On the Price of Heterogeneity in Parallel Systems

P. Brighten Godfrey and Richard M. Karp

SPAA'06 - July 31, 2006

Introduction

- Consider a parallel system in which each node has a *capacity*

processing speed, Bottleneck bandwidth to Internet, memory, ...

- Does increasing heterogeneity of the capacity distribution help or hurt?

Introduction

System A

All nodes have equal capacity

System B

Same *total* capacity; higher variance

Does A or B perform better?

Yes

...can do either, depending on what system we're talking about AND the conditions under which we're running the system

Example

Minimum Makespan Scheduling

- Set of jobs, each with a *length*
- *n* processors, each with *speed* c_i
- Assign jobs to processors to minimize *makespan*: time until last processor completes its jobs
- Completion time of processor *i*: sum of job lengths given to it, divided by c_i

Example 1

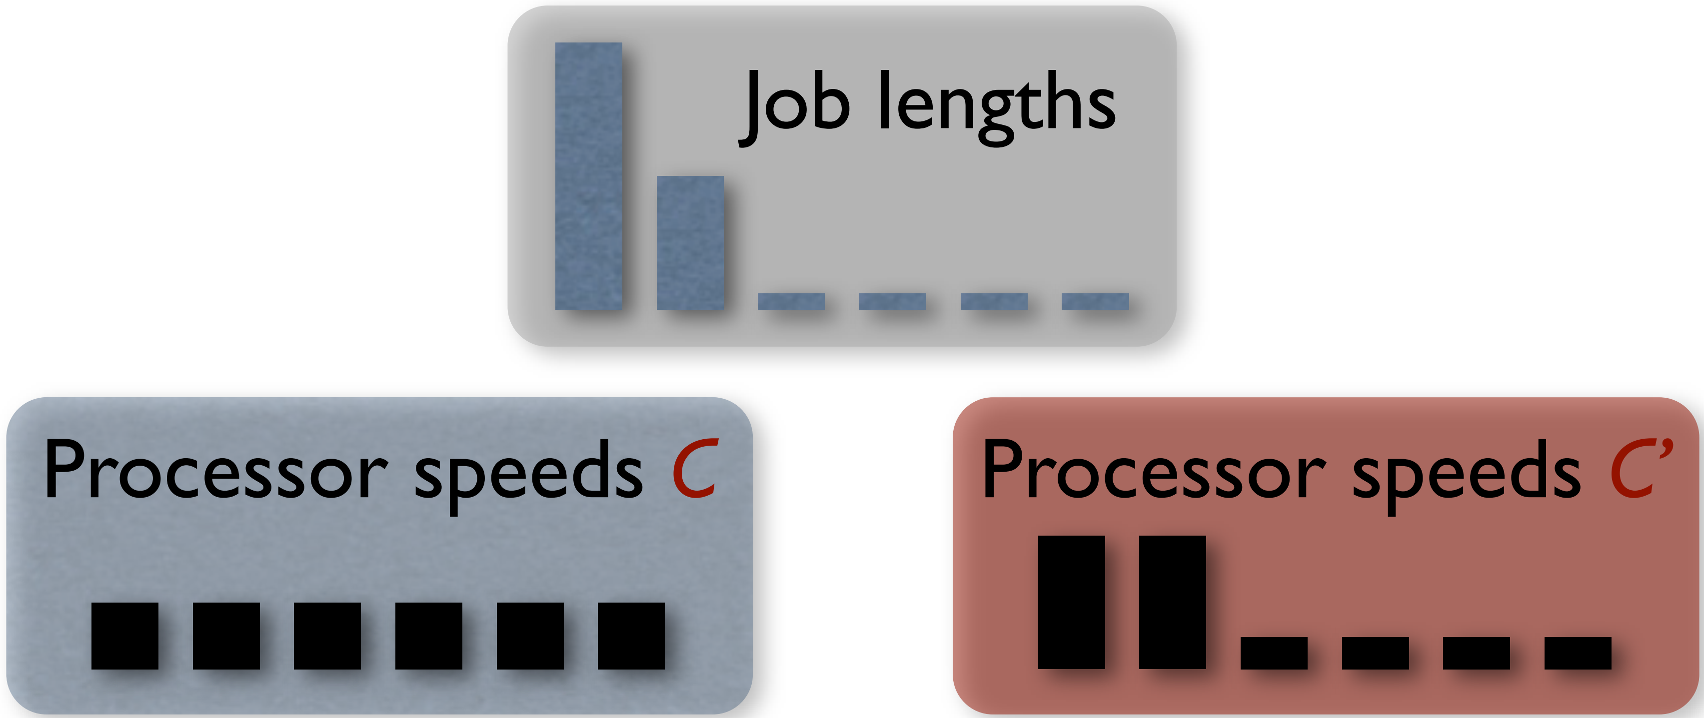

Completion time $4\ \mathrm{sec}$

Completion time $2\ \mathrm{sec}$

Example 2

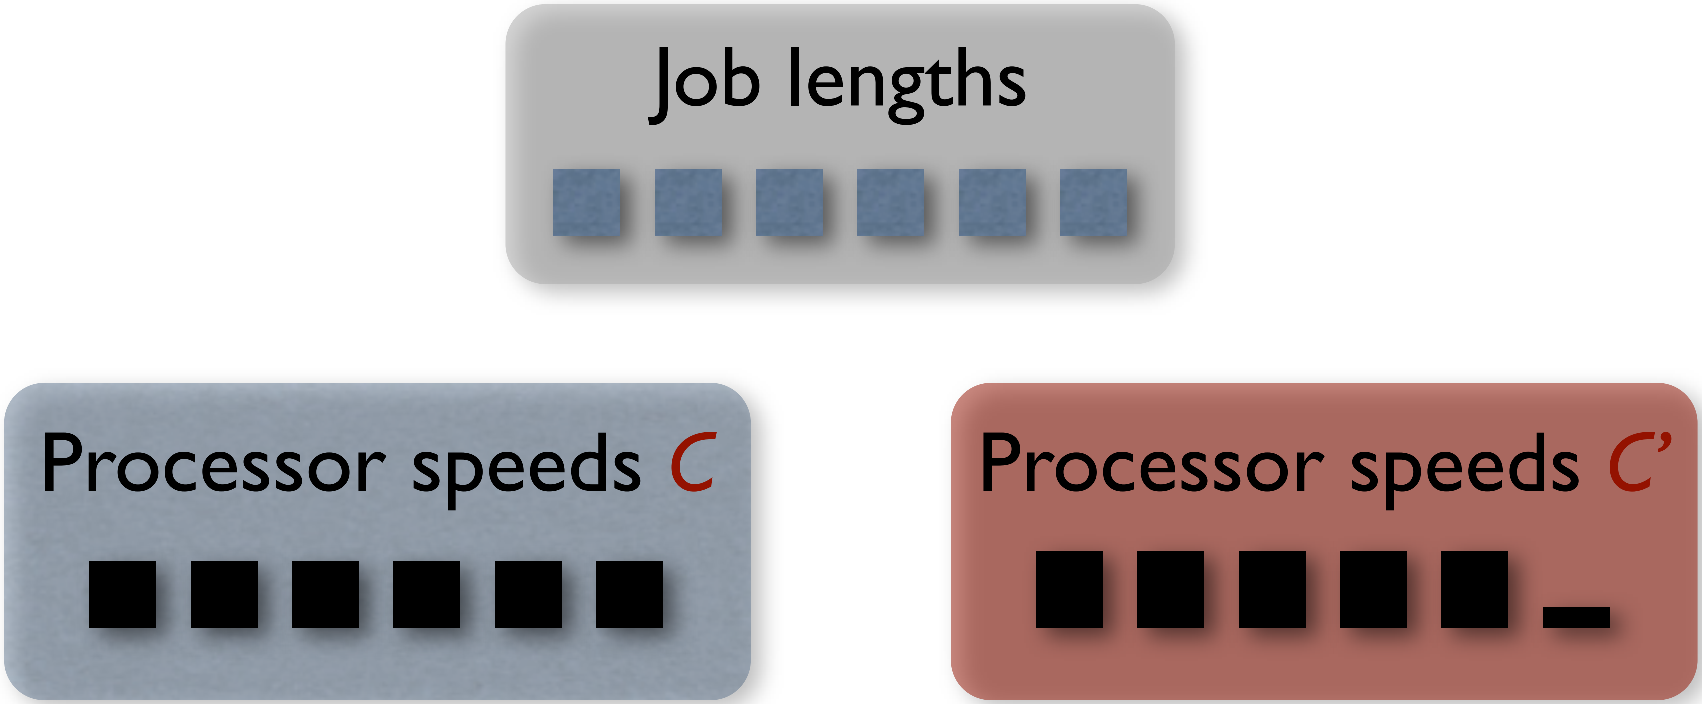

Completion time	Completion time
1 sec	≈ 2 sec

So increasing heterogeneity
can help or hurt.

