Models and Heuristics for Robust Resource Allocation in Parallel and Distributed Computing Systems

D. Janovy, J. Smith, H. J. Siegel, and A. A. Maciejewski

Colorado State University

Outline

- models and metrics for robust resource allocation
- deterministic robustness
- stochastic robustness
- current research

Simple Example: Cluster Computing

- cluster of *M* heterogeneous machines
- *T* independent tasks
  - production environment
- given estimated time to compute task *i* on machine *j*
  - for each task on each machine
- need to allocate resources to tasks
  - static, off-line allocation
- makespan: time to complete all tasks
  - minimize estimated makespan
- want to be robust with respect to uncertainty of execution time estimates
  - actual makespan $\leq$ 1.2 $\times$ estimated makespan
  - 1.2 is a user specified constraint

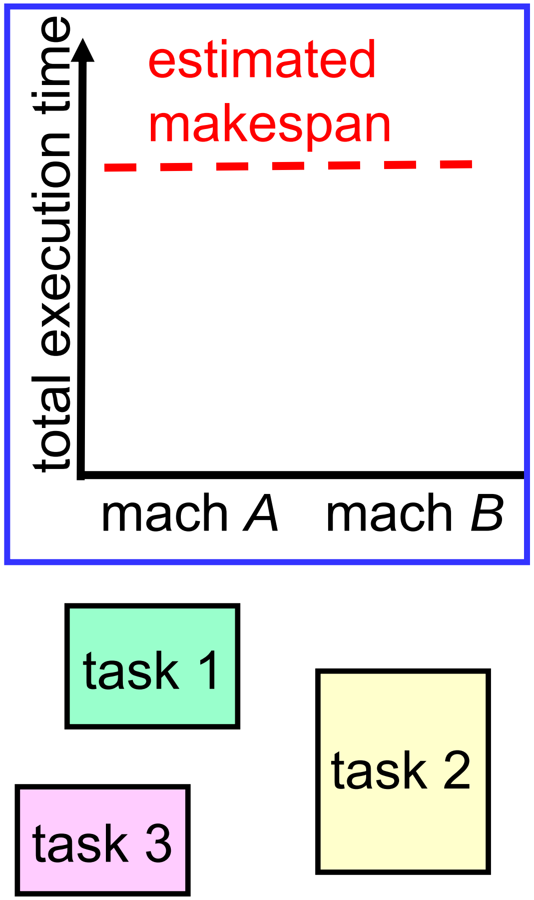

Deterministic Robust Resource Allocation

- reference to our group's work

  "Measuring the Robustness of a Resource Allocation," *IEEE Transactions on Parallel and Distributed Systems*, July 2004

The Three Robustness Questions

- what behavior of the system makes it robust?
  - ex. actual makespan ≤ 1.2 estimated makespan
- what uncertainties is the system robust against?
  - ex. variations in actual execution times
- quantitatively, exactly how robust is the system?
  - ex. largest collective increase (Euclidian measure) in actual task completion times before a makespan constraint violation occurs

Cluster Computing Example: Robustness Metric

FePIA procedure for cluster computing example

- performance **Fe**atures
  - ex. F_j: finishing time of machine j for an allocation
- **P**erturbation parameter
  - ex. vector of actual task computation times for given allocation
- **I**mpact of perturbation parameter on performance features
  - ex. F_j = sum of computation times of tasks on machine j
- **A**nalysis
  - ex. robustness radius for machine j
    - smallest collective increase between estimated and actual computation times that will cause F_j to be > 1.2 × estimated makespan
  - robustness metric = minimum robustness radius

