Toward Understanding Heterogeneity in Computing

Arnold L. Rosenberg
Ron C. Chiang
Electrical & Computer Engineering
Colorado State University
Fort Collins, CO, 80523, USA

Heterogeneity in Computing

One encounters *HETEROGENEITY* in virtually all modern computing systems

Heterogeneity in Computing

One encounters heterogeneity in virtually all modern computing systems

- Computers in clusters/grids differ in power (*NODE-HETEROGENEITY*).

Heterogeneity in Computing

One encounters heterogeneity in virtually all modern computing systems

- Computers in clusters/grids differ in power (*node-heterogeneity*).
- Computers intercommunicate across varied networks (*LINK-HETEROGENEITY*).

Heterogeneity in Computing

One encounters heterogeneity in virtually all modern computing systems

- Computers in clusters/grids differ in power (*node-heterogeneity*).
- Computers intercommunicate across varied networks (*link-heterogeneity*).

WE FOCUS ON *NODE-HETEROGENEITY*.

“Big” Questions about Heterogeneity

Heterogeneity complicates the efficient use of multicomputer platforms

“Big” Questions about Heterogeneity

Heterogeneity complicates the efficient use of multicomputer platforms

— *BUT CAN IT ENHANCE THEIR PERFORMANCE?*

“Big” Questions about Heterogeneity

Heterogeneity complicates the efficient use of multicomputer platforms

— but can it enhance their performance?

HOW DOES ONE STUDY THIS QUESTION RIGOROUSLY?

Detailed Questions about Heterogeneity

- *WHAT MAKES ONE CLUSTER MORE POWERFUL THAN ANOTHER?*

Detailed Questions about Heterogeneity

- What makes one cluster more powerful than another?

- *ARE YOU BETTER OFF . . .*

  — WITH ONE SUPER-FAST COMPUTER AND MANY "AVERAGE" ONES?

Detailed Questions about Heterogeneity

- What makes one cluster more powerful than another?

- *ARE YOU BETTER OFF . . .*

  — WITH ONE SUPER-FAST COMPUTER AND MANY "AVERAGE" ONES?

  — WITH ALL COMPUTERS "MODERATELY" FAST?

Detailed Questions about Heterogeneity

- What makes one cluster more powerful than another?
- Are you better off with

  — one super-fast computer and many "average" ones

  — or with all computers "moderately" fast?
- *IF YOU COULD "SPEED UP" JUST ONE COMPUTER . . . WHICH ONE WOULD YOU CHOOSE?*

Detailed Questions about Heterogeneity

- What makes one cluster more powerful than another?
- Are you better off with

  — one super-fast computer and many "average" ones

  — or with all computers "moderately" fast?

- *IF YOU COULD "SPEED UP" JUST ONE COMPUTER . . . WHICH ONE WOULD YOU CHOOSE?*

  — THE FASTEST ONE?

Detailed Questions about Heterogeneity

- What makes one cluster more powerful than another?
- Are you better off with

  — one super-fast computer and many "average" ones

  — or with all computers "moderately" fast?

- *IF YOU COULD "SPEED UP" JUST ONE COMPUTER . . .*
  WHICH ONE WOULD YOU CHOOSE?

  — *THE FASTEST ONE?*

  — *THE SLOWEST ONE?*

A Formal Framework for Studying the Questions

Cluster $\mathcal{C}$ has computers $C_1, C_2, \ldots, C_n$

A Formal Framework for Studying the Questions

Cluster $\mathcal{C}$ has computers $C_1, C_2, \ldots, C_n$

C_i completes one unit of work in ρ_i time units.

A Formal Framework for Studying the Questions

Cluster $\mathcal{C}$ has computers $C_1, C_2, \ldots, C_n$

C_i completes one unit of work in ρ_i time units.

$\mathcal{C}$'s *heterogeneity profile:*

$$\mathsf{P}_c \;=\; \langle \rho_1, \rho_2, \ldots, \rho_n \rangle$$

A Formal Framework for Studying the Questions

Cluster $\mathcal{C}$ has computers $C_1, C_2, \ldots, C_n$

C_i completes one unit of work in ρ_i time units.

$\mathcal{C}$'s *heterogeneity profile:* $\mathsf{P}_{\mathcal{C}} = \langle \rho_1, \rho_2, \ldots, \rho_n \rangle$

One finds in

> M. Adler, Y. Gong, A.L. Rosenberg (2008): On "exploiting" node-heterogeneous clusters optimally. *Theory of Computing Systems 42*, 465–487

a solution to the *CLUSTER-EXPLOITATION PROBLEM* ...

— a search for a schedule that maximizes $\mathcal{C}$'s rate of completing work

A Formal Framework for Studying the Questions

Cluster $\mathcal{C}$ has computers $C_1, C_2, \ldots, C_n$

C_i completes one unit of work in ρ_i time units.

$\mathcal{C}$'s *heterogeneity profile:* $\mathsf{P}_{\mathcal{C}} = \langle \rho_1, \rho_2, \ldots, \rho_n \rangle$

One finds in

> M. Adler, Y. Gong, A.L. Rosenberg (2008): On "exploiting" node-heterogeneous clusters optimally. *Theory of Computing Systems 42*, 465–487

a solution to the *CLUSTER-EXPLOITATION PROBLEM*

THE OPTIMAL SCHEDULE FOR $\mathcal{C}$ *DEPENDS ONLY ON* $\mathsf{P}_{\mathcal{C}}$

A Formal Framework for Studying the Questions

Cluster $\mathcal{C}$ has computers $C_1, C_2, \ldots, C_n$

C_i completes one unit of work in ρ_i time units.

$\mathcal{C}$'s *heterogeneity profile:* $\mathsf{P}_{\mathcal{C}} = \langle \rho_1, \rho_2, \ldots, \rho_n \rangle$

One finds in

> M. Adler, Y. Gong, A.L. Rosenberg (2008): On "exploiting" node-heterogeneous clusters optimally. *Theory of Computing Systems 42*, 465–487

a solution the *CLUSTER-EXPLOITATION PROBLEM*

The optimal schedule for $\mathcal{C}$ depends only on $\mathsf{P}_{\mathcal{C}}$

THE WORK COMPLETED UNDER THIS SCHEDULE
IS OUR MEASURE OF
$\mathcal{C}$'s *"POWER"*

A Formal Framework for Studying the Questions

Cluster $\mathcal{C}$ has computers $C_1, C_2, \ldots, C_n$

C_i completes one unit of work in ρ_i time units.

$\mathcal{C}$'s *heterogeneity profile:* $\mathsf{P}_{\mathcal{C}} = \langle \rho_1, \rho_2, \ldots, \rho_n \rangle$

$\mathcal{C}$'s *"power":* the work completed by the optimal solution to the *CLUSTER-EXPLOITATION PROBLEM*

The expression for this work is complicated ...

— so we also measure $\mathcal{C}$'s *"power"* by its

HECR: Homogeneous Equivalent Computing Rate

A Formal Framework for Studying the Questions

Cluster $\mathcal{C}$ has computers $C_1, C_2, \ldots, C_n$

C_i completes one unit of work in ρ_i time units.

$\mathcal{C}$'s *heterogeneity profile:* $\mathsf{P}_c = \langle \rho_1, \rho_2, \ldots, \rho_n \rangle$

$\mathcal{C}$'s *HECR* (*Homogeneous Equivalent Computing Rate*) ...

the computing rate $\rho^{(c)}$ such that

> the *HOMOgeneous* cluster with profile $\langle \rho^{(c)}, \rho^{(c)}, \ldots, \rho^{(c)} \rangle$
>
> completes work at the same rate as $\mathcal{C}$.

ON TO OUR QUESTIONS!

Which ONE Computer Should You Speed UP?

Which Computer to Speed Up: *Additive Speedup*

Speeding up computer C_i *additively* by the amount φ ...

replaces profile

$$\mathsf{P}_c \;=\; \langle \rho_1, \ldots, \rho_{i-1}, \boxed{\rho_i}, \rho_{i+1}, \ldots, \rho_n \rangle$$

by profile

$$\mathsf{P}_c \;=\; \langle \rho_1, \ldots, \rho_{i-1}, \boxed{\rho_i - \varphi}, \rho_{i+1}, \ldots, \rho_n \rangle$$

Say that $0 < \varphi < \min_i\{\rho_i\}$, so every C_i can be sped up.

