

# NOSSDAV 2009

The 19th International Workshop on Network and Operating Systems Support for Digital Audio and Video  
Williamsburg, Virginia. June 3-5, 2009

[Home](#)

[Organization](#)

[Call For Papers](#)

[Submit Papers](#)

[Registration](#)

[Program](#)

[Scribes](#)

[Abstracts, Papers,  
Slides](#)

[Travel Logistic](#)

[Social Networks](#)



[Past NOSSDAV  
Workshops](#)

- [2008](#)
- [2007](#)
- [2006](#)
- [2005](#)
- [more..](#)

**Sponsored by:**



Association for  
Computing Machinery

*Advancing Computing as a Science & Profession*



**In cooperation with:**



SIGOPS SIGCOMM

## Announcement

- **07 Jan 2009:** [Submission site](#) is opened and running.
- **29 Jan 2009:** The submission deadline is extended to 16 February, 2009.
- **19 Mar 2009:** [Registration](#) is open.
- **20 Mar 2009:** List of [accepted papers](#) is now available.

## Important Dates

**Paper Deadline:** ~~9 February~~ 16 February, 2009, 23:59 PST (Firm)

**Notification:** 20 March, 2009

**Camera Ready Due:** 6 April, 2009

**Advance Registration  
Deadline :** ~~1 May~~ 10 May, 2009



Photos by [Jeff Kubina](#)

NOSSDAV 2009 will be held in Williamsburg, Virginia. Williamsburg is part of the Historic Triangle of Virginia, and is known for Colonial Williamsburg, the restored Historic Area of the City. Williamsburg is also home to the College of William and Mary, the second-oldest university in the United States.

NOSSDAV 2009 will continue the workshop's long tradition of focusing on emerging topics, controversial ideas, and future

research directions in the area of multimedia systems research. The workshop will also continue to be held in a setting that stimulates lively discussions among the senior and junior participants.

It is also an established practice for NOSSDAV to encourage experimental research based on real systems and data sets. Further, public availability of source code and data sets is highly encouraged.

The scope of NOSSDAV has evolved over the years. For NOSSDAV 2009, we would like to especially highlight two new topics of interest: novel use of GPU for multimedia and multi-core processors support for multimedia. NOSSDAV 2009 continues to welcome submissions in the traditional topics of networked multimedia systems, operating system support for multimedia, and multimedia security and rights management.

This website is mirrored at <http://www.nossdav.org/2009/> and <http://nemesys.comp.nus.edu.sg/nossdav09/>.

# NOSSDAV 2009

The 19th International Workshop on Network and Operating Systems Support for Digital Audio and Video  
Williamsburg, Virginia. June 3-5, 2009

[Home](#)

[Organization](#)

[Call For Papers](#)

[Submit Papers](#)

[Registration](#)

[Program](#)

[Scribes](#)

[Abstracts, Papers,  
Slides](#)

[Travel Logistic](#)

[Social Networks](#)



[Past NOSSDAV  
Workshops](#)

- [2008](#)
- [2007](#)
- [2006](#)
- [2005](#)
- [more..](#)

**Sponsored by:**



**In cooperation with:**



## Announcement

- **07 Jan 2009:** [Submission site](#) is opened and running.
- **29 Jan 2009:** The submission deadline is extended to 16 February, 2009.
- **19 Mar 2009:** [Registration](#) is open.
- **20 Mar 2009:** List of [accepted papers](#) is now available.

## Important Dates

**Paper Deadline:** ~~9 February~~ 16 February, 2009, 23:59 PST (Firm)

**Notification:** 20 March, 2009

**Camera Ready Due:** 6 April, 2009

**Advance Registration Deadline :** ~~4 May~~ 10 May, 2009



Photos by [Jeff Kubina](#)

NOSSDAV 2009 will be held in Williamsburg, Virginia. Williamsburg is part of the Historic Triangle of Virginia, and is known for Colonial Williamsburg, the restored Historic Area of the City. Williamsburg is also home to the College of William and Mary, the second-oldest university in the United States.

NOSSDAV 2009 will continue the workshop's long tradition of focusing on emerging topics, controversial ideas, and future

research directions in the area of multimedia systems research. The workshop will also continue to be held in a setting that stimulates lively discussions among the senior and junior participants.

It is also an established practice for NOSSDAV to encourage experimental research based on real systems and data sets. Further, public availability of source code and data sets is highly encouraged.

The scope of NOSSDAV has evolved over the years. For NOSSDAV 2009, we would like to especially highlight two new topics of interest: novel use of GPU for multimedia and multi-core processors support for multimedia. NOSSDAV 2009 continues to welcome submissions in the traditional topics of networked multimedia systems, operating system support for multimedia, and multimedia security and rights management.

This website is mirrored at <http://www.nossdav.org/2009/> and <http://nemesys.comp.nus.edu.sg/nossdav09/>.

# NOSSDAV 2009

The 19th International Workshop on Network and Operating Systems Support for Digital Audio and Video

Williamsburg, Virginia. June 3-5, 2009

[Home](#)

[Organization](#)

[Call For Papers](#)

[Submit Papers](#)

[Registration](#)

[Program](#)

[Scribes](#)

[Abstracts, Papers,  
Slides](#)

[Travel Logistic](#)

[Social Networks](#)



[Past NOSSDAV  
Workshops](#)

[2008](#)

[2007](#)

[2006](#)

[2005](#)

[more..](#)

[Sponsored by:](#)



Association for  
Computing Machinery

*Advancing Computing as a Science & Profession*



[In cooperation with:](#)



SIGOPS SIGCOMM

## Organization

**Co-Chair:** Dongyan Xu, Purdue University  
Wei Tsang Ooi, National University of Singapore

**Local Arrangement:** Ketan Mayer-Patel, University of North Carolina at Chapel Hill

**Publicity:** Kuan-Ta Chen, Academia Sinica, Taiwan  
Kang Li, University of Georgia, USA  
Tristan Henderson, University of St. Andrews, UK

**Technical Program Committee** Kevin Almeroth, University of California, Santa Barbara, USA  
Grenville Armitage, Swinburne University of Technology, Australia  
Ernst Biersack, Institute Eurecom, France  
Surendar Chandra, University of Notre Dame, USA  
Kuan-Ta Chen, Academia Sinica, Taiwan  
Songqing Chen, George Mason University, USA  
Mark Claypool, Worcester Polytechnic Institute, USA

Wu-chang Feng, Portland State University, USA

Wu-chi Feng, Portland State University, USA

Romulus Grigoras, IRIT-Toulouse, France

Carsten Griwodz, University of Oslo, Norway

Pål Halvorsen, University of Oslo, Norway

JongWon Kim, Gwangju Institute of Science and Technology, Korea

Baochun Li, University of Toronto, Canada

Kang Li, University of Georgia, USA

Yong Liu, Polytechnic University, USA

Andreas Mauthe, Lancaster University, UK

Klara Nahrstedt, University of Illinois at Urbana-Champaign, USA

Ketan Mayer-Patel, University of North Carolina at Chapel Hill, USA

Cristina Nita-Rotaru, Purdue University, USA

Sanjay Rao, Purdue University, USA

Reza Rejaie, University of Oregon, USA

Karsten Schwan, Georgia Institute of Technology, USA

Henning Schulzrinne, Columbia University, USA

Prashant Shenoy, University of Massachusetts, Amherst, USA

Shervin Shirmohammadi, University of Ottawa, Canada

Ralf Steinmetz, TU Darmstadt, Germany

Lars Wolf, TU Braunschweig, Germany

Zhenyu Yang, Florida International University

Roger Zimmermann, National University of

Singapore

This website is mirrored at <http://www.nossdav.org/2009/> and <http://nemesys.comp.nus.edu.sg/nossdav09/>.

---

# NOSSDAV 2009

The 19th International Workshop on Network and Operating Systems Support for Digital Audio and Video

Williamsburg, Virginia. June 3-5, 2009

[Home](#)

[Organization](#)

[Call For Papers](#)

[Submit Papers](#)

[Registration](#)

[Program](#)

[Scribes](#)

[Abstracts, Papers,  
Slides](#)

[Travel Logistic](#)

[Social Networks](#)



[Past NOSSDAV  
Workshops](#)

[2008](#)  
[2007](#)  
[2006](#)  
[2005](#)  
[more..](#)

[Sponsored by:](#)



Association for  
Computing Machinery

*Advancing Computing as a Science & Profession*



[In cooperation with:](#)



SIGOPS SIGCOMM

## Call For Papers

NOSSDAV 2009 will continue the workshop's long tradition of focusing on emerging topics, controversial ideas, and future research directions in the area of multimedia systems research, held in a setting that stimulates lively discussions among the senior and junior participants.

It is also an established practice for NOSSDAV to encourage experimental research based on real systems and data sets. Public availability of source code and data sets is highly encouraged.

The scope of NOSSDAV has evolved over the years. For NOSSDAV 2009, we would like especially highlight two new topics of interest: unconventional use of GPU for multimedia and multi-core processors support for multimedia. NOSSDAV 2009 continues to welcome submissions in the traditional topics of networked multimedia systems, operating system support for multimedia, and multimedia security and rights management.

The topics of interest include, but are not limited to:

- OS, middleware, and network support for multimedia
- Overlay networks for multimedia
- Media streaming, distribution, and storage
- Web 2.0 systems and social networks for multimedia
- Wireless and mobile multimedia systems
- Media sensor and ad hoc networks
- Grid/cloud computing support for multimedia
- Embedded systems for multimedia
- Multicore architecture support for multimedia
- GPU for multimedia
- Network processor support for multimedia
- Tele-immersion and mixed-reality systems
- Networked graphics and virtual environments

- Networked games
- Multimedia communications and system security
- Digital rights management

A broad view will be taken in deciding what topics are within scope. Please contact the workshop co-chairs if you are unsure and wish to check if a particular topic is within the scope of NOSSDAV.

Submissions should be at most SIX pages in length using standard ACM proceedings style. We expect these submissions to be the kernel of what will eventually lead to full-length papers at high-quality conferences or journals.

Authors of selected, high quality papers from NOSSDAV 2009 will be invited to submit an extended version of their papers to a special issue of *ACM/Springer Multimedia Systems Journal (MMSJ)*.

### Important Dates

**Paper Deadline:** ~~9 February~~, 16 February, 2009, 23:59 PST (Firm)

**Notification:** 20 March, 2009

**Camera Ready Due:** 6 April, 2009

**Workshop:** 3-5 June, 2009

Download this CFP as [txt](#) or [pdf](#)

This website is mirrored at <http://www.nossdav.org/2009/> and <http://nemesys.comp.nus.edu.sg/nossdav09/>.



# NOSSDAV 2009

The 19th International Workshop on Network and Operating Systems Support for Digital Audio and Video  
Williamsburg, Virginia. June 3-5, 2009

[Home](#)  
[Organization](#)

[Call For Papers](#)  
[Submit Papers](#)

[Registration](#)  
[Program](#)  
[Scribes](#)  
[Abstracts, Papers,  
Slides](#)  
[Travel Logistic](#)

[Social Networks](#)  


[Past NOSSDAV  
Workshops](#)  
[2008](#)  
[2007](#)  
[2006](#)  
[2005](#)  
[more..](#)

**Sponsored by:**



**In cooperation with:**



**Click to register for NOSSDAV 2009**

## Registration Deadline

Advance Registration Deadline: ~~1 May~~, 10 May, 2009.

## Registration Fees

There are two types of registrations.

- **Regular:** Each paper presented at the workshop must be associated with at least one regular registration.
- **Student:** Full time students are eligible for this rate. Student (co-)authors are eligible for this rate after one regular rate registration is associated with their paper.

Special rates are available to ACM members (for regular registration) and to SIGMM student members (for student registration). Membership must be valid at the time of registration and membership number must be provided during registration in order to register using the special rate.

Each registration includes coffee breaks, lunches, banquet, and electronic workshop proceedings. For non-SIGMM student member, a one-year SIGMM student membership is included.

Additional banquet tickets are available for purchase at USD 50 per ticket.

