Islands of Near-Perfect Self-Prediction

Stephen F. Bush, bushsf@crd.ge.com
General Electric Corporate Research and Development
One Research Circle, Niskayuna, NY 12309, USA

KEYWORDS:

Self-Prediction, Active Networks, Network Management, Brittle Systems

ABSTRACT

In the course of efforts to more fully utilize the power of active networks to build a self-managing communications network, the nature of entanglement and the relationship between modeling and communication become of utmost importance. This paper provides a very brief introduction to Active Networks and the Active Virtual Network Management Prediction Project whose goal is a self-managing communications network. The focus of the paper is upon the effects of near-infinite resources; that is, how will such a self-predictive system behave as processing and bandwidth become ever larger and more powerful. An attempt is made to identify new theories required to understand such highly self-predictive systems.

1 ACTIVE NETWORKS AND SELF-MANAGEMENT

The problem this paper addresses is the complexity of managing large and rapidly growing communication networks with new more powerful technology. Network management consists of a wide variety of responsibilities including configuration management, performance management, fault management, accounting management, and security management. A network management system must be able to monitor, control, and report upon the status of all of these areas. Today, this is usually performed using a standards based management protocol such as the Common Management Information Protocol **(CMIP)** or the Simple Network Management Protocol **(SNMP 1991)**. A goal of network management is to pro-actively detect problems in each of these areas. This means detecting potiential events such as performance problems and faults before they occur. This is a primary goal of Active Virtual Network Management Prediction (AVNMP) **(Bush 1999)**.

Active networks **(Tennenhouse 1997)** are a relatively recent concept in communication networks. Active networks are capable of executing general purpose code within packets as the packets are transmitted through intermediate network nodes. A framework for supporting the execution of general purpose code within packets as they travel through a network is an on-going research effort. Thus active networks differ from today's communications networks because active networks offer a computational service in addition to a data transport service. In current communication networks non-executable data is passively forwarded through the traditional communication layers; intermediate devices such as bridges and routers only access the data link or network headers of packets. In active networks, intermediate devices can execute generic code within active packets as they travel through the network. The ability for communication networks to perform such computation offers opportunities for great advantages in such areas as efficiency, rapid protocol development and deployment, and network flexibility. However, active networks also add additional complexity, particularly in network management and security. The goal of Active Virtual Network Management Prediction is to use the advantages active networks provide in order to handle the additional complexity in network management.

A detailed discussion of AVNMP is outside the scope of this paper, however, AVNMP's relevant characteristics are that it attempts to become a closed predictive system by injecting a simple model of the open ends of the network back into the network as shown in Figure 4. Another characteristic is that AVNMP attempts to utilize the inherent parallel nature of communications networks in an optimal manner, by using a Time Warp like mechanism to ensure causality among virtual messages within the system. The network uses information thus generated about its likely future state to improve its current performance.

2 NEAR-INFINITE RESOURCES

Now, imagine stepping across a discontinuity into a world where computing power, bandwidth, and computational ubiquity are nearly infinite. Our vision focuses on effects that near-perfect self-prediction would have upon such a world. First we would have near-perfect optimization of resources since local minima could be pushed far into the horizon. Second, currently wasted effort could be avoided, since the outcome of any action could be determined with very precise limits. Critical missing elements are a theory and applications involving highly predictive systems and components. Study is needed to explore the exciting new world of near-perfect self-prediction and the relationship between highly predictive systems and communications in particular. Figure 1 shows an abstract view of computers embedded within almost all devices. Current engineering organizes computing devices in such a way as to optimize communications performance. In our hypothetical world of near-perfect predictive capabilities, direct communication is less important and, in many cases, no longer required, as discussed later. Instead, computational organization is based on forming systems or islands of near-perfect self-prediction. As shown in Figure 2, self-predictive capability is used to enhance the performance of the system, which in turn improves the predictive capability, which again improves the performance of the system, ad infinitum, driving the error towards zero.

