Genetically Modified Software: Realizing Viable Autonomic Agency

A.G. Laws, A. Taleb-Bendiab & S.J. Wade

School of Computing & Mathematical Sciences
Liverpool John Moores University
United Kingdom, L3 3AF
`{A.Laws, A.Talebbendiab, S.J.Wade}@livjm.ac.uk`

Abstract: Inspired by the autonomic aspects of the human central nervous system, the vision of "*autonomic computing*" arrived with a fully-formed wish list of characteristics that such systems should exhibit, essentially those self-referential aspects required for effective self-management. Although much progress has been made, a unifying approach or indeed an underlying theoretical foundation to support such work has not as yet emerged. Here, the authors contend that the biologically-inspired managerial cybernetics of Beer's Viable System Model (VSM) provides significant conceptual guidance for the development of a general architecture for the operation and management of such complex, evolving, adaptive systems that arguably extends the concept of autonomic systems to cognitive systems. Consequently, the VSM has been used as the basis of a unifying reference model that provides the "blueprint" for an extensible intelligent agent architecture that readily scales to a polyarchical agency of autonomic systems. Furthermore and with recourse to the classical cybernetics that underpin the VSM, the authors demonstrate that survival in a changing environment requires that such systems should develop capabilities to identify, specify, develop and deploy appropriately, a repertoire of actions that would guide their adaptation to changing circumstances/environments. The authors then show that the use of Holland's Genetic Algorithms (GA's) can, in specific circumstances, provide a means to provide tailored responses to environmental change and when coupled with the associated Learning Classifier Systems (LCS) approach allow the system to develop an adaptive environmental model of appropriate, optimized responses. Of course, although these approaches appear to offer a profitable route forward, the general applicability and scalability of GAs in particular, are open to question, therefore the paper concludes with some speculation on the possible contributions that associated approaches like genetic programming may have to offer.

Introduction

Classical cybernetics, rooted in the study of adaptation in a changing environment and the later development of managerial cybernetics, underpinned by classical cybernetics and structured around the human central nervous system have long been used to study the viability of human systems. More recently and with the advent of the possibility of "*autonomic computing systems"* realizable by *adaptive software agents*, the need for a unifying theory to inform and guide the development of such systems has become apparent. It would appear that cybernetic studies have much to contribute in terms of a well-established foundation for the development of such systems. In this paper, the authors attempt to show the value of assuming such a position by demonstrating the contribution that a cybernetic viewpoint can bring to the development of such systems.

The remainder of this position paper is organized as follows. The next section provides a brief introduction to the managerial cybernetics of Beer's Viable System Model (VSM) (Beer 1981). This is followed by the development of a robust and theoretically defensible adaptive software architecture realizable by means of multi-agent software paradigm. The next section reviews the classical cybernetics that underpin the VSM to demonstrate the appropriateness of the use of genetically inspired approaches to provide our target system with elements of "creativity" and "learning". Such approaches are considered in the penultimate section of the paper, which then concludes by both acknowledging the inherent limitations of such approaches and speculating on the further research opportunities these nevertheless afford.

The Viable System Model – An Architectural Model for Self-Managing Systems

The Viable System Model (Beer 1981), founded on classical cybernetic studies and modelled on the human central nervous system, provides a theoretically supported cybernetic model of organization.

Viable systems may be defined as being robust against internal malfunction and external disturbances and have the ability to continually respond and adapt to unexpected stimuli. The model specifically attempts to imbue the system with the ability to adapt to circumstances not foreseen by the original designer and identifies the *necessary* and *sufficient* communication and control systems that must exist for any organization to remain *viable* in a changing environment. The major systems (i.e. S1s, S3, S4 and S5) are structured hierarchically and connected by a central ‘spine’ of communication channels passing from the higher-level systems through each of the S1 management elements (Figure 1). These provide high priority communication facilities to determine resource requirements, accounting for allocated resources, alerts indicating that a particular plan is failing and re-planning is necessary and the provision of the "legal and corporate requirements" or policies of the system.

The systems shown in Figure 1 concern the management structure at one level of the system, and consequently specifies the communication and control structures that must exist to manage a set of S1 units. However, the power of the model derives from its recursive nature. Each S1, consisting of an operational element and it's management unit, is expected to develop a similar VSM structure, consequently, the structure of systems is open ended in both directions and may be pursued either upwards to ever wider encompassing systems or downwards to ever smaller units. However, at each level the same structure of systems would occur although their detail would necessarily differ depending on context.

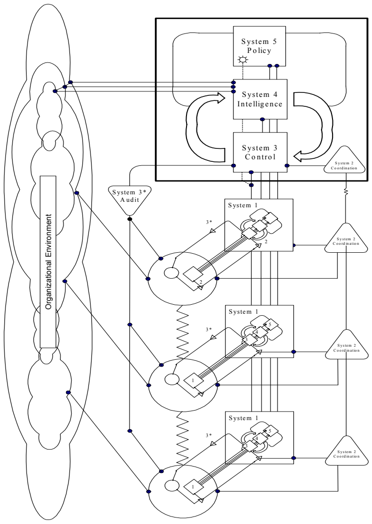

Figure 1. The Viable System Model.

