Business as a Control System
the essence of an intelligent enterprise

Richard Taylor, Chris Tofts
HP Laboratories Bristol
HPL-2003-247
December 12th , 2003*

E-mail: {richard.taylor, chris.tofts}@hp.com

modeling, representation, analysis, control

In searching for mechanisms to automate the monitoring and management of information systems in support of businesses, we often lose sight of the need to tie together an understanding of the overall systems dynamics, control points and the measurements that must be made to best exploit them. We propose that one useful (but not exclusive) analogy that can be exploited is that of the control system with a sense of 'self'. By comparing a proposed measurement and control regime for a business to that of a traditional control system, valuable lessons can be learnt about the relevance, practicality and efficacy of measurement and control regimes. This short paper, the first of two, introduces such an analogy and its broad implications.

* Internal Accession Date Only Approved for External Publication

Business as a Control System
the essence of an intelligent enterprise

Richard Taylor, Chris Tofts
{richard.taylor, chris.tofts}@hp.com

Abstract

In searching for mechanisms to automate the monitoring and management of information systems in support of businesses, we often lose sight of the need to tie together an understanding of the overall systems dynamics, control points and the measurements that must be made to best exploit them. We propose that one useful (but not exclusive) analogy that can be exploited is that of the control system with a sense of 'self'. By comparing a proposed measurement and control regime for a business to that of a traditional control system, valuable lessons can be learnt about the relevance, practicality and efficacy of measurement and control regimes. This short paper, the first of two, introduces such an analogy and its broad implications.

Introduction

The term intelligence is often used and as frequently abused by marketing organisations. In their defence, it should be said that a single definition of intelligence is difficult to come by - psychologists, philosophers, neuroscientists, computer scientists, chess players and many more have argued the toss for over a hundred years, with arguments over definitions, metrics (such as IQ) and construction still bitterly contested. In the case of the enterprise however, a more rigid definition that allows us to place our work in the context of all of the warm feelings that come from prepending a noun with the adjective 'intelligent' should help in explaining and assessing the scope of work intended to enhance the ability of large organisations to self-manage.

One very workable (as we intend to demonstrate) definition of intelligence is based upon a combination of classic control theory (an ability to observe the current state of a system, an ability to observe outcome, compute a difference between the measured and the desired and then stimulate appropriate transducers to initiate changes which are predicted to bring the measured closer to the desired), and a distinction between self and non-self[1]. The key properties this (very informal) intelligent system possesses are prediction, observation, control and boundary identification (where boundaries may be fluid over time).

[1] And anyone who attempts to claim biological inspiration for the principles and further development of theories of intelligent enterprises will be taken round the back of the building and kicked until it hurts - and then kicked again.

