Slow Down or Sleep?
That is the Question

NICTA

Etienne Le Sueur and Gernot Heiser

USENIX ATC, June 2011
elesueur@nicta.com.au

Australian Government
Department of Broadband, Communications and the Digital Economy
Australian Research Council

NICTA Funding and Supporting Members and Partners

Energy management is important

Dynamic voltage and frequency scaling (DVFS)

NICTA

$$P_{processor} = \alpha C f V^2 + I_{leakage} V$$

- Reduced effectiveness
  - Increasing static power
  - Smaller voltage ranges
  - Better memory performance makes single threaded workloads less memory-bound, thus DVFS affects their performance more.

Low power idle states

- Idle states are becoming **more** effective
  - Processors now have numerous idle states.
  - They have low power consumption, and
  - Low transition overhead and latency.

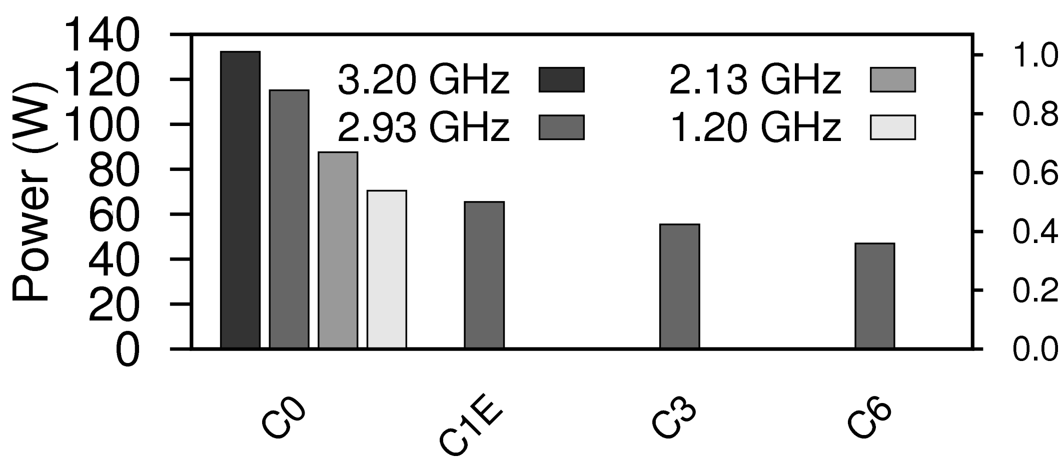

Power draw, non-uniform C-states

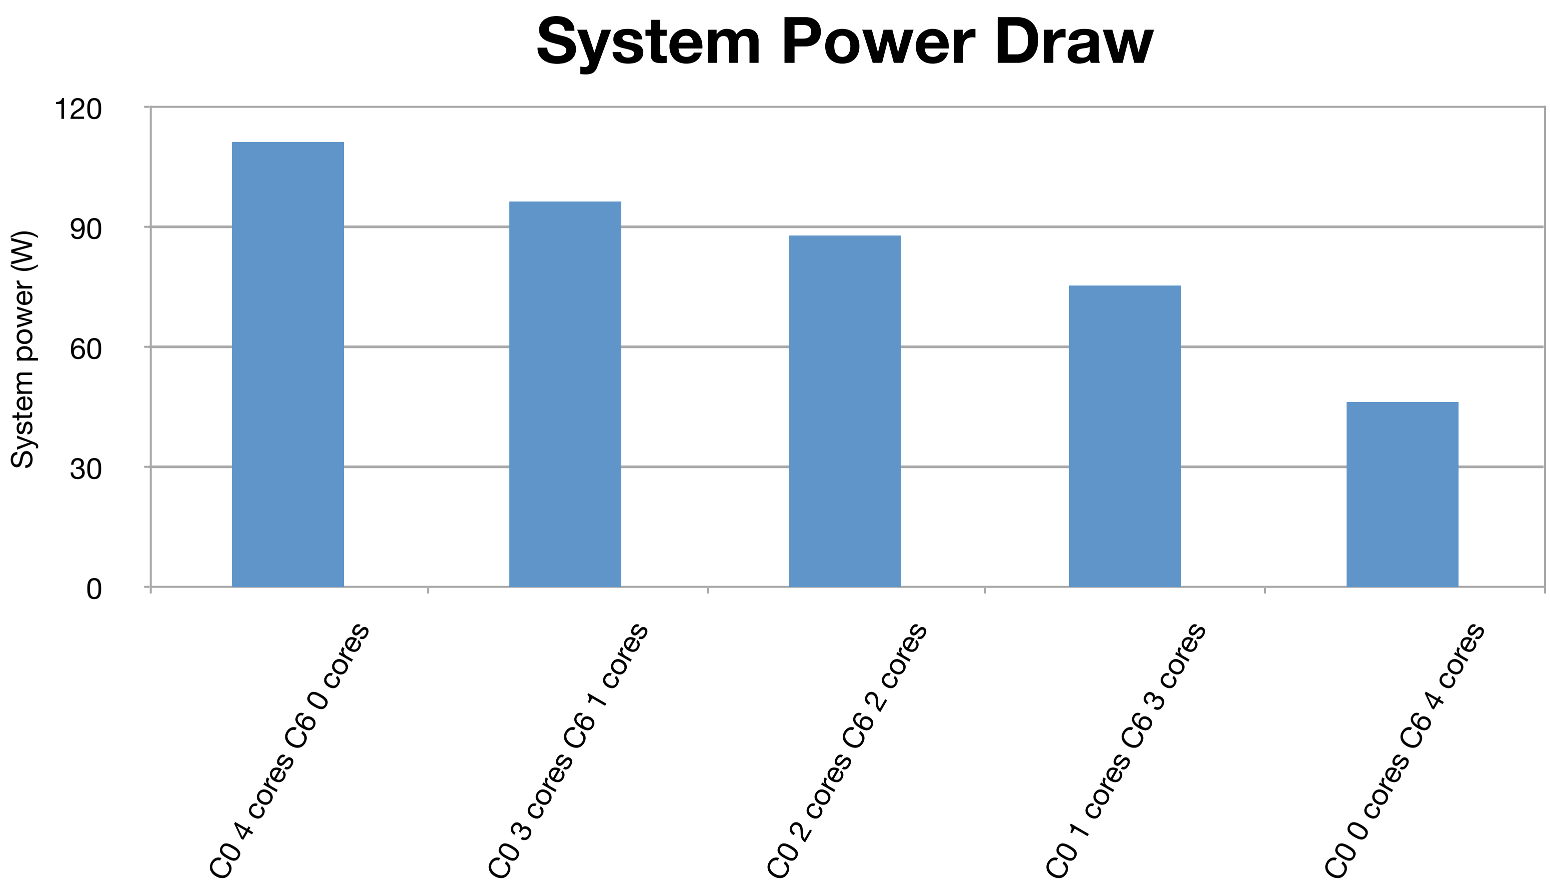

A trade-off is presented...

