

Conference Participant Report: OPENARCH 2002, New York, USA

Carl Cook

June 29, 2002

Overview

The Fifth Annual IEEE Open Architectures and Network Programming Conference presented 15 papers from 56 submissions. Whilst programmable and active networks are traditionally presented in this conference, this year's conference also hosted papers on emerging topics such as overlay and peer-to-peer networks. The conference was held in the Hilton Manhattan, in the heart of the city. Approximately 80 people attended the conference, which is co-shared with IEEE Infocom.

Conference Welcoming

The conference started with a welcome talk by Larry Peterson, co-chair of the programme committee. Larry noted that this year yielded the largest number of submissions, and also commented positively on the quality of the accepted papers.

Keynote Speech

The keynote speaker was Bill St. Arnaud, director of Advance Networks. Bill talked about the control of high-bandwidth optical networks by the end users—eg the university campus, the hospital, etc. He discussed the problems associated with this, such as configuration protocols, inter-domain issues, and issues with the current client/server network control.

The speaker then suggested that the internet peer-to-peer model will be suitable to control switched networks, as opposed to using conventional routing. As an example, for a large file transfer, a dedicated path should be established, based on the fact that large overheads exist using traditional packet forwarding and TCP. The underlying problem to solve is how to build frameworks to support dedicated paths, and how to reduce the bottlenecks caused by software-based routing.

The topic then changed to a related Canadian research project, originally called the Spanish Inquisition, but now funded by Cisco and renamed to Scalable Architectures and Frameworks. The project focused on user-controlled high speed optical networks, mainly for the purpose of large-scale file transfers. The storage will be located in space, reducing government regulation of peer-to-peer networking such as napster systems. Many questions were raised by the audience relating to security and granularity of control for customers.

Technical Programme

The following is a summary of the conference proceedings. These notes were taken during the presentation of the papers. A copy of the official proceedings (OPENARCH 2002) is available from my office.

1. Primitives in Support of Overlay Networks

Network Layer Support for Overlay Networks: Whilst overlay networks are useful for configuration without consensus, many overheads exist. This paper discussed the problems in designing and configuring overlay networks, and proposed a framework. The design goals of the framework included flexibility, simplicity, incremental in deployment, and limited for security reasons.

Topology Discovery Service for Router-Assisted Multicast Transport: This paper works towards scalable multicast transport in an active network, even in the presence of several passive, conventional routers. A simplified architecture is presented, where the network consists of active nodes connected within an overlay network, reducing complexity and security risks. Emphasis was placed on active signaling between adjacent routers to determine reverse paths for multicast routing, allowing topology discovery, etc.

PAMcast—Programmable Any-Multicast for Scalable Message Delivery: This paper presented another framework for multicast and anycast, based upon message delivery modes and architectures. This architecture allows packet delivery to any number of individuals out of a multicast group, with reasonable computational and state requirements.

2. Panel Session: Tech Transfer, a Path to Relevance

This session involved a discussion about how to get network research out into the real world, and was presented by both researchers and industry representatives. The general message was: simple APIs, permanent test-beds, and actual next-generation networked applications are regarded as keys to establishing strong public usage.

As application-specific networks are too expensive to develop (using cable TV and phone systems as examples), the case for an open service delivery system. Hiding complexity, providing flexibility, and platform independence are all key goals for a framework that supports open service delivery.

The general idea is that applications and customers rule the development of networks. Therefore, if an infrastructure is relatively open, consumers can be exposed to new research ideas, with the potential for market acceptance and possible funding.

A focus on active services rather than active networks was called for. Instead of trying to replace the entire network infrastructure, the injection of services at key route points or end services was seen as important. The limitation of network down-time was also emphasized. This led to an open discussion with the audience. The conclusion was that active networks will be replaced with overlay networks, and that this is effectively the second birth of active network related research.

3. Constructing Services from Components

Modular Components for Network Address Translation: This paper described a general-purpose toolkit for network address translation. The system, Click, offers facilities to allow users to define where and how NAT takes place, based upon an array of selectable system components. The system has been in use for 18 months; examples of its use and a performance evaluation are presented in this paper.

A Comparative Study of Extensible Routers: As the title suggests, this paper presented a survey of three software-based (and programmable) routers, as used within active network test-beds. The paper also gave a fairly simple model of extensible routers, based upon similarities between all three architectures. The model defined four key components of extensible routers: Queues, classifiers, schedulers, and forwarders. The paper argued that any extensible or active router must support user configuration of these components.

