# Statistical Predictors of Computing Power in Heterogeneous Clusters

Ron C. Chiang, Anthony A. Maciejewski,
Arnold L. Rosenberg, H. J. Siegel

Colorado State University

# Motivation

- if we can predict the performance of a cluster by its statistical measures
    - mean computing rates
    - variance in computing rate

- then we can
    - quickly compare different clusters' performance
    - understand how to construct a high performance cluster

# Questions Investigated

- does cluster $\boldsymbol{C}_1$ outperform cluster $\boldsymbol{C}_2$?
  - when cluster $\boldsymbol{C}_1$ has a faster mean computing rate than cluster $\boldsymbol{C}_2$
  - when cluster $\boldsymbol{C}_1$ and cluster $\boldsymbol{C}_2$ have the same mean computing rate, but $\boldsymbol{C}_1$ has a higher variance in computing rate

we answer these questions within the framework of the Cluster-Exploitation Problem

# The Cluster-Exploitation Problem

- server $c_0$ must complete as many units of work as possible on cluster ***C*** within a given lifespan of *L* time units

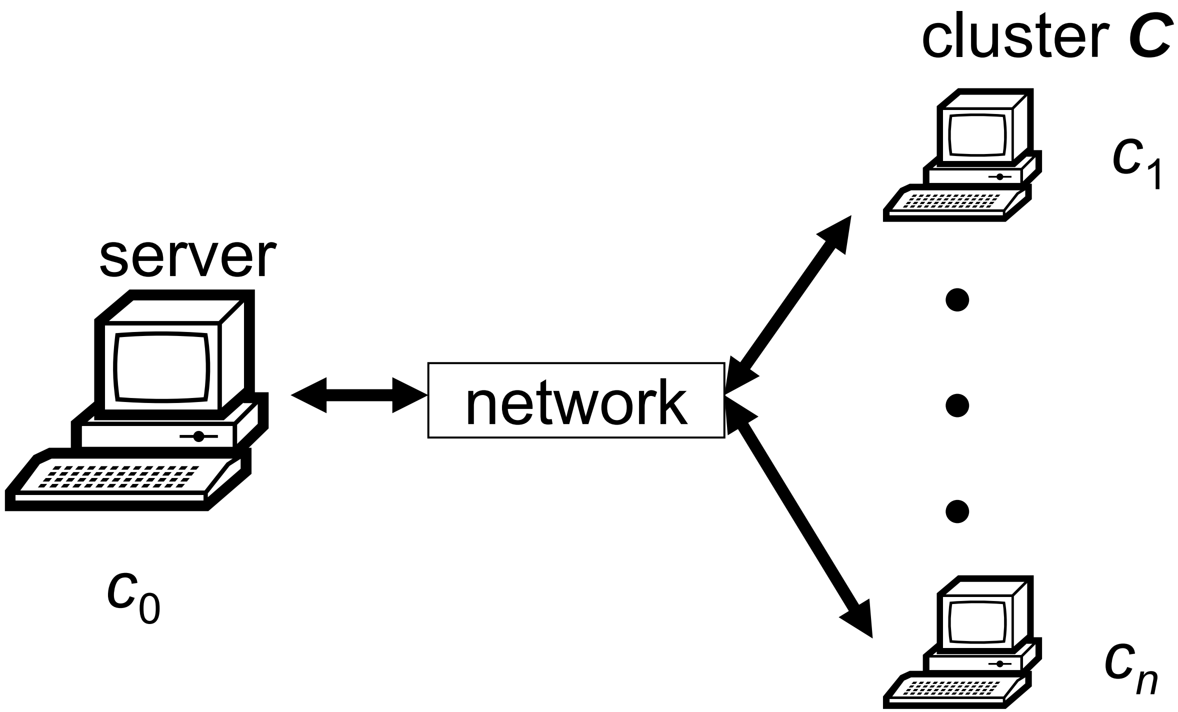

# The Architectural Model

- the server $c_0$
- a "cluster" ***C*** with $n$ computers
  - $c_1, \ldots, c_n$
- $c_i$ completes one unit of work in time $\rho_i$
- heterogeneity profile of ***C***
  - $\rho_1, \ldots, \rho_n$