Slow down or **go to sleep?**

How do we optimise energy?

NICTA

- Slowing down
  - Reduces CPU *power* draw, but **only** the CPU.
  - Increases execution time i.e. reduced performance.
  - Energy consumption is **not** necessarily reduced.
  - Reduces time spent in idle states.
- Going to sleep
  - Keep the processor running fast.
  - Execution time is not significantly affected.
  - Maximises time spent in idle states.
  - Energy consumption depends on idle states used.

Benchmarks

- Real world - server, desktop, mobile
- Bursty
  - Allows the CPU to idle/sleep

Benchmark method

NICTA

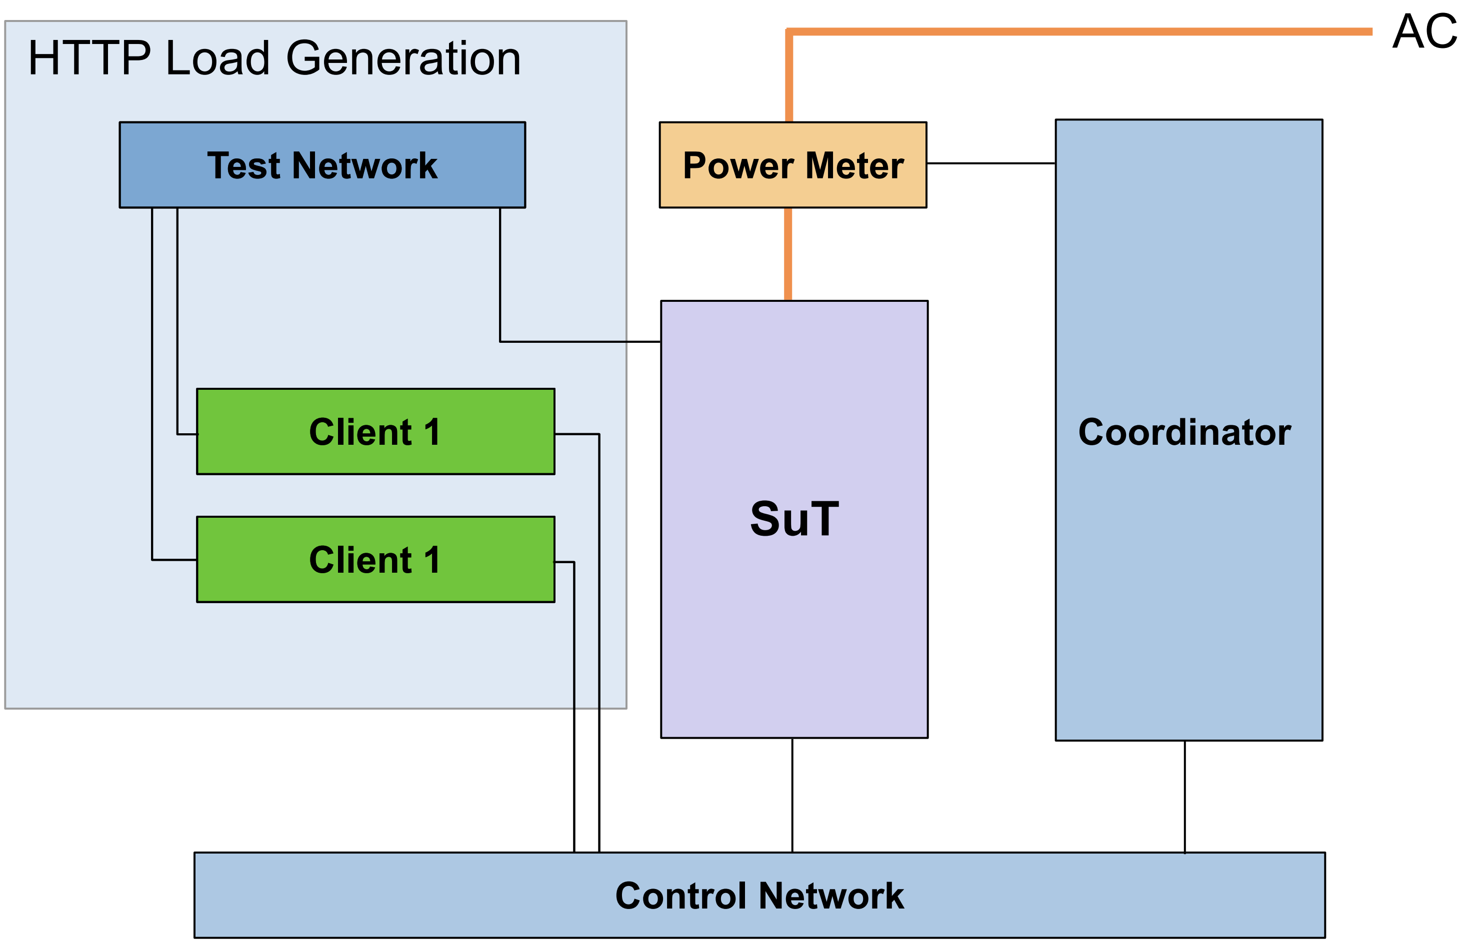

