Hot Topics: Operating Systems

CSE Job Hunters 2010

Michael Vrable

University of California, San Diego

February 9, 2010

What is (Operating) Systems Research?

- Operating System design
- File and storage systems
- Reliability/security/fault tolerance
- Power efficiency
- Systems management
- Distributed systems

What is an Operating System

- Software that manages resources, provides infrastructure for other code as it executes

What is an Operating System

- Software that manages resources, provides infrastructure for other code as it executes
- Traditional operating systems

What is an Operating System

- Software that manages resources, provides infrastructure for other code as it executes

- Traditional operating systems
- Virtual machine monitors
- Web browsers
- Sensor mote OSes
- Distributed software layers

What are Hot Research Areas in OS?

- Designing OSes for multicore/heterogenous multicore systems
- Scalability, especially to datacenter- or Internet-scale
- Reliability: Finding/avoiding bugs, making software more tolerant of failures
- Support for debugging production software
- New non-volatile storage (flash, PCM)

“Real Systems” Papers

- Google File System (GFS), MapReduce, Chubby, BigTable (Google)
- Dynamo (Amazon)
- DryadLINQ (Microsoft)

Recent Work

From SOSP 2009:

- FAWN: Fast Array of Wimpy Nodes (CMU)

Recent Work

From SOSP 2009:

- FAWN: Fast Array of Wimpy Nodes (CMU)
- RouteBricks (EPFL, Intel Research)

Recent Work

From SOSP 2009:

- FAWN: Fast Array of Wimpy Nodes (CMU)
- RouteBricks (EPFL, Intel Research)
- Multikernel (ETH Zurich, MSR Cambridge)

Recent Work

From SOSP 2009:

- FAWN: Fast Array of Wimpy Nodes (CMU)
- RouteBricks (EPFL, Intel Research)
- Multikernel (ETH Zurich, MSR Cambridge)
- PRES: Probabilistic Replay with Execution Sketching (UIUC/UCSD)
- ODR: Output-Deterministic Replay (UC Berkeley)

Recent Work

From SOSP 2009:

- FAWN: Fast Array of Wimpy Nodes (CMU)
- RouteBricks (EPFL, Intel Research)
- Multikernel (ETH Zurich, MSR Cambridge)
- PRES: Probabilistic Replay with Execution Sketching (UIUC/UCSD)
- ODR: Output-Deterministic Replay (UC Berkeley)
- Tolerating Hardware Device Failures in Software (University of Wisconsin Madison)

What is UCSD Working On?

- Yuanyuan (YY) Zhou and her group have done quite a bit of work on system reliability:
  - Replay debugging (e.g., PRES)
  - Code analysis to find bugs in software (iComment, MUVI, etc.) system reliability:
- Voeler, Savage, Vahdat, Snoeren, Varghese: Xen hypervisor (Difference Engine, Neon)
- Rajesh Gupta, Tajana Rosing, Amin Vahdat: MuSyC (Multiscale Systems Center), energy-efficiency in large datacenters
- Vahdat: Building balanced systems