Energy Corollaries to Amdahl's Law

Sangyeun Cho, Rami Melhem, Michael Moeng
University of Pittsburgh

Presented By: Michael Moeng

Outline

- Amdahl's Law
- Power Applications
- Problem Formulation
- Derivations
- Conclusions

Amdahl's Law

- Maximum speedup for a program given
  - Sequential fraction s, parallel fraction $p=1-s$
  - N processors
- Speedup $= \dfrac{1}{s + p/N}$
- Assumes parallel portion can be perfectly parallelized

Power Applications

- Keep execution time constant and improve dynamic energy
- Assumes $p \sim f^{\alpha}$
- Find optimal serial and parallel frequency by taking derivative

Power Applications

- Keep execution time constant and improve energy

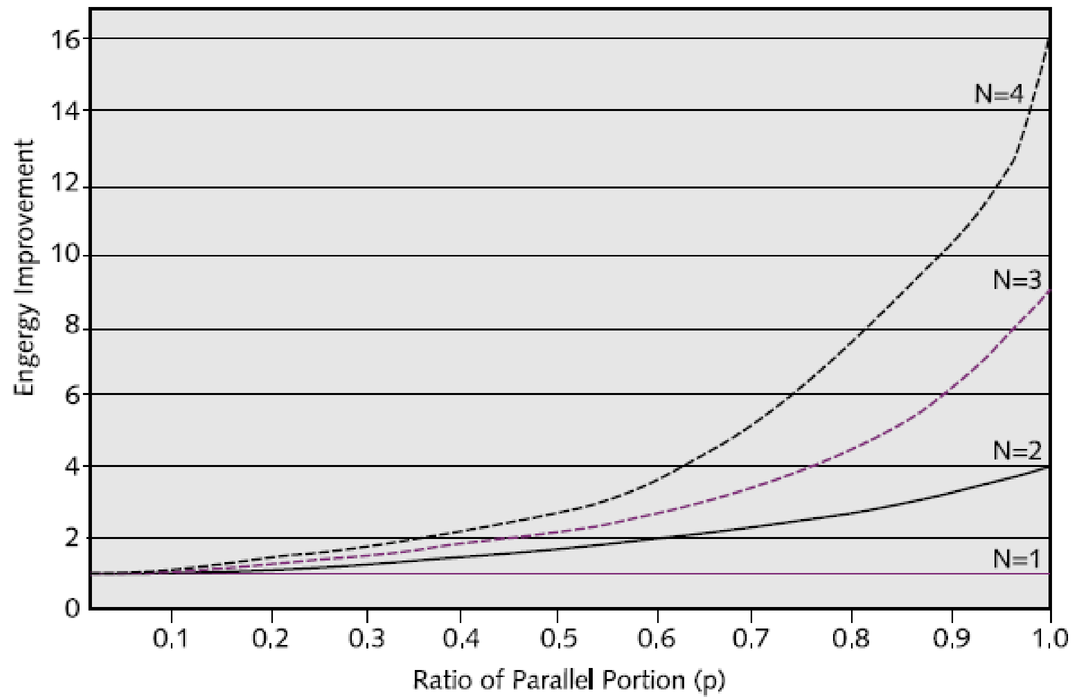

Figure 1. Achievable dynamic energy improvement assuming $\alpha = 3$ and using 1, 2, 3, and 4 processors, given the ratio of serial and parallel work in a program.

Generalized Power

- Now we set speedup to $x \leq \dfrac{1}{s + p/N}$
- Set f_s, f_p
- What are optimal serial and parallel frequencies?

Generalized Power

- Now we set speedup to $x \leq \dfrac{1}{s + p/N}$
- Set f_s, f_p
- What are optimal serial and parallel frequencies?

- Trivial – minimize frequencies
- Must consider static power!

