

# Effects of Low-Quality Computation Time Estimates in Policed Schedulers



Justin M. Wozniak, Yingxin Jiang, and Aaron Striegel

March 28, 2007



# Outline



- x Outline
- x Motivation
- x Observations
- x Estimation
- x Simulation
- x Policy
- x Results
- x Conclusion

# Motivation



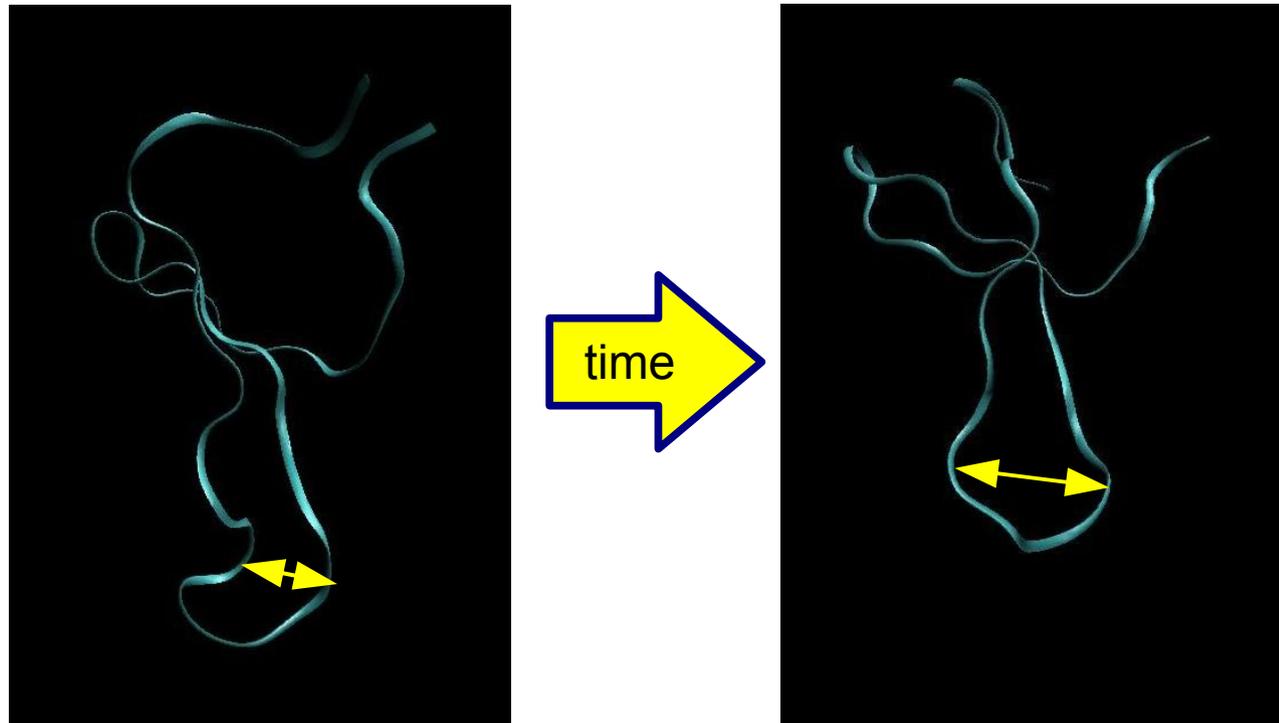
xOutline  
xMotivation  
xObservations  
xEstimation  
xSimulation  
xPolicy  
xResults  
xConclusion

- Ordinary scientific applications may involve computation time constraints
  - Performance/utilization constraints
  - User constraints
  - Economic constraints
- Numerical scientific codes may be predictable
  - Invariant I/O requirements
  - Well-understood underlying algorithms

# Motivation

- Example: Transition Path Sampling (Chemistry)

xOutline  
xMotivation  
xObservations  
xEstimation  
xSimulation  
xPolicy  
xResults  
xConclusion



# Motivation

- xOutline
- xMotivation
- xObservations
- xEstimation
- xSimulation
- xPolicy
- xResults
- xConclusion

- Biomolecular simulations may be managed as deadline-driven jobs to enhance system utilization and user experience
- Whole (Monte Carlo) batches must complete to obtain a result