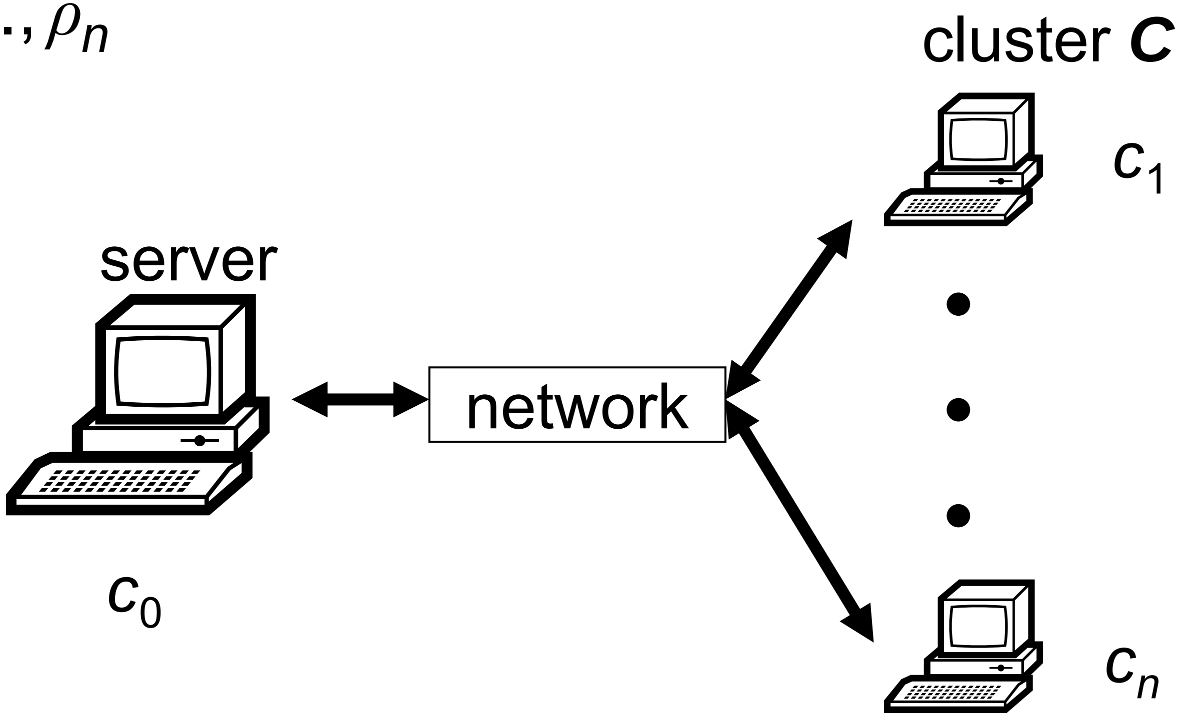

# Worksharing Protocol

- a schedule that solves the Cluster-Exploitation Problem

- three steps
  - $c_0$ transmits work to $c_i$ in a single message
  - $c_i$ computes the work immediately after receiving it
  - as soon as $c_i$ completes its work, it transmits results to $c_0$ in a single message

- FIFO (First-In-First-Out) worksharing protocol
  - coincident startup ordering and finish orderings
  - optimal schedules for the Cluster-Exploitation Problem under all startup orderings
  - we use it to study node-heterogeneity in clusters

# Cluster Performance Comparison

- if cluster $C_1$ completes more work than cluster $C_2$
  - cluster $C_1$ outperforms cluster $C_2$

# Cluster Performance Comparison

- if cluster $C_1$ completes more work than cluster $C_2$
  - cluster $C_1$ outperforms cluster $C_2$
- we predict which cluster has a better performance by comparing clusters'

# Cluster Performance Comparison

- if cluster $\boldsymbol{C}_1$ completes more work than cluster $\boldsymbol{C}_2$
  - cluster $\boldsymbol{C}_1$ outperforms cluster $\boldsymbol{C}_2$
- we predict which cluster has a better performance by comparing clusters'
  - mean computing rate
    - $$\bar{\rho} = \frac{(\rho_1 + \cdots + \rho_n)}{n}$$

recall:

$\rho_i$ is the time for computer $c_i$ to complete one unit of work

# Cluster Performance Comparison

- if cluster $\boldsymbol{C}_1$ completes more work than cluster $\boldsymbol{C}_2$
  - cluster $\boldsymbol{C}_1$ outperforms cluster $\boldsymbol{C}_2$
- we predict which cluster has a better performance by comparing clusters'
  - mean computing rate
    - $$\bar{\rho} = \frac{(\rho_1 + \cdots + \rho_n)}{n}$$
  - variance in computing rate
    - $$VAR = \sum_{i=1}^{n} \frac{(\rho_i - \bar{\rho})^2}{n}$$

recall:

$\rho_i$ is the time for computer $c_i$ to complete one unit of work

# Simulation Procedure

- evaluate and compare clusters' productivity under different scenarios

- generate sample clusters with
  - different distributions of mean computing rates
  - different distributions of variance in computing rate

# Mean Computing Rates: Uniform Distribution

- different distributions of **mean computing rates** (between 0.01 and 1 time units per task)
  - **uniform distribution**
    - assume equal numbers of clusters with different mean computing rates

# Mean Computing Rates: Normal Distribution

- different distributions of **mean computing rates** (between 0.01 and 1 time units per task)
  - uniform distribution
  - **normal distribution**
    - assume most clusters have moderate mean computing rates

# Mean Computing Rates: Household Income Dist.

- different distributions of **mean computing rates** (between 0.01 and 1 time units per task)
  - uniform distribution
  - normal distribution
  - **household income distribution**
    - assume the computing power of a cluster reflects its owner's income

# Mean Computing Rates: Household Income Dist.

- different distributions of **mean computing rates** (between 0.01 and 1 time units per task)

