



# Improving LUC in Cloud Environments


Kirk W. Cameron & Ali R. Butt  
Computer Science  
Virginia Tech

This material is based upon work supported by the National  
Science Foundation under CNS Grant No. 1422788.

stack@cs.vt.edu



THE GREEN  
500™



spec



**MiserWare**  
Saving Energy, Saving Money, Saving the World.


# Performance according to...

- Greg Linden, Amazon Source: [highscalability.com/](http://highscalability.com/)
  - “Every 100 ms of latency costs them 1% in sales.”
- Marissa Meyer, Google
  - “An extra .5 seconds in search page generation time dropped traffic by 20%.”



# Conventional wisdom

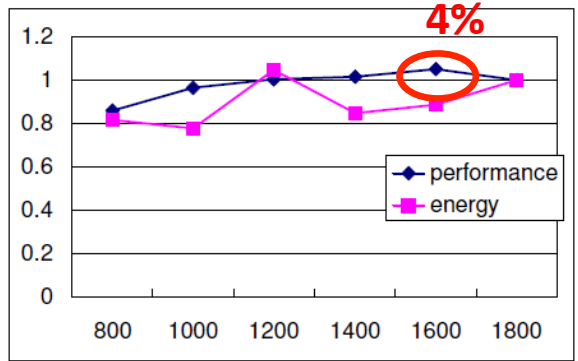
Higher frequency = higher performance



# Not-so-conventional results


Over 30 peer-reviewed publications observe slowdowns at higher speeds.