# Motivation

- Real-world job unpredictability:

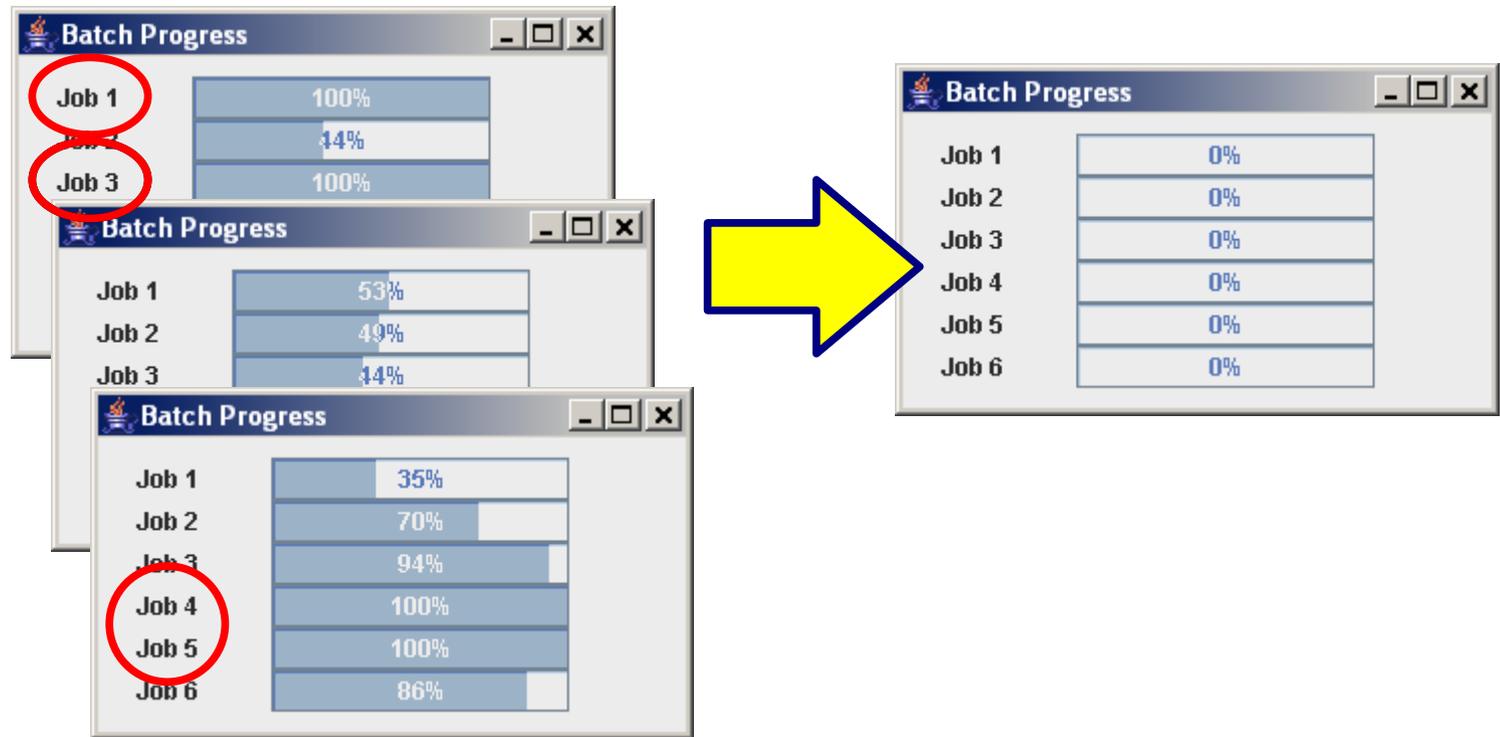
- xOutline
- xMotivation
- xObservations
- xEstimation
- xSimulation
- xPolicy
- xResults
- xConclusion