		<b>Advance</b>		<b>Late/On-Site</b>	
Student	SIGMM Member	USD 180	USD 230		
	Otherwise	USD 200	USD 250		
Regular	ACM Member	USD 460	USD 510		
	Otherwise	USD 510	USD 560		

This website is mirrored at <http://www.nossdav.org/2009/> and <http://nemesys.comp.nus.edu.sg/nossdav09/>.

---

# NOSSDAV 2009

The 19th International Workshop on Network and Operating Systems Support for Digital Audio and Video  
Williamsburg, Virginia. June 3-5, 2009

[Home](#)

[Organization](#)

[Call For Papers](#)

[Submit Papers](#)

[Registration](#)

[Program](#)

[Scribes](#)

[Abstracts, Papers,  
Slides](#)

[Travel Logistic](#)

[Social Networks](#)



[Past NOSSDAV  
Workshops](#)

[2008](#)  
[2007](#)  
[2006](#)  
[2005](#)  
[more..](#)

**Sponsored by:**



Association for  
Computing Machinery

*Advancing Computing as a Science & Profession*



**In cooperation with:**



SIGOPS SIGCOMM

## Technical Program

---

### 03 June 2009

1800 **Registration Opens**

1900 **Reception**

---

### 04 June 2009

0800 **Registration and Breakfast**

0830 - 0845 **Welcome Address**

0845 - 1000 **Keynote Session**

- **Experiences Building Telepresence Systems**  
**Keynote Speaker: Prof. Henry Fuchs**

Federico Gil Professor of Computer Science  
Adjunct Professor of Biomedical Engineering

University of North Carolina at Chapel-Hill

1000 - 1030 **Coffee Break**

1030 - 1200 **Session 1: New Applications and Services**

*Discussion Lead: Wei Tsang Ooi (National University of Singapore)*

*Scribe: Amir Hassan Rasti Ekbatani (University of Oregon)*

- **Dynamic Overlay Multicast in Multi-Stream/Multi-Site 3D Collaborative Systems**

*Wanmin Wu (University of Illinois at Urbana-Champaign)*

*Zhenyu Yang (Florida International University)*

*Klara Nahrstedt (University of Illinois at Urbana-Champaign)*

- **Web 2.0 Traffic Measurement: Analysis on Online Map Applications**

*Song Lin (Tsinghua University)*

*Zhiguo Gao (IBM China Research Laboratory)*

*Ke Xu (Tsinghua University)*

- **Peer-Assisted Online Storage and Distribution: Modeling and Server Strategies**

*Ye Sun (Hong Kong University of Science & Technology)*

*Fangming Liu (Hong Kong University of Science & Technology)*

*Bo Li (Hong Kong University of Science & Technology)*

*Baochun Li (University of Toronto)*

1200 - 1330 **Lunch**

1330 - 1500 **Session 2: P2P Streaming I**

*Discussion Lead: Ketan Mayer-Patel (University of North Carolina at Chapel Hill)*

*Scribe: Remo Meier (ETH Zurich)*

- **InstantLeap: Fast Neighbor Discovery in P2P VoD Streaming**

*Xuanjia Qiu (Sun Yat-Sen University)*

*Chuan Wu (The University of Hong Kong)*

*Xiaola Lin (Sun Yat-Sen University)*  
*Francis C. M. Lau (The University of Hong Kong)*

- **Overlay Monitoring and Repair in Swarm-based Peer-to-Peer Streaming**

*Nazanin Magharei (University of Oregon)*  
*Reza Rejaie (University of Oregon)*

- **Adaptive Overlay Topology for Mesh-Based P2P-TV Systems**

*Richard John Lobb (University of Canterbury)*  
*Ana Paula Couto da Silva (Federal University of Juiz de Fora)*  
*Emilio Leonardi (Politecnico di Torino)*  
*Marco Mellia (Politecnico di Torino)*  
*Michela Meo (Politecnico di Torino)*

1500 - 1530 **Coffee Break**

1530 - 1730 **Session 3: OS and End Systems**

*Discussion Lead: Kevin Almeroth (UC Santa Barbara)*

*Scribe: Ishan Vaishnavi (Centrum voor Wiskunde en Informatica)*

- **Random Network Coding on the iPhone: Fact or Fiction?**

*Hassan Shojania (University of Toronto)*  
*Baochun Li (University of Toronto)*

- **SLIPstream: Scalable Low-latency Interactive Perception on Streaming Data**

*Padmanabhan S. Pillai (Intel Research Pittsburgh)*  
*Lily B. Mummert (Intel Research Pittsburgh)*

*Steven W. Schlosser (Intel Research  
Pittsburgh)*

*Rahul Sukthankar (Intel Research  
Pittsburgh)*

*Casey J. Helfrich (Intel Research  
Pittsburgh)*

- **Server-Efficient High-Definition  
Media Dissemination**

*Philip W. Frey (IBM Research GmbH)*

*Andreas Hasler (IBM Research GmbH)*

*Bernard Metzler (IBM Research GmbH)*

*Gustavo Alonso (ETH Zurich)*

- **Power Efficient Real-Time Disk  
Scheduling**

*Damien Le Moal (Hitachi Ltd.)*

*Donald Molaro (Hitachi Global Storage  
Technologies, San Jose Research Center)*

*Jorge Campello (Hitachi Global Storage  
Technologies, San Jose Research Center)*

1900

## **Banquet**

---

# **05** June 2009

0800

## **Registration and Breakfast**

0830 - 1000

## **Session 4: Virtual Environments and Games**

*Discussion Lead: Kuan-Ta Chen (Academia  
Sinica)*

*Scribe: Pengpeng Ni (Simula Research Lab and  
University of Oslo)*

- **A Delaunay Triangulation  
Architecture Supporting Churn and  
User Mobility in MMVE-s**

*Mohsen Ghaffari (Sharif University of  
Technology)*

*Behnoosh Hariri (Sharif University of*

*Technology and University of Ottawa)*  
*Shervin Shirmohammadi (Sharif University*  
*of Technology and University of Ottawa)*

- **Probabilistic Event Resolution with the Pairwise Random Protocol**

*John L. Miller (Microsoft Research and*  
*University of Cambridge)*  
*Jon Crowcroft (University of Cambridge)*

- **Cross-Tree Adjustment for Spatialized Audio Streaming over Networked Virtual Environments**

*Ke Liang (National University of Singapore)*  
*Roger Zimmermann (National University of*  
*Singapore)*

1000 - 1030 **Coffee Break**

1030 - 1200 **Session 5: Security**

*Discussion Lead: Klara Nahrstedt (UIUC)*

*Scribe: Liang Ke (National University of*  
*Singapore)*

- **End-to-End Secure Delivery of Scalable Video Streams**

*Kianoosh Mokhtarian (Simon Fraser*  
*University)*  
*Mohamed Hefeeda (Simon Fraser*  
*University)*

- **Secure Multimedia Content Delivery with Multiparty Multilevel DRM Architecture**

*Tony Thomas (Nanyang Technological*  
*University)*  
*Sabu Emmanuel (Nanyang Technological*  
*University)*  
*Amitabha Das (Nanyang Technological*  
*University)*

*Mohan S. Kankanhalli (National University of Singapore)*

- **Rapid Identification of Skype Traffic**

*Philip A. Branch (Swinburne University of Technology)*

*Amiel Heyde (Swinburne University of Technology)*

*Grenville J. Armitage (Swinburne University of Technology)*

1200 - 1330 **Lunch**

1330 - 1500 **Session 6: Understanding and Improving User Experience**

*Discussion Lead: Mohamed Hafeeda (Simon Fraser University)*

*Scribe: Wei Cheng (National University of Singapore)*

- **An Empirical Evaluation of VoIP Playout Buffer Dimensioning in Skype, Google Talk, and MSN Messenger**

*Chen-Chi Wu (National Taiwan University)*

*Kuan-Ta Chen (Academia Sinica)*

*Chun-Ying Huang (National Taiwan Ocean University)*

*Chin-Laung Lei (National Taiwan University)*

- **Fine-Grained Scalable Streaming from Course-Grained Video**

*Pengpeng Ni (Simula Research Laboratory and University of Oslo)*

*Alexander Eichhorn (Simula Research Laboratory)*

*Carsten Griwodz (Simula Research Laboratory and University of Oslo)*

*Pål Halvorsen (Simula Research Laboratory and University of Oslo)*



- **Estimate and Serve: Scheduling Soft Real-Time Packets for Delay Sensitive Media Applications on the Internet**

*Ishan Vaishnavi (Centrum voor Wiskunde en Informatica)*

*Dick C. A. Bulterman (Centrum voor Wiskunde en Informatica)*

1500 - 1530 **Coffee Break**

1530 - 1700 **Session 7: P2P Streaming II**

*Discussion Lead: Dongyan Xu (Purdue University)*

*Scribe: Saurabh Ratti (University of Ottawa)*

- **Zebroid: IPTV Data Mining to Support Peer-Assisted VoD Content Delivery**

*Yih-Farn Robin Chen (AT&T Laboratories - Research)*

*Rittwik Jana (AT&T Laboratories - Research)*

*Daniel Stern (AT&T Laboratories - Research)*

*Bin Wei (AT&T Laboratories - Research)*

*Mike Yang (AT&T Laboratories - Research)*

*Hailong Sun (Beihang University)*

- **Robust Live Media Streaming in Swarms**

*Thomas Locher (ETH Zurich)*

*Remo Meier (ETH Zurich)*

*Roger Wattenhofer (ETH Zurich)*

*Stefan Schmid (TU Munich)*

- **Providing Statistically Guaranteed Streaming Quality for Peer-to-Peer Live Streaming**

*Miao Wang (University of Nebraska-Lincoln)*

*Lisong Xu (University of Nebraska-Lincoln)*

*Byrav Ramamurthy (University of Nebraska-  
Lincoln)*

1700 - 1715 **Concluding Remarks**

This website is mirrored at <http://www.nossdav.org/2009/> and <http://nemesys.comp.nus.edu.sg/nossdav09/>.

---

# NOSSDAV 2009

The 19th International Workshop on Network and Operating Systems Support for Digital Audio and Video  
Williamsburg, Virginia. June 3-5, 2009

[Home](#)  
[Organization](#)

[Call For Papers](#)  
[Submit Papers](#)

[Registration](#)  
[Program](#)  
[Scribes](#)  
[Abstracts, Papers,  
Slides](#)

[Travel Logistic](#)

[Social Networks](#)



[Past NOSSDAV  
Workshops](#)

[2008](#)  
[2007](#)  
[2006](#)  
[2005](#)  
[more..](#)

[Sponsored by:](#)



Association for  
Computing Machinery

*Advancing Computing as a Science & Profession*



[In cooperation with:](#)



SIGOPS SIGCOMM

## Scribes From NOSSDAV 2009

---

# 04 June 2009

## Keynote Session

- **Experiences Building Telepresence Systems**  
**Keynote Speaker: Prof. Henry Fuchs**

Federico Gil Professor of Computer Science  
Adjunct Professor of Biomedical Engineering  
University of North Carolina at Chapel-Hill

*Scribe: Wei Cheng (National University of Singapore)*

Professor Henry Fuchs from University of North Carolina at Chapel Hill gave the keynote speech titled "Experiences Building Tele-presence System". The keynote first explained the term tele-presence, of which the main aim of is to give users convincing immersive experience. Henry then introduced tele-presence systems that have been developed since 1997, underlining the main challenges in developing such systems. The keynote concluded by highlighting two key system challenges for immerse environment: (i) a middleware, and (ii) a set of integrated OS components, for common, lower level tele-presence functions, such as display management.

The keynote was interleaved with many questions from the audience. Kevin Almeroth (UC Santa Barbara) asked why a new term ``tele-presence" is coined (Answer from Henry: to emphasize human "sense"). Yih-Fan Chen (Robin, AT and T Research) argued that boundary exists naturally between the virtual world and the real world and true immersion cannot be achieved, to which Henry

agreed. Current solutions to virtual world are either not fully immersive or require wearing of special devices.