Which Computer to Speed Up: *Additive Speedup*

Speeding up computer C_i *additively* by the amount φ:

$$\langle \rho_1, \ldots, \rho_{i-1}, \boxed{\rho_i}, \rho_{i+1}, \ldots, \rho_n \rangle \quad \longrightarrow \quad \langle \rho_1, \ldots, \rho_{i-1}, \boxed{\rho_i - \varphi}, \rho_{i+1}, \ldots, \rho_n \rangle$$

Theorem.

Under the additive-speedup scenario, the most advantageous single computer to speed up is $\mathcal{C}$'s fastest *computer.*

Which Computer to Speed Up: *Additive Speedup*

Speeding up computer C_i *additively* by the amount φ:

$$\langle \rho_1, \ldots, \rho_{i-1}, \boxed{\rho_i}, \rho_{i+1}, \ldots, \rho_n \rangle \quad \longrightarrow \quad \langle \rho_1, \ldots, \rho_{i-1}, \boxed{\rho_i - \varphi}, \rho_{i+1}, \ldots, \rho_n \rangle$$

Theorem.

Under the additive-speedup scenario, the most advantageous single computer to speed up is $\mathcal{C}$'s fastest *computer.*

Initial profile: $\langle 1,\ 1/2,\ 1/3,\ 1/4 \rangle$

Speedup amount: $\varphi = 1/16$

i	Speed up computer C_i	Work ratio OLD $\div$ NEW
1	$\langle 15/16,\ 1/2,\ 1/3,\ 1/4 \rangle$	1.008
2	$\langle 1,\ 7/16,\ 1/3,\ 1/4 \rangle$	1.014
3	$\langle 1,\ 1/2,\ 13/48,\ 1/4 \rangle$	1.034
4	$\langle 1,\ 1/2,\ 1/3,\ 3/16 \rangle$	1.159

Which Computer to Speed Up: *Additive Speedup*

Speeding up computer C_i *additively* by the amount φ:

$$\langle \rho_1, \ldots, \rho_{i-1}, \boxed{\rho_i}, \rho_{i+1}, \ldots, \rho_n \rangle \quad \longrightarrow \quad \langle \rho_1, \ldots, \rho_{i-1}, \boxed{\rho_i - \varphi}, \rho_{i+1}, \ldots, \rho_n \rangle$$

Theorem.

Under the additive-speedup scenario, the most advantageous single computer to speed up is $\mathcal{C}$*'s* <u>fastest</u> *computer.*

i	Speed up computer C_i	Work ratio OLD $\div$ NEW
1	$\langle 15/16,\ 1/2,\ 1/3,\ 1/4 \rangle$	1.008
2	$\langle 1,\ 7/16,\ 1/3,\ 1/4 \rangle$	1.014
3	$\langle 1,\ 1/2,\ 13/48,\ 1/4 \rangle$	1.034
4	$\langle 1,\ 1/2,\ 1/3,\ 3/16 \rangle$	1.159

INTUITION: *MORE BANG FOR THE BUCK*

Which Computer to Speed Up: *Multiplicative Speedup*

Speeding up computer C_i *multiplicatively* by factor ψ ...

replaces profile

$$\mathsf{P}_c \;=\; \langle \rho_1, \ldots, \rho_{i-1}, \boxed{\rho_i}, \rho_{i+1}, \ldots, \rho_n \rangle$$

by profile

$$\mathsf{P}_c \;=\; \langle \rho_1, \ldots, \rho_{i-1}, \boxed{\psi\rho_i}, \rho_{i+1}, \ldots, \rho_n \rangle$$

Say that $0 < \psi < 1$, so every C_i can be sped up.

Which Computer to Speed Up: *Multiplicative Speedup*

Speeding up computer C_i *multiplicatively* by factor ψ:

$$\langle \rho_1, \ldots, \rho_{i-1}, \boxed{\rho_i}, \rho_{i+1}, \ldots, \rho_n \rangle \quad \longrightarrow \quad \langle \rho_1, \ldots, \rho_{i-1}, \boxed{\psi\rho_i}, \rho_{i+1}, \ldots, \rho_n \rangle$$

Say that $0 < \psi < 1$, so every C_i can be sped up finitely.

"**Theorem**."

Under the multiplicative-speedup scenario:

The most advantageous single computer to speed up is $\mathcal{C}$'s fastest *computer . . .*

Which Computer to Speed Up: *Multiplicative Speedup*

Speeding up computer C_i *multiplicatively* by factor ψ:

$$\langle \rho_1, \ldots, \rho_{i-1}, \boxed{\rho_i}, \rho_{i+1}, \ldots, \rho_n \rangle \quad \longrightarrow \quad \langle \rho_1, \ldots, \rho_{i-1}, \boxed{\psi\rho_i}, \rho_{i+1}, \ldots, \rho_n \rangle$$

Say that $0 < \psi < 1$, so every C_i can be sped up finitely.

"**Theorem**."

Under the multiplicative-speedup scenario:

The most advantageous single computer to speed up is $\mathcal{C}$'s fastest *computer …*

— *UNLESS*

Which Computer to Speed Up: *Multiplicative Speedup*

Speeding up computer C_i *multiplicatively* by factor ψ:

$$\langle \rho_1, \ldots, \rho_{i-1}, \boxed{\rho_i}, \rho_{i+1}, \ldots, \rho_n \rangle \quad \longrightarrow \quad \langle \rho_1, \ldots, \rho_{i-1}, \boxed{\psi\rho_i}, \rho_{i+1}, \ldots, \rho_n \rangle$$

Say that $0 < \psi < 1$, so every C_i can be sped up finitely.

"**<u>Theorem</u>**."

Under the multiplicative-speedup scenario:

The most advantageous single computer to speed up is $\mathcal{C}$'s <u>fastest</u> *computer ...*

— *UNLESS* <u>either</u> *this computer is already "very fast"*

<u>or</u> *the speedup factor ψ is "very small."*

Which Computer to Speed Up: *Multiplicative Speedup*

At least one computer is not "very fast":

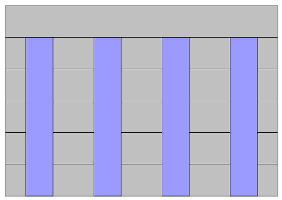

- A 4-computer cluster

  — HOMOgeneous (before any speedups)

- Bar height is ρ-value ...

  — a lower bar is a faster computer

Which Computer to Speed Up: *Multiplicative Speedup*

At least one computer is not "very fast":

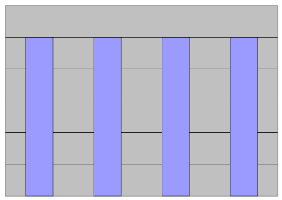

- A 4-computer cluster

  — HOMOgeneous (before any speedups)

- Bar height is ρ-value ...

  — a lower bar is a faster computer

START SPEEDING UP ONE COMPUTER OPTIMALLY ...

— *BY THE FACTOR* $\psi = 1/2$

Which Computer to Speed Up: *Multiplicative Speedup*

At least one computer is not “very fast”:

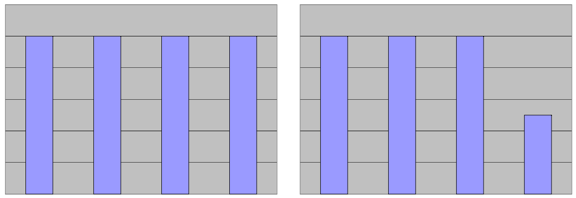

Which Computer to Speed Up: *Multiplicative Speedup*

At least one computer is not "very fast":

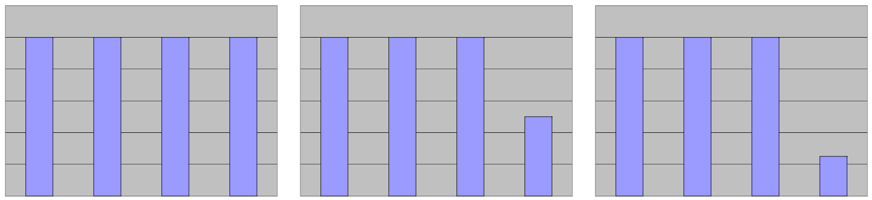

Which Computer to Speed Up: *Multiplicative Speedup*

At least one computer is not "very fast":

