

Cooperative Checkpointing: A Robust Approach to Large- Scale Systems Reliability



Adam J. Oliner, Stanford University



Larry Rudolph, MIT, CSAIL



Ramendra K. Sahoo, IBM, T.J. Watson

June 28th

ICS 2006, Cairns, Australia



The Goal

- We want checkpointing that is practical, efficient, and robust against varied and varying system conditions and failure distributions




Motivation

- Two trends with increasing scale:
 - Increasing failure rate
 - Increasing checkpoint overhead



Motivation





Cooperative Checkpointing

Dynamically skip checkpoint requests made by applications at runtime:

1. User inserts checkpoint requests
2. Compiler optimizes
3. Runtime gatekeeper grants/denies



Key Idea

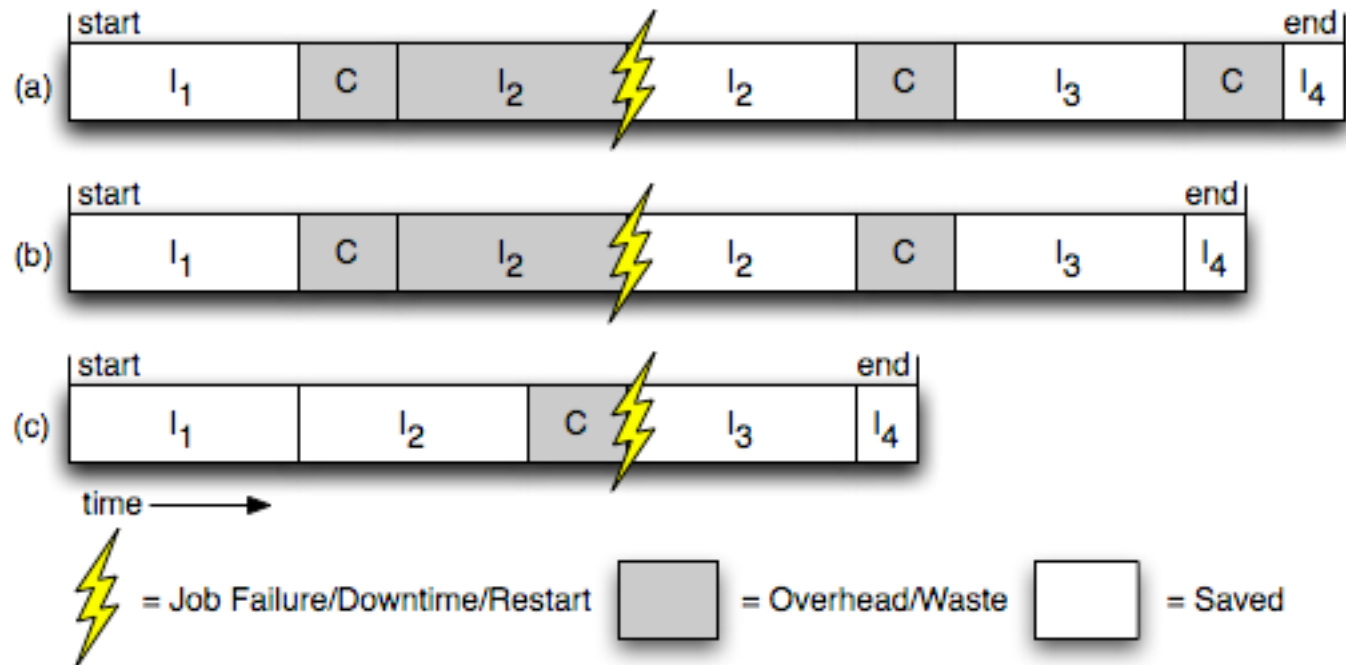
Flexible checkpoint times
+ Good heuristics



= Robustness



CC Example



- Intuition: skip checkpoints less likely to be used for rollback




Experiments: Overview

- Two sets of simulations
- Varied parameters include:
 - Job log, failure trace, prediction accuracy, interconnect topology, checkpoint request interval, checkpoint overhead, cooperative checkpointing heuristic...
- 30 million jobs, 600K machine hours, hundreds of millions of failures



Experiments: Failure Traces

- Real:
 - AIX cluster [Sahoo, KDD 2003]
 - BG/L prototype [Sahoo, DSN 2005]
- Artificial:
 - Exponential
 - Weibull Bathtub
 - Uniform
 - Split Uniform