job	progress	last heard from	
17f	xxxxxxx	03/19/2007	08:41PM
17b	xxxxxxxxxxxxxxxxxxxxxxxxxxxx	03/18/2007	12:33PM
18f	xxxxxxxxxxxxxxxxxxxxxxxxxxxx	03/19/2007	10:29PM
18b			
19f	xxxxxxxxxxxxxxxxxxxxxxxxxxxx	03/19/2007	09:43PM
19b	xxxxxxxxxxxxxxxxxxxxxxxxxxxx	03/19/2007	08:10PM
20f	xxxxxxxxxxxxxxxxxxxxxxxxxxxx	03/19/2007	06:56PM
20b	xxxxxxxxxxxxxxxxxxxxxxxxxxxx	03/19/2007	09:47PM
21f	xxxxxxxxxxxxxxxxxxxxxxxxxxxx	03/19/2007	10:48PM
21b	xxxxxxxxxxxxxxxxxxxxxxxxxxxx	03/19/2007	10:48PM
22f	xxxxxxxxxxxxxxxxxxxxxxxxxxxx	03/19/2007	09:19PM
22b	xxxxxxxxxxxxxxxxxxxxxxxxxxxx	03/19/2007	10:17PM
23f	xxxxxxxxxxxxxxxxxxxxxxxx	03/19/2007	06:16AM
23b	xxxxxxxxxxxxxxxxxxxxxxxxxxxx	03/19/2007	08:34PM
24f	xxxxxxxxxxxxxxxxxxxx	03/19/2007	06:33PM
24b	xxxxxxxxxxxxxxxxxxxxxxxxxxxx	03/18/2007	11:41AM

# Motivation

- xOutline
- xMotivation
- xObservations
- xEstimation
- xSimulation
- xPolicy
- xResults
- xConclusion

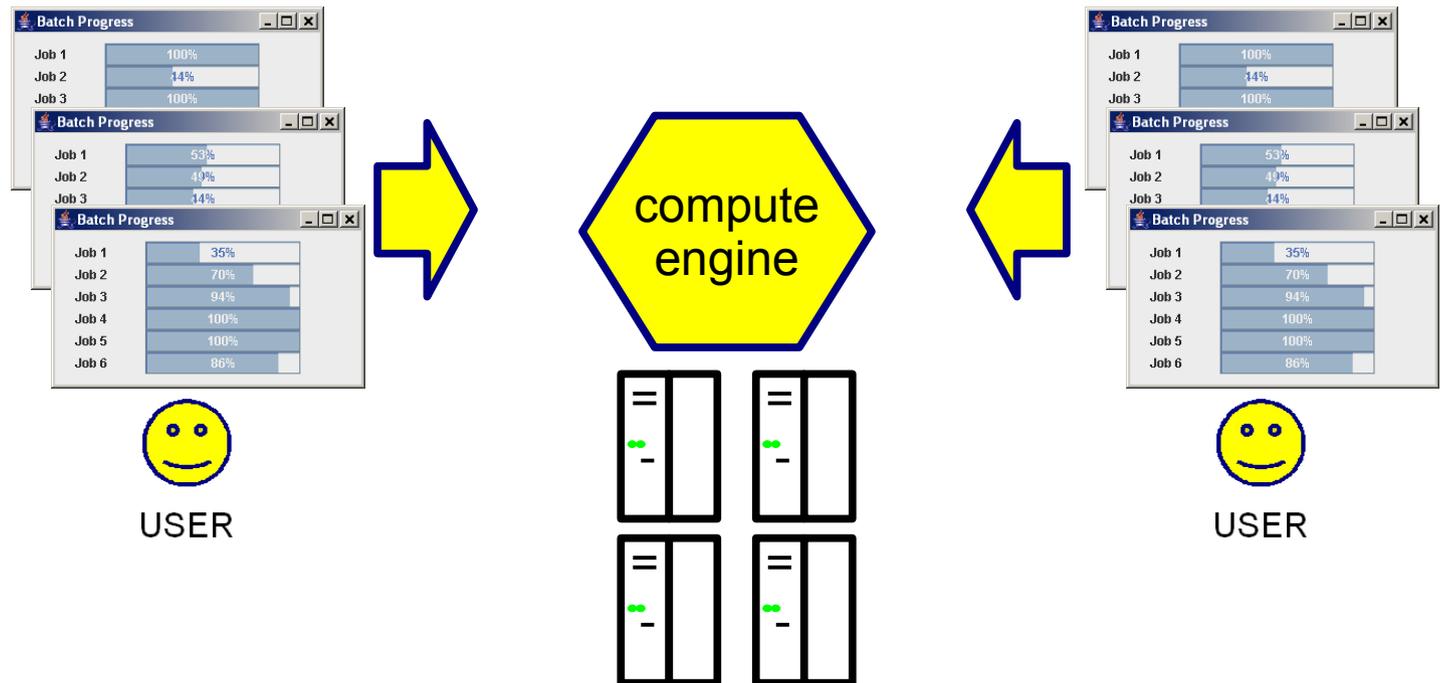
- Phrasing biomolecular simulations as deadline-driven jobs