Problem Formulation

- Serial work *s* completes in time t
- Parallel work p=1−s completes in time 1/x−t
  - Recall that x is speedup, so 1/x is total time
- $f_s = \frac{s}{t}$
- $f_p = \frac{1-s}{(1/x - t) \cdot N}$
- Energy is composed of
  - Serial dynamic energy $= t \cdot f_s^a$
  - Parallel dynamic energy $= N \cdot \left(\frac{1}{x} - t\right) \cdot f_p^a$
  - Static energy $= N \cdot \lambda \cdot \frac{1}{x}$

Case: x=1

- First consider special case, x=1 (same execution time as serial case)
- Take derivative with respect to *t*

Case: x=1

- First consider special case, x=1 (same execution time as serial case)
- Take derivative with respect to *t*
- $t^* = \frac{s}{s + p/N^{(a-1)/a}}$

$$f_s^* = \frac{s}{t^*} = s + \frac{p}{N^{(a-1)/a}}$$

$$f_p^* = \left(s + \frac{p}{N^{(\alpha-1)/\alpha}}\right)/N^{\frac{1}{\alpha}}$$

Case: x=1

- First consider special case, x=1 (same execution time as serial case)
- Take derivative with respect to *t*
- $t^* = \frac{s}{s + p/N^{(a-1)/a}}$

$$f_s^* = \frac{s}{t^*} = s + \frac{p}{N^{(a-1)/a}}$$

$$f_p^* = \left(s + \frac{p}{N^{(\alpha-1)/\alpha}}\right)/N^{\frac{1}{\alpha}} = f_s^*/N^{\frac{1}{a}}$$

- Relation between f_s^* and f_p^* is dependent on N, but not *s*!
- Note that static energy has no effect in this case

