



**ACM SIGCOMM 2002 CONFERENCE  
PITTSBURGH PA ~ AUGUST 19-23**

# **ADVANCE PROGRAM**

**Applications, Technologies, Architectures and  
Protocols for Computer Communication**

**August 19-23, 2002**

**Omni William Penn Hotel**



[www.acm.org/sigcomm/sigcomm2002](http://www.acm.org/sigcomm/sigcomm2002)



# WELCOME TO PITTSBURGH

## Welcome to Pittsburgh

Welcome to Pittsburgh and ACM SIGCOMM 2002. We have an outstanding technical program and exciting tutorials. The program committee selected 25 presentation papers from 300 submissions. In addition to the SIGCOMM award keynote, this year's program also includes an invited talk by Mike O'Dell, Workshops, Tutorials, a Student Poster Session, and two Position Papers.

We would like to acknowledge the generous support of AOL, Intel Research, Agilent Technologies, AT & T, Cisco Systems, Microsoft Research, Qwest and the NSF.

We hope that you have a pleasant time in Pittsburgh.



## Registration

**Early registration deadline:** July 26, 2002.

On-line registration:  
[www.acm.org/sigcomm/sigcomm2002](http://www.acm.org/sigcomm/sigcomm2002).

**Special requests:** If you have special dietary requirements or need other special hotel arrangements, please indicate that on the registration form.

**Cancellation policy:** Cancellations must be received in writing by contacting the ACM member services department (POB 11405, New York, NY 10286-1405, USA). A \$50 cancellation fee will be charged.

## TABLE OF CONTENTS

Welcome/Registration Information	Page 2
Program Overview	Page 3
Workshops	Page 4
Tutorials	Pages 5-7
Technical Program	Pages 7-8
Conference Details	Page 8
Travel Information	Page 10
Local Information	Page 11
Sponsors	Page 11

# PROGRAM OVERVIEW

## Monday, August 19

### Tutorials:

IP Control of Optical Networks, and GMPLS (half day) - *Jennifer Yates and Albert Greenberg, AT&T Labs-Research*

Detecting Packet Patterns at High Speeds (half day) - *George Varghese, University of California, San Diego*

How to Use the Emulab Public Network Testbeds

*Jay Lepreau, Mac Newbold, Rob Ricci (University of Utah) and Chris Alfield (University of Wisconsin, Madison)*

**Note:** This tutorial is offered as both a half-day morning lecture and a full-day version that includes an afternoon lab session.

## Tuesday, August 20

### Tutorials:

Data-mining the Internet: What we know, what we don't, and how we can learn more (full day)

*Michalis Faloutsos, University of California, Riverside and Christos Faloutsos, Carnegie Mellon University*

### **New** Workshops:

Computer Networking: Curriculum Designs and Educational Challenges

**CANCELLED** - Retrospective Workshop on Asynchronous Transfer Mode (ATM)

### Opening Reception

## Wednesday, August 21

### Introductions, Awards, Keynote

*Keynote Address - Dr. Scott Shenker of the ICSI Center for Internet Research, Sigcomm 2002 Award Winner*

**AM Technical Program: BGP analysis**

**PM Technical Program: Overlay Networks**

**PM Technical Program: Congestion Control**

**Banquet**

## Thursday, August 22

### **New** Invited Talk - Mike O'Dell

*"Reflections on the Future of Very Large Data Networks"*

Large-scale engineering is an economic optimization problem. The tools required are indeed technological, but their evolution must be driven by the economic nature of the problem. This talk will offer lively, contrarian, and perhaps even curmudgeonly reflections on those economic matters and what they portend if we are to see the successful deployment of very large data networks.

