

# NSDI 2014 Conference Report

Tim Wood

A decorative border at the bottom of the slide consisting of a row of stylized, overlapping clouds in various shades of light blue and white.

# SSL Updates

Next week: Henning Schulzrinne CTO of FCC

- “Transitioning the Public Switched Telephone Network to IP”
- **at noon not 12:30!**

SSL Website: [\*\*http://cloudlab.seas.gwu.edu/ssl\*\*](http://cloudlab.seas.gwu.edu/ssl)

SSL Student Organizer - volunteer?

- Organize speakers for next semester
- Maintain web site
- Add a line to your resume

# Why go to a conference?

To present your own ideas

- could just make a you tube video

To hear other people's ideas

- could just read their papers

To network and meet other people

- connections do matter
- practice introducing yourself, giving an elevator pitch, and asking questions after talks

# What to do at a conference?

Listen

Make notes

Ask questions

Meet people

Think

The screenshot shows the Zotero application window. The search bar at the top right contains 'NSDI 14'. The left sidebar shows a folder structure under 'Group Libraries' with 'gwcloud' selected. The main pane displays a list of papers with columns for Title, Creator, Publication, Year, and Date. The right pane shows a rich text editor with a paragraph of text about 'TIM: Key Value store' and a 'Paragraph' dropdown menu.

Title	Creator	Pub...	Year	Da...
NetVM: High Performance and Flexible Net...	Hwang et al.	Sym...	2014	4/...
MICA: A Holistic Approach to Fast In-Mem...	Lim et al.	Pro...	2014	4/...
TIM: Key Value store focused on small it...				4/...
nsdi14-paper-lim.pdf				4/...
nsdi14-paper-lim.pdf				4/...
Blizzard: Fast, Cloud-scale Block Storage f...	Mickens et al.	Pro...	2014	4/...
TIM: Goal is to provide a cloud data stor...				4/...
nsdi14-paper-mickens-james.pdf				4/...
High Throughput Data Center Topology De...	Singla et al.	Pro...	2014	4/...
Facebook has more inter-cluster networ...				4/...
nsdi14-paper-singla.pdf				4/...
SENIC: Scalable NIC for End-Host Rate Limi...	Radhakrishna...	Pro...	2014	4/...
TIM:				4/...
nsdi14-paper-radhakrishnan.pdf				4/...
FaRM: Fast Remote Memory	Dragojević et ...	Pro...	2014	4/...
TIM: remote access to memory over RD...				4/...
nsdi14-paper-dragojevic.pdf				4/...
Exalt: Empowering Researchers to Evaluate...	Wang et al.	Pro...	2014	4/...
TIM:				4/...
nsdi14-paper-wang_yang.pdf				4/...
Aggregation and Degradation in JetStream:...	Rabkin et al.	Pro...	2014	4/...
TIM: Sending a stream of sensor data to...				4/...
nsdi14-paper-rabkin.pdf				4/...
Operational Experiences with Disk Imaging...	Atkinson et al.	Pro...	2014	4/...
TIM: Paper on disk contents of VMs in t...				4/...
nsdi14-paper-atkinson.pdf				4/...
Network Virtualization in Multi-tenant Dat...	Koponen et al.	Pro...	2014	4/...
TIM: Paper describes VMware's virtual n...				4/...
nsdi14-paper-koponen.pdf				4/...
mTCP: a Highly Scalable User-level TCP Sta...	Jeong et al.	Pro...	2014	4/...
TIM:				4/...
nsdi14-paper-jeong.pdf				3/...
ClickOS and the Art of Network Function Vi...	Martins et al.	Pro...	2014	4/...
TIM:				4/...

Paragraph

TIM: Key Value store focused on small items (fit in one packet). Other systems are slower, and are really slow if there is a high degree of writes. They get 60M+ req/sec even with 50/50 read/write workloads. They require the clients to do request direction (to a specific core). Optimize system to support high write rates.

Log based allocation is also used for eviction--does that mean they don't do LRU at all? Paper says that use FIFO, but could re-add entries on access to emulate LRU.

Tested memcached+DPDK and had similar/worse performance to regular memcached. Why?

Could you use jumbo frames to allow larger values to be sent in one packet?