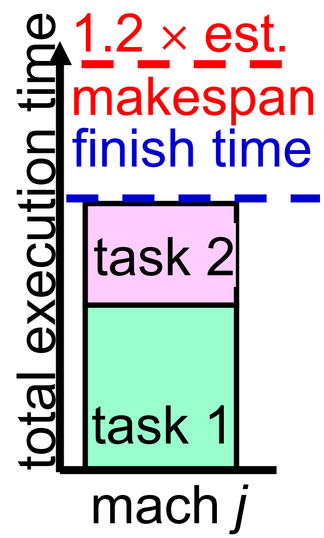

Results of FePIA Analysis Step for Machine j

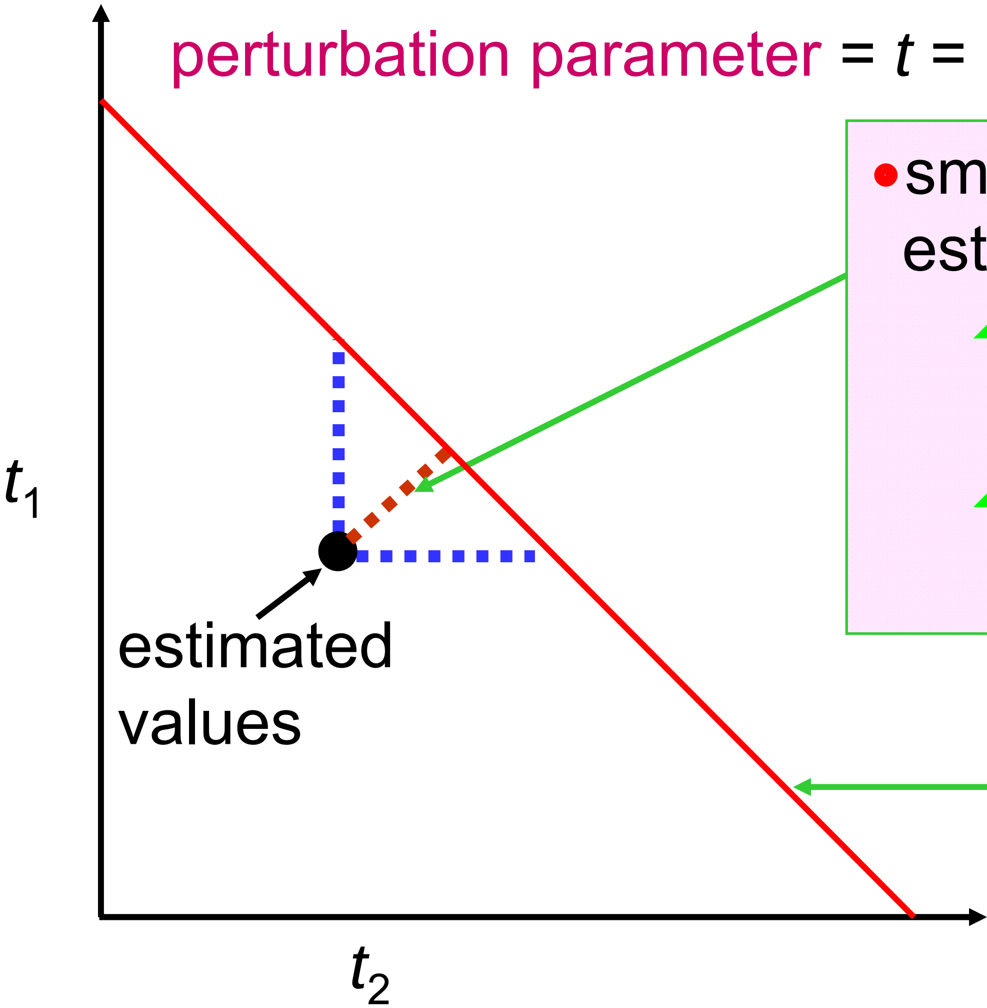

- smallest change in t from its estimated value
  - that causes violation of specified constraint
  - robustness radius r_{μ} for resource allocation μ

plot for resource allocation μ where $\{t| F_j(t) = 1.2 \times \text{estimated makespan}\}$

- robustness metric: minimum robustness radii over all machines

Colorado State University

Cluster: Robustness Versus Estimated Makespan

- 1000 random resource allocations for 20 tasks and 5 machines
- normalized robustness = robust metric / estimated makespan