MPEG decode/playback

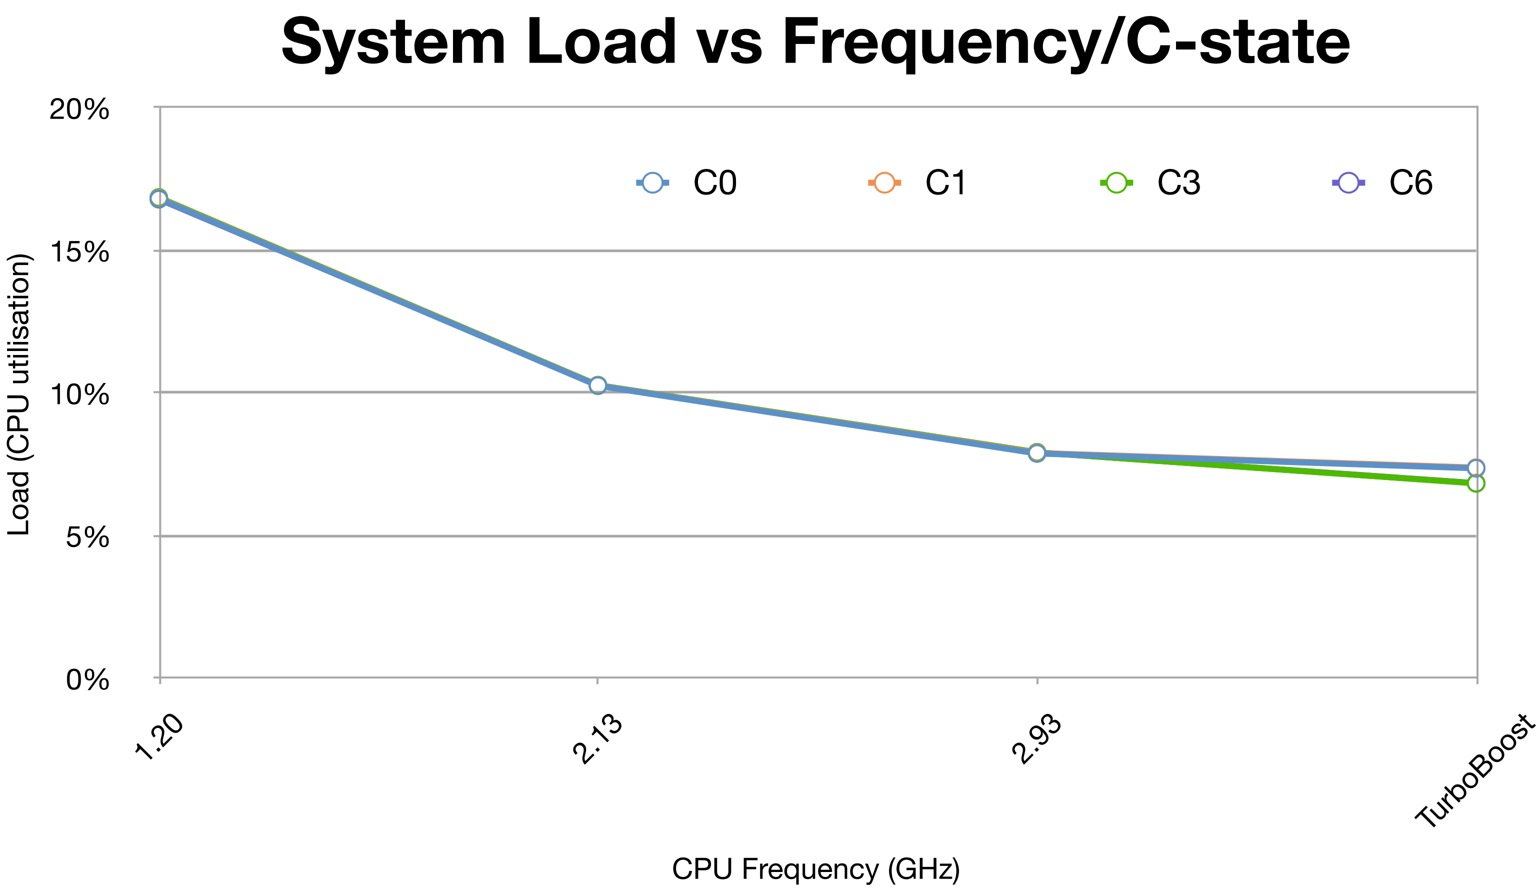

System Energy vs Frequency/C-state

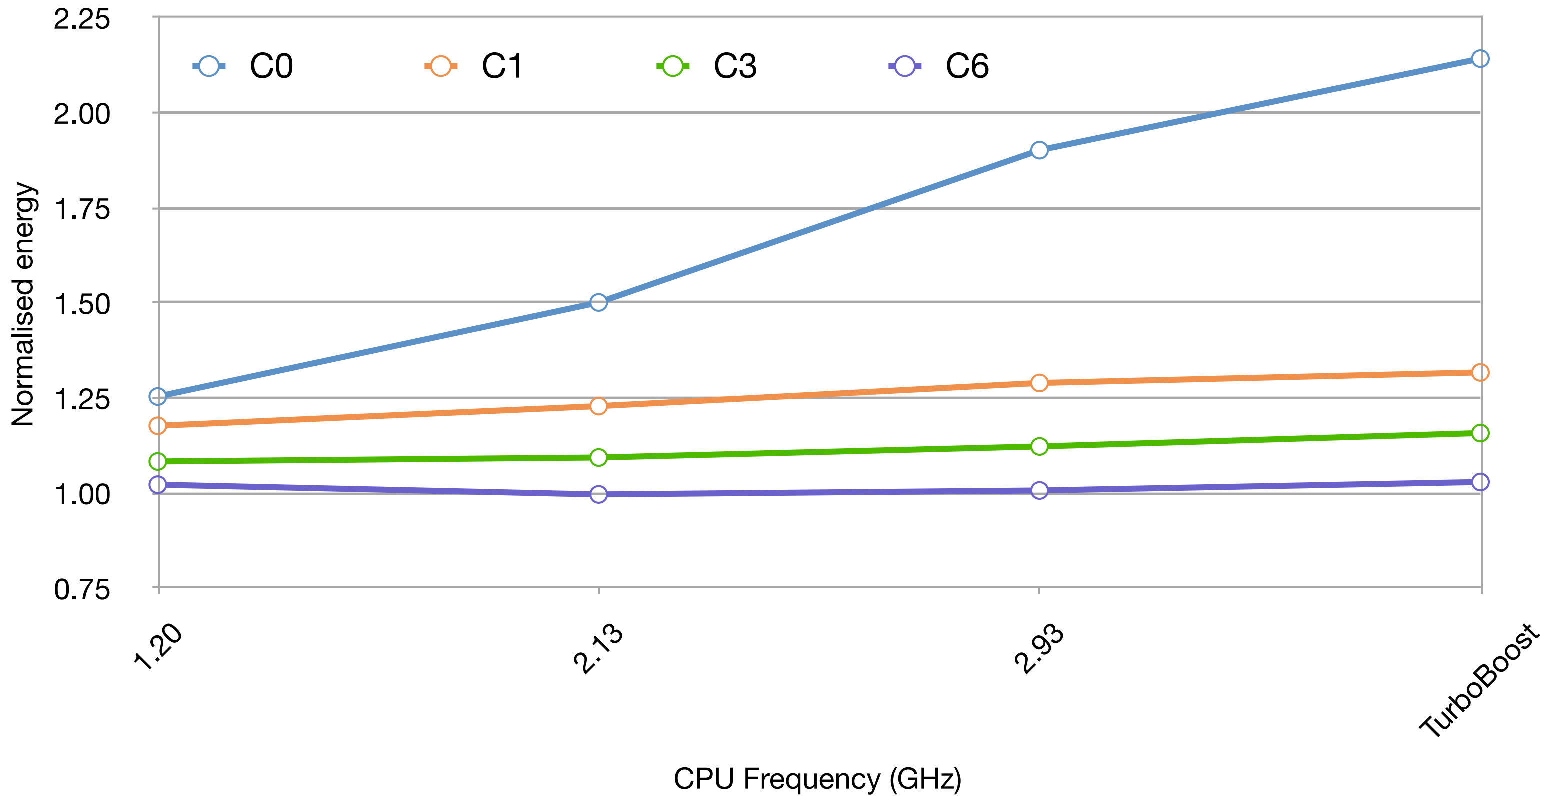

Server load provisioning

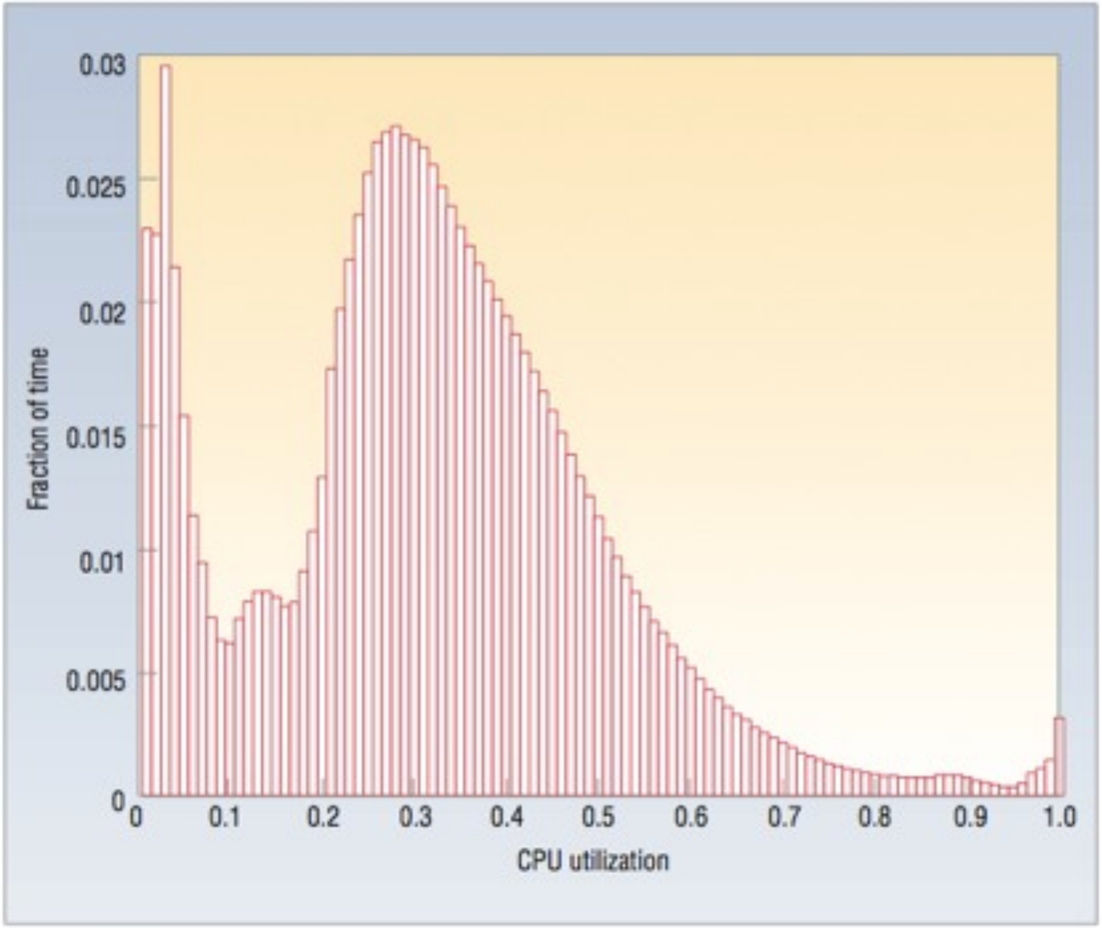

L. A. Barroso and U. Hölzle. The case for energy-proportional computing. IEEE Computer, 40(12): 33–37, Dec. 2007.