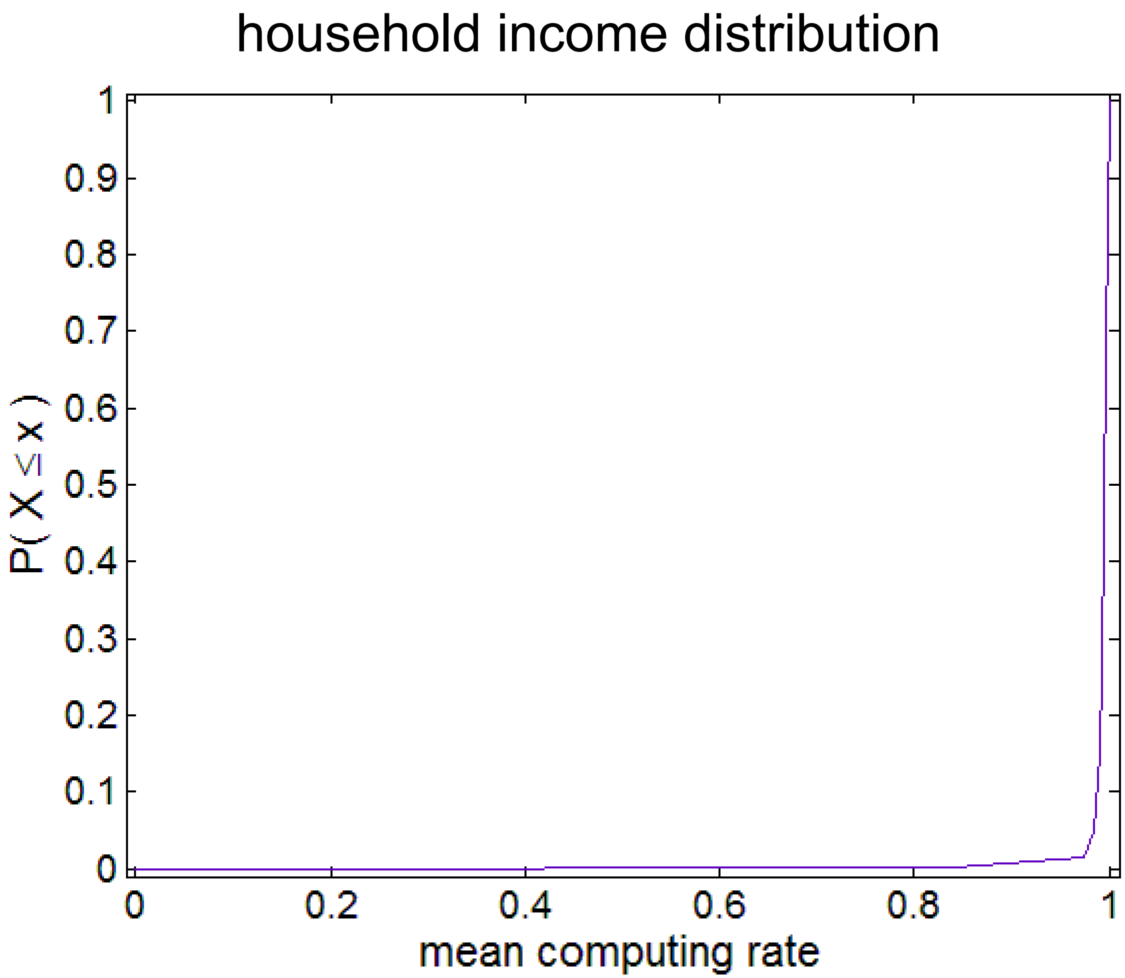




Colorado State University

# Variances: Uniform Distribution

- different distributions of **variances** in computing rate
  - **uniform distribution**
    - assume equal numbers of clusters with different variances in computing rate

# More Sample Profiles Have Small Variances

- different distributions of **variances** in computing rate
  - **"small variance" distribution**: assume more clusters have small variances

$$P_1(X \geq x) = \frac{e^{1-x} - 1}{e - 1}$$

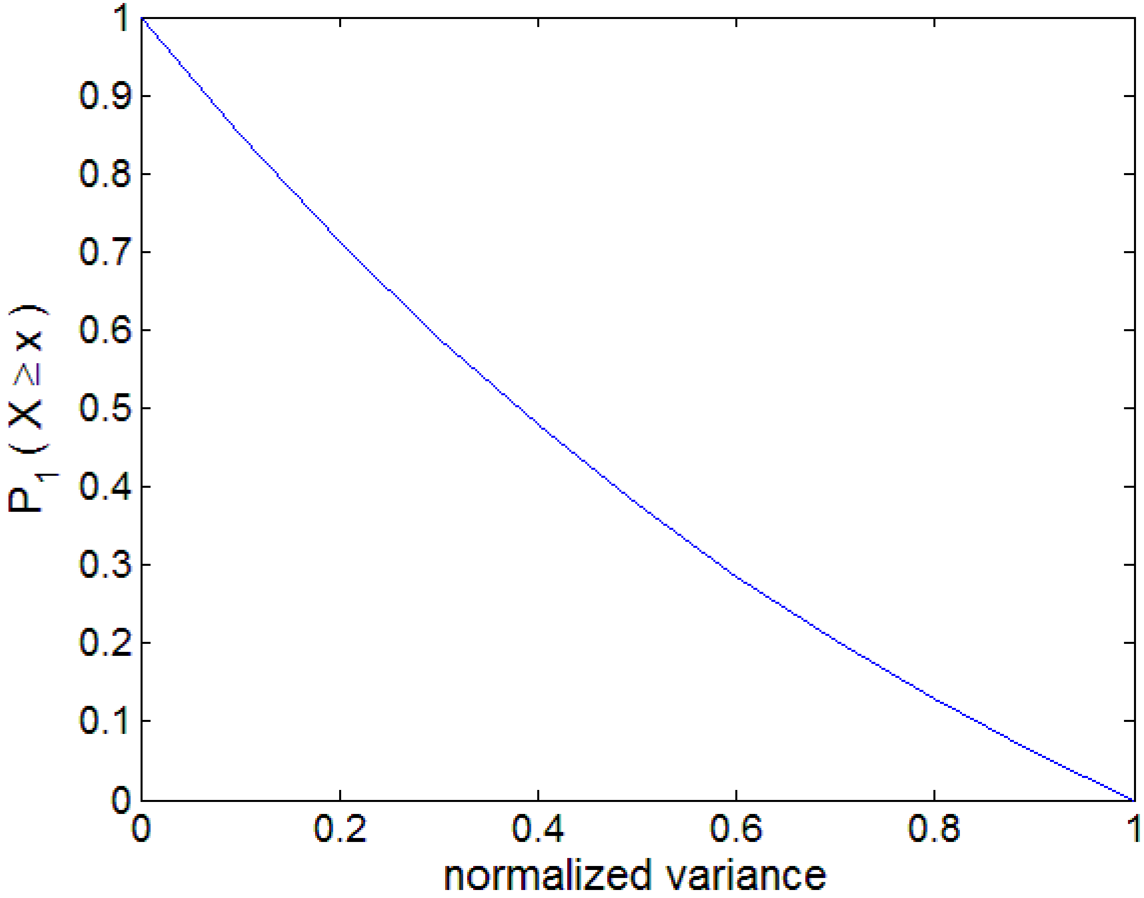

# Even More Sample Profiles Have Small Variances

- different distributions of **variances** in computing rate
  - **"more small variance" distribution**: assume even more clusters have small variances

