Paper-Reading Group

JouleGuard: Energy Guarantees for Approximate Computations

Henry Hoffmann, University of Chicago

SOSP'15

Motivation

- Energy is important
  - To users (mobile, electricity bill)
  - To operators (data center)
- Hardware has many trade-offs
  - Big/little
  - Frequencies …
- Approximate Computing is getting more relevant
  - Search, Big Data, Financial Analysis ...

Motivation

What is the best approximation we can achieve on an energy budget?

Example

- Search Engine (swish++)
  - 90 mJ / Query
  - How can we get to 90mJ / Query?

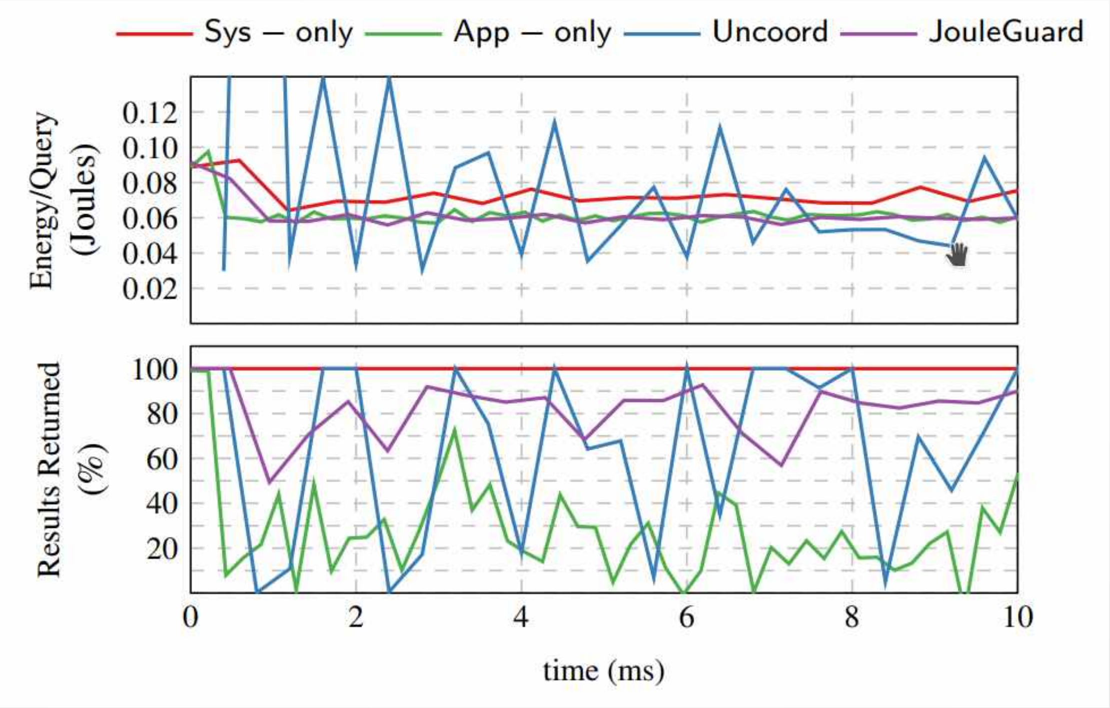

Coordination

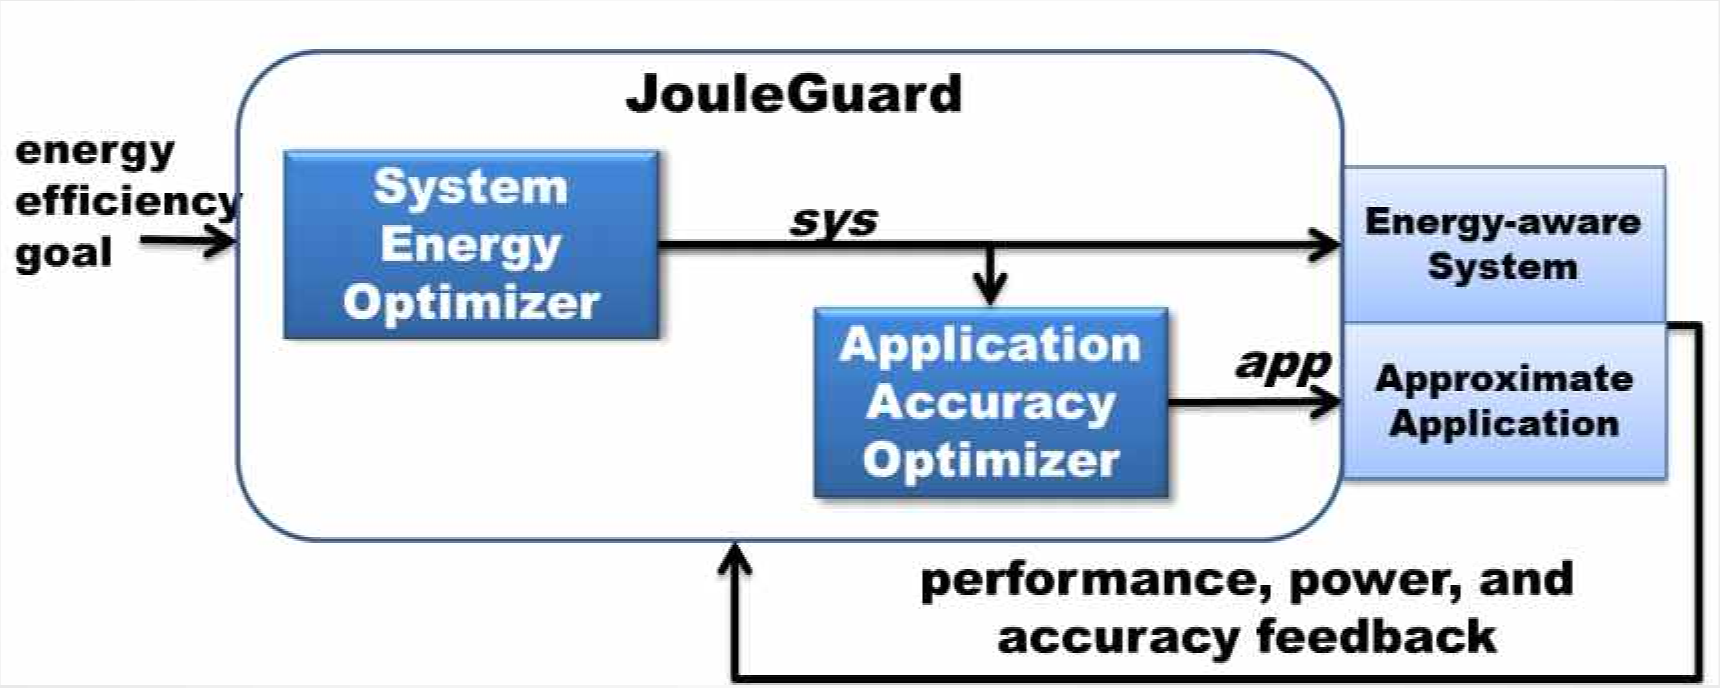

Part 1 - SEO

- Value-Difference Based Exploration
- e(t) is a measure for the difference between estimated energy efficiency and estimated energy efficiency
  - If *rand* < e(t)
    - Select a random system configuration
  - Else:
    - Select the one with highest est EE so far (?)

Part 2: Application Accuracy Opt.

- Find speedup for app. config that saves remaining energy
- Continually adjust this
  - To avoid interference slow down on high error:

$$s(t) = s(t-1) + \frac{(1 - pole(t)) \cdot error(t)}{\hat{r}_{bestsys}(t)}$$

  - Select highest accuracy that can satisfy speedup

Implementation

Application	Configs	Speedup	Acc. Loss (%)	Accuracy Metric
x264	560	4.26	6.2	Peak Signal to Noise Ratio (PSNR) [25]
swaptions	100	100.35	1.5	swaption price [25]
bodytrack	200	7.38	14.4	track quality [25]
swish++	6	1.52	83.4	precision and recall [25]
radar	26	19.39	5.3	signal to noise ratio [21]
canneal	3	1.93	7.1	wire length [56]
ferret	8	1.24	18.2	similarity [56]
streamcluster	7	5.52	0.55	quality of clustering [56]

Table 2. Approximate Application configurations.

Implementation (2)

System	Configuration	Settings	Speedup	Powerup
Mobile	clock speed	8	2.72	1.94
	big cores	4	4.52	2.00
	big core speeds	19	10.23	10.42
	LITTLE cores	4	4.52	1.32
	LITTLE core speeds	13	7.11	2.62
Tablet	clock speed	8	2.72	1.94
	core usage	2	1.81	1.22
	hyperthreading	2	1.10	1.03
	idle	n/a	1.00	1.00
Server	clock speed	16	3.23	2.05
	core usage	16	15.99	2.03
	hyperthreading	2	1.92	1.11
	idle	n/a	1.00	1.00
	mem controllers	2	1.84	1.11

Table 3. System configurations.