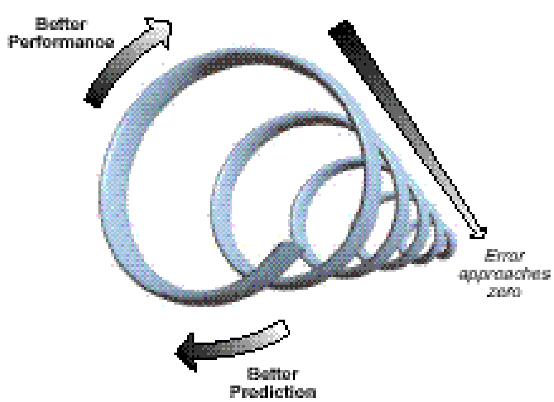

Figure 2: This predictive capability is used to drive the error toward zero.

Why do we assume near-perfect prediction rather than perfect prediction and why do we assume islands rather than perfect prediction everywhere? Clearly, perfect prediction everywhere would take us into a deterministic world where the final outcome of all choices would be known to everyone and the optimal choice could be determined in all cases. In this project it is assumed that limits, however small, exist, such as lack of knowledge about quantum state or of the depths of space. In order to study near-perfect self-predictive islands, the characteristics of such islands need to be identified. It would appear that closed self-predictive islands would be the easiest to understand. The scope of closed self-predictive islands includes all driving forces acting upon the system. Imagine that one has full knowledge of the state of a room full of ping-pong balls and their elasticity. This information can be used to predict the position of the balls at any point in time. However, one is external to the room. The goal is for the balls to predict their own behavior as illustrated in the inner sphere of Figure 3. If elasticity represents the dynamics of communication endpoint entities A and B, and movement of the ping-pong balls represents communication, then any exchange of information between A and B is unnecessary since it can be perfectly predicted. Instead of transmitting messages between A and B, an initial transmission of the dynamics of A and B are transmitted to each other, perhaps as active packets within an active network environment. Thus a near-perfect self-predictive island is turned inward upon itself as shown in Figure 4. In an active network environment, an executable model can be included within an active packet. When the active packet reaches the target intermediate device, the load model provides virtual input messages to the logical process and the payload of the virtual message is passed to the actual device, as described in the Active Virtual Network Management Protocol (Active Networks) project [DARPA ITO F30602-98-C-0230] **(Bush 1999)**. A streptichron is an active packet facilitating prediction in our self-adjusting Time Warp System. GE-CRD is currently experimenting with streptichrons in an active network environment and this proposal takes that idea as close as possible to the extreme limit.

Open self-predictive islands will contain inaccuracies in prediction because, by definition, open self-predictive islands include the effects of unknown driving forces upon entities within the scope of the system. Figure 3 shows a force (F_1) acting on the inner system. F_1 is external to the inner system because it is not included within the system itself or in the virtual messages passed into the system. The system could become closed by either enlarging the scope to include the driving forces within the system, as shown in the figure, or by accepting a level of inaccuracy in the system. Thus we can imagine many initial points of near-perfect self-predictive islands, each attempting to improve prediction fidelity by expanding to incorporate more elements. These are the islands of near-perfect self-prediction.

