



Center for Networked Systems

UC San Diego

October 2014 Research Review

Director: Stefan Savage

Associate Director: George Porter

Welcome!

- Completing our 10th year!
- Over 40+ projects completed
 - 20+ PIs
 - 100+ students
 - 12 sponsors
 - Seeded ~\$40M in federal funds



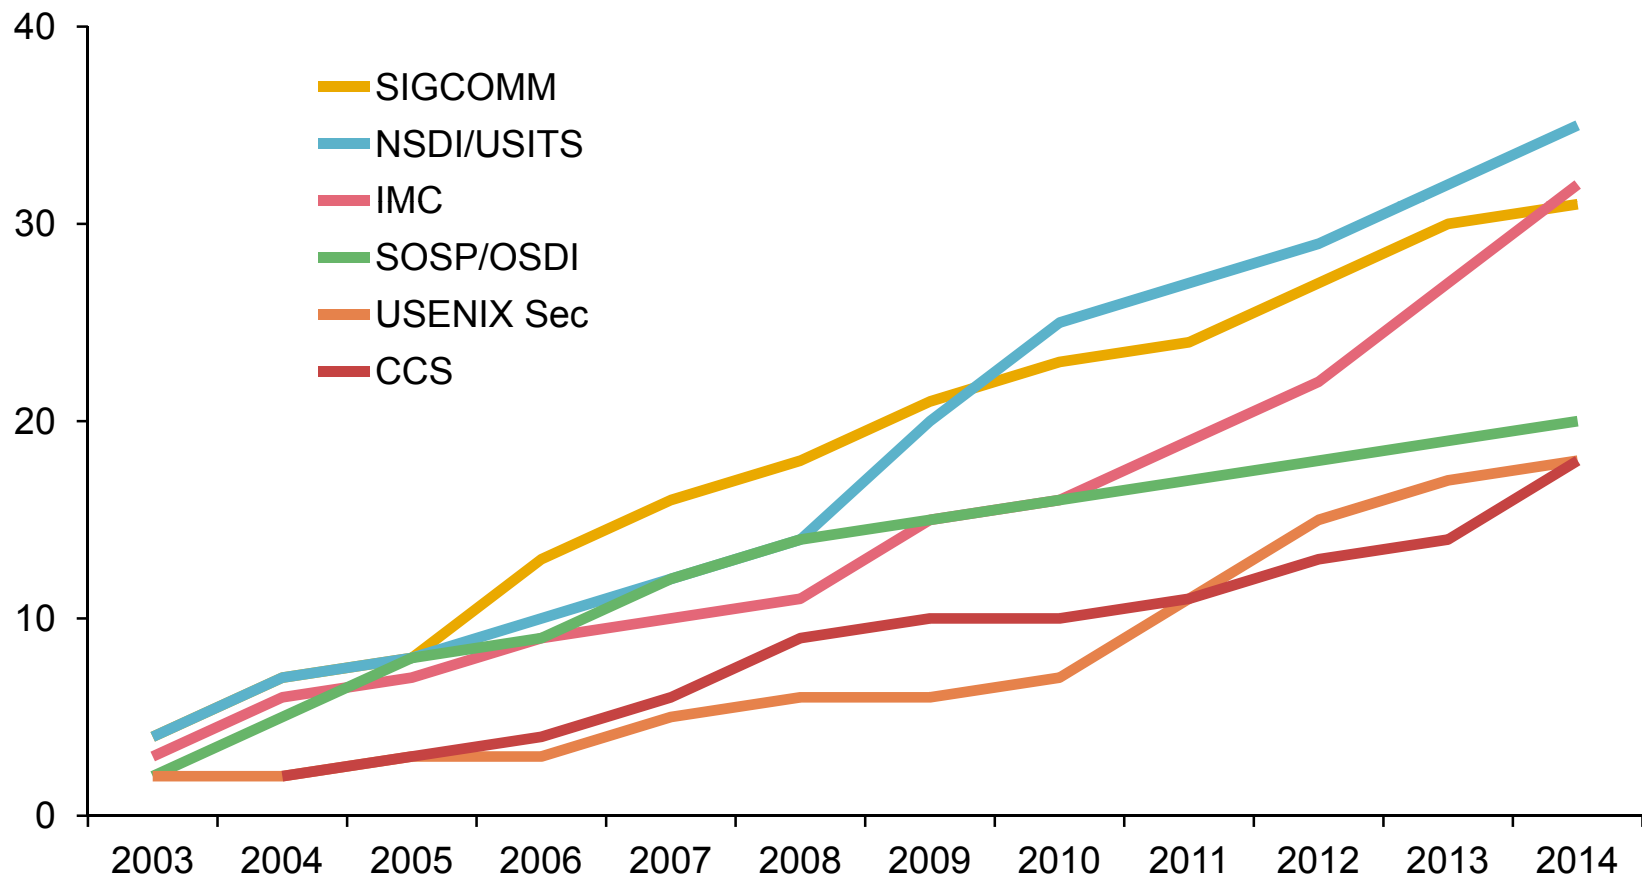
What is CNS about?