The value of assuming such a viewpoint is in the immediate provision not only of the outline architecture that the autonomic software system itself must assume, namely that of the Viable System Model, but also the identification of the requisite communication links to bind the system to the organization.

An Intelligent Agent Architecture

We now extend and apply this cybernetic approach and consider an S1 of the VSM in terms of an autonomic software system. To demonstrate, a conceptual, architectural outline of such a system is determined, using both the principles of the VSM and the terminology and design of a classical Artificial Intelligence design, namely Bratman *et al.*'s Intelligent Resource-Bounded Machine Architecture (IRMA) (Bratman, Israel et al. 1988) as a constructional guide. As shown in Figure 2, the developed VAAA/J-Reference architecture embeds a Beliefs, Desires, Intentions (BDI) unit at the S5 level representing;

- Desires - or what the agent wants to do and is taken as a given for the moment.
- Beliefs - or what the system currently knows and is represented by two structures. A model of the external world and a model of the current internal status of the architecture.
- Intentions - or what will actually be done, is determined by a process of deliberation, which interprets desires in the light of current beliefs about both the environment and the 'stance' of the system.

S3, using a reasoning process supported by a plan library and the capacity to audit the current status of operational S1 units, structures the intentions into plans, these are then passed to a scheduling process. The scheduling process, in cooperation with a resource bargaining process, responsible for negotiating resource deployment and usage monitoring, schedule the enactment of the plan. The schedule passes to the coordinating S2 channel for dissemination to participating S1 elements.

Environmental change is addressed by S4, which equipped with an Opportunity Analyzer and guided by the S5 desires model, scans the environment for detrimental events or beneficial opportunities. There are two outcomes of this process; the first is the formulation of a view of the outside world which is provided to S5 in the form of the World model. The second outcome is the production of development plans for the future of the system, either exploiting advantageous opportunities or avoiding detrimental occurrences. Plans are then passed to the deliberation process to begin the intention forming cycle again.

As noted above, the power of this approach lies in the recursivity of the underlying model. Figure 2, indicates that the entire architecture described above is repeated in the client S1 unit in the next layer. Consequently, the intentions channel at one recursion informs the desires model in the next, thus allowing an autonomous response to local conditions at each level while remaining within the purpose of the overall organization.

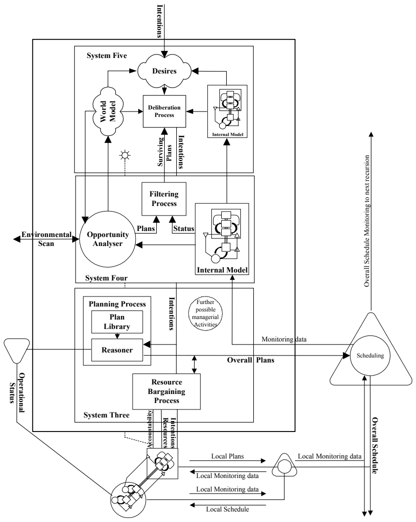

Figure 2. The J-Reference Model

Classical Cybernetics Revisited

While the managerial architecture presented above provides an overall structure for the system, classical cybernetics provides the theoretical foundation for the detail needed to fill out that framework. Here, we consider the more general lessons learned from the cybernetic study of adaptive organisms surviving in a changing environment. Central to this approach is the notion of system state, where both the organism under consideration and the environment in which it exists are described by sets of variables that indicate the current status of each party. Consequently, the environment of the organism can be defined as:

> *"...those variables whose changes affect the organism and those variables that are changed by the organism's behaviour."* (Ashby 1954).

Similarly, the environmental changes that the system is capable of responding to effectively is measured by the number of states that the system can adopt, this measure is termed *variety,* where variety is defined as:

> *"the number of possible states of a system."* (Espejo and Harnden 1989).

This leads to the view that one form of adaptation is aimed at the maintenance of the internal stability of the organism as represented by the state of its critical variables or as Ashby puts it:

> *"...a form of behaviour is adaptive if it maintains the essential variables within physiological limits."* (Ashby 1954).

To accomplish this, the system must have the ability to influence or cause change in other elements that make up the environment. In effect, the system must attempt to exert a form of "control" over some part of the environment. However, control can only be attained if the variety of the controller is at least as great as the situation to be controlled. This is Ashby's *Law of Requisite Variety*, which simply stated indicates that for every environmental action there is an equal and opposite response (Ashby 1956). Moreover, given a range of possible responses to an environmental disturbance, the system must "know" which action to select from a variety of available actions; this is the "*Law of Requisite Knowledge"* (Heylighen and Joslyn 1993). Furthermore, in order to ensure effective control or regulation of a controlled situation requires that the controller model the situation to be controlled. This requirement, known as the Conant-Ashby theorem, states:

> *"Every good regulator of a system must be a model of that system."* (Ashby 1970).

So ideally, an optimal regulator will be an isomorphic model of the situation to be controlled. However, when such isomorphism is not possible as in highly complex systems, then the regulator must be, i.e. contain, a strongly homomorphic model of the situation.

