# On the Energy Consumption and Performance of Systems Software

Z. Li, R. Grosu, P. Sehgal, S. Smolka,
S. Stoller, and ***E. Zadok***

http://www.fsl.cs.sunysb.edu

# Motivation

- Optimize energy *and* performance
- Compression study [SYSTOR 2009]
  - File type, hardware, compression algs.
  - 10x better, to 200x worse
- Server workload study [FAST 2010]
  - Web/DB/Email/file server workloads
  - F/S mount/format params., hardware
  - 50% to 9x variation in perf./energy

**How can we predict and control these savings?!**

# Server Energy: Text File

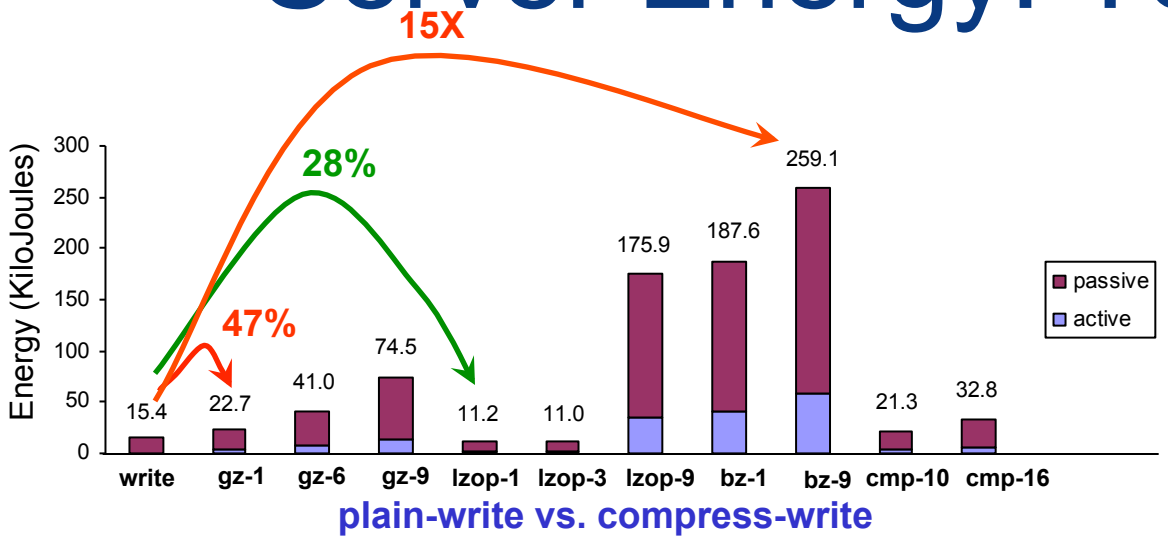


plain-write vs. compress-write

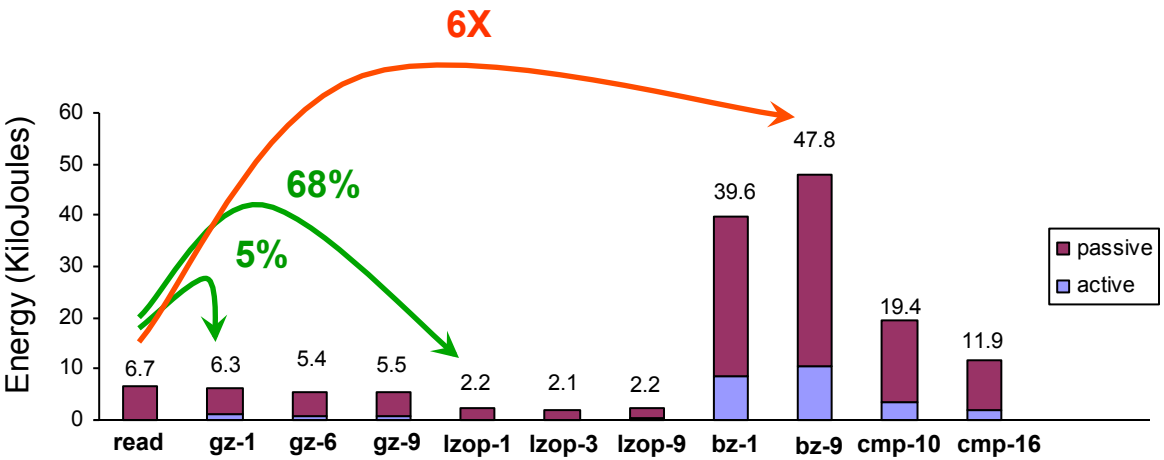


plain-read vs. decompress-read

| | |
|---|---|
| lzop-1,3 | ∀ |
| bzip | ✗ |
| gzip | **n reads/writes** |

# Web Server Results

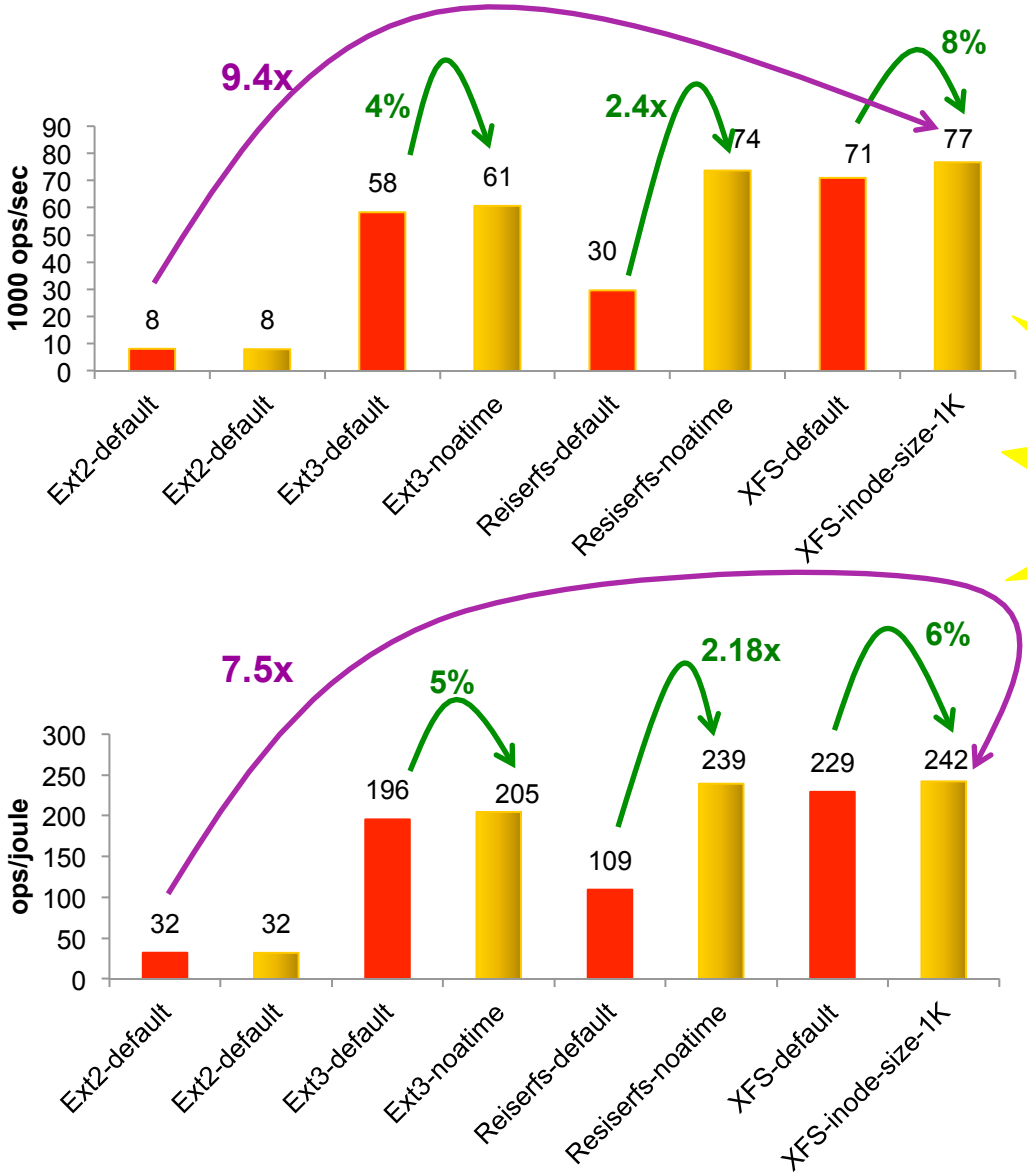


**Performance**

**Is XFS the best for all workloads?**

**Energy Efficiency**

**Default FS Config.**

**Best FS Config.**



STONY BROOK
STATE UNIVERSITY OF NEW YORK

# Overview

- Motivation
- *Related*
- Background
- Methodology
- Evaluation
- Conclusion
- Future

# Related

- Energy Saving Tech
  - Virtualization Techniques
  - Energy-aware cache replacement algorithm, task and interrupt management
  - File systems pruning techniques
  - Predictive data grouping and replication techniques
  - Modeling for optimal use
  - Etc.

# Related (cont.)

- Control theory in Computing system
  - Database Systems
  - Storage Systems
  - Web Servers
  - Data Centers
  - Etc.

***QoS (power and performance) requirements***

# Overview

- Motivation
- Related
- *Background*
- Methodology
- Evaluation
- Conclusion
- Future

# We favor Control Theory

- Control Theory Steps:
  - ***System Identification***
  - Controller Design
  - Controller implementation