- We focus on emerging challenges in **networked systems** platforms and environments
 - Build, test, measure, evaluate **real systems**
 - Deliver results that have **impact**
- Play to strengths of members & partners
 - Networking, data centers, measurement
 - Security, data analysis

How CNS works?

- We succeed through collaboration with **you**
- CNS is really a facilitator
 - Help to match faculty interests/expertise for sponsor interests/needs
 - These research reviews are one mechanism
 - Place to find out what we've *done* and talk about what we might do **next**
- Projects funded directly with PIs through CNS
 - Membership info:
<http://cns.ucsd.edu/memberbenefits.shtml>

Measures of success: Major CNS publication venues



Measures of success: Recent CNS Faculty Awards

POPL Test of Time Award



Ranjit Jhala

IEEE Internet Award



kc Claffy



Factoid: Six CNS faculty have been named ACM Fellows

Measures of success: CNS Grad placement!



123 CNS internships over the last decade

69 in the last five years!



Measures of success: Grooming new researchers

- Projects in Security and DC hosting UCSD's new "**Early Research Scholars Program**"
- Embedded into groups as **freshmen!**
- **Security**
 - Aaron Hurtado
 - Tiffany Allen
 - Kelsey Ma
 - Andrew Jabasa
- **Data centers**
 - Luis Sanchez
 - Edgar Lopez
 - Huayin Zhou
 - Mingshan Wang

Measures of success: research accomplishments

- Some highlights from the last decade

- **Data Centers**


- Network scalability
- Data processing scalability

- **Security**

- Evidence-based interventions
- Embedded security

Building the next generation of data center networks: FatTrees


- Scalable packet switching to 100k+ servers



Sigcomm 2008, Sigcomm 2009, NSDI 2010, IEEE Micro

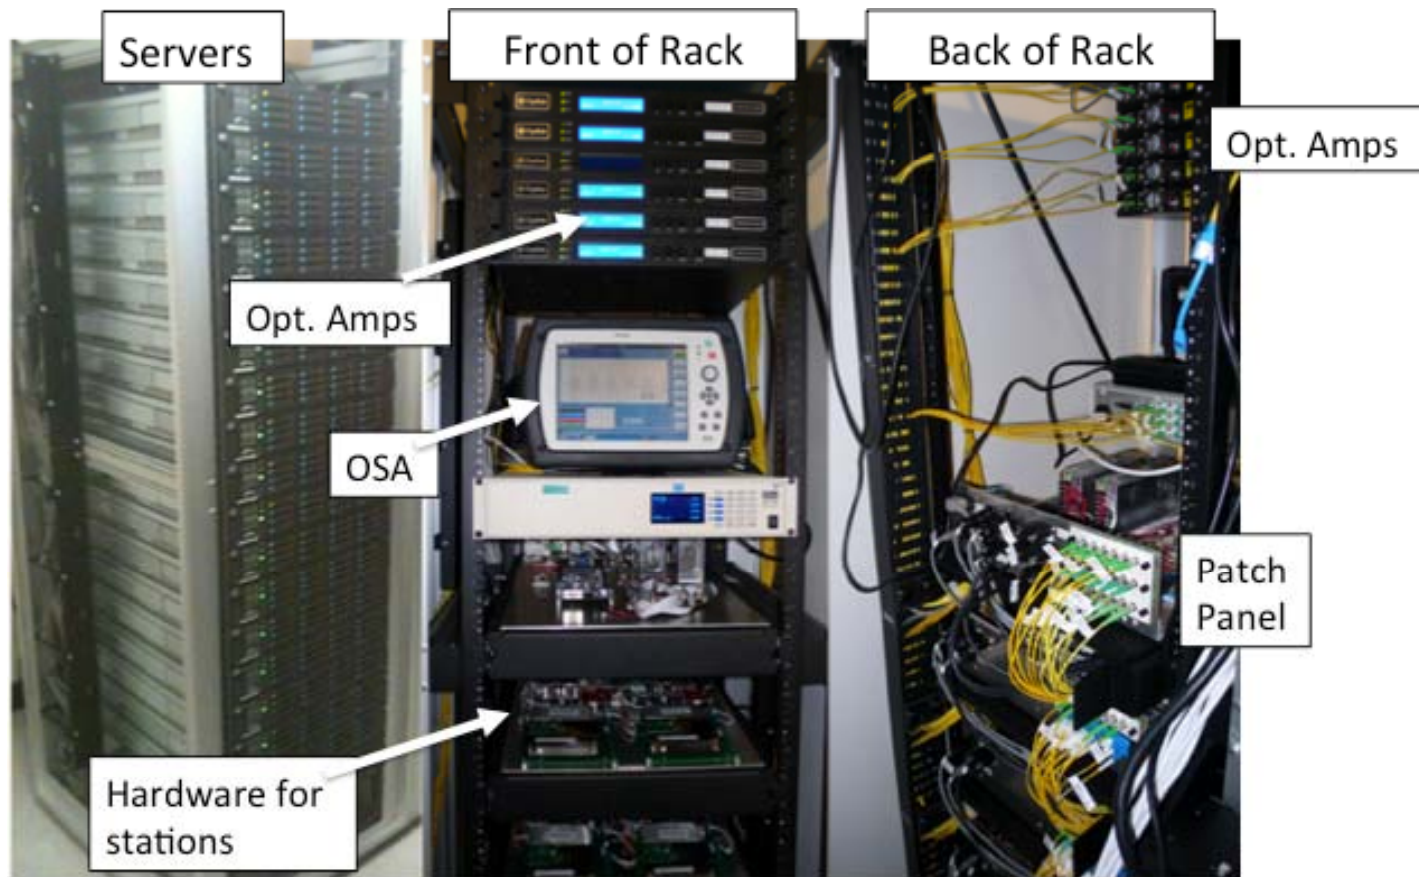
Building the next generation of data center networks: Helios