Genetic Algorithms, Genetic Programming & Learning Classifier Systems

Equipping our system to conform to the Law of Requisite Variety requires either the provision of a response for every possible environmental change, the provision of the ability to generate a response or perhaps a mixture of the two. The first option is evidently not scalable/viable/feasible as it requires the software designer to predict and cater for every eventuality the system may encounter and adherence to a consistent terminology/ontology, etc. The second option, whilst even more challenging, offers the potential for significant advances in the nature of systems development.

In constructing an approach to this second option, a range of genetically/biologically inspired approaches have been considered. Beginning simply, a repetition with some variation of Hillis's well-known experiments in designing an optimal $n = 16$ sorting network using a genetic algorithm (Holland 1992) and (incidentally) a Connection Machine (Hillis 1990), was undertaken. Here, the environment is considered to provide an infinite supply of n element arrays of integers for sorting, an approach that allows for a highly controllable "environment" that can range from smooth, i.e. minimal change between arrays, to highly discontinuous where arrays may switch between almost sorted to completely reversed. The task is to evolve an appropriate/optimal sorting mechanism using a genetic algorithm. In our experiments $n = 8$ was used although a Connection Machine was not. Our experiments showed that a sorting algorithm could be indeed be derived and in an acceptable time scale of minutes rather than hours, but to a subset of the possible range of values in $n = 8$, rather than a universal sort algorithm. This, of course, raises the possibility of automatically deriving a portfolio of algorithms each optimized for a specific area of the environment and offering enhanced sorting performance. Clearly, except in specific circumstances, such an approach is unsuitable for general application, although valuable lessons may be derived in determining the limits of applicability and lays the foundation for further work in associated fields such as genetic programming (Koza 1994) or Artificial Immune Systems (AIS).

Having generated a repertoire of responses, it is still necessary for the system to develop the appropriate knowledge to use each response appropriately, i.e. conform to the Law of Requisite Knowledge. Holland's Learning Classifier Systems (LCS) (Holland 1992) provide a means whereby such knowledge may be developed. Such systems are constructed using; detectors to convey the state of the environment to the system, a rule and message system where messages are used both for internal processing and to direct the system's effectors that act upon the environment, an apportionment of credit system that, using some measure of performance provided by the environment, apportions that credit to participating rules and a genetic algorithm to develop new rules. Rules are formulated in `if <condition> then <action>` sets and those that match the condition reported by effectors compete to contribute in the final action reported by effectors. Rules accrue or lose credit depending on the success or otherwise of the outcome reported by the environment. As a consequence, more effective rules gain strength whilst less effective rules diminish in importance. New rules can then be genetically derived from already proven effective rule sets. What emerges is an adaptive model of the environment of the system encapsulated in effective rules, essentially the system *learns* to adapt to the environment in which it resides.

Conclusions

In this paper, the authors have attempted to demonstrate the value of adopting a cybernetic viewpoint as a unifying conceptual framework for the development of autonomic/agent systems. The breadth of experience gained in the study of organisms/organizations surviving in a changing environment provides a solid foundation on which to ground such systems.

References

Ashby, W. R. (1954). Design for a Brain. London, Chapman & Hall.
Ashby, W. R. (1956). An Introduction To Cybernetics. London, Chapman & Hall.
Ashby, W. R. (1970). "Every good regulator of a system must be a model of that system." International Journal of System Science **1**(2): 89-97.
Beer, S. (1981). Brain of the Firm. Chichester, John Wiley & Sons.
Bratman, M. E., D. J. Israel, et al. (1988). "Plans and Resource-Bounded Practical Reasoning." Computational Intelligence **4**(4): 349-355.
Espejo, R. and R. Harnden (1989). The Viable Systems Model - Interpretations and Applications of Stafford Beer's VSM. Chicester, John Wiley & Sons.
Heylighen, F. and C. Joslyn (1993). The Law of Requisite Knowledge, Principia Cybernetica Web. **2005**.
Hillis, W. D. (1990). "Co-evolving parasites improve simulated evolution as an optimization procedure." Physica D(42): 228-234.
Holland, J. H. (1992). Adaptation in Natural and Artificial Systems. Cambridge, Mass, The MIT Press.
Koza, J. R. (1994). Genetic Programming II: Automatic Discovery of Reusable Programs. Cambridge, Mass, The MIT Press.

Author Contact Details

Mr A.G. Laws Primary Contact

Address: School of Computing & Mathematical Sciences
Liverpool John Moores University
James Parsons Building
Byrom St
LIVERPOOL
Merseyside
United Kingdom
L3 3AF

Telephone +44 (0)151 231 2270

Professor A. Taleb_Bendiab

Address: School of Computing & Mathematical Sciences
Liverpool John Moores University
James Parsons Building
Byrom St
LIVERPOOL
Merseyside
United Kingdom
L3 3AF

Telephone +44 (0)151 231 2480

Dr. S.J. Wade

Address: School of Computing & Mathematical Sciences
Liverpool John Moores University
James Parsons Building
Byrom St
LIVERPOOL
Merseyside
United Kingdom
L3 3AF

Telephone +44 (0)151 231 2111