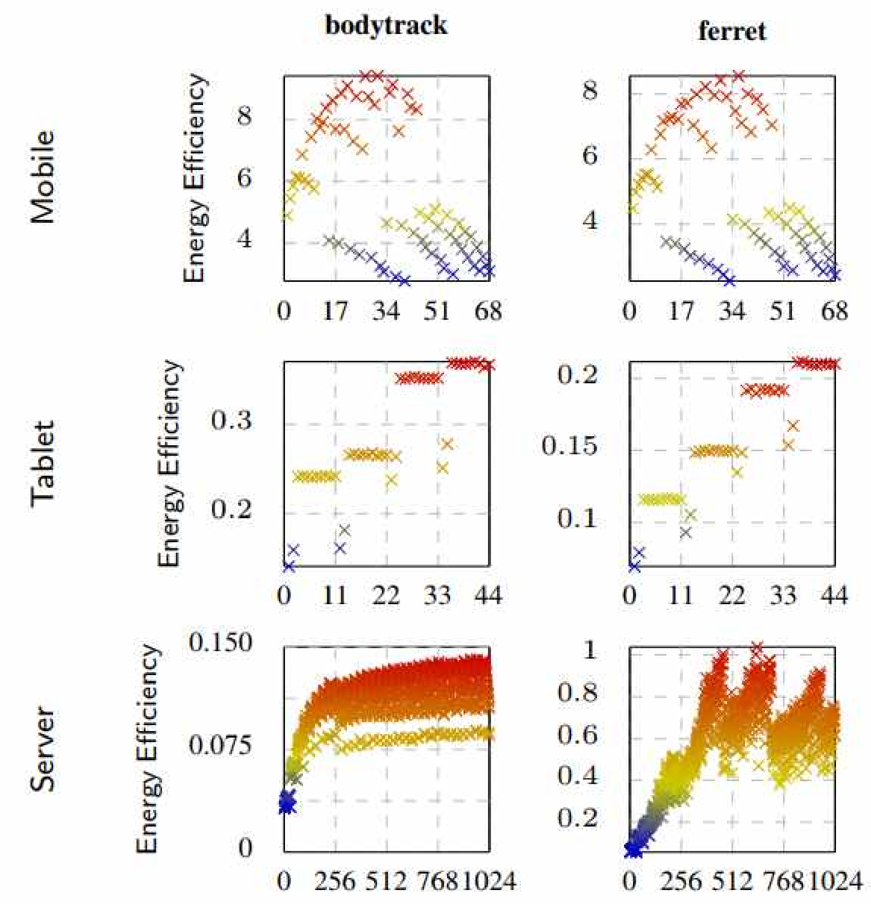

Figure 3. Example energy efficiencies.