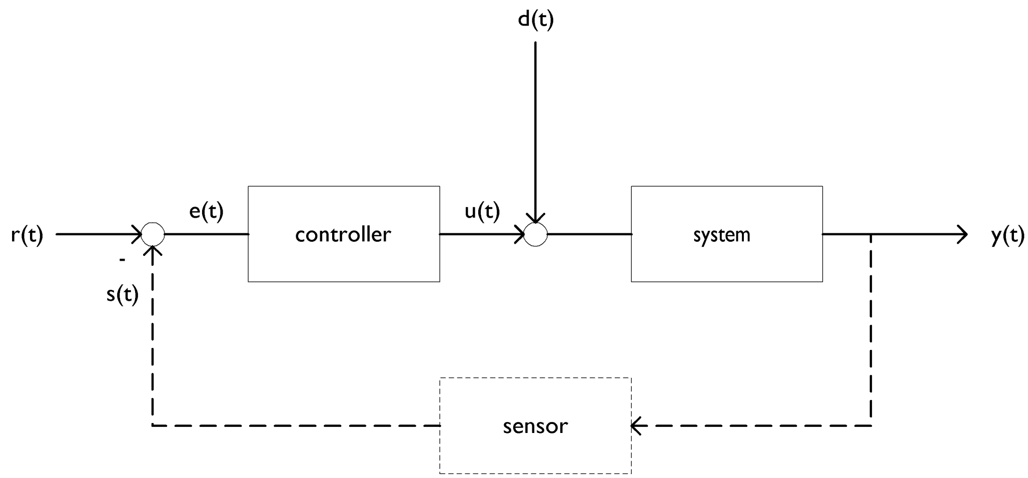

Figure 1: There are two forms of control system - open and closed loop. We have a system which we wish to control (with a desired outcome of $y(t)$) based upon some reference trajectory, $r(t)$. In an open loop system, the controller bases the behaviour of the system on an idealised model, taking the reference trajectory and generating a control or actuation $u(t)$. In most systems however, disturbances to the system ($d(t)$) can cause the output or behaviour of the system to deviate significantly from the reference (or desired) trajectory. In addition, innacuracies in the model we use to predict the effect of $u(t)$ on the system can also cause significant deviations. In this case, a feedback loop (shown as a dotted line) is added to the system, enabling measurements taken by a sensor from the controlled system $s(t)$ to be subtracted from the reference trajectory $r(t)$ to provide error $e(t)$. The controller can then make decisions based upon the difference between the current state of the system and it's desired state (plus extensions of such direct measurement, including first and second derivatives). In practice, most complex systems make use of some form of feedback as well as multiple sensors and actuators.

The purpose of the rest of this document is to explore the effect of such a model on the way we think about the opportunities, the requirements and the technologies that would enable effective information system based intelligent enterprises to thrive. We have been somewhat cavalier (bold sloppiness) in some of our descriptions of control systems to make a point while attempting not to confuse the issues and analogies.

The abstract control system

A control system typically has two groups of input signals - one of which describe the desired state, and the other, key properties of the existing state of the system. These are compared and a model is used to compute the most appropriate response that will move the measured towards the desired. In a steady state system, a series of such measurements, comparisons, computations and actuations will eventually move the system from its current to its desired state. The rate at which measurements are made and the actuators fired is governed by the response time of the system - the rate at which change can occur and the nature of the trajectory that is being planned by the controlling algorithm. From this (very simple) description of a closed loop control system there are several observations that can be made

1. the parameters measured and the driven actuators are determined by the model of the system under control; typically this is an abstraction of the real system being controlled (the ‘drive by wire’ accelerator in a normal car only takes as inputs the position of the carburettor and accelerator pedal, and their respective rates of change, despite the wide range of other sensory inputs available - cabin temperature, time of day, position of sunroof etc...);
2. for the same set of inputs and outputs, different (functionally correct) control algorithms can have very different outcomes, from the rate at which they converge on the desired state, to the ‘violence’ of oscillation that occurs before that state is reached;
3. observations may well not be directly on data that is required by the control system, this must often be inferred from measurements of multiple sources;
4. through a combination of shared observations and shared resources (the position of an aircraft plus it’s vectors, combined with ailerons and wings), many control systems, while they can be considered as individual units for some forms of analysis are both coupled and occupy a position in a hierarchy of such systems;
5. the rate at which change can be completed, and naturally the ability of the system to track changes in a dynamic environment is primarily limited by the time constants of the underlying system, and the quality of the controller.

The business control system

The issues faced by the designers of engineering control systems all have direct analogues in the business systems we wish to design and operate;

1. the question of what comes first - model or observations/actuators is not moot. Without a model that describes the relationship between observations and actuation, even given an extant system, the choice of observations, their frequency, the choice of actuators and their frequency can not be made. Top down (in the abstraction sense) design of a system from model to instantiation has been overwhelmingly demonstrated to result in more reliable, robust designs than those that are retrofitted.

   A striking example of this is given by SAP based ERP systems. In order to manage the real time constraints of the critical transactions in such a system it is necessary to identify, track and profile not only the critical tasks but also their by-blows. Current transaction architectures have not been designed to gather the appropriate information at the appropriate point in the stack, which means that the best that can be done (without massive retro-fitting) is to gather massive amounts of low level data and direct data mining exercises based upon a fuzzy notion of the relationship between business processes and the transaction stack.