# Controller

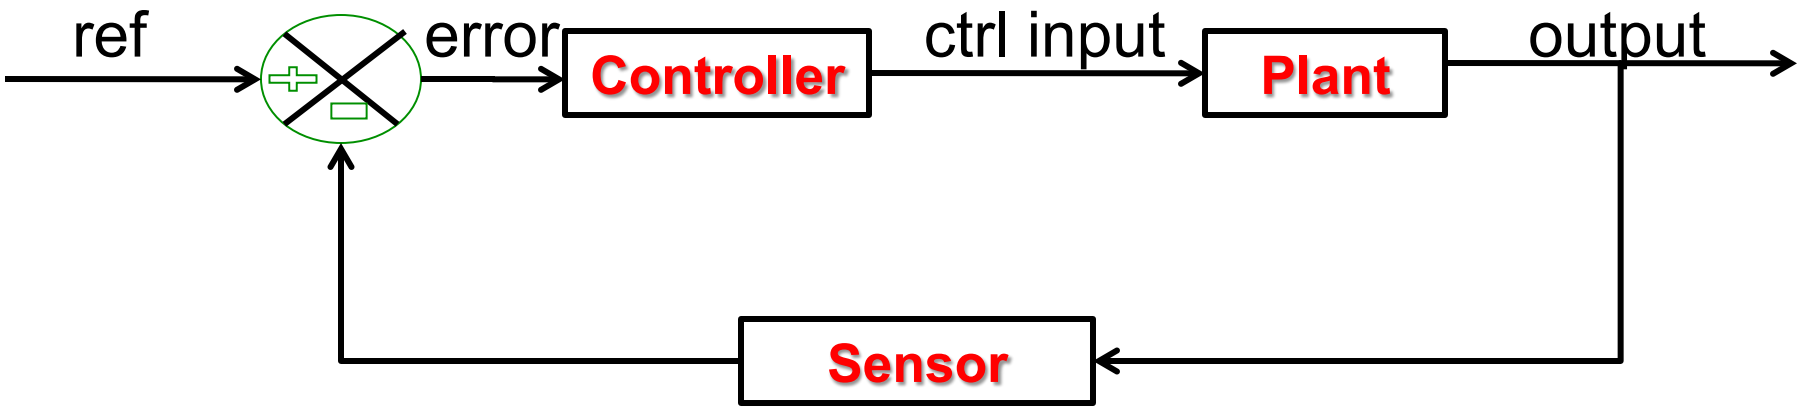


Plant with feedback controller

# System Identification

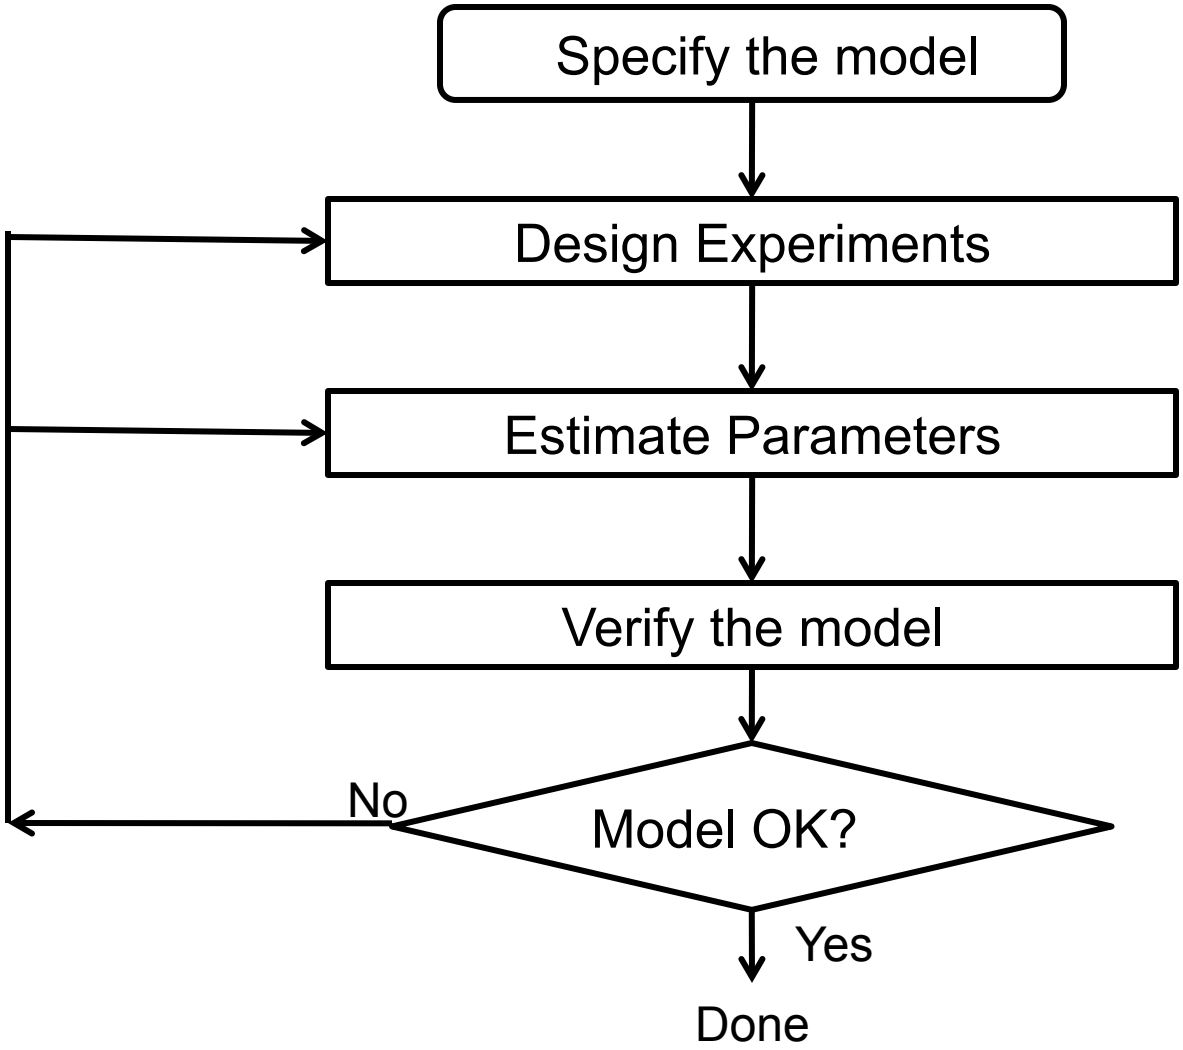


*Understanding the modeling complexity*

# Desired System Plant

- E.g., centralized backup system
  - Receiving multiple streams

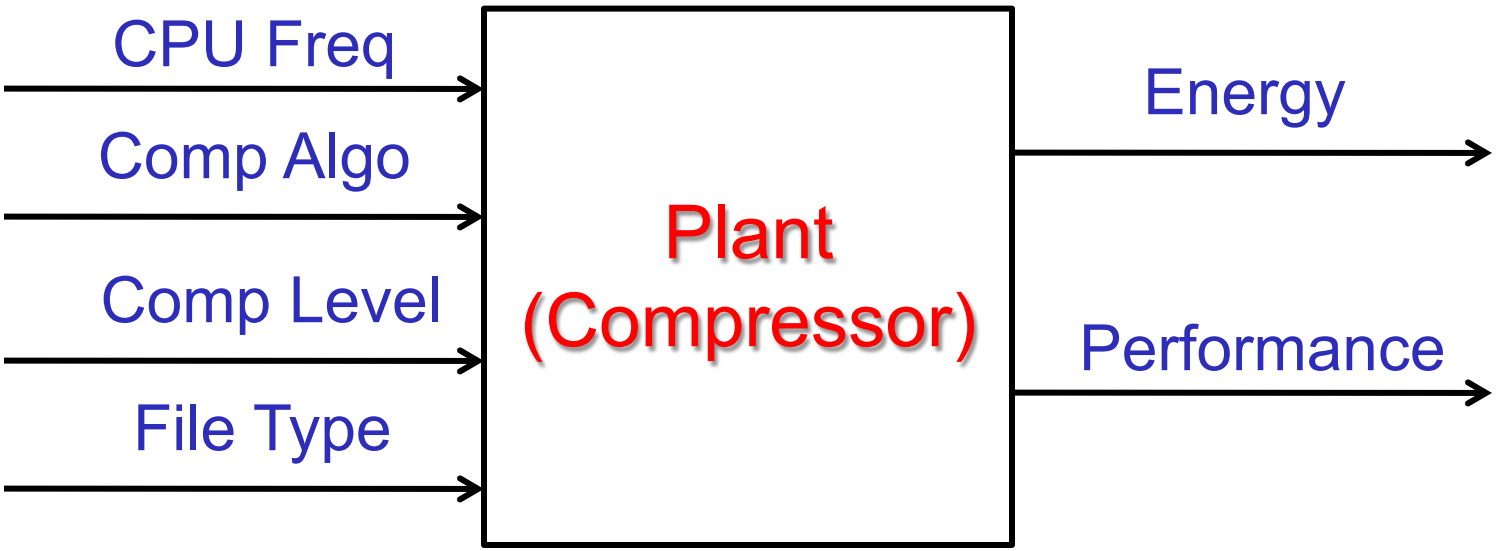

# Results (take 1)

- Low Accuracy (17—24%)

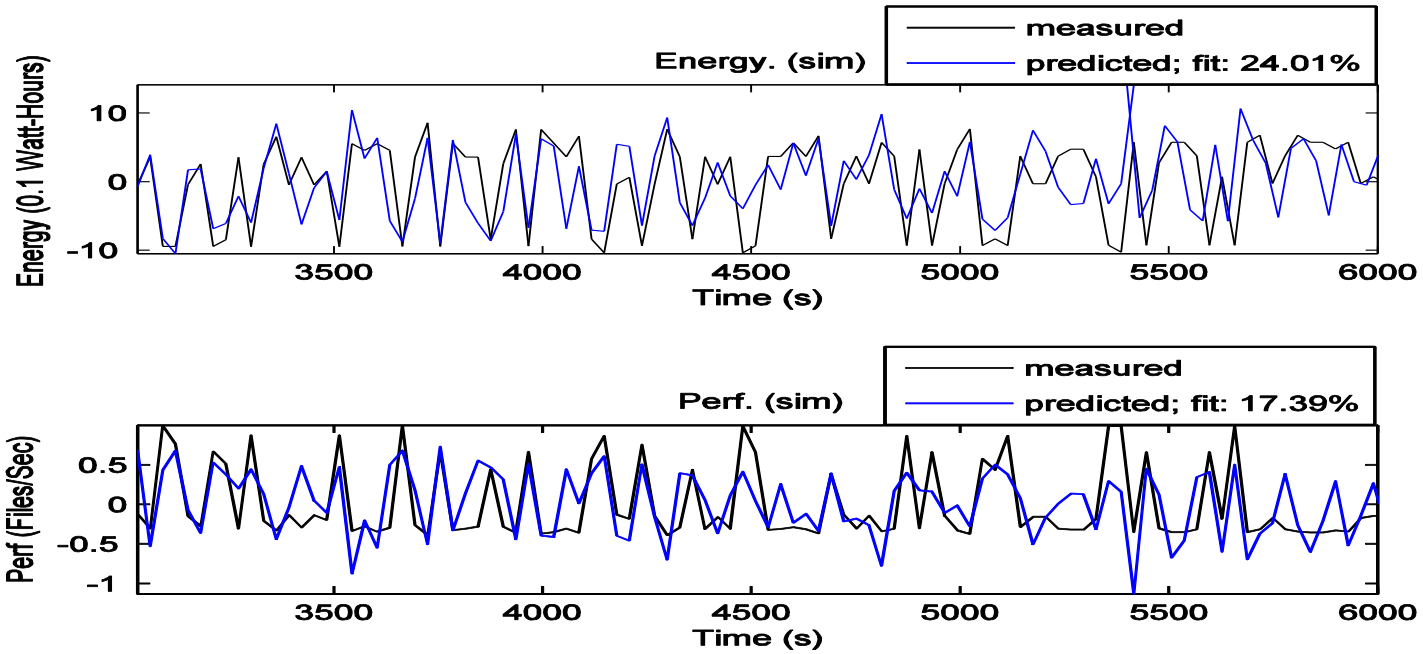


Inputs: File Type & Freq; Gzip + level 9

**Lead to study power consumption and performance in more detail**

# Overview

- Motivation
- Related
- Background
- *Methodology*
- Evaluation
- Conclusion
- Future

# Methodology

- Hardware
  - Dell PowerEdge R710, Wattsup meter
- Benchmarks, vary factors
  - Disk type, Scheduler, File type, DVFS, Compression algorithm & level
- Experiments
  - 4,810,320 data points per run
  - 15 clock days single run

# Hardware Setup

**Linux Server**

Server Power
Readings (USB)

**WattsUP Pro ES (server)**

**A/C Power Supply**

# Overview

- Motivation
- Related
- Background
- Methodology
- *Evaluation*
- Conclusion
- Future

# Evaluation

- Nonlinearity
- Instability
- Multi-dimensionality
  - CPU Frequency
  - I/O Schedulers
  - File Types
  - Disk Types
  - Compression Algorithm + Level
- Non-numeric labels