**AM Technical Program: Measuring and Simulating Networks**

**PM Technical Program: P2P and Multicast**

**Student Poster Session** (with refreshments)

**PM Technical Program: Routing Dynamics**

## Friday, August 23

**AM Technical Program: Router Design and Analysis**

**AM Technical Program: Measuring Paths and Flows**

### **New** Position Papers: (the final presentation will end at 2:30 p.m.)

"An End-to-End Approach to Globally Scalable Network Storage"

"Tussle in Cyberspace: Defining Tomorrow's Internet"

# WORKSHOPS

## Computer Networking: Curriculum Designs and Educational Challenges (9:00 a.m. to 5:00 p.m., Tuesday, August 20, 2002)

The goal of this workshop is to bring together faculty from a broad spectrum of academic institutions, industry engineers, scientists, and others with an interest in education to discuss curriculum design and teaching practices in the field of computer networks. Through a series of highly interactive panel discussions and breakout sessions, the workshop will address:

- **Undergraduate Curriculum.** What "core" topics should be covered in a first networking course? What are the roles of labs? What is the relationship to ACM and IEEE 2001 curriculum recommendations? Are there "model curricula"?
- **Graduate Curriculum.** What should graduate courses in our field look like? What multi-course sequences are possible?
- **Lab courses.** What approaches can be taken in developing "hands-on" laboratory-based courses at the undergraduate/graduate level?

The workshop will provide an opportunity to discuss these and other related issues; to share viewpoints, experiences, and materials; and to learn what others are doing, and considering doing. The workshop will produce a report detailing the attendees' discussions and may serve as a springboard for future activity.

**Organizing Committee:** *J. Kurose (U. Massachusetts), J. Liebeherr (U. Virginia), E. Nemeth (U. Colorado), S. Ostermann (Ohio University), T. Ott Boisseau (CAIDA)*

Participants are encouraged to submit a one-or-two-page paper for distribution to workshop attendees by June 14. See <http://www.acm.org/sigs/sigcomm/sigcomm2002/workshop1.html> for more details, including possible travel grant opportunities.

## Retrospective Workshop on Asynchronous Transfer Mode (ATM)

### **THIS WORKSHOP HAS BEEN CANCELLED**

In the late 1980s and early 1990s, Asynchronous Transfer Mode (ATM) was the new, exciting protocol. Based on work by some of the leading researchers in data communications, it progressed from an idea to a standard in a very short time, and because of its wide applicability, ATM products were built for a wide variety of markets. Now, a little over a decade later, ATM has settled into middle age. It is playing an important role in some markets, and a niche role in some others.

This workshop seeks to capture as best we can, some of the important lessons or experiences from ATM's history to date. The goal of the workshop is to give its participants, and the SIGCOMM community (through a published workshop report) a greater perspective on how promising technologies evolve both in the research laboratory and the marketplace. Potential topics include: the ATM standards process and how it worked, ATM's experience in different markets, and technology perspectives on ATM.

Participation in the workshop will be based on 3 page position papers that are due June 14. See <http://www.acm.org/sigs/sigcomm/sigcomm2002/workshop2.html> for more details.

# TUTORIALS

## IP Control of Optical Networks, and GMPLS

*Jennifer Yates and Albert Greenberg*

*AT&T Labs- Research*

*Half Day – Day 1 AM*

### Content:

Advances in optical networking hardware, particularly optical cross-connects, promise huge, rapidly configurable bandwidth – enabling rapid provisioning and cost-effective restoration. The challenge is to manage the bandwidth – i.e., to design a scalable, inter-operable control plane for addressing and routing (intra-domain and inter-domain), topology and resource discovery, provisioning and restoration, and management of optical resources and lightpaths. In this tutorial, we will cover these control plane topics and more, focusing on mechanisms and protocols derived from IP counterparts, and importantly, *what networking people need to know about optical transport to contribute to this rapidly evolving area.*

This tutorial is a broad introduction to IP control of Optics and the developing Generalized Multi-Protocol Label Switching (GMPLS) specifications. The audience will not be expected to have any prior knowledge of GMPLS or optical networks. The first part of the tutorial will provide a detailed description of the GMPLS architecture and the protocols and mechanisms defined. In particular, the tutorial will examine the properties and requirements of optical networks and compare them to the properties and requirements of packet networks. The tutorial will also examine the impact that these different properties and needs have on the design of the GMPLS architecture and protocols. A particular emphasis will be placed on the control plane and protocols and mechanisms used for rapid protection and restoration. The second part of the tutorial will also discuss work and competing ideas related to how transport networks might best be designed to support IP, including supporting IP over a rapidly reconfigurable optical network.