The role of using large displays in tele-presence systems were discussed with questions from Dick Bulterman (CWI) and Damien Le Moal (Hitachi). Henry pointed out that using large display helps users in knowing the relative position and context within the data displayed, and allows users to move around more freely. Dick also asked about the relative importance of audio and visual in tele-presence system. Henry pointed out that the importance of audio/visual information is a psychological issue and depends on the applications.

Damien asked why OS integration is important in large display system when we can just develop dedicate applications. Henry argued that system tasks (such as windows management) need to be adapted to the new display system, so only an OS supporting the large display system can give users uniform experience in both system tasks and dedicate applications.

Padmanabhan Pillai (Intel Research) and Damien asked about the role of haptic and smell. Henry answered that the sense of touch and smell are important in some cases (e.g., medical applications) but the current technology is limited. Haptic devices confine movement to a small distance, and the basis of smell (like RGB for color) are yet to be found.

## **Session 1: New Applications and Services**

*Discussion Lead: Wei Tsang Ooi (National University of Singapore)*

*Scribe: Saurabh Ratti (University of Ottawa)*

- **Dynamic Overlay Multicast in Multi-Stream/Multi-Site 3D Collaborative Systems**

*Wanmin Wu (University of Illinois at Urbana-Champaign)*

*Zhenyu Yang (Florida International University)*

*Klara Nahrstedt (University of Illinois at Urbana-Champaign)*

Wanmin Wu from UIUC presented a 3D video collaborative

system that uses an P2P overlay for dissemination of content streams. Each peer has the capability of producing and rendering multiple 3D video streams (depth as the added dimension). The main contribution of this work is an algorithm that accounts for the contributing factor of streams to the overall view when selecting peers. Experimental results demonstrated that the algorithm performs more efficiently when compared to using view proximity only.

The ensuing discussion on 3D video capturing systems revealed that the current camera technology is very limiting in capturing larger areas, as those currently being developed are being aimed at desktop usage, with ranges of 1 - 3 meters. Wei Tsang Ooi (National University of Singapore) asked about the evaluation of visual quality of 3D video, to which Wanmin replied that designing an objective visual quality metric for 3D video remains an open challenge.

- **Web 2.0 Traffic Measurement: Analysis on Online Map Applications**

*Song Lin (Tsinghua University)*

*Zhiguo Gao (IBM China Research Laboratory)*

*Ke Xu (Tsinghua University)*

Song Lin (Tsinghua University) presented analysis of characteristics of captured traffic generated by four online map applications. The findings of note are (i) the traffic follows diurnal cycles, (ii) the majority of map application traffic consists of image traffic (68.7% for Sogou), (iii) 40% of Google Maps's traffic is from mash-up sources. Due to map hotspots and the applicability of Zipf's law, the study suggests that CDN servers can improve the performance of online map application at these hotspots, especially when the hotspots would differ according to physical location of users.

Wei Cheng (National University of Singapore) asked whether there was any data supporting the conclusion that requests from different physical locations have different map hotspots. Lin replied that data supporting this is

difficult to gather, but the conclusion can be inferred from the visual overlay of requests distribution on a geographical map. Wei Tsang inquired what could online map service providers learned from this work. Lin mentioned two lessons. First, providing good API is critical in encouraging other website to mash-up the map data and drive traffic to the map server. Second, CDN and other caching methods are still very relevant in such online applications.

- **Peer-Assisted Online Storage and Distribution: Modeling and Server Strategies**

*Ye Sun (Hong Kong University of Science & Technology)*

*Fangming Liu (Hong Kong University of Science & Technology)*

*Bo Li (Hong Kong University of Science & Technology)*

*Baochun Li (University of Toronto)*

Yuen Feng (University of Toronto) presented, on behalf of the authors, work on modeling a real world peer-assisted online storage system called FS2You, which has become popular in China with 500 GB to 1 TB of data transferred daily. The goal of the work was to derive a lower bound for the average file download time and evaluate the performance of different bandwidth allocation strategies. Using average file download time and peer satisfaction level as metrics for evaluation, it was found that FS2You's current strategy of allocating bandwidth inversely to file popularity has favorable performance, but not when bandwidth is unconstrained.

Wei Cheng asked whether the file popularity in FS2You follows the Zipf distribution, to which Feng replied that it does. Robin asked about the assumption that peers have uniform upload/download rate. Feng explained that this assumption is a simplification to the model and is an area for future work.

## **Session 2: P2P Streaming I**

*Discussion Lead: Ketan Mayer-Patel (University of North Carolina at Chapel Hill)*

*Scribe: Remo Meier (ETH Zurich)*

Instead of the usual format where each presentation is followed by a Q and A session, Session 2 is organized as a mini-panel, where each presenter takes turn to present their work, followed by a 30 minutes discussion on the work.

- **InstantLeap: Fast Neighbor Discovery in P2P VoD Streaming**

*Xuanjia Qiu (Sun Yat-Sen University)*

*Chuan Wu (The University of Hong Kong)*

*Xiaola Lin (Sun Yat-Sen University)*

*Francis C. M. Lau (The University of Hong Kong)*

Hassan Shojanian (University of Toronto) presented the paper on behalf of the authors. An on-demand P2P streaming systems have to handle peers with different playback positions and, consequently, a low block overlap among these peers. Hassan described InstantLeap, a mechanism to quickly discover appropriate peers for a given playback position, improving current systems with a typical  $O(\log(\#\text{peers}))$  communication overhead to  $O(1)$  with high probability using a lightweight indexing protocol. The protocol groups peers by their playback locality into segments. Peers then maintain some connections to other peers from the same segment and to some peers from other segments using random exchanges of known peers.

- **Overlay Monitoring and Repair in Swarm-based Peer-to-Peer Streaming**

*Nazanin Magharei (University of Oregon)*

*Reza Rejaie (University of Oregon)*

Nazanin Magharei (University of Oregon) presented her work on clustering of peers and their impact in peer-to-peer live streaming systems using a random mesh as P2P overlay. Clusters occur, for example, due to good connectivity among peers having the same Internet Service Provider (ISP). Peers in a clusters may suffer from poor content availability in their neighborhood as the cluster's peers are mostly connected among each other and only to few outside peers. The authors provide a

mechanism to detect and break-up such clusters. The proposed protocol groups peers by their distance to the source and increases the number of connections to other peers and swaps parents in case of problems. This allows reduction of the buffer size, to improve playback quality, and to reduce the distance to the source.

- **Adaptive Overlay Topology for Mesh-Based P2P-TV Systems**

*Richard John Lobb (University of Canterbury)*

*Ana Paula Couto da Silva (Federal University of Juiz de Fora)*

*Emilio Leonardi (Politecnico di Torino)*

*Marco Mellia (Politecnico di Torino)*

*Michela Meo (Politecnico di Torino)*

Michela Meo (Politecnica di Torino) presented their work addressing the problems of heterogeneous bandwidth availability in P2P live streaming systems. Like in the preceding talks of this session, a random mesh is used as overlay. The proposed protocol better exploits high-bandwidth peers by increasing their connectivity and by positioning them closer to the source. This way, the high-bandwidth peers provide a backbone to quickly disseminate new data blocks, benefiting all the peers. The bandwidth of a peer is indirectly estimated based on the fraction of used neighbor connections and without the peer having to explicitly know or measure its bandwidth.

- **Discussion**

A short panel discussion with all the presenters was held at the end of this session. The main topic of the discussion was the locality-awareness of the proposed protocol, i.e., whether they are aware of the underlying physical topology in order to avoid long distance connections and to keep most traffic within the same ISP or close-by ISPs. Ketan Mayer-Patel (UNC) compared the two work presented by Michela and Nazanin -- one tries to bring structure into the overlay while the other tries to remove structure. Mohamed Hefeeda pointed out that increasing randomness in overlay topology seems to be counter



intuitive. Nazanin explained that, without random links among peers, content might not flow to other ISPs.

Padmanabhan asked how else can P2P applications become location aware if ISP information is not available. The audience and presenters mentioned several alternatives, including Network Weather Service, CDN-based Relative Network Positioning, and the P4P project.

Padmanabhan also asked whether the overlays account for correlated churn, e.g, when a large number of peers join and leave simultaneously. Such churn, however, are not addressed by the papers.

### Session 3: OS and End Systems

*Discussion Lead: Kevin Almeroth (UC Santa Barbara)*

*Scribe: Ishan Vaishnavi (Centrum voor Wiskunde en Informatica)*

- **Random Network Coding on the iPhone: Fact or Fiction?**

*Hassan Shojania (University of Toronto)*

*Baochun Li (University of Toronto)*

Hassan Shojania presented their experience of implementing random network coding on the Apple iPhone and iPod Touch platform. The presentation overviewed network coding technologies and the challenges encountered in migrating the algorithm on a mobile device, with special attention to the lack of SIMD instruction set on the iPod. The presentation compared the coding efficiency over iPod Touch and iPhone, and presented the results of the feasibility study. The presentation also compared CPU-based and a GPU-based implementations.

Robin inquired about the effects of network coding on battery usage. Hassan replied that the interface to power management on the current iPhone is not exposed. He added that the CPU usage was, however, non linear. Hassan said that they were looking forward to the new processor in the forthcoming iPhone, which he believed had better SIMD instruction set support.

- **SLIPstream: Scalable Low-latency Interactive Perception on Streaming Data**

*Padmanabhan S. Pillai (Intel Research Pittsburgh)*

*Lily B. Mummert (Intel Research Pittsburgh)*

*Steven W. Schlosser (Intel Research Pittsburgh)*

*Rahul Sukthankar (Intel Research Pittsburgh)*

*Casey J. Helfrich (Intel Research Pittsburgh)*

Padmanabhan started his presentation with the question "What can one do in a room with 1000 cameras?" The presentation then focused on the scale of data and the amount of processing required to extract some constructive information from data from these cameras. The core of the presentation was a runtime system called Sprout, used to identify and split parallel tasks across hundreds of machines with a final goal of achieving low latency for interactive applications with large data processing. The presentation focused on processing event/gesture recognition in videos in "interactive time" and identified three steps in the process: (i) identification of low-level features, (ii) matching low-level data against a training set to extract events of interest, (iii) aggregation of these results into appropriate events/gestures. Padmanabhan noted that the second step involves a sweep over the stream in both time and space, which can be parallelized (even within each frame), providing significant reduction in latency.

Klara Nahrstedt (UIUC) asked how close they were to achieving interactive time. Padmanabhan replied in short "very far" and said that the acquiring time from the cameras was huge since they used off-the-shelf systems. The presenter also added that the networks are not optimized as yet to handle such sudden bursts of packets from multiple synchronized cameras. Dongyan Xu then inquired if they had looked into any compiler techniques and generic code optimization techniques. Padmanabhan said that they had not done any low-level parallelization and had focused on parallelization at a higher level.

- **Server-Efficient      High-Definition      Media**

## Dissemination

*Philip W. Frey (IBM Research GmbH)*

*Andreas Hasler (IBM Research GmbH)*

*Bernard Metzler (IBM Research GmbH)*

*Gustavo Alonso (ETH Zurich)*

Philip W. Frey from IBM Research started with an assumption that "the network is no longer the bottleneck, but the memory to NIC bus is". The presentation focused on reducing the data fetch request processing time on server by removing the current 4-way copy (two DMA copies and two CPU copies) and two context switches required to send data to an incoming request. The authors proposed new protocol based on RDMA (Remote Direct Memory Access). The presentation showed the CPU usage in HTTP (with and without kernel `sendfile()` support) and RTP-based servers. They then presented their new client-driven protocol based on RDMA. Their impressive results showed a zero CPU load using the protocol. They also presented results showing lower context switches and lower interrupt rates. For live-streaming applications, direct memory transfer made their solution a zero-copy solution.

There are several clarifying questions from the audience. Philip explained that RDMA works over TCP and the address space used in the experiments is virtual. A follow-up question was on the DMA transfer that maybe required for virtual memory access. The speaker replied that such memory access is a one-time cost.

