distance less than  $\varepsilon$  on the emergent state. Typically,  $C(\varepsilon)$  scales with  $\varepsilon$  as  $C(\varepsilon) \sim \varepsilon^{-D_2}$ , which provides a computational way to evaluate  $D_2$  [Grassberger and Procaccia, 1983; Ding, Ott, and Grebogi, 1993; Lai, Lerner, and Hayden, 1996; Lai and Lerner, 1998].

Hence, a designated agent in a regional or global system:

- Collects the time-series measurements of aggregate traffic flows at the observation points
- Detects emergent states
- Evaluates these states as to whether or not they are desirable
- Informs individual agents in the system of these desirable or undesirable emergent states.

To evaluate whether or not an emergent state is desirable, we can use Statistical Process Control (SPC) techniques [Ye, Borrer, and Zhang, 2002; Ye and Chen, 2003; Ye, Emran, Chen, and Vilbert, 2002; Emran and Ye, 2002] to compare the time-series measurements of traffic flows resulting in the emergent state against the “normal” time-series measurements of traffic flows before the emergent state. The emergent state is desirable if the measurements of aggregate traffic flows are significantly better in the emergent state than those in the “normal” state. The emergent state is undesirable if the measurements of aggregate traffic flows are significantly worse in the emergent state than those in the “normal” state.

To learn and control a complex system, the recent concept of “attribution of credit” is interesting [Axelrod and Cohen, 2000]. To explain the idea, we quote directly from Axelrod and Cohen [2000]: “If we are designing interventions, improvement on some measure is what we want to promote. For a system to exhibit adaptation that enhances survival (or another measure of success), it must increase the likelihood of effective strategies and reduce the likelihood of ineffective strategies. We call such a process attribution of credit if an agent uses a performance criterion to increase the frequency of successful strategies or to decrease the frequency of unsuccessful ones.” In a broad sense, our method can be regarded as a process of attribution of

credit, as follows: We detect many possible emergent states and use a performance criterion to select the desirable emergent states; efforts are then devoted to stabilize some of these emergent states; and finally, the system is then more likely to operate under effective strategies. This is apparent because if no intervention is introduced, then the system will visit the undesirable and desirable emergent states in an unpredictable way.

A key feature in the control of complex systems is that, quite counterintuitively, complexity can in fact facilitate the process of attribution of credit. Suppose we wish to change the performance criterion, which naturally results in a different set of desirable emergent states. Due to the existence of a large number of emergent states, it is relatively easy to determine the new set of desirable states and to apply small controls to steer the system to visit them more (this is equivalent to increasing the frequency of successful strategies). That is, in response to a change in the performance criterion, only a slight change in the system is necessary to achieve the attribution of credit. This is to be contrasted to the case of non-complex systems where a change in the performance criterion often requires a large change to the system for it to generate some new, desirable states.

## 7. CONTROL OF THE SYSTEM ABOUT DESIRABLE AND UNDESIRABLE EMERGENT STATES

Let us now consider the control strategies: stabilizing the system about a desirable emergent state or driving the system out of an undesirable emergent state. In the phase space, the dynamics of the Complex Adaptive System interplays among the large number of emergent states and chaotic sets on the basin boundary. Under the influence of noise, it can be expected that the system moves among different emergent states in time, in contrast to simple, low-dimensional nonlinear systems where the evolution is always towards one dominating attracting set. In other words, the dynamics of a Complex Adaptive System tend to alternate among different emergent states, each of which is sensitive to small perturbations. The key concept underlying our control strategy at the regional or global level is that the two attributes of the Complex Adaptive System in a dynamic environment, accessibility to many states and sensitivity to small perturbation, present us with an opportunity to manipulate the system’s dynamics.