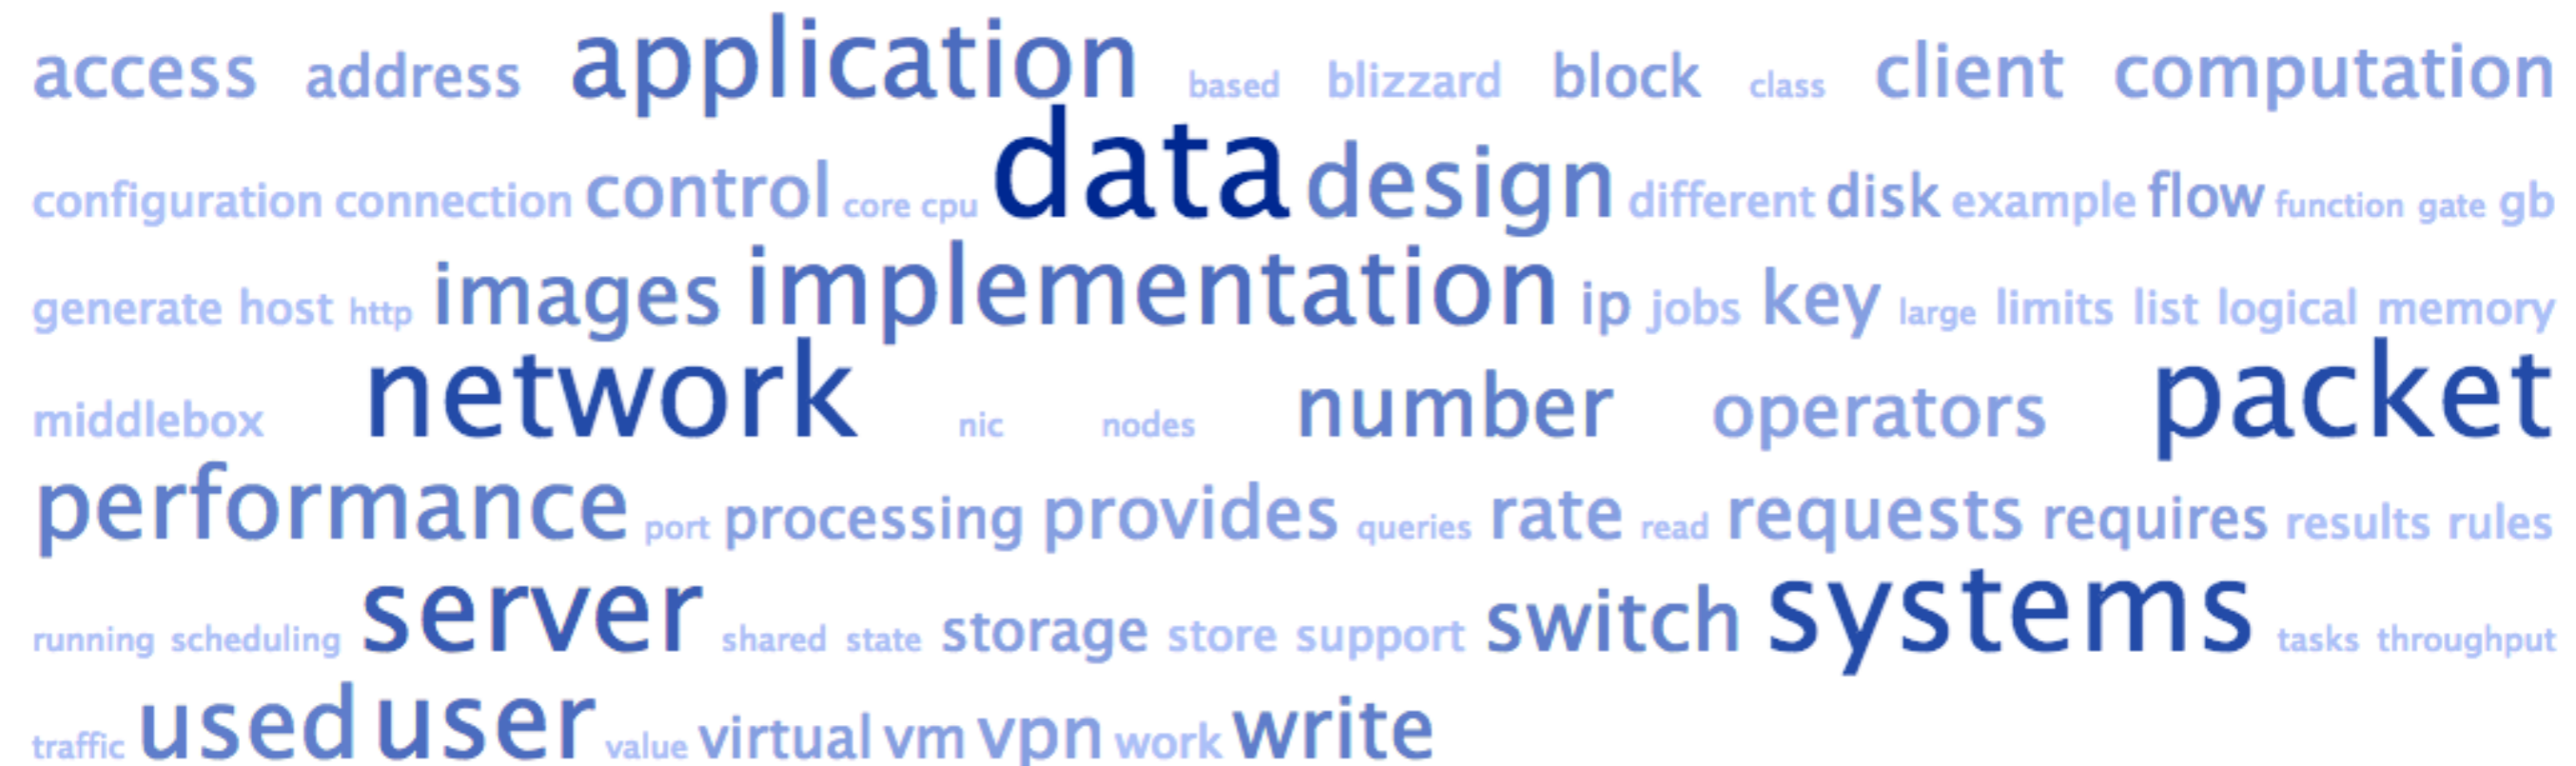
Use YCSB for workload generation

Related: [click here]  
Tags: [click here]  
Edit in a separate window

# What is NSDI?

Networking papers with a strong systems component

- Targets: Data centers, WAN, mobile, sensor networks
- Rank: OSDI/SOSP/Sigcomm, NSDI/area focused
- Papers: 12 pages + references, 38 of 223 accepted
- Size: 3 days (single track), ~250 attendees
- History: 10 years old, first conference has H-index ~20



# Sessions

## Data Center Networks

- Switches and topologies

## Scalable Networking

- Software-based network services

## Improving Throughput and Latency

- Sensor networks, protocols

## Operational Systems

- Networks and disks in data centers, VPN privacy

## Data Storage and Analytics

- scalable storage and data analysis

## In-Memory Computing and Caching

- RDMA, caching semantics, scalable key value stores

## Debugging Complex Systems

- Root cause analysis, mobile, scalable packet trace analysis

## Software Verification and Testing

- Tracking correctness and performance of networks and storage

## Security and Privacy

- mobile, web apps, clouds, physical network

## New Programming Abstractions

- programming languages, APIs

## Interpreting Signals

- body and motion tracking

# Data Center Networks

## Circuit Switching Under the Radar with REACToR

- He Liu, Feng Lu, Alex Forencich, Rishi Kapoor, Malveeka Tewari, Geoffrey M. Voelker, George Papen, Alex C. Snoeren, and George Porter, University of California, San Diego
- **Mixing optical and ethernet network HW to support both very large and very small network flows**

## Catch the Whole Lot in an Action: Rapid Precise Packet Loss Notification in Data Center

- Peng Cheng, Fengyuan Ren, Ran Shu, and Chuang Lin, Tsinghua University

## High Throughput Data Center Topology Design

- Ankit Singla, P. Brighten Godfrey, and Alexandra Kolla, University of Illinois at Urbana-Champaign
- **Properties of network topologies, benefits of randomness**