- Lowering costs via optics in the network core



Building the next generation of data center networks: Mordia


- Lowering costs via optics throughout network



Sigcomm 2013

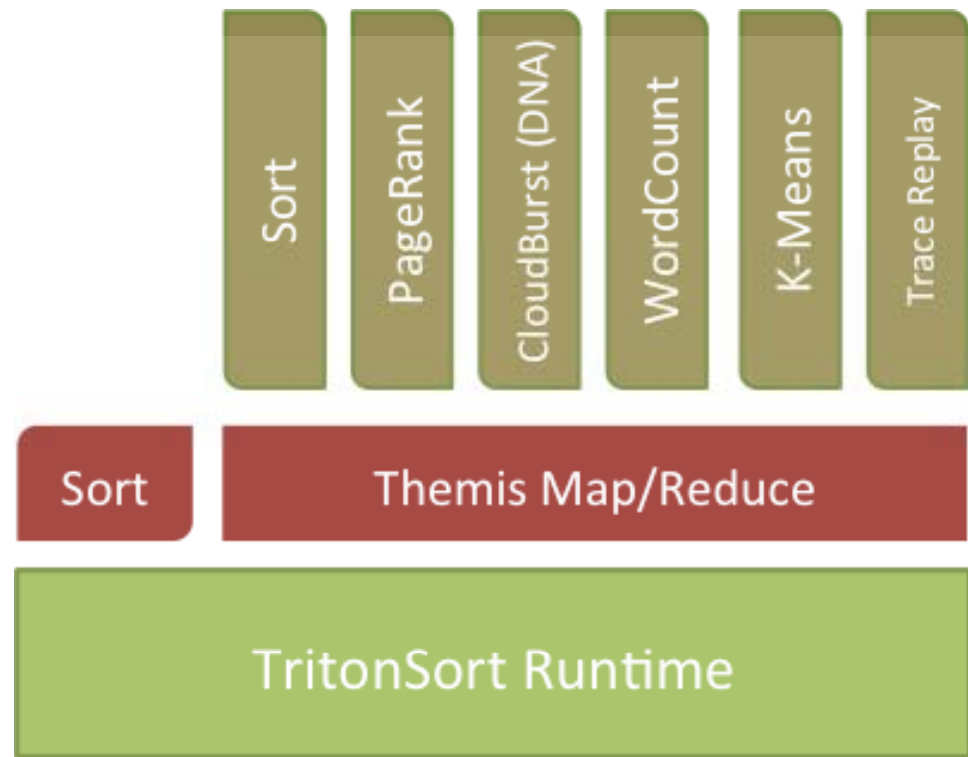
Building the next generation of data center networks: REACToR

- Circuit-switch capable top-of-rack switches




Building *efficient* scalable data processing systems

- Existing software stacks horribly inefficient
- *Balanced hardware* drives all resources as close to 100% as possible
- *Balanced software* to fully exploit this hardware
- Focus on sorting, then generalizing to MapReduce