Our control strategy of stabilizing the system about a desirable emergent state is then as follows. When the system comes to a desirable emergent state, we will apply small control perturbations to stabilize the system



about it to prolong the desirable emergent state as much as possible. Since control is applied only when the system falls in the vicinity of the desirable emergent state, the required magnitude of the perturbations can be made arbitrarily small. That is, the desirable system performance can be achieved at low cost. Suppose that, at a later time, the system moves out of the desirable emergent state due to the interplay among the emergent states, the chaotic sets on the basin boundaries and noise. We can then stop perturbing the system so that the system can evolve by itself ergodically, and wait until the system gets close to the new desirable state. This strategy offers a great deal of flexibility: Due to the accessibility of the system to many emergent states, any desirable state can easily be stabilized with little cost.

The feasibility of controlling complex dynamics using the above concept has been demonstrated successfully in both theory [Dressler and Nitsche, 1992; Romeiras et al., 1992] and experiments in numerous fields including physics, chemistry, and biology [Boccaletti et al., 2000; Garfinkel et al., 1992]. Here we draw on just one example: the successful control of a chaotically beating rabbit heart [Garfinkel et al., 1992]. A heart consists of a large number of interconnected cells, which is very similar to a Complex Adaptive System in the information infrastructure. A chaotically beating heart is certainly undesirable, just as a randomly behaving system in the information infrastructure is not as desirable as emergent states producing good traffic flow performance. Normally, the only accessible information about the heart is a time series, such as the heartbeat rate as a function of time. From such a time series, a mathematical phase space can be reconstructed by using the delay-coordinate embedding technique. It was demonstrated, experimentally, that by delivering small electrical pulses at times calculated from the dynamics of the heart in the reconstructed phase space, a chaotic heart can be converted into a periodically beating one [Garfinkel et al., 1992]. Because of the common feature, complexity, between a cardiac network and a Complex Adaptive System in the information infrastructure, it is hopeful that the latter may be harnessed in a similar way.

To calculate the required control perturbations, we can adopt the idea of controlling chaos [Ott, Grebogi, and Yorke, 1990a, 1990b; Lai and Grebogi, 1997a, 1997b; Grebogi, Lai, and Hayes, 1997; Boccaletti et al., 2000] for the control of Complex Adaptive Systems [Poon and Grebogi, 1995]. The key ingredient for the control of chaos [Ott, Grebogi, and Yorke, 1990a, 1990b] is the observation that a chaotic set, on which the trajectory of the chaotic process lives, has embedded within it a large number of unstable low-periodic

orbits. In addition, because of ergodicity, the trajectory visits or accesses the neighborhood of each one of these periodic orbits. Some of these periodic orbits may correspond to a desired system performance. The second ingredient is the realization that chaos, while signifying sensitive dependence on small changes to the current state and hence rendering the system unpredictable over a long period of time, also implies that the system's behavior can be altered by using small perturbations [Ott, Grebogi, and Yorke, 1990a, 1990b]. Then, the accessibility of the chaotic system to many different periodic orbits combined with its sensitivity to small perturbations allows for the control and the manipulation of the chaotic process. Specifically, the Ott-Grebogi-Yorke (OGY) approach is then as follows. One first determines some of the unstable low-periodic orbits that are embedded in the chaotic set. One then examines the location and the stability of these orbits and chooses one that yields a desirable system performance. Finally, one applies small control to stabilize this desired periodic orbit. However, all this can be done from data [Ott, Grebogi, and Yorke, 1990a, 1990b] by using nonlinear time series analysis for the observation, understanding and control of the system. This is particularly important since chaotic systems are rather complicated and detailed knowledge of the equations of the process is often unknown.

In our information infrastructure, the emergent states correspond to unstable periodic orbits and the wandering among these states corresponds to chaotic motion. Let us say that we have the following  $M$ -dimensional discrete-time dynamical systems:  $x_{n+1} = \mathbf{F}(x_n, p_n)$ , where  $\mathbf{F}$  is a smooth vector function, and  $p_n$  is an accessible parameter that can be externally perturbed. For example,  $p_n$  can be the number of open inbound and outbound queues of the key routers in a dynamic coalition at a given time. The agent for each router can open or close a queue at any time. Hence,  $p_n$  results from the individual agents' local actions of opening or closing queues, rather than a top-down command.