### Speakers' Biographies:

JENNIFER YATES is a member of the IP Network Management and Performance Dept. at AT&T Labs – Research. She has been focusing on issues relating to IP control of optical networks. She is a co-author on numerous papers on using IP to manage optical networks and rapid shared mesh restoration. She is also extensively involved in the development of a prototype GMPLS implementation and in the OIF/IETF GMPLS-related standardization activities.

ALBERT GREENBERG heads the IP Network Management and Performance Dept. at AT&T Labs-Research. His research interests include Internet traffic measurement, modeling and engineering, network management, and optical networking. In collaboration with several others in AT&T Labs, he is developing a unified toolkit to manage IP networks. In recent years, he has also worked with Jennifer Yates and others at AT&T on the *Smart Router Simple Optics (SRSO)* project, which initially proposed IP control of optical resources, and which has produced a complete GMPLS prototype.

## Detecting Packet Patterns at High Speeds

*George Varghese*

*University of California, San Diego*

*Half Day – Day 1 PM*

### Content:

This tutorial provides an introduction to a set of old and new problems that must be solved for router line cards to operate at wire speeds. Line cards need to be able to detect important patterns (e.g., Internet lookups, packet classification, QoS enforcement, detecting Denial-of-Service Attacks and port scans, maintaining packet counters) on arriving packets. Any such processing must be completed within a packet interarrival time (8 nsec at the highest link speeds today) and hence must take a small number of memory references, and also store state in limited-size high speed memories (analogous to cache or register memory). Many of these problems may be easily solved if one had memory for each flow but the number of flows appears to be much larger than the amount of available SRAM. Therefore, these pattern detection algorithms must use a small constant number of operations and a relatively modest amount of state.

This tutorial will start with a set of models intended to introduce hardware design issues (on chip SRAMs, buses, pins, interleaved memories etc.) and a set of principles intended to help the audience think about new router problems. This tutorial is intended for students and implementers who would like to understand what can and cannot be implemented in router line cards, the state of the art for many traditional line card processing tasks, and potential new processing tasks. We will then delve into details of the best known solutions for detecting a variety of traditional router processing patterns (IP lookups, packet classification, QoS) and then move on to important new processing patterns in security (e.g., DOS attack detection, accounting). We will also briefly contrast network and traditional processor architectures. We will not cover switch design.

## Speaker's Biography:

GEORGE VARGHESE worked at DEC for several years designing DECNET protocols before obtaining his Ph.D in 1992 from MIT. He joined Washington University in 1993 as an Associate Professor where he won the ONR Young Investigator Award in 1996. He is currently a Professor at the University of California, San Diego where he works on efficient protocol implementation and protocol design. Several of the algorithms he has helped develop (e.g., IP Lookups, timing wheels, DRR) have found their way into commercial systems that range from HotMail to the Cisco GSR Router.

## How to Use the Emulab Public Network Testbeds

*Jay Lepreau, Mac Newbold, Rob Ricci, University of Utah  
Chris Alfeld, University of Wisconsin, Madison*

**Note:** This tutorial is offered as both a half-day morning lecture and a full-day version that includes an afternoon lab session. For more information see [www.emulab.net/tutorial.html](http://www.emulab.net/tutorial.html).

## Content:

**Half-Day Tutorial:** This tutorial focuses on how to effectively use Emulab (<http://www.emulab.net>), a new, public experimental resource for research, education, or development in networking and distributed systems. Over 35 research projects, 25 institutions and 2 remote networking classes have used Emulab.