Results

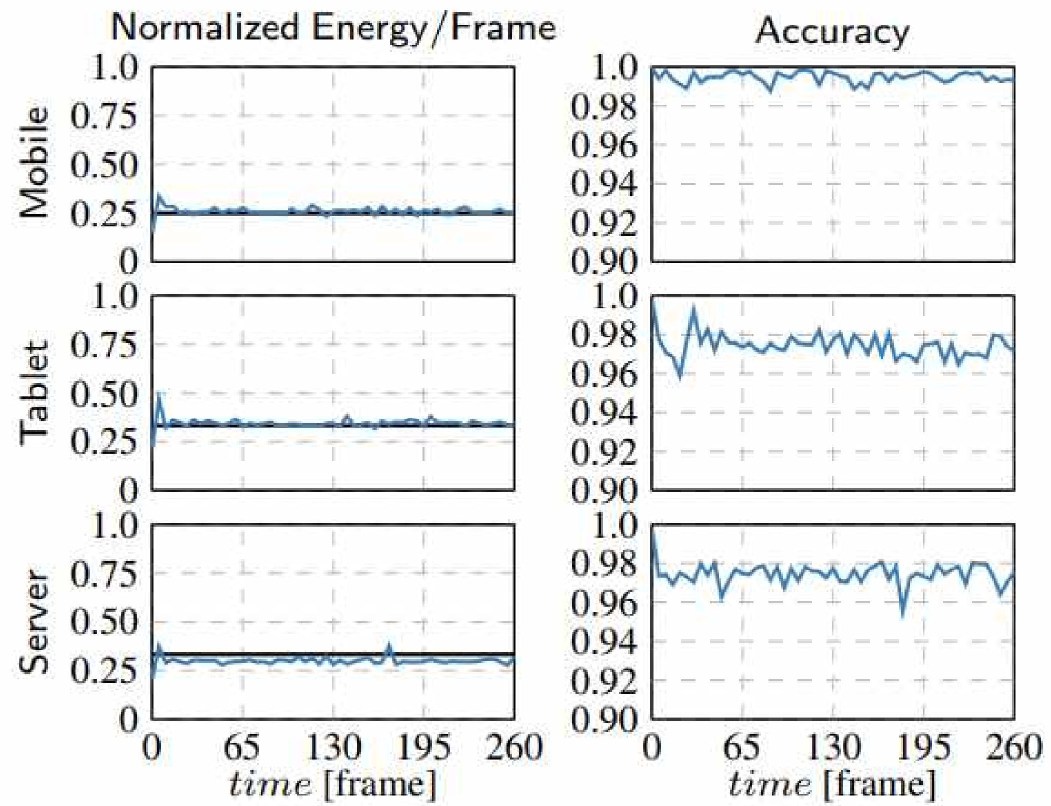

Figure 4. JouleGuard stabilizes to a consistent energy consumption subject to application and system noise.