# Scalable Networking

## NetVM: High Performance and Flexible Networking Using Virtualization on Commodity Platforms

- Jinho Hwang, The George Washington University; K. K. Ramakrishnan, Rutgers University; Timothy Wood, The George Washington University
- **High speed network services inside VMs**

## ClickOS and the Art of Network Function Virtualization

- Joao Martins and Mohamed Ahmed, NEC Europe Ltd.; Costin Raiciu and Vladimir Olteanu, University Politehnica of Bucharest; Michio Honda, Roberto Bifulco, and Felipe Huici, NEC Europe
- **High speed network services inside VMs**

## mTCP: a Highly Scalable User-level TCP Stack for Multicore

- EunYoung Jeong, Shinae Woo, Muhammad Jamshed, and Haewon Jeong, (KAIST); Sunghwan Ihm, Princeton; Dongsu Han and KyoungSoo Park, (KAIST)
- **High speed network services inside user space apps**

## SENIC: Scalable NIC for End-Host Rate Limiting



# Improving Throughput & Latency

Enabling Bit-by-Bit Backscatter Communication in Severe Energy Harvesting Environments

- Pengyu Zhang and Deepak Ganesan, University of Massachusetts Lowell

Full Duplex MIMO Radios

- Dinesh Bharadia and Sachin Katti, Stanford University

Recursively Cautious Congestion Control

- Radhika Mittal, Justine Sherry, and Sylvia Ratnasamy, University of California, Berkeley; Scott Shenker, University of California, Berkeley, and International Computer Science Institute

How Speedy is SPDY?

- Xiao Sophia Wang, Aruna Balasubramanian, Arvind Krishnamoorti, and David Wetherall, University of Washington

Hardware focused

Protocol focused

# Operational Systems

## Network Virtualization in Multi-tenant Datacenters

- Teemu Koponen, Keith Amidon, Peter Balland, Martín Casado, Anupam Chanda, Bryan Fulton, Igor Ganichev, Jesse Gross, Natasha Gude, Paul Ingram, Ethan Jackson, Andrew Lambeth, Romain Lenglet, Shih-Hao Li, Amar Padmanabhan, Justin Pettit, Ben Pfaff, and Rajiv Ramanathan, VMware; Scott Shenker, International Computer Science Institute and the University of California, Berkeley; Alan Shieh, Jeremy Stribling, Pankaj Thakkar, Dan Wendlandt, Alexander Yip, and Ronghua Zhang, VMware

## Operational Experiences with Disk Imaging in a Multi-Tenant Datacenter

- Kevin Atkinson, Gary Wong, and Robert Ricci, University of Utah

## VPN Gate: A Volunteer-Organized Public VPN Relay System with Blocking Resistance for Bypassing Government Censorship Firewalls

- Daiyuu Nobori and Yasushi Shinjo, University of Tsukuba

# Data Storage & Analytics

## Blizzard: Fast, Cloud-scale Block Storage for Cloud-oblivious Applications

- James Mickens, Edmund B. Nightingale, Jeremy Elson, and Darren Gehring, Microsoft Research; Bin Fan, Carnegie Mellon University; Asim Kadav and Vijay Chidambaram, University of Wisconsin—Madison; Osama Khan, Johns Hopkins University
- **Scalable IO for unmodified apps. Reorder writes, but still give consistency**

## Aggregation and Degradation in JetStream: Streaming Analytics in the Wide Area

- Ariel Rabkin, Matvey Arye, Siddhartha Sen, Vivek S. Pai, and Michael J. Freedman, Princeton
- **How to do filtering/summarization of sensor data at network edge**

## GRASS: Trimming Stragglers in Approximation Analytics

## Bolt: Data Management for Connected Homes

# In-memory Computing & Caching

## FaRM: Fast Remote Memory

- Aleksandar Dragojević, Dushyanth Narayanan, Orion Hodson, and Miguel Castro; Microsoft Research
- **RDMA with some nice data structure tricks**