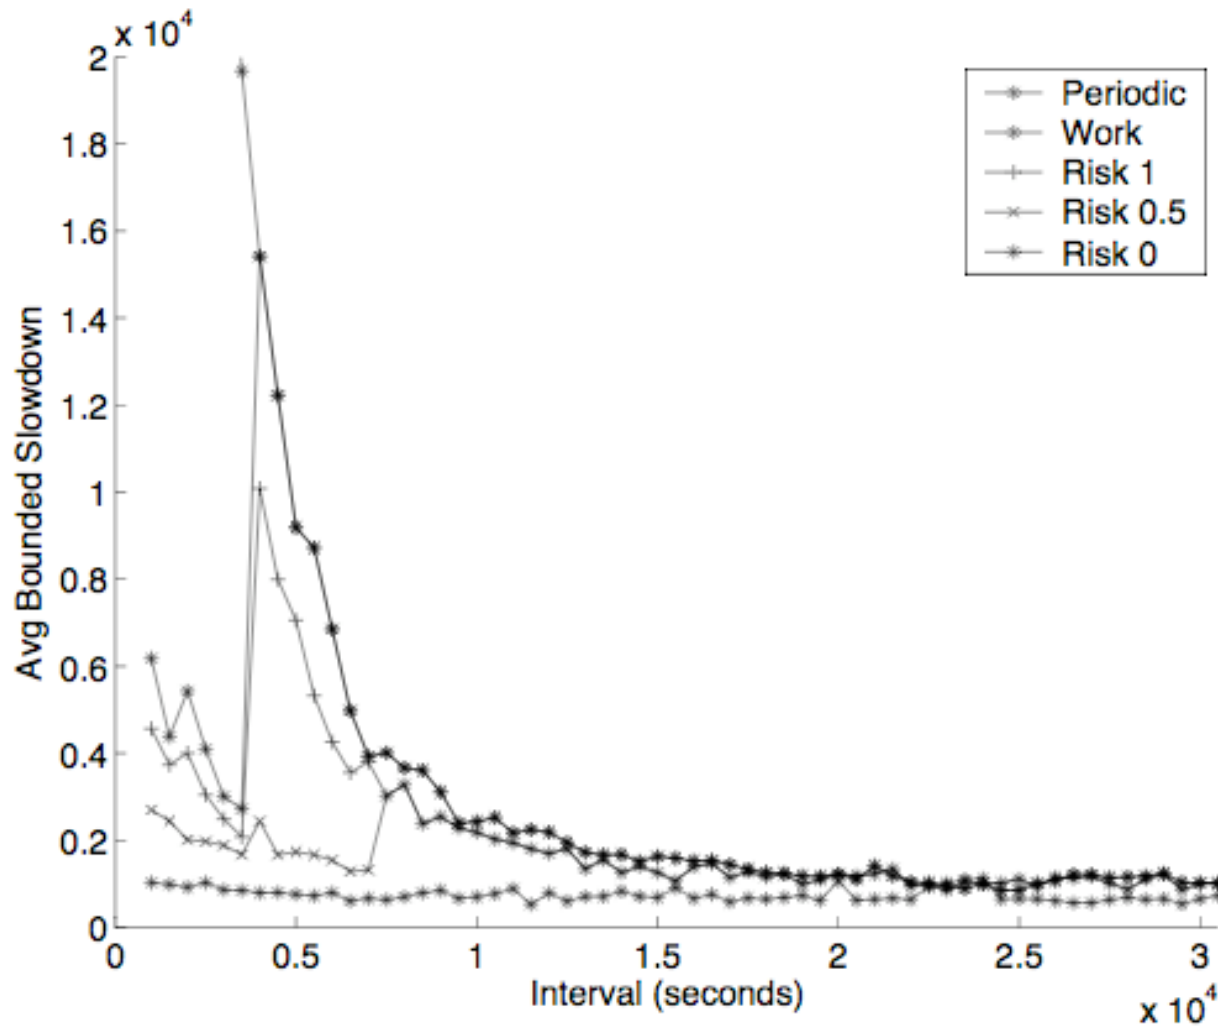


Experiments: Heuristics

- OPT
- Periodic
- Revised Periodic
- Backoff
- Risk-based

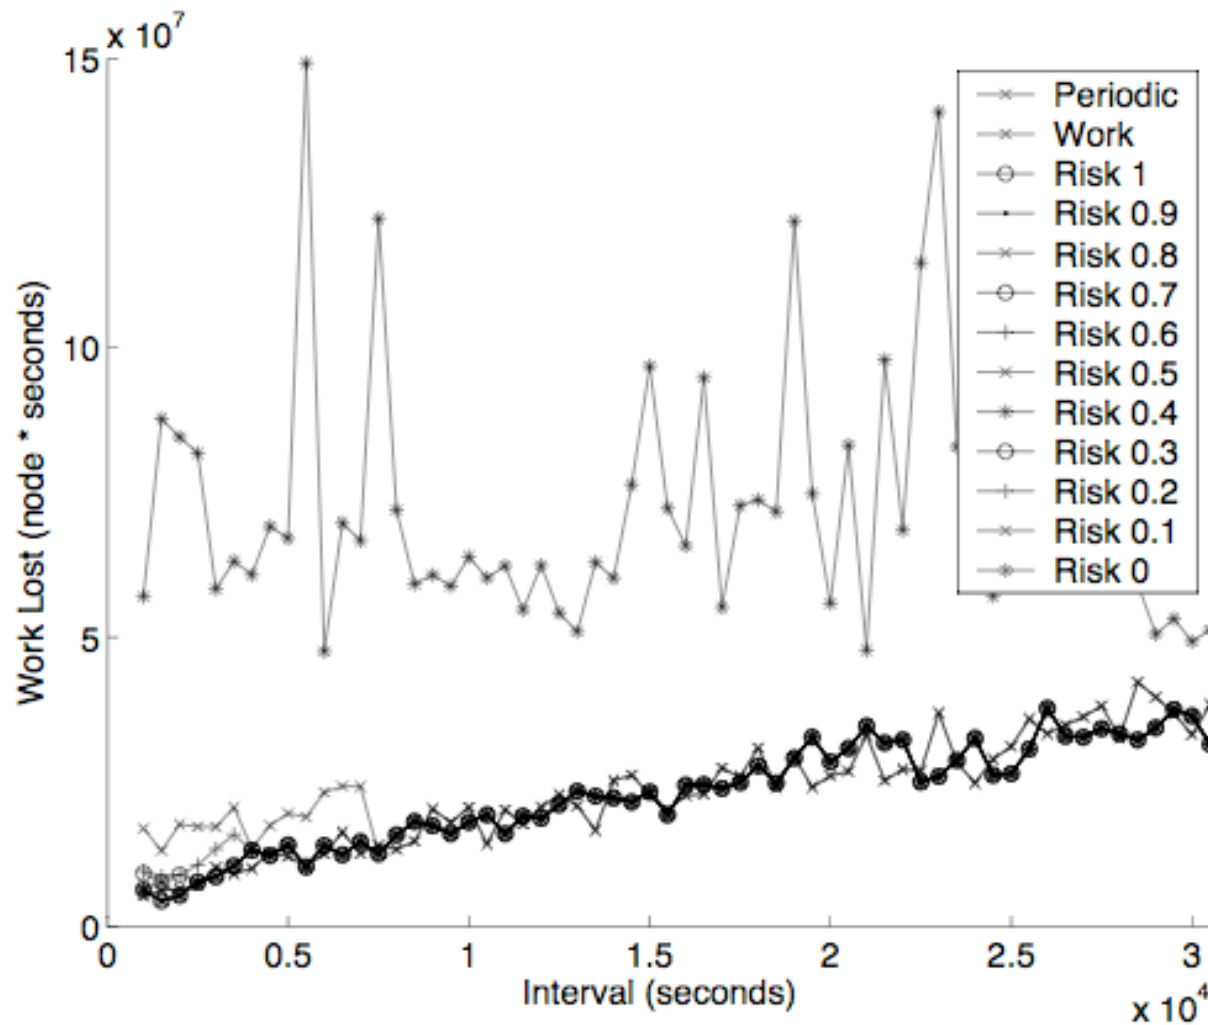


Results: Controlling Scale





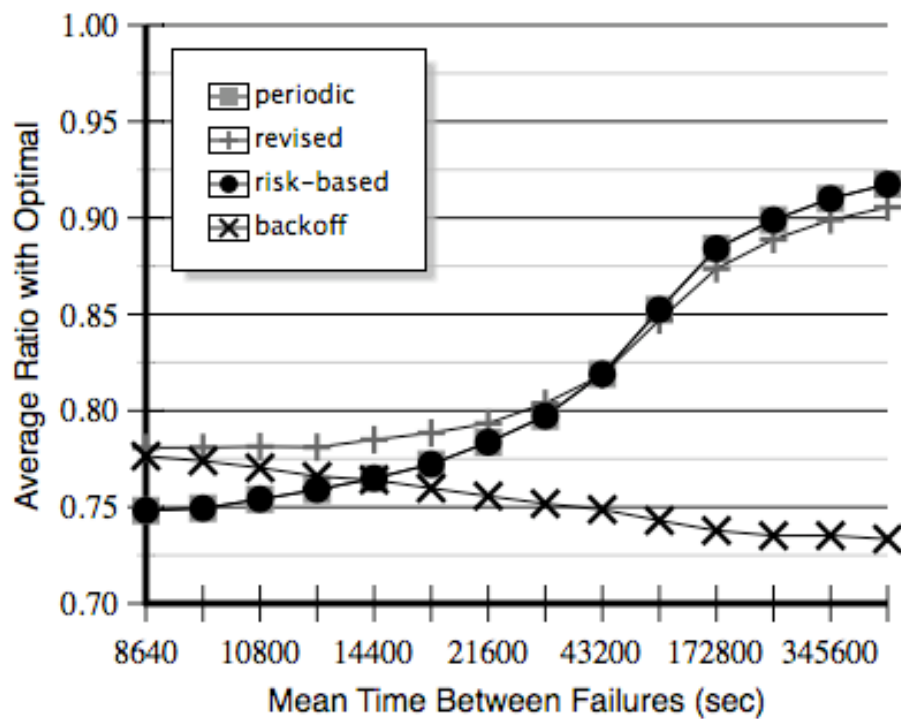
Results: Prediction Accuracy






Results: Ratio with OPT

