

# EXPLORATORY TESTING AT SCALE


Jiri Simsa, Randy Bryant, Garth Gibson (Carnegie Mellon University), Jason Hickey (Google)

## PROBLEM

- Concurrency manifests test non-determinism
- Systematic testing controls the order of concurrent events and systematically enumerates different concurrent scenarios
- To mitigate combinatorial explosion, state space reduction techniques avoid exploration of equivalent scenarios
- For large-scale distributed programs, parallelize testing: efficient & load-balanced exploration

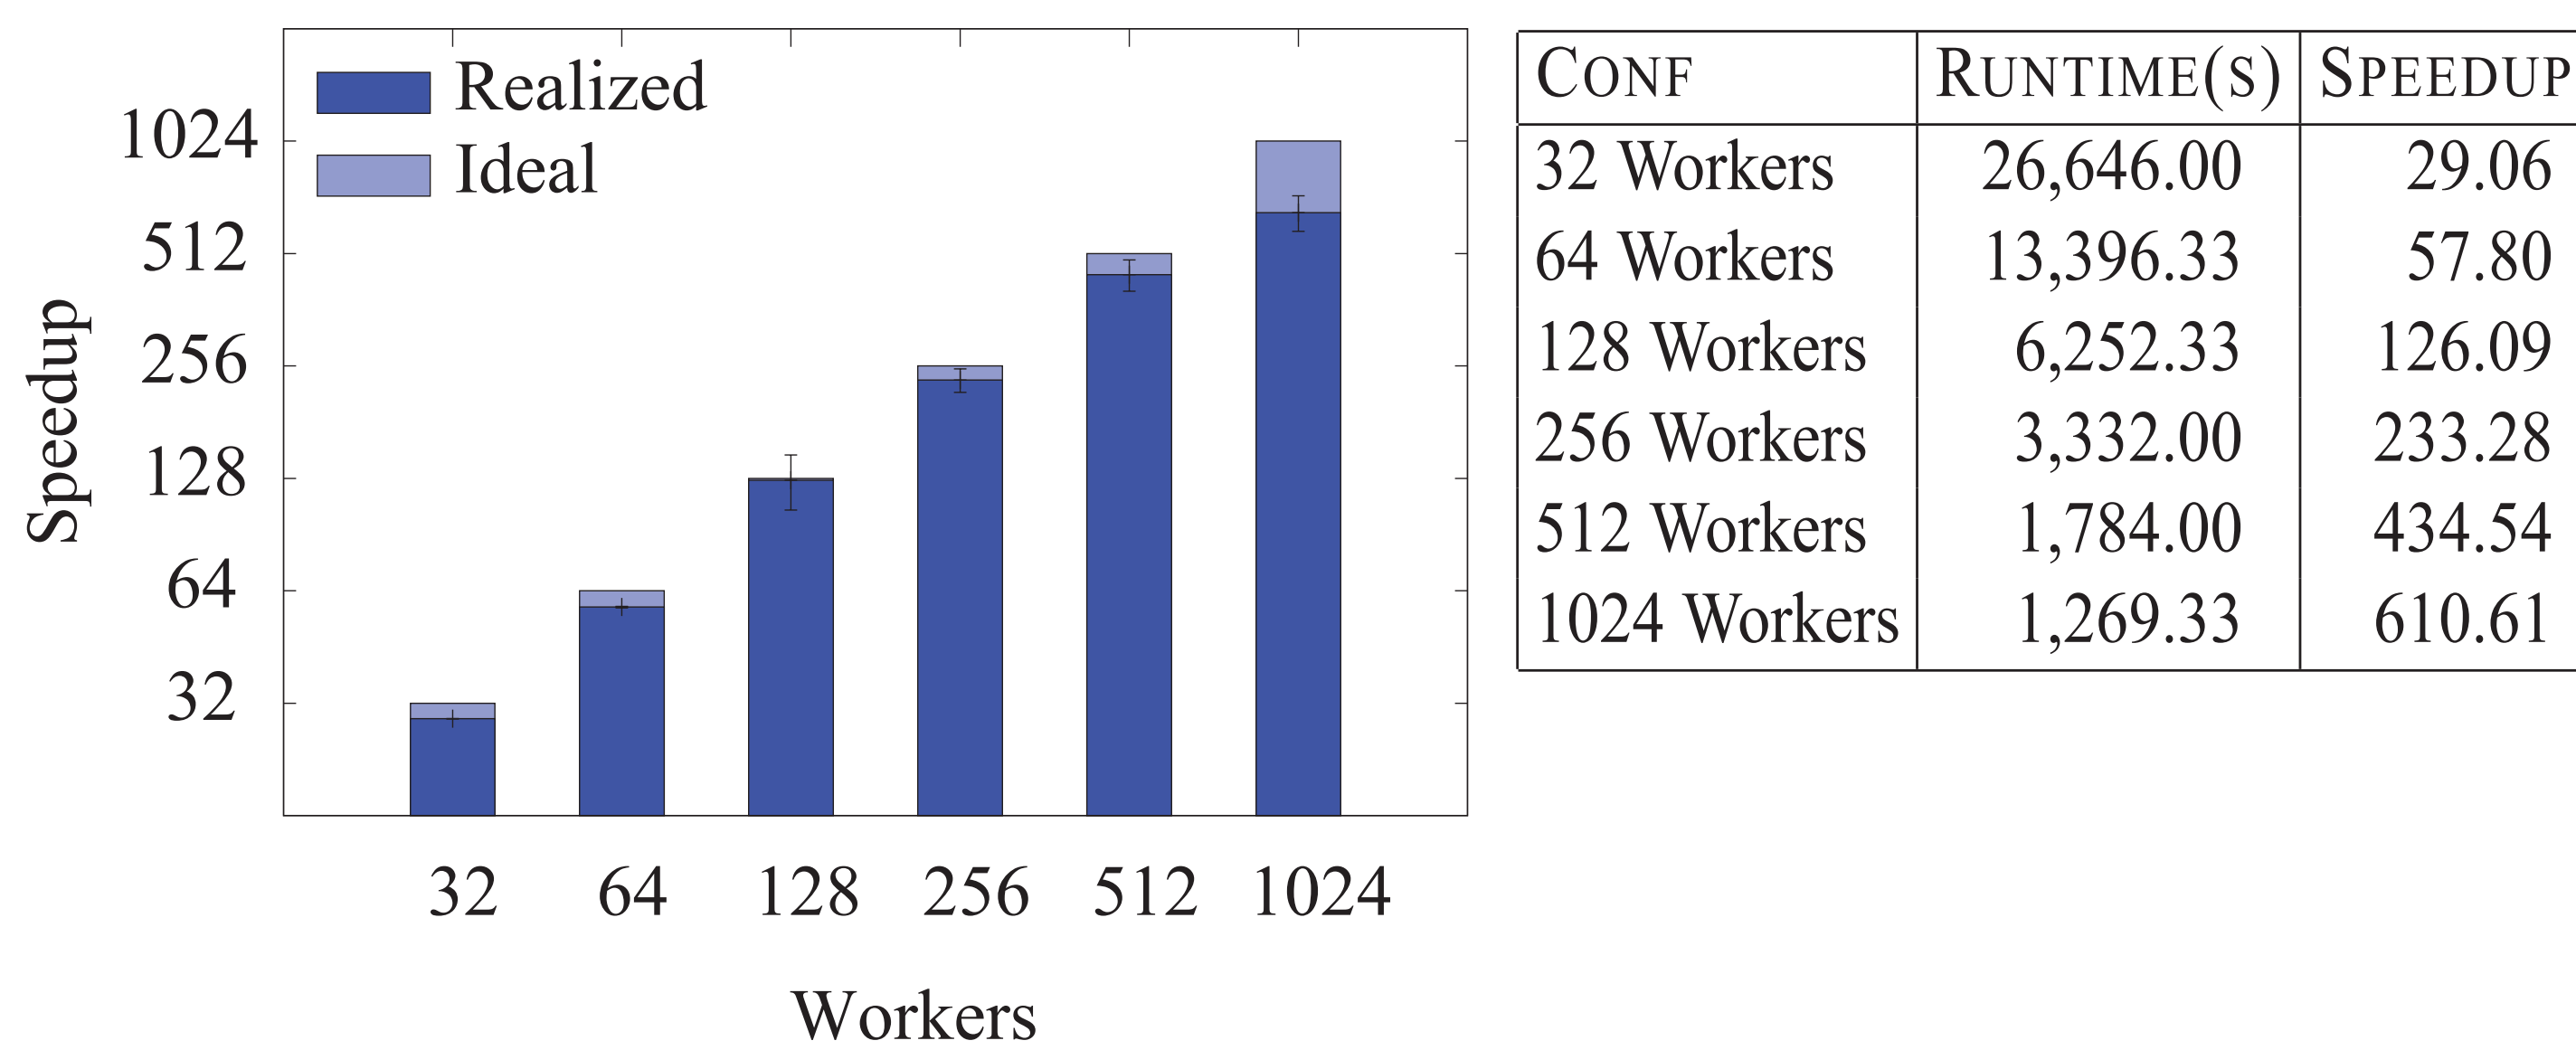
## SCALABILITY

Resource(6,6)\* with concurrent DPOR




For this example, dynamic partial order reduction explores on the order of 18.5 million branches and the sequential implementation is expected to require 209 hours to finish.