# Definitions

- Contention

- xOutline
- xMotivation
- xObservations
- xEstimation
- xSimulation
- xPolicy
- xResults
- xConclusion



# Definitions

---

---

- xOutline
- xMotivation
- xObservations
- xEstimation
- xSimulation
- xPolicy
- xResults
- xConclusion

- Admission control:  
“Given the available resources, can the system guarantee a result by a given deadline?”
- Estimated computation cost must be available to request admission.
- Enforcement:  
“Given the estimates and current state of the system, what must be done to satisfy outstanding guarantees?”

# Definitions

---

---

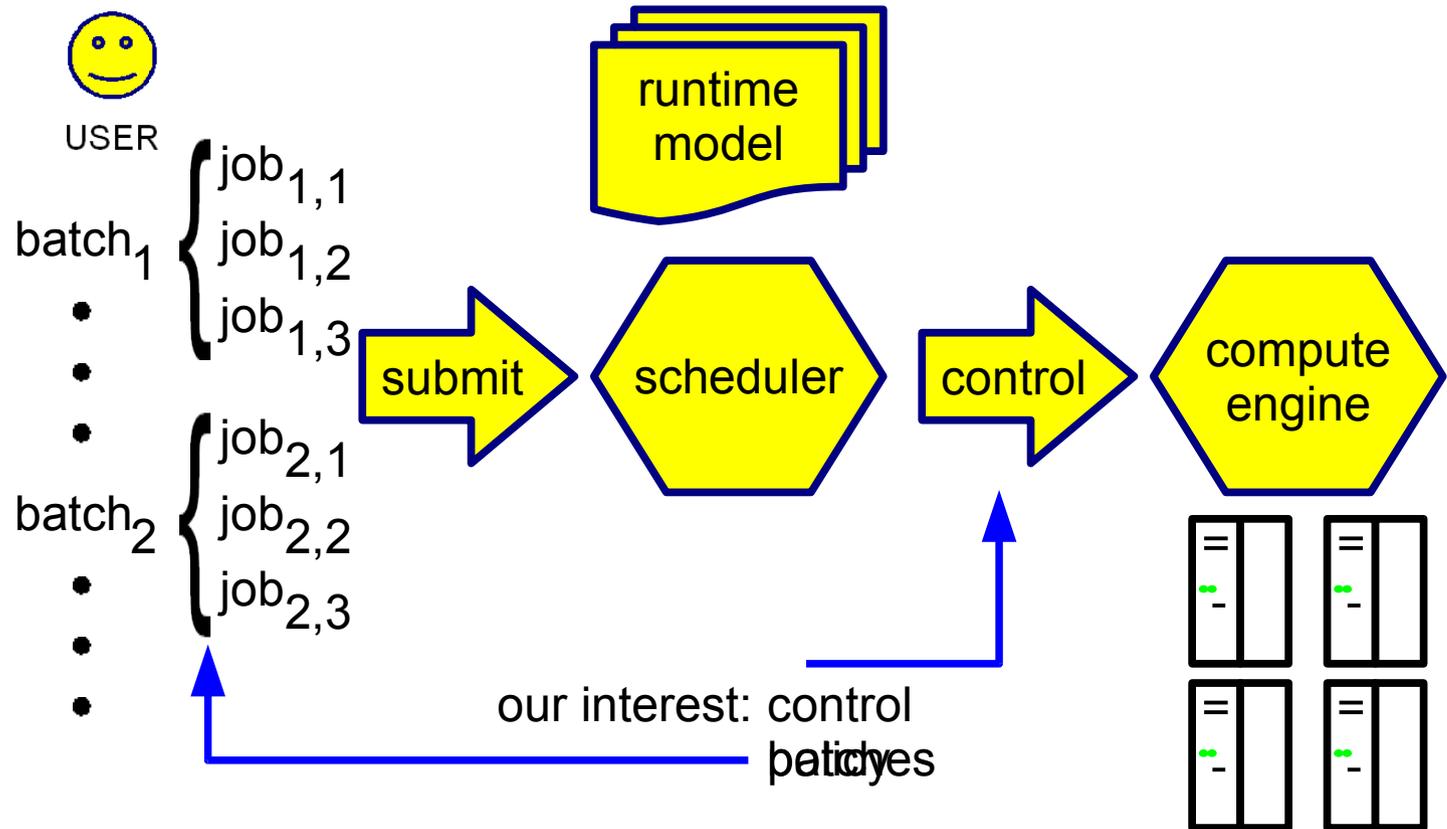
- xOutline
- xMotivation
- xObservations
- xEstimation
- xSimulation
- xPolicy
- xResults
- xConclusion