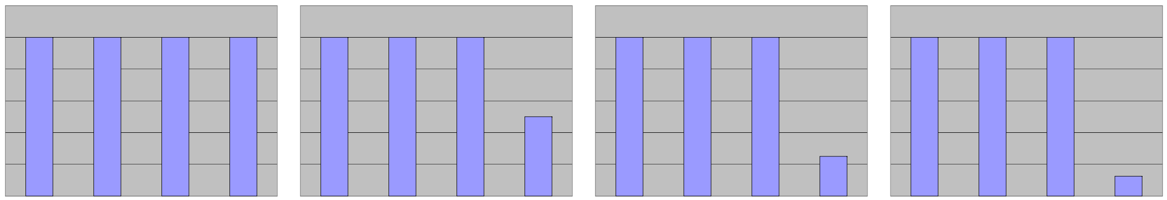

Which Computer to Speed Up: *Multiplicative Speedup*

At least one computer is not "very fast":

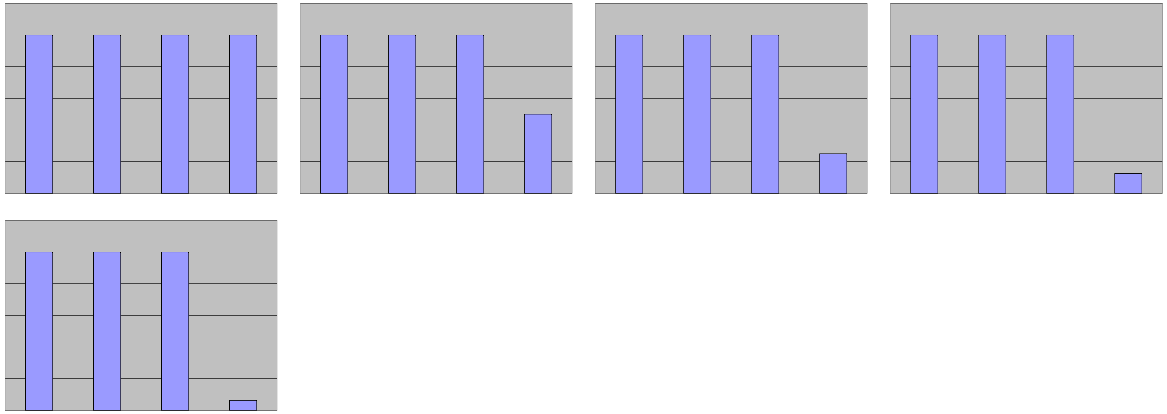

Which Computer to Speed Up: *Multiplicative Speedup*

At least one computer is not "very fast":

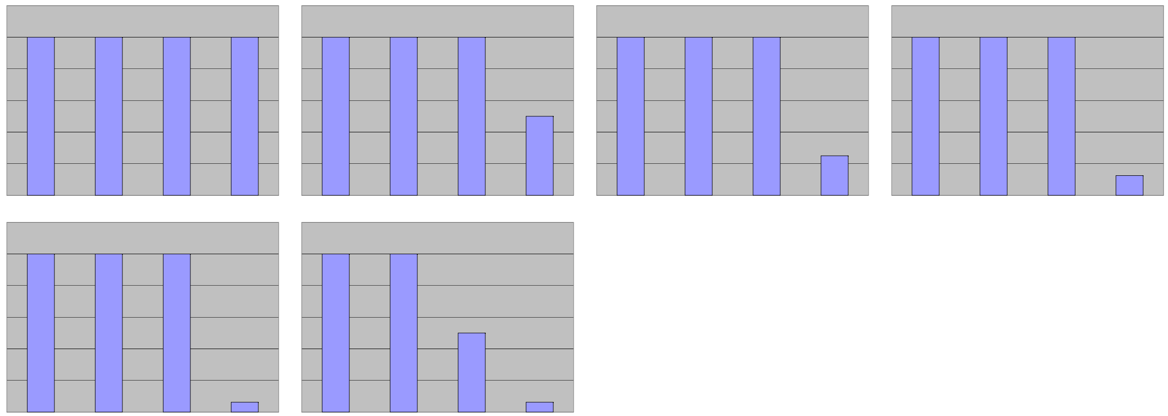

Which Computer to Speed Up: *Multiplicative Speedup*

At least one computer is not “very fast”:

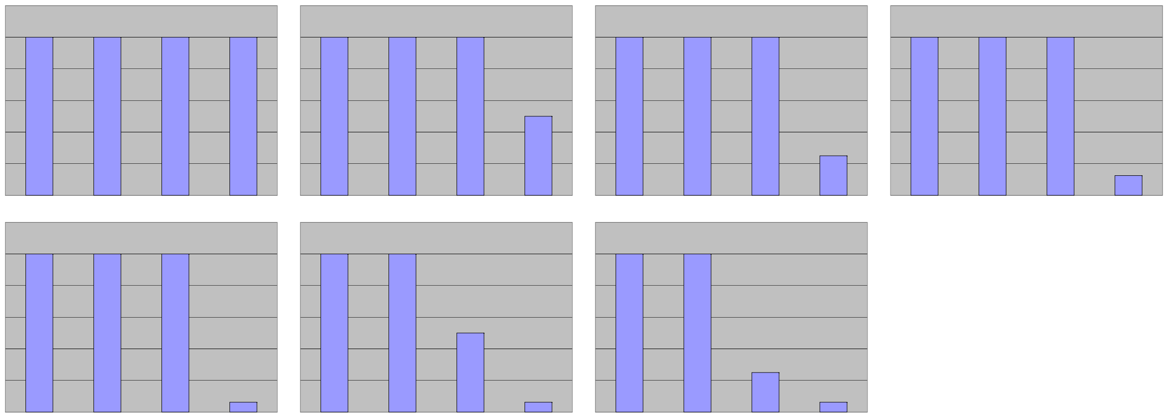

Which Computer to Speed Up: *Multiplicative Speedup*

At least one computer is not "very fast":

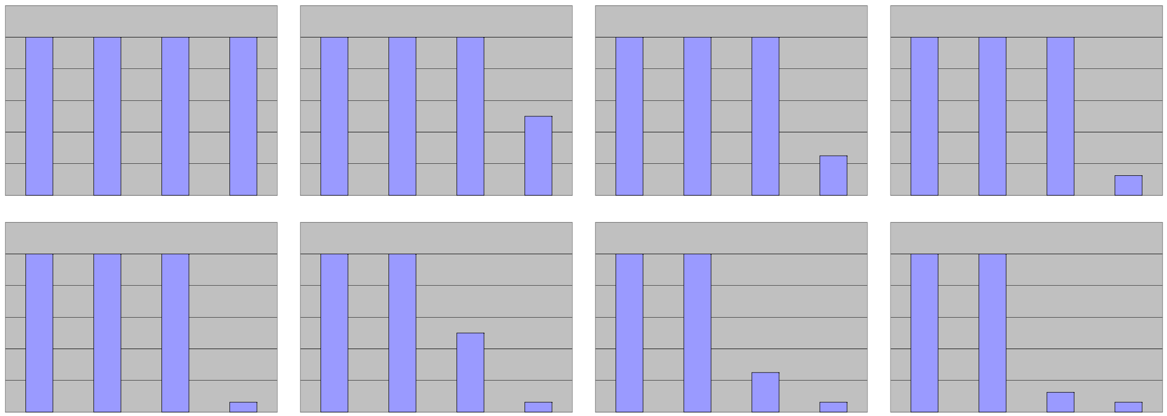

Which Computer to Speed Up: *Multiplicative Speedup*

At least one computer is not "very fast":