Exponential Distribution




Split Uniform Distribution



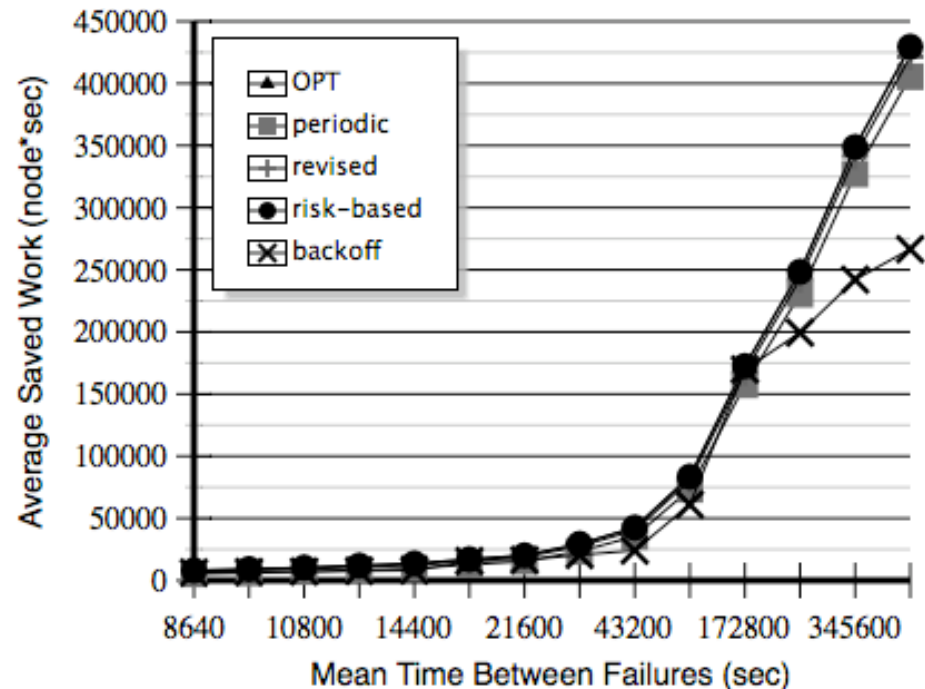


Results: Measuring Progress

Exponential Distribution



Split Uniform Distribution





Surprises

- Even low accuracy prediction can be a powerful resource
- Common metrics may be contradictory
- Skipping checkpoints can improve performance and reliability



Contributions

- Motivates the need for a new approach to checkpointing
- Presents cooperative checkpointing as a practical, efficient, and robust solution
- Demonstrates robustness under a variety of reliability conditions