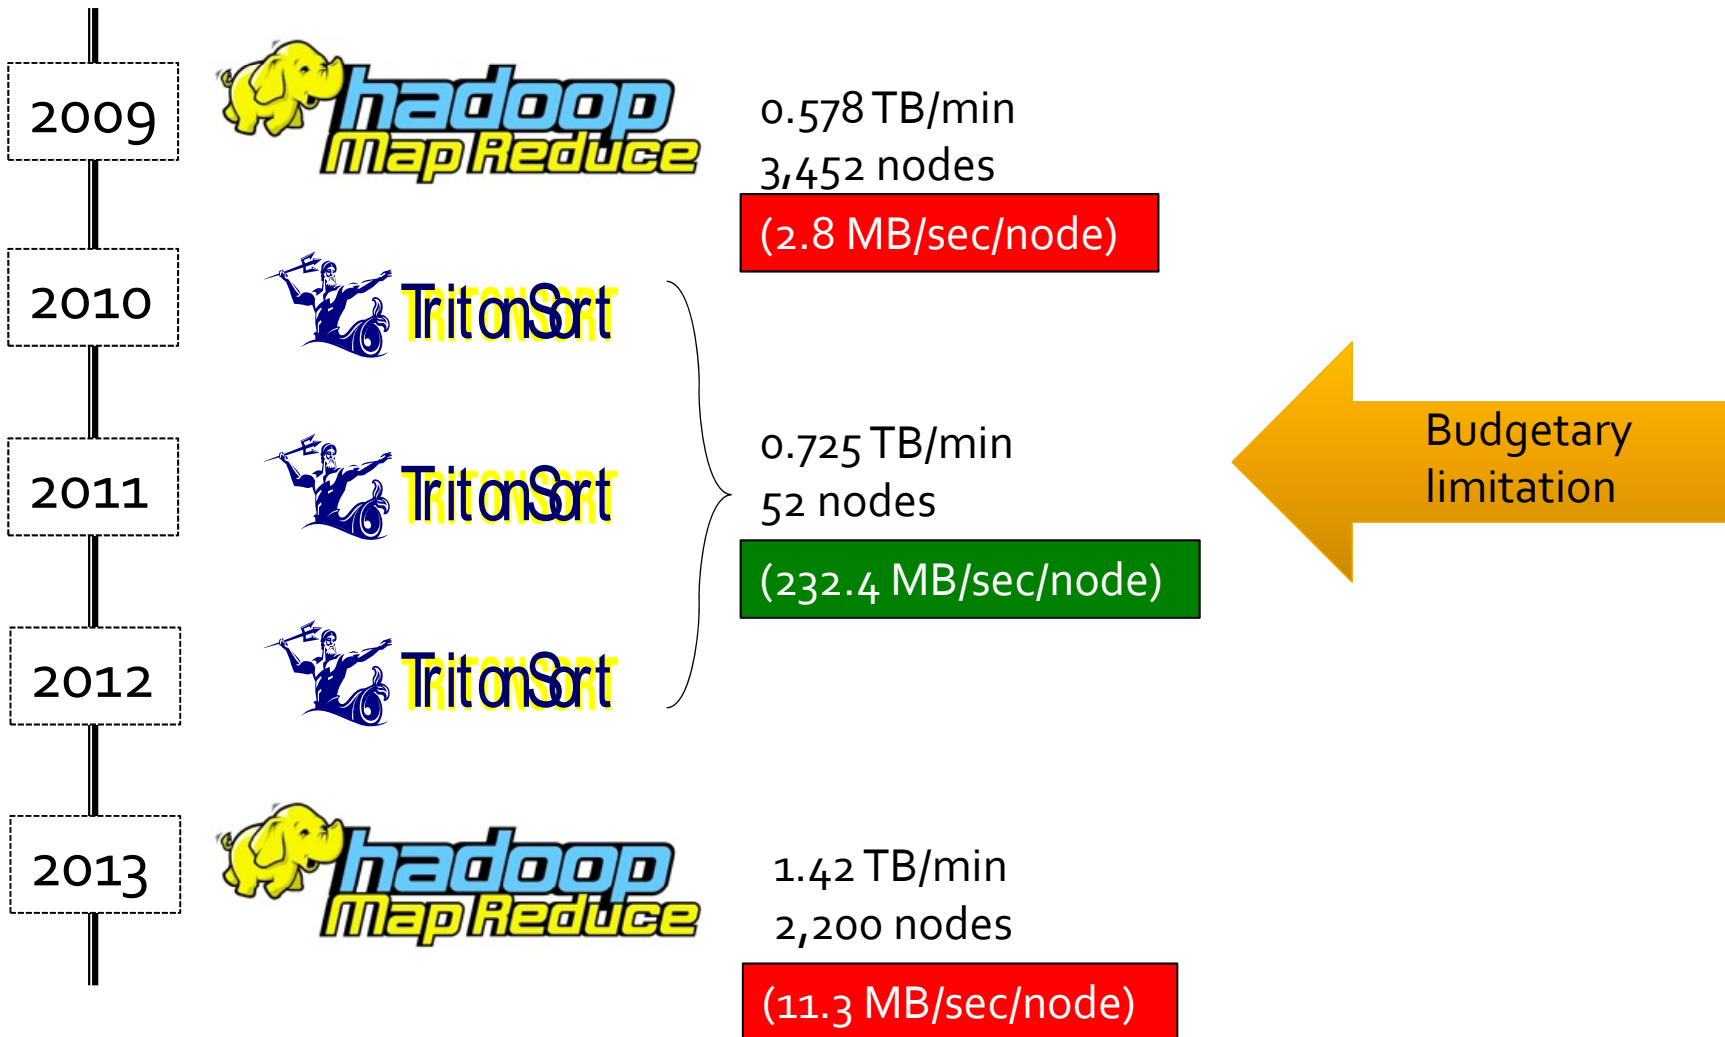
How to measure Data Processing efficiency?