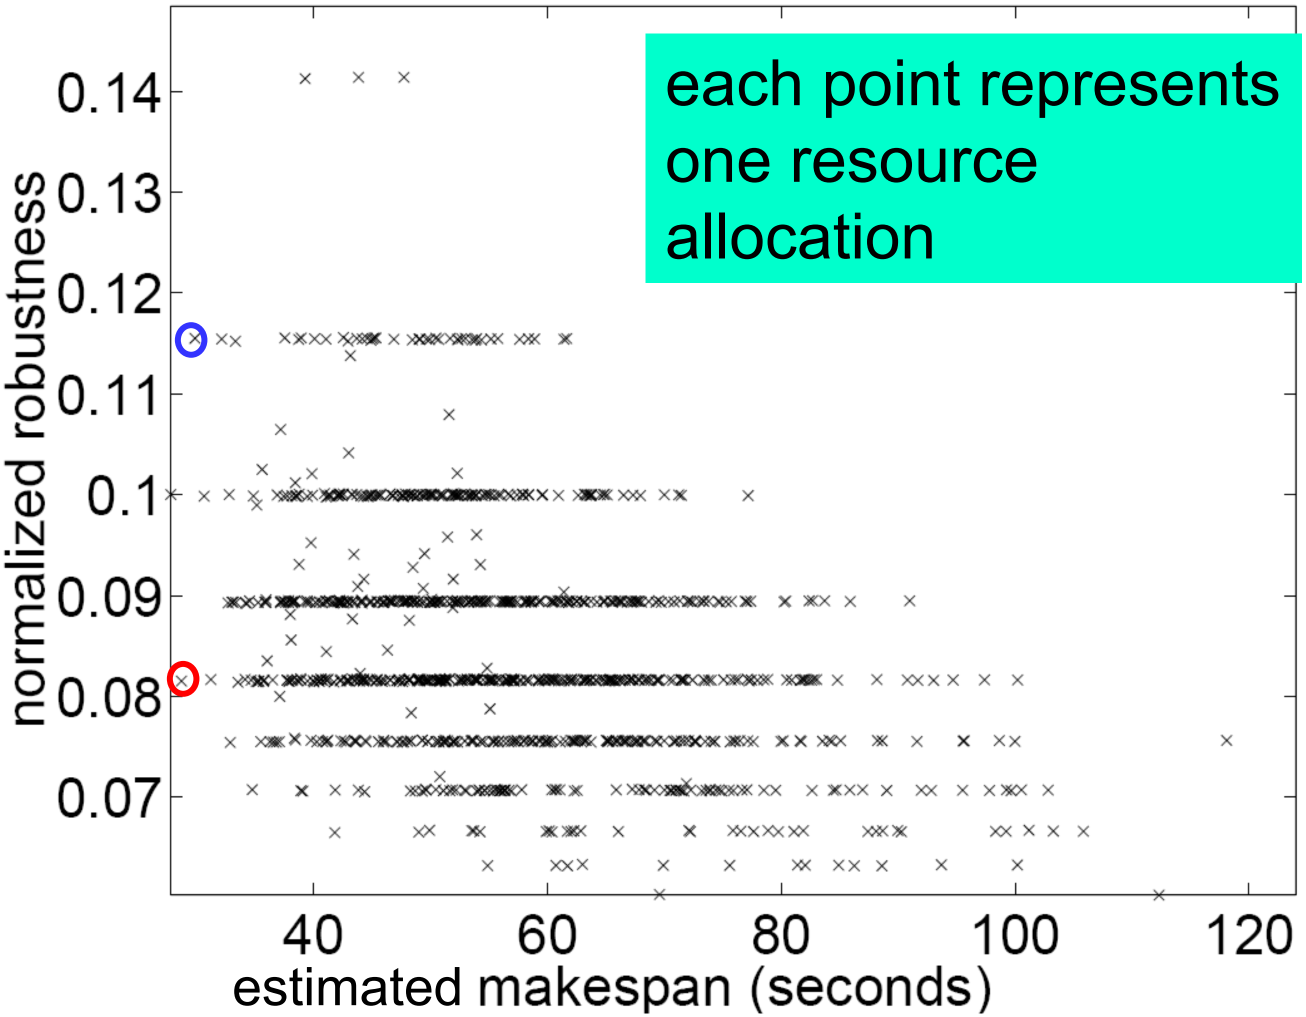

Deterministic Robustness: Static Mapping

- reference to our group's work
  "Robust Static Allocation of Resources for Independent Tasks under Makespan and Dollar Cost Constraints," *Journal of Parallel and Distributed Computing*, accepted, to appear
- robust resource allocations with uncertain task execution times
- two problem variations were considered
  - fixed machine suite in a production environment
    - task execution times based on estimates
    - maximize robustness within a given makespan constraint
  - machine selection with a purchasing cost constraint
    - machines vary in performance and cost
    - select a suite of machines that maximizes robustness while meeting makespan constraints
- six heuristics were designed and evaluated