- **Power Efficient Real-Time Disk Scheduling**

*Damien Le Moal (Hitachi Ltd.)*

*Donald Molaro (Hitachi Global Storage Technologies, San Jose Research Center)*

*Jorge Campello (Hitachi Global Storage Technologies, San Jose Research Center)*

Damien Le Moal (Hitachi Development Laboratory, Japan) spoke about a new method of optimizing disk access in an effort to save power consumption. The presentation distinguished between normal data access and audio-

visual (AV) data access targeted for set-top boxes, but general enough for use in normal computing systems. The idea is to have a separate file system for AV access, where each AV data request would be associated with a deadline it must meet. Thus, the system can hold all such requests in a queue and schedule them together in a more efficiently re-ordered fashion (according to seek distances). This increases the efficiency of the system by reducing the seek distance between two consecutive requests. The disk can also be spun down when the deadlines are far enough.

Mohamed Hefeeda inquired if scheduling these real time events in such a manner would adversely affect interactivity of multimedia application. Damien replied that as long as the deadlines are selected correctly, it will actually improve the interactivity since the event is now certain to be scheduled before its deadline expires.

---

## 05 June 2009

### Session 4: Virtual Environments and Games

*Discussion Lead: Kuan-Ta Chen (Academia Sinica)*

*Scribe: Pengpeng Ni (Simula Research Lab and University of Oslo)*

- **A Delaunay Triangulation Architecture Supporting Churn and User Mobility in MMVE-S**

*Mohsen Ghaffari (Sharif University of Technology)*

*Behnoosh Hariri (Sharif University of Technology and University of Ottawa)*

*Shervin Shirmohammadi (Sharif University of Technology and University of Ottawa)*

Saurabh Ratti (University of Ottawa) presented, on behalf of the authors, a distributed algorithm for dynamic construction of an overlay topology that supports greedy routing in massively multi-user virtual environments. The idea is to partition the topology into two non-overlapping

sets, where each set is updated in a separate phase using the other set as a reference. After explaining the theoretical background, Saurabh illustrated the algorithm steps and concluded with simulation results.

Dongyan Xu asked how to ensure the robustness of the algorithm against when nodes leave. Saurabh answered that the reliability would be guaranteed if TCP is used. If UDP was used, node leave would be detected in the next update circle. Robin suggested installing a scanner in each geometrical region for detecting the mobility of avatars. Saurabh argued that a decentralized solution is more preferable than centralized approach. Wei Tsang asked if there is any drawback of the algorithm. The response from Saurabh was that the resumption time presented in the simulation results is considerable. Although increasing the number of nodes has little effect on the resumption time, the baseline value of 1.1 seconds is not optimal, especially if the game has multiple updates within a second.

- **Probabilistic Event Resolution with the Pairwise Random Protocol**

*John L. Miller (Microsoft Research and University of Cambridge)*

*Jon Crowcroft (University of Cambridge)*

John Miller (MSR and University of Cambridge) proposed the use of Pairwise Random protocol (PRP), which uses secure coin flipping to fairly determine the resolution of competition between two avatars in distributed virtual environments. He started his talk with an introduction to some related DVE security researches and explained subsequently the three-way message exchange included in PRP protocol. Performance analysis of three different variations of PRP was presented. John concluded that the choice of the PRP variations was a trade-off between overhead and security compromise.

In response to Dongyan Xu's question about the reliability of the proposed approach, John said that the PRP protocol assumed reliable transmission between the avatars. Wei Tsang asked about the possibility of extending the PRP

protocol to more than two parties. The response was that it was possible to decompose the message exchanges between multiple parties, but the extension would be challenging due to the increased complexity and performance overhead. Ardalan Kangarlou (Purdue University) asked whether there was any mechanism that could validate the hash value used in PRP. John answered the exchanged message itself can be used to prove the correctness of the delivered information.

- **Cross-Tree Adjustment for Spatialized Audio Streaming over Networked Virtual Environments**

*Ke Liang (National University of Singapore)*

*Roger Zimmermann (National University of Singapore)*

Ke Liang presented an algorithm that constructs and adjusts multicast trees for spatialized audio streaming in a peer-to-peer overlay network. The algorithm has two objectives: Maximize total number of audio receivers, subject to nodes' bandwidth limits; while minimizing the average latency of those receivers. Ke interpreted their solution to the problem as achieving a compromise between the two objectives via maximizing the number of receiver that have the minimum latency. Extending from their previous work, Ke proposed the cross-tree adjusting (CTA) algorithm that can re-allocate the upload bandwidth incrementally for nodes with bandwidth conflicts in all existing multicast trees. Simulation results showed that CTA can achieve good performance in terms of fairness and the total number of successful receivers that have low latency.

Klara Nahrstedt asked if the cocktail effect had been considered in the proposed algorithm. The cocktail party effect describes the human's ability to focus one's listening attention on a single talker among a mixture of conversations and background noises. Ke agreed on that the cocktail party effect could be something for further investigation, although it was not covered by their current work. A suggestion to their future work from Klara was that the distance between avatars in the virtual world

could be taken into account when constructing the multicast trees.

## Session 5: Security

*Discussion Lead: Klara Nahrstedt (UIUC)*

*Scribe: Liang Ke (National University of Singapore)*

- **End-to-End Secure Delivery of Scalable Video Streams**

*Kianoosh Mokhtarian (Simon Fraser University)*

*Mohamed Hefeeda (Simon Fraser University)*

Mohamed Hefeeda presented their work on the problem of securing the delivery of scalable video streams so that receivers can ensure the authenticity (originality and integrity) of the video. Their focus in the paper is on recent scalable video coding techniques, e.g., H.264/SVC, which can provide three scalability types at the same time: temporal, spatial, and quality (or PSNR). Mohamed presented an efficient authentication scheme that accounts for the full scalability of video streams, enabling verification of all possible substreams that can be extracted and decoded from the original stream.

Dick Bulterman (CWI) wondered if there is any standard fee needed to be charged to scalable video coding techniques. The presenter responded that they are using the open source standard and there are a lot of companies (e.g. Google) that have already use the standard and the reference source code.

- **Secure Multimedia Content Delivery with Multiparty Multilevel DRM Architecture**

*Tony Thomas (Nanyang Technological University)*

*Sabu Emmanuel (Nanyang Technological University)*

*Amitabha Das (Nanyang Technological University)*

*Mohan S. Kankanhalli (National University of Singapore)*

Tony Thomas (Nanyang Technological University) presented a joint watermarking and cryptographic mechanisms for securely delivering multimedia content on multiparty multilevel digital rights management (DRM)

architecture, where content are delivered by an owner to a consumer through several levels of distributors. The presenter emphasised that license is more important than content since the content will be encrypted before delivery from the owner or distributors. The authors proposed a mechanism that minimizes the possible degradation of the quality of a content due to embedding of watermark signals. Furthermore, in case the owner or a distributor finds an unauthorized copy, they can identify the traitor with the help of a judge.

Philip Frey (IBM Research GmbH) asked what is the advantage of multi-layer design. Tony answered that the multimedia content is distributed by multiple distributors which will encrypt the contents before distribution. From the client's point of view, it could track the distributor according to the content it received. Alexander Eichhorn (Simula Research) asked why multiple layers of license servers are used. The presenter responded that it is due to business issues. Since clients may be geographically distributed, multiple distributors that may reencrypt the content are used. Thus, client need to request different licenses from different license servers.

- **Rapid Identification of Skype Traffic**

*Philip A. Branch (Swinburne University of Technology)*

*Amiel Heyde (Swinburne University of Technology)*

*Grenville J. Armitage (Swinburne University of Technology)*

In this talk, the presenter presented results of experimental work using machine learning techniques (decision tree) to rapidly identify Skype traffic. They use a number of feature classes to classify IP flows as being Skype or non-Skype flows. The feature classes they found most effective were the interarrival times between IP packets frequently occurring IP packet lengths less than 80 bytes. Their results show that using three feature classes in a single classifier provided 98 percent precision and 99 percent recall when using a window duration of 5 seconds or more. The presenter emphasizes that their classifiers do not rely on observing any particular part of a



flow. They also report on the performance of classifiers built using combinations of two of these features and of each feature in isolation.

Amir Hassan Rasti Ekbatani (University of Oregon) asked the presenter how to detect audio flows. The presenter answered that their approach works on small sequence of any part of flow. Since they make use of techniques of machine learning, they can identify Skype traffic reliably with only a few seconds of traffic, whose characteristics are extracted from short sliding windows and used to train a classifier. Amir also asked if the authors tried using larger packet size, to which Philip answered that they only used SVOPC codec so far with packet size less than 80 bytes. Michela asked what happened when Skype traffic has both audio and video data (i.e. packet size is increased). The presenter answered they still do not have a good solution since their classifier is trained from pure audio flows.

## **Session 6: Understanding and Improving User Experience**

*Discussion Lead: Mohamed Hafeeda (Simon Fraser University)*

*Scribe: Wei Cheng (National University of Singapore)*

- **An Empirical Evaluation of VoIP Playout Buffer Dimensioning in Skype, Google Talk, and MSN Messenger**

*Chen-Chi Wu (National Taiwan University)*

*Kuan-Ta Chen (Academia Sinica)*

*Chun-Ying Huang (National Taiwan Ocean University)*

*Chin-Laung Lei (National Taiwan University)*

In VoIP applications, buffer size needs to be carefully chosen to trade-off between audio quality and delay. Finding optimal buffer size is challenging because many factors are involved. According to their experimental results, real-life VoIP applications, such as Skype, MSN, and gTalk, do not adjust the buffer size well. They proposed a regression-based lightweight algorithm to compute the optimal buffer size in real time.

Wei Tsang pointed out that some VoIP applications may apply time compression to drop the silent period. He asked whether time compression might affect buffer size estimation. Kuan-Ta Chen said that estimating the buffer size according to time shifting is the only thing that can be done currently. Finding out whether time compression is used by reverse engineering may work better but needs large number of experiments. Mohamed asked which factor affects the user experience most as many factors exist. Kuan-Ta answered that it is still an open question and his personal opinion is that the codec may be the most important factor.

- **Fine-Grained Scalable Streaming from Course-Grained Video**

*Pengpeng Ni (Simula Research Laboratory and University of Oslo)*

*Alexander Eichhorn (Simula Research Laboratory)*

*Carsten Griwodz (Simula Research Laboratory and University of Oslo)*

*Pål Halvorsen (Simula Research Laboratory and University of Oslo)*

In scalable video streaming, the granularity of bit rate adaptation with CGS is limited, while using MGS brings high signal overhead. Pengpeng Ni proposed that by switching between layers quickly, arbitrary bandwidth could be achieved without high overhead. To understand and compare the user experience when different switching patterns are used, the authors conducted a user study. The result shows that frequent layer switching is effective and could be better than quality downscaling, with the switching period as a crucial design parameter.

Damien Le Moal asked whether the nominal frame rate is considered. Pengpeng answered that only 25 and 12 fps is used in user study as covering different frame rate variations in the user study is time-consuming. Saurabh asked whether the consistency of the user response is considered. Pengpeng answered that during binary measurement, two sequence may appear twice with different order (A, B and B, A), so it could be used to test

the consistency of user response. Padmanabhan asked whether experiments are done on CRT since LCD, which has lower refresh rate, may reduce the negative effect of low frame rate. Pengpeng pointed out since the field study is done on iPhone, only LCD is used. Mohamed asked whether the lighting condition is considered and whether the experiments are done following the ITU regulation. According to Pengpeng, they find field study more interesting than experiments in a controlled lab since the former is closer to the real world usage.

- **Estimate and Serve: Scheduling Soft Real-Time Packets for Delay Sensitive Media Applications on the Internet**

*Ishan Vaishnavi (Centrum voor Wiskunde en Informatica)*

*Dick C. A. Bulterman (Centrum voor Wiskunde en Informatica)*

Ishan highlighted that scheduling at intermediate nodes is essential in soft real-time applications. After he briefly introduced the traditional IntServ and DiffServ and their disadvantages, Ishan proposed a new method that compares the estimated transmission time and deadline in each node and schedules the packet with shortest per-hop time first. The main advantage of this method is (i) bursts are better handled, (ii) packets cannot meet the deadline are discarded to avoid wasting bandwidth, and (iii) the server is stateless.

