The Energy/Frequency Convexity Rule of Energy Consumption for Programs:

Modeling, Thermosensitivity, and Applications

Karel De Vogeleer

Ph.D. defense

September 4th, 2015

Special thanks to **Fondation TELECOM** for funding this research


PARIS 2050
VINCENT CALLEBAUT
ARCHITECTURES


PHONE BATTERIES
Y U NO LAST LONGER
MEMEFUL.COM

A Green IT Thinking

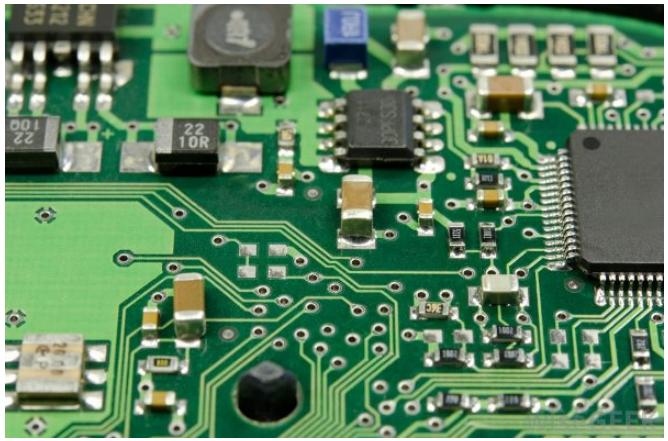

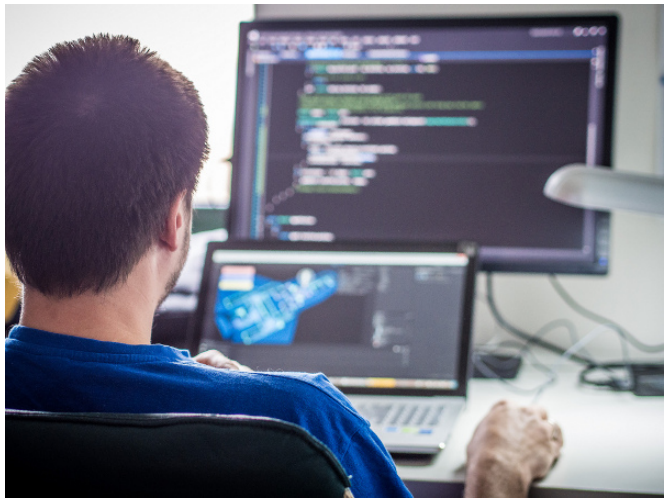

- Off-line, including
  - transistor design,
  - circuit design,
  - architecture,
  - software design,
  - software coding,
  - compiler optimization;
- on-line, including
  - system reconfiguration,
  - compiler optimization,
  - context placement.

Image source jiji.ng and wisegeek.com

Contributions

- Energy consumption analysis for computer systems:
  - analytical model,
  - Energy/Frequency Convexity Rule,
  - supportive measurement data;
- temperature/power relationship demystified:
  - supportive measurement data,
  - guidelines for power measurement;
- transient thermal model for microprocessors:
  - analytical model including radiation,
  - approximations,
  - applicability analysis.

Contributions

- Energy consumption analysis for computer systems:
  - analytical model,
  - Energy/Frequency Convexity Rule,
  - supportive measurement data;
- temperature/power relationship der
  - supportive measurement data,
  - guidelines for power measurement
- transient thermal model for micropı
  - analytical model including radiati
  - approximations,
  - applicability analysis.

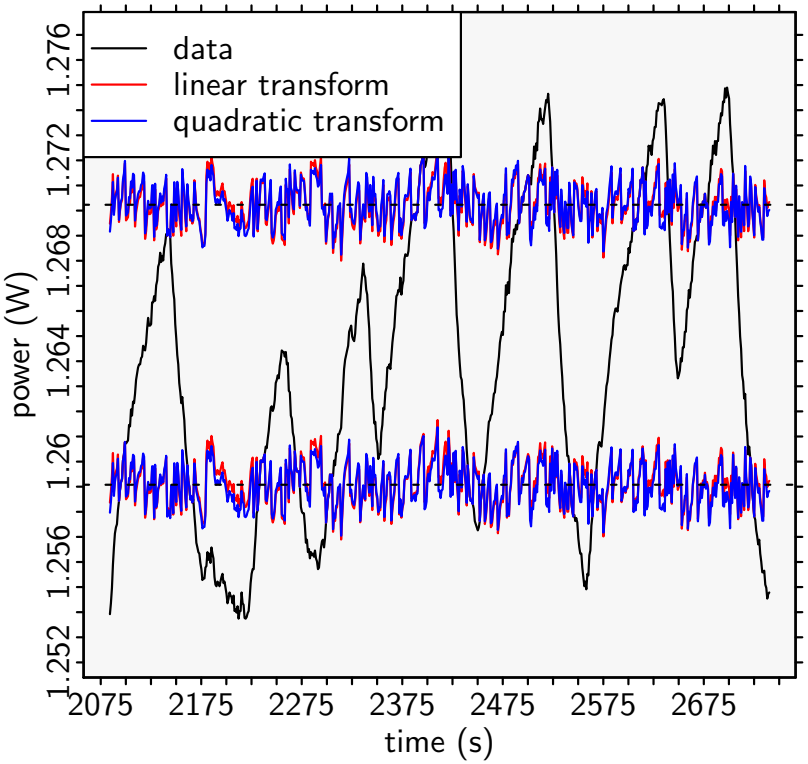

Contributions

- Energy consumption analysis for computer systems:
  - analytical model,
  - Energy/Frequency Convexity Rule,
  - supportive measurement data;
- temperature/power relationship demystified:
  - supportive measurement data,
  - guidelines for power measurement;
- transient thermal model for microprocessors:
  - analytical model including radiation,
  - approximations,
  - applicability analysis.

Contributions

- Energy consumption analysis for computer systems:
  - analytical model,
  - Energy/Frequency Convexity Rule,
  - supportive measurement data;
- temperature/power relationship der
  - supportive measurement data,
  - guidelines for power measurement
- transient thermal model for micropı
  - analytical model including radiati
  - approximations,
  - applicability analysis.

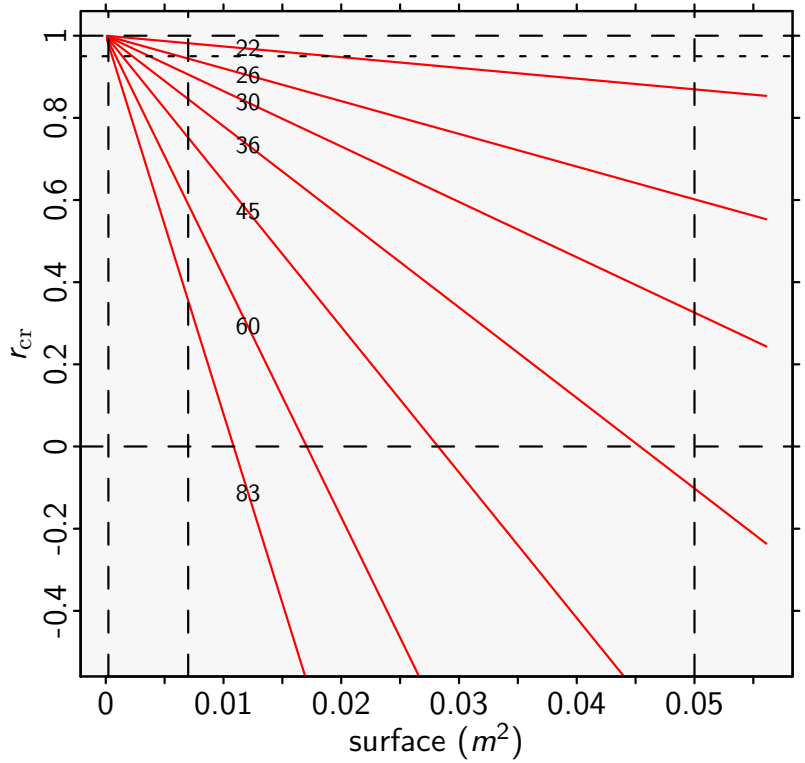

Contributions

- energy consumption analysis for computer systems:
  - analytical model,
  - Energy/Frequency Convexity Rule,
  - supportive measurement data;
- temperature/power relationship demystified:
  - supportive measurement data,
  - guidelines for power measurement;
- transient thermal model for microprocessors:
  - analytical model,
  - approximations,
  - applicability analysis.