$$P_2(X \geq x) = \frac{e^{1-\sqrt{x}} - 1}{e - 1}$$

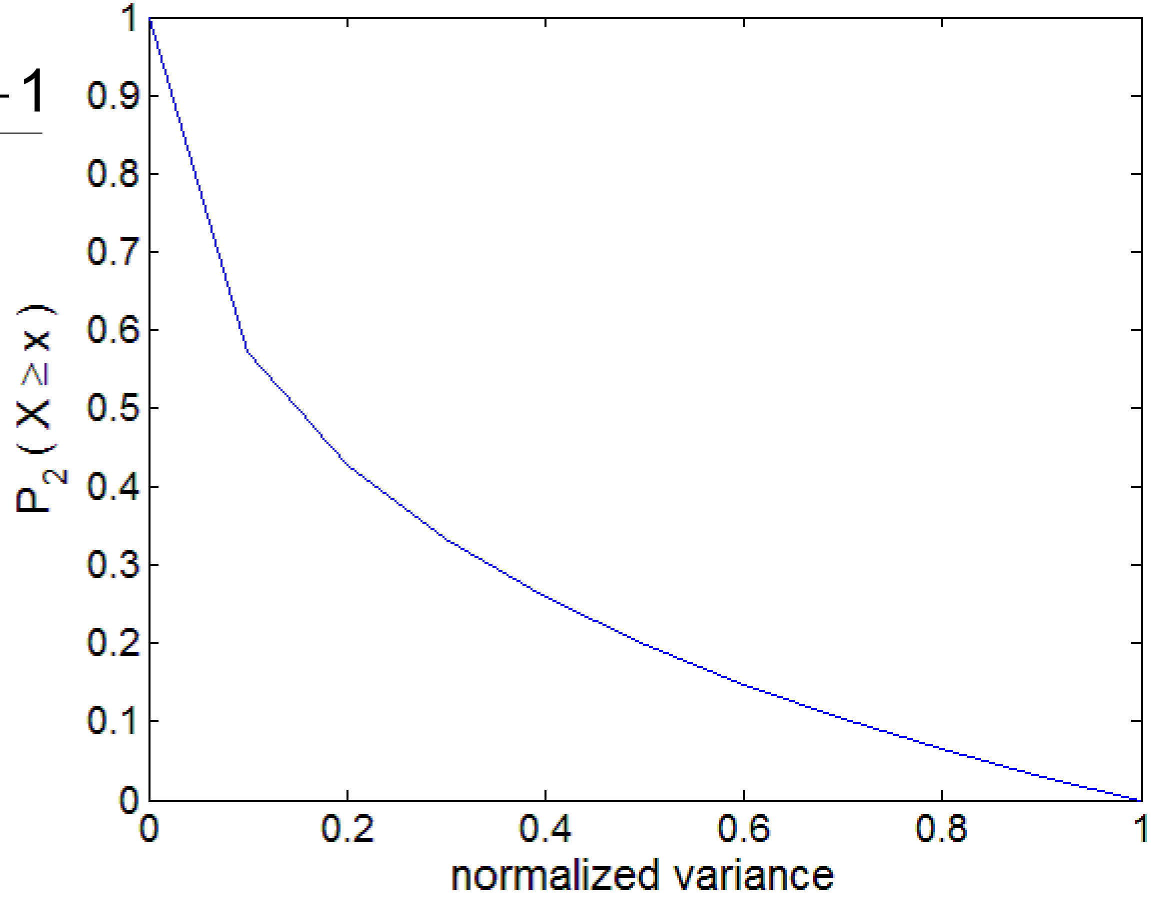

# Comparing Mean Computing Rate

- compare cluster pairs with different mean computing rates
- failed prediction
  - if the cluster with a smaller mean computing rate does not outperform the other cluster
- % of failed predictions
  - failed predictions / all possible cluster pairs
  - 10,000 sample cluster profiles for each different cluster sizes ($2^1$ to $2^{12}$)

# Comparing Mean Computing Rate – Uniform

- % of failed predictions decreases when more clusters have small variances in computing rate
  - cluster performance is closer to a homogeneous cluster when a cluster has a small variance