We conceive using only small controls, so we restrict  $p$  to lie in some small interval:  $|p_n - \bar{p}| < \delta$ , where  $\bar{p}$  is a nominal parameter value, e.g., the number of the open inbound and outbound queues of the routers in a dynamic coalition while in the desirable emergent state. If  $p_n$  is outside this interval, we set  $p_n = \bar{p}$ . Assuming that the dynamical system  $\mathbf{F}(x_n, \bar{p})$  generates behaviors with various emergent states, our goal is to vary the parameter  $p_n$  within the range  $(\bar{p} - \delta, \bar{p} + \delta)$  in such a way that for almost all initial conditions, the system dynamics converges onto a desirable emergent state. To do this, we consider a small neighborhood of size comparable to  $\delta$  of the desirable emergent state. In this neighborhood, the dynamics is approximately linear.

Since linear systems can be stabilized based on the controllability law, it is reasonable to assume that the desirable emergent state can be stabilized by feedback control. The ergodic nature of the complex dynamics guarantees that the state trajectory enters the neighborhood. Once inside, we apply the stabilizing feedback control law [Ogunnaike and Ray, 1994] to keep the trajectory in the neighborhood of the desirable emergent state.

A key feature of the above control strategy is that we actually make use of a complex system's sensitivity to small changes to stabilize it around a desirable state by using only small perturbations. While feedback control is applied in a small linear neighborhood of the desirable state, the control perturbations are small and time-dependent. Insofar as the control is present, the system can be stabilized around the desirable state. When small noise is present, the required strength of the perturbation should be greater than the noise amplitude to ensure control. In the rare event when the noise becomes large (as for Gaussian noise), it is possible that control will be lost. In this case, when the system leaves the linear neighborhood of the desirable state, we turn off the control. Because of ergodicity, at a later time the system will enter the linear neighborhood again, reactivating control. In this sense, linear feedback control law, which works ideally for linear system, can be robust for nonlinear systems as well. It is possible that nonlinear control law can be derived for complex systems, which, however, may require perturbations of large magnitude. Our strategy emphasizes on taking the advantage of a complex system's sensitivity to small changes to achieve control by small perturbations, for which a linear control law appears to suffice. This is usually not possible even for linear systems, where large control is required to bring the system to a desirable state.

The OGY idea of controlling chaos has been demonstrated to be capable of stabilizing a desired random trajectory, which has potential applications to problems such as synchronization of chaotic systems [Lai and Grebogi, 1993], conversion of transient chaos into sustained chaos [Lai and Grebogi, 1994], communication with chaos [Hayes, Grebogi, and Ott, 1993; Bollt, Lai, and Grebogi, 1997] and selection of a desired chaotic phase [Nagai and Lai, 1995].

Let us consider the control strategy of driving the system out of an undesirable emergent state. A crisis is an undesirable emergent state characterized by a high degree of synchronization among agents in the network. For example, in a transportation network near the scene of a major traffic accident, all vehicles are driving at approximately the same speed. It is thus important to design control that will desynchronize the dynamical system to drive the system out of the undesirable emer-

gent state. A physically observable state of synchronization implies the existence of (1) an underlying mathematical object, the so-called synchronization manifold which is stable with respect to perturbations that are transverse to it, and (2) dynamically invariant sets in the manifold [Nagai and Lai, 1997; Lai and Grebogi, 1999]. Embedded in the invariant sets are periodic orbits that can be either transversely stable or unstable. The synchronization state being stable means that the majority of these periodic orbits are transversely stable. Generally, however, there is still a subset of periodic orbits that are transversely unstable, even when the manifold itself is transversely stable. Nagai and Lai [1997] investigated how the dynamical interplay between these two distinct groups of periodic orbits lead to synchronization in a network of coupled complex elements that exhibit chaotic behaviors by calculating their weighted transverse stability coefficients. The existence of this particular set of periodic orbits provides an opportunity to draw the system out of the synchronization state. In particular, we envision that a number of these transversely unstable periodic orbits can be observed and recorded from analyzing the time series using the delay-coordinate embedding technique. When the time-evolved state of the system falls in the neighborhood of one of these "desynchronized" periodic orbits, small controls are applied to stabilize the system state about it so that the system is no longer in a synchronized state. The idea of control is thus similar to that utilized by OGY to stabilize a desirable unstable periodic orbit. To reduce the time required for the system state to fall into a desynchronized orbit, one can induce perturbations that break the symmetry, which leads to the existence of the synchronization manifold in the first place. Such symmetry-breaking perturbations could be, for instance, a serious mismatch between agents coupled in the system or external forces applied at a selected set of locations in the system.