Alexander asked which router is used in the implementation since this method requires clock synchronization and other computation. Ishan answered that currently only Linux machines with IP forwarding enabled are used, but they plan to use Cisco routers in the future. Padmanabhan asked about misbehaving node that reports false deadline requirement to gain advantage. Ishan argued that packets that cannot meet deadline will be dropped, so a misbehavior flow may be dropped to benefit other flows. John Miller asked how to obtain the value of TTL and deadline. Ishan answered that TTL value can be measured by nodes in the first round trip transmission and the deadline is decided according to the

requirement of the application.

## Session 7: P2P Streaming II

*Discussion Lead: Dongyan Xu (Purdue University)*

*Scribe: Amir Hassan Rasti Ekbatani (University of Oregon)*

- **Zebroid: IPTV Data Mining to Support Peer-Assisted VoD Content Delivery**

*Yih-Farn Robin Chen (AT&T Laboratories - Research)*

*Rittwik Jana (AT&T Laboratories - Research)*

*Daniel Stern (AT&T Laboratories - Research)*

*Bin Wei (AT&T Laboratories - Research)*

*Mike Yang (AT&T Laboratories - Research)*

*Hailong Sun (Beihang University)*

Robin opened his talk by emphasizing on the difference between Internet TV and IPTV. The former is applied to the best effort delivery of video content over a broadband connection and the latter is defined by ITU-T to provide security, reliability, and interactivity. He showed AT&T's architecture for IPTV, where optical media connects the video source to the DSLAM switch, which are connected to models using high bandwidth DSL links. The uplink bandwidth of a DSLAM switch is 1 Gbps and thus only supports a small number of IPTV users concurrently. The speaker then explained their proposed peer-assisted system for IPTV, Zebroid, where popular video-on-demand content is pre-stripped on user set-top boxes (STB) during idle hours. The STBs will then form a peer-assisted streaming overlay to deliver the VoD content to each interested user during the peak hours. Zebroid also tries to predict the probability of a node failure which happens when an STB is turned off and provision a level of redundancy accordingly. The simulation-based evaluation shows that their proposed peer-assisted delivery is only possible for the 8Mbps and 12 Mbps neighborhoods and only when the number of requesting peers is at most 8.

- **Robust Live Media Streaming in Swarms**

*Thomas Locher (ETH Zurich)*

*Remo Meier (ETH Zurich)*

*Roger Wattenhofer (ETH Zurich)*

*Stefan Schmid (TU Munich)*

Remo Meier started by talking about the challenges introduced by the P2P application model, including peer heterogeneity, robustness against randomness and selfishness, as well as system's fairness among peers. The presenter mentioned the features of a hypercubic overlay including simplicity of construction and maintenance as well as neighbor selection flexibility. Remo then presented the proposed P2P streaming mechanism in which the peers form multiple hypercubic trees and push the content blocks down the tree to reach a quarter of the peers. In the next phase, with a BitTorrent-like pull-based mechanism, the content is distributed among all peers. Using simulation-based evaluation, Meier showed that their proposed system successfully limits free riders. Peers who share too few packets often suffer from underflows. He suggested that by using layered audio/video codecs, it is possible to accommodate heterogeneous peer groups.

- **Providing Statistically Guaranteed Streaming Quality for Peer-to-Peer Live Streaming**

*Miao Wang (University of Nebraska-Lincoln)*

*Lisong Xu (University of Nebraska-Lincoln)*

*Byrav Ramamurthy (University of Nebraska-Lincoln)*

Lisong Xu opened his talk by positing that a common problem for P2P streaming was the "best-effort" quality. Due to user dynamics and system complexity, achieving a guaranteed quality is hard. The speaker defined the paper's goal as providing statistical guarantee on the overall upload bandwidth of a P2P streaming system using admission control. The paper uses a queueing model used based on heterogeneous upload bandwidth and includes user dynamics as well as admission control, and defines the statistical guarantee problem as deciding whether to admit a new peer to maintain the probability of having enough upload bandwidth. Toward this end, they show that a dynamic admission control algorithm that exploit state information leads to lower rejection rate with the same level of guarantee. The speaker also explore

admission control's sensitivity to user lifetime distribution and user arrival process. and concluded that there is a fundamental tradeoff between user rejection rate and user behavior insensitivity.

## • Discussion

The session continued with a question/answer panel including all three speakers.

Mohamed asked the last speaker where the admission control algorithms were to be executed. Lisong pointed out that their paper had assumed a central entity responsible for admission control. Mohamed then asked the first speaker whether Zebroid provides multiple encoding rates or scalable coding for the same TV channel to accommodate users who want to receive HD or SD versions of each channel. Robin Chen was not sure whether multiple encoding rates were provided, yet he pointed out that in future such schemes would be necessary to provide IPTV service to mobile devices.

Ardalan asked Remo about the possibility of unbalanced delivery tree. Remo responded that due to randomness of IDs, unbalanced tree is rare, however, their algorithm periodically ensure that the tree is balance.

Michela raised the concern of energy consumption at the set-top boxes in Zebroid during the idle hours when the boxes would be otherwise off or stand-by. Robin Chen explained that the set-top boxes have the potential to be the next battlefield for a variety of applications similar to cell phones and could stay on most of the time.

Hassan asked Robin about the mention of erasure code in his presentation. Robin clarified that erasure code is necessary to compensate for the peers that are not up during the content push time or at the peak hour. However, such mechanism was not implemented into Zebroid yet.

Saurabh asked Robin about the level of peer outage considered. Robin explained that they assume 80% of the peers are up at all times, and this rate is based on

measurement data.

Mohamed asked Robin if data pushed to each client has any correlation with the usage history of that user. Chen responded that they plan to include a mechanism to predict the usage of each user and consider that prediction in pushing the data and then refining the predictions gradually with the user's choices. Discussion ensued about the privacy issues on this approach.

Philip asked Remo about ID conflict between clients since IDs are assigned randomly. Remo Meier explains that although the probability is low, they had proposed a mechanism to ensure that does not happen in another paper.

Dongyan then asked the panelists how they assessed their proposed systems' robustness. Robin mentioned that they could configure the system to ensure that the user had no access to the set-top box, while Remo pointed out that their approach designed each peer's behavior to be as greedy as possible, therefore they would not expect any manipulation to change the code to make it more selfish.

This website is mirrored at <http://www.nossdav.org/2009/> and <http://nemesys.comp.nus.edu.sg/nossdav09/>.

# NOSSDAV 2009

The 19th International Workshop on Network and Operating Systems Support for Digital Audio and Video  
Williamsburg, Virginia. June 3-5, 2009

[Home](#)

[Organization](#)

[Call For Papers](#)

[Submit Papers](#)

[Registration](#)

[Program](#)

[Scribes](#)

[Abstracts, Papers,  
Slides](#)

[Travel Logistic](#)

[Social Networks](#)



[Past NOSSDAV  
Workshops](#)

- [2008](#)
- [2007](#)
- [2006](#)
- [2005](#)
- [more..](#)

**Sponsored by:**



**In cooperation with:**



## Technical Program

---

### 03 June 2009

1800 **Registration Opens**

1900 **Reception**

---

### 04 June 2009

0800 **Registration and Breakfast**

0830 - 0845 **Welcome Address**

0845 - 1000 **Keynote Session**



*Discussion Lead: Dongyan Xu*

**Experiences Building  
Telepresence Systems**  
**Keynote Speaker: Prof. Henry  
Fuchs**

Federico Gil Professor of Computer  
Science  
Adjunct Professor of Biomedical  
Engineering  
University of North Carolina at Chapel-  
Hill

1000 - 1030 **Coffee Break**

1030 - 1200 **Session 1: New Applications  
and Services**



*Discussion Lead: Wei Tsang Ooi (National University of Singapore)*

*Scribe: Saurabh Ratti (University of Ottawa)*



Presented by:  
 Wanmin Wu  
 (University of Illinois at Urbana-Champaign)



## **Dynamic Overlay Multicast in Multi-Stream/Multi-Site 3D Collaborative Systems**

*Wanmin Wu (University of Illinois at Urbana-Champaign)*

*Zhenyu Yang (Florida International University)*

*Klara Nahrstedt (University of Illinois at Urbana-Champaign)*

Multi-stream/multi-site 3D video collaborative systems are promising as they enable remote users to interact in a 3D virtual space with a sense of co-presence. However, the decentralized content dissemination remains a challenge. In this work, we explore approaches to construct adaptive overlay based on the users' visual interest in the collaborative space. Particularly, we consider the practical challenge that the user's interest might change dynamically. We propose, compare, and evaluate three algorithms to handle the view dynamics. With extensive experiments, we demonstrate that an algorithm that exploits view locality can achieve efficient bandwidth utilization, high topology stability, and great scalability.




Presented by:

## **Web 2.0 Traffic Measurement: Analysis on Online Map Applications**

*Song Lin (Tsinghua University)*


*Zhiguo Gao (IBM China Research)*

Song Lin  
(Tsinghua  
University)



Laboratory)  
Ke Xu (Tsinghua University)

In recent years, web based online map applications have been getting more and more popular, such as Google Maps, Yahoo Maps. Many new Web 2.0 techniques such as mash-up and AJAX were adopted in these map applications to improve user experiences. But few researches have been done on traffic analysis of the Web 2.0 based online map applications. In this paper, we introduced our research on features of online map applications that previous studies hadn't cover. In our research, we captured map application related HTTP traffic in a campus network while not violating user privacy. We introduced the traffic overview, mash-up and web caching characteristics of four map web sites (Google Maps, Yahoo Maps, Sogou Maps and Baidu Maps). For the first time, the mash-up characteristics of Google map traffic were analyzed using a new method proposed in this paper. The same method could be applied to other mash-up analysis works. These results can help us optimize the future web application designs and CDN based accelerating solution designs.



Presented by:  
Yuen Feng  
(University of  
Toronto)

## Peer-Assisted Online Storage and Distribution: Modeling and Server Strategies

Ye Sun (Hong Kong University of Science & Technology)

Fangming Liu (Hong Kong University of Science & Technology)

Bo Li (Hong Kong University of Science &



*Technology)*

*Baochun Li (University of Toronto)*

Peer-assisted online storage and distribution systems have recently enjoyed large-scale deployment gaining increased popularity for multimedia content sharing in the Internet. Such systems typically deploy dedicated servers while effectively leveraging peer bandwidth in a complementary fashion, in order to guarantee adequate levels of service quality and minimize server cost. In this paper, motivated by our recent empirical study on a real-world system, FS2You, we develop a mathematical model to characterize and understand peer-assisted online storage systems serving multiple files of different popularity. Specifically, we examine and compare representative server bandwidth allocation strategies, and investigate the critical performance metrics and factors. We demonstrate that different server strategies may lead to remarkably different service qualities in terms of average downloading times, peer satisfaction levels and service quality differentiation. In particular, the current server strategy in FS2You is able to offer system-wide average downloading times comparable to the theoretical bound derived from our model.

1200 - 1330

**Lunch**

1330 - 1500

**Session 2: P2P Streaming I**

*Discussion Lead: Ketan Mayer-Patel  
(University of North Carolina at Chapel Hill)*



Presented by:  
Hassan  
Shojanian  
(University of  
Toronto)



Scribe: Remo Meier (ETH Zurich)

## InstantLeap: Fast Neighbor Discovery in P2P VoD Streaming

Xuanjia Qiu (Sun Yat-Sen University)

Chuan Wu (The University of Hong Kong)

Xiaola Lin (Sun Yat-Sen University)

Francis C. M. Lau (The University of Hong Kong)

A fundamental challenge in peer-to-peer (P2P) Video-on-Demand (VoD) streaming is to quickly locate new supplying peers whenever a VCR command is issued, in order to achieve smooth viewing experiences. For most existing commercial systems which resort to tracking servers for such neighbor discovery, the increasing scale of P2P VoD systems has brought heavy load onto the dedicated servers. To avoid overloading the servers and achieve instant neighbor discovery over the self-organizing P2P overlay, we design a novel method of organizing peers watching the same video, that constitutes a light-weighted indexing structure to support efficient streaming and fast neighbor discovery at the same time. InstantLeap achieves an  $O(1)$  neighbor discovery efficiency upon any playback "leaps" across the media stream in streaming overlays of any sizes, with a low messaging cost for the overlay maintenance. We support our design with rigorous analysis and extensive simulations.