- Accept ratio:  
“What fraction of my batch requests will be admitted to the system?”
- Guarantee ratio:  
“What fraction of the admitted batches will complete on time?”

# Simulation

## Fast simulated framework

- xOutline
- xMotivation
- xObservations
- xEstimation
- xSimulation
- xPolicy
- xResults
- xConclusion



# Estimation

---

---

xOutline  
xMotivation  
xObservations  
xEstimation  
xSimulation  
xPolicy  
xResults  
xConclusion

- Inherent Unpredictability
  - Grid resources are volatile in availability and variable in performance
  - User estimates are pretty bad
- The Fundamental Enforcement Policy Question:

When should the system kill jobs  
to satisfy guarantees?

# Estimation

---

---

- xOutline
- xMotivation
- xObservations
- xEstimation
- xSimulation
- xPolicy
- xResults
- xConclusion

- Quality of Estimate:  
“QoE = 50% if a job is guaranteed to never exceed its estimate by 50%.”
- If QoE = 0% and the compute system is totally predictable, we can make perfect guarantees. *This is not the case, and some batches will not complete on time.*

# Estimation

---

---

- xOutline
- xMotivation
- xObservations
- xEstimation
- xSimulation
- xPolicy
- xResults
- xConclusion

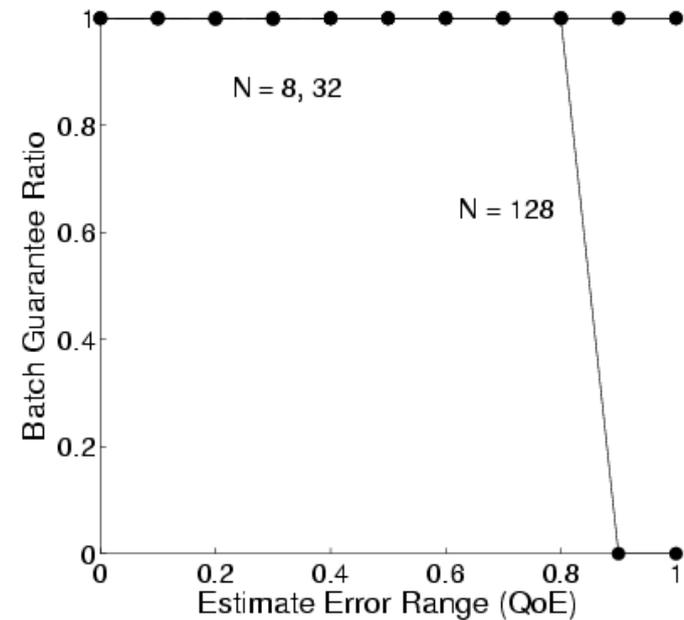
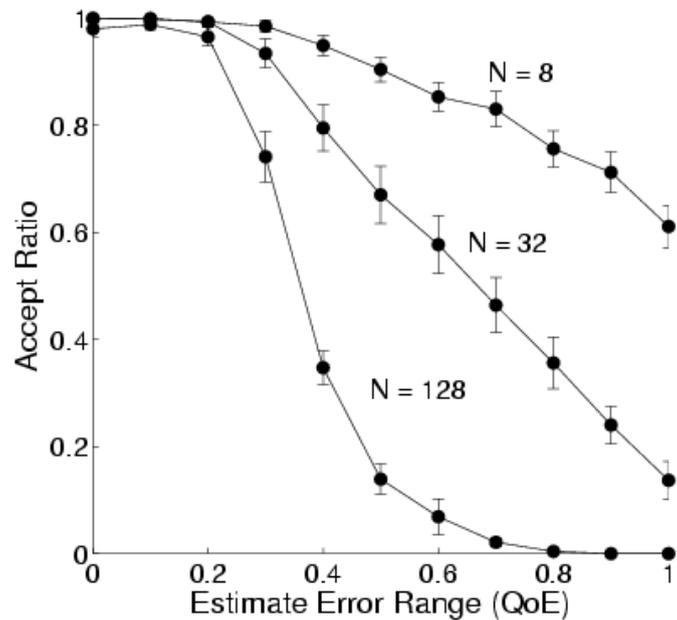
## The Game:

- Users will try to give low estimates to gain admission to the system.
- Underestimates will cause deadline misses. Worse,
- Malicious users will intentionally underestimate to obtain admission.
- We can enforce estimates by killing jobs that have exceeded their estimate.
- This will have unfortunate results.

# Simulation Results

## Non-enforcement of estimates

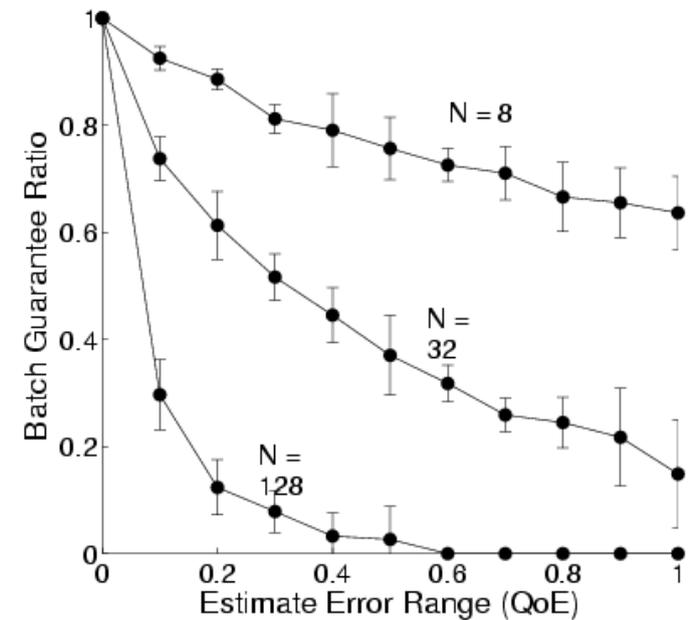
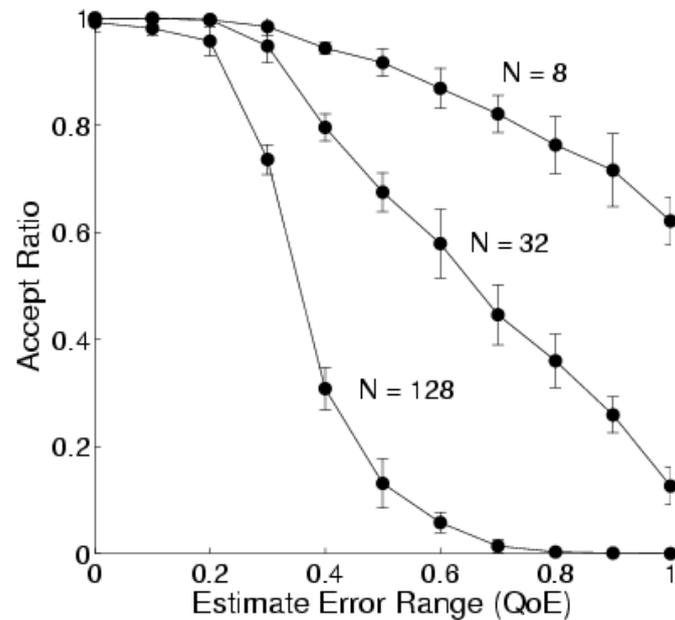
- × Outline
- × Motivation
- × Observations
- × Estimation
- × Simulation
- × Policy
- × Results
- × Conclusion