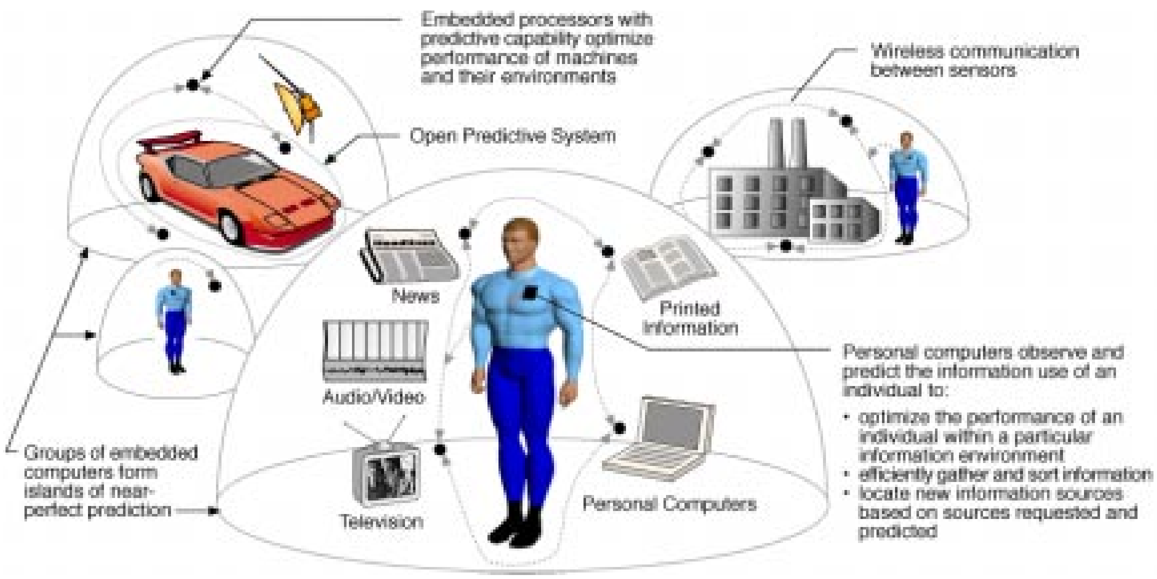

Figure 1: Computational organization is based on forming systems or islands of near-perfect self-prediction.

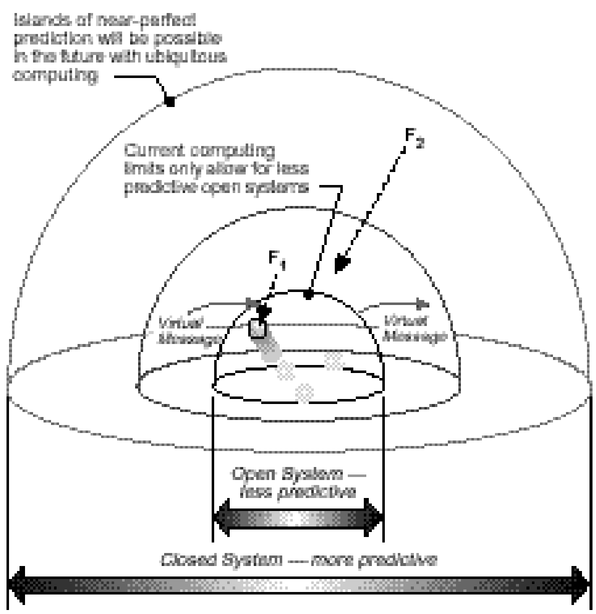

Figure 3: Self-predictive islands can improve prediction fidelity by expanding to incorporate more elements.

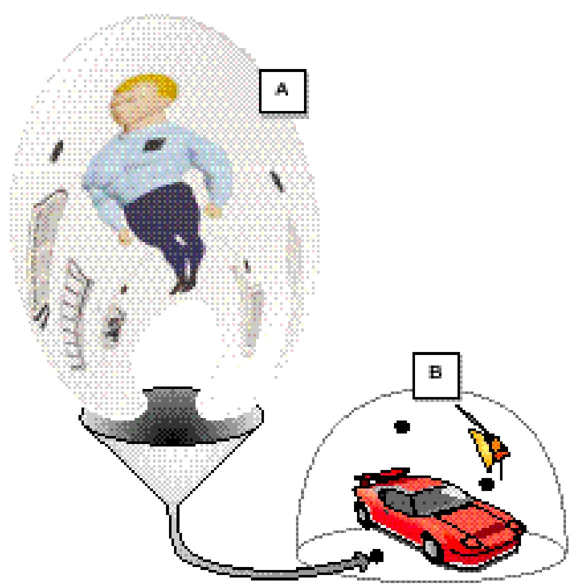

Figure 4: Direct communication between A and B is unnecessary as the dynamics of A can be transmitted to B, allowing B to interact with a near-perfect model of A.