Can we make any
generalizations?

In This Paper

a general framework
to quantify the worst-case effect
of increasing heterogeneity

Contents

- **Model**
- Results
- Conclusion

Model

Price of Heterogeneity of g

g usually optimal value to some combinatorial optimization problem.

$$\sup_{W,C,C' : C \preceq C'} \frac{g(C', W)}{g(C, W)}$$

cost function
(processing time in optimal schedule)

node capacities
(CPU speed)

workload
(jobs to run)

Defining Heterogeneity

- Capacity vectors (sorted decreasing)

$$C = (c_1, \ldots, c_n)$$
$$C' = (c'_1, \ldots, c'_n)$$

- $C' \succeq C$ when

$$\forall k \; \sum_{i=1}^{k} c'_i \geq \sum_{i=1}^{k} c_i \quad \text{and} \quad \sum_{i=1}^{n} c'_i = \sum_{i=1}^{n} c_i$$

- Majorization partial order

$$C' \succeq C \implies \begin{cases} \mathrm{var}(C') \geq \mathrm{var}(C) \\ -H(C') \geq -H(C) \end{cases}$$

So majorization is consistent with both variance and negative entropy.

Majorization example

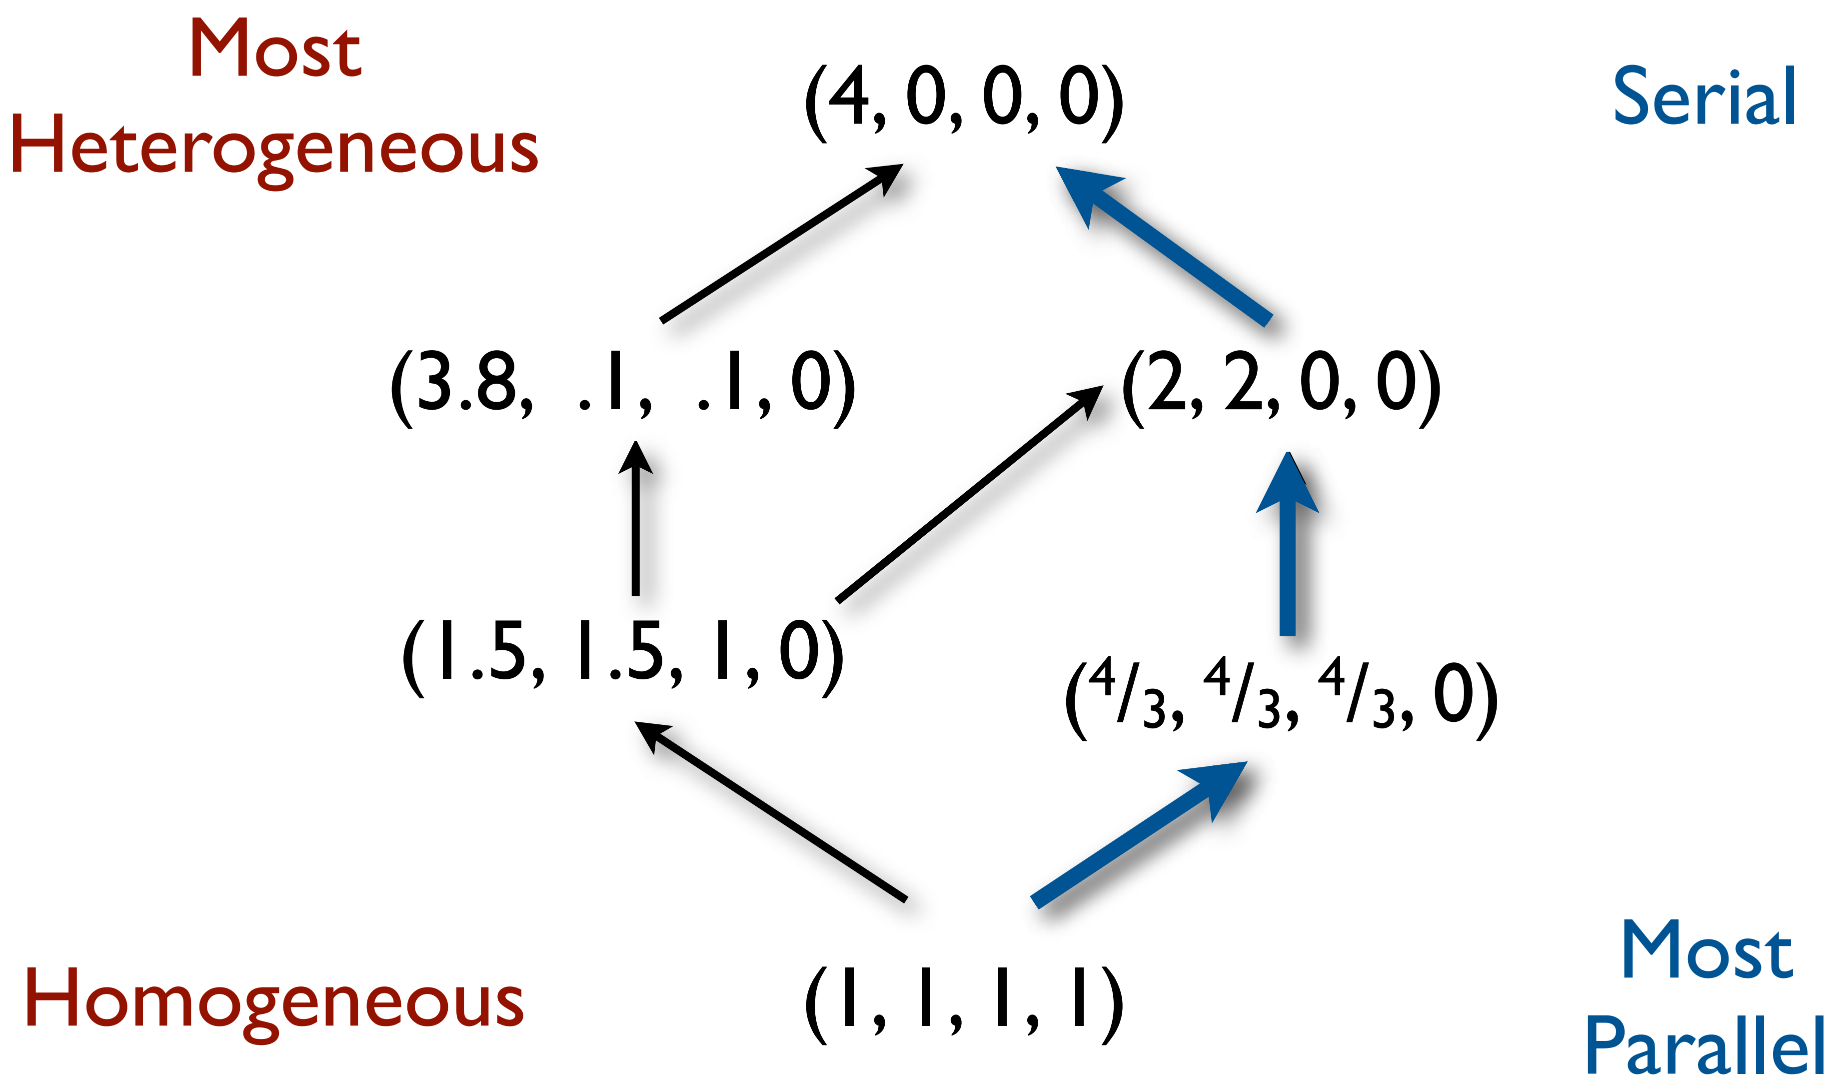

So Price of Heterogeneity

also bounds the

the Value of Parallelism!

Using the Price of Heterogeneity

- Justified generalizations

  (Constant vs. unbounded PoH)

- Comparison across systems

- Worst cases for testing

What characteristics place a cost function in one or the other category?