The base Emulab is a time-and space-shared reconfigurable "network emulator," comprised of hundreds of PCs, miles of cable, and lots of code. Emulab allows full "root" access and lets anyone run their own custom software, including OS. Using a Web interface and the "ns" networking simulation language or a Java GUI, remote users can configure Emulab experiments with networks of arbitrary topology and link characteristics, running arbitrary programs. Experiments can use pure simulation, pure emulation, or a mixture of both. Emulab also provides the doorway to a growing set of machines physically distributed around the Internet. This recent work gives flexible access to a testbed whose links are the live Internet.

This tutorial is intended for researchers interested in new methods of experimentation, validation, or instruction. It would be helpful, but not required, to have some familiarity with the "ns" simulator. In addition to teaching how to use Emulab, some time will be spent outlining how Emulab works internally, how to build your own and how to join the distributed testbed. Finally, we will discuss how Emulab is used in education today and future plans for the project.

**Full-Day Tutorial:** Following the morning's introduction, the afternoon will primarily be a hands-on "lab session" in which attendees will pair up at laptops to use remote Emulabs. We will lead you through a series of mini-projects that exercise many Emulab features. The afternoon will also include more detail on Emulab internals and issues in building your own. *Full-Day Tutorial enrollment is limited. You are strongly encouraged to bring a laptop with a browser, ssh client and Ethernet/802.11 NIC installed. Users with a laptop will be given enrollment priority.*

## Speaker's Biographies:

JAY LEPREAU is a faculty member in the University of Utah's School of Computing, where he heads the Flux Research Group. While his primary research interest is operating systems, he is also interested in networking, distributed systems, programming languages, component-based systems, information and resource security, and even a pinch of software engineering and formal methods. In 1994, he founded the prestigious OSDI conference series and was its first program chair. He conceived and leads the Emulab project.

Jay will be joined by Mac Newbold, Rob Ricci and Chris Alfeld who are experienced instructors and key developers of the Emulab. The instructors will act as roving lab assistants during the afternoon session. In addition, an educator using Emulab remotely and a director of another Emulab will outline their experiences.

## Data-mining the Internet: What we know, what we don't, and how we can learn more?

*Michalis Faloutsos, U.C. Riverside and  
Christos Faloutsos, CMU  
Full Day – Day 2*

## Content:

What do we know about the Internet? How can we learn more about it? Despite significant research efforts, we actually know very little about the Internet. Furthermore, commonly used data analysis techniques based on averages, standard deviation and Poisson processes have exhausted their capabilities. Unfortunately, most network researchers are unaware of the wealth of data-mining and statistical analysis tools available.

First, we present the state of the art of WHAT we know about modeling and simulating the Internet. Second, we present cutting edge techniques on HOW to further our understanding of the network. The following two scenarios describe the type of questions this tutorial will answer:

- Scenario 1 (WHAT): You want to simulate your new protocol. What topology should you use? What is the distribution of sources and destinations? What kind of background traffic should you use?
- Scenario 2 (HOW): You just obtained large amount of measured data regarding round trip delays among several node pairs over a few hours. How can you characterize it? How do you compare the delays between different end-points? How do you cluster "similar" round-trip behavior? How can you identify abnormal behavior such as a distributed denial-of-service attack (DDoS)?

(speaker biographies start on next page)

## Speaker's Biographies:

MICHALIS FALOUTSOS received his Ph.D. degree in Computer Science from the University of Toronto, Canada (1999) and is currently an assistant professor at the University of California Riverside. He has received an NSF CAREER award (2000) and two major DARPA grants. His interests include Internet measurements, multicast protocols, real-time communications, and wireless networks. Michalis co-authored, with Christos and Petros Faloutsos, the highly cited paper, "On Powerlaws of the Internet Topology" (SIGCOMM'99), which has renewed interest in modeling Internet topology.