Presentation's Outline

Preliminary Evidence of Energy/Frequency Curves

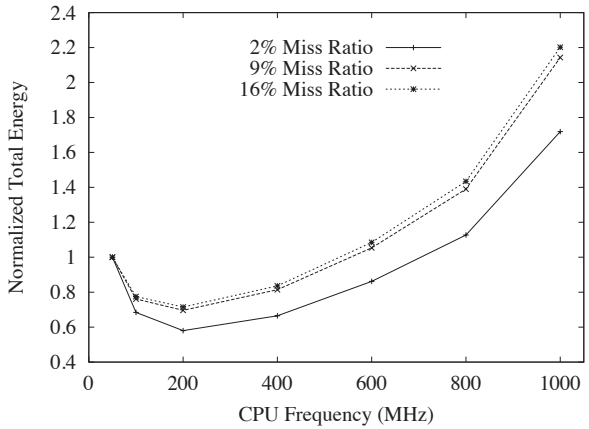

(a) Fan et al. [1]

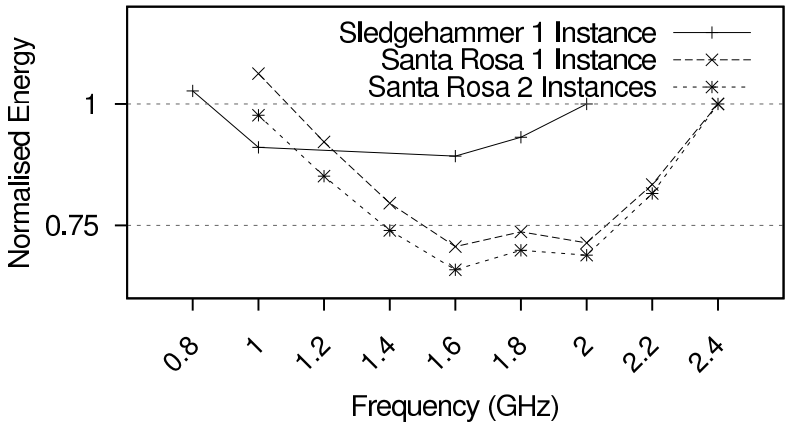

(b) Le Sueur and Heiser [3]

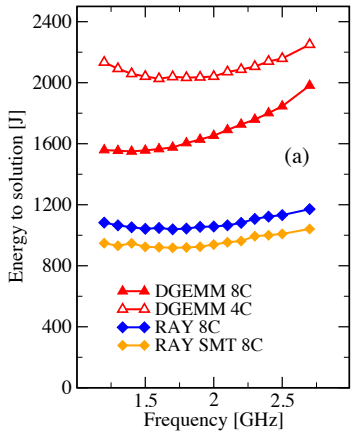

(c) Hager et al. [2]

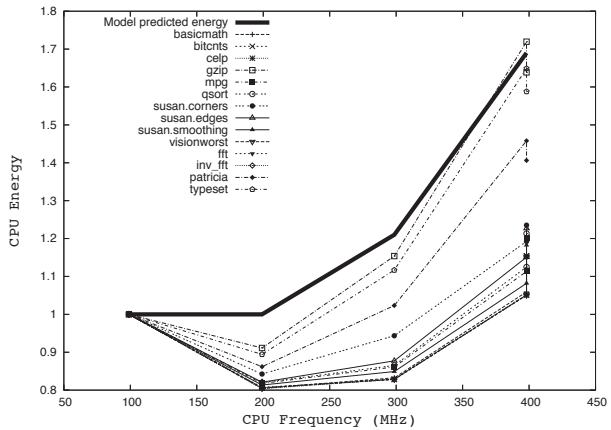

(d) Snowdon et al. [4]

System Energy Consumption Model (E_{sys})

- System's energy consumption E_{sys} definition

$$E_{\text{sys}} = \int_0^{\Delta t} P_{\text{sys}}\, dt$$

$$= \int_0^{\Delta t} (P_{\text{cpu}} + P_{\text{back}})\, dt;$$

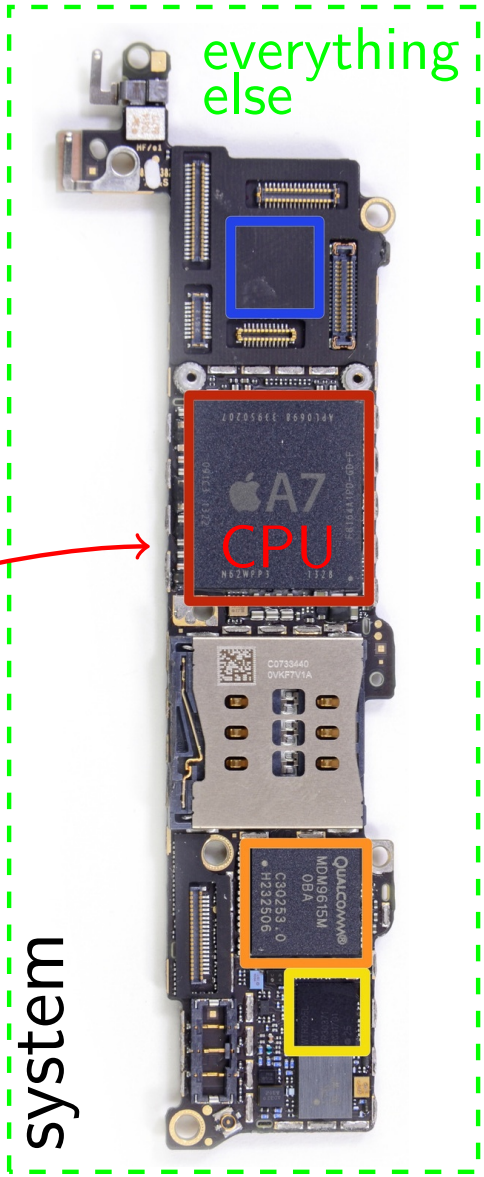

- Examples of P_{back} include:
  - LCD screen,
  - radio interface,
  - sensors (e.g. GPS);
- If P_{cpu} and P_{back} are constant over Δt:

$$E_{\text{sys}} = (P_{\text{cpu}} + P_{\text{back}}) \cdot \Delta t.$$

Microprocessor Power Model

CPU power P_{cpu} consists of:

- dynamic power P_{dyn},
- leakage current P_{leak},
- short-circuit current P_{sc},

$$\begin{aligned} P_{\text{cpu}} &= P_{\text{dyn}} + P_{\text{leak}} + P_{\text{sc}} \\ &= (1 + \gamma V) \cdot \eta \alpha C V^2 f \\ &= (1 + \gamma V) \cdot \xi V^2 f. \end{aligned}$$

Execution Time Model

Execution time Δt depends on:

- cc_{b} code size in clock cycles,
- f CPU clock frequency,
- f_{k} frequency thieves,
- β slack time per clock cycle,

$$\Delta t = cc_{\text{b}} \left(\frac{1}{f - f_{\text{k}}} + \beta \right).$$

Optimal Clock Frequency (f_{opt})

- System's energy consumption model

$$\begin{aligned} E_{\text{sys}}(f) &= (P_{\text{cpu}} + P_{\text{back}}) \cdot \Delta t \\ &= ([1 + \gamma V]\xi V^2 f + P_{\text{back}}) \cdot cc_{\text{b}} \left(\frac{1}{f - f_{\text{k}}} + \beta \right), \end{aligned}$$

  where $\{\gamma, \xi, P_{\text{back}}, cc_{\text{b}}, f_{\text{k}}, \beta\} \in \mathbb{R}^+$;

- a single minimum for $E_{\text{sys}}(f)$ exists at f_{opt} when

$$\left(\frac{\partial E_{\text{sys}}}{\partial f} \right)_{f = f_{\text{opt}}} = 0, \quad \text{and} \quad \frac{\partial^2 E_{\text{sys}}}{\partial f^2} > 0 \quad \text{holds;}$$

- V is approximately an affine map of f: $V \rightarrow m_2 f + m_1$.