## Easy Freshness with Pequod Cache Joins

- Bryan Kate, Eddie Kohler, and Michael S. Kester, Harvard; Neha Narula, Yandong Mao, and Robert Morris, MIT/CSAIL

## MICA: A Holistic Approach to Fast In-Memory Key-Value Storage

- Hyeontaek Lim, Carnegie Mellon University; Dongsu Han, KAIST; David G. Andersen, CMU; Michael Kaminsky, Intel Labs
- **User-level networking + very parallelizable data structure**

# Debugging Complex Systems

Adtributor: Revenue Debugging in Advertising Systems

- MS

DECAF: Detecting and Characterizing Ad Fraud in Mobile Apps

- USC, MSR

I Know What Your Packet Did Last Hop: Using Packet Histories to Troubleshoot Networks

- Nikhil Handigol, Brandon Heller, Vimalkumar Jeyakumar, David Mazières, and Nick McKeown, Stanford University
- **Store everything, diagnose later using map reduce style processing**

Libra: Divide and Conquer to Verify Forwarding Tables in Huge Networks

- Hongyi Zeng, Stanford University; Shidong Zhang and Fei Ye, Google; Vimalkumar Jeyakumar, Stanford University; Mickey Ju and Junda Liu, Google; Nick McKeown, Stanford University; Amin Vahdat, Google and UCSD
- **Verification using map reduce style processing**

# SW Verification & Testing

## Software Dataplane Verification

- Mihai Dobrescu and Katerina Argyraki, École Polytechnique Fédérale de Lausanne. *Awarded Best Paper!*
- **Tricks for reducing complexity when verifying networks**

## NetCheck: Network Diagnoses from Blackbox Traces

- Yanyan Zhuang, Polytechnic Institute of New York University and University of British Columbia; et al
- **Gather sys trace data and use to help diagnose**

## Exalt: Empowering Researchers to Evaluate Large-Scale Storage Systems

- Yang Wang, Manos Kapritsos, Lara Schmidt, Lorenzo Alvisi, and Mike Dahlin, The University of Texas at Austin
- **Allow large scale performance testing by faking it**

# Security and Privacy

## ipShield: A Framework For Enforcing Context-Aware Privacy

- Supriyo Chakraborty, Chenguang Shen, Kasturi Rangan Raghavan, Yasser Shoukry, Matt Millar, and Mani Srivastava, University of California, Los Angeles
- **Mobile phone security preferences**

## Building Web Applications on Top of Encrypted Data Using Mylar

- Raluca Ada Popa, MIT/CSAIL; Emily Stark, Meteor, Inc.; Steven Valdez, Jonas Helfer, Nickolai Zeldovich, and Hari Balakrishnan, MIT/CSAIL
- **Key management for client/server encryption coordination**

## PHY Covert Channels: Can you see the Idles?

- Ki Suh Lee, Han Wang, and Hakim Weatherspoon, Cornell University

## cTPM: A Cloud TPM for Cross-Device Trusted Applications

- Chen Chen, Carnegie Mellon University; Himanshu Raj, Stefan Saroiu, and Alec Wolman, Microsoft Research
- **Key management, cloud-based TPM services: storage, clock**

# Programming Abstractions

## Warranties for Faster Strong Consistency

- Jed Liu, Tom Magrino, Owen Arden, Michael D. George, and Andrew C. Myers, Cornell University

## Tierless Programming and Reasoning for Software-Defined Networks

- Tim Nelson, Andrew D. Ferguson, Michael J.G. Scheer, and Shriram Krishnamurthi, Brown University

## Enforcing Network-Wide Policies in the Presence of Dynamic Middlebox Actions using FlowTags

- Seyed Kaveh Fayazbakhsh, Carnegie Mellon University; Luis Chiang, Deutsche Telekom Labs; Vyas Sekar, Carnegie Mellon University; Minlan Yu, University of Southern California; Jeffrey C. Mogul, Google



# Overall

Lots of low-level systems work to support distributed systems or networking infrastructure

Some higher level work on verification, privacy, security

Use a centralized control plane to manage distributed, stupid elements

- network verification, cloud TPM, cloud storage (from SDNs, MapReduce)

Push intelligence to end-nodes

- User space networking, client-side encryption, smarter KV store clients