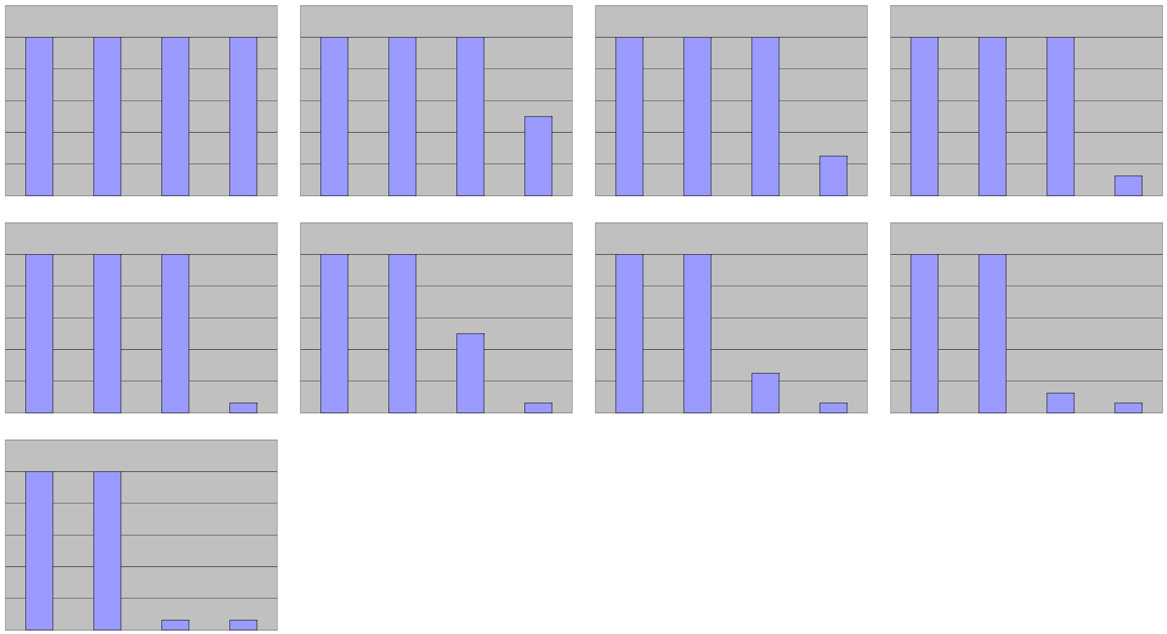

Which Computer to Speed Up: *Multiplicative Speedup*

At least one computer is not "very fast":

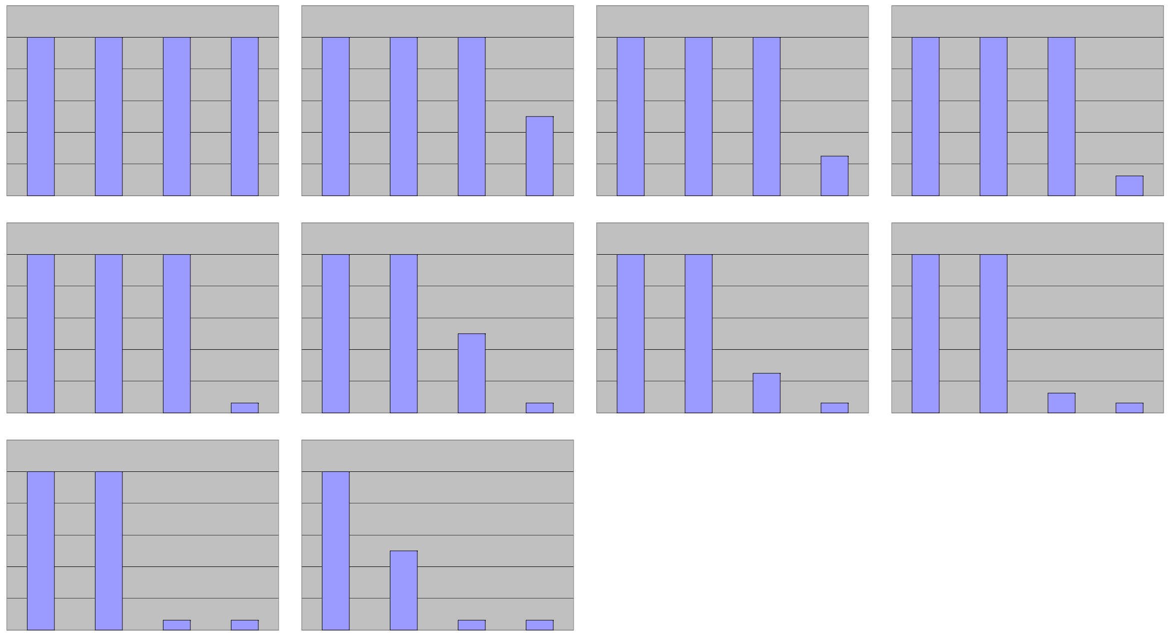

Which Computer to Speed Up: *Multiplicative Speedup*

At least one computer is not “very fast”:

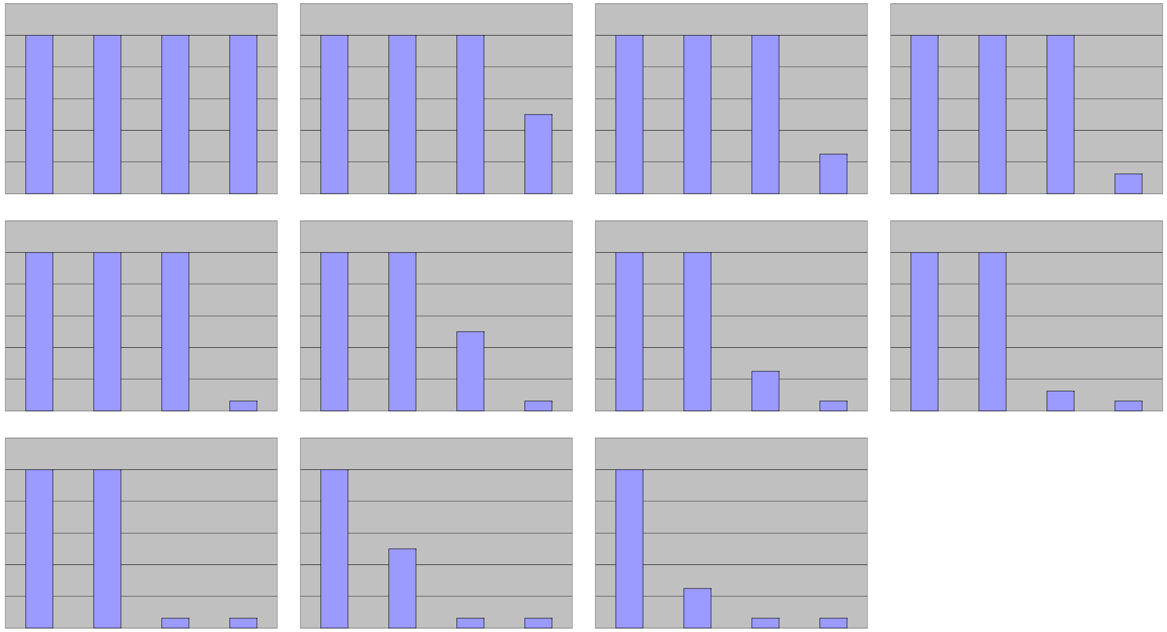

Which Computer to Speed Up: *Multiplicative Speedup*

At least one computer is not "very fast":

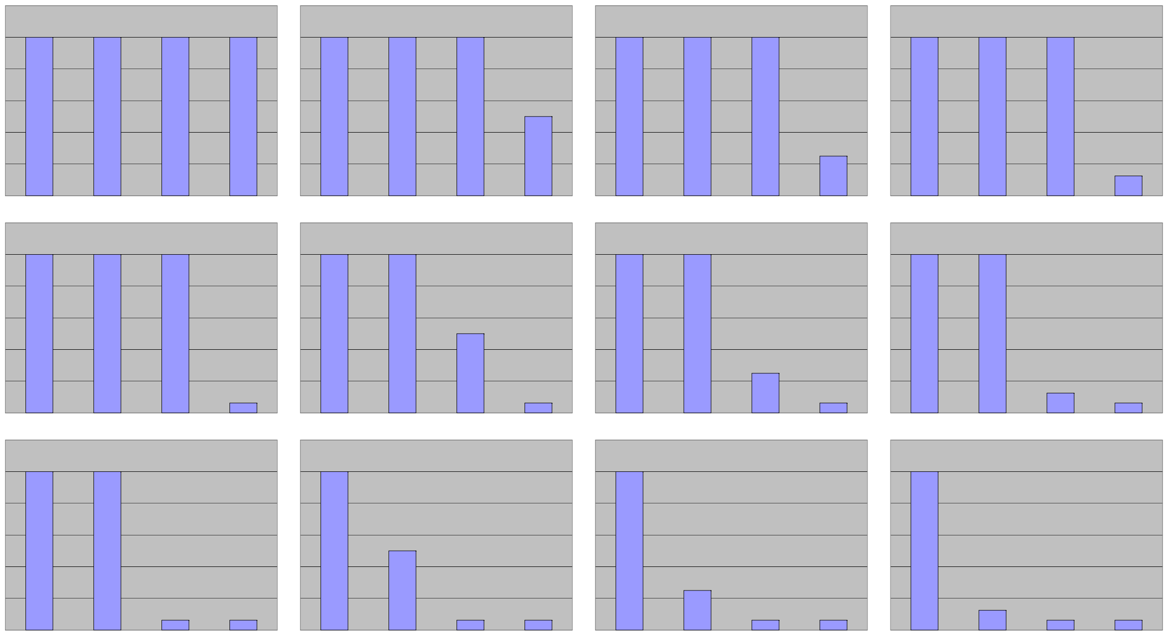

Which Computer to Speed Up: *Multiplicative Speedup*

At least one computer is not "very fast":