Case: x=1

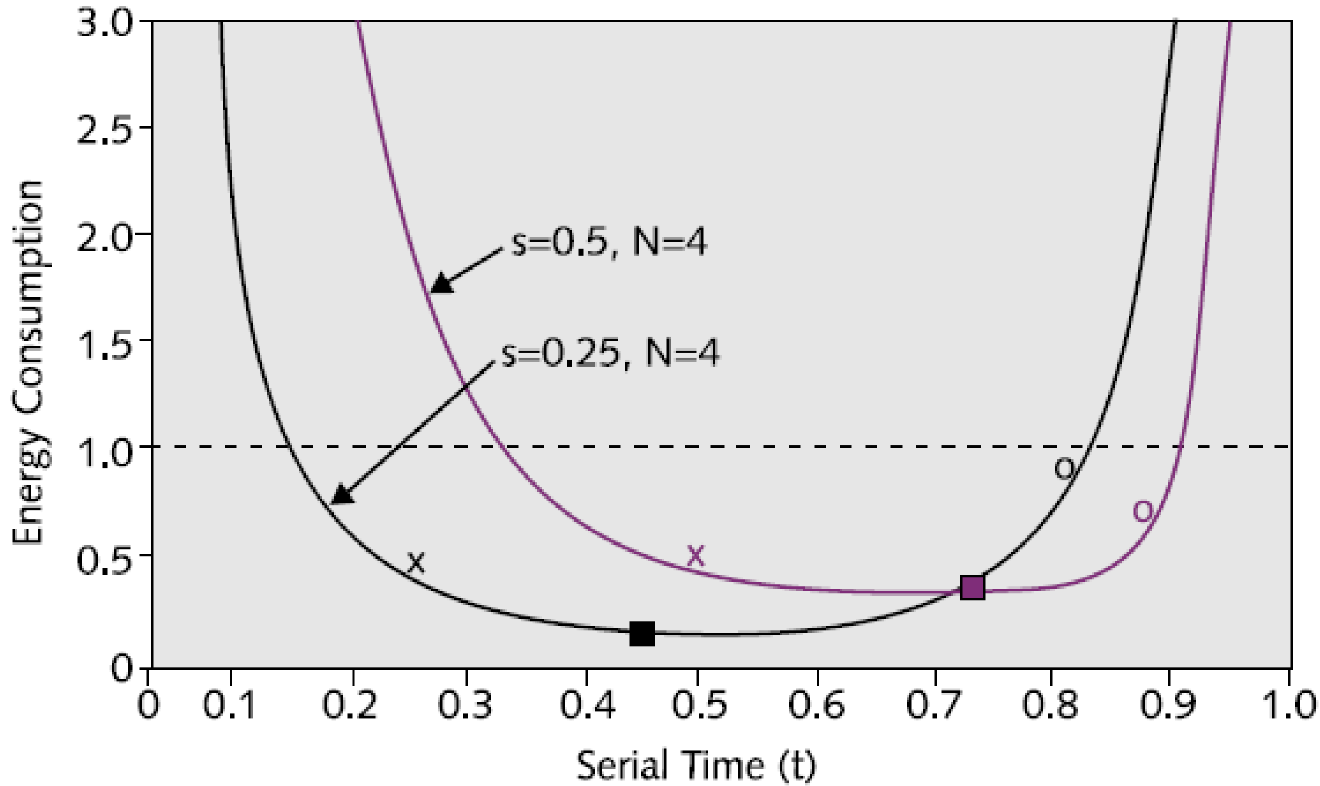

3.0
2.5
2.0
1.5
1.0
0.5
0
Energy Consumption
s=0.5, N=4
s=0.25, N=4
x
x
o
o
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1.0
Serial Time (t)

Case: x=1

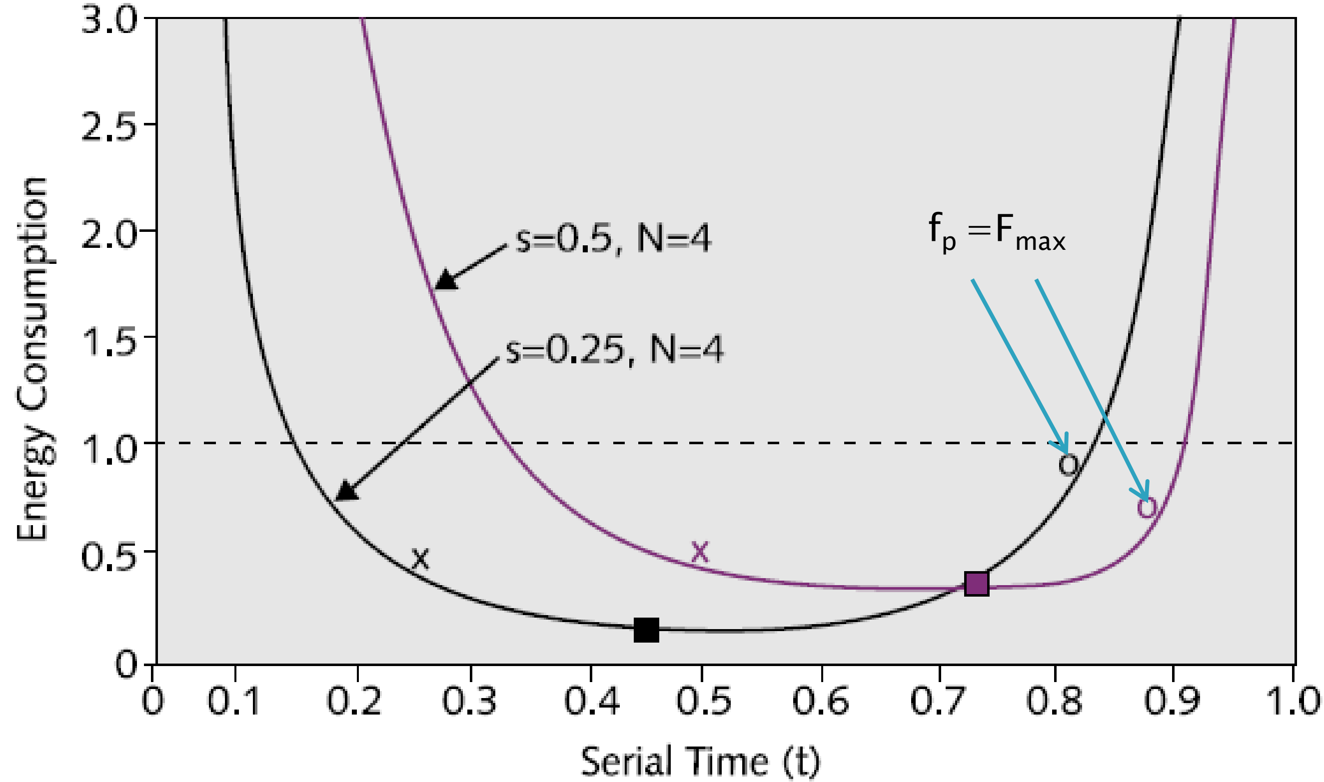

3.0
2.5
2.0
1.5
1.0
0.5
0
Energy Consumption
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1.0
Serial Time (t)
s=0.5, N=4
s=0.25, N=4
$f_p = F_{max}$
x
x
o
o