Store(12,3,3)\* with concurrent DPOR



For this example, dynamic partial order reduction explores on the order of 21 million branches and the sequential implementation is expected to require 215 hours to finish.

Scheduling(10)\* with concurrent DPOR




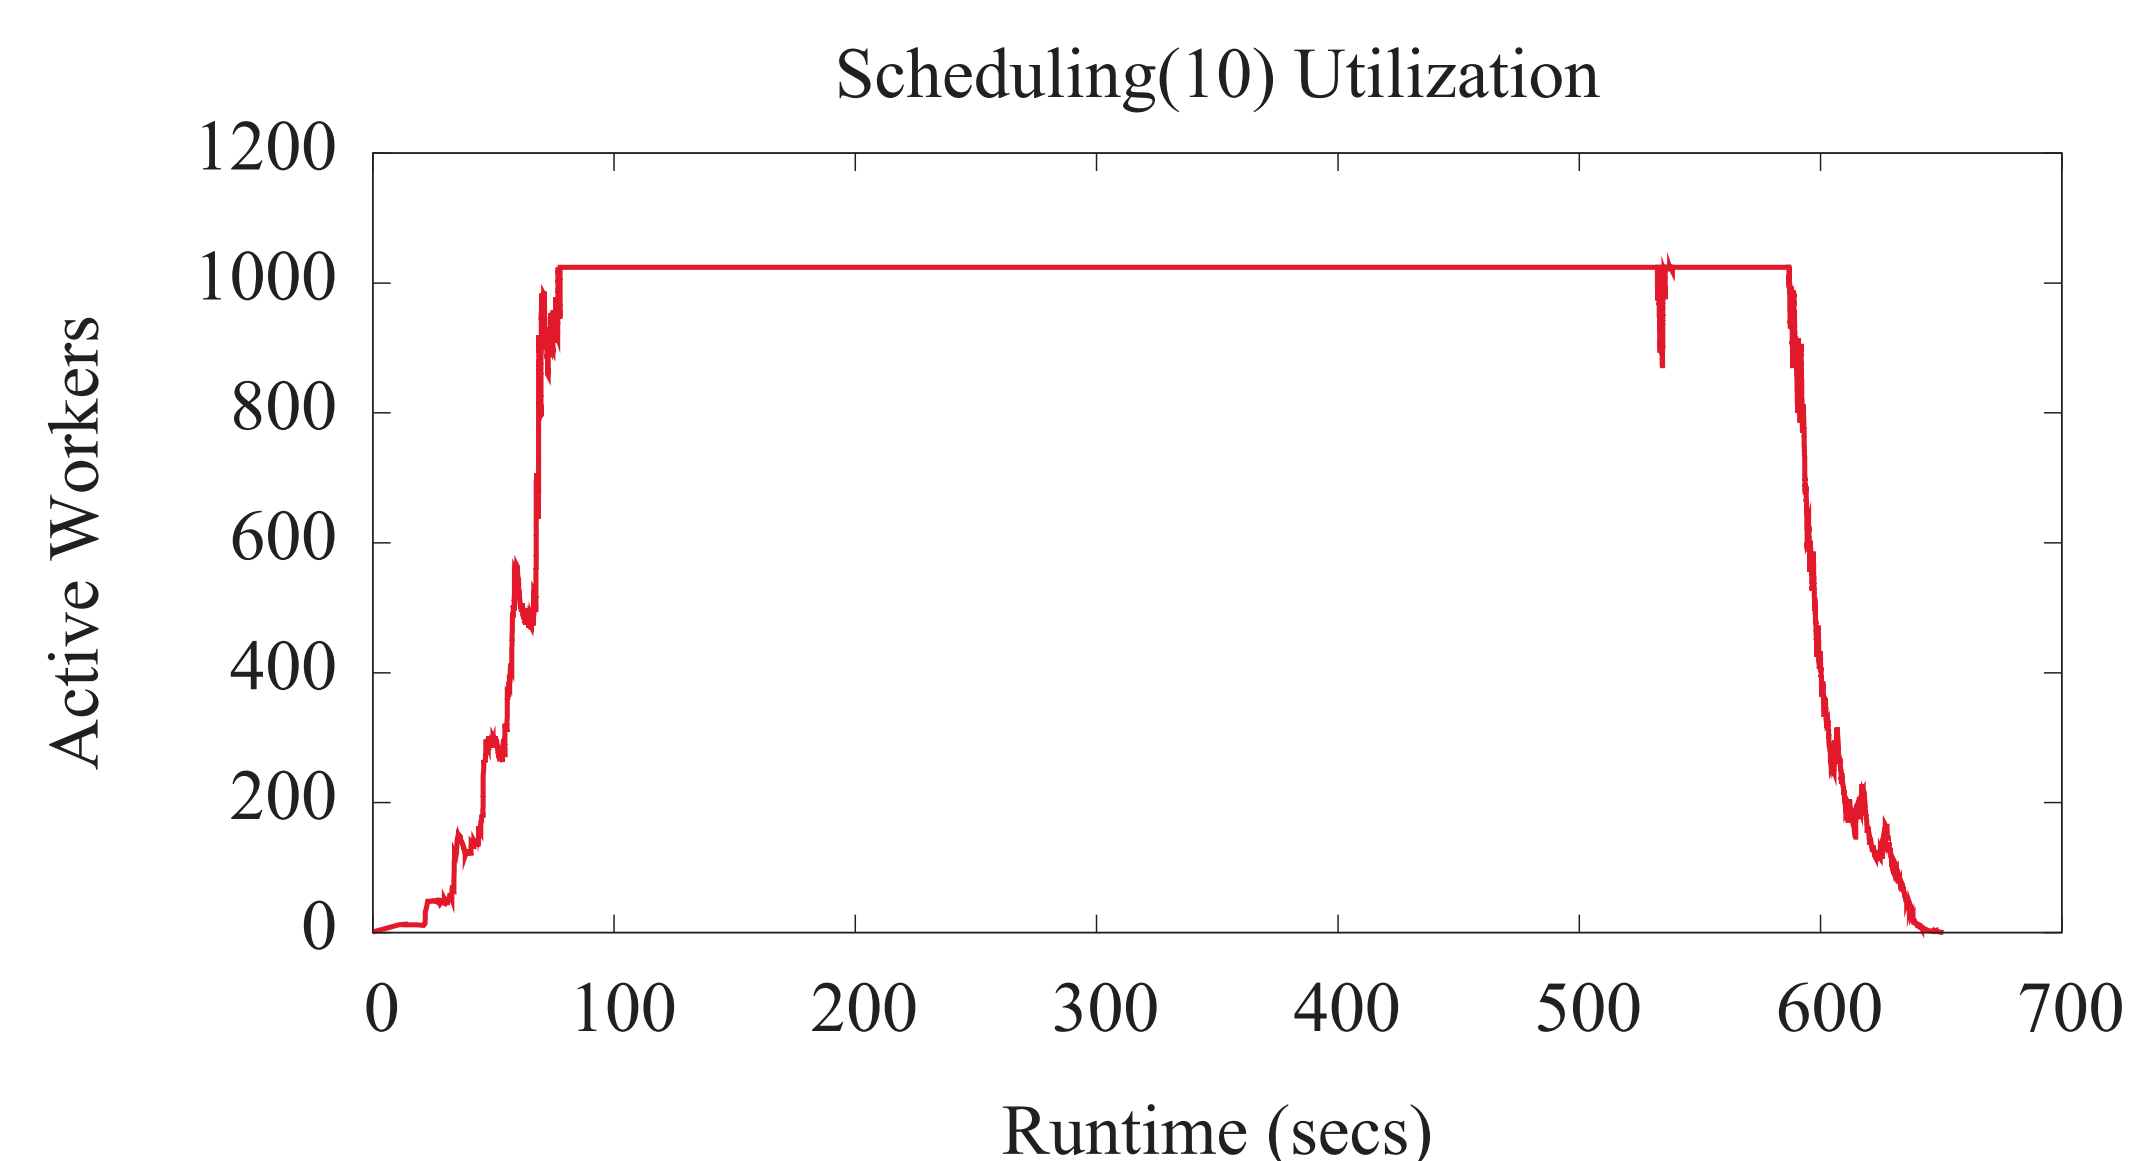
For this example, dynamic partial order reduction explores on the order of 3.6 million branches and the sequential implementation is expected to require 126 hours to finish.

## SOLUTION

- Distributed dynamic partial order reduction (DPOR) based on a novel state space exploration algorithm: n-partitioned depth-first search
- Dynamic deadline for fault tolerance; trade space complexity for parallelism in testing
- Implementation demonstrated to scale up to cluster of 1,024 nodes (speed up upwards of 750x)

## IMPROVING UTILIZATION

- Initial and final phase of exploration underutilizes worker fleet at large scale
- Two techniques to improve utilization: variable time budget, redundant exploration



- Improved speed up from 538x to 916x, from 610x to 759x, and from 696x to 865x for the respective tests

## THEORETICAL LIMITS

- The centralized master is a potential bottleneck
- At 1024 workers, memory overhead of the master 4MB → possible to scale to hundreds of thousands of workers
- At 1024 workers and 10-second worker time budget, master processes roughly 100 RPCs per second at 20% CPU utilization
- To scale beyond thousands of workers → scale worker time budget, better hardware, optimize software stack



\*instances of actor program tests from the test suite of Google's next-generation cluster management system

