Lazy Shadowing:
power-aware resilience for extreme scale systems

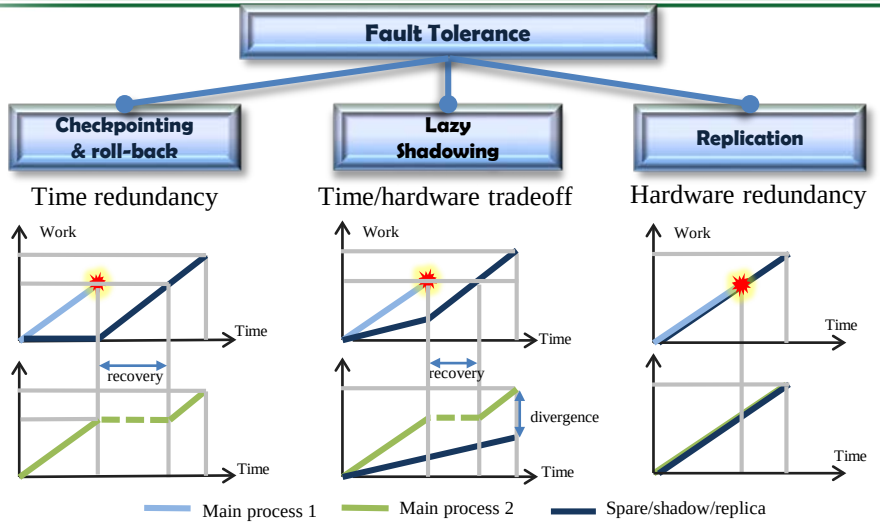

Novel Ideas

- A form of replication with slow replicas (low overhead)
- A form of checkpointing with roll forward (efficient recovery)
- Partial replicas through program slicing (faster error detection)
- Optimize a combined function of resilience, power and time
- Shadows/replicas can be slowed down by a factor of K by
  - Reducing voltage/frequency by a factor of K (DVFS)
  - Collocating K shadows on the same node
- Shadow leaping reduces the recovery time
- Rejuvenation (shadow as a rescuer rather than a hot spare)

Impact

- A new concept to add resilience at extreme scale with affordable overhead in performance and power.
- MPI implementation allows transparent recovery from faults

  1) Brings fault-tolerance to main stream high performance computing.

  2) Provides resilience at extreme scale with a low overhead relative to replication and checkpointing.
- Combination with program slicing for faster fault detection and better resilience at lower energy overhead

PI: Rami Melhem (University of Pittsburgh)
CoPI: Taieb Znati (University of Pittsburgh)
CoPI: Krishna Kant (Temple University)

Milestones/Dates/Status

Analytical modeling	Dec. 2015
Preliminary evaluation of shadowing and slicing	June 2016
Apply shadowing to map/reduce	Dec. 2016
Implement lazy shadowing in MPI	Dec. 2016
Implement/analyze fault-induced leaping	June 2017
Evaluate shadowing Vs. checkpointing	June 2017
Integration of slicing and shadowing	Dec. 2017
Implement/analyze rejuvenation	Dec. 2017
Implement/analyze forced leaping	June 2018
Finding the optimal forced leaping interval	June 2018

U.S. DEPARTMENT OF ENERGY | Office of Science

March 13, 2017