Energy Error

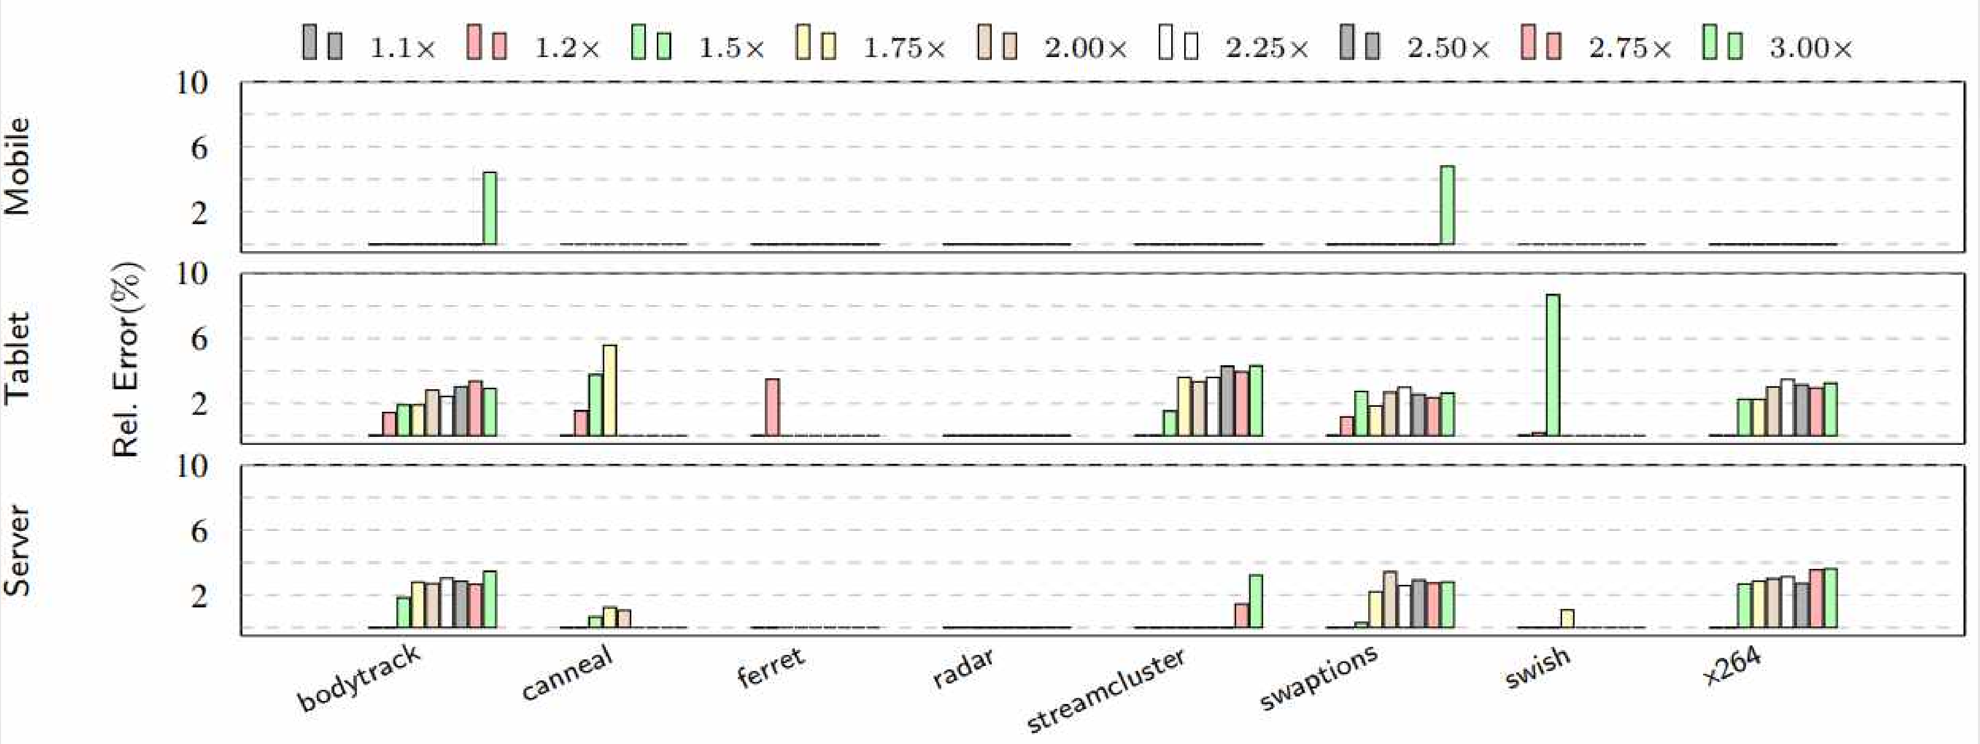

Figure 5. JouleGuard's low relative error shows it is within a few % of the desired energy.

Accuracy (vs. optimal)