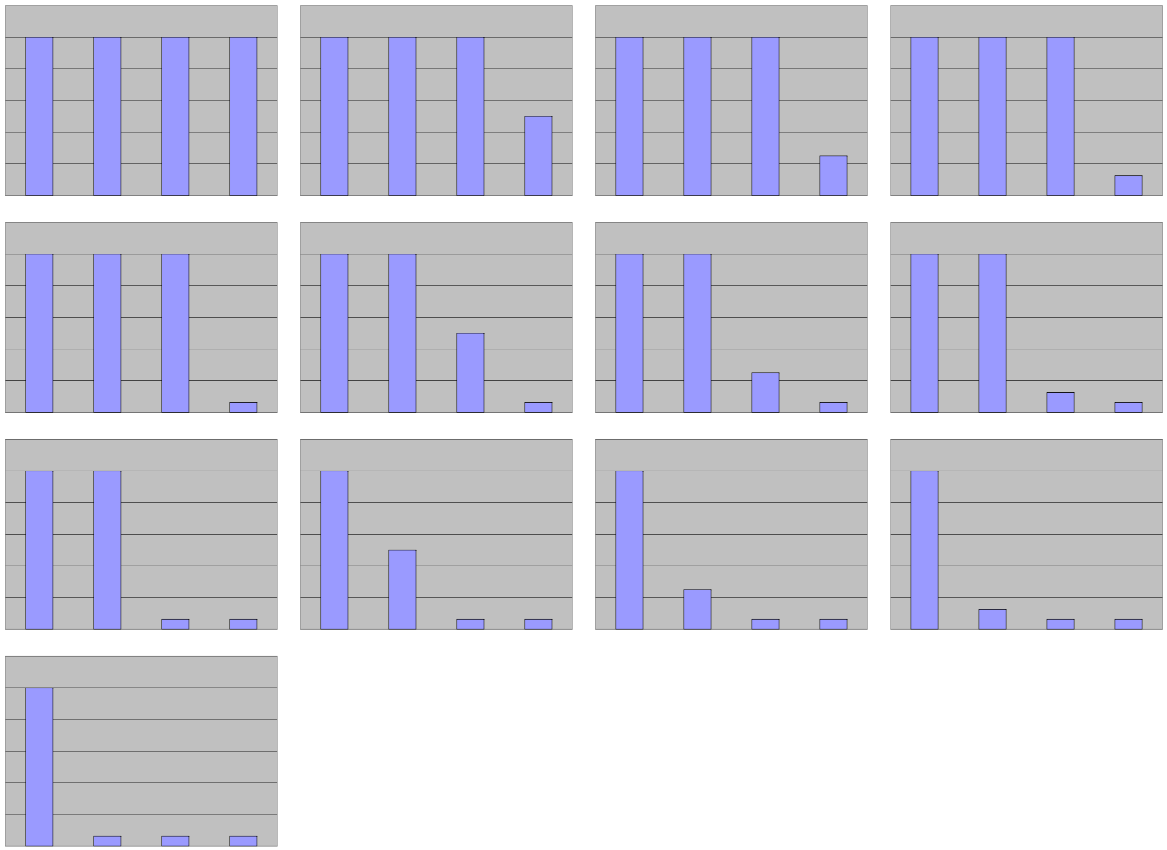

Which Computer to Speed Up: *Multiplicative Speedup*

At least one computer is not "very fast":

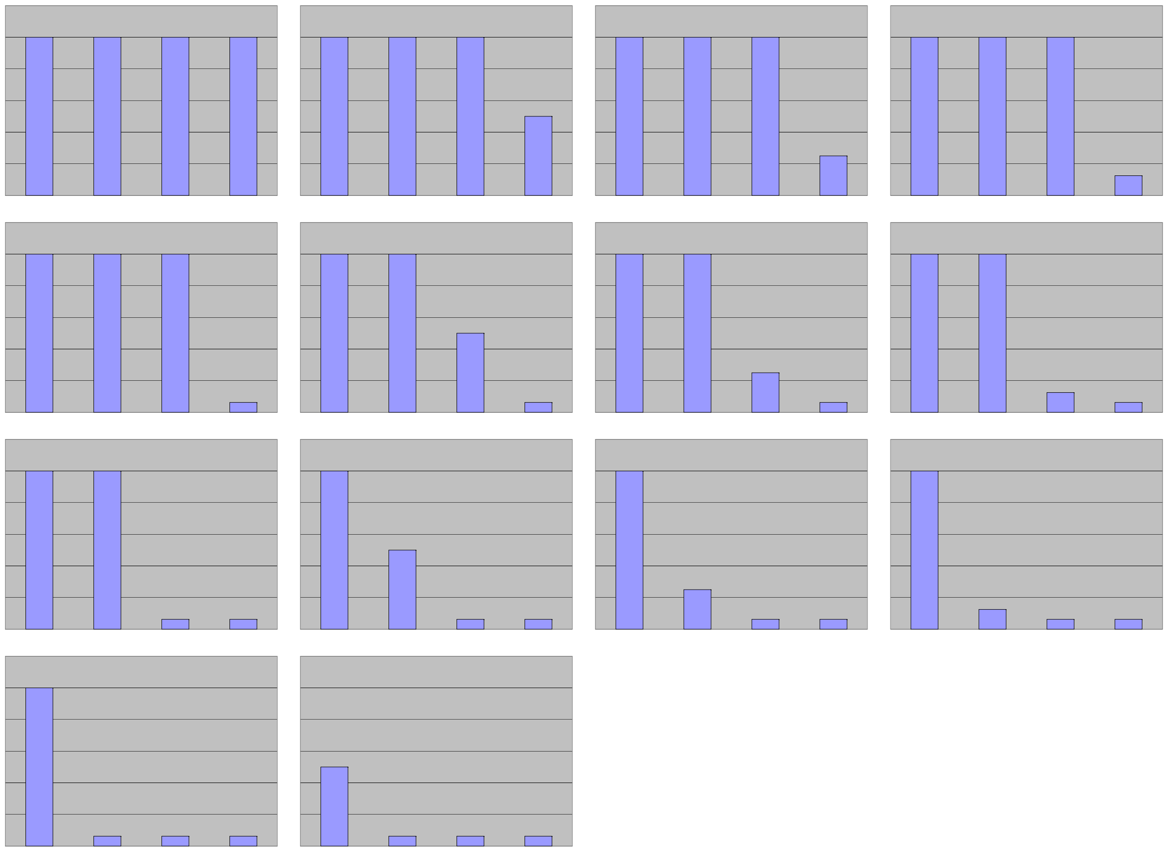

Which Computer to Speed Up: *Multiplicative Speedup*

At least one computer is not "very fast":