Deterministic Robustness: Dynamic Mapping

- reference to our group's work
  "Dynamic Resource Allocation Heuristics that Manage Tradeoff between Makespan and Robustness,"
  Journal of Supercomputing, accepted, to appear
- establish the notion of deterministic dynamic robustness
- task arrival times are not known in advance
- fixed suite of machines
- robust resource allocations with uncertain task execution times
- two problem variations were considered
  - minimize makespan while maintaining a specified level of robustness over all mapping events
  - maximize the minimum robustness over time within a given makespan constraint
- ten heuristics were designed and evaluated

Stochastic Robustness Model: Cluster Example

- reference to our group's work
  A Stochastic Approach to Measuring the Robustness of Resource Allocations in Distributed Systems," *2006 International Conference on Parallel Processing*, Aug. 2006
- stochastic robustness metric is probability that completion time of all machines meets requirements

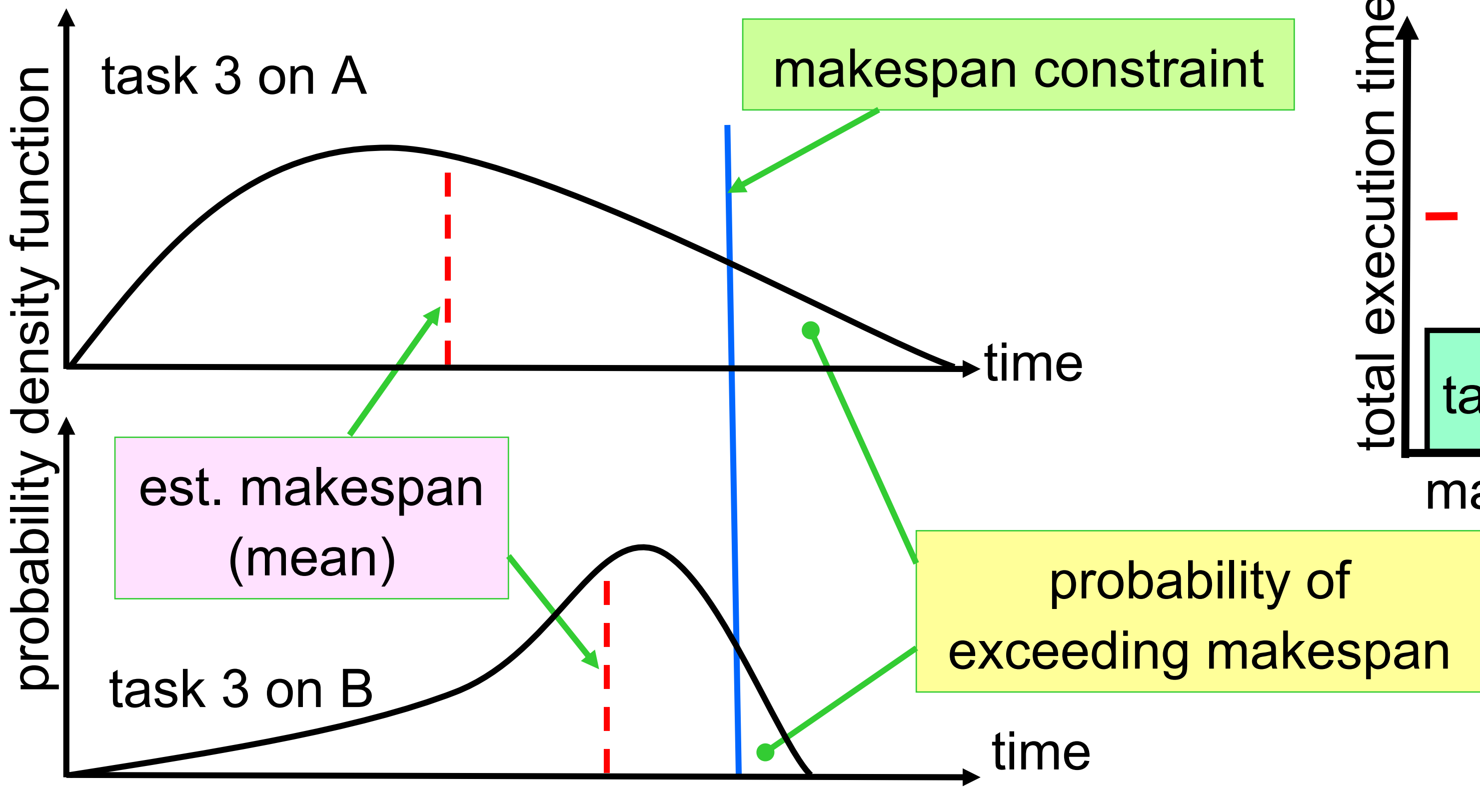

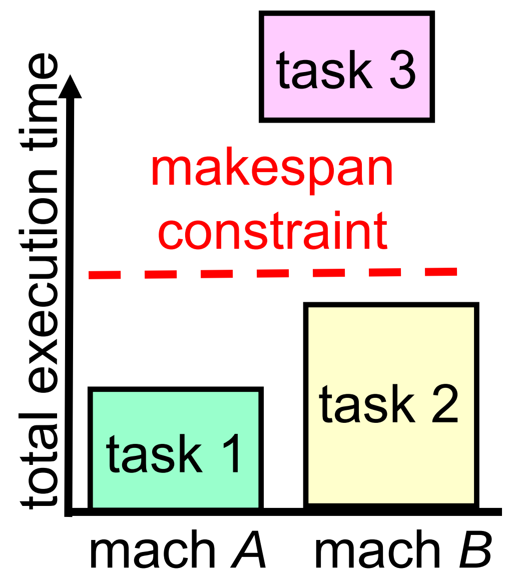

Stochastic Robustness: Mapping Example

- periodic data sets received from sensors
- changes in input data cause variability in execution time
- data needs to be processed before next data set arrives
- goal: minimize the period Λ, between data sets
  to allow more sets to be processed