University of Tsukuba, IPDPS'06




Memory bound

Beijing Computing Center, PDCAT'11



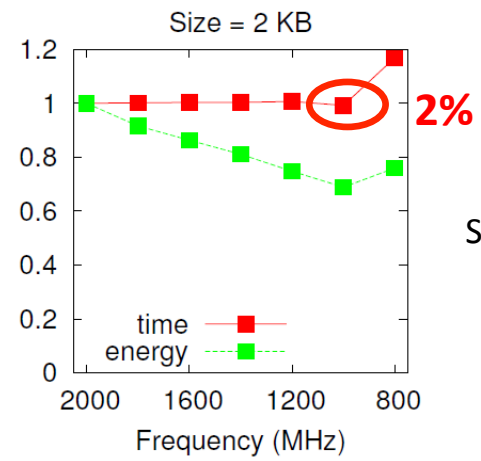
Communication bound

IBM, CF'05, IPDPS'05



CPU bound

North Carolina State University, SC'06



MPI library call

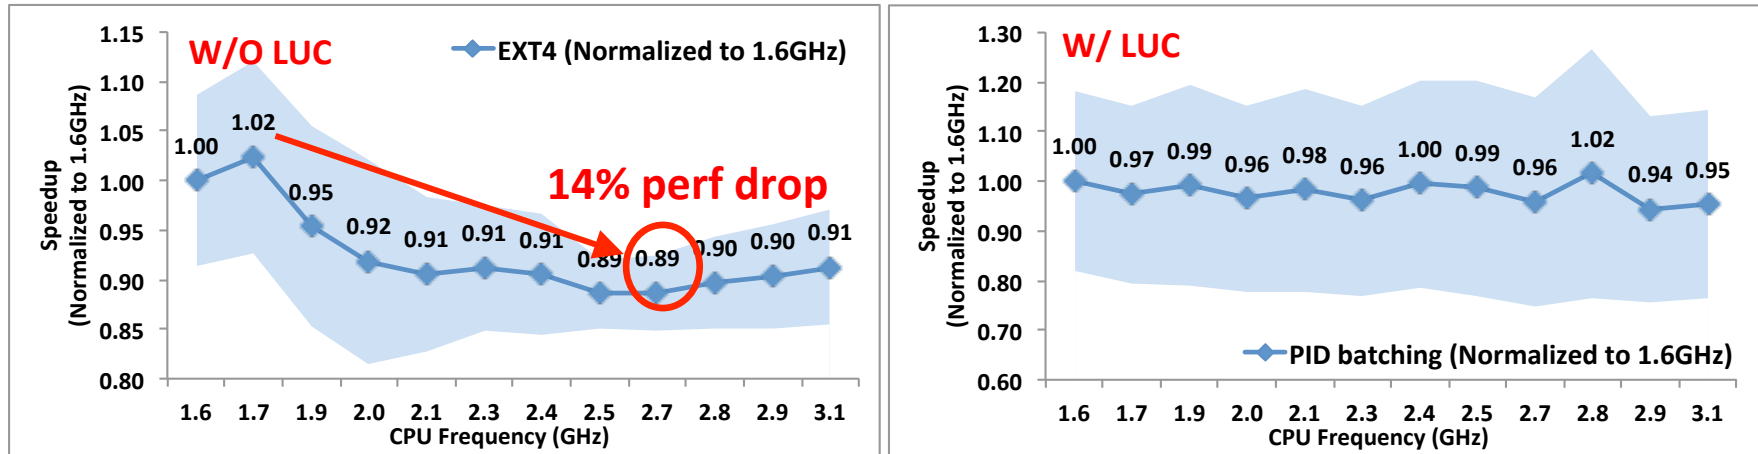
# Unintended consequences

Slowdown from power scaling is pervasive and significant in I/O.

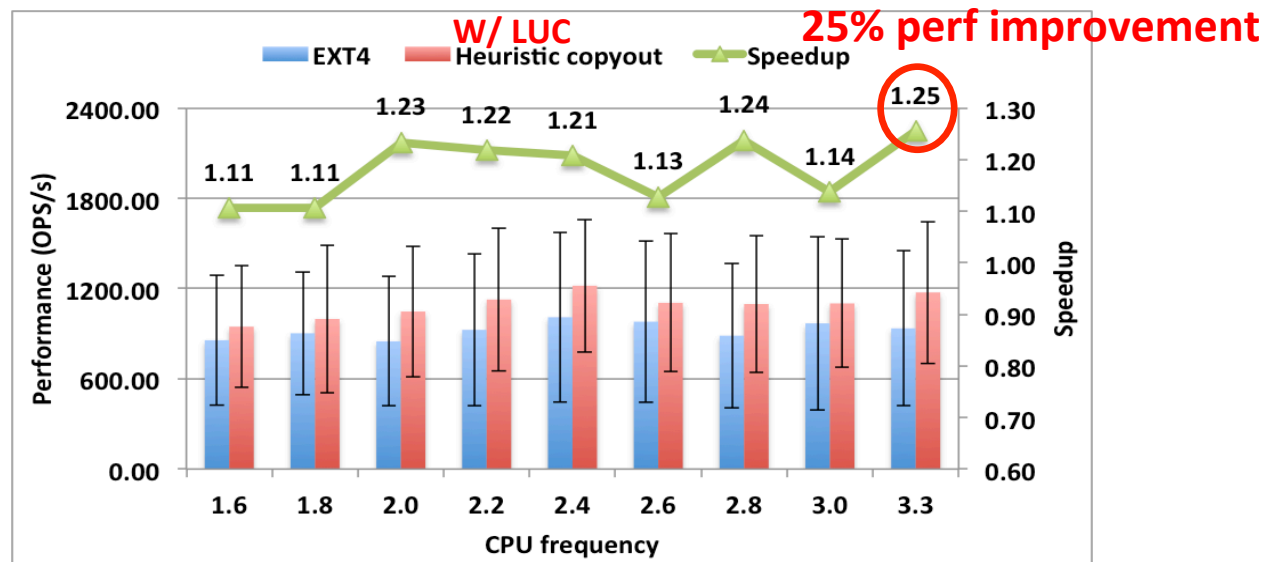
	HW	Comm	SW	I/O
Intel: ICAC'11	2%			
IBM: CF'05, IPDPS'05	3%			
Barcelona Supercomputing Center: JPDC' 12		4%		
North Carolina State University; SC'06		2%		
Virginia Tech: SC'05			12%	
North Carolina State University: PPOPP'05			1%	
NICTA, University of South Wales: HotPower'10			5%	
Marquette Univ. : GreenCom'10, IPDPSW'10, CCGrid'12				40%
Results from OUR WORK				
IOZone (simple)				68%
varMail (complex)				14%