## Overlay Monitoring and Repair in Swarm-based Peer-to-Peer Streaming



Presented by:

Nazanin  
Magharei  
(University of  
Oregon)

*Nazanin Magharei (University of Oregon)*

*Reza Rejaie (University of Oregon)*



In Swarm-based Peer-to-Peer Streaming (SPS) mechanisms, participating peers form a randomly connected mesh over which they incorporate swarm-like content delivery. In practice, a subset of participating peers may form clusters in the overlay due to various reasons such as localization of connectivity within edge ISPs. Despite the commonly held assumptions, the appearance of such clusters could significantly degrade the delivered quality to participating peers in SPS mechanisms. This paper examines the effect of overlay clustering on the performance of SPS mechanisms for live content. Leveraging the notion of two-phase content delivery in SPS mechanisms, we illustrate the effect of overlay clustering on content delivery. We propose the Overlay Monitoring and Repair (OMR) mechanism as a distributed and scalable approach to maintain proper overlay connectivity in SPS mechanisms. The key idea is to use delivered quality to individual peers as an indication of poor connectivity from other regions of the overlay. OMR employs a probabilistic approach to ensure an adequate number of properly-positioned peers reacts to detected clustering in the overlay without any coordination. Reacting peers rewire a small number of carefully-selected connections in the overlay to significantly improve the performance of

content delivery. Our preliminary evaluations demonstrate that OMR mechanism can achieve its goals.



Presented by:  
Michela Meo  
(*Politecnico di Torino*)



## Adaptive Overlay Topology for Mesh-Based P2P-TV Systems

*Richard John Lobb (University of Canterbury)*  
*Ana Paula Couto da Silva (Federal University of Juiz de Fora)*

*Emilio Leonardi (Politecnico di Torino)*  
*Marco Mellia (Politecnico di Torino)*  
*Michela Meo (Politecnico di Torino)*

In this paper, we propose a simple and fully distributed mechanism for constructing and maintaining the overlay topology in mesh-based P2P-TV systems. Our algorithm optimizes the topology to better exploit large bandwidth peers, so that they are automatically moved close to the source. This improves the chunk delivery delay so that all peers benefit, not just the high bandwidth ones. A key property of the proposed scheme is its ability to indirectly estimate the upload bandwidth of peers without explicitly knowing or measuring it. Simulation results show that our scheme significantly outperforms overlays with homogeneous properties, achieving up to 50% performance improvement. Moreover, the algorithm is robust to both parameter setting and changing conditions, e.g., peer churning.

1500 - 1530

**Coffee Break**


## Session 3: OS and End Systems

*Discussion Lead: Kevin Almeroth (UC Santa Barbara)*

*Scribe: Ishan Vaishnavi (Centrum voor Wiskunde en Informatica)*



Presented by:  
Hassan  
Shojania  
(University of  
Toronto)



### Random Network Coding on the iPhone: Fact or Fiction?

*Hassan Shojania (University of Toronto)*

*Baochun Li (University of Toronto)*

In multi-hop wireless networks, random network coding represents the general design principle of transmitting random linear combinations of blocks in the same "batch" to downstream relays or receivers. It has been recognized that random network coding in multi-hop wireless networks may improve unicast throughput in scenarios when multiple paths are simultaneously utilized between the source and the destination. However, the computational complexity of random network coding, and its energy consumption implications, may potentially limit its applicability and practicality in mobile devices. In this paper, we present our real-world implementation of random network coding on the Apple iPhone and iPod Touch mobile platforms, and offer an in-depth investigation with respect to the difficulties towards such an implementation, the limitations of the ARM processor and the hardware platform, as well as our hand-tuning efforts to maximize coding performance on the iPhone platform. With our implementation deployed on both the iPhone 3G and the second-generation iPod Touch, we report its coding performance, energy consumption rates,



as well as CPU usage with multimedia streaming.



Presented by:  
Padmanabhan  
Pillai  
(Intel  
Research  
Pittsburg)

## SLIPstream: Scalable Low-latency Interactive Perception on Streaming Data

*Padmanabhan S. Pillai (Intel Research Pittsburgh)*

*Lily B. Mummert (Intel Research Pittsburgh)*

*Steven W. Schlosser (Intel Research Pittsburgh)*

*Rahul Sukthankar (Intel Research Pittsburgh)*

*Casey J. Helfrich (Intel Research Pittsburgh)*



A critical problem in implementing interactive perception applications is the considerable computational cost of current computer vision and machine learning algorithms, which typically run one to two orders of magnitude too slowly to be used interactively. Fortunately, many of these algorithms exhibit coarse-grained task and data parallelism that can be exploited across machines. The SLIPstream project focuses on building a highly-parallel runtime system called Sprout that can harness the computing power of a cluster to execute perception applications with low latency. This paper makes the case for using clusters for perception applications, describes the architecture of the Sprout runtime, and presents two compute-intensive yet interactive applications.

**Server-Efficient High-Definition Media Dissemination**





Presented by:  
Philip Frey  
(IBM Research  
GmbH)



*Philip W. Frey (IBM Research GmbH)*  
*Andreas Hasler (IBM Research GmbH)*  
*Bernard Metzler (IBM Research GmbH)*  
*Gustavo Alonso (ETH Zurich)*

Internet usage has changed dramatically in the past few years. Content is no longer dominated by static websites, but comprises an increasing number of multimedia streams. With the widespread availability of broadband connections, the quality of the media provided by video-on-demand as well as streaming services increases constantly. Even though today most videos are still encoded with a rather low bit rate, large Internet service providers already foresee high definition media becoming the predominant format in the near future. However, a larger number of clients requesting media at high bit rates poses a challenge for the server infrastructure. Conventional stream dissemination methods, such as RTP over UDP or HTTP over TCP, result in high server loads due to excessive local data copy, context switching, and interrupt processing overhead. In this paper, we illustrate and discuss this problem in detail through extensive experiments with existing solutions. We then present a new approach based on zero-copy protocol stack implementations in software as well as dedicated RDMA hardware. Our performance experiments indicate that these optimizations allow servers to scale better and remove most of the overhead caused by current approaches.



Presented by:  
Damien Le  
Moal  
(Hitachi Ltd.)



## Power Efficient Real-Time Disk Scheduling

*Damien Le Moal (Hitachi Ltd.)*

*Donald Molaro (Hitachi Global Storage Technologies, San Jose Research Center)*

*Jorge Campello (Hitachi Global Storage Technologies, San Jose Research Center)*

Hard-disk drive power consumption reduction methods focus mainly on increasing the amount of time the disk is in standby mode (disk spun down) by implementing aggressive data readahead and caching at the operating system and/or application level. However, these methods cannot be applied efficiently to systems with limited memory and high bit-rate requirements such as digital video recorders handling high-definition video. In this paper, we introduce the Audio/Video File System (AVFS), composed of a file system and a disk I/O scheduler. Compared to traditional methods, the proposed scheduler reduces seek overhead by processing real-time requests to video files using batches built dynamically depending on the requests deadlines. Evaluation results show an important reduction in disk utilization rates and a reduction of up to 20 % of the disk power consumption with only 4 MB of data buffer per video stream.

1900

**Banquet**

**05** June 2009

0800

**Registration and Breakfast**

0830 - 1000

**Session 4: Virtual Environments and Games***Discussion Lead: Kuan-Ta Chen (Academia Sinica)**Scribe: Pengpeng Ni (Simula Research Lab and University of Oslo)*

Presented by:  
Saurabh Ratti  
(University of  
Ottawa)

**A Delaunay Triangulation Architecture Supporting Churn and User Mobility in MMVE-s***Mohsen Ghaffari (Sharif University of Technology)**Behnoosh Hariri (Sharif University of Technology and University of Ottawa)**Shervin Shirmohammadi (Sharif University of Technology and University of Ottawa)*

This article proposes a new distributed architecture for update message exchange in massively multi-user virtual environments (MMVE). MMVE applications require delivery of updates among various locations in the virtual environment. The proposed architecture here exploits the location addressing of geometrical routing in order to alleviate the need for IP-specific queries. However, the use of geometrical routing requires careful choice of overlay to achieve high performance in terms of minimizing the delay. At the same time, the MMVE is dynamic, in sense that users are constantly moving in the 3D virtual space. As such, our architecture uses a distributed topology control scheme that aims at maintaining the requires QoS to best support the greedy geometrical routing, despite user mobility or churn. We will further prove the functionality and performance of the proposed scheme through both theory and

simulation.



Presented by:  
John Miller  
(*Microsoft  
Research and  
University of  
Cambridge*)



## Probabilistic Event Resolution with the Pairwise Random Protocol

*John L. Miller (Microsoft Research and University of Cambridge)*

*Jon Crowcroft (University of Cambridge)*

Peer-to-peer distributed virtual environments (DVE's) distribute state tracking and state transitions. Many DVE's - such as online games - require ways to fairly determine the outcome of probabilistic events. While trivial when a trusted third party is involved, resolving these actions fairly between adversaries without a trusted third party is much more difficult. This paper proposes the Pairwise Random Protocol (PRP), which uses secure coin flipping to enable adversaries to fairly determine the result of a probabilistic event without a trusted third party. Three different variations of PRP are presented, and the time impact and network overhead are examined. We conclude that PRP enables DVE's to distribute the work of determining probabilistic events between adversaries without loss of security or fairness, and with acceptable overhead.



Presented by:  
Ke Liang  
(*National*


## Cross-Tree Adjustment for Spatialized Audio Streaming over Networked Virtual Environments

*Ke Liang (National University of Singapore)*

*Roger Zimmermann (National University of Singapore)*

In recent years, integrated spatialized

University of  
Singapore)



voice services have become an appealing application for networked virtual environments (NVE), e.g., Second Life. With a spatialized voice service, people can identify who is talking if there are several participants in the vicinity. The key challenge in a spatialized audio streaming application is how to disseminate audio streams while observing bandwidth limits of end-user computers and tight latency constraints, and it can be modeled as NP-complete problem. In this paper, we propose a heuristic algorithm called CTA for spatialized audio streaming over NVEs in a peer-to-peer manner. The proposed algorithm was applied to real avatar mobility traces collected from Second Life, and the simulation results demonstrate that (a) CTA can achieve a high ratio of successful to candidate receivers and (b) CTA enables most of the successful receivers to enjoy minimum latency.

1000 - 1030

## Coffee Break

1030 - 1200

## Session 5: Security

*Discussion Lead: Klara Nahrstedt (UIUC)*

*Scribe: Liang Ke (National University of Singapore)*



Presented by:

Mohamad  
Hefeeda  
(Simon Fraser


## End-to-End Secure Delivery of Scalable Video Streams

*Kianoosh Mokhtarian (Simon Fraser University)*


*Mohamed Hefeeda (Simon Fraser University)*

We investigate the problem of securing the delivery of scalable video streams so

University)



that receivers can ensure the authenticity (originality and integrity) of the video. Our focus is on recent scalable video coding techniques, e.g., H.264/SVC, that can provide three scalability types at the same time: temporal, spatial, and quality (or PSNR). This three-dimensional scalability offers a great flexibility that enables customizing video streams for a wide range of heterogeneous receivers and network conditions. This flexibility, however, is not supported by current stream authentication schemes in the literature. We propose an efficient authentication scheme that accounts for the full scalability of video streams: it enables verification of all possible substreams that can be extracted and decoded from the original stream. Our evaluation study shows that the proposed authentication scheme is robust against packet losses, adds low communication and computation overheads, and is suitable for live streaming systems as it has short delay.



