

---

# Partial Redundancy and In-memory checkpoints Under Heterogeneous Failure Likelihoods

Rami Melhem\*<sup>1</sup>

<sup>1</sup>The University of Pittsburgh – United States

## Abstract

Partial redundancy and in-memory checkpoint placement are studied for HPC systems where individual node failure distributions are not identical. First, we show that partial redundancy may provide the best performance when nodes have different reliabilities and we derive the replicas-to-nodes assignment that maximizes system reliability. Next, we explore the problem of placing in-memory checkpoints among nodes with different reliabilities and provide results on the optimal way to place in-memory checkpoints such that the probability of occurrence of a catastrophic failure is minimized. We support the theoretical results by numerical experiments as well as simulations that use failure logs spread over 5 years of a 49,152 node supercomputer.

**Keywords:** HPC, reliability, checkpoints, redundancy

---

\*Speaker