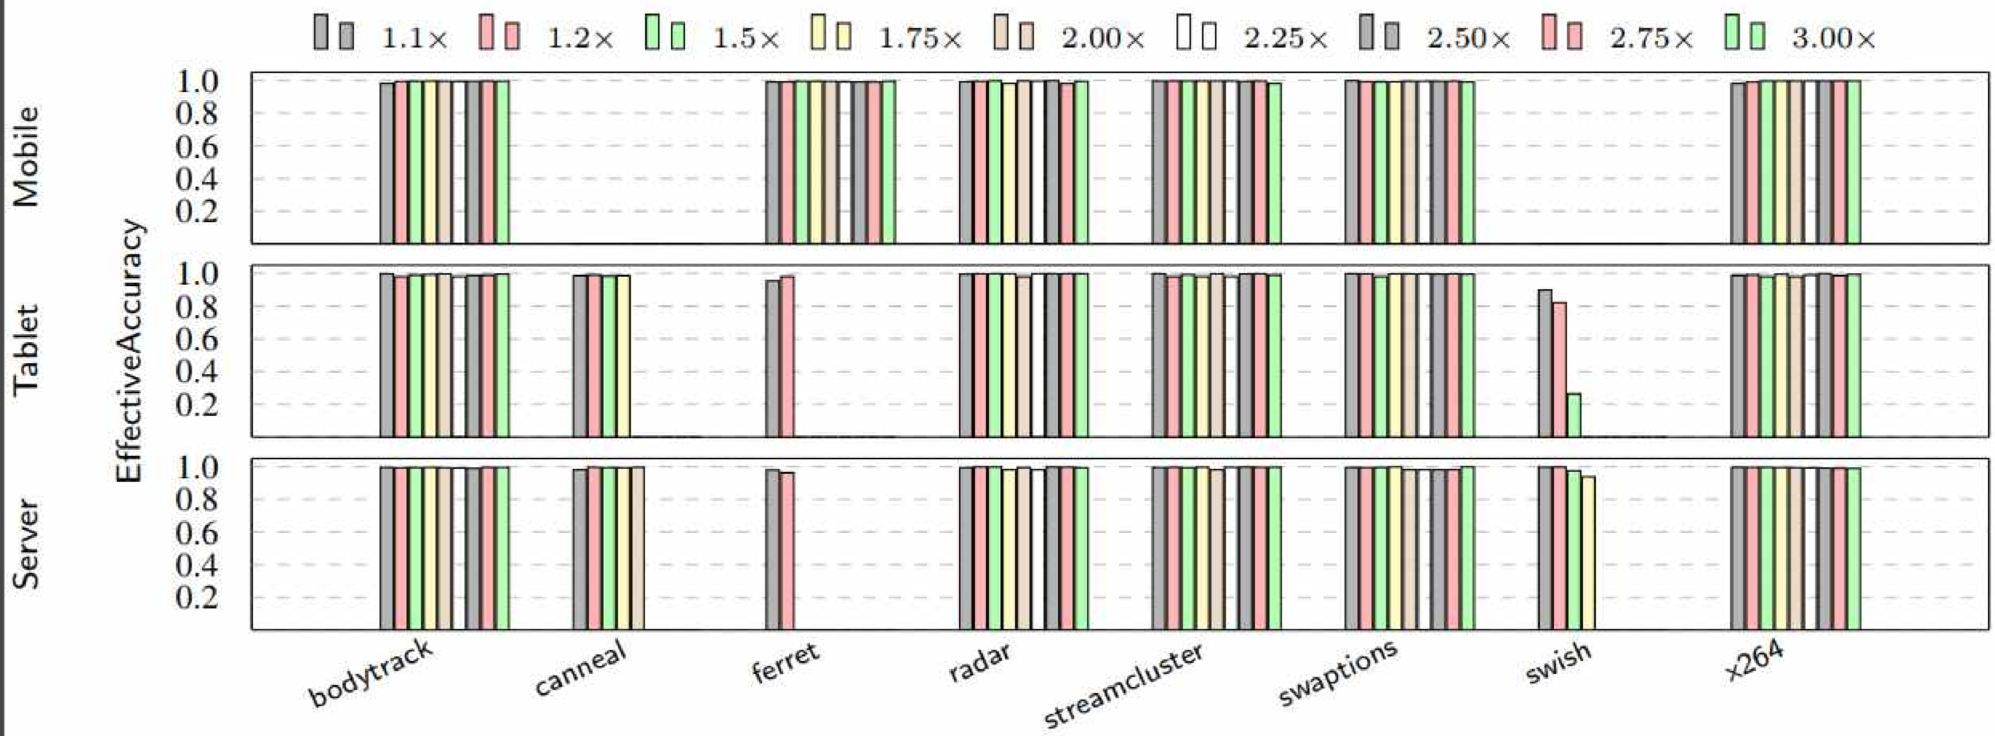

Figure 6. JouleGuard achieves near optimal accuracy for different energy goals.