Presented by:  
Tony Thomas  
(Nanyang  
Technological  
University)

## Secure Multimedia Content Delivery with Multiparty Multilevel DRM Architecture

*Tony Thomas (Nanyang Technological University)*

*Sabu Emmanuel (Nanyang Technological University)*

*Amitabha Das (Nanyang Technological University)*

*Mohan S. Kankanhalli (National University of Singapore)*




For scalability of business, multiparty multilevel digital rights management

(DRM) architecture, where a multimedia content is delivered by an owner to a consumer through several levels of distributors has been suggested as an alternative to the traditional two party (buyer-seller) DRM architecture. A combination of cryptographic and watermarking techniques are usually used for secure content delivery and protecting the rights of seller and buyer in the two party DRM architecture. In a multiparty multilevel DRM architecture the cryptographic and watermarking mechanism need to ensure the secure delivery of the content as well as the security concerns of the owner, multiple levels of the distributors and the consumer. In this paper, we propose a mechanism which takes care of the above security issues, for delivering multimedia content through multiparty multilevel DRM architecture.



Presented by:  
Philip Branch  
(Swinburne  
University of  
Technology)



## Rapid Identification of Skype Traffic

*Philip A. Branch (Swinburne University of Technology)*

*Amiel Heyde (Swinburne University of Technology)*

*Grenville J. Armitage (Swinburne University of Technology)*

In this paper we present results of experimental work using machine learning techniques to rapidly identify Skype traffic. We show that Skype traffic can be identified by observing 5 seconds of a Skype traffic flow, with recall and precision better than 98%. We found the most effective features for classification were characteristic packet lengths less



than 80 bytes, statistics of packet lengths greater than 80 bytes and inter-packet arrival times. Our classifiers do not rely on observing any particular part of a flow. We also report on the performance of classifiers built using combinations of two of these features and of each feature in isolation.

1200 - 1330

**Lunch**

1330 - 1500

## Session 6: Understanding and Improving User Experience

*Discussion Lead: Mohamed Hefeeda (Simon Fraser University)*

*Scribe: Wei Cheng (National University of Singapore)*



Presented by:  
Kuan-Ta Chen  
(Academia Sinica,  
Taiwan)

### An Empirical Evaluation of VoIP Playout Buffer Dimensioning in Skype, Google Talk, and MSN Messenger

*Chen-Chi Wu (National Taiwan University)*

*Kuan-Ta Chen (Academia Sinica)*

*Chun-Ying Huang (National Taiwan Ocean University)*

*Chin-Laung Lei (National Taiwan University)*



VoIP playout buffer dimensioning has long been a challenging optimization problem, as the buffer size must maintain a balance between conversational interactivity and speech quality. The conversational quality may be affected by a number of factors, some of which may change over time. Although a great deal of research effort has been expended in trying to solve the problem, how the research results are applied in practice is unclear. In this paper, we investigate the playout buffer



dimensioning algorithms applied in three popular VoIP applications, namely, Skype, Google Talk, and MSN Messenger. We conduct experiments to assess how the applications adjust their playout buffer sizes. Using an objective QoE (Quality of Experience) metric, we show that Google Talk and MSN Messenger do not adjust their respective buffer sizes appropriately, while Skype does not adjust its buffer at all. In other words, they could provide better QoE to users by improving their buffer dimensioning algorithms. Moreover, none of the applications adapts its buffer size to the network loss rate, which should also be considered to ensure optimal QoE provisioning.



Presented by:  
Pengpeng Ni  
(Simula  
Research Lab  
and University  
of Oslo)

## Fine-Grained Scalable Streaming from Course-Grained Video

*Pengpeng Ni (Simula Research Laboratory and University of Oslo)*

*Alexander Eichhorn (Simula Research Laboratory)*

*Carsten Griwodz (Simula Research Laboratory and University of Oslo)*

*Pål Halvorsen (Simula Research Laboratory and University of Oslo)*



Scalable video is an attractive option for adapting the bandwidth consumption of streaming video to the available bandwidth. Fine-grained scalability can adapt most closely to the available bandwidth, but this comes at the cost of a high compression penalty. In the context of VoD streaming to mobile end systems, we have therefore explored whether a similar adaptation to the available bandwidth can be achieved by performing

layer switching in coarse-grained scalable videos. In this approach, enhancement layers of a video stream are switched on and off to achieve any desired longer-term bandwidth. We performed user studies to evaluate the idea, and came to the far-from-obvious conclusion that layer switching is viable way for bit-rate savings and fine-grained bitrate adaptation even for rather short times between layer switches.




Presented by:  
Ishan  
Vaishnavi  
(Centrum voor  
Wiskunde en  
Informatica)

## Estimate and Serve: Scheduling Soft Real-Time Packets for Delay Sensitive Media Applications on the Internet

*Ishan Vaishnavi (Centrum voor Wiskunde en Informatica)*

*Dick C. A. Bulterman (Centrum voor Wiskunde en Informatica)*



This paper presents a new scheduling algorithm for real time network delivery of packets over Diffserv networks for delay sensitive applications. We name the networks that support this algorithm as Estimated Service (Estserv) networks. These networks, for real time packets, estimate the probability of the packet meeting its deadline and schedule it according to this estimation. This paper validates, given this estimation mechanism, the better performance of the scheduling algorithm over traditional solutions. We show that using Estserv for delay sensitive applications, we can provide out-of-band scheduling, save bandwidth on packets with expired deadlines and handle bursts without losing the scalability of Diffserv. We

show with the help of an implementation in the Linux kernel's ip-forwarding part, that, given the estimation value, Estserv performs better than Diffserv in terms of deadlines, while still saving bandwidth.

1500 - 1530 **Coffee Break**

1530 - 1700 **Session 7: P2P Streaming II**

*Discussion Lead: Dongyan Xu (Purdue University)*

*Scribe: Amir Hassan Rasti Ekbatani (University of Oregon)*



Presented by:  
Yih-Farn Robin Chen  
(AT&T Laboratories - Research)

### **Zebroid: IPTV Data to Support Peer-Assisted VoD Content Delivery**

*Yih-Farn Robin Chen (AT&T Laboratories - Research)*


*Rittwik Jana (AT&T Laboratories - Research)*

*Daniel Stern (AT&T Laboratories - Research)*

*Bin Wei (AT&T Laboratories - Research)*


*Mike Yang (AT&T Laboratories - Research)*

*Hailong Sun (Beihang University)*



P2P file transfers and streaming have already seen a tremendous growth in Internet applications. With the rapid growth of IPTV, the need to efficiently disseminate large volumes of Video-on-Demand (VoD) content has prompted IPTV service providers to consider peer-assisted VoD content delivery. This paper describes Zebroid, a VoD solution that uses IPTV operational data on an on-going basis to determine how to pre-position popular content in customer set-top boxes during idle hours to allow these peers to assist the VoD server in content delivery during peak hours. Latest VoD request distribution, set-top box

availability, and capacity data on network components are all taken into consideration in determining the parameters used in the striping algorithm of Zebroid. We show both by simulation and emulation on a realistic IPTV testbed that the VoD server load can be significantly reduced by more than 50-80% during peak hours by using Zebroid.



Presented by:  
Remo Meier  
(ETH Zurich)

## Robust Live Media Streaming in Swarms

*Thomas Locher (ETH Zurich)*

*Remo Meier (ETH Zurich)*

*Roger Wattenhofer (ETH Zurich)*

*Stefan Schmid (TU Munich)*



Data dissemination in decentralized networks is often realized by using some form of swarming technique. Swarming enables nodes to gather dynamically in order to fulfill a certain task collaboratively and to exchange resources (typically pieces of files or packets of a multimedia data stream). As in most distributed systems, swarming applications face the problem that the nodes in a network have heterogeneous capabilities or act selfishly. We investigate the problem of efficient live data dissemination (e.g., TV streams) in swarms. The live streams should be distributed in such a way that only nodes with sufficiently large contributions to the system are able to fully receive it -- even in the presence of freeloading nodes or nodes that upload substantially less than required to sustain the multimedia stream. In contrast, uncooperative nodes cannot properly receive the data stream

as they are unable to fill their data buffers in time, incentivizing a fair sharing of resources. If the number of selfish nodes increases, our emulation results reveal that the situation steadily deteriorates for them, while obedient nodes continue to receive virtually all packets in time.



Presented by:  
Lisong Xu  
(University of  
Nebraska-  
Lincoln)



## Providing Statistically Guaranteed Streaming Quality for Peer-to-Peer Live Streaming

*Miao Wang (University of Nebraska-Lincoln)*  
*Lisong Xu (University of Nebraska-Lincoln)*  
*Byrav Ramamurthy (University of Nebraska-Lincoln)*

Most of the literature on peer-to-peer (P2P) live streaming focuses on how to provide best-effort streaming quality by efficiently using the system bandwidth; however, there is no guarantee about the provided streaming quality. This paper considers how to provide statistically guaranteed streaming quality to a P2P live streaming system. We study a class of admission control algorithms which statistically guarantee that a P2P live streaming system has sufficient overall bandwidth. Our results show that there is a tradeoff between the user blocking rate and user-behavior insensitivity (i.e., whether the system performance is insensitive to the fine statistics of user behaviors). We also find that the system performance is more sensitive to the distribution change of user inter-arrival times than to that of user lifetimes.

This website is mirrored at <http://www.nossdav.org/2009/> and <http://nemesys.comp.nus.edu.sg/nossdav09/>.



# NOSSDAV 2009

The 19th International Workshop on Network and Operating Systems Support for Digital Audio and Video  
Williamsburg, Virginia. June 3-5, 2009

[Home](#)

[Organization](#)

[Call For Papers](#)

[Submit Papers](#)

[Registration](#)

[Program](#)

[Scribes](#)

[Abstracts, Papers,  
Slides](#)

[Travel Logistic](#)

[Social Networks](#)



[Past NOSSDAV  
Workshops](#)

[2008](#)

[2007](#)

[2006](#)

[2005](#)

[more..](#)

**Sponsored by:**



Association for  
Computing Machinery

*Advancing Computing as a Science & Profession*



**In cooperation with:**



SIGOPS SIGCOMM

## Accomodation at Conference Hotel

Advance Booking Deadline: 4 May, 2009.

NOSSDAV 2009 will be held at [Williamsburg Marriot](#). The hotel is offering a room rate of USD 129 per night to NOSSDAV 2009 attendees. To enjoy this rate, book online with special code [ACAACAA](#) or contact the in-house Group Housing Coordinator [Marry Woodhouse](#) at [\(757\) 259-5693](#) or [mwoodhouse@columbiasussex.com](mailto:mwoodhouse@columbiasussex.com) no later than 4 May 2009.

[View Larger Map](#)

## Getting to Williamsburg

The website [visitwilliamsburg.com](http://visitwilliamsburg.com) contains useful information about how to get to Williamsburg by car, air, and train.

## Airport Shuttle

[Williamsburg Airport Shuttle](#) from [Newport News \(PHF\)](#), [Richmond \(RIC\)](#), and [Norfolk \(ORF\)](#).

## Visa Applications

The sites <http://www.unitedstatesvisas.gov> and <http://travel.state.gov> have information about obtaining a visa for those traveling to the United States. Both sites have links to websites for U.S. embassies and consulates worldwide. The embassy and consulate websites have very helpful information about procedures, timelines, communities served, required documentation, and fees.

International registrants should be particularly aware and careful about visa requirements, and should plan travel well in advance. All visa inquiries must be handled by ACM HQ. Please send your request for a letter in support of a visa application to [supportletters@acm.org](mailto:supportletters@acm.org) include *your name, mailing address, fax number, and the name of the conference you are attending.* (*Authors of papers/posters should include the title*). Visa support letters cannot be issued via email. Please note that ACM does not issue formal "letters of invitation" to any of its conferences.

This website is mirrored at <http://www.nossdav.org/2009/> and <http://nemesys.comp.nus.edu.sg/nossdav09/>.