CHRISTOS FALOUTSOS received his Ph.D. degree in Computer Science from the University of Toronto, Canada and is currently a professor at Carnegie Mellon University. He has received an NSF Presidential Young Investigator Award (1989), three "best paper" awards (SIGMOD 94, VLDB 97, KDD01 (runner-up)), and four teaching awards. He has published over 100 refereed articles, one monograph, and holds four patents. His research interests include data-mining, network analysis, indexing in relational and multimedia databases.

# TECHNICAL PROGRAM

## Opening Session (Wednesday, August 21 9:00 a.m. to 10:00 a.m.)

### BGP Analysis (Wednesday, August 21 10:30 a.m. to 12:00 p.m.)

#### Understanding BGP Misconfigurations

*Ratul Mahajan, David Wetherall and Tom Anderson  
University of Washington*

#### On the Correctness of IBGP Configuration

*Timothy G. Griffin and Gordon Wilfong  
AT&T Labs-Research / Bell Laboratories, Lucent Technology*

#### Realistic BGP Traffic for Test Labs

*Olaf Maennel and Anja Feldmann  
Saarland University*

## Overlay Networks (Wednesday, August 21 1:30 p.m. to 3:00 p.m.)

#### Informed Content Delivery Across Adaptive Overlay Networks

*John Byers, Jeffrey Considine, Michael Mitzenmacher and Stanislav Rost  
Boston University / Harvard University / MIT*

#### SOS: Secure Overlay Services

*Angelos Keromytis, Vishal Misra and Dan Rubenstein  
Columbia University*

#### Internet Indirection Infrastructure

*Ion Stoica, Dan Adkins, Shelley Zhuang, Scott Shenker and Sonesh Surana  
U.C. Berkeley / ICSI*

**Congestion Control (Wednesday, August 21 3:30 p.m. to 5:00 p.m.)**

Internet Congestion Control for Future High Bandwidth-Delay Product Environments

*Dina Katabi, Mark Handley and Charles Rohrs*  
*MIT / ICSI / Tellabs*

On the Long-Run Behavior of Equation-Based Rate Control

*Milan Vojnovic and Jean-Yves Le Boudec*  
*EPFL*

Selfish Behavior and Stability of the Internet: A Game-Theoretic Analysis of TCP

*Aditya Akella, Richard Karp, Christos Papadimitriou, Srinivasan Seshan and Scott Shenker*  
*CMU / ICSI / UC Berkeley*

**Invited Talk - Mike O'Dell (Thursday, August 22 9:00 a.m. to 10:00 a.m.)**

*"Reflections on the Future of Very Large Data Networks"*

**Measuring and Simulating Networks (Thursday, August 22 10:30 a.m. to 12:00 p.m.)**

Measuring ISP Topologies with Rocketfuel

*Neil Spring, Ratul Mahajan and David Wetherall*  
*University of Washington*

Network Topology Generators: Degree-Based vs Structural

*Hongsuda Tangmunarunkit, Ramesh Govindan, Sugih Jamin, Scott Shenker and Walter Willinger*  
*USC/Information Sciences Institute / ICSI / University of Michigan / AT&T Labs-Research*

Traffic Matrix Estimation: Existing Techniques Compared and New Directions

*Alberto Medina, Nina Taft, Kave Salamatian, Supratik Bhattacharyya and Christophe Diot*  
*Boston University / Sprint / University of Paris VI*

**P2P and Multicast (Thursday, August 22 1:30 p.m. to 3:00 p.m.)**

Replication Strategies in Unstructured Peer-to-Peer Networks

*Edith Cohen and Scott Shenker*  
*AT&T Labs-Research / ICSI*

Wave and Equation Based Rate Control Using Multicast Round Trip Time

*Michael Luby, Vivek Goyal, Simon Skaria and Gavin B. Horn*  
*Digital Fountain / UC Irvine / Pulsent*

Scalable Application Layer Multicast

*Suman Banerjee, Bobby Bhattacharjee and Christopher Kommareddy*  
*University of Maryland College Park*

