A new Approach to Active Network Management

Ian W Marshall and Chris Roadknight
BT Labs, Adastral Park, Martlesham Heath Ipswich IP5 3RE

Abstract Future multiservice networks will be multifractal, and thus impossible to control accurately using conventional finite-state-machine models. Active networking can reduce the cost to operators by delegating management to network users, but this will increase the overall system complexity. Management architectures for active networks must therefore use a new approach. An open-ended alternative to existing architectures is proposed, based on delivering homeostasis using environmental control and feedback hierarchies.

1. Introduction

There has been a long sequence of attempts to create multiservice networks, including ISDN, IN and ATM. Historically multiservice offerings have not won the expected interest from consumers, as the available services were not sufficiently attractive. To a large extent this was a result of the enormous difficulties encountered in creating, deploying and maintaining large numbers of new mass-market services on the available infrastructures. On the other hand, consumer take-up of Internet based services has been extremely rapid, primarily due to the fact that end-users are free to add whatever features they like in the end systems. Despite the fact that existing Internet performance is inadequate for "real-time" communication, the number of services offered on the existing Internet is growing exponentially, along with the associated management costs. Current research efforts focus on improving the Internet by extending router capability with features such as QoS control, to support real-time applications, and multicast to support multi-party interactions. Embedding these new value-added features in the network will most likely make the costs and complexity escalate even faster.

Active networking [1] takes a different approach. The aim is to enable network clients to easily add new service components to network devices. Often the resulting service performance will be significantly better than if the component executed in the user's end systems. For example, a network based conference gateway can be located so as to minimise the latency of the paths used in the conference, whereas an end system based gateway will usually be in a sub-optimal location. Users of an active network supply the programs and policies required for their custom services in transport packets alongside their data. Active networking therefore reduces network operating costs by delegating management functions to the user's network terminals.

However, it is already apparent that Internet usage is multifractal [2,3]. This implies that the number of accessible states is almost infinite. In addition the structure of the Internet is so complex that its 'State' cannot even be estimated accurately [4]. Therefore, unlike previous "multiservice" networks, the Internet cannot easily be managed using an information modelling approach based on finite-state-machine control theory. Active networking will increase this complexity, since it deliberately removes restrictions on what network devices can be used for, and therefore on their possible state. In addition active network research has shown that, for optimal performance, some of the programs must execute in the data path, where they must operate at line speeds. This carries a heavy penalty in performance or cost unless the device hardware is made dynamically configurable (e.g. using programmable logic arrays [5]). In other words active networking will introduce significant new degrees of freedom.

2. Management in a Fractal Environment

Clearly active networks will need a new management approach that is designed to constrain a system with chaotic properties, and ensure that the system will always respond quickly, whilst allowing adaptation to deal with new types of demand. It is well known [6] that biological systems can quickly adapt to fractal variability in the environment, are resource efficient, and are often stable for very long periods of time. Therefore, in the absence of any good control theory for chaotic systems, we have looked for guidance from biology. Our approach is based on observations of prokaryote ecology [7,8], which we believe offers a good analogy with an active network. We model a set of programs executing within a small unit of isolated

programming space as a bacterium (isolated within a peptidoglycan cell wall). An important aspect of bacteria is their ability to exchange genetic information, and adopt new genes that are locally advantageous. This means that there are no distinct species and competitive pressure acts directly on the genes. In the model, service components (programs) are analogous to genes, and can be supplied by users or obtained from other nearby bacteria, in response to the user's service demands. A service may be performed by a single bacterium, a homogeneous colony, or a heterogeneous team. In the latter case the steps are normally executed on diverse network nodes. Demand is the analogue of food and revenue is the analogue of energy extracted. Different types of demand require different service components (genes) in order to extract revenue and prolong life. Inefficient service components tend to be replaced when better versions become available (either through mutation or external input). Each bacterium can execute up to six service components concurrently (real bacteria have up to six mobile genes), and will pass demand it cannot handle to neighbours. The service components are written in the form of policies [9], and can be abstracted into a high level language accessible to users. The bacterium also has a more static genome consisting of policies supplied by the operator that are mostly concerned with when to import or export genes, when to reproduce and when to die. There are currently around 20 genes of this type in the model (again similar to the number in real bacteria). The bulk of the 200 or so genes in a real bacterial genome are related to the structure of the bacterium. Since we do not yet allow the bacteria access to hardware configuration, this type of gene is not required in our model. To enable long term adaptation without intervention all the genes are allowed to mutate and can move from the static part of the genome to the active part of the genome (and vice versa).

3. Initial Results

Although the stability improves as the numbers of bacteria increase, we have found that even small numbers exhibit very desirable properties without requiring intervention. Figure 1 shows the response of a colony of up to 400 bacteria to random cumulative increases in demand (at 50 second intervals) for 4 network services. The service latency degrades at the start of each demand increase, but is then autonomously restored to a low level.
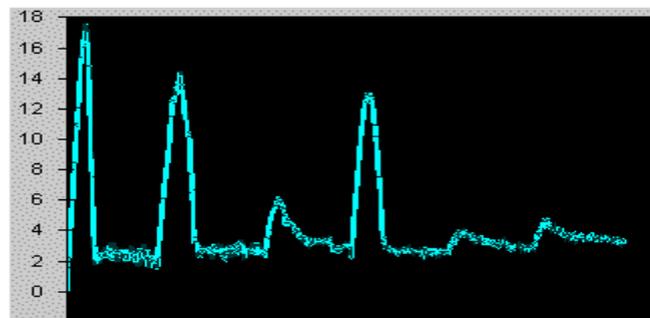
Figure 1. Service latency in a self regulating active network

References
[1] Covaci S (ed) "Active Networks", LNCS 1653, Springer (1999)
[2] Feldman A, Gilbert A, Huang P Willinger W "Dynamics of IP traffic: A study of the role of variability and the impact on control" Comp Comm Rev 29, 4 pp301-313 (1999)
[3] Marshall I and Roadknight C "Linking cache performance to user behaviour" Computer Networks 30 pp2123-2131 (1999)
[4] Allman M and Paxson V, "On estimating end to end network path properties" Proc Sigcomm 1999 pp263-274
[5] Sanchez E, Sipper M et.al. "Static and dynamic configurable systems" IEEE trans on computers, 48, 6, pp586-564June 1999
[6] Kauffman S "Origins of order" New York OUP 1993
[7] Sonea S and Pannisset M, "A new bacteriology" Jones and Bartlett 1983
[8] Caldwell D et.al "Do bacterial communities transcend Darwinism?" Adv Microb Biol 15 pp105-191 (1997)
[9] Sloman M., "Policy Driven Management for Distributed Systems", Journal of Network and Systems Management, 2, 4, pp333-360 (1994)