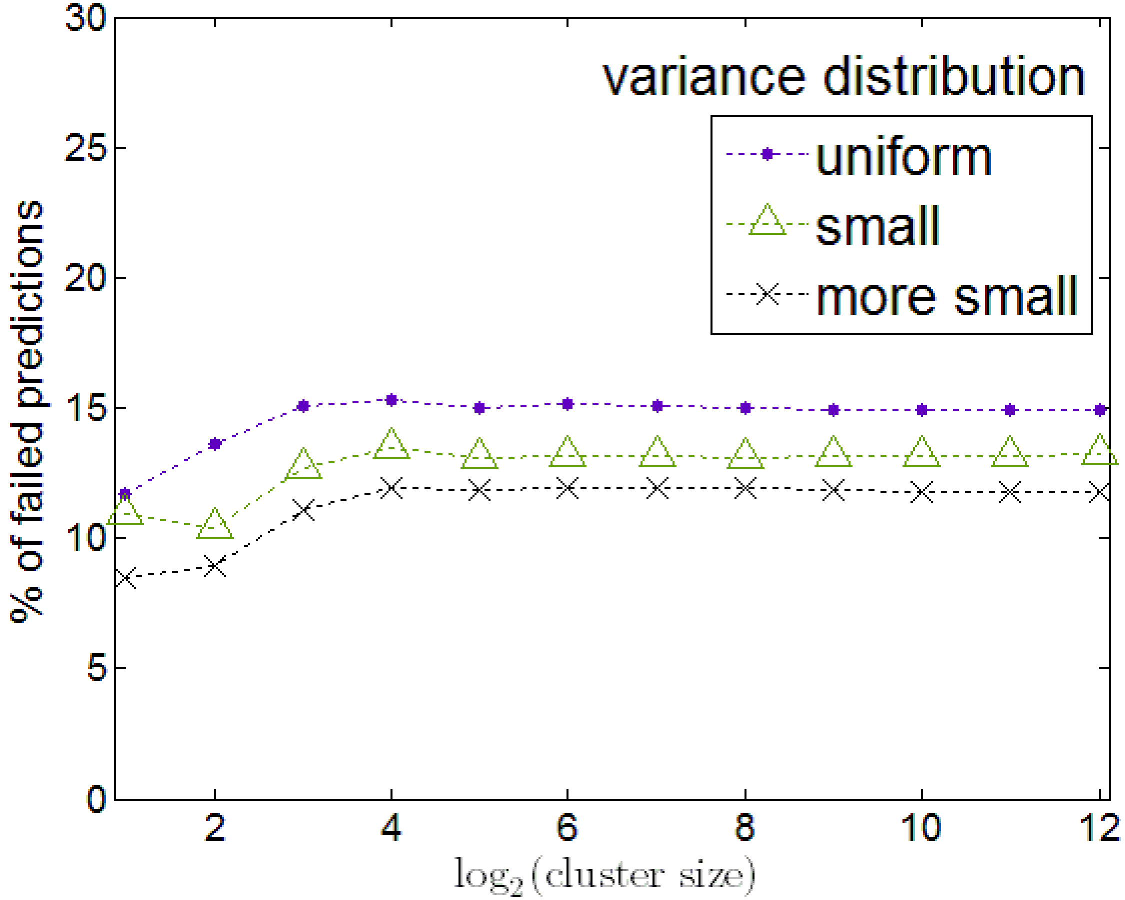


reminder:

- a failed prediction: cluster with a faster mean computing rate does not outperform the other cluster

Colorado State University

# Comparing Mean Computing Rate – Normal

- % of failed predictions decreases when more clusters have small variances in computing rate
- higher % of failed predictions than uniformly distributed
  - because more clusters have close mean computing rates

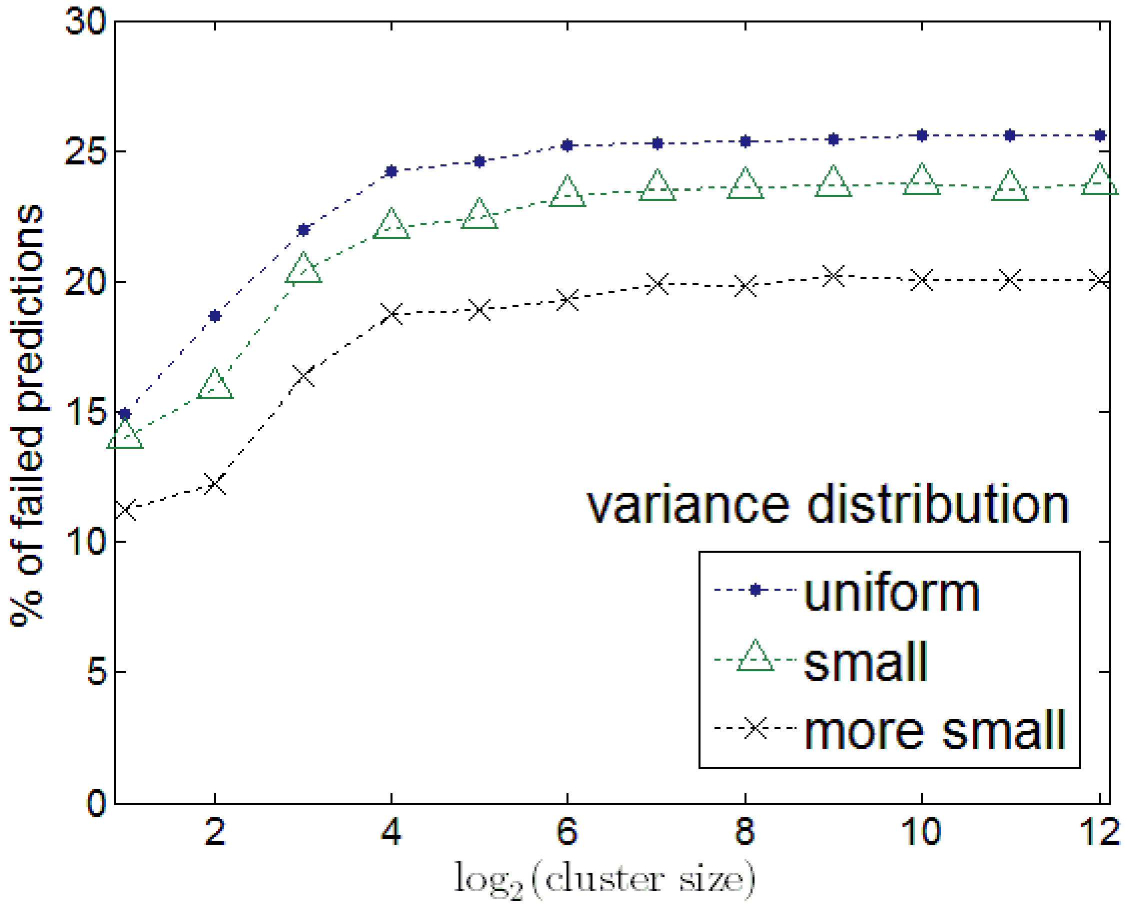


reminder:

- a failed prediction: cluster with a faster mean computing rate does not outperform the other cluster

# Comparing Mean Computing Rate – Household Income Dist.

- mean computing rates have a household income distribution
- % of failed predictions decreases when more clusters have small variances in computing rate