   One might argue that this represents a great opportunity for consultancy, since only experts with a broad understanding of not only the systems but also the original (often undocumented)

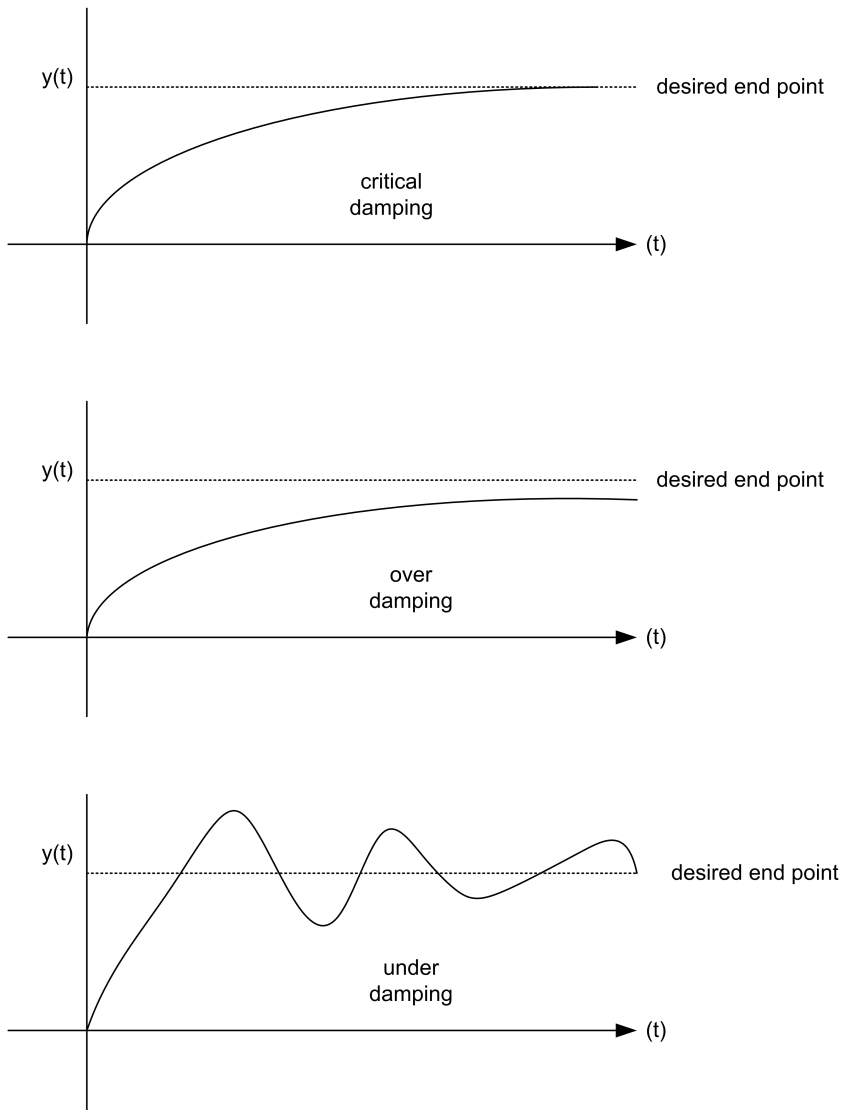

Figure 2: Characteristic responses of a control system : (a) a critically damped system reaches the desired state as rapidly as possible; (b) an overdamped system overshoots before reaching the desired state; (c) an under damped system oscillates around the desired state before reaching equilibrium.

ERP systems analysis can hope to make a contribution. It is however highly debateable whether such a contribution can ever be scaled, in general, the capacity of such a consultancy activity will be roughly linear with the number of experience consultants available. *moral: design the system with control in mind, design the measurement system with control in mind.*

2. competing businesses with apparently similar structures and control points demonstrate very different dynamic characteristics. While excellence in execution is an obvious discriminator, the ability of an organisation to measure and control appropriately requires a fundamental understanding of the relationships between inputs and outputs.

   This might simply be a matter of cost - over engineering a measurement system when the information can not be applied is at best an expensive luxury. Alternatively it might even be harmful. Overcompensated systems that are automated tend to oscillate (too much product, too little product, too much. . .) Manually controlled systems (and the micro management of finances within a very large organisation is a good example of this), give the controllers an exaggerated and potentially dangerously misleading impression of what they are capable of controlling.