# Nonlinearity

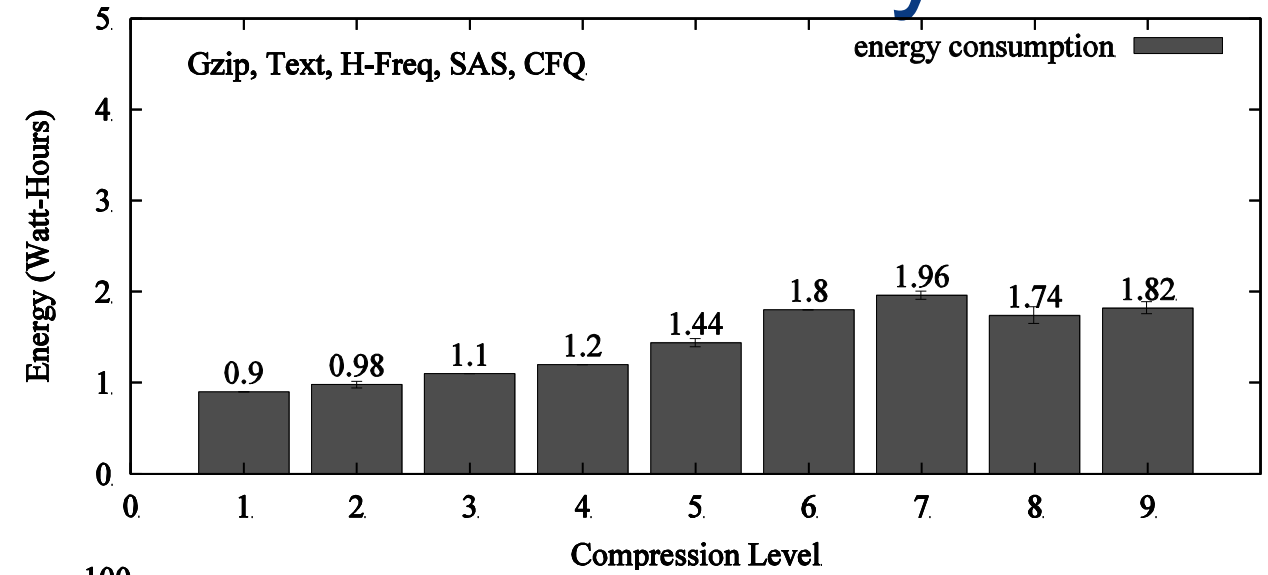


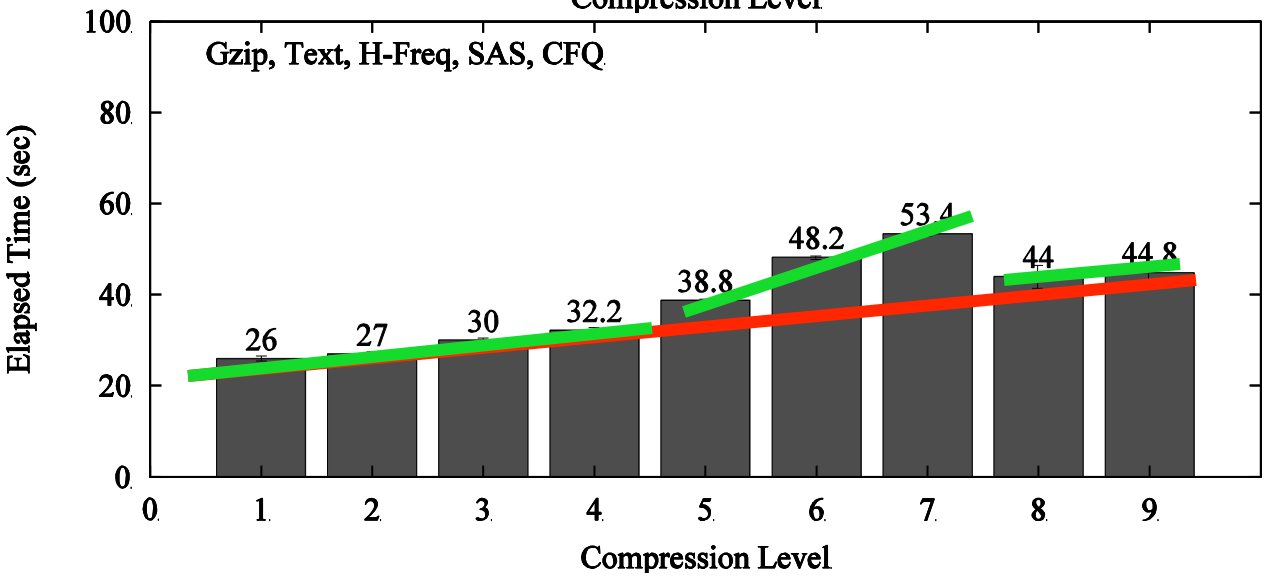




STONY BROOK
STATE UNIVERSITY OF NEW YORK

# Instability

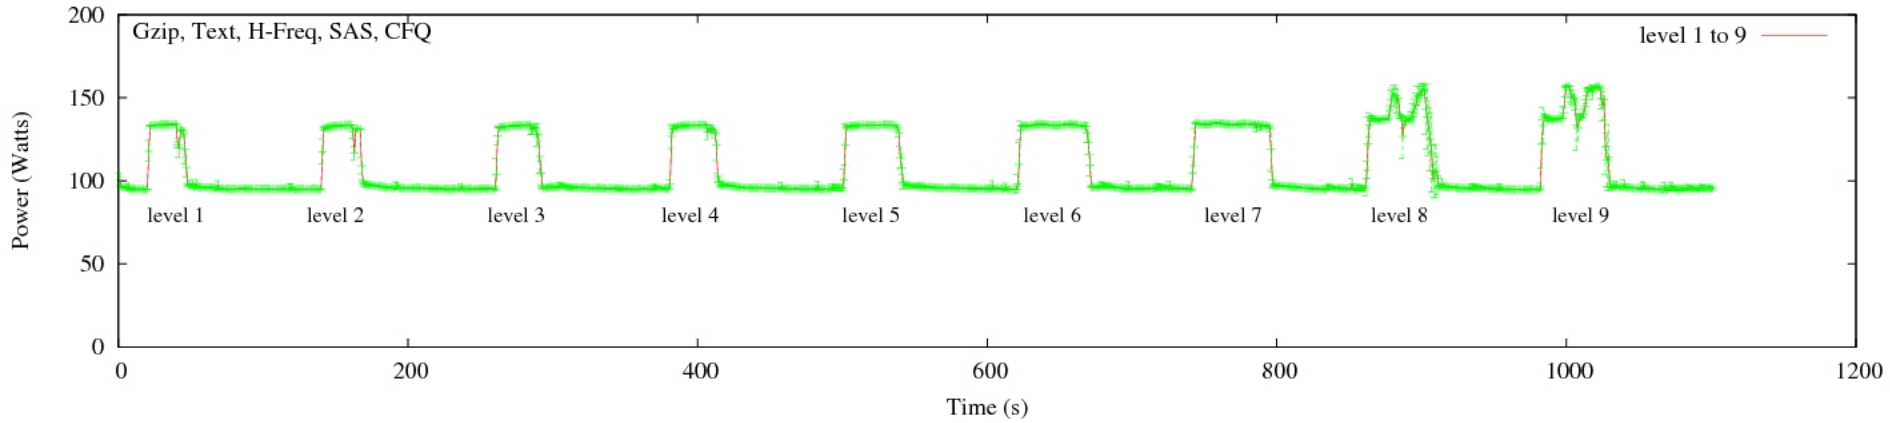


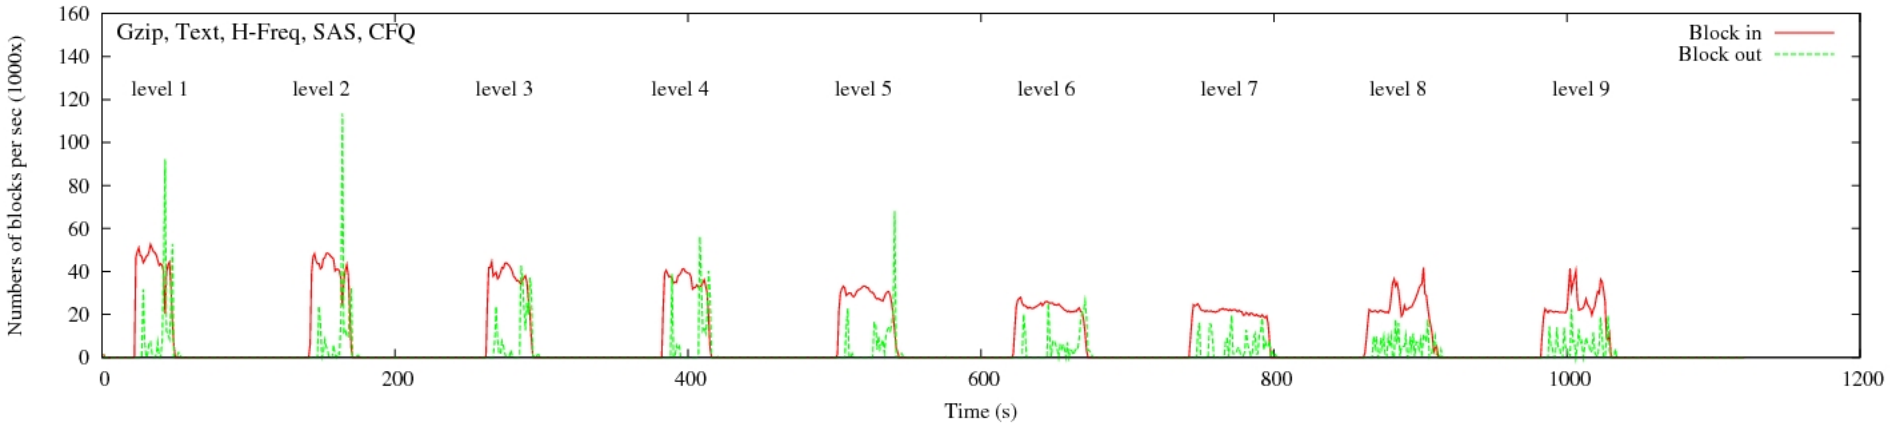




STONY BROOK
STATE UNIVERSITY OF NEW YORK

# Multi-dimensionality

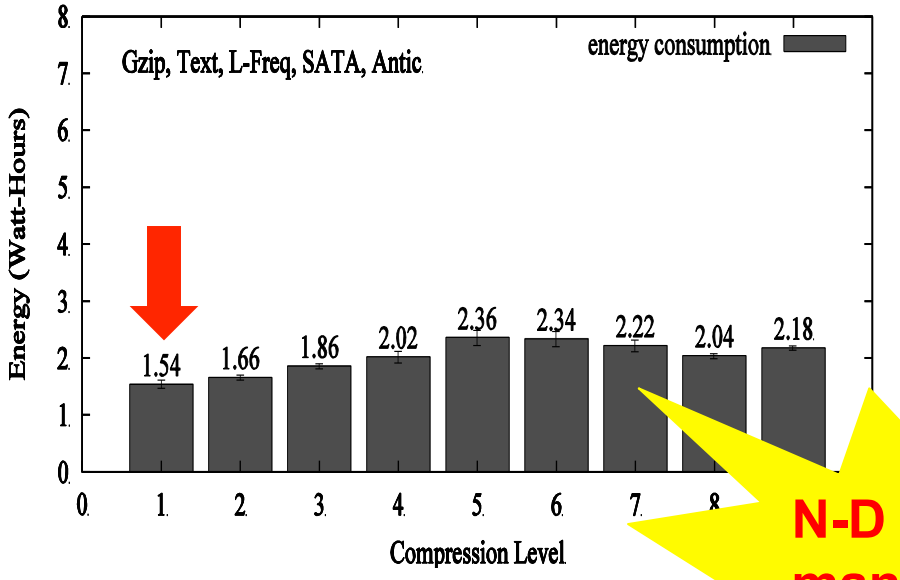


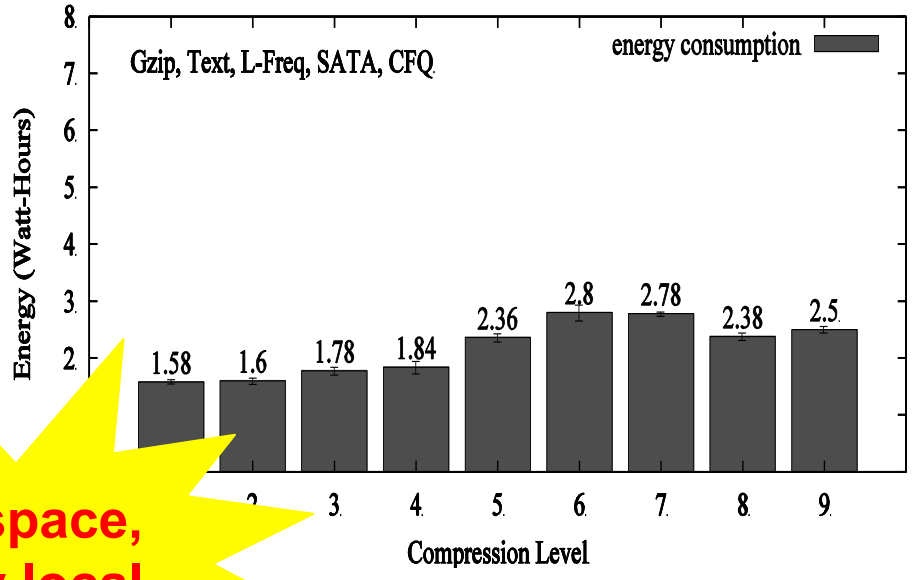


N-D space, many local minima