Load range

NICTA

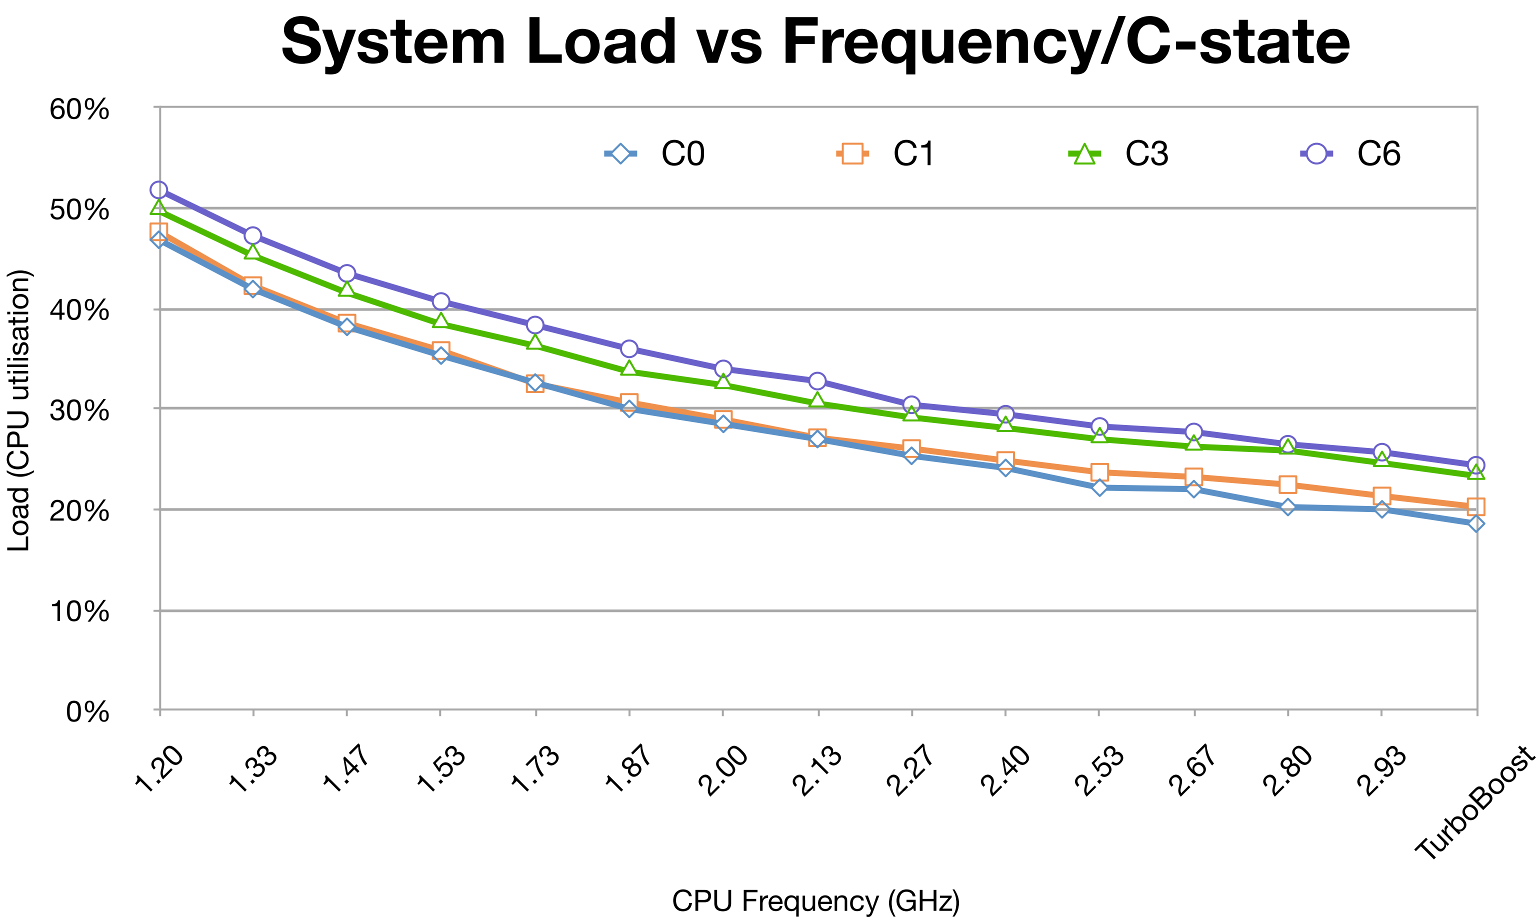

System energy

NICTA

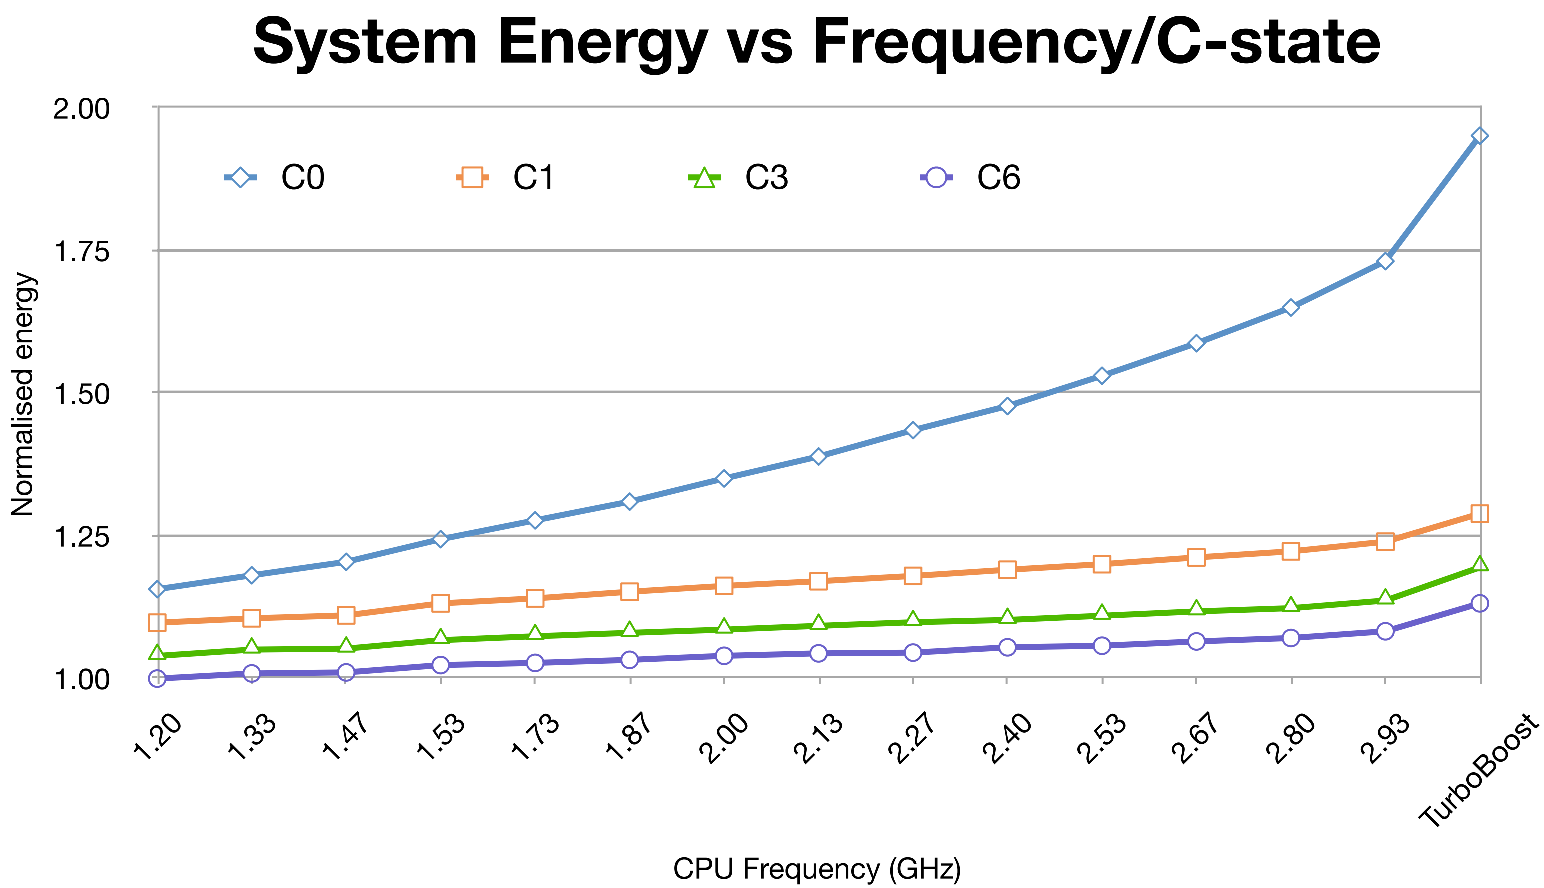

Throughput

NICTA

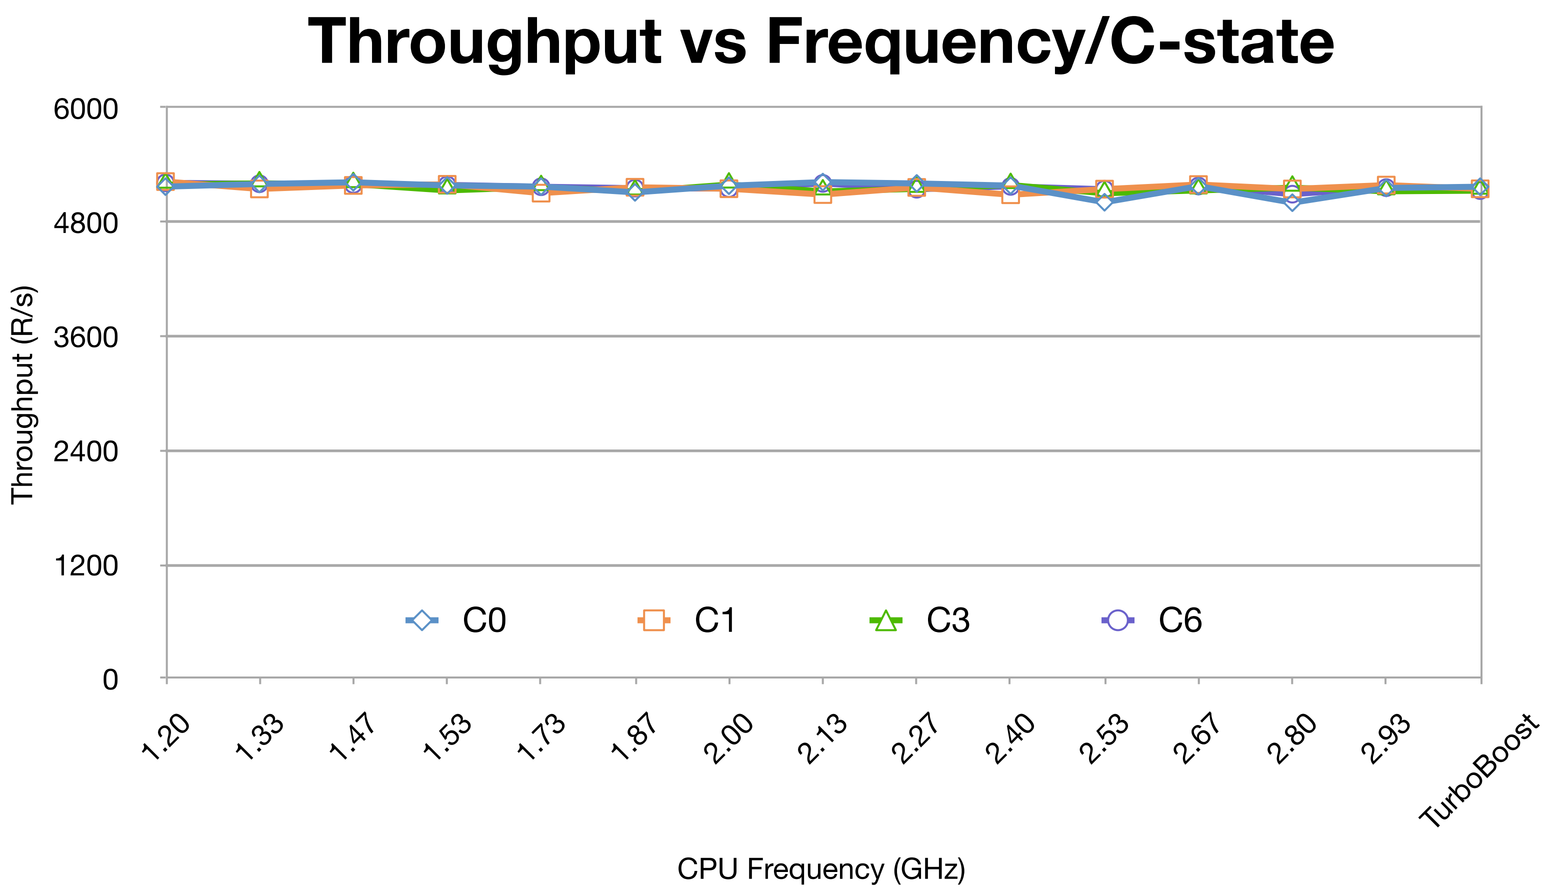

Request efficiency

NICTA

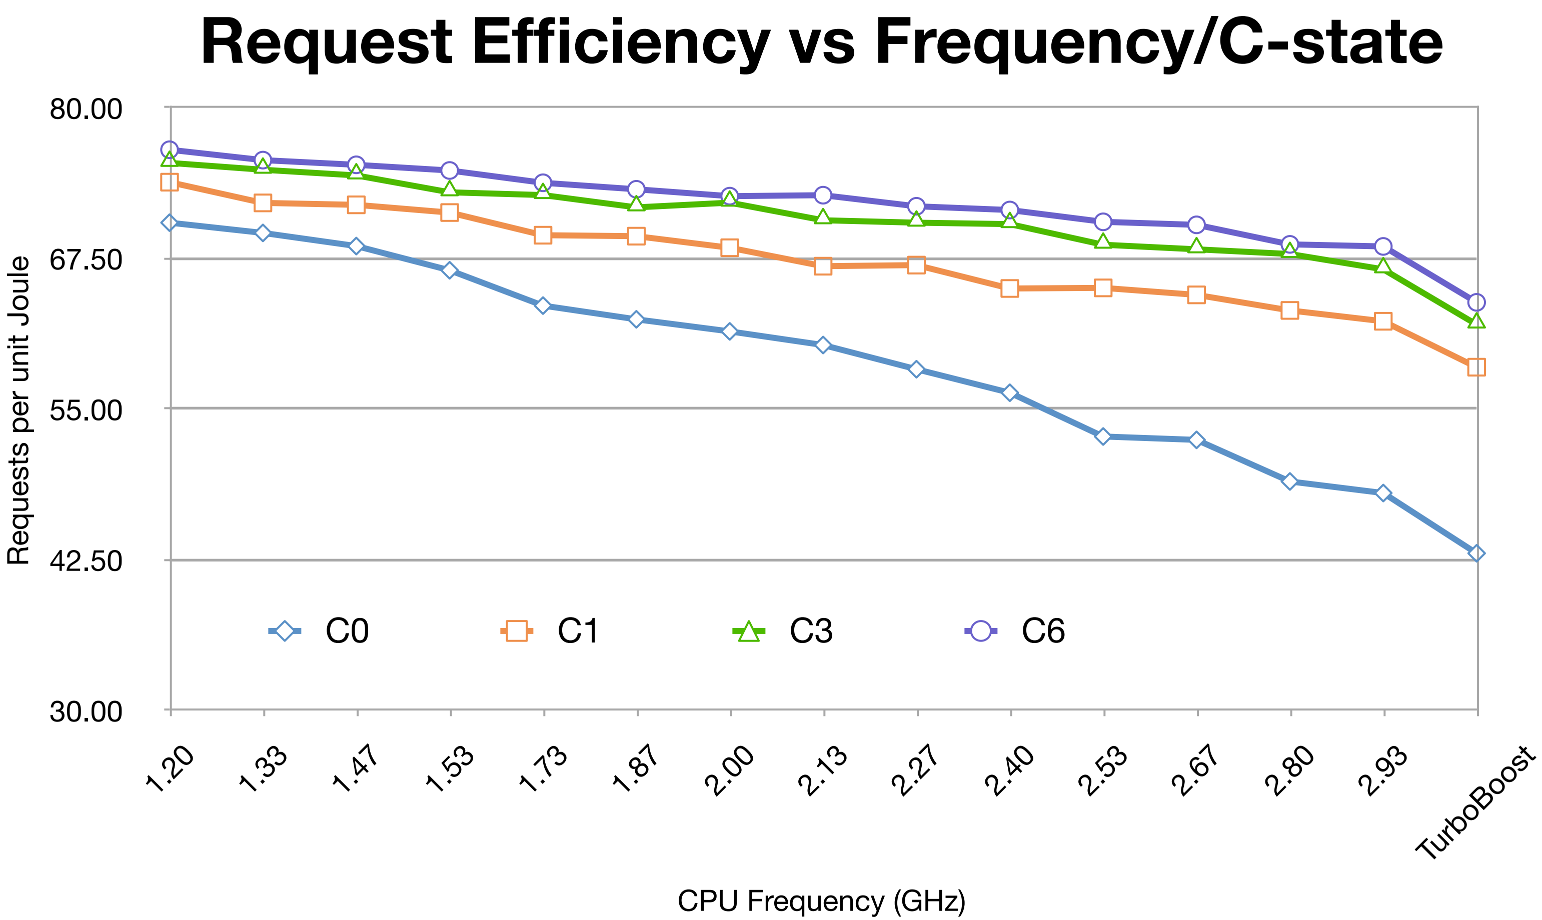

Response time

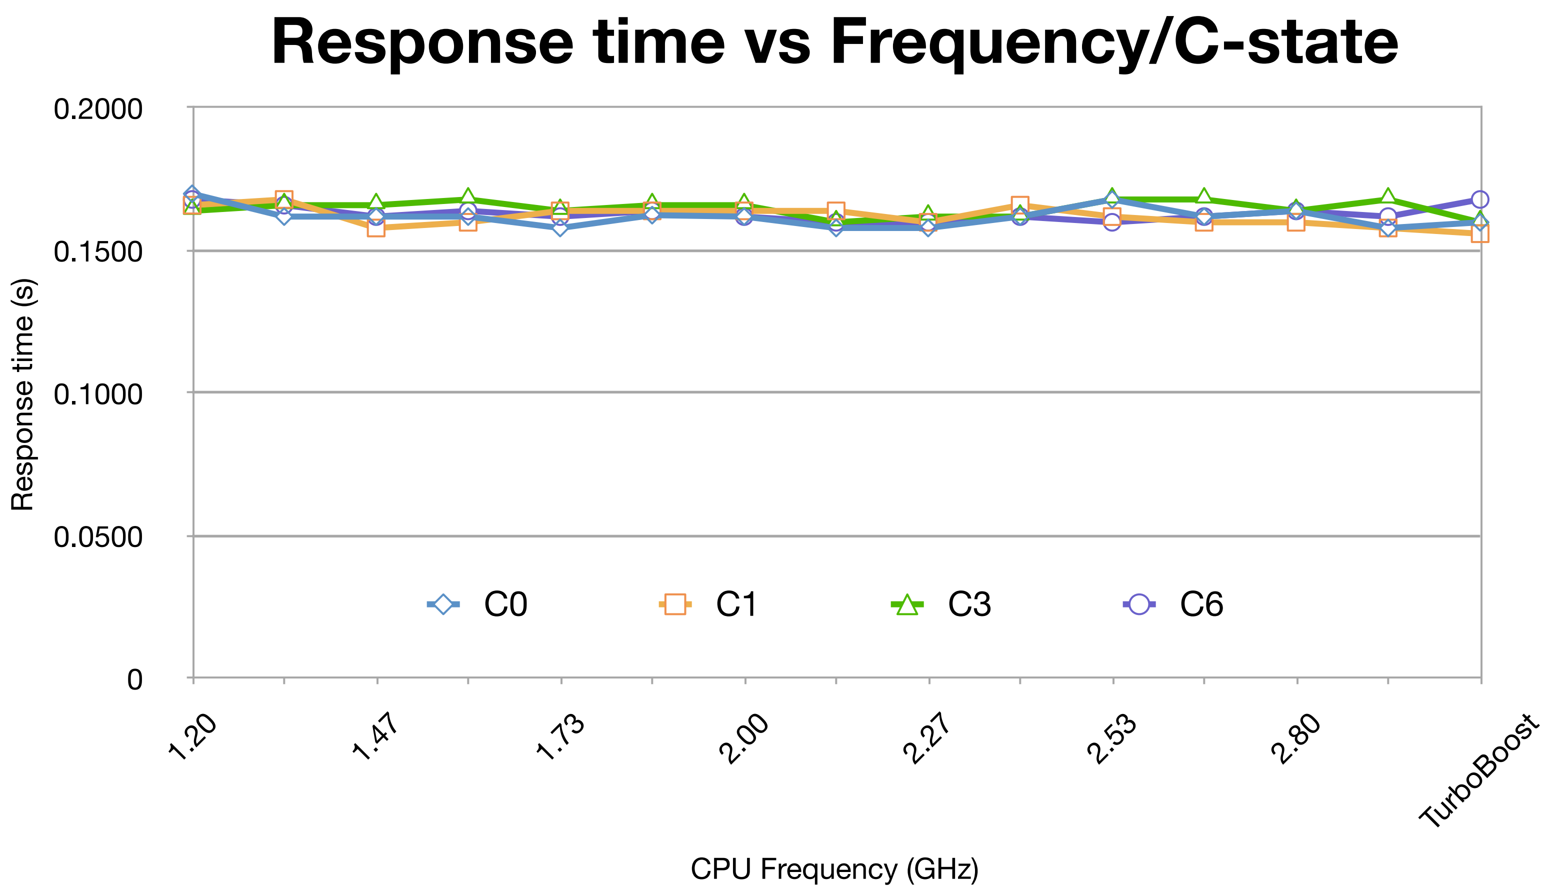