**Routing Dynamics (Thursday, August 22 4:30 p.m. to 5:30 p.m.)**

Route Flap Damping Exacerbates Internet Routing Convergence

*Zhuoqing Mao, Ramesh Govindan, Randy Katz and George Varghese*  
*UC Berkeley / ICSI / UC San Diego*

Route Oscillations in I-BGP with Route Reflection

*Anindya Basu, C.-H. Luke Ong, April Rasala, F. Bruce Shepherd and Gordon Wilfong*  
*Bell Laboratories / Oxford University / MIT*

**Router Design and Analysis (Friday, August 23 9:00 a.m. to 10:30 p.m.)**

Routers with a Single Stage of Buffering

*Sundar Iyer, Rui Zhang and Nick McKeown*  
*Stanford University*

Lightweight Network Support for Scalable End-to-End Services

*Kenneth L. Calvert, James Griffioen and Su Wen*  
*University of Kentucky*

On Fundamental Tradeoffs between Delay Bounds and Computational Complexity in Packet Scheduling Algorithms

*Jun Xu and Richard Lipton*  
*Georgia Institute of Technology*



**Measuring Paths and Flows (Friday, August 23 11:00 a.m. to 12:30 p.m.)**

End-to-end Available Bandwidth: Measurement Methodology, Dynamics, and Relation with TCP Throughput  
*Manish Jain and Constantinos Dovrolis*  
*University of Delaware*

On the Characteristics and Origins of Internet Flow Rates  
*Yin Zhang, Lee Breslau, Vern Paxson and Scott Shenker*  
*AT&T Labs-Research / ICSI*

New Directions in Traffic Measurement and Accounting  
*Cristian Estan and George Varghese*  
*UC San Diego*

**Position Papers (Friday, August 23 1:30 p.m. to 2:30 p.m.)**

An End-to-End Approach to Globally Scalable Network Storage  
*Micah Beck, Terry Moore and James S. Plank*  
*University of Tennessee*

Tussle in Cyberspace: Defining Tomorrow's Internet  
*David D. Clark, John Wroclawski, Karen Sollins and Robert Braden*  
*MIT / USC/Information Sciences Institute*

# CONFERENCE DETAILS

## CONFERENCE VENUE

SIGCOMM 2002 will be held at the prestigious Omni William Penn Hotel, 530 William Penn Way, Mellon Square (412.281.7100). The hotel room block is called ACM SIGCOMM 2002, and the rate is \$128.00/night.

## SIGCOMM AWARD

The SIGCOMM Award is given annually to a person whose career and technical achievements demonstrate a long-term commitment to the field of data communications. ACM SIGCOMM is pleased to announce that the 2002 SIGCOMM Award is being given to Dr. Scott Shenker of the ICSI Center for Internet Research. Dr. Shenker will receive the award and give the conference keynote address in the opening session on Wednesday.

## ACM SIGCOMM

The Association of Computing Machinery (ACM) Special Interest Group on Data Communications (SIGCOMM) is a professional forum for the discussion of topics in the field of communications and networks, including technical design and engineering, regulations and operations, and the social implications of computer networking. The SIG's members are particularly interested in the systems engineering and architectural questions of communication. For more information on membership, please contact ACM Membership Services Department at +1 212 626 0500 or via [www.acm.org](http://www.acm.org).

ACM SIGCOMM 2003 will be held in Karlsruhe, Germany on 25-29 August, 2003. Please visit [www.acm.org/sigcomm](http://www.acm.org/sigcomm) for more information.