Supply Voltage/Frequency Relationship

A linear trend between V and f is observed: $V = m_2 f + m_1$.

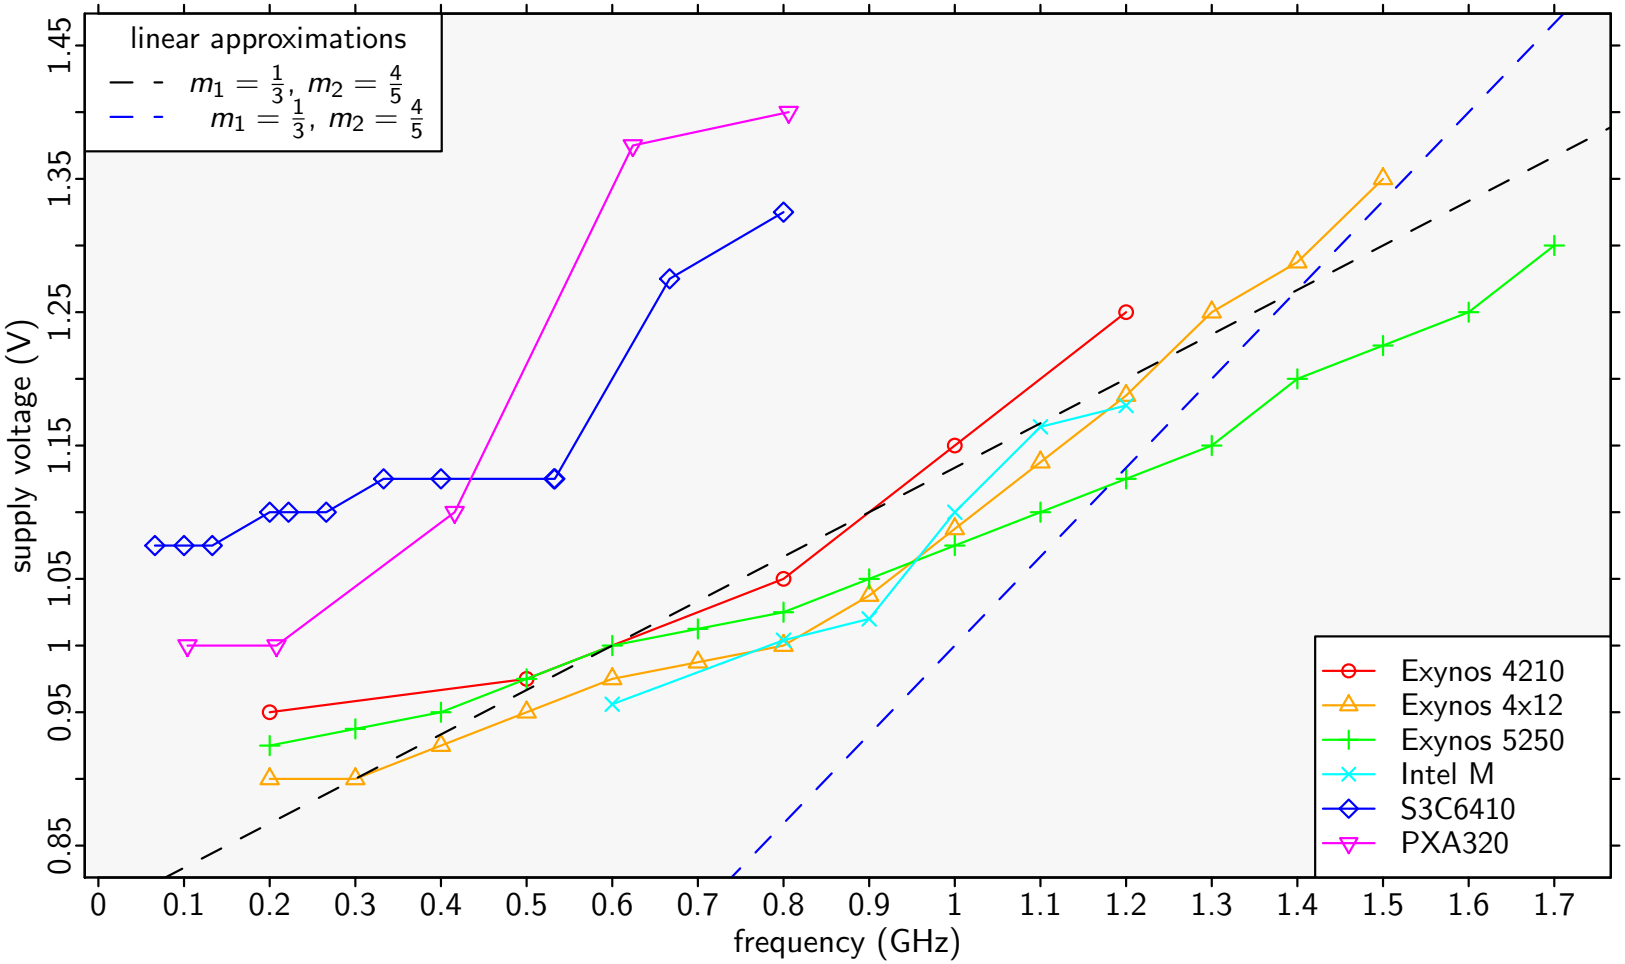

Benchmark and Testbed

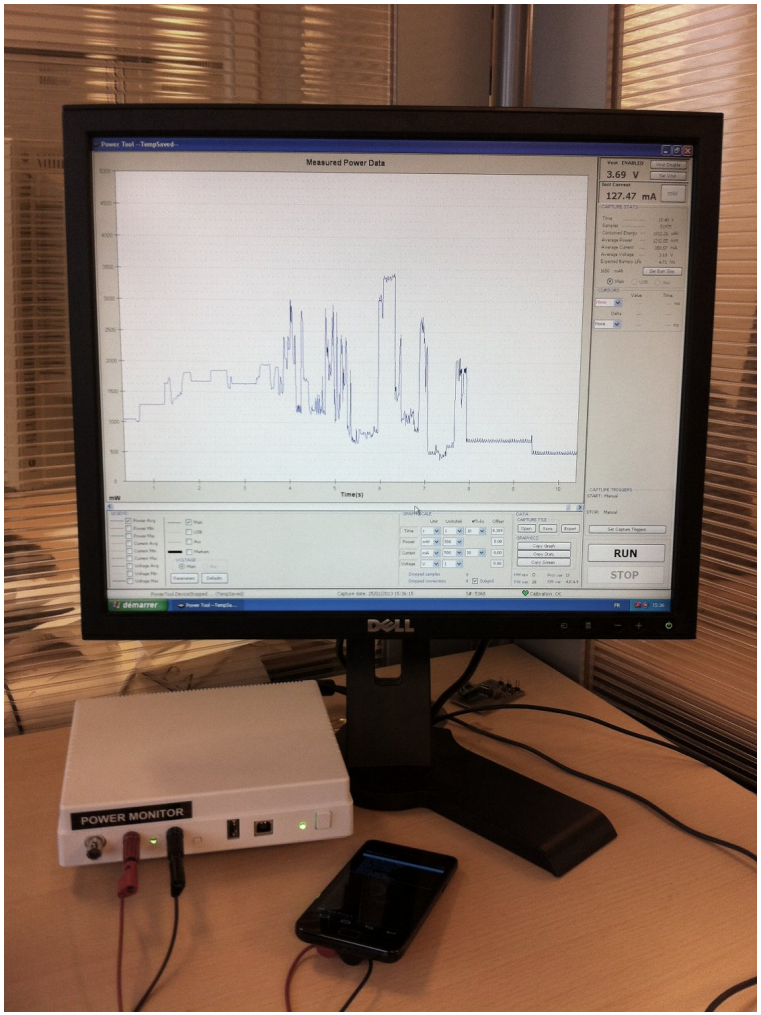

- **Benchmark**: bit-reverse algorithm, part of the DFT algorithm:

```
void bitreverse_gold_rader
     (int N, complex *data) {
  int n = N, nm1 = n-1;
  int i = 0, j = 0;
  for (; i < nm1; i++) {
    int k = n >> 1;
    if (i < j) {
      complex temp = data[i];
      data[i] = data[j];
      data[j] = temp;}
    while (k <= j) {
      j -= k; k >>= 1;}
    j += k ;
  }
}
```

- **testbed**: Samsung Galaxy SII;
- **power Measurement**: Monsoon.

Benchmark and Testbed

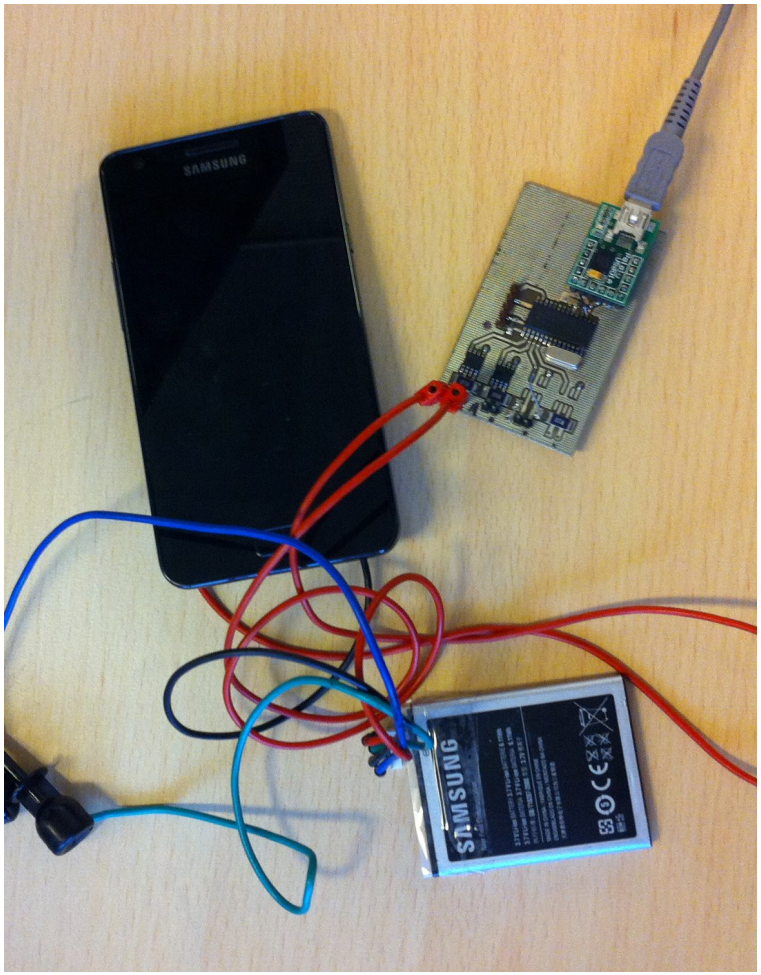

- **Benchmark**: bit-reverse algorithm, part of the DFT algorithm:

```
void bitreverse_gold_rader
    (int N, complex *data) {
  int n = N, nm1 = n-1;
  int i = 0, j = 0;
  for (; i < nm1; i++) {
    int k = n >> 1;
    if (i < j) {
      complex temp = data[i];
      data[i] = data[j];
      data[j] = temp;}
    while (k <= j) {
      j -= k; k >>= 1;}
    j += k ;
  }
}
```

- **testbed**: Samsung Galaxy SII;
- **power Measurement**: Monsoon.

Benchmark and Testbed

- **Benchmark**: bit-reverse algorithm, part of the DFT algorithm:

```
void bitreverse_gold_rader
    (int N, complex *data) {
  int n = N, nm1 = n-1;
  int i = 0, j = 0;
  for (; i < nm1; i++) {
    int k = n >> 1;
    if (i < j) {
      complex temp = data[i];
      data[i] = data[j];
      data[j] = temp;}
    while (k <= j) {
      j -= k; k >>= 1;}
    j += k ;
  }
}
```

- **testbed**: Samsung Galaxy SII;
- **power Measurement**: Monsoon.

The Energy/Frequency Convexity Rule

Energy consumption versus CPU clock frequency shows convex properties.