For example, when a crisis of aggregate traffic flows occurs in a dynamic coalition, all flows in the dynamic coalition may move at the same pace. The individual agents of the key routers in the dynamic coalition can then reset the on-off switches of inbound and outbound queues in those routers so that traffic flows in and out of important host machines will move faster, whereas traffic flows in and out of unimportant host machines will move slower. The individual agents of the routers can also decrease the speed with which they forward messages to a congested area.

Overall, the control strategy at the regional and global levels aims to let individual agents continue what they are doing when a regional or global system gets into a desirable emergent state as a result of the self-synchronization of its agents in a favorable environment

in order to prolong self-synchronization, and to let individual agents change what they are doing when the system gets into an undesirable emergent state as a result of self-synchronization in an unfavorable environment in order to break the self-synchronization or desynchronize.

We wish to remark that our control strategy is suitable for stabilizing relatively local, suboptimal states. There is no guarantee that it can yield a globally optimized state. The main reason is that a complex system is characterized by the presence of many emergent states at different spatial and time scales. It is not clear whether a complex system possesses any globally optimal state. Even if it does, to our knowledge, at the present there appears to be no method to determine or control a global state that is optimal under certain performance criterion. Our strategy is based on the idea of “using complexity to harness complexity,” which is similar to the idea of controlling chaos [Ott et al., 1990a]. Nevertheless, we are fully aware of the fact that our approach will in general not optimize, but rather improve, network performance. The control of suboptimal states, as we have described, can be an effective way to cope with complex systems. Our hope is that investigating and implementing the control approach proposed in this paper may help generate better ideas for global control and optimization of complex systems.

## 8. SUMMARY

This paper provides a theoretical discussion on performance goals and control strategies at the regional and global levels of an information infrastructure based on Complexity Theory for Complex Adaptive Systems. We outline the conceptual design of the bottom-up self-synchronization approach to QoS assurance and stateful resource management for a dependable information infrastructure. This approach has the potential to overcome problems with existing information infrastructures, for example, the top-down centralized resource management with Computational Grids and the stateless resource management with the Internet.

## ACKNOWLEDGMENT

This work is sponsored by the Department of Defense (DoD) and the Air Force Office of Scientific Research (AFOSR) under Grant Number F49620-01-1-0317. The U.S. government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily

representing the official policies or endorsements, either express or implied, of, DoD, AFOSR, or the U.S. Government.