Contents

- Model
- **Results**
- Conclusion

Results

Problem	PoH
Minimum makespan scheduling	2-1/n
Scheduling on related machines	$O(1)$
PCS, unit length jobs	≤ 16
Precedence Constrained Sched.	$O(\log n)$
Sched. with release times	*Unbounded*
Minimum network diameter	≤ 2

One way to bound PoH

- Goal: show *C'* is as good as more homogeneous capacities *C*

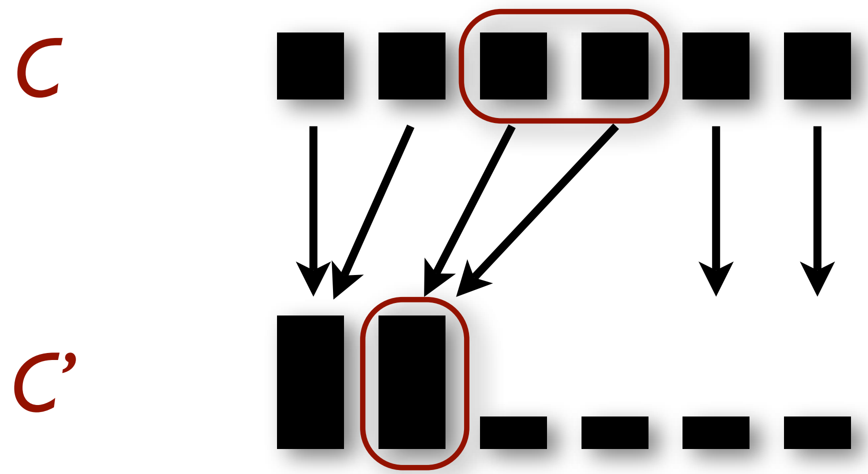

- Total capacity "simulated" by each *C'* node must be not much more than its own capacity

Simulation Lemma

- A *ß*-simulation is a mapping from *C* to *C'* such that no *C'*-node gets more than *ß* times its capacity.
- **Lemma:** a (2-1/*n*)-simulation always exists for any *C* and more heterogeneous *C'*

Lay of the land

- Minimum Makespan Scheduling & a class of generalizations: $O(1)$ PoH

- Precedence Constrained Scheduling (PCS): $O(\log n)$ PoH

- Scheduling with release times: *unbounded* PoH

PCS

- Like Min. Makespan Scheduling, except...
- Given set of precedence constraints: "Job *i* must finish before job *k* starts"

PCS

- Simulation technique cannot succeed

C-machines Time

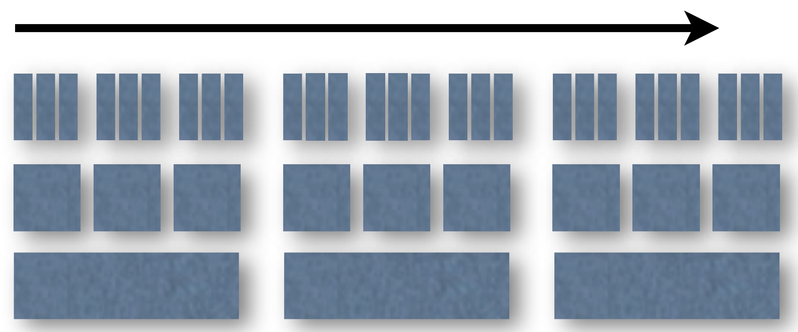

1
2
3
... ...

- Design capacity distributions such that some *C'*-machines simulate multiple *C*-machines

- Factor *n*/4 increase in schedule length!

PCS

- Instead, use LP relaxation of PCS due to Chudak and Shmoys
- Can apply Simulation Lemma to optimal values of the LP
  - Key relaxed constraint: machine can only execute one job at a time
- LP is within $O(\log n)$ of optimal => PoH of PCS is $O(\log n)$

Contents

- Model
- Results
- **Conclusion**

Conclusion

- Introduced a framework to characterize worst-case effect of increasing heterogeneity
- "Batch" scheduling problems have low PoH
  - Even PCS has $O(\log n)$ PoH, while release times cause unbounded PoH
- Does PCS have $O(1)$ PoH?