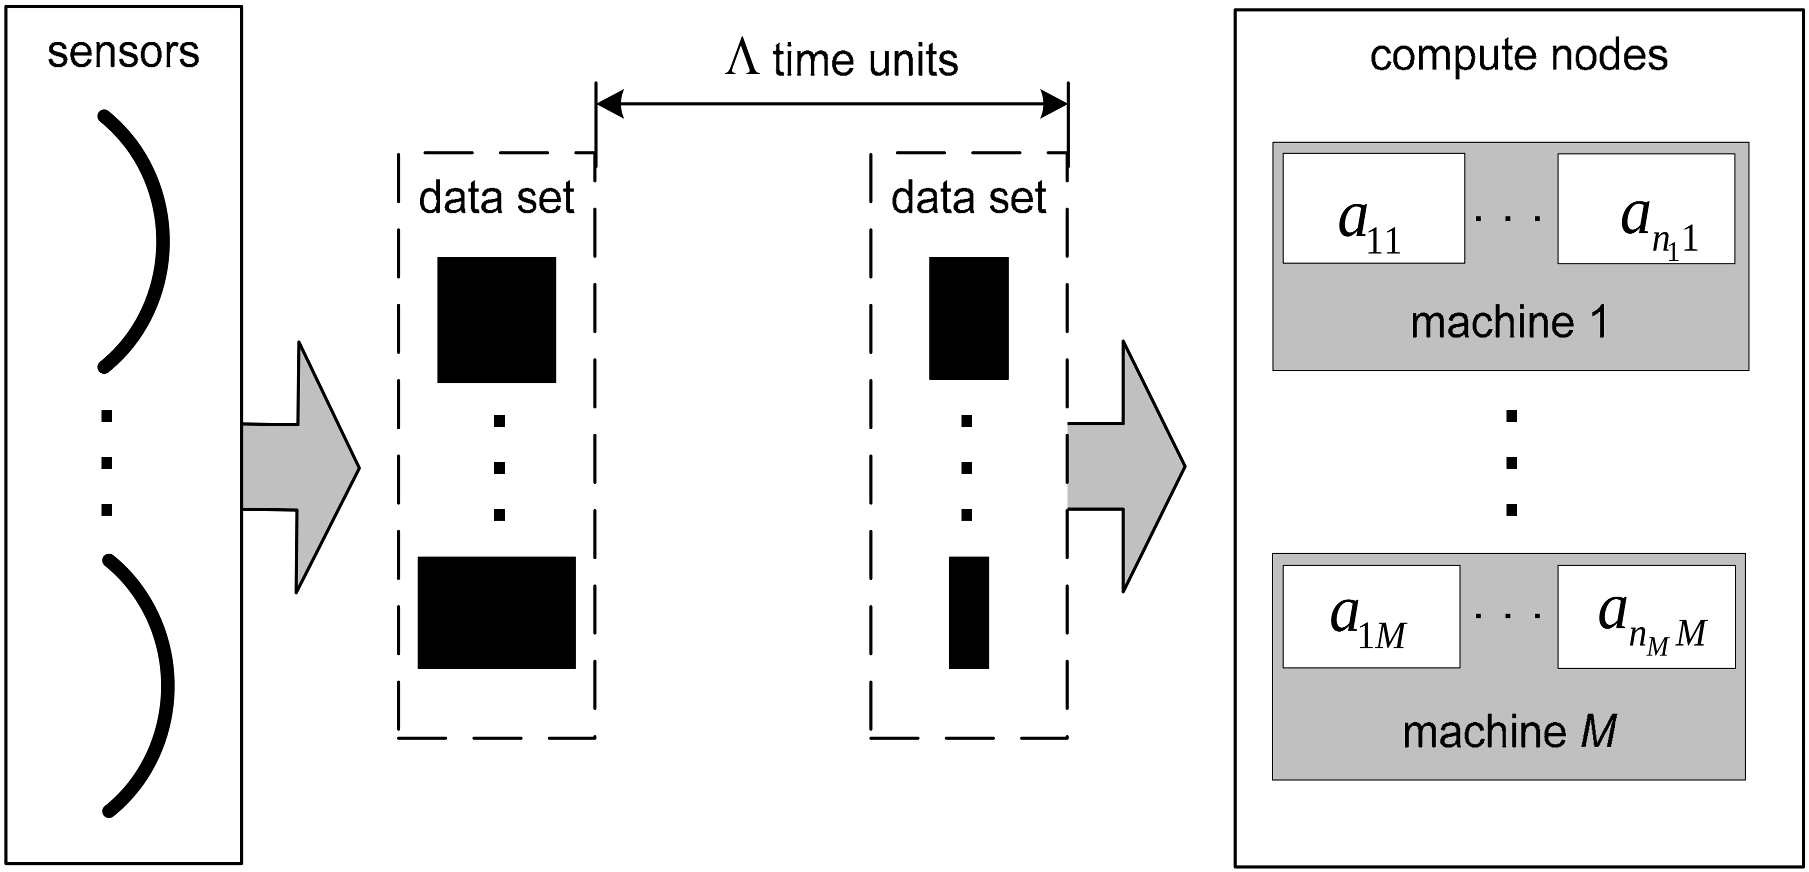

Stochastic Robustness: Greedy & Iterative Static Mapping

- reference to our group's work
  "Greedy Approaches to Static Stochastic Robust Resource Allocation for Periodic Sensor Driven Distributed Systems," *2006 International Conference on Parallel and Distributed Processing Techniques and Applications*, June 2006
  - ex. two-phase min-min type approach

- reference to our group's work
  "Iterative Algorithms for Stochastically Robust Static Resource Allocation in Periodic Sensor Driven Clusters," *18th IASTED International Conference on Parallel and Distributed Computing and Systems*, Nov. 2006
  - iterative static heuristics (versus greedy) look at entire resource allocation during each iteration
  - ex. genetic algorithm type approach

Stochastic Robustness: Dynamic Mapping

- reference to our group's work
  "Measuring the Robustness of Resource Allocations in a Stochastic Dynamic Environment,"
  21st International Parallel and Distributed Processing Symposium, Mar. 2007, Session 26, Thursday ~ 4:30
- establish the notion of stochastic dynamic robustness
- environment consisted of a heterogeneous, distributed computing system designed for a high volume web site
- tasks arrival times are not known in advance
- task execution times described by probability mass functions
- dynamic stochastic robustness metric is defined as the average over all mapping events of the instantaneous stochastic robustness metric values
- used to validate the dynamic robustness as a predictor of performance

Current Robustness Research

- design of robust resource allocation in distributed systems under random machine failures using stochastic data
  - reallocation of resources is allowed
  - redundant task assignment to ensure high priority task complete when resource reallocation is not allowed
- resource allocations for virtual world environments
  - where the number of users is uncertain
  - response time is robust to the number of users being added to the system
  - also applies to calculations being distributed across a P2P network
- how sensitive are resource allocation algorithms to errors in the models of uncertainty
- use of dynamic stochastic robustness measure to guide resource allocations in a dynamic environment

Colorado State University