## REFERENCES

- R. Axelrod and M.D. Cohen, *Harnessing complexity: Organizational implications of a scientific frontier*, Free Press, New York, 2000 (paperback edition: Basic Books, New York, 2001).
- R. Badii and A. Politi, *Complexity—Hierarchical structures and scaling in physics*, Cambridge University Press, Cambridge, 1997.
- S. Boccaletti, C. Grebogi, Y.-C. Lai, H. Mancini, and D. Maza, The control of chaos: Theory and applications, *Phys Rep* 329 (2000), 103–197.
- E.M. Bollt, Y.-C. Lai, and C. Grebogi, Coding, channel capacity, and noise resistance in communicating with chaos, *Phys Rev Lett* 79 (1997), 3787–3790.
- S.J. Chatterjee, B.S. Sydir, M. Davis, and T.F. Lawrence, Modeling applications for adaptive QoS-based resource management, *Proc IEEE High Assurance Systems Engineering Workshop*, 1997, pp. 194–201.
- M. Coombs and D. Birx, Czerwinski’s order regimes: Where should we begin to address the IO risk? *Proc IEEE SMC Information Assurance and Security Workshop*, 2000, pp. 202–205.
- M. Ding, E. Ott, and C. Grebogi, Crisis control: Preventing chaos-induced capsizing of a ship, *Phys Rev E* 50 (1994), 4228–4230.
- U. Dressler and G. Nitsche, Controlling chaos using time delay coordinates, *Phys Rev Lett* 68 (1992), 1–4.
- S.M. Emran and N. Ye, Robustness of chi-square and Canberra techniques in detecting intrusions into information systems, *Quality and Reliability Engineering International*, 18 (2002), 19–28.
- I. Foster and C. Kesselman, *The grid: Blueprint for a new computing infrastructure*, Morgan Kaufmann, San Francisco, 1999.
- A. Garfinkel, M.L. Spano, W.L. Ditto, and J.N. Weiss, Controlling cardiac chaos, *Science* 257 (1992), 1230–1233.
- N. Giroux and S. Ganti, *Quality of service in ATM networks*, Prentice Hall, Upper Saddle River, NJ, 1999.
- P. Grassberger and I. Procaccia, Measuring the strangeness of strange attractors, *Phys D* 9 (1983), 189–208.
- C. Grebogi, Y.-C. Lai, and S. Hayes, Control and applications of chaos, *Internat J Bifur Chaos Appl Sci Engrg* 7 (1997), 2175–2198.
- S. Hayes, C. Grebogi, and E. Ott, Communicating with chaos, *Phys Rev Lett* 70 (1993), 3031–3034.
- Y.-C. Lai, Unstable dimension variability and complexity in chaotic systems, *Phys Rev E* 59 (1999), R3807–R3810.
- Y.-C. Lai and C. Grebogi, Synchronization of chaotic trajectories using control, *Phys Rev E* 47 (1993), 2357–2360.
- Y.-C. Lai and C. Grebogi, Converting transient chaos into sustained chaos by feedback control, *Phys Rev E* 46 (1994), 1094–1098.