Recursion is a recurring theme in this work. For example, assume that the inner near-perfect self-predictive island in Figure 3 is a wireless mobile communications system and F_1 is the weather. Now assume that ubiquitous computing can be used to include weather observation and prediction, for example, computers within planes, cars, space craft, etc. The heat from the circuitry of the wireless system, even though negligible, could have an impact on the weather. This is known as the butterfly effect in Chaos Theory. In recent years the study of chaotic nonlinear dynamical systems has lead to diverse applications where chaotic motions are described and controlled into some desirable motion. Chaotic systems are sensitive to initial condition. Researchers now realize that this sensitivity can also facilitate control of system motion. For example, in communications, chaotic lasers have been controlled, as have chaotic diode resonator circuits **(Aronson, 1994)** **(DiBernardo, 1996)**. Hence, studying the effects of external forces controlling a chaotic system has become a very important goal and should be a subject for research.

A fascinating perspective on the topic of near-perfect self-predictive islands is found in Reference **(Hofstadter, 1980)**, which is a study of the nature of human and artificial intelligence. A central point in this study is that intelligence is a tangled hierarchy. When two hierarchical levels are folded together, a tangled hierarchy results. Near-perfect prediction as presented in this proposed work is a tangled hierarchy on several levels: simulation-reality and also present-future time, each modifying the other in a near-perfect prediction environment. By allowing for a given tolerance in the amount of error and assuming accuracy in prediction that increases as real time approaches the actual time of an event, this study assumes that a useful near-perfect self-predictive island can be implemented. In the Active Virtual Network Management project, an attempt is being made to embed predictive capability within an active network using a self-adjusting Time Warp based mechanism for prediction propagation. This self-adjusting property has been found to be useful in prediction and is referred to as autoanaplasis. In addition to autoanaplasis, it is well known that such systems can exhibit super-criticality, faster than critical path execution. However, due to limited and non-ubiquitous computational power in current technology, prediction inaccuracy causes rollbacks to occur. In a world of near-infinite bandwidth and computing power, the cost of a rollback to a “safe” time becomes infinitesimal. This is one of the many new ideas this project will explore involving the relationship between bandwidth, computing power, and prediction. Given near-infinite bandwidth, the system state can be propagated nearly instantaneously. With nearly infinite and ubiquitous computing, driving processes can be developed with near-perfect accuracy. Let us define near-perfect accuracy of our self-adjusting Time Warp based system in the presence of rollback as the characteristic that a predicted state value (V_v) approaches the real value (V_r) as t approaches $GVT(t_1)$ very quickly, where $GVT(t_1)$ is the Global Virtual Time of the system at time t_1. Explicitly, this is, $\forall \epsilon > 0, \exists \delta > 0$ s.t. $|f(t) - f(GVT(t_1))| < \epsilon \rightarrow 0 < |GVT(t_1) - t|$ where $f(t) = V_r$ and $f(GVT(t_1)) = V_v$. $f(t)$ is the prediction function. The effect of extreme parameter values should not be ignored.

3 PERFORMANCE OF NEAR-PERFECT SELF-PREDICTIVE ISLANDS

One focus of study could be on the interfaces between systems with various levels of predictive capability. The self-predictive islands formed in Figure 1 will have various degrees of prediction capability. Our recent theoretical results from the Active Virtual Network Management project indicates that self-predictive islands can exhibit high degrees of performance when prediction is accurate, but are brittle when the tolerance for inaccuracy is reached. With respect to network performance as enhanced with AVNMP, systems with little or no prediction capability appear to be ductile, as they are much better able to tolerate prediction inaccuracy, as shown in Figure 6. In other words, performance is moderate, but there are no sudden degradations in performance. Compare this to a system with a large lookahead and sudden, near catastrophic degradations in performance.