# ORGANIZERS

## **Program Committee:**

Tom Anderson, University of Washington  
Hari Balakrishnan, MIT  
Steve Bellovin, AT&T Labs - Research  
Dave Clark, MIT  
Mark Crovella, Boston University and Network Appliance  
Jon Crowcroft, Cambridge University  
Peter Druschel, Rice University  
Anja Feldmann, University of Saarbruecken  
Matt Grossglauser, AT&T Labs - Research  
Jim Kurose, University of Massachusetts  
Craig Labovitz, Arbor Networks  
Nick McKeown, Stanford University  
Greg Minshall, Redback Networks  
Robert Morris, MIT  
Craig Partridge, BBN  
Vern Paxson, ICSI  
Jennifer Rexford, AT&T Labs - Research  
Jim Roberts, France Telecom  
Stefan Savage, University of California San Diego  
Scott Shenker, ICSI

Ion Stoica, University of California Berkeley  
George Varghese, University of California San Diego  
David Wetherall, University of Washington  
Ellen Zegura, Georgia Institute of Technology

## **Organization Committee:**

### **General Co-Chairs:**

Peter Steenkiste, Carnegie Mellon University  
Matt Mathis, Pittsburgh Supercomputing Center

### **Program Co-Chairs:**

Vern Paxson, ICSI  
Hari Balakrishnan, MIT

### **Tutorials Chair:**

Srinivasan Seshan, Carnegie Mellon University

### **Publicity Chair:**

Vivian Benton, Pittsburgh Supercomputing Center

### **Local Arrangements Chair:**

Janet Brown, Pittsburgh Supercomputing Center

### **Student Travel Grant Chair:**

Ion Stoica, UCB

### **Student Poster Session Chair:**

Karen Sollins, MIT

# TRAVEL INFORMATION

## **Travel**

You will be flying into Pittsburgh International Airport (PIT). The conference hotels are 30-35 minutes away by car. The cheapest way to get to downtown from the airport is via the Airport Flyer (28X) bus service. The Airport Flyer is run by Port Authority Transit of Allegheny County. You can board the bus in the ground transportation area outside of the baggage claim. One-way fare is \$12.00. There are also Taxis available at the airport, Yellow Cab charges a rate of \$2.00 (to step in the cab) and \$1.70 per mile. There is no flat rate.

## **Climate**

Expect the temperature in late August to be approximately 80°F and humid during the day, and approximately 60°F at night.

## **Computing Facilities**

There will be a laptop bar and wireless facilities available at the conference.

## **Student Travel Grants**

SIGCOMM 2002 is pleased to announce its Student Travel Grant Program. The purpose of the program is to encourage graduate student participation at the conference by partially or fully funding the travel costs of students who would otherwise be unable to attend. SIGCOMM 2002 thanks CISCO, NSF and SIGCOMM for funding the program this year. For more information, please go to:

<http://www.acm.org/sigcomm/sigcomm2002/travel.html>

The deadline for applications is June 21, 2002.

## **Conference Hotel**

SIGCOMM has arranged special room rates with The Omni William Penn Hotel. The room block is “**ACM SIGCOMM 2002**” and the rate is \$128.00/night.

# LOCAL ATTRACTIONS

## Local Attractions

For general information about the City of Pittsburgh, go to <http://www.city.pittsburgh.pa.us/>. Popular sights and activities in the Pittsburgh area:



- **The Andy Warhol Museum**  
117 Sandusky Street, North Side  
(412.237.8300)
- **The Carnegie Museum of Art and  
The Carnegie Museum of Natural History**  
4400 Forbes Avenue, Oakland  
(412.622.3131)
- **The Duquesne Incline**  
1220 Grandview Avenue, Mt. Washington;  
1197 West Carson Street (412.381.1665)
- **The Frick Art & Historical Center**  
7227 Reynolds Street, Point Breeze  
(412.371.0606)
- **Nationality Classrooms**  
University of Pittsburgh, Cathedral of  
Learning, Fifth & Bigelow, Oakland  
(412.724.6000)
- **Pittsburgh Zoo**  
One Wild Place (412.665.3640)
- **Phipps Conservatory**  
Oakland-University area (412.255.2376)
- **National Aviary of Pittsburgh**  
Allegheny Commons West, North Side  
(412.323.7234)
- **Station Square**  
One Station Square (412.471.5808)

**SIGCOMM 2002 THANKS THE FOLLOWING COMPANIES FOR  
THEIR SUPPORT:**

Intel Research



Microsoft Research

AOL Time Warner



Agilent Technologies