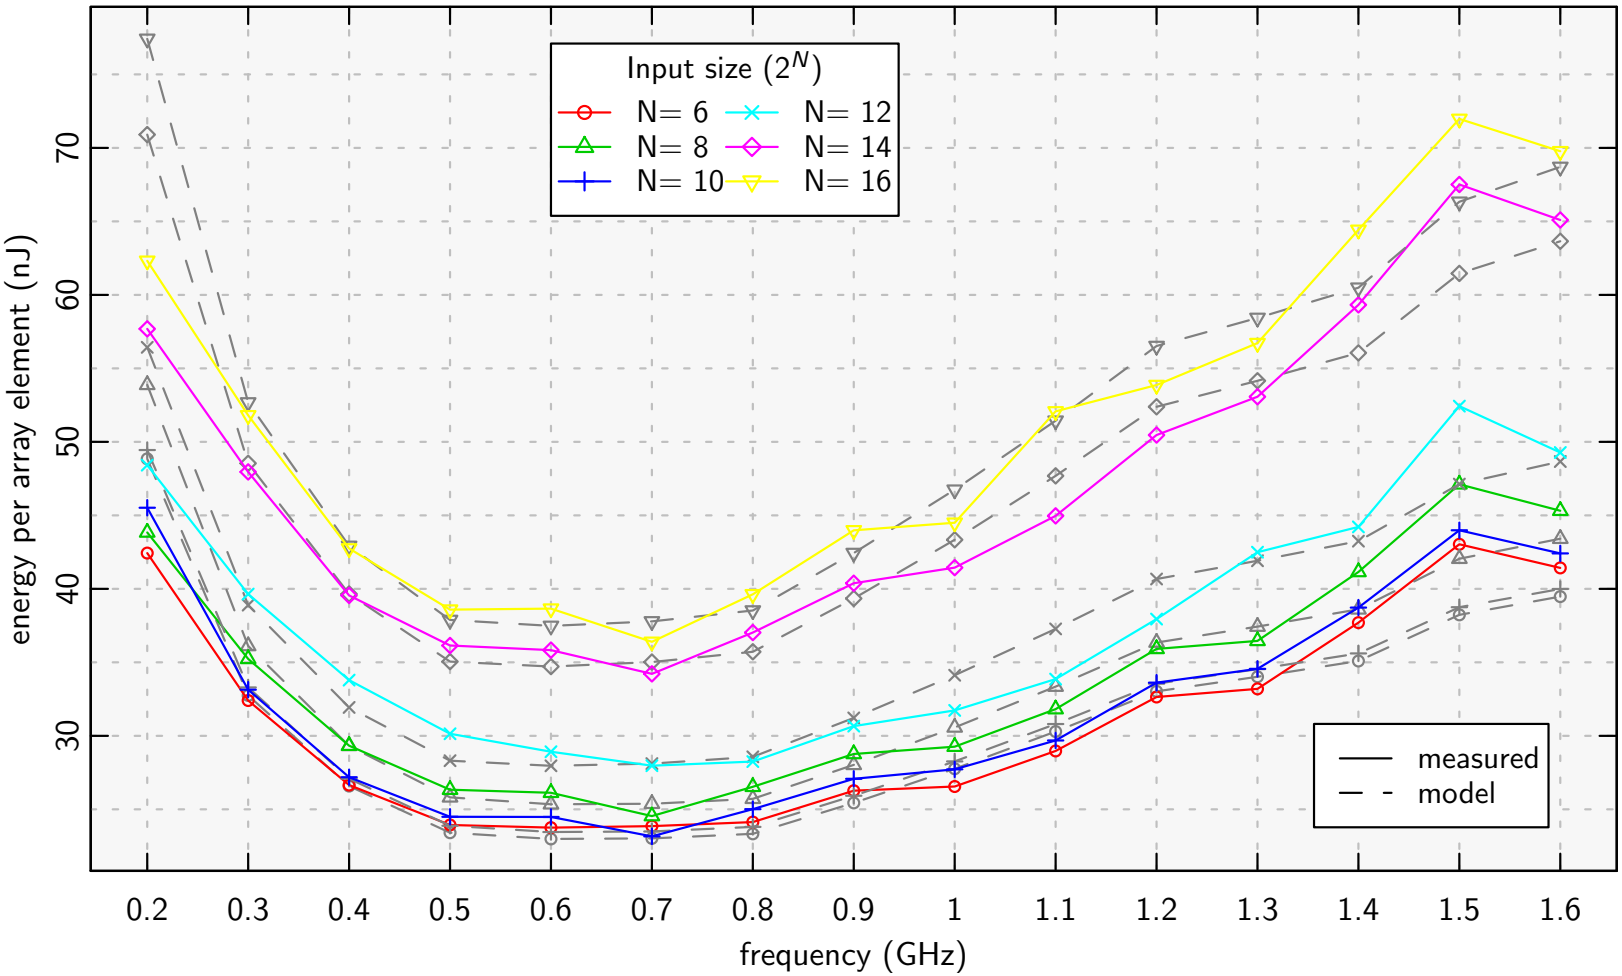

Energy Model's Parameter Sensitivity Analysis

- Energy consumption model under analysis:

$$E_{\text{sys}} = ([1 + \gamma V] \cdot \xi V^2 f + P_{\text{back}}) \cdot cc_{\text{b}} \left(\frac{1}{f - f_{\text{k}}} + \beta \right),$$

$$\left(\frac{\partial E_{\text{sys}}}{\partial f} \right)_{f = f_{\text{opt}}} = 0;$$

- The aim is to find the conditions under which f_{opt} is exploitable;
- The following parameters will be looked at in more detail:
  - frequency thieves (overhead) f_{k},
  - background power P_{back},
  - power gain ξ,
  - temperature $\gamma(T)$;
- Analysis based on energy profile of the Exynos 4210.

Influence of *frequency thieves* f_k on f_{opt}

$$E_{sys} = ([1 + \gamma V] \cdot \xi V^2 f + P_{back}) \cdot cc_b \left(\frac{1}{f - f_k} + \beta \right)$$

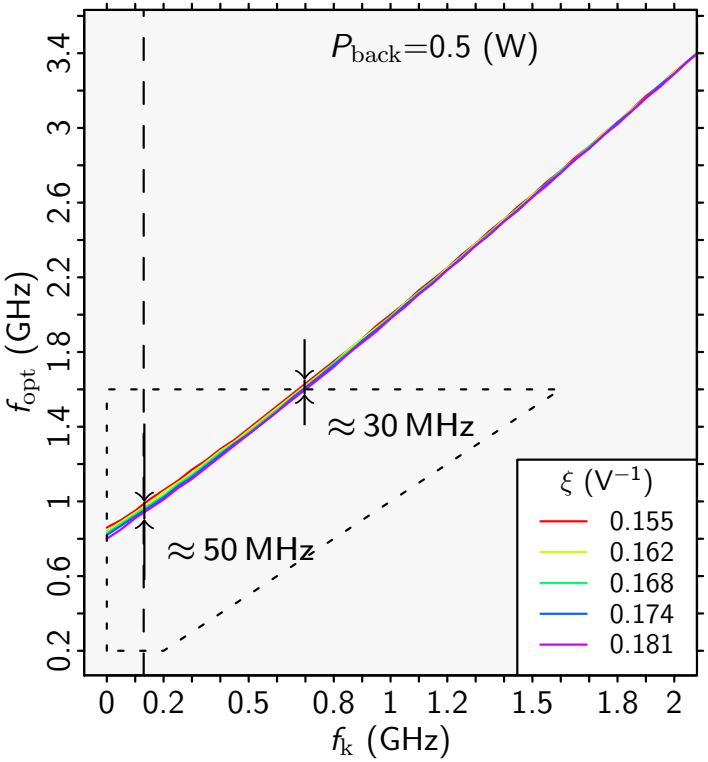

(e) $f_{opt}(f_k, \xi)$

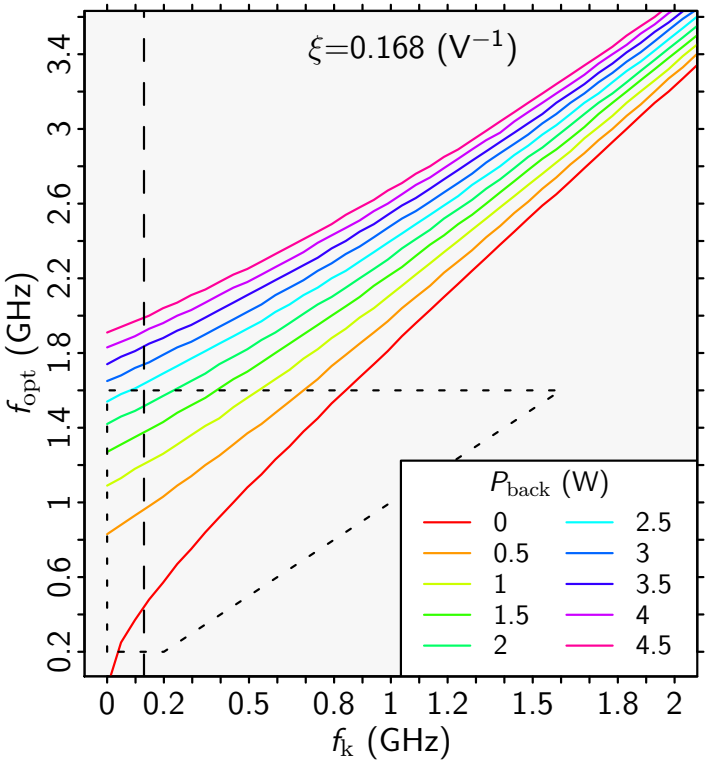

(f) $f_{opt}(f_k, P_{back})$

Influence of *Background Power* P_{back} on f_{opt}

$$E_{\text{sys}} = ([1 + \gamma V] \cdot \xi V^2 f + P_{\text{back}}) \cdot cc_{\text{b}} \left(\frac{1}{f - f_{\text{k}}} + \beta \right)$$