   Examples of this include notorious (at least amongst US listed companies) 'end of quarter' financial controls (the stop-start nature of which can be seen to be excessively destructive to the deep pipelines that exist for many activities). Over a period of time, external compensating effects (such as employees anticipating shutdown and making appropriate buffering decisions) may well come into play and further damage the global behaviour of the organisation.

   Nestlés real time ERP for factory systems provides another example of such a system - a controller in Switzerland can observe the behaviour of production lines in York at levels of granularity that are several orders of magnitude greater than their ability to control them.

   While a futile attempt to impose high frequency changes on a system with a long time constant might simply be irritating, the cost of such a system must be exaggerated to cope with more onerous real time constraints and the associated data volumes. *moral: the nature of the control system as well as the structure of the underlying system are necessary for optimising the control for the environment.*

3. indirect inference of control information is not always trivial. Managing the real time experiments that might be required to disambiguate management information generally requires a formal (computable) statement of the relationship between views. Without this, automation of drill down (for example) in measurement systems for root cause analysis becomes unnecessarily difficult; *moral : a statement of the computable relationship between data sources is required for disambiguation and also for generalisation.*

4. coupling in business systems, especially as they always have shared resources, (which might 'simply' be the capital available to an organisation, but is more likely to be a complex mix of investment resources, personnel, acceptable risk etc. . .) will always be with us. For all control systems, the old software engineering principle of no coupling without cohesion (no taxation without representation?) holds good.

   Assessing the level and the type of coupling between the different parts of an organisation (if for no other reason that to work out whether they can be ignored) is impossible without representations of the resources each hold, the flow of these resources and their relative importance in the life of the organisation.

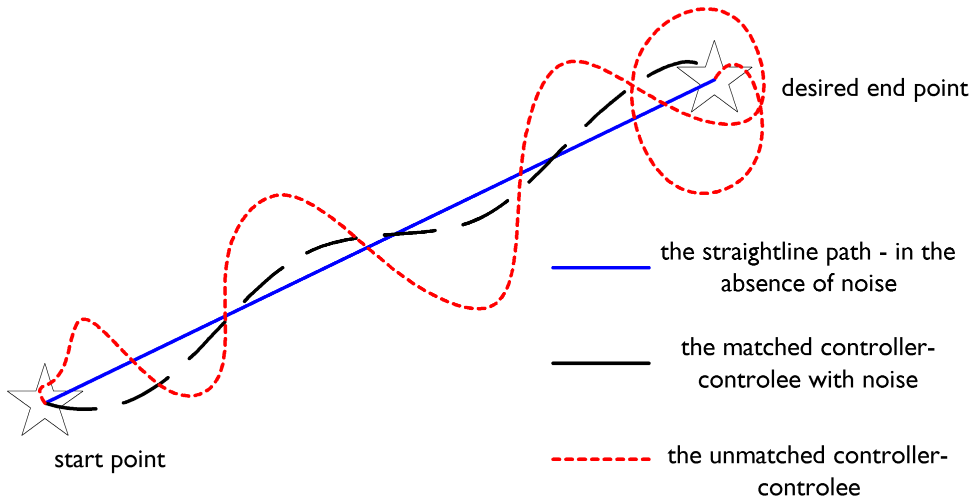

Figure 3: Attempting to control a system with a long time constant through occasional high frequency perturbations rarely has the desired effect - as quickly as the operator would like - the analogy of the tug boat and the super-tanker is a good one here, despite the fact that the tug *could* manoeuver faster than the tanker, it matches its response to the time constant of the tanker.

Neglecting treasury functions in favour of a manufacturing report, for example, in a resource constrained ERP system can be extremely expensive, very quickly. The representation of such couplings and their investigation is a requirement for any serious holistic business systems analysis. Control systems with their models of phenomenon-observation, actuation-phenomena provide a sound first point of call in assessing potential couplings and their nature. *moral : treating a business system as a set of uncoupled control issues is likely to lead to unpredictable global outcomes. Using the control systems and the phenomena and actuators that they observe/control can provide insight into potential couplings.*

5. briefly returning to the issue of time constants - most organisations are represented by groups of coupled control systems with very different time constants. Quite apart from the need to customise the controllers to each, the detection of both short and long term trends for capacity planning and forward risk analysis requires an understanding of the nature of these time constants, as well as an appreciation for how and why they might vary. Examples of real life systems in which a combination of time constant and resonance has been (sometimes fatally) misunderstood are common.