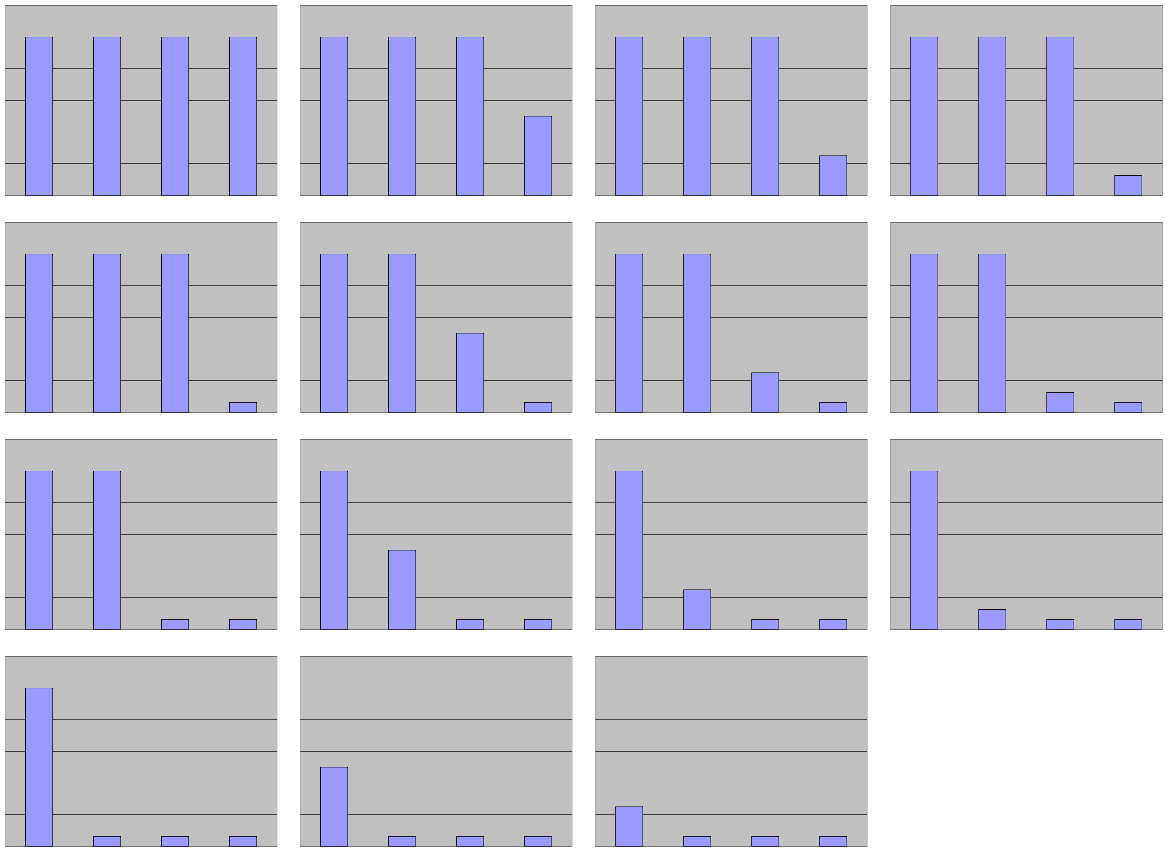

Which Computer to Speed Up: *Multiplicative Speedup*

At least one computer is not "very fast":

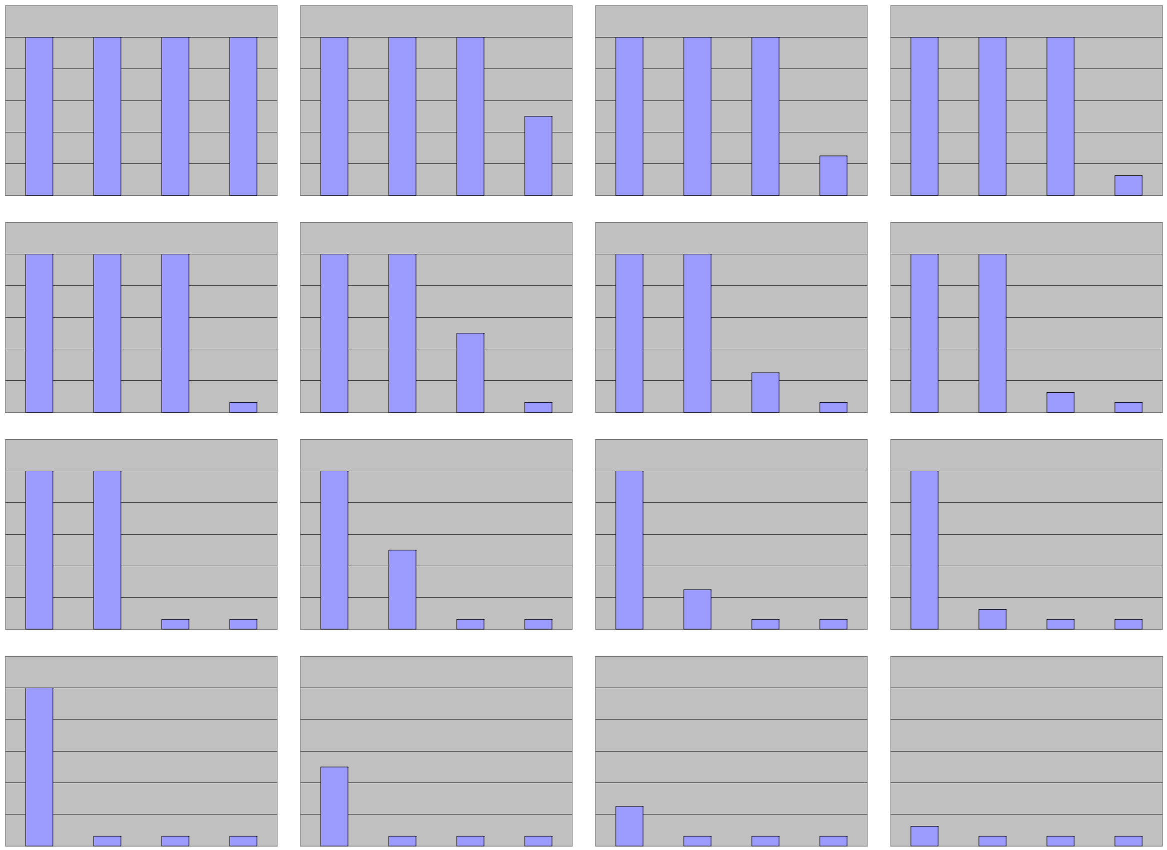

Which Computer to Speed Up: *Multiplicative Speedup*

At least one computer is not "very fast":

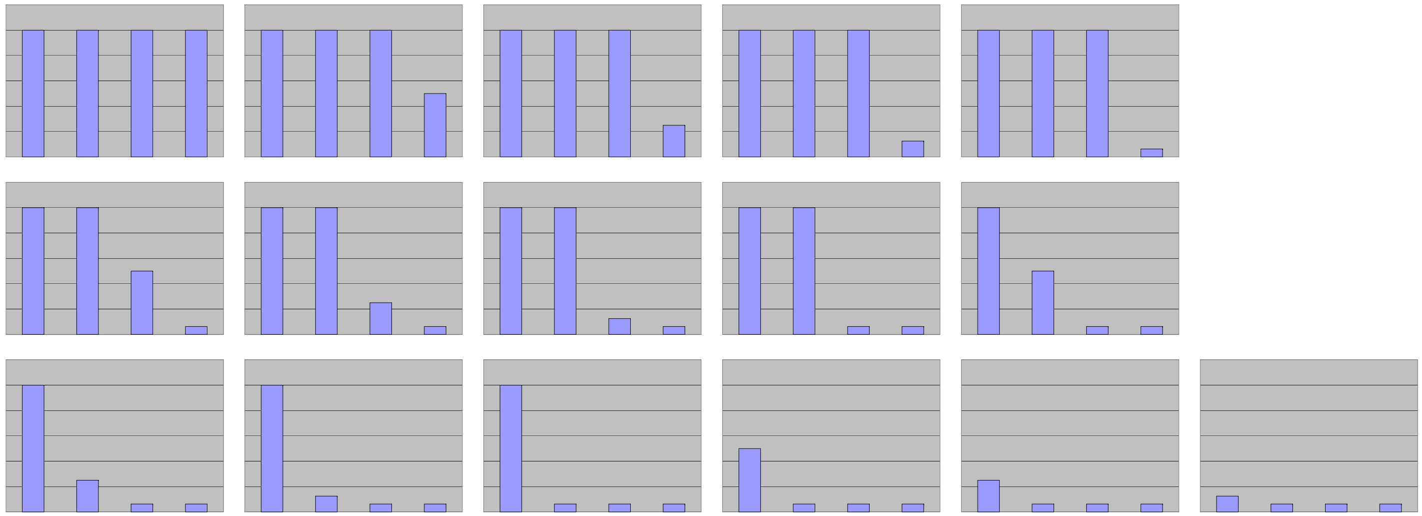

When all computers are very fast:

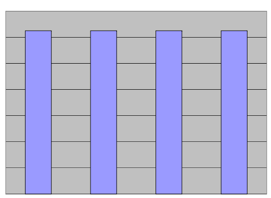

Which Computer to Speed Up: *Multiplicative Speedup*

At least one computer is not "very fast":