Comparison with embedded platforms

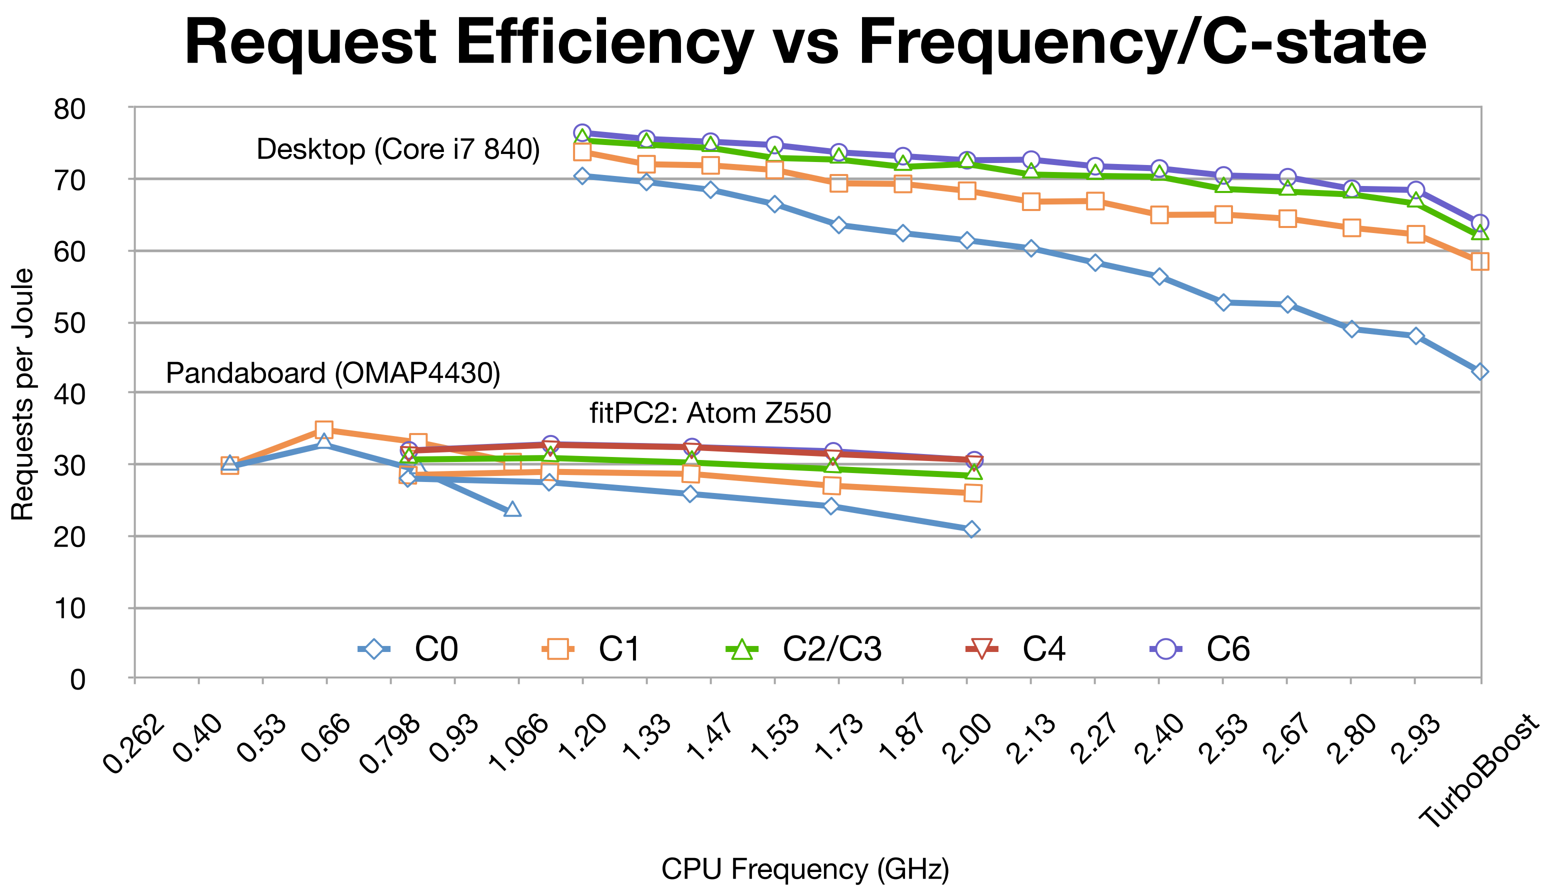

Take homes

NICTA

- DVFS of **cores** is becoming less effective with better C-states.
- Lightly loaded systems lose no throughput or latency with reduced frequency and deep C-states.
- High C-state transition rates increase the overhead from deep C-states.
- The energy-efficiency of a simple HTTP server is improved at lower frequencies with the current state-of-the-art C-state.
- Energy-per-request is minimised on the desktop-class, not the embedded-class.
- TurboBoost is **not** energy-efficient.

NICTA

Conclusions

- CPU power management is changing direction.
  - Maybe we need more control (L3 caches, memory/memory controller, QPI/HT links, etc)
- Minimising idle and static power losses is much more important for deep sub-micron processors.
  - Critical for mobile embedded devices which spend most of their time asleep.

Questions?

NICTA

NICTA

From **imagination** to **impact**