   Characterising technologies by their impact on time constant, and then assessing the impact of that on the overall system provides a far more sophisticated view of the impact of technology, analogous to the recognition of the von-Neumann bottleneck. *moral : time constants are important, they determine the most appropriate mechanisms of control and may also effect other coupled systems with apparently very different dynamics.*

The response of the control system to error

When a business system fails to meet expectations the investigation into the reason for the failure often appears unfocussed. The control system framework provides a means of structuring the questions that must be asked of both the assumptions and the execution.

- are the observations we made of the real world consistent with what we expected to observe (and tested the control system against)? If they differ, what is the form of the difference - rate of change? absolute values? ability to collect the data? quality of observation? observational obfuscation?
- if the observations are consistent with our expectations, is the controller generating actuations that are consistent with the stated goals of the organisation (tested through the internal model from which the controller was derived)?
- are the actuations implemented effectively - are the predicted time constants consistent with what was expected by the controller?
- is the world model that is being used to test controllers against adequate? What needs to be changed - assumed structure? response times? actuation probabilities?
- do we need to modify the form that the actuation mechanisms take?

To conclude, the control system based view of our activities, their cross-coupling, observation and actuation provides a means of placing technologies and their relative values into the context of the overall (business) system goals. As well as providing a means of representing the overall system, the functional and non functional requirements for new sub systems, their measurement requirements, the relationships between measurements and control, and long term trend and forecasting can all be achieved within this framework.

Self and non self

All intelligent systems that are not closed require the ability to distinguish between self and non-self. Distinguishing in this way is an important part of the controllable - non controllable spectrum. There are several distinct issues that need to be addressed

1. what we can(not) measure (or estimate)?
2. what we can(not) control?
3. what would be the benefit of converting non-self to self (through strategic alliance for example) in terms of additional control and monitoring opportunities?
4. what would the benefits be of divesting ourselves of self?

5. what id the effect of noise - both in the measurement and the actuation paths?

One of the advantages that the 'intelligent enterprise' possesses over the it's natural analogue is the potential to construct itself in such a manner as to be able to control what is self and non self in a more rational[2] and constructed manner. While history and culture will combine with environment to restrict the choices that an organisation can make, exploiting the properties of observation and control in making decisions about the degree of self, non-self can improve the quality of decision making that an organisation makes. When combined with *agility metrics and models*, decisions about sourcing and capacity development for example, can be assessed in terms of the impact they have on the organisations ability to react to change. Bringing external partners 'into the fold', or alternatively restricting their role within the organisation can be seen as contributing in a measurable and at least (partically) predictable manner to the organisations agility. Further, the observation that many, apparently sophisticated, control systems can be reduced to a relatively simple continuous representation allows some estimation of the potential instabilities in coupling such systems to be made.

Conscious vs autonomic

It is worthwhile briefly touching on the difference between the conscious and the autonomic in this discussion. Autonomic functions, functions that are carried out without thinking about them, extend the ability of an intelligent system to operate. They are however generally considered to be subsidiary to the intelligent system. As touched upon in the introductory discussion, there is unlikely to be a hard line dividing autonomic and intelligent - what looked intelligent five years ago is likely to be considered simply supportive automation tomorrow. The reason that we might wish to place technologies along this continuum is to enable us to discuss their role and the required abstraction translations that are necessary to exploit them.

Conclusions

The augmented control system provides a useful analogy to the intelligent business enterprise. Questions of measurement (requirements, type, quality, frequency and detail) are placed in the context of acutation or control - something that has been frequently forgotten, or at least sidelined within HP with it's traditions of measurement as opposed to control technologies.

Exploiting notions of control in a determination of self and non-self can prove critical differentiation, especially for organisations that are attempting to place value on the construction, management and maintenance of relationships with other organisations. Finally, placing technologies (extant, under development or proposed) on the autonomic-conscious continuum may provide insights into relative values as well as requirements for integration.

[2] Often honoured more in the breach than observation - but one can hope.