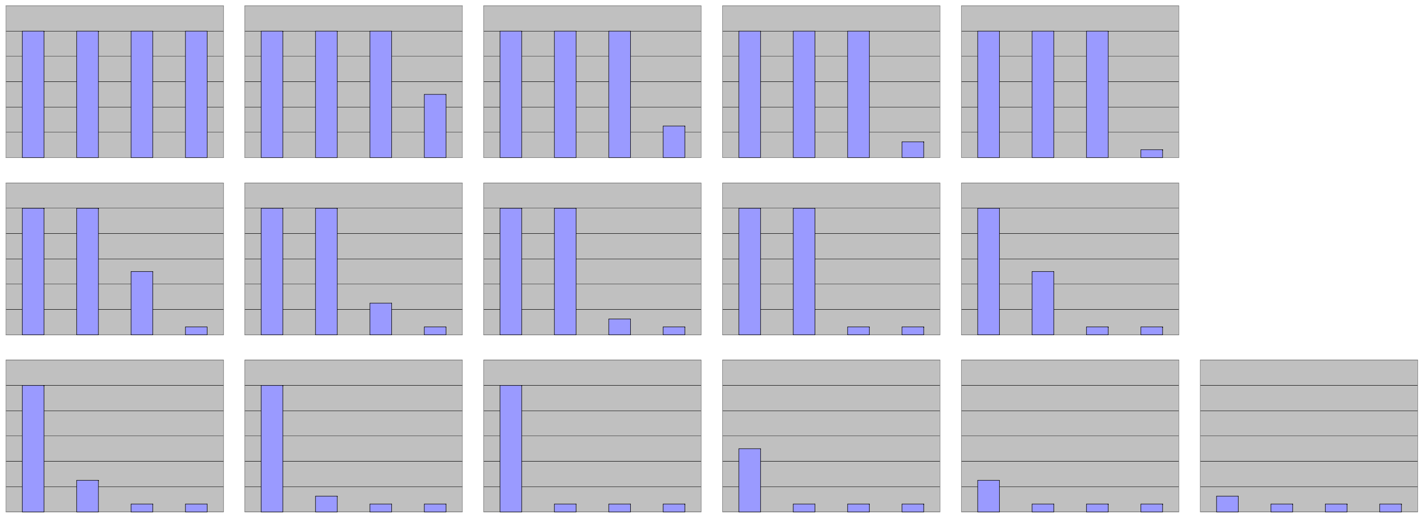

When all computers are very fast:

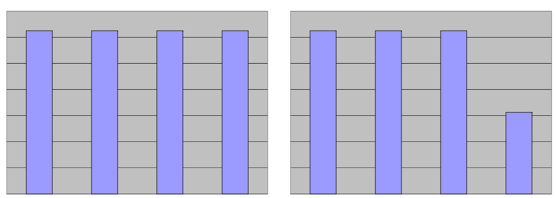

Which Computer to Speed Up: *Multiplicative Speedup*

At least one computer is not “very fast”:

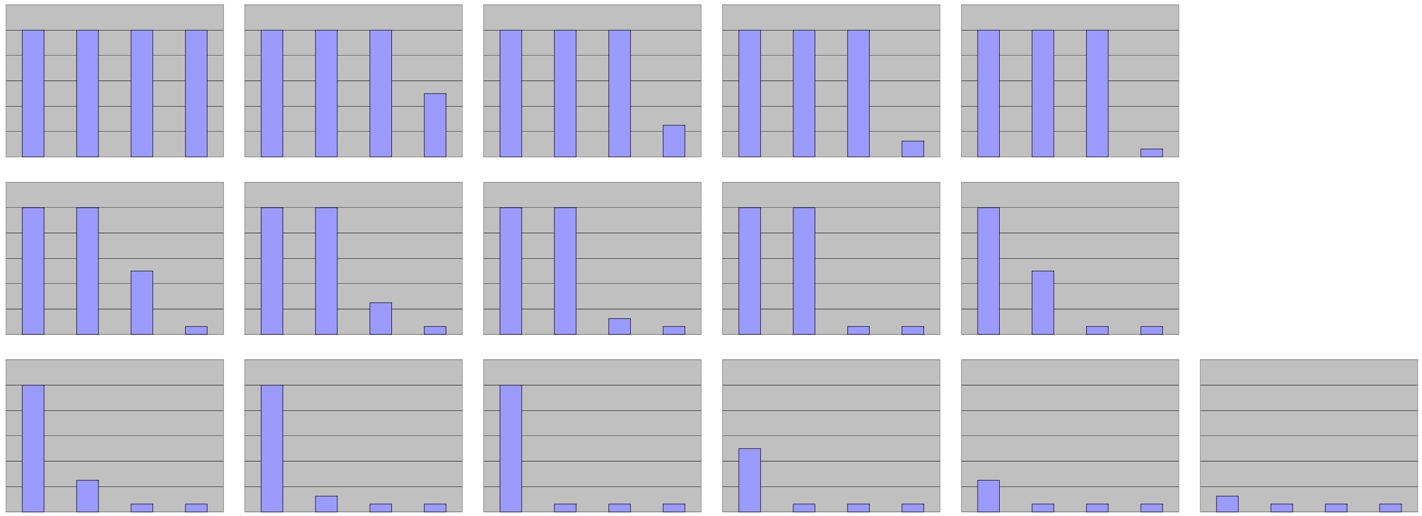

When all computers are very fast:

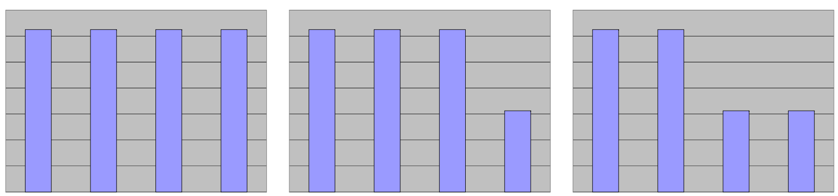

Which Computer to Speed Up: *Multiplicative Speedup*

At least one computer is not "very fast":

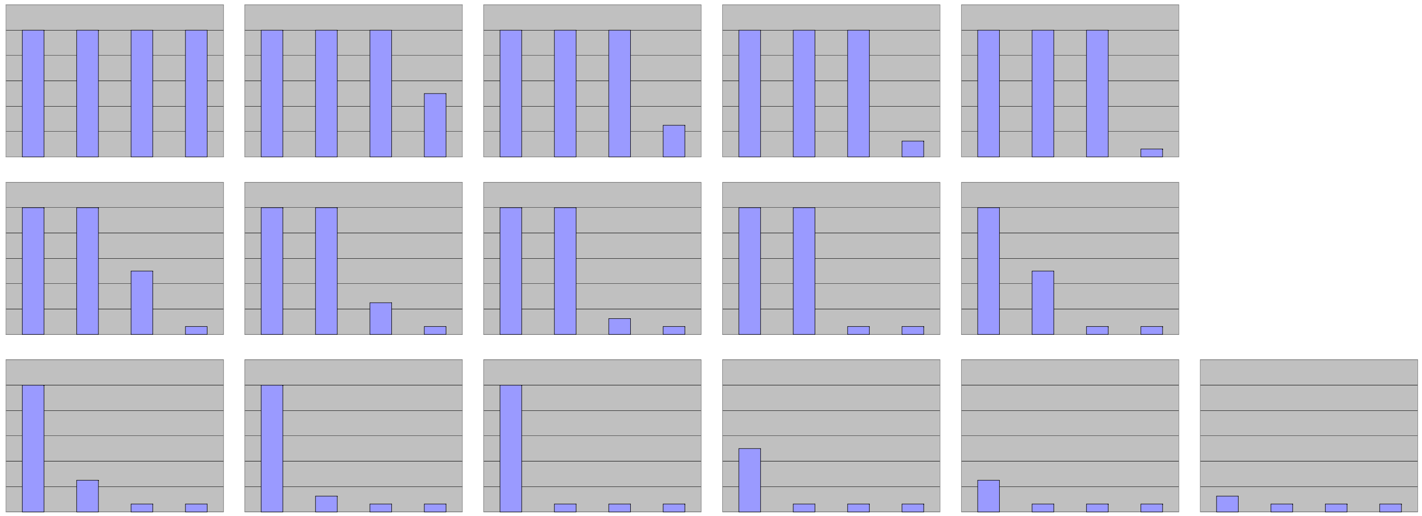

When all computers are very fast:

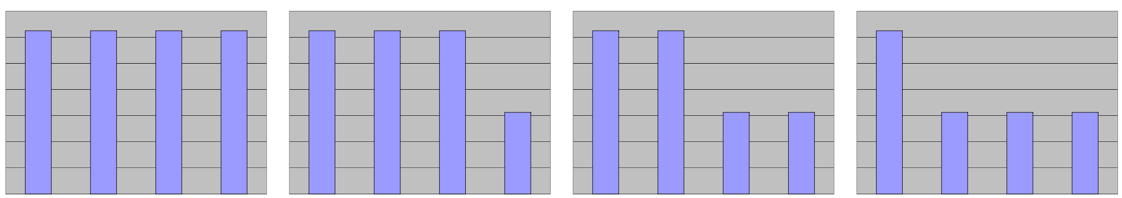

Which Computer to Speed Up: *Multiplicative Speedup*

At least one computer is not "very fast":

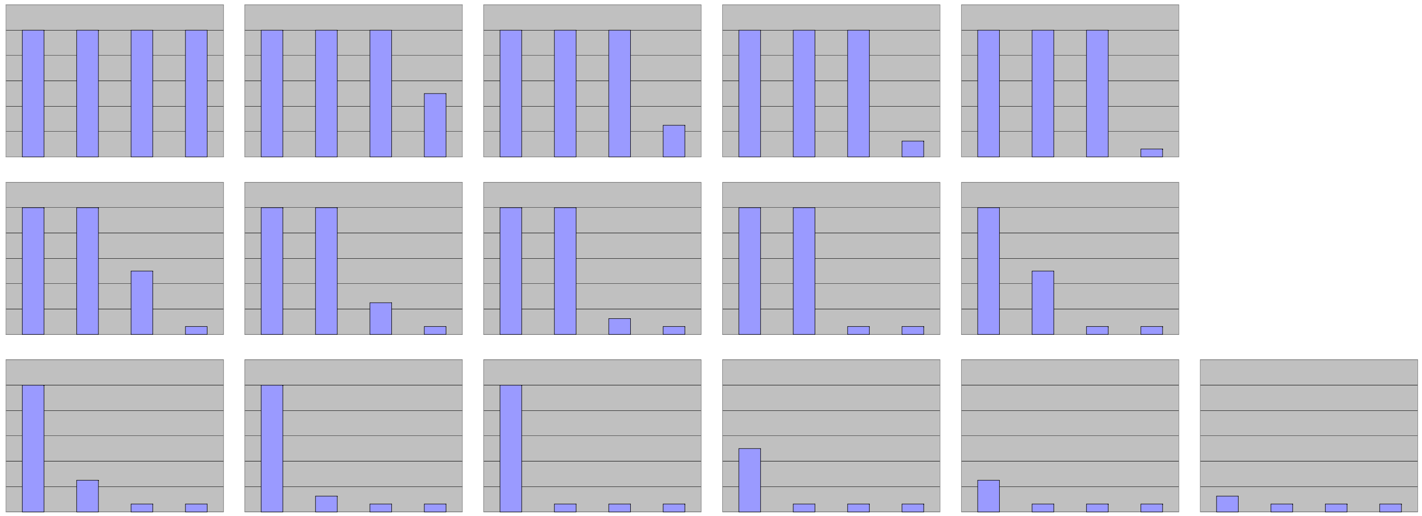

When all computers are very fast:

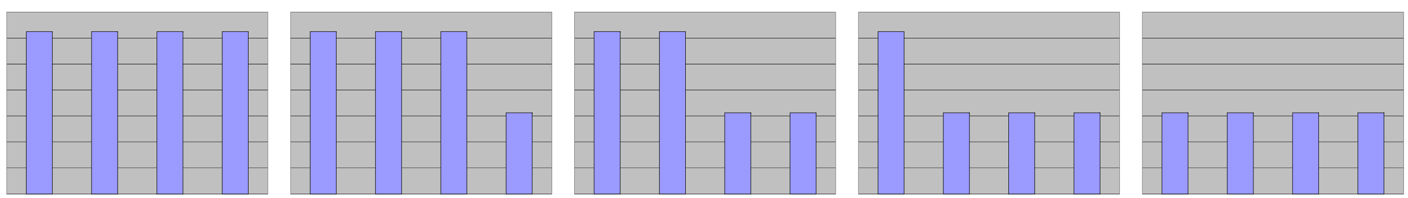

What Makes Clusters Powerful?

Absolute and Relative Answers

What Makes Clusters Powerful: *Variance in Computer Speeds*

Say that cluster $\mathcal{C}_1$, with profile P_1, and cluster $\mathcal{C}_2$, with profile P_2, share the same mean speed.

Theorem.

Say that $\mathcal{C}_1$ and $\mathcal{C}_2$ each has 2 computers.

Then

$\mathcal{C}_1$ outperforms $\mathcal{C}_2$

if and only if

$$VAR(\mathsf{P}_1) > VAR(\mathsf{P}_2).$$

What Makes Clusters Powerful: *Variance in Computer Speeds*

Say that cluster $\mathcal{C}_1$, with profile P_1, and cluster $\mathcal{C}_2$, with profile P_2, share the same mean speed.

Say that $\mathcal{C}_1$ and $\mathcal{C}_2$ each has 2 computers. Then $\mathcal{C}_1$ outperforms $\mathcal{C}_2$ if and only if $VAR(\mathsf{P}_1) > VAR(\mathsf{P}_2)$.

Corollary.

HETEROGENEITY CAN ACTUALLY LEND POWER TO A CLUSTER ...

if 2-*computer clusters $\mathcal{C}_1$ and $\mathcal{C}_2$ share the same mean speed*
and $\mathcal{C}_1$ *is* heterogeneous, *while* $\mathcal{C}_2$ *is* homogeneous
then $\mathcal{C}_1$ *outperforms* $\mathcal{C}_2$.

What Makes Clusters Powerful: *Variance in Computer Speeds*

Say that cluster $\mathcal{C}_1$, with profile P_1, and cluster $\mathcal{C}_2$, with profile P_2, share the same mean speed.

Say that $\mathcal{C}_1$ and $\mathcal{C}_2$ each has 2 computers. Then $\mathcal{C}_1$ outperforms $\mathcal{C}_2$ if and only if $VAR(\mathsf{P}_1) > VAR(\mathsf{P}_2)$.

Unfortunately:

THIS RESULT DOES NOT EXTEND TO 3-COMPUTER CLUSTERS

What Makes Clusters Powerful: *Variance in Computer Speeds*

Say that cluster $\mathcal{C}_1$, with profile P_1, and cluster $\mathcal{C}_2$, with profile P_2, share the same mean speed.

Say that $\mathcal{C}_1$ and $\mathcal{C}_2$ each has 2 computers. Then $\mathcal{C}_1$ outperforms $\mathcal{C}_2$ if and only if $VAR(\mathsf{P}_1) > VAR(\mathsf{P}_2)$.

Unfortunately:

This result does not extend to 3-computer clusters

BUT . . .

What Makes Clusters Powerful: *Variance in Computer Speeds*

Say that cluster $\mathcal{C}_1$, with profile P_1, and cluster $\mathcal{C}_2$, with profile P_2, share the same mean speed.

Theorem.

Say that $\mathcal{C}_1$ and $\mathcal{C}_2$ each has 3 computers.

There exists a threshold $\theta > 0$ *such that:*

if $VAR(\mathsf{P}_1) \geq VAR(\mathsf{P}_2) + \theta$
then $\mathcal{C}_1$ *outperforms* $\mathcal{C}_2$.

What Makes Clusters Powerful: *Variance in Computer Speeds*

Say that cluster $\mathcal{C}_1$, with profile P_1, and cluster $\mathcal{C}_2$, with profile P_2, share the same mean speed.

Theorem.

Say that $\mathcal{C}_1$ and $\mathcal{C}_2$ each has 3 computers.

There exists a threshold $\theta > 0$ *such that:*

if $VAR(\mathsf{P}_1) \geq VAR(\mathsf{P}_2) + \theta$
then $\mathcal{C}_1$ *outperforms* $\mathcal{C}_2$.

This result seems (based on simulations) to extend to big clusters.