Thus, an obvious question arises as to what is the optimal grouping of predictive components within a system. What happens when the slope shown in Figure 6 becomes nearly vertical? The lookahead into the future will be tremendously large in some self-predictive islands and small in others. If the lookahead is small in a self-predictive island that feeds into a large lookahead system, then large rollbacks are likely to occur. One focus of study should be on the interfaces between systems with various levels of predictive capability and the associated "index of refraction" of performance through the interfaces between islands of near-perfect self-prediction.

Near-perfect self-predictive islands' brittle behavior is shown by point D along curve P_h in Figure 7. P_h is the performance curve for a high-performance system with brittle characteristics, P_l is a lower-performance system with ductile characteristics. Clearly, the slope from point D along curve P_h is much steeper than that of point E along curve P_l. The steep decline of performance along P_h can be caused by input parameters that exceed a specified tolerance, or by environmental conditions that exceed specified operating boundaries.

Materials Science	*Near-Perfect Prediction Systems*
Brittle Behavior	Sudden steep decline in performance
Ductile Behavior	Graceful degradation in performance
Stress	Amount parameter exceeds its tolerance
Toughness	System robustness
Hardness	Level of performance within tolerance
Ductility	Level of performance outside of tolerance
Plastic Strain	Degradation from which system cannot recover
Elastic Strain	Degradation from which system can recover
Brittleness	Ratio of hardness over ductility
Deformation	Degradation in performance
Young's Modulus	Amount tolerance is exceeded over degradation

Figure 5: Near-perfect predictive systems' characteristics and behavior can be described in terms borrowed from materials science.

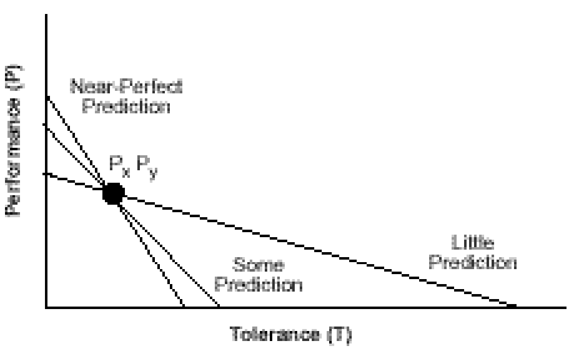

Figure 6: Performance of Self-predictive islands.

4 BRITTLE ISLANDS OF NEAR-PERFECT SELF-PREDICTION

Consider a system whose self-predictive islands exhibit various degrees of ductility as defined above. Just as adding impurities to a pure metal causes it to become stronger but more brittle, the addition of more efficient

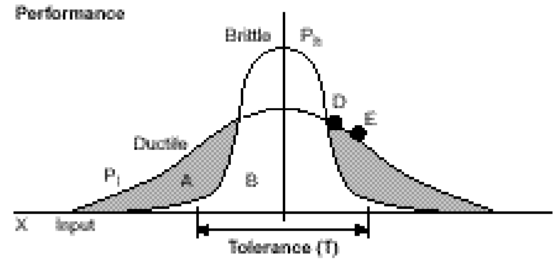

Figure 7: A brittle vs. ductile system.

but also more sensitive components to a system such as near-perfect self-prediction systems, causes the system to increase performance within its operating range, but become less ductile. How do the effects of ductility propagate among the self-predictive islands to influence the ductility of the entire system? Assume the performance response curve is known for each self-predictive island and that the output from one component feeds into the input of the next component as shown in Figure 8. The self-predictive islands are labeled S_n and the performance curves as a function of tolerance for error are shown above each island. We wish to carry forward this analogy and deliver a theory and models of the relationships among computing, communications, and near-perfect self-prediction.