# Limit Unintended Consequences

varMail

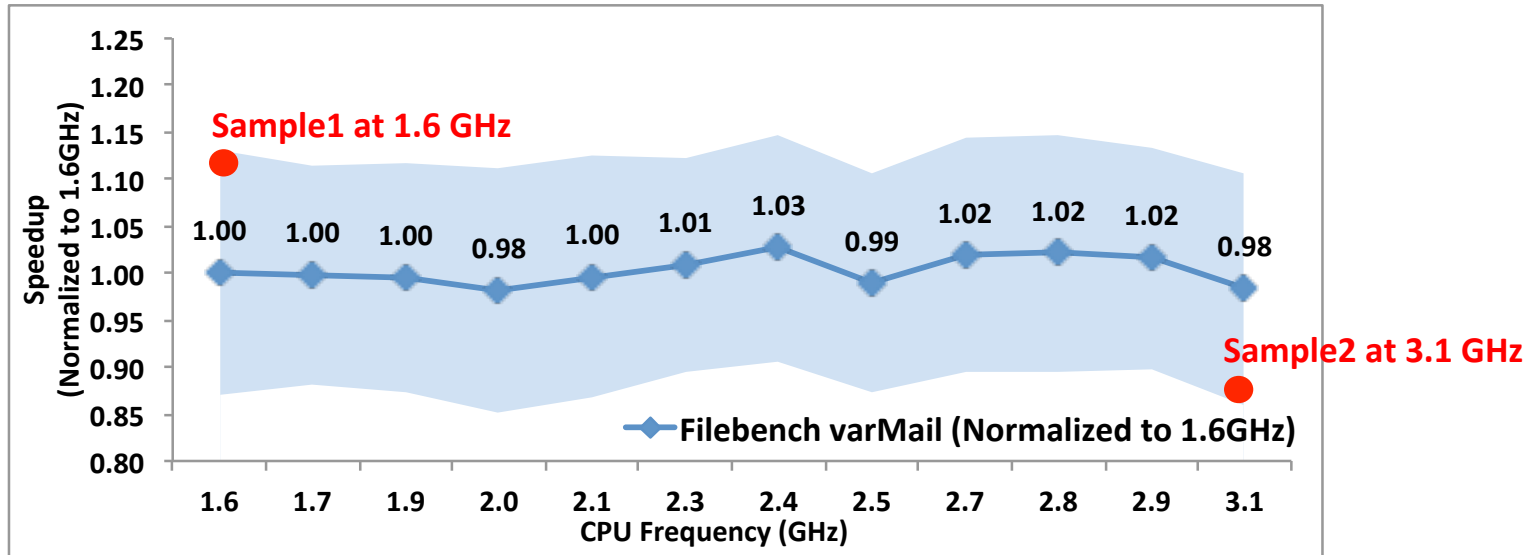


oltp





# Challenge 1: High variance




- Issue: Nondeterminism
  - False positive/negative
  - Magnitude of slowdown varies
- A possible solution: large runs, but ...
  - Takes weeks/months to collect data



# Challenge 2: Complexity

- Systems include
  - Software
  - Middleware →
  - Hardware
- Hard to isolate
  - Which layer?
  - Where in layer?
  - How to find?
- Cloud makes everything harder



Source: linuxfoundation.org

# Experimental wish list

- Design “better” cloud systems

## Performance, Power, Energy

- Measurement
  - Power, performance, thermals (in-band, out-of-band)
  - Fine-grain and course grain
- Hardware
  - Variety (heterogeneity)
  - Scale (more than 256 cores at a time)
- Software
  - Variety (OSs – bare metal, VMs)
  - Tools (MP APIs, Instrumentation)
- Access & Control
  - Time (>90% confidence in results)
  - Flexible resource allocation

cameron@cs.vt.edu



Thanks!