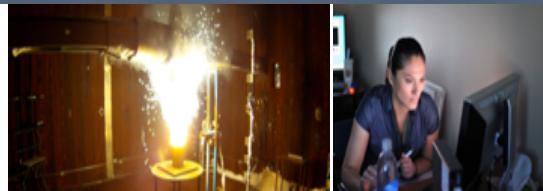




Co-design of System Software for Compute Accelerators and SmartNICs



PRESENTED BY

Ryan E. Grant
Center for Computing Research
Sandia National Laboratories
regrant@sandia.gov

Challenge

How can we effectively exploit the compute resources on SmartNICs and DPUs to accelerate scientific and engineering applications?

What is a SmartNIC/DPU?

- Can perform general compute
- May be a host or kernel interface
- Has local memory
- Capable of performing network operations independently

Opportunity

- Currently not clear where/how SmartNICs are best used
- Obvious use:
 - offload network processing and “system noise” tasks
- Application use:
 - Leverage for accelerating application performance
- Innovative use:
 - Entirely new types of distributed system software
 - E.g. self-learning and tuning for network performance and scheduling

Use determines Architecture

- “Blank silicon” due to EE concerns for signaling
- For obvious use:
 - Put general purpose cores on NIC – ARM cores either high performance or mainstream low power solutions
 - Build in accelerators for purpose-built network tasks
 - E.g Cryptography, compression, etc.
- For Application use:
 - Depends on highest demand from application space
 - Matrix multiply units
 - Data pack/unpacking
 - High throughput cores (GPU-like)



No Killer App - yet

- Successful performance improvement through noise reduction or application enhancement is a solid motivator
 - Most likely minimal additional cost to SmartNICs over existing high performance networks
- Innovative use could be clear killer app
 - AI-enabled networks could have broad impact
 - Use TPUs to run useful optimizations even without application awareness – global resource optimization
 - Enhance existing apps with strong AI acceleration
 - Support new workflows

Unique Architectural Impact

- Rare opportunity to co-design HPC SmartNICs
- Open silicon area that's easy to use
- Open questions:
 - What do we use this silicon area for?
 - Largest impact drives utilization
 - Re-configurable? Slower but flexible
 - Combination of fixed and reconfigurable? What proportion

Where do we go from here

- Need to explore potential co-design areas
- Co-design may not need to be with applications themselves
- Could co-design with systems software
- Explore accelerator options that are not common in current systems
- Think about entirely new system architecture
 - Does the node of the future look like a GPU with attached SmartNIC and no CPU?



Thank you

Questions?



Acknowledgments:

This work was funded through the Computational Systems and Software Environment sub-program of the Advanced Simulation and Computing Program funded by the National Nuclear Security Administration