

The Network Stack (1)

Lecture 5, Part 3: Network Stack Research Taster


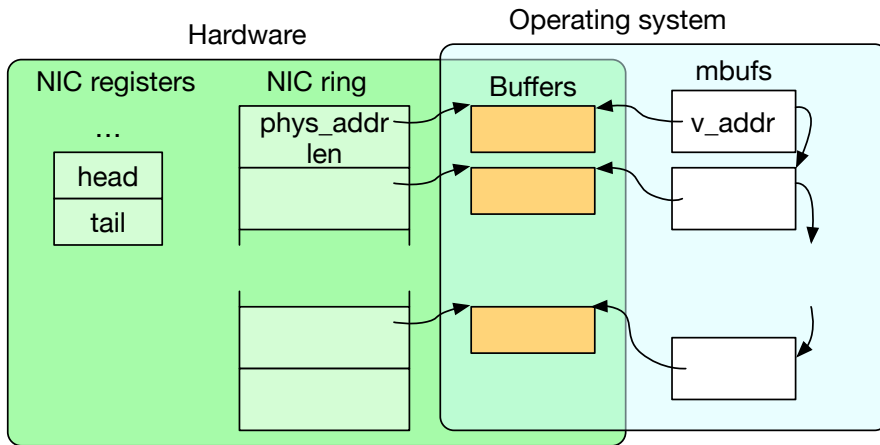
Dr Robert N. M. Watson

2020-2021



Netmap: a novel framework for fast packet I/O

Luigi Rizzo, USENIX ATC 2012 (best paper).

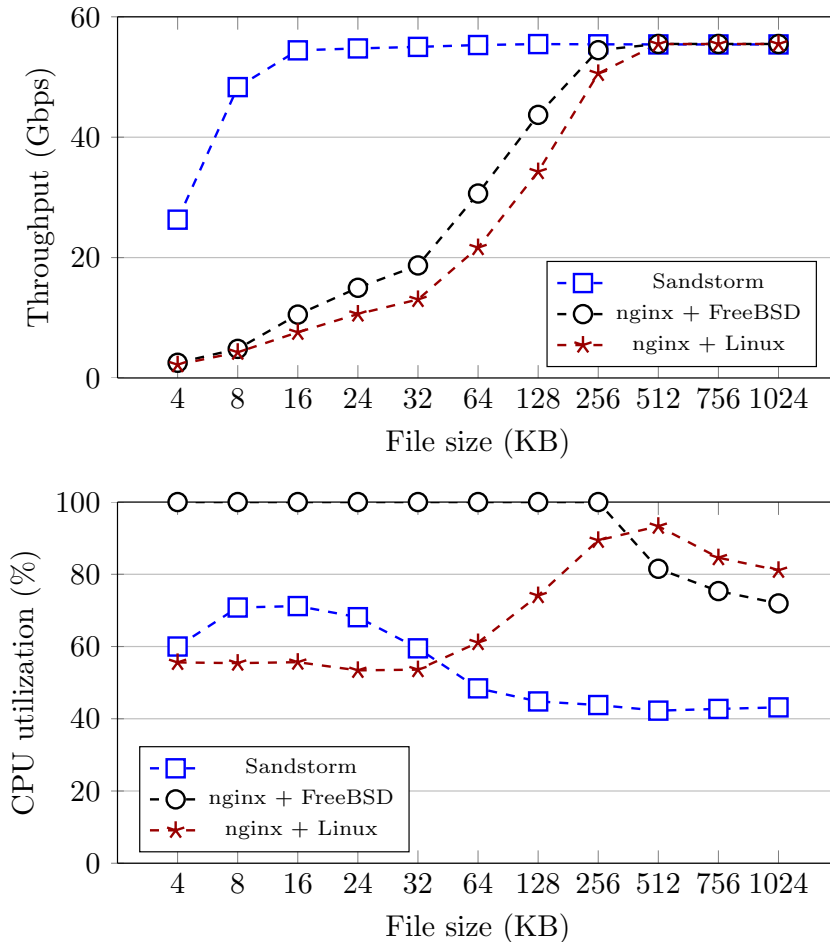


- Map NIC buffers directly into user process memory
- Not the sockets API: Zero copy to/from application
- System calls initiate DMA, block for NIC events
- Packets can be reinjected into normal stack
- Ships in FreeBSD; patch available for Linux
- Userspace network stack can be **specialised** to task (e.g., packet forwarding)



Network stack specialisation for performance

Ilias Marinos, Robert N. M. Watson, Mark Handley, SIGCOMM 2014, 2017.



- 30 years since the network-stack design developed
- Massive changes in architecture, micro-architecture, memory...
 - Optimising compilers
 - Cache-centered CPUs
 - Multiprocessing, NUMA
 - DMA, multiqueue
 - 10 Gigabit/s Ethernet
- Performance lost to 'generality' throughout stack
- Revisit fundamentals through clean-slate stack
- Orders-of-magnitude performance gains