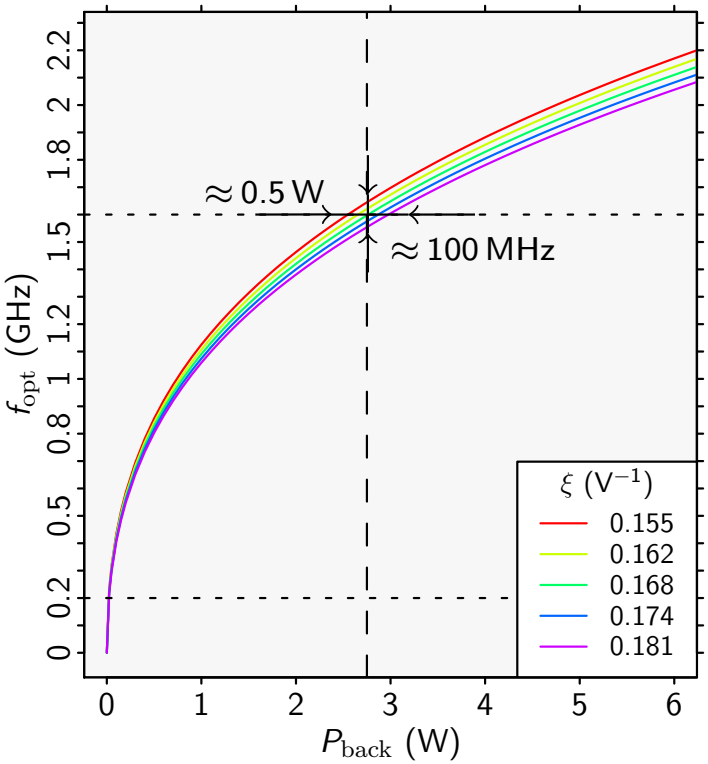

(g) $f_{\text{opt}}(P_{\text{back}}, \xi)$

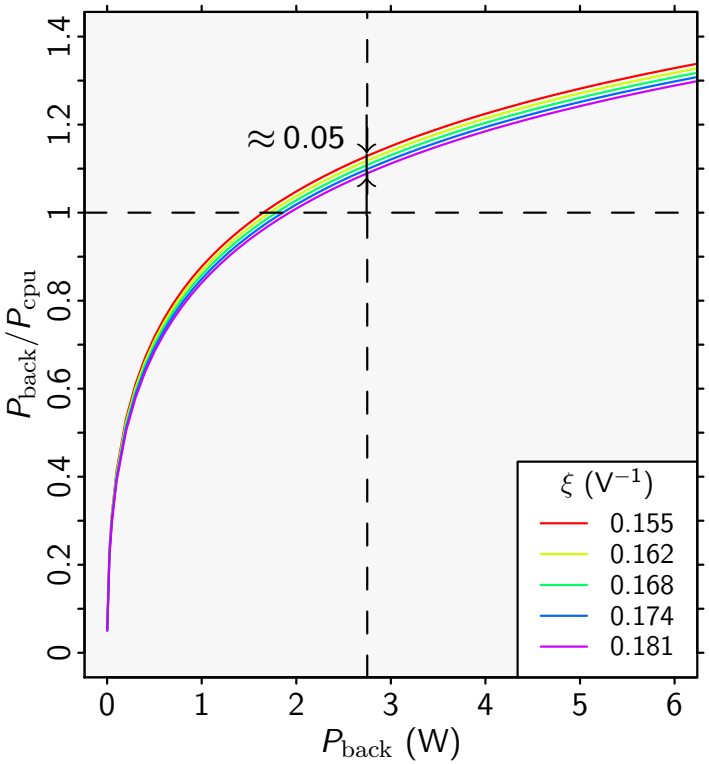

(h) $P_{\text{back}}/P_{\text{cpu}}$ ratio at f_{opt}

Influence of *Power Gain* ξ(s) on f_{opt}

$$E_{\text{sys}} = ([1 + \gamma V] \cdot \xi V^2 f + P_{\text{back}}) \cdot cc_{\text{b}} \left(\frac{1}{f - f_{\text{k}}} + \beta \right)$$

- Cooperative microprocessors on the same die:
  1. power-efficient: Cortex A7,
  2. high-performance: Cortex A15;
- ξ is scaled by s between its lower and upper bound: $s \in \{1, 2\}$;
- Exynos 5410 power model.

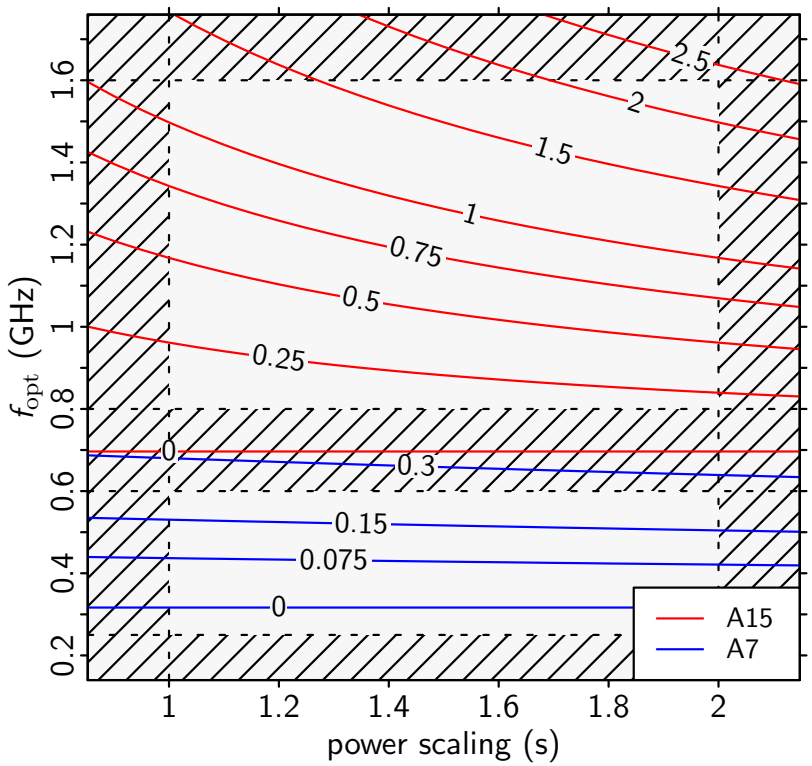

Numbers on the lines represent the background power for that line.

Influence of *Temperature* $\gamma(T)$ on f_{opt}

$$E_{\text{sys}} = ([1 + \gamma V] \cdot \xi V^2 f + P_{\text{back}}) \cdot cc_{\text{b}} \left(\frac{1}{f - f_{\text{k}}} + \beta \right)$$

- γ is a function of temperature;
- temperature/power model of Exynos 5410 is used;
- temperature/power shows exponential behavior;

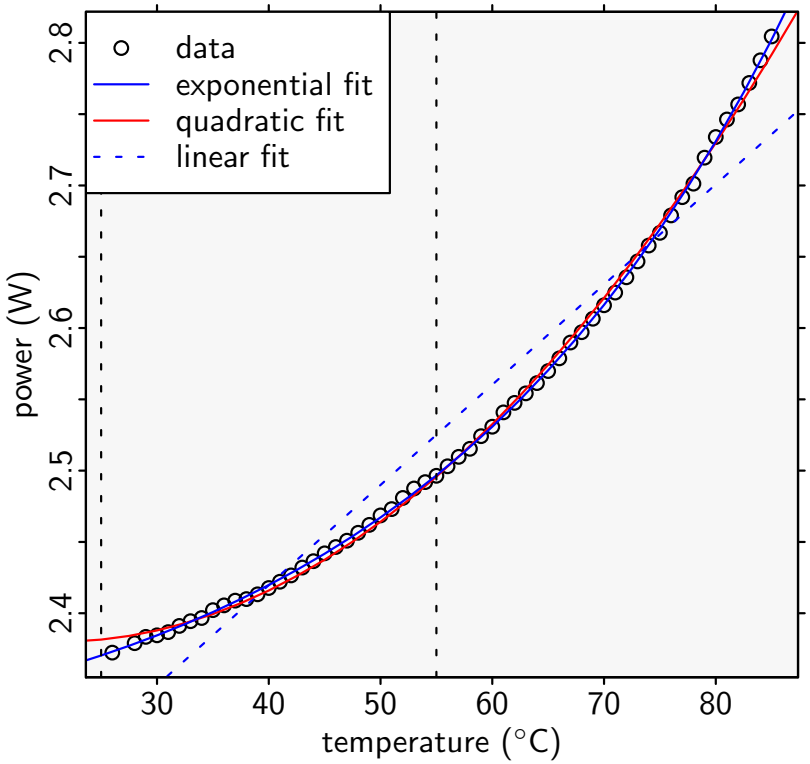

$\Delta f_{\text{opt}} \approx 200\,\text{MHz}$ when $25°\text{C} < T < 85°\text{C}$.

