University of St Andrews

School of Computer Science

Distinguished Lecture Series 2003/04

Tuesday 9th December 2003
Lecture Theatre C, Mathematical Institute,
North Haugh St Andrews

Towards Automated Management of Large-Scale Distributed Systems

by

Professor J. Sventek
University of Glasgow

As the scale of distributed systems continues to grow dramatically, it is becoming increasingly difficult to manage such systems manually. Automating the management of large-scale systems is not trivial, requiring closed-loop activities at many levels of abstraction.

Some researchers think that management and control of large-scale system will only occur when we understand the emergent behaviour of such systems, as epitomized by IBM's "Autonomic Computing" challenge to computer scientists. What is clear is that we need to think "out-of-the-box" with regards to how to measure, analyze, and control such systems in order to automatically manage their activities.

This series of lectures will describe how researchers at Hewlett-Packard Laboratories and Agilent Laboratories have attempted to address some of the issues that stand in the way of automated management of large-scale communication systems.

Biography

Professor Sventek obtained his B.A. in Mathematics from the University of Rochester and his PhD in Nuclear Chemistry from the University of California. He is currently the Professor of Communication Systems in the Department of Computing Science at the University of Glasgow. Prior to joining Glasgow, he had a distinguished career pursuing research into networked and distributed systems and managing research teams at Lawrence Berkeley Laboratory (1979-1986), Hewlett-Packard (1987-1999), and Agilent Technologies (1999-2002). His research interests include programmable networks, embedded systems, closed-loop network management, and distributed system architectures. He has several publications on these topics, as well as holds four patents (with three other patents pending) in these particular areas. Professor Sventek was the principal author of the original OMG CORBA specification as well as several of the Common Object Services (Trading, Events, Naming); he also was the rapporteur for the TeleManagement Forum's most recent release of the Technology Neutral Architecture document. He has been the general chair for TINA99 and Middleware 2001, programme chair for COOTS98, TINA99, and Middleware 2000, and a member of programme committees too numerous to mention. He is an advisor to the TeleManagement Forum Board, is an adviser to the Wiley Series in Communications Networking and Distributed Systems, and was on the editorial board of the IEE/BCS/IOP Distributed System Engineering Journal.

Programme

10.15 – 11.15 Lecture 1: Traditional Network Management Systems

Network management systems have been with us for many years, for both telephone and data networks. In order to understand what innovations are needed for network management systems to automatically manage large-scale environments, it is necessary to understand how traditional network management systems (called Operational Support Systems, or OSS's for short) are constructed and used.

This lecture will describe the traditional structure of OSS's, how they are typically used, and their scaling characteristics. It will conclude with a discussion of a network management pattern that is observed at several levels of abstraction in modern network management systems, and how these pattern instances must be related to automate management of the system.

11.15 – 11.45 ***Coffee in John Honey Building***

11.45 – 13.00 Lecture 2: A Scalable Control Plane Architecture

As described in the 1st lecture, one aspect of traditional OSS's that scales poorly is the control plane, since it assumes centralized control AND explicit communication between OSS Central and the components of the network when configurations need to change. As the number of components making up the network increases, this centralized structure simply implodes.

This lecture describes one, out-of-the-box, approach to re-architecting the control plane to eliminate this problem. It assumes that responsibility for being in the appropriate configuration is delegated to the individual components, and focuses on making sure that the system asymptotically approaches the correct configuration over a bounded time span. As such, this control plane architecture takes advantage of emergent behaviour.

14.30 – 15.30 Lecture 3: An Always-On Active Measurement Approach for IP Networks

As described in the 1st lecture, another aspect of traditional OSS's that causes problems in large-scale IP configurations is the lack of built-in diagnostic measurement information that relates to a user's traffic. The measurements available to the OSS are low-level and primitive; additionally, they are usually taken after a problem has been detected, usually by out-of-band means.

This lecture describes an active measurement mechanism that we have devised to permit continuous measurements of the behaviour of a user's flows, such that the measurement load is bounded. It will work

in IPv4 networks, but is a natural for IPv6 networks since the measurement payloads can be embedded in IPv6 extension headers. Such measurements are introduced in programmable network elements, and can vary in complexity and sophistication to suit measurement needs.