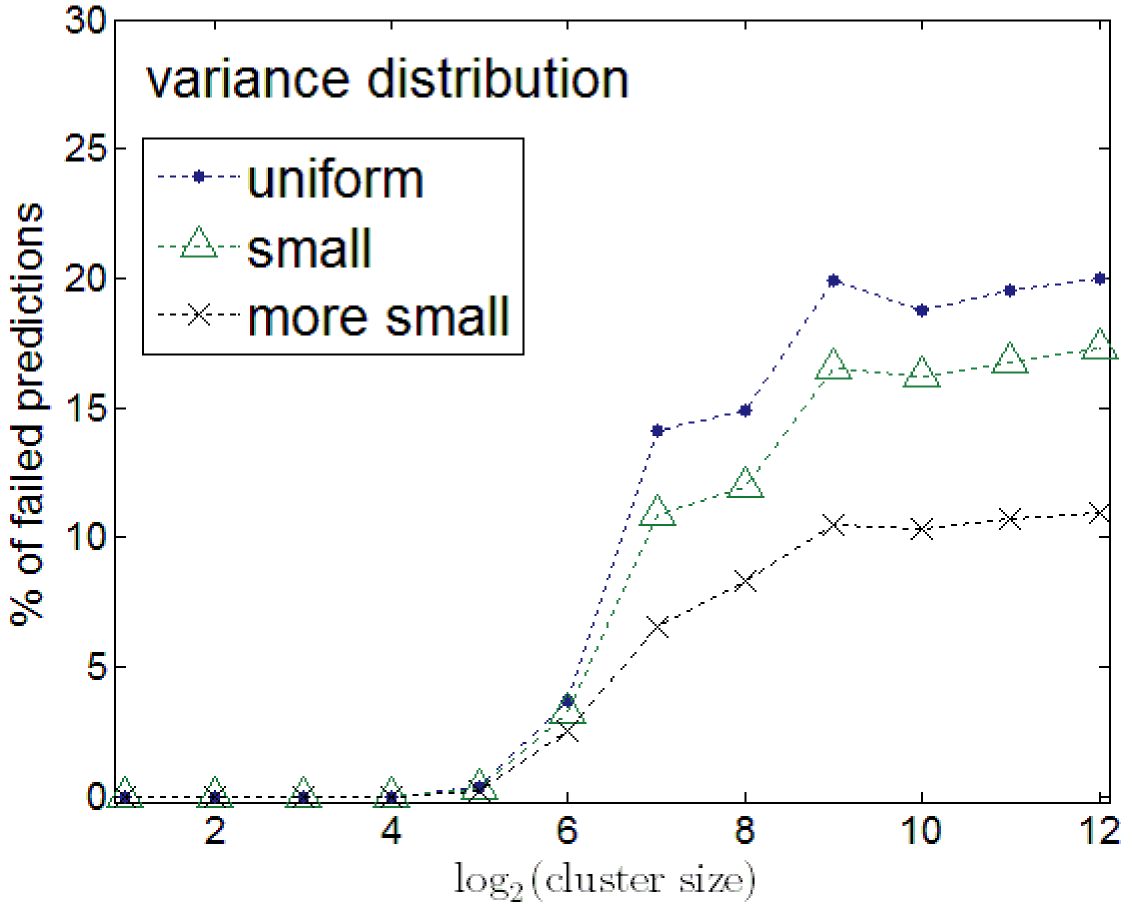


reminder:

- a failed prediction: cluster with a faster mean computing rate does not outperform the other cluster

# Comparing Variance in Computing Rate

- compare cluster pairs with the same mean computing rate but different variances in computing rate

- failed prediction
  - if the cluster with higher variance in computing rate does not outperform the other cluster

- % of failed predictions
  - failed predictions / all possible cluster pairs
  - 10,000 sample cluster profiles for each different cluster sizes ($2^1$ to $2^{12}$)

# Comparing Variance in Computing Rate - Uniform

- % of failed predictions decreases when more clusters have small variances in computing rate
  - some clusters that have big variances do not perform as expected

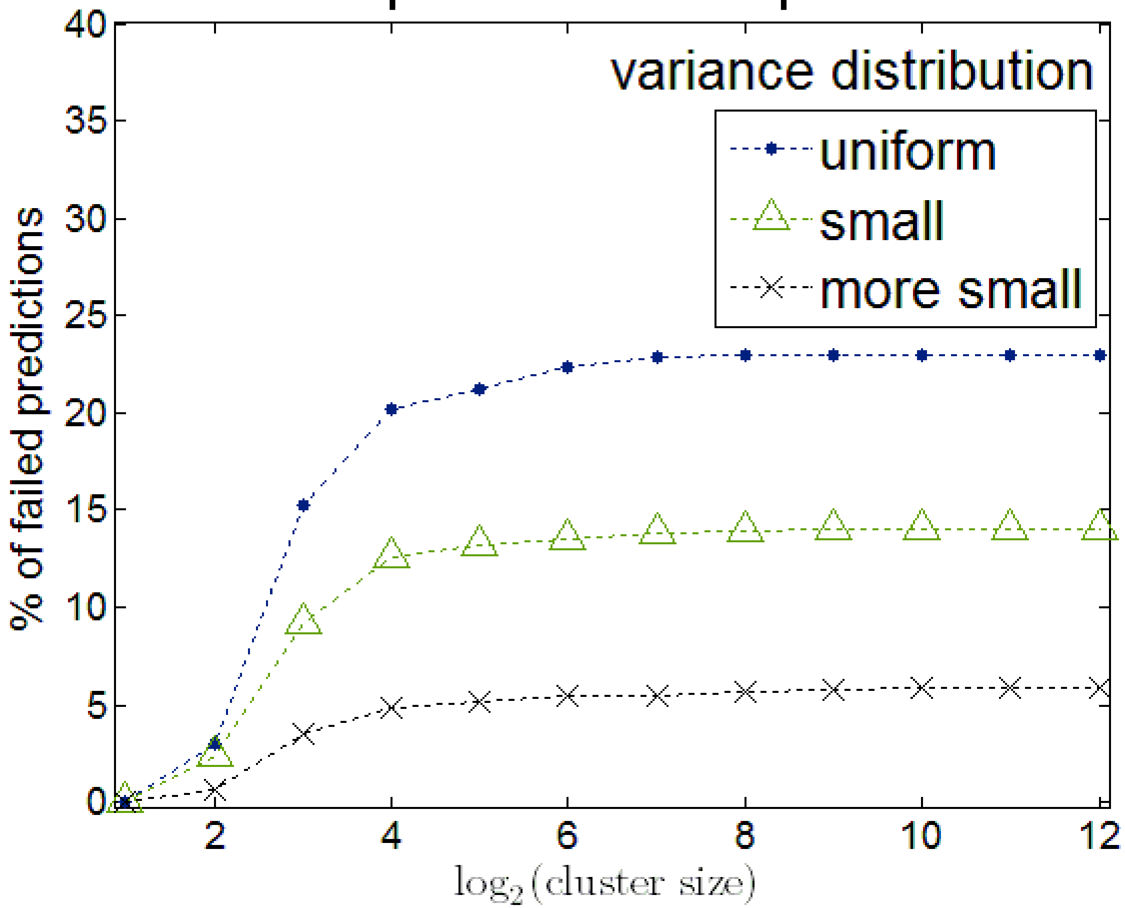


reminder:

- a failed prediction: cluster with a higher variance in computing rate does not outperform the other cluster

# Comparing Variance in Computing Rate - Normal

- lower % of failed predictions than uniformly distributed
- variance is a better measure of cluster performance
  - at a slower mean computing rate than a faster one

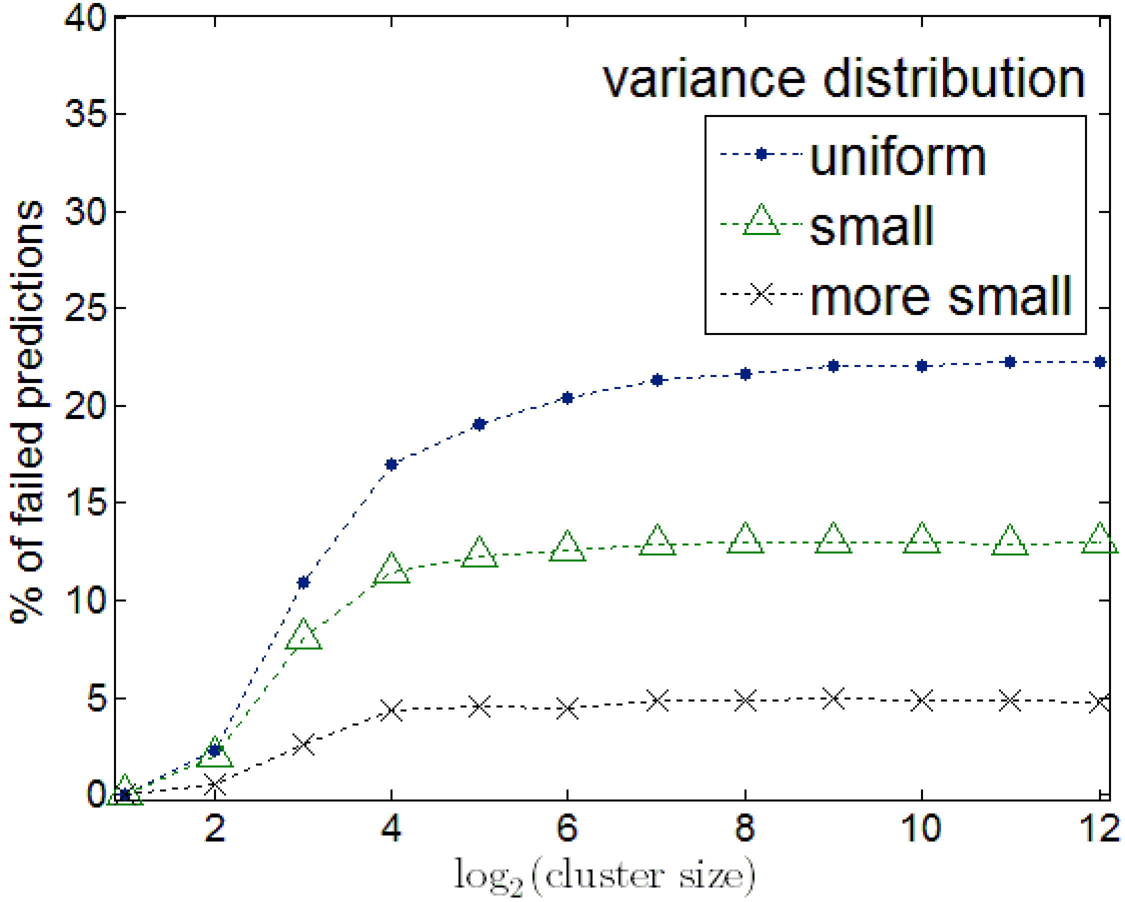


reminder:
- a failed prediction: cluster with a higher variance in computing rate does not outperform the other cluster

# Comparing Variance in Computing Rate – Household Income

- mean computing rates have a household income distribution
- percentages of failed predictions are all lower than 5%