Case Study 1: f_{opt} Classification

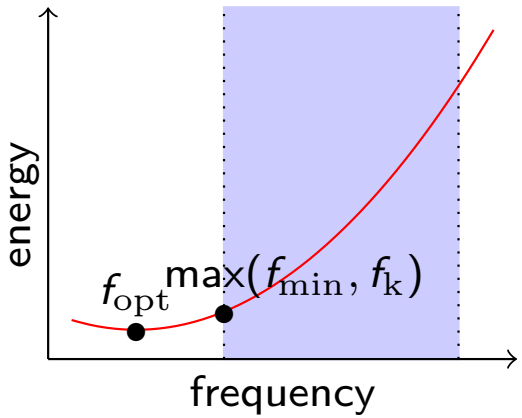

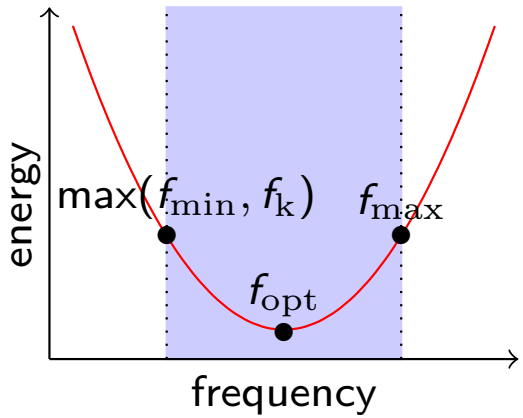

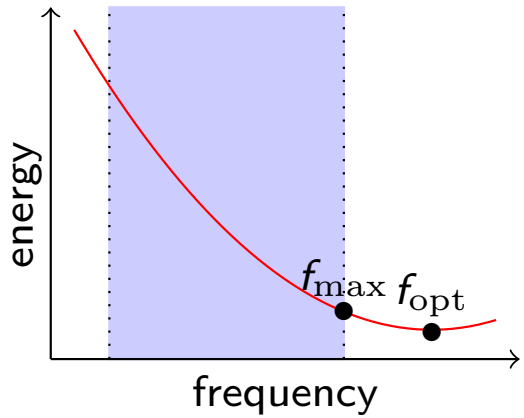

1			f_{opt}	$<$	$\max(f_{\text{min}}, f_{\text{k}})$	the slower, the better
2	$\max(f_{\text{min}}, f_{\text{k}})$	$\leq$	f_{opt}	$\leq$	f_{max}	chase f_{opt}
3	f_{max}	$<$	f_{opt}			race-to-halt

Case Study 2: f_{opt} and Multi-core Code Execution

Clock frequency scheduling schemes:

1. **on-demand**: binary (high/low) as work arrives;
2. **selfish**: each core is individually energy optimized;
3. **thread-cooperation**: all cores are collectively energy optimized.

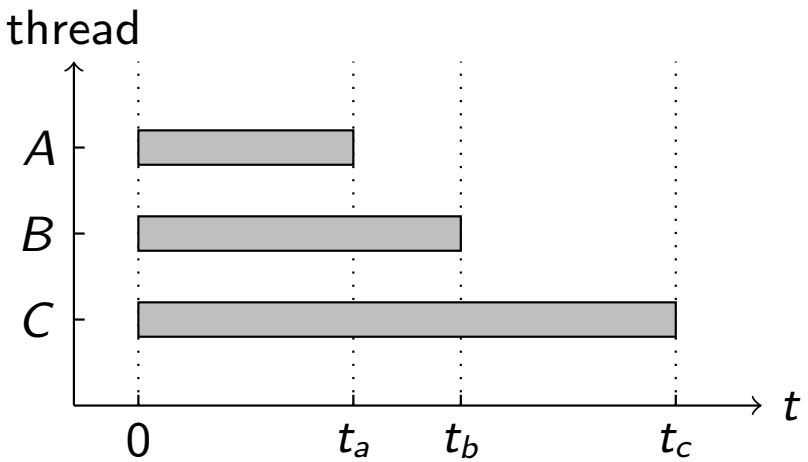

(a) default clock scaling

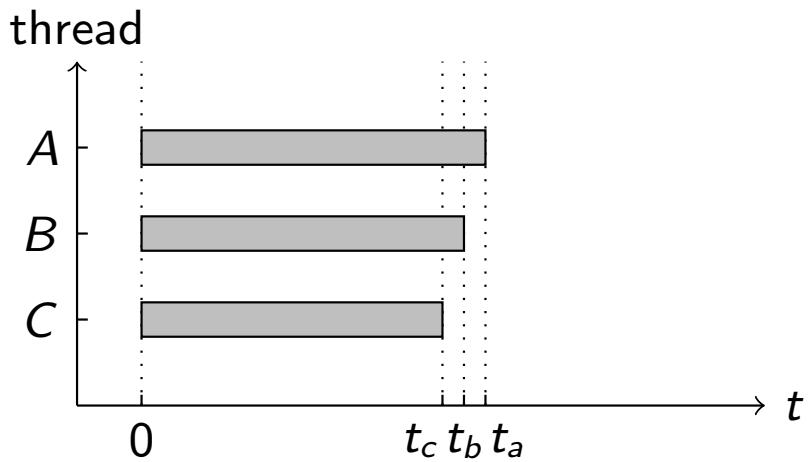

(b) cooperative clock scaling

Case Study 2: f_{opt} and Multi-core Code Execution contd.

Problem statement:

- n threads executed in parallel with common deadline $t_{\max}$;
- threads individually clock frequency f_i scalable;
- $E_{\text{tot}}(f_i) : \mathbb{R}^m \to \mathbb{R}$ to be minimized over f_i:

$$E_{\text{tot}}(f_i) = P_{\text{back}}\, t_{\max} + \sum_{i=0}^{n} \left[\frac{cc_{\text{b},i}}{f_i} P^+ + \left(t_{\max} - \frac{cc_{\text{b},i}}{f_i} \right) P^\circ \right],$$

$$\text{subject to} \quad \forall i, \quad \frac{cc_{\text{b},i}}{f_i} \leq t_{\max} \quad \text{and} \quad f_{\min} \leq f_i \leq f_{\max};$$

- $\{cc_{\text{b}}, f_i, t_{\max}, P_{\text{back}}, P^\circ, P^+\} \in \mathbb{R}^+$;
- active power (P^+) and idle power (P°) are generated by the Exynos 5410 power model.

Case Study 2: f_{opt} and Multi-core Code Execution contd.

Performance evaluation of 4 clock frequency scalable parallel threads.