- Y.-C. Lai and C. Grebogi, Complexity in Hamiltonian-driven dissipative chaotic dynamical systems, *Phys Rev E* 54 (1996), 4667–4676.
- Y.-C. Lai and C. Grebogi, Controlling chaotic dynamical systems, *Systems Control Lett* 31 (1997a), 307–312.
- Y.-C. Lai and C. Grebogi, Controlling chaos in high dimensions, *IEEE Trans Circuits Systems* 44 (1997b), 971–975.
- Y.-C. Lai and C. Grebogi, Modeling of coupled chaotic oscillators, *Phys Rev Lett* 82 (1999), 4803–4806.
- Y.-C. Lai and D. Lerner, Effective scaling regime for computing the correlation dimension in chaotic time series analysis, *Phys D* 115 (1998), 1–18.
- Y.-C. Lai, D. Lerner, and R. Hayden, An upper bound for the proper delay time in chaotic time series analysis, *Phys Lett A* 218 (1996), 30–34.
- T.F. Lawrence, The quality of service model and high assurance, *Proc IEEE High Assurance Systems Engineering Workshop*, 1997, pp. 38–39.
- R. Lewin, *Complexity: Life at the edge of chaos*, University of Chicago Press, Chicago, 1999.
- S. Liu, Y. Xiong, and P. Sun, On defense of the distributed denial of service attacks: An on-off control approach, *Proc 2000 IEEE Workshop on Information Assurance and Security*, 2000, pp. 282–293.
- Y. Nagai and Y.-C. Lai, Periodic-orbit theory of the blowout bifurcation, *Phys Rev E* 56 (1997), 4031–4041.
- Y. Nagai and Y.-C. Lai, Selection of desirable chaotic phase using small feedback control, *Phys Rev E* 51 (1995), 3842–3848.
- National Science Foundation, New grid portal to improve U.S. researchers' access to advanced computing resources, NSF, Washington, DC, 2000, <http://www.nsf.gov/cgi-bin/getpub?pr0088>.
- B.A. Ogunnaike and W.H. Ray, *Process dynamics, modeling and control*, Oxford University Press, New York, 1994.
- E. Ott, (2002). *Chaos in Dynamical Systems*. Cambridge: Cambridge University Press (2nd edition).
- E. Ott, C. Grebogi, and J.A. Yorke, Controlling chaos, *Phys Rev Lett* 64 (1990a), 1196–1199.
- E. Ott, C. Grebogi, and J.A. Yorke, "Controlling chaotic dynamical systems," *CHAOS/XAOC*, Soviet-American perspective on nonlinear science, D. Campbell (Editor), AIP Press (Springer-Verlag), New York, 1990b, pp. 153–172.
- L. Poon and C. Grebogi, Controlling complexity, *Phys Rev Lett* 75 (1995), 4023–4026.
- F.J. Romeiras, C. Grebogi, E. Ott, and W.P. Dayawansa, Controlling chaotic dynamical systems, *Phys D* 58 (1992), 165–192.
- B. Sabata, S. Chatterjee, M. Davis, J. Sydir, and T.F. Lawrence, Taxonomy for QoS specification, *Proc IEEE Comput Soc 3rd Int Workshop Object-Oriented Real-Time Dependable Systems*, 1997, pp. 100–107.
- T. Sauer, J.A. Yorke, and M. Casdagli, *Embedology*, *J Stat Phys* 65 (1991), 579–616.
- F. Takens, "Detecting strange attractors in turbulence," *Dynamical systems and turbulence*, D.A. Rand and L.-S. Young (Editors), *Lecture Notes in Mathematics* 898, Springer-Verlag, New York, pp. 366–381.
- J. Theiler, Spurious dimension from correlation algorithms applied to limited time series data, *Phys Rev A* 34 (1986), 2427–2432.
- N. Ye, C. Borrer, and Y. Zhang, EWMA techniques for computer intrusion detection through anomalous changes in event intensity, *Quality and Reliability Engineering International*, 18 (2002), 443–451.
- N. Ye and Q. Chen, Computer intrusion detection through EWMA for auto-correlated and uncorrelated data, *IEEE Transactions on Reliability* 52 (2003), 73–82.
- N. Ye, S.M. Emran, Q. Chen, and S. Vilbert, Multivariate statistical analysis of audit trails for host-based intrusion detection, *IEEE Transactions on Computers* 51 (2002), 810–820.



Nong Ye is a Professor of Industrial Engineering and an Affiliated Professor of Computer Science and Engineering at Arizona State University. She holds a Ph.D. in Industrial Engineering from Purdue University, an M.S. in Computer Science from the Chinese Academy of Sciences, and a B.S. in Computer Science from Peking University. Her research interests are in assuring process quality and preventing faults and errors in information systems. Dr. Ye is an Associate Editor for *IEEE Transactions on Reliability* and *IEEE Transactions on Systems, Man, and Cybernetics*. She is a senior member of the Institute of Industrial Engineers and a senior member of IEEE.



Ying-Cheng Lai is a Professor of Mathematics and of Electrical Engineering at Arizona State University. His honors include an Air Force Presidential Early Career Award for Scientists and Engineers (PECASE) in 1997, an NSF Faculty Career Award in 1997, and election as a Fellow of the American Physical Society in 1999. His current research interests are applied chaotic dynamics, complex networks, signals processing, and computational biology.



Toni Farley is the Assistant Director of the Information and Systems Assurance Laboratory, and a Ph.D. student of Computer Science at Arizona State University (ASU), Tempe, Arizona. She has been awarded a graduate fellowship from the AT&T Labs Fellowship Program (ALFP). She received her B.S. in Computer Science from ASU in 2003. Her research interests include computer networks, network security, and information assurance. She is a member of IEEE and IEEE Computer Society.