Case: x=1

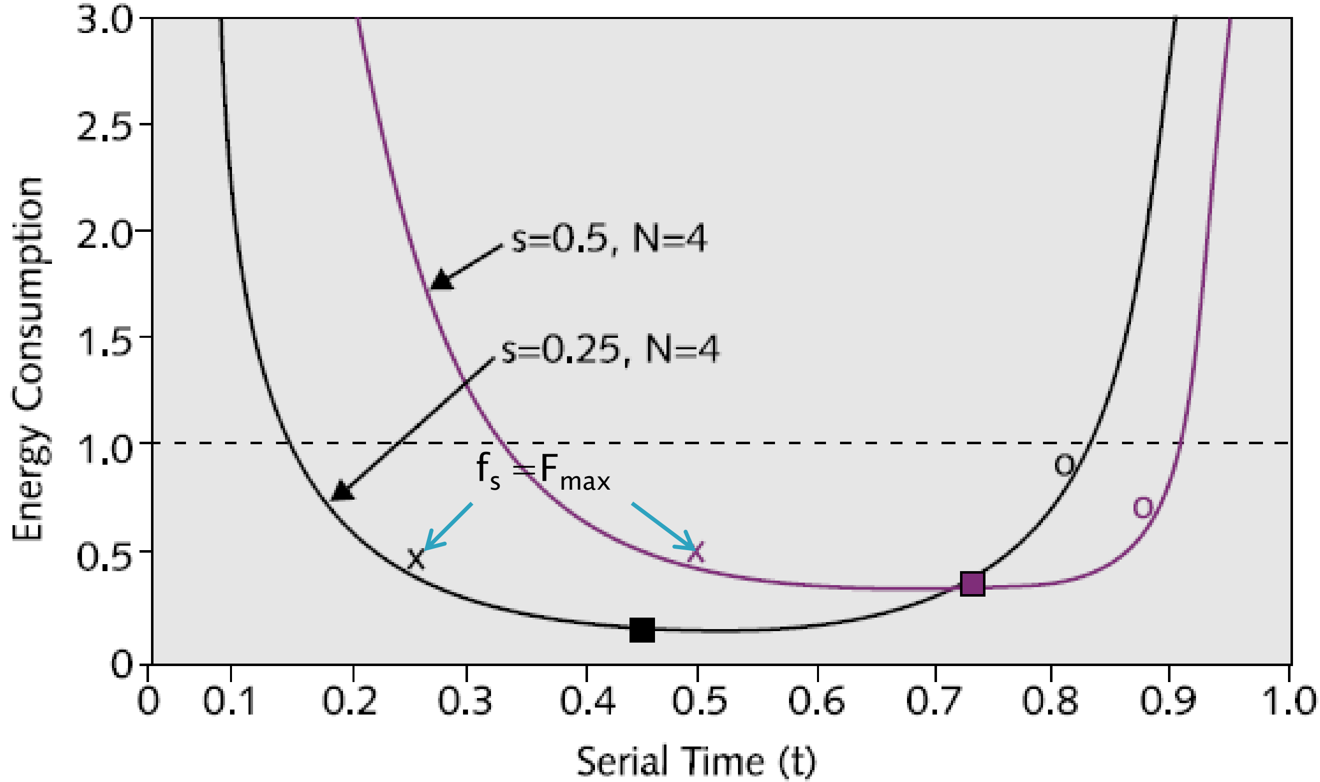

3.0
2.5
2.0
1.5
1.0
0.5
0
Energy Consumption
s=0.5, N=4
s=0.25, N=4
$f_s = F_{max}$
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1.0
Serial Time (t)

