Autonomous Infrastructure for Science Panel 2018

Peter Dinda
Prescience Lab
Department of EECS
Northwestern University

pdinda.org
presciencelab.org

My Possibly Relevant Background

- DARPA Quorum Project (as grad student)
  - Adaptive distributed computing on COTS warship
- Relational Grid Information Services (RGIS) Project
- Virtuoso Project
  - Adaptive virtual infrastructure ("IaaS cloud")
  - Inference+prediction of parallel application resource demand + distributed system (grid) resource supply
  - Matching of one to other using virtualization-based mapping, scheduling, and resource reservation
- Empathic Systems Project
  - Optimize for measured user satisfaction with performance (human-in-the-loop)
  - Biometric sensor hardware
- V3VEE Project
  - Infer/predict/match/implement within a VMM driving a NUMA node
- ABSYNTH Project
  - Sensor network programming languages and systems for scientists and engineers, not embedded systems experts

General Thoughts / Opinions

- Adaptive computing is pretty old
  - Late 1990s at least
  - 1980s if we consider load sharing / load balancing
- It has lots of names and lots of domains
  - Distributed systems, parallel systems, embedded systems, networking, performance analysis, autonomic computing, …
- It has lots of approaches
  - Queuing theory, predictive signal processing, control theory, …
- AI-Science should focus on what is different...
- … and consider some possible lessons-learned
  - Systems-level very different from application-level
  - Problems should be abstract and generalizable
    - And their implementations too! [See virtualization]
  - Problems should be in P
  - Mechanism is fun, Policy is hard

In terms of adding autonomous capabilities, how is science infrastructure different from commercial infrastructure?

- Virtual infrastructure?
- Sensor/actuator networks?
- Collective operation
  - Parallelism, data parallel pipelines, DAGs…
- Correctness
  - What can the system (software and hardware) do to make programs “more correct”?
- Possibly scientific constraints map to systems-level constraints that can be tractably implemented
  - Most adaptation/optimization problems in systems often go NP-Hard very fast, and you’re left arguing heuristics or having some ML black box you don’t understand

Provide 3 examples of autonomous capabilities that you want to see in science infrastructure.

- Safe deployment schemes
  - Trustless computing (in both directions)
- Global time and Real-time
  - Thinking in terms of a distributed system is not natural for many scientific domains
- Correctness enhancement
  - How can we improve existing codes?

- Marketplace of Independent Consultants
  - How to find capable people: numerical methods, software development, parallelization, performance analysis…

For More Information

- Peter Dinda
  - pdinda@northwestern.edu
  - http://pdinda.org
- Prescience Lab
  - http://presciencelab.org
- Acknowledgements
  - NSF, DOE