# Simulation Results

## Hard enforcement of estimates

- ×Outline
- ×Motivation
- ×Observations
- ×Estimation
- ×Simulation
- ×Policy
- ×Results
- ×Conclusion



# A Probabilistic Method

---

---

## Probabilistic enforcement of estimates

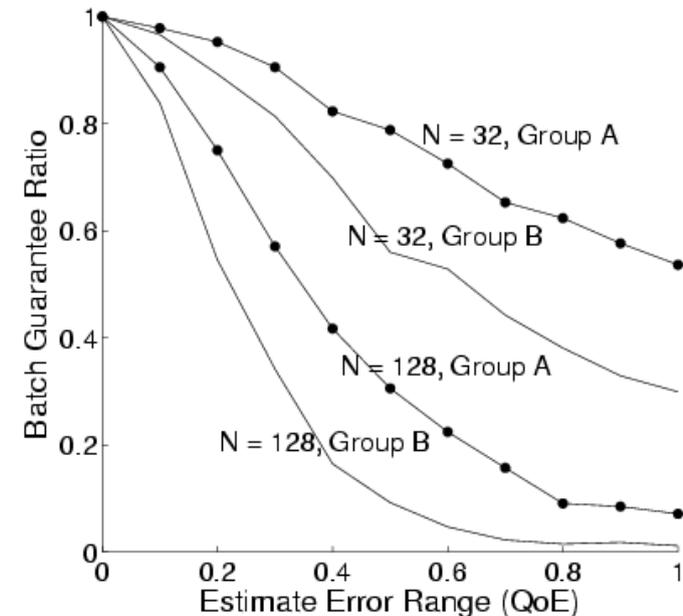
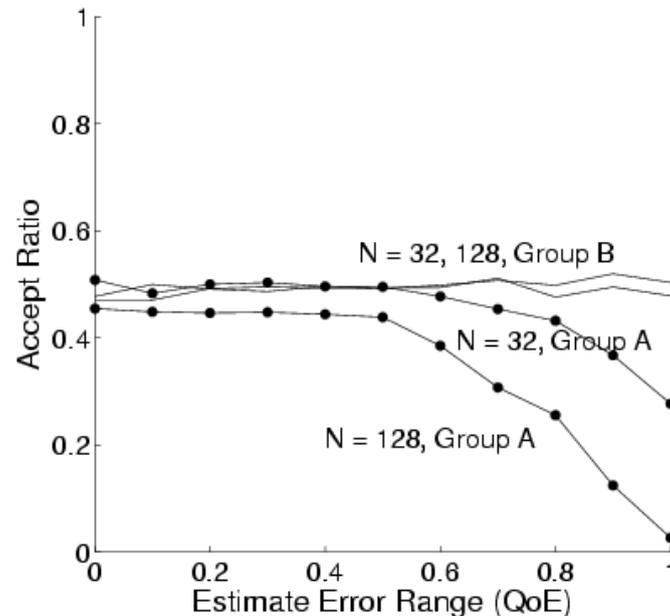
- xOutline
- xMotivation
- xObservations
- xEstimation
- xSimulation
- xPolicy
- xResults
- xConclusion

- Kill misbehaving job with probability  $p$ :  
“If a job has exceeded its estimate by a factor of  $1 + p$ , kill it with probability  $p$ .”
- Attempt to forgive mildly misbehaving jobs while discouraging gross underestimation. Give the best results to the best users.

# Simulation Results

## Probabilistic enforcement of estimates

- × Outline
- × Motivation
- × Observations
- × Estimation
- × Simulation
- × Policy
- × Results
- × Conclusion



- Group A: reality-centered estimates
- Group B: deliberate under-estimates

# Conclusion

---

---

xOutline  
xMotivation  
xObservations  
xEstimation  
xSimulation  
xPolicy  
xResults  
xConclusion

- Hard job runtime estimate enforcement may be overkill in unpredictable environments, especially for batches of interdependent jobs.
- New enforcement policies should be based on the reality of unpredictable computing environments and the possibility of contention for resources.