The Effect of Static Power

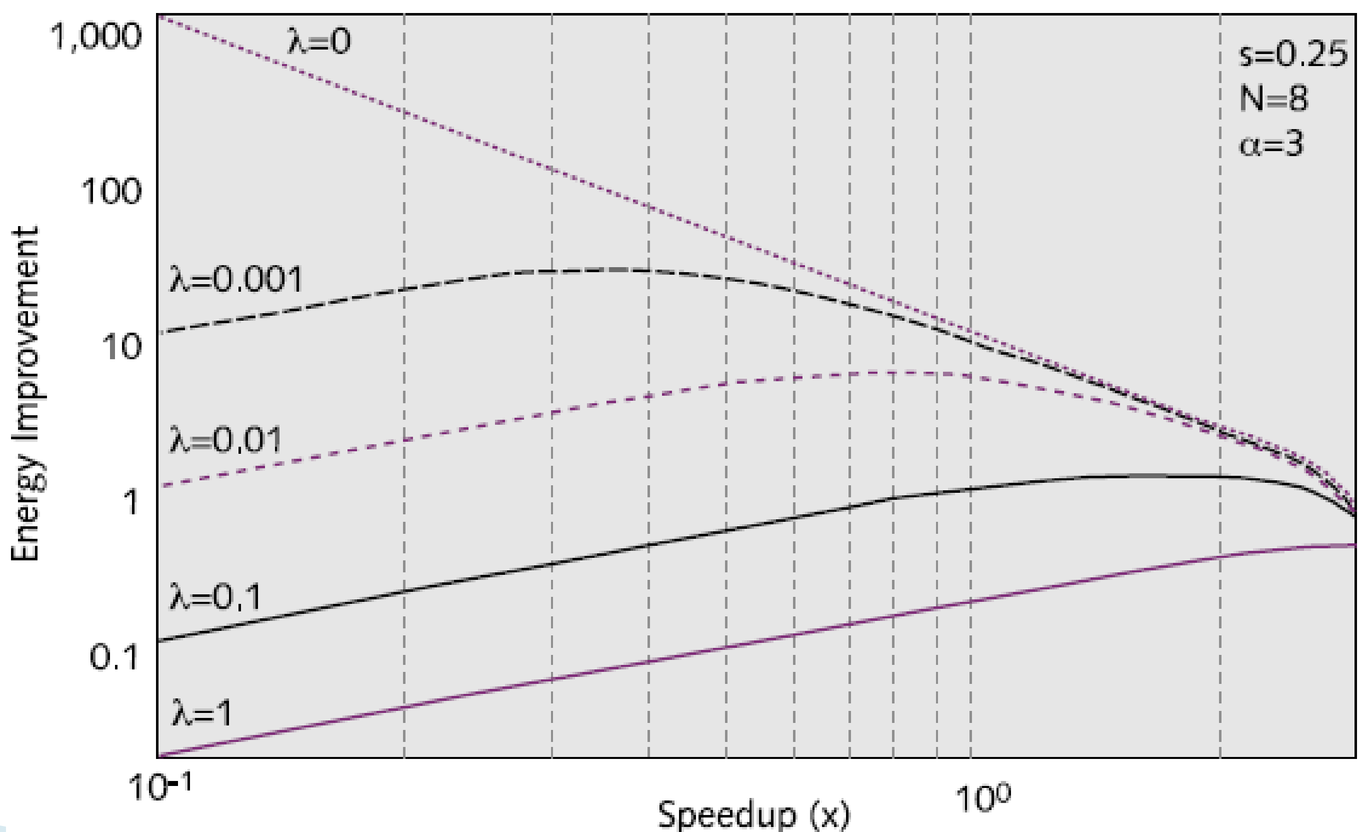

The Effect of Static Power

- Relax execution time constraint
  - Set derivatives with respect to both x and t to 0
  - Solve equations

The Effect of Static Power

- Relax execution time constraint
  - Set derivatives with respect to both *x* and *t* to 0
  - Solve equations

$$t^* = \left(\frac{\alpha - 1}{\lambda N}\right)^{\frac{1}{\alpha}} \cdot s$$

$$x^* = \left(\frac{\lambda N}{\alpha - 1}\right)^{\frac{1}{\alpha}} \cdot \left(\frac{1}{s + \dfrac{p}{N^{(\alpha-1)/\alpha}}}\right)$$