- Sorting benchmark
(sortbenchmark.org)
- How long to sort a big dataset?
(i.e., 100TB)
- Various versions
 - Daytona, Indy, Minute, etc
 - Traditionally on dedicated clusters
 - Big clusters = big winners?



Past World Records in per-node processing efficiency

100TB GraySort





Rumor has it...

- 2014 Themis TritonSort (now with 100% cloud)
 - Indy Graysort: <15m w/178 EC2 nodes
 - 6.76 TB/min, ~633MB/sec/node
 - Daytona Graysort: <23min, w/186 EC2 nodes
 - 4.35 TB/min, ~390MB/sec/node
 - Indy/Daytona Cloudsort, <\$450/100TB w/330 EC2 nodes
- Fastest, cheapest and most efficient?
- Mike will explain later...

Best measure of success: research accomplishments

- Some highlights from the last decade

- **Data Centers**

- Network scalability
- Data processing scalability

- **Security**

- Evidence-based interventions
- Embedded security

Evidence-based intervention

- Lots of money spent to counter online threats
 - But which interventions work?
 - Which are cost effective?
 - Which undermine criminal business models?
- We've been working for 10 years to place these questions on an **empirical footing**
- Some important results:
 - Blacklisting/takedown/etc limited
 - Payment systems are key

Four years of “following the money”

- **2011:** our group identifies the “payment bottleneck” in spam/abuse
- **2012:** First large-scale demonstration of payment intervention by Microsoft DCU
 - Subsequent efforts with pharma sector
- **2013:** Obama administration endorses IACC payment portal as IP best practices
 - Interpol adds payment to Pangea V
- **2014:** Demonstrate deanonymization of Bitcoin flows; help frame policy



Embedded security

- The “Internet of Things” -> Insecurity of Things?
- How should we think about security in embedded devices?
 - Side effects
 - Software engineering impacts
 - Supply chain issues
 - Disclosure, updating
 - Technology vs cost structure

Automotive security

- Automotive security work (w/UW)
 - Identified range of **internal vulnerabilities** in automotive systems [Koscher et al, Oakland '10]
 - **External vulnerabilities** allowing remote compromise/control [Checkoway et al, USEC '11]
- Major impacts on US automotive industry



SAE *International*®

USCAR

★★★★★
NHTSA
www.nhtsa.gov

Broad range of efforts

- Voting machines [EVT '2009]
- Peripherals [WOOT '14]
- Embedded PRNG [Oakland '13]
- Aviation [CCS '14 & more to come]
- Later today:
 - X-ray body scanners
 - Provable control systems



Tip of iceberg for CNS

- We do work in
 - Storage design and NV memories
 - NIC design
 - Ad fraud/abuse
 - Automated detection of configuration errors
 - Post-mortem fault diagnosis
 - Energy efficiency
 - Network measurement & inference
 - And then some...

For the next two days:

- Listen, kibitz, critique, suggest, network...
- Continuing event: Student lighting round!
- Two poster sessions!
 - Tonight
 - Tomorrow at lunch
- The goal of all this is not for us to show off (although we'll do that if we can) but to find points of mutual interest



Center for Networked Systems

UC San Diego

Thank you!

