

XStack PI Meeting Resilience Panel 2014

Moderators: Andrew Chien, Mattan Erez

Panelists: Michael Carbin, Kath Knobe,
Vivek Sarkar, Martin Schulz,
and Sriram Krishnamoorthy

Resilience Panel Report out (1/3)

- Hardware Estimates of error rates have wide uncertainties (1-100x), with projected rates as high as 10^{-2} uncorrected-errors/second
- Unaddressed, such error rates would render exascale machines unusable for DOE-scale applications
- Mainstream computing is exploring the opportunity in “high error rate” operating points (approximate computing); this may spill over into HPC
- => We must create a set of composable resilience approaches to ensure that DOE-scale computations can be accomplished (correctly and efficiently)

Resilience Panel Report out (2/3)

- Applications need to specify critical resilience requirements (data and computation)
- Resilience approaches should not require major new efforts by application scientists (gentle slope)
- Resilience approaches should have controllable cost, and applications/systems should be able to manage this cost to match to their requirements
- Fail-stop errors must be converted to signaled errors with opportunity for recovery (x-layer)
- Architectural approaches (across layers of software and hardware) are needed to ensure that the resilience challenge in the form it eventually arises can be successfully addressed

Resilience Panel Report out (3/3)

We discussed a number of promising approaches

Hierarchical and task-based models

Idempotence

Data-oriented Resilience

Programming model semantics

Application-semantics based resilience

Tools to manage resilience

but there are many more...