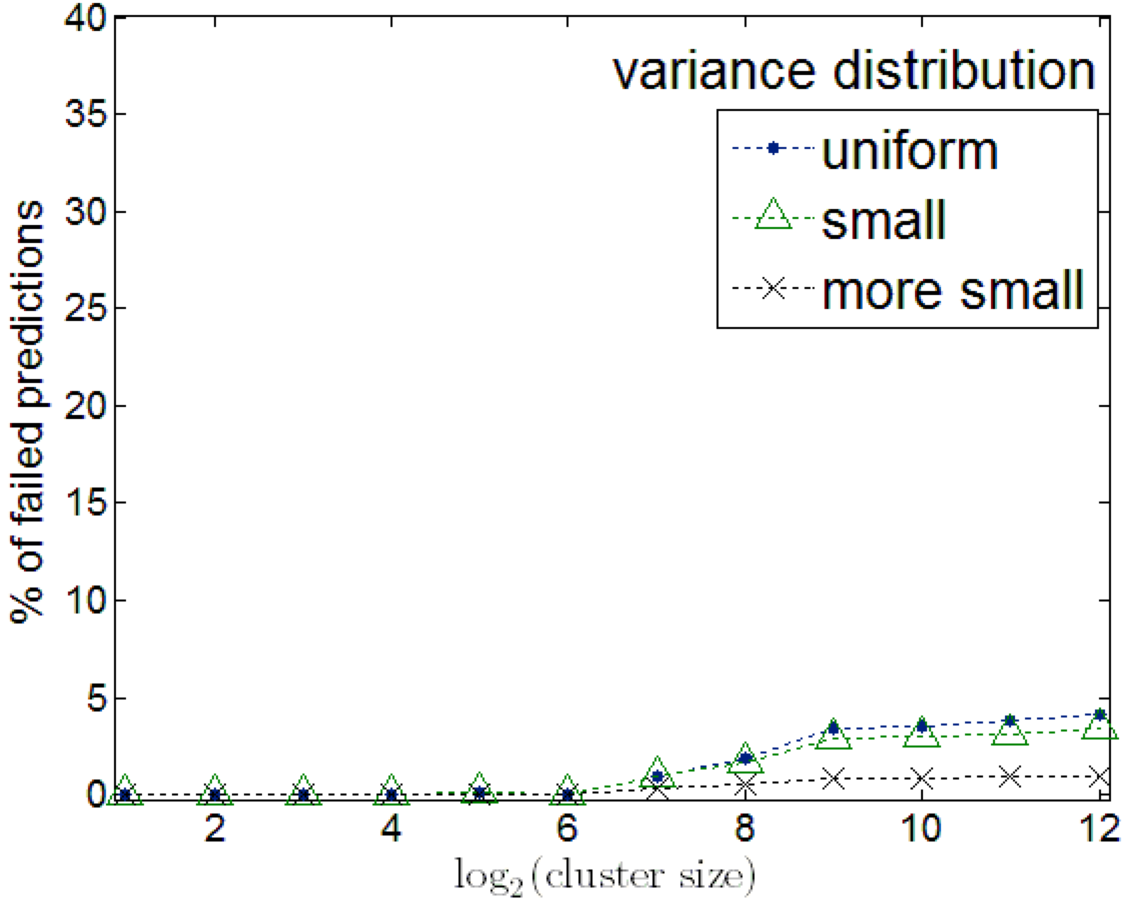


reminder:

- a failed prediction: cluster with a higher variance in computing rate does not outperform the other cluster

# Comparing Variance in Computing Rate

- big mean computing rate has small % of failed predictions

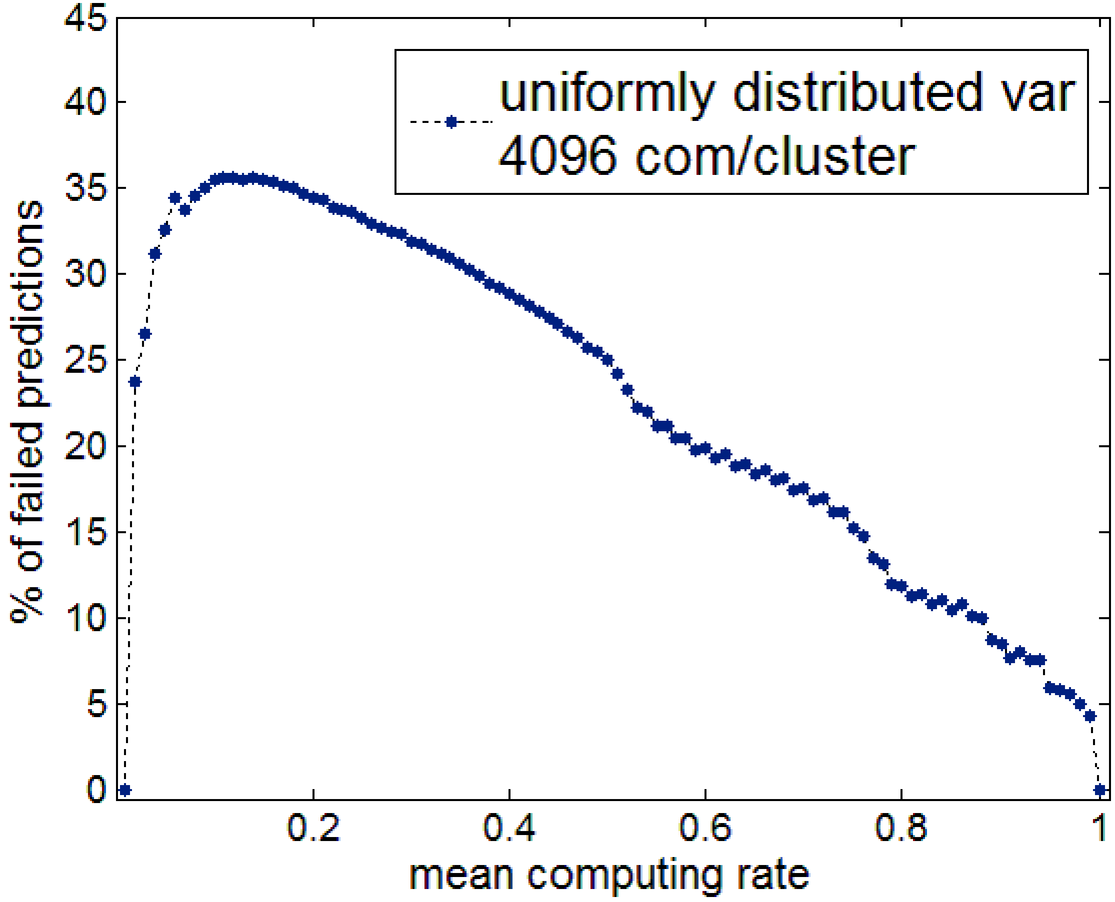


other cluster sizes and variance distributions show similar results

# Conclusions – Using Mean to Predict Performance

- % of failed predictions decreases when more clusters have small variances in mean computing rate
  - because cluster performance is close to a homogeneous cluster when a cluster has a small variance

- mean computing rate is a better measure of performance
  - when mean computing rates are uniformly distributed than normally distributed
  - because more clusters have close mean computing rates

# Conclusions – Using Variance to Predict Performance

- % of failed predictions decreases when more clusters have small variances in mean computing rate
  - some clusters that have big variances do not perform as expected

- variance is a better measure of cluster performance
  - at a slower mean computing rate than a faster one
  - when mean computing rates have a household income distribution as opposed to other distributions

## Questions?