Forwarding without Loops in Icarus: An active network-based implementation of a loop prevention filter was presented in this paper. Icarus uses Bloom filters to detect loops, and is implemented within the ANTS active network toolkit. It turns out that such a filter is easy to implement, and relatively efficient. Icarus can be used in detecting multicasting loops as well as bad unicast traffic.

4. Network Processors and Configurable Hardware

The Potential of Run-Time Compilation in Active Networks based on Network Processors: As the title suggests, this paper presented a method for compiling active packets in byte-code form to native machine code for network processors. The SNAP active network framework was used to show the potential of this technique. Each SNAP operator was tested, both before and after compilation.

NetBind—A Binding Tool for Constructing Data Paths in Network Processor-Based Routers: This paper also takes an in-depth look at network processors for active networks. A dynamic binding algorithm is introduced, providing better binding performance between the processor and the software components than existing network processor systems. The design and implementation of the system, NetBind, was presented in detail.

APE—Fast and Secure Active Networking Architecture for Active Packet Editing: This paper presents a node operating system for active networks. The system has two main facilities: fast software-based routines for packet processing, and firmware-based packet classifier and editor. Upon receipt of active packets, the software components dynamically configure the packet editor. Security is also addressed in this system, by way of key allocation for each packet flow.

5. Architectures and Frameworks

A Framework for Efficient and Programmable Sensor Networks: A framework called SensorWare was introduced in this paper. This is an event-based architecture that allows user-defined scripts to control sensor networks. SensorWare is customised for ad hoc networks.

ComAN—A Multiple-Language Active Network Architecture Enabled via Middleware: The paper presentation went quite well. I ran out of time, as per usual, the feedback from the audience was positive. One question at the end could have been paraphrased as “so what?!” , but I managed to handle that question. . . I have had that one before.

Safe Kernel Programming in the OKE: This paper presents a system for loading user-defined modules into the linux kernel. It is not network specific, but is likely to be suitable for active networks, and members of the audience stressed that they would like a copy of the source code to try the system out with existing active networks. Security in the system is managed by a trusted compiler, restricted and security-customisable languages, plus security policies.

6. Short Papers Session

Several short papers were presented in this system. Papers included surveys (with the intention of making development into discovered shortcomings/issues), security frameworks for active networks, peer-to-peer support, programmable service deployment, active distributed caching, and an update on the FAIN project. A real assortment of research groups were present, with groups coming from Japan, Germany, the U.K., and Australia, as well as the United States.

7. Resource Management

Global Flow Control for Wide Area Overlay Networks—A Cost-Benefit Approach: This paper presented a flow-control protocol for multicast and unicast active/overlay networks. The protocol does not focus on flows of packets, but rather the individual links between participating routers. Simulations were presented, along with results from live internet testing. The protocol attempts to achieve simplicity, fairness, and cost minimisation.

Opus—An Overlay Peer Utility Service: This paper presented a model for distributed computing where a collection of servers, distributed throughout the network, support the requests of all networked applications. Several facilities are available to distributed applications to avoid constant re-implementation of network-based routines. Resource management, topology consideration, and quality of service mechanisms are discussed in this paper.

CALM—Congestion-Aware Layered Multicast: Congestion awareness was addressed in this paper. The system provides congestion information to end-systems, such as location of congested links, and the levels of congestion. It was shown that this type of information is particularly useful for layered multicast systems, allowing reductions in the router state information, and supporting the coordination of receivers. A focus on a lightweight system was emphasised, and the paper then gave a description of the proposed CALM architecture.

Conference Dinner

The conference dinner was held at the Playwright Tavern and Restaurant, located on Broadway and Times Square. This was a surprisingly quiet event, probably due to the large number of computer scientists :)

Other Comments

Holding conferences in New York is a particularly good idea—you are guaranteed to find something interesting as soon as stepping out of the hotel door. New York has attractions in all directions: Central park, the Empire States building, Times Square, the Statue of Liberty, Wall Street, Broadway, Soho, the Rockefeller Centre, just to name a few. The city has a very vibrant atmosphere. This is not just due to the sheer size and number of sky-scrapers. . . there are crazy people everywhere. A final few words of advice: avoid crossing the road if a taxi is in the area, and always yield to cyclists under the influence of illegal substances.