The Effect of Static Power

- Relax execution time constraint
  - Set derivatives with respect to both *x* and *t* to 0
  - Solve equations

$$t^* = \left(\frac{\alpha - 1}{\lambda N}\right)^{\frac{1}{\alpha}} \cdot s$$

$$x^* = \left(\frac{\lambda N}{\alpha - 1}\right)^{\frac{1}{\alpha}} \cdot \left(\frac{1}{s + \frac{p}{N^{(\alpha-1)/\alpha}}}\right)$$

$$f_s^* = \left(\frac{\lambda N}{\alpha - 1}\right)^{\frac{1}{\alpha}}$$

$$f_p^* = \left(\frac{\lambda}{\alpha - 1}\right)^{\frac{1}{\alpha}} = f_s^* / N^{\frac{1}{\alpha}}$$

The Effect of Static Power

- Relax execution time constraint
  - Set derivatives with respect to both *x* and *t* to 0
  - Solve equations

$$t^* = \left(\frac{\alpha - 1}{\lambda N}\right)^{\frac{1}{\alpha}} \cdot s$$

$$x^* = \left(\frac{\lambda N}{\alpha - 1}\right)^{\frac{1}{\alpha}} \cdot \left(\frac{1}{s + \frac{p}{N^{(\alpha-1)/\alpha}}}\right)$$

$$f_s^* = \left(\frac{\lambda N}{\alpha - 1}\right)^{\frac{1}{\alpha}}$$

$$f_p^* = \left(\frac{\lambda}{\alpha - 1}\right)^{\frac{1}{\alpha}} = f_s^* / N^{\frac{1}{\alpha}}$$

The Effect of Static Power

- Relax execution time constraint
  - Set derivatives with respect to both *x* and *t* to 0
  - Solve equations

$$t^* = \left(\frac{\alpha - 1}{\lambda N}\right)^{\frac{1}{\alpha}} \cdot s$$

$$x^* = \left(\frac{\lambda N}{\alpha - 1}\right)^{\frac{1}{\alpha}} \cdot \left(\frac{1}{s + \frac{p}{N^{(\alpha-1)/\alpha}}}\right)$$

$$f_s^* = \left(\frac{\lambda N}{\alpha - 1}\right)^{\frac{1}{\alpha}}$$

$$f_p^* = \left(\frac{\lambda}{\alpha - 1}\right)^{\frac{1}{\alpha}} = f_s^* / N^{\frac{1}{\alpha}}$$

The Effect of Static Power

- Relax execution time constraint
  - Set derivatives with respect to both *x* and *t* to 0
  - Solve equations

$$t^* = \left(\frac{\alpha - 1}{\lambda N}\right)^{\frac{1}{\alpha}} \cdot s$$

$$x^* = \left(\frac{\lambda N}{\alpha - 1}\right)^{\frac{1}{\alpha}} \cdot \left(\frac{1}{s + \frac{p}{N^{(\alpha-1)/\alpha}}}\right)$$

$$f_s^* = \left(\frac{\lambda N}{\alpha - 1}\right)^{\frac{1}{\alpha}}$$

$$f_p^* = \left(\frac{\lambda}{\alpha - 1}\right)^{\frac{1}{\alpha}} = f_s^* / N^{\frac{1}{\alpha}}$$

$$E_{dynamic} = \boxed{\frac{1}{\alpha - 1}} \cdot \frac{N\lambda}{x^*} \qquad E_{static} = \frac{N\lambda}{x^*}$$

