UC San Diego

CSE 291: Operating Systems in Datacenters

Amy Ousterhout

Nov. 9, 2023

UC San Diego

Agenda for Today

- Announcements
- SmartNICs overview
- AccelNet discussion
- iPipe discussion

Announcements

- Project check ins next week
  - I will post sign ups on Canvas
- Be prepared to talk about:
  - What have you been working on so far?
  - What are your next steps?
  - What obstacles or challenges are you currently facing?
- Feel free to share diagrams or graphs, or draw on the whiteboard
  - No need for a formal presentation
- All group members should participate

UC San Diego

SmartNICs

Incentives to Offload

- End of Moore's Law
  - Can we increase compute capacity with accelerators?
- Short tasks (e.g., 1-2 µs GET in a key-value store)
  - Can we reduce software overheads?
- Increasing network speeds
  - How can we support high bandwidth links?

SmartNICs!

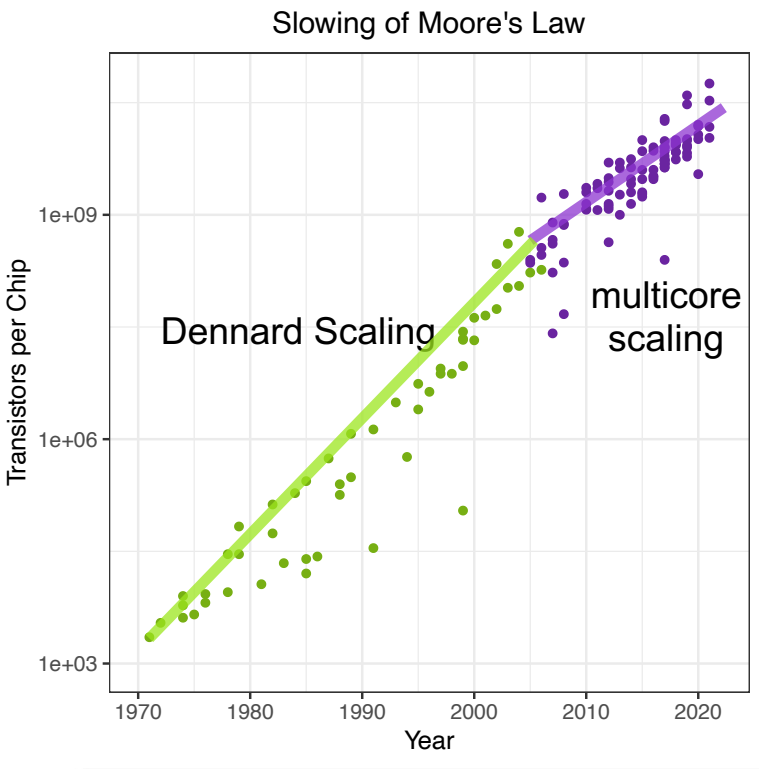

Different Varieties of SmartNICs

- ASIC-based SmartNICs
  - Limited flexibility (e.g., match-action tables)
  - Best performance
- FPGA-based SmartNICs
  - More flexibility but harder to program
- Multicore SoC-based SmartNICs → iPipe
  - Embed cores on the NIC → Intel's IPU
  - Most flexibility, worse performance
- Codesign the NIC and CPU

AccelNet →

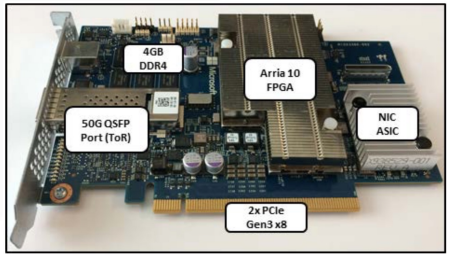

(b) Azure SmartNIC Gen2, 50GbE w/ on-board NIC

FPGA-based SmartNIC

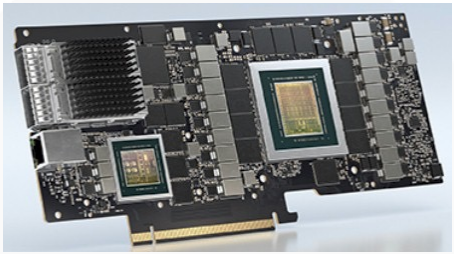

SoC-based SmartNIC

"Azure Accelerated Networking: SmartNICs in the Public Cloud", NSDI 2018
https://www.fierceelectronics.com/electronics/nvidia-claims-dpu-performance-record

Research Questions Raised by SmartNICs

- What tasks should we offload to SmartNICs?
  - OS-level functionality: checksums, virtualization, transport layer, scheduling
  - App-level functionality: serialization, encryption, compression, part of an app, complete app
  - What is the “killer app”?
- How dynamic should this be?
- How should the SmartNIC and CPU communicate?
  - DMA
  - RDMA
  - Directly connected

UC San Diego

AccelNet Discussion
iPipe Discussion