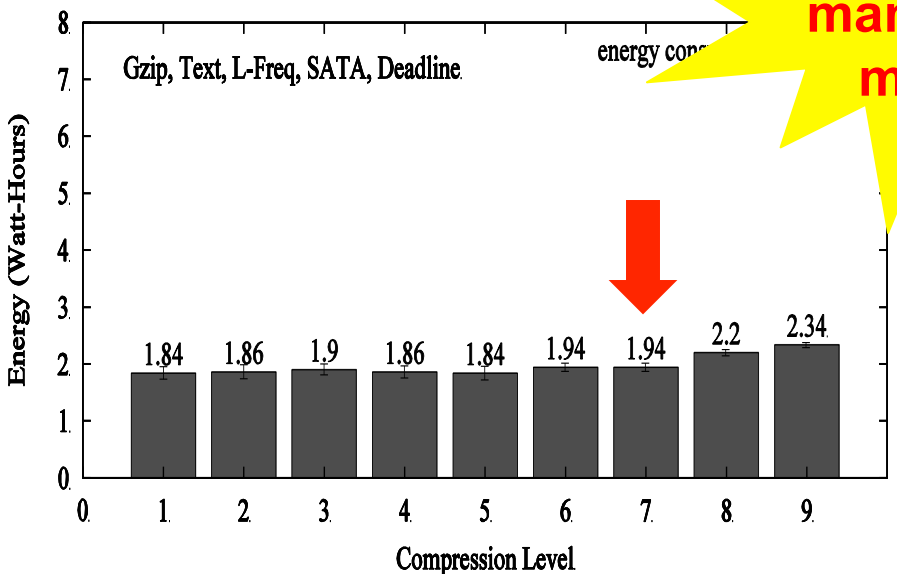


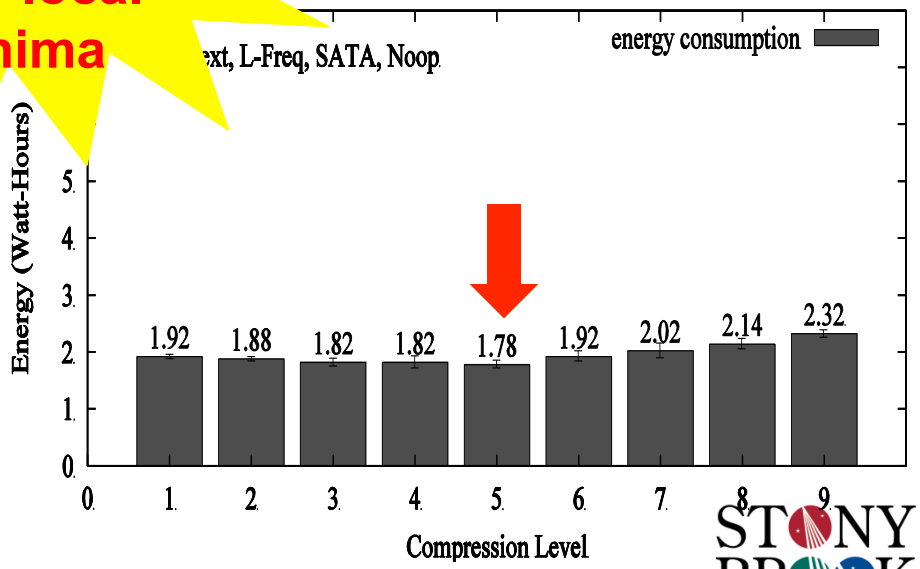




STONY BROOK
STATE UNIVERSITY OF NEW YORK

# Multi-dimensionality (cont.)

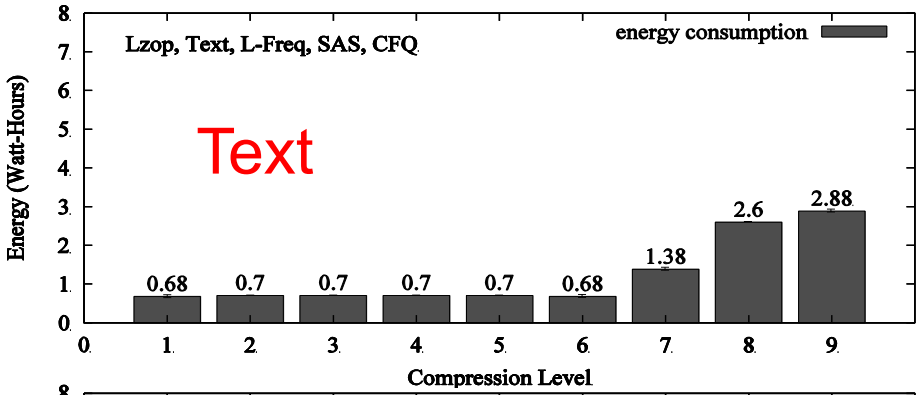


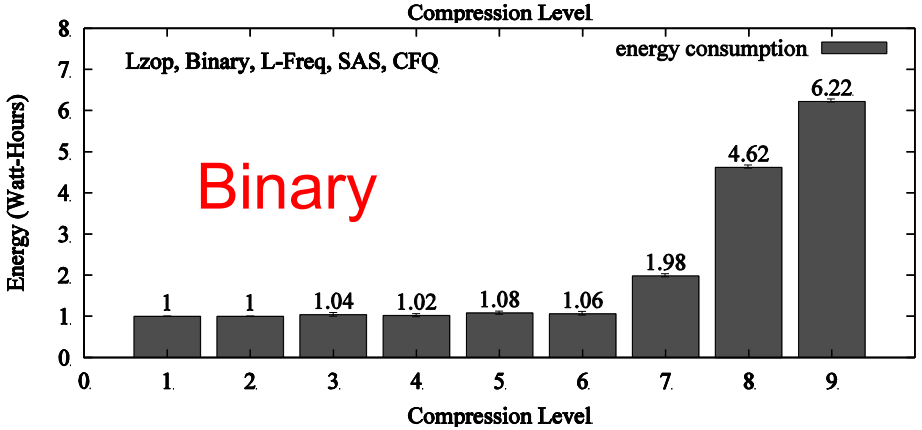


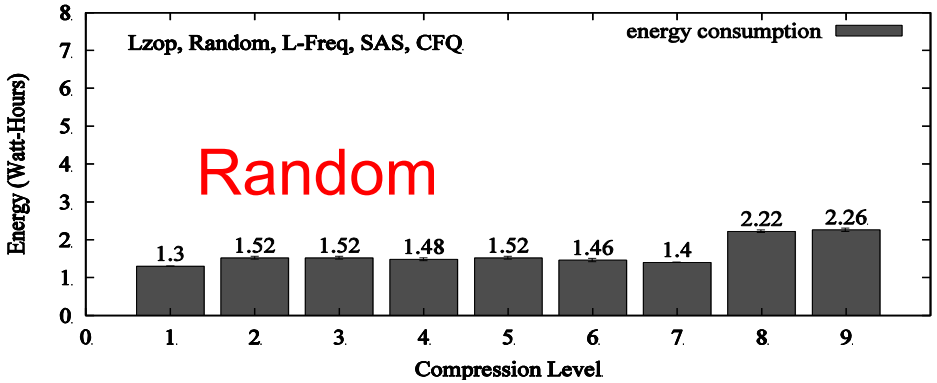


# Labels as Inputs

- Label Issue
  - Compression Algorithm
  - Compression Level
  - File Type

The numerical value should not impose arbitrary quantitative relationships