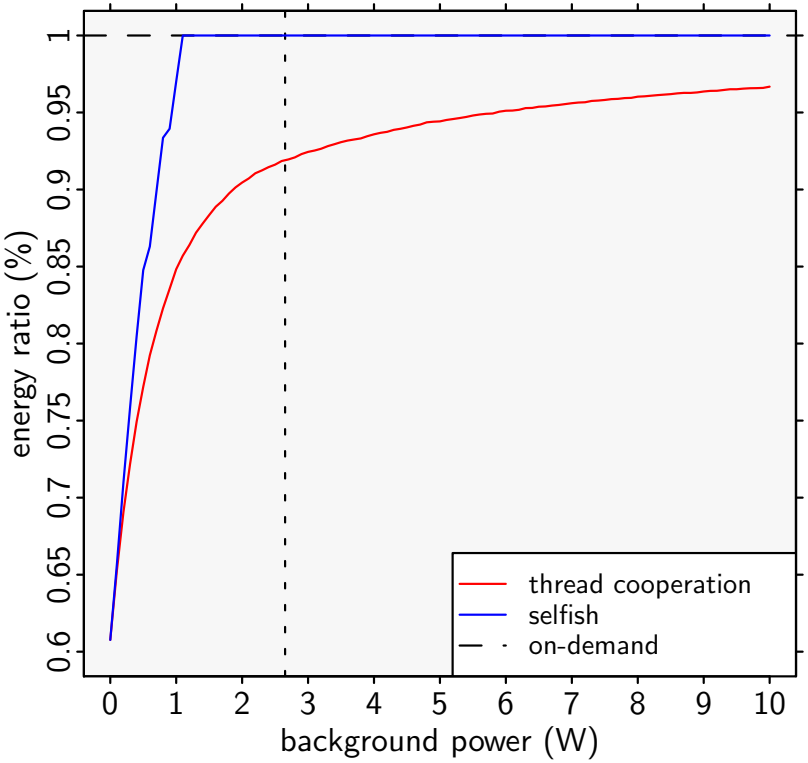

(a) energy

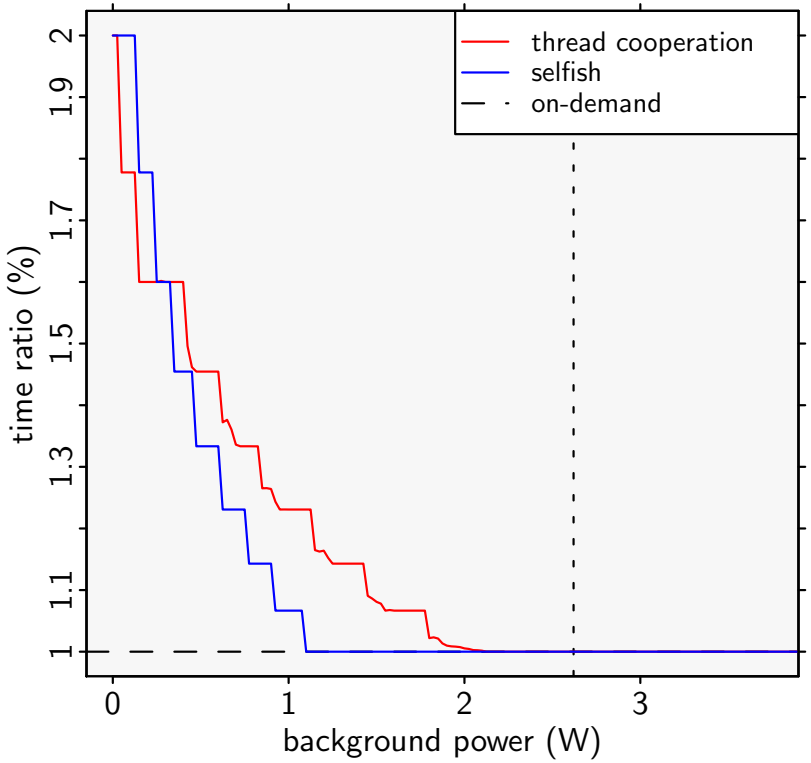

(b) time

Case Study 3: big-LITTLE Heterogeneous Computing

- Optimal clock frequency for cooperative microprocessors:
  1. power-efficient: Cortex A7,
  2. high-performance: Cortex A15;
- f_{opt} is chosen on the core yielding best efficiency;
- Exynos 5410 power model used.

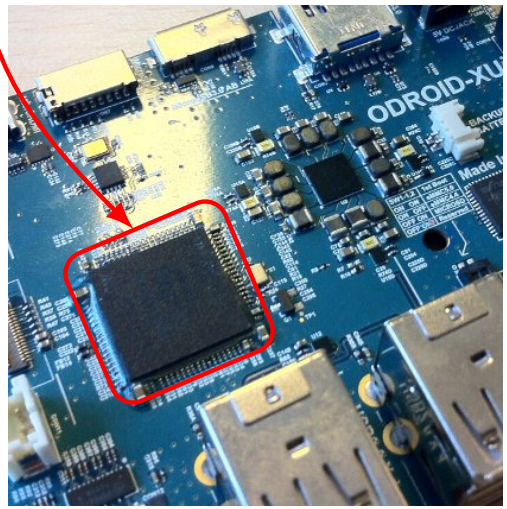

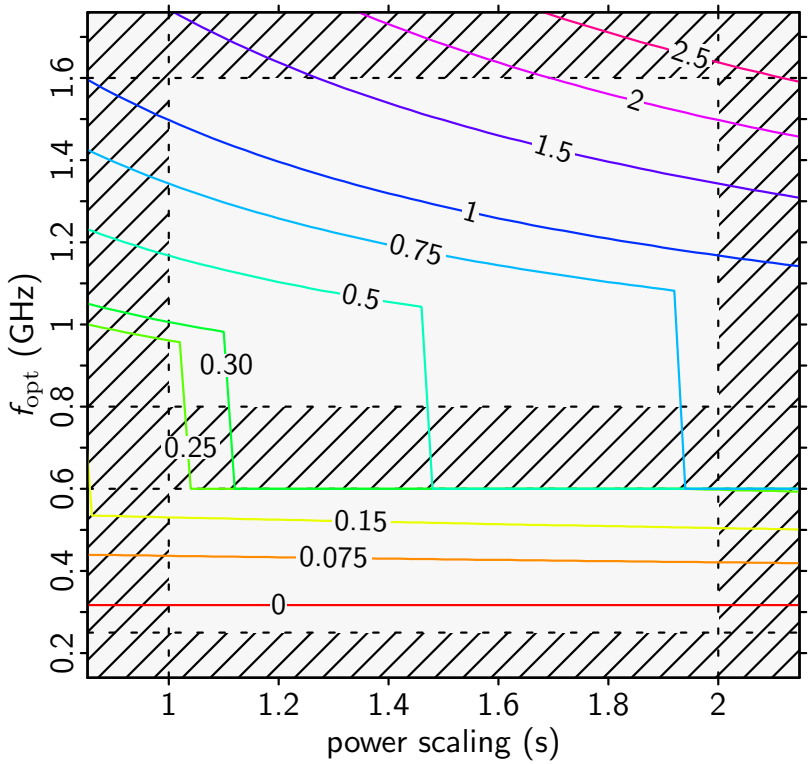

Numbers on the lines represent the background power for that line.

Conclusion

- System's energy consumption shows convex properties over f;
- rules of thumb for an exploitable f_{opt}:
  - P_{back} should be smaller than P_{cpu},
  - overhead (f_{k}) should be limited,
  - slack time β should be limited,
  - power profile (ξ) has *minimal effect*,
  - code size (cc_{b}) has *no effect*;

- energy gains could be from 10% up to 50% at fixed temperature;
- temperature/Power relationship shows exponential behavior;
- radiation can be omitted for small devices.

Future Work

Including:

- apply results to other domains:
  - multi-core,
  - HPC,
  - clock modulation,
  - interactive/performance;
- exploit the thermal behavior;
- better understanding of how much energy can practically be gained.

The Energy/Frequency Convexity Rule of Energy Consumption for Programs:

Modeling, Thermosensitivity, and Applications

Karel De Vogeleer

Ph.D. defense

September 4th, 2015

Publications

- K. De Vogeleer, G. Memmi, P. Jouvelot, and F. Coelho, "The Energy/Frequency Convexity Rule: modeling and experimental validation on mobile devices," in *Proceedings of the 10th Conference on Parallel Processing and Applied Mathematics*. Springer Verlag, Sep. 2013.
- K. De Vogeleer, G. Memmi, P. Jouvelot, and F. Coelho, "Modeling the temperature bias of power consumption for nanometer-scale CPUs in application processors," in *14th International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation*, Jul. 2014, pp. 172-180.
- K. De Vogeleer, P. Jouvelot, and G. Memmi, "The impact