The Effect of Static Power

- Previous solution necessitates $\lambda N \leq \alpha - 1$
  - $f_s = F_{max}$

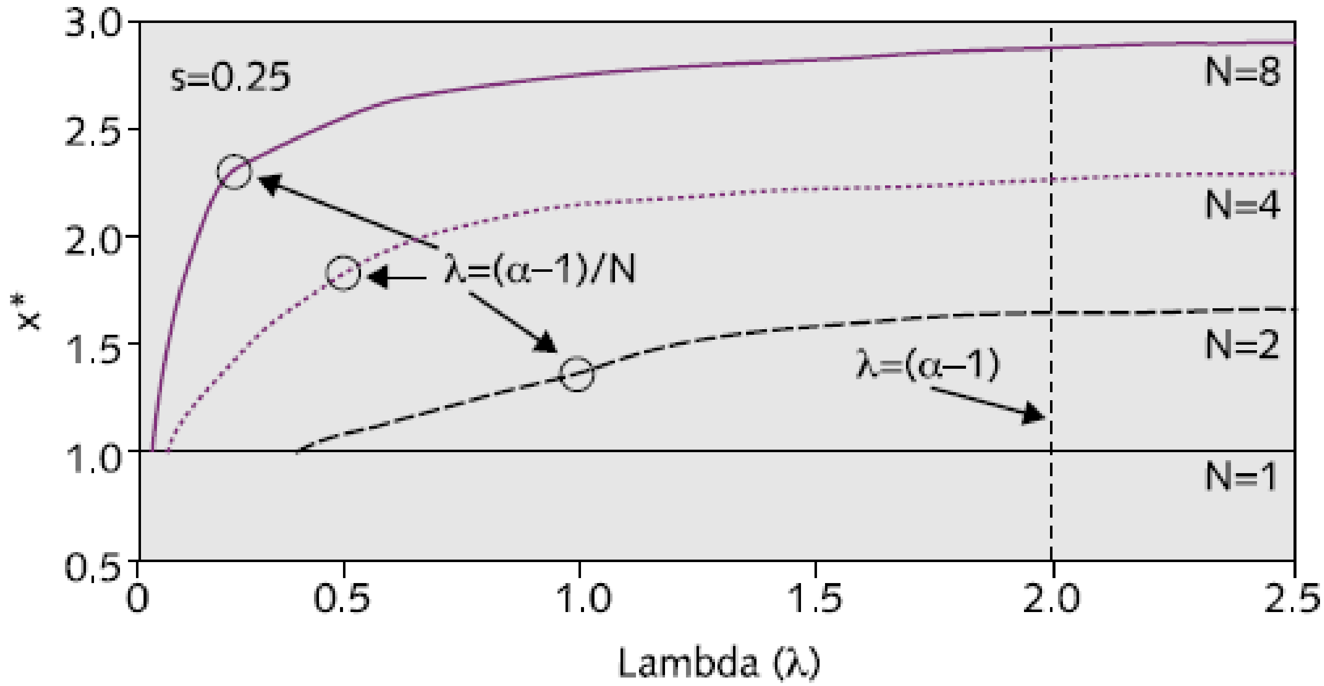

Dynamic and Static Energy

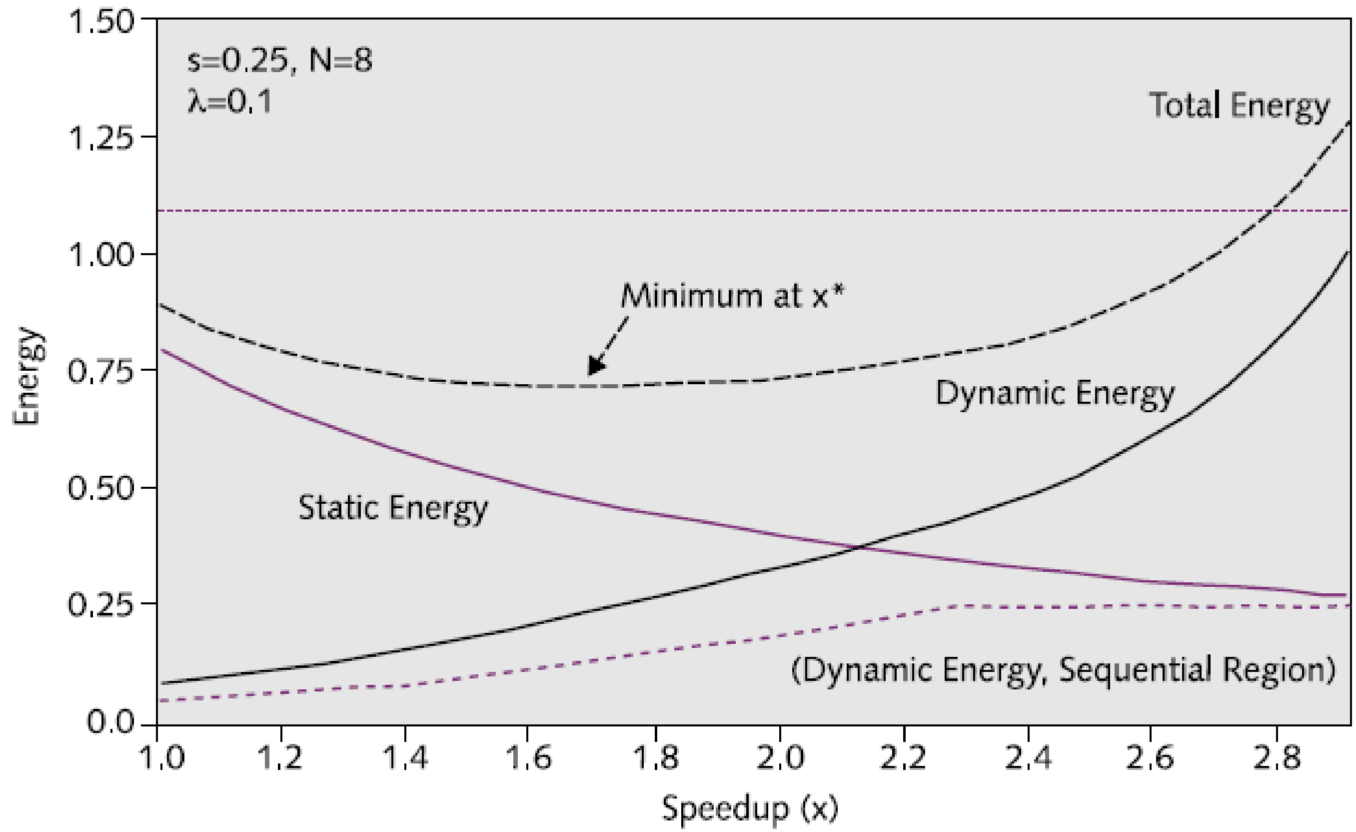

Other Results

- Energy Delay

$$ED(t,x) = (t \cdot f_s^{\alpha} + N \cdot \left(\frac{1}{x} - t\right) \cdot f_p^{\alpha} + N \cdot \lambda \cdot \frac{1}{x}) \cdot \frac{1}{x}$$

Other Results

- Energy Delay

$$ED(t,x) = (t \cdot f_s^{\alpha} + N \cdot \left(\frac{1}{x} - t\right) \cdot f_p^{\alpha} + N \cdot \lambda \cdot \frac{1}{x}) \cdot \frac{1}{x}$$

$$f_s^* = \left(\frac{2N\lambda}{\alpha - 2}\right)^{1/\alpha}$$

$$f_p^* = \left(\frac{2\lambda}{\alpha - 2}\right)^{1/\alpha}$$

Other Results

- Energy Delay

$$ED(t,x) = (t \cdot f_s^{\alpha} + N \cdot \left(\frac{1}{x} - t\right) \cdot f_p^{\alpha} + N \cdot \lambda \cdot \frac{1}{x}) \cdot \frac{1}{x}$$

$$f_s^* = \left(\frac{2N\lambda}{\alpha - 2}\right)^{1/\alpha}$$

$$f_p^* = \left(\frac{2\lambda}{\alpha - 2}\right)^{1/\alpha}$$

- Same relation between f_s^* and f_p^*

- Implications?

Other Results

- Synchronization cost
  - Parallel architectures have some overhead for communication
  - As N increases, more total work is required, which can be expressed as a function
  - Total work is now $s + p(1 + \sigma(N))$

Other Results

- Synchronization cost
  - Parallel architectures have some overhead for communication
  - As N increases, more total work is required, which can be expressed as a function
  - Total work is now $s + (p + \sigma(N))$

- Optimal serial and parallel frequencies have the same relationship
  - Still differ by $N^{1/\alpha}$

Andahl's Law in the Multicore Era

Mark D. Hill, University of Wisconsin-Madison and
Michael R. Marty, Google

Presented By: Michael Moeng

Asymmetric Cores

- Chip has a total resource budget, defined in the number of 'baseline' cores that can be supported
- Assume performance for an *r*–BCE core grows by sqrt(*r*)

Trade-off power vs performance