---



## Session 1: New Applications and Services

Dynamic Overlay Multicast in Multi-Stream/Multi-Site 3D Collaborative Systems  
W. Wu, Z. Yang, K. Nahrstedt

Web 2.0 Traffic Measurement: Analysis on Online Map Applications  
S. Lin, Z. Gao, K. Xu

Peer-Assisted Online Storage and Distribution: Modeling and Server Strategies  
Y. Sun, F. Liu, B. Li, B. Li

## Session 2: P2P Streaming I

InstantLeap: Fast Neighbor Discovery in P2P VoD Streaming  
X. Qiu, C. Wu, X. Lin, F. Lau

Overlay Monitoring and Repair in Swarm-based Peer-to-Peer Streaming  
N. Magharei, R. Rejaie

Adaptive Overlay Topology for Mesh-Based P2P-TV Systems  
R. Lobb, A. Silva, E. Leonardi, M. Mellia, M. Meo

## Session 3: OS and End Systems

Random Network Coding on the iPhone: Fact or Fiction?  
H. Shojania, B. Li

SLIPstream: Scalable Low-latency Interactive Perception on Streaming Data  
P. Pillai, L. Mummert, S. Schlosser, R. Sukthankar, C. Helfrich

Server-Efficient High-Definition Media Dissemination  
P. Frey, A. Hasler, B. Metzler, G. Alonso

Power Efficient Real-Time Disk Scheduling  
D. Moal, D. Molaro, J. Campello

## Session 4: Virtual Environments and Games

A Delaunay Triangulation Architecture Supporting Churn and User Mobility in MMVE-s  
M. Ghaffari, B. Hariri, S. Shirmohammadi

Probabilistic Event Resolution with the Pairwise Random Protocol  
J. Miller, J. Crowcroft

Cross-Tree Adjustment for Spatialized Audio Streaming over Networked Virtual Environments  
K. Liang, R. Zimmermann

## Session 5: Security

End-to-End Secure Delivery of Scalable Video Streams  
Kianoosh, M. Hefeeda

Secure Multimedia Content Delivery with Multiparty Multilevel DRM Architecture  
T. Thomas, S. Emmanuel, A. Das, M. Kankanhalli

Rapid Identification of Skype Traffic  
P. Branch, A. Heyde, G. Armitage

## Session 6: Understanding and Improving User Experience

An Empirical Evaluation of VoIP Playout Buffer Dimensioning in Skype, Google Talk, and MSN Messenger  
C. Wu, K. Chen, Y. Chang, C. Lei

Fine-Grained Scalable Streaming from Course-Grained Video  
P. Ni, A. Eichhorn, C. Griwodz, P. Halvorsen

Estimate and Serve: Scheduling Soft Real-Time Packets for Delay Sensitive Media Applications on the Internet  
I. Vaishnavi, D. Bulterman

## **Session 7: P2P Streaming II**

Zebroid: IPTV Data Mining to Support Peer-Assisted VoD Content Delivery  
Y. Chen, R. Jana, D. Stern, B. Wei, M. Yang, H. Sun

Robust Live Media Streaming in Swarms  
T. Locher, R. Meier, S. Schmid, R. Wattenhofer

Providing Statistically Guaranteed Streaming Quality for Peer-to-Peer Live Streaming  
M. Wang, L. Xu, B. Ramamurthy

Call For Papers  
NOSSDAV 2009

The 19th International Workshop on  
Network and Operating Systems Support for Digital Audio and Video

Williamsburg, Virginia, 3-5 June 2009

SPONSORED BY:  
ACM SIGMM  
in cooperation with  
ACM SIGCOMM and ACM SIGOPS

IMPORTANT DATES

Paper Deadline: 16 February 2009, 23:59 PST (Firm)

Notification: 20 March 2009

Camera Ready Due: 6 April 2009

<http://www.nossdav.org/2009/>

<http://nemesys.comp.nus.edu.sg/nossdav09/>

NOSSDAV 2009 will continue the workshop's long tradition of focusing on emerging topics, controversial ideas, and future research directions in the area of multimedia systems research, held in a setting that stimulates lively discussions among the senior and junior participants.

It is also an established practice for NOSSDAV to encourage experimental research based on real systems and data sets. Public availability of source code and data sets is highly encouraged.

The scope of NOSSDAV has evolved over the years. For NOSSDAV 2009, we would like especially highlight two new topics of interest: unconventional use of GPU for multimedia and multi-core processors support for multimedia. NOSSDAV 2009 continues to welcome submissions in the traditional topics of networked multimedia systems, operating system support for multimedia, and multimedia security and rights management.

The topics of interest include, but are not limited to:

- o OS, middleware and network support for multimedia
- o Overlay networks for multimedia
- o Media streaming, distribution and storage
- o Web 2.0 systems and social networks for multimedia
- o Wireless and mobile multimedia systems
- o Media sensor and ad hoc networks
- o Grid/cloud computing support for multimedia
- o Embedded systems for multimedia
- o Multicore architecture support for multimedia
- o GPU for multimedia
- o Network processor support for multimedia
- o Tele-immersion and mixed-reality systems

- o Networked graphics and virtual environments
- o Networked games
- o Multimedia communications and system security
- o Digital rights management

A broad view will be taken in deciding what topics are within scope. Please contact the workshop co-chairs if you are unsure and wish to check if a particular topic is within the scope of NOSSDAV.

Submissions should be at most SIX pages in length using standard ACM proceedings style. We expect these submissions to be the kernel of what will eventually lead to full-length papers at high-quality conferences or journals.

Authors of selected, high quality papers from NOSSDAV 2009 will be invited to submit an extended version of their papers to a special issue of ACM/Springer Multimedia Systems Journal.

## ORGANIZING COMMITTEE

### Workshop Co-Chairs

Dongyan Xu, Purdue University, USA  
Wei Tsang Ooi, National University of Singapore

### Local Arrangement Chair

Ketan Mayer-Patel, UNC Chapel Hill, USA

### Publicity Co-Chairs

Tristan Henderson, University of St. Andrews, UK  
Kang Li, University of Georgia, USA  
Kuan-Ta Chen, Academia Sinica, Taiwan

### Technical Program Committee

Kevin Almeroth, UC Santa Barbara, USA  
Grenville Armitage, Swinburne U. of Tech, Australia  
Ernst Biersack, Institute Eurecom, France  
Surendar Chandra, University of Notre Dame, USA  
Kuan-Ta Chen, Academia Sinica, Taiwan  
Songqing Chen, George Mason University, USA  
Mark Claypool, WPI, USA  
Wu-chang Feng, Portland State University, USA  
Wu-chi Feng, Portland State University, USA  
Romulus Grigoras, IRIT-Toulouse, France  
Carsten Griwodz, University of Oslo, Norway  
Pål Halvorsen, University of Oslo, Norway  
JongWon Kim, GIST, Korea  
Baochun Li, University of Toronto, Canada  
Kang Li, University of Georgia, USA  
Yong Liu, Polytechnic University, USA  
Andreas Mauthe, Lancaster University, UK

Klara Nahrstedt, UIUC, USA  
Ketan Mayer-Patel, UNC Chapel Hill, USA  
Cristina Nita-Rotaru, Purdue University, USA  
Sanjay Rao, Purdue University, USA  
Reza Rejaie, University of Oregon, USA  
Karsten Schwan, Georgia Tech, USA  
Henning Schulzrinne, Columbia University, USA  
Prashant Shenoy, UMass, Amherst  
Shervin Shirmohammadi, University of Ottawa, Canada  
Ralf Steinmetz, TU Darmstadt, Germany  
Lars Wolf, TU Braunschweig, Germany  
Zhenyu Yang, Florida International University, USA  
Roger Zimmermann, National University of Singapore

Workshop Co-Chairs

**Dongyan Xu**  
Purdue University, USA

**Wei Tsang Ooi**  
National University of Singapore

Local Arrangement Chair

**Ketan Mayer-Patel**  
UNC Chapel Hill, USA

Publicity Co-Chairs

**Tristan Henderson**  
University of St. Andrews, UK

**Kang Li**  
University of Georgia, USA

**Kuan-Ta Chen**  
Academia Sinica, Taiwan

Technical Program Committee

**Kevin Almeroth**  
UC Santa Barbara, USA

**Grenville Armitage**  
Swinburne U. of Tech, Australia

**Ernst Biersack**  
Institute Eurecom, France

**Surender Chandra**  
University of Notre Dame, USA

**Kuan-Ta Chen**  
Academia Sinica, Taiwan

**Songqing Chen**  
George Mason University, USA

**Mark Claypool**  
WPI, USA

**Wu-chang Feng**  
Portland State University, USA

**Wu-chi Feng**  
Portland State University, USA

**Romulus Grigoras**  
IRIT-Toulouse, France

**Carsten Griwodz**  
University of Oslo, Norway

**Pål Halvorsen**  
University of Oslo, Norway

**JongWon Kim**  
GIST, Korea

**Baochun Li**  
University of Toronto, Canada

**Kang Li**  
University of Georgia, USA

**Yong Liu**  
Polytechnic University, USA

**Andreas Mauthe**  
Lancaster University, UK

**Klara Nahrstedt**  
UIUC, USA

**Ketan Mayer-Patel**  
UNC Chapel Hill, USA

**Cristina Nita-Rotaru**  
Purdue University, USA

**Sanjay Rao**  
Purdue University, USA

**Reza Rejaie**  
University of Oregon, USA

**Karsten Schwan**  
Georgia Tech, USA

**Henning Schulzrinne**  
Columbia University, USA

**Prashant Shenoy**  
UMass, Amherst

**Shervin Shirmohammadi**  
University of Ottawa, Canada

**Ralf Steinmetz**  
TU Darmstadt, Germany

**Lars Wolf**  
TU Braunschweig, Germany

**Zhenyu Yang**  
Florida International University

**Roger Zimmermann**  
National University of Singapore

# NOSSDAV 2009

## Call For Papers

The 19th International Workshop on  
Network and Operating Systems Support for Digital Audio and Video  
Williamsburg, Virginia, 3-5 June 2009

NOSSDAV 2009 will continue the workshop's long tradition of focusing on emerging topics, controversial ideas, and future research directions in the area of multimedia systems research, held in a setting that stimulates lively discussions among the senior and junior participants.

It is also an established practice for NOSSDAV to encourage experimental research based on real systems and data sets. Public availability of source code and data sets is highly encouraged.

The scope of NOSSDAV has evolved over the years. For NOSSDAV 2009, we would like especially highlight two new topics of interest: unconventional use of GPU for multimedia and multi-core processors support for multimedia. NOSSDAV 2009 continues to welcome submissions in the traditional topics of networked multimedia systems, operating system support for multimedia, and multimedia security and rights management.

The topics of interest include, but are not limited to:

- OS, middleware, and network support for multimedia
- Overlay networks for multimedia
- Media streaming, distribution, and storage
- Web 2.0 systems and social networks for multimedia
- Wireless and mobile multimedia systems
- Media sensor and ad hoc networks
- Grid/cloud computing support for multimedia
- Embedded systems for multimedia
- Multi-core architecture support for multimedia
- GPU for multimedia
- Network processor support for multimedia
- Tele-immersion and mixed-reality systems
- Networked graphics and virtual environments
- Networked games
- Multimedia communications and system security
- Digital rights management

A broad view will be taken in deciding what topics are within scope. Please contact the workshop co-chairs if you wish to check if a particular topic is within the scope of NOSSDAV.

Submissions should be at most SIX pages in length using standard ACM proceedings style. We expect these submissions to be the kernel of what will eventually lead to full-length papers at high-quality conferences or journals.

Authors of selected, high quality papers from NOSSDAV 2009 will be invited to submit an extended version of their papers to a special issue of ACM/Springer Multimedia Systems Journal.

Paper Deadline **16 February 2009, 23:59 PST (Firm)**  
Notification **20 March 2009**  
Camera Ready Due **6 April 2009**

Sponsored By:



In Cooperation With :



SIGCOMM



SIGOPS