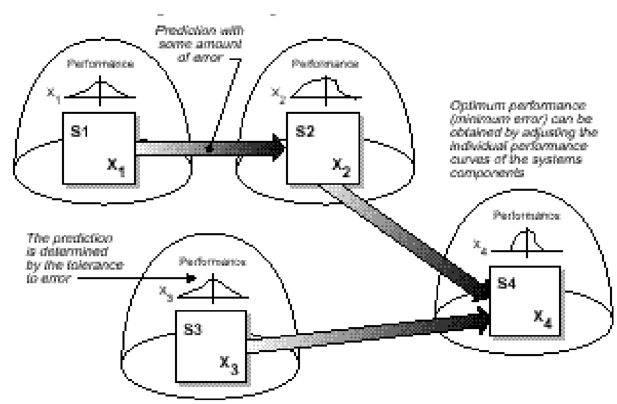

Figure 8: Brittle subsystem components.

5 CONCLUSIONS

The high level conclusion for this paper is that more research is required to understand the nature of entanglement, causality, and the relationship between modeling and communications. For example, Active Network Management Prediction (AVNMP) uses a model within a network to enhance the network performance to enhance the model's own performance which thus enhances the network's performance thus enhancing the model's performance ad infinitum as shown in Figure 2. Futhermore, the Active Virtual Network Management Prediction mechanism uses a Time Warp-like method to ensure causality, yet there is something non-causal about the way AVNMP uses future events to optimize current behavior. This entanglement issue resonated with physicists and those studying the nature of agent autonomy at the conference. Clearly, this needs to be explored in a much deeper manner. Also, formation of islands of near-perfect self-prediction and the need to study the interfaces between those islands was discussed. The idea of wrappers and integration spaces as introduced in **(Landauer 1996a)** is likely to provide insight into bringing together complex system components in a self-organizing manner. Another suggestion for the study of predictive interfaces is in a tolerance interaction space **(Landauer 1996b)**.

6 REFERENCES

(Aronson, 1994) I. Aronson, J. Levine and L. Tsimring; "Controlling Spatio-Temporal Chaos"; "Phys. Rev. Lett." Vol. 72, pp 2561-2564; 1994.

(Bush 1999) Stephen F. Bush, "Active Virtual Network Management Prediction", Parallel and Discrete Event Simulation (PADS) Conference, 1999, http://www.crd.ge.com/people/bush/an.

(Bush 2000) Stephen F. Bush, "Islands of Near-Perfect Self-Prediction", Virtual Worlds Conference, 2000, Slides in http://www.crd.ge.com/people/bush/an.

(CMIP) Open Systems Interconnection - Management Protocol Specification - Part 2: Common Management Information Protocol.

(DiBernardo, 1996) M. DiBernardo, "An Adaptive Approach to the Control and Synchronization of Continuous-Time Chaotic Systems"; "Int. J. of Bifurcation and Chaos" Vol. 6, No. 3, pp 557-568; World Scientific, 1996.

(Hofstadter, 1980) Douglas R. Hofstadter, "Godel, Escher, Bach: An Eternal Golden Braid", Vintage Books, 1980, ISBN 0-394-74502-7.

(Landauer 1996a) Christopher Landauer, Kirstie L. Bellman, "Semiotics of Constructed Complex Systems", pp. 35-40 in Alex Meystel, Jim Albus, R. Quintero (eds.), Intelligent Systems: A Semiotic Perspective, Proceedings of the 1996 International Multidisciplinary Conference, Volume I: Theoretical Semiotics, Workshop on Intelligence in Constructed Complex Systems, 20-23 October 1996, NIST, Gaithersburg, Maryland (1996).

(Landauer 1996b) Christopher Landauer, Kirstie L. Bellman, "Integration Systems and Interaction Spaces", pp. 161-178 in Proceedings of FroCoS'96: The First International Workshop on Frontiers of Combining Systems, 26-29 March 1996, Munich (March 1996).

(SNMP, 1991) M. T. Rose, "The Simple Book, An Introduction to the Management of TCP/IP Based Internets", Prentice Hall, 1991.

(Tennenhouse, 1997) D. L. Tennenhouse, J. M. Smith, W. D. Sincoskie, D. J. Wetherall, and G. J. Minden, "A Survey of Active Network Research", IEEE Communications Magazine 35(1):80-86, Jan. 1997.