JouleGuard ./. App only ./. Sys only

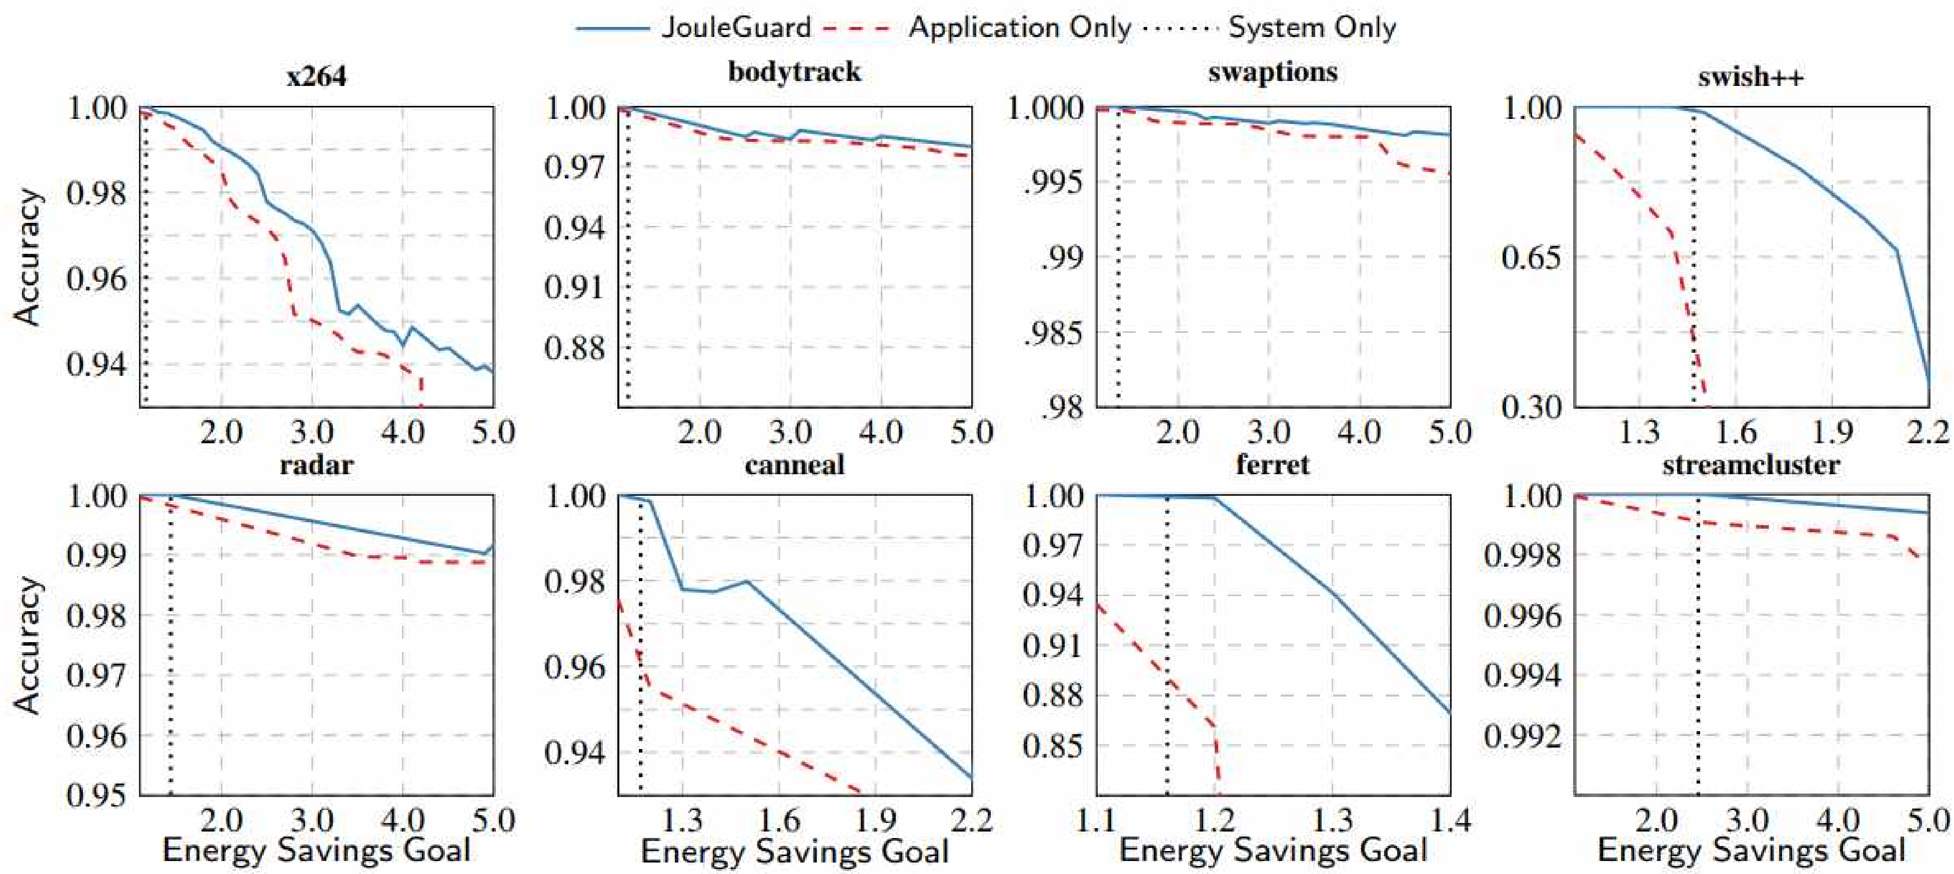

Figure 7. Comparison of JouleGuard and application- or system-only approaches on Server (higher is better).

Discussion

- Energy measurement
  - RAPL / built in sensors ignore large parts of system
- Still unsure about the SEO approach
- This assumes that approximation is only achieved through time. How about double/float?