# Overview

- Motivation
- Related
- Background
- Methodology
- Evaluation
- *Conclusion*
- Future

# Conclusion

- Software systems are very complex
- Great savings *are* possible
- Experimental foundation for further research

# Overview

- Motivation
- Related
- Background
- Methodology
- Evaluation
- Conclusion
- *Future*

# Models

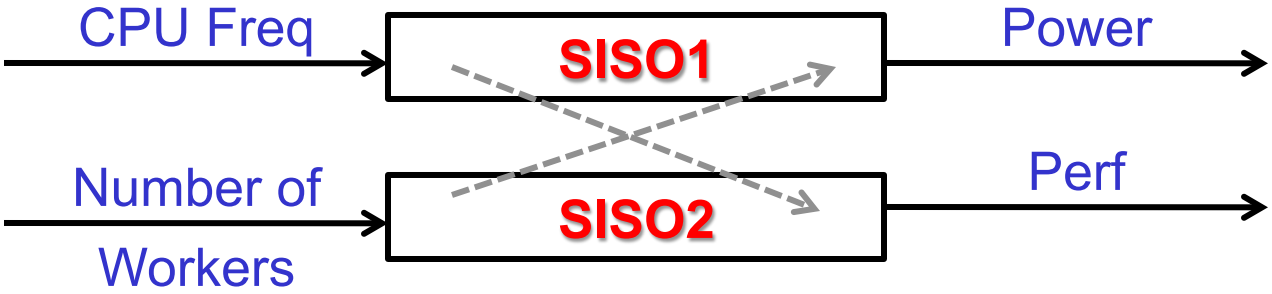



MIMO model and two SISO models

# MIMO vs. SISO

| Model | Fixed Input | Order | Accuracy |
|---|---|---|---|
| MIMO | N/A | 3 | Power: 77%<br>Perf: 76% |
| SISO1 | 1 worker | 1 | Power: 73% |
| SISO1 | 2 workers | 1 | Power: 73% |
| SISO1 | 3 workers | 1 | Power: 73% |
| SISO1 | 4 workers | 1 | Power: 71% |
| SISO2 | 2395MHz Freq | 1 | Perf: 43% |
| SISO2 | 1995MHz Freq | 2 | Perf: 61% |
| SISO2 | 1596MHz Freq | 1 | Perf: 44% |

Evaluation of MIMO and SISO models with MIMO data [ERSS'11]

# למה מה?
# Q&A