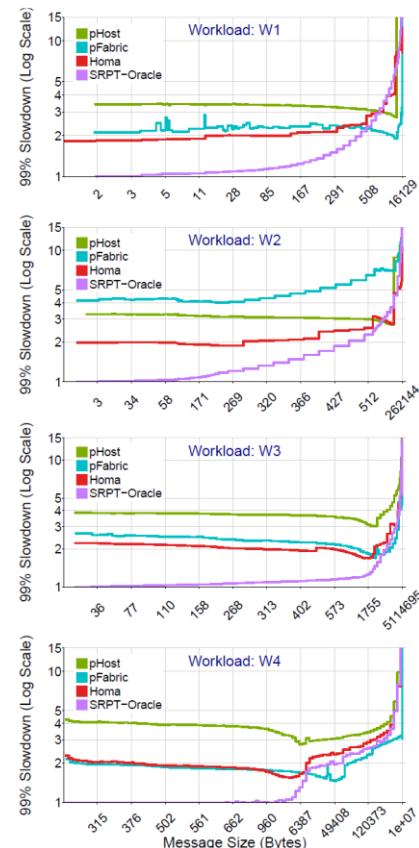


Panel on Granular Computing



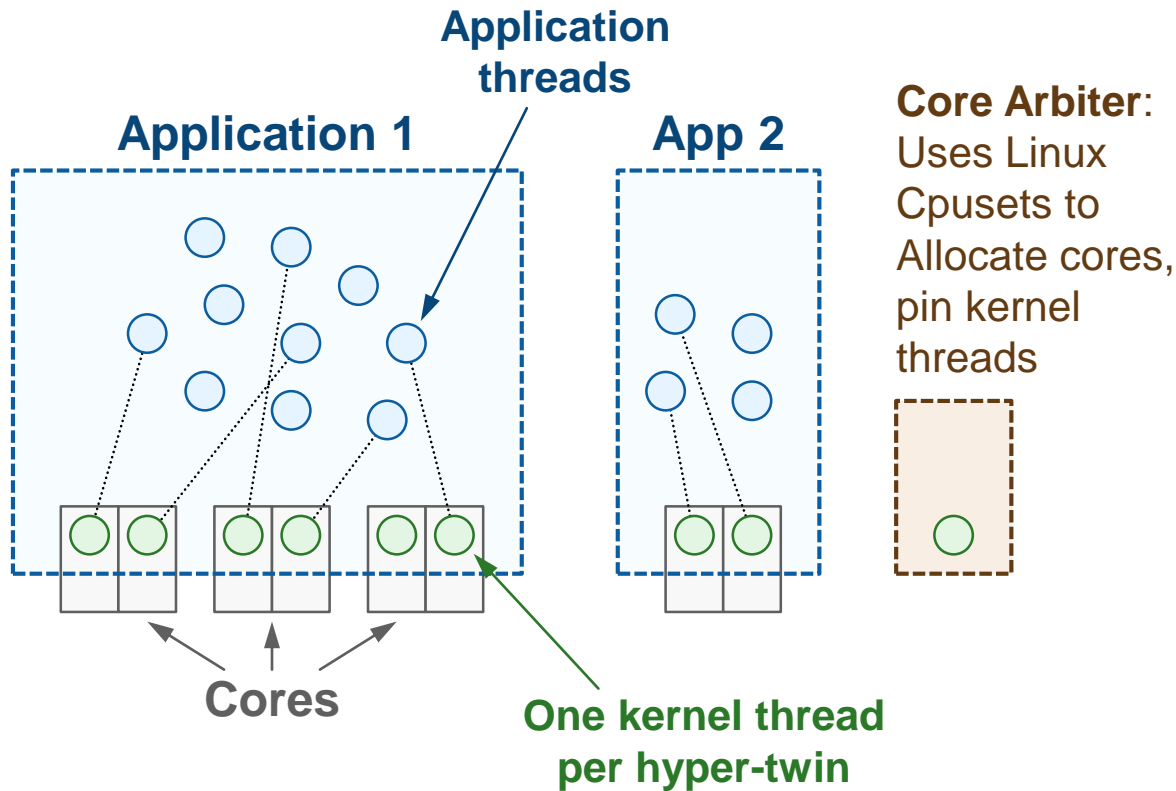
Homa: New Datacenter Transport

- **Consistent low latency for short messages:**
 - 99%-ile slowdown < 2.5x at 80% network load
- **Supports high network loads: 90%**
- **Outperforms pFabric at high loads (and works with existing switches)**
- **Key techniques:**
 - Receivers schedule incoming packets
 - Minimize buffer occupancy
 - Use 4-8 network priority levels for preemption
 - New solution to bipartite matching problem: controlled buffering + priorities
- **Talk with Behnam Montazeri, Yilong Li**



Arachne: New Thread/Core Mgmt

- Start thread on new core: 130 ns
- Wake up sleeping thread: 80 ns
- Better core utilization for granular apps
- Thread/core allocation entirely at user level
- Application “owns” cores for long intervals (10-100 ms)?
- Talk with Henry Qin, Jacqueline Speiser



NanoLog

- **High-performance logging for low-latency systems**
- **Supports log messages at granularity of 10-50ns**
- **Stay tuned for Stephen Yang's talk**

Granular Notifications

- **Notifications: dual of remote procedure call (RPC)**
 - RPC: sender-driven
 - Notifications: receiver-driven (subscribe to certain kinds of events)
 - Notifications are also potentially multicast
- **Likely to be important for Big Control (event driven)**
- **New project to create high-performance notification mechanism:**
 - Ultra-low latency
 - High throughput
- **Talk with Collin Lee**

Audience Questions

- **What would you do with:**
 - 20 μ s persistent remote task
 - 500 ns local task
- **What advice would you give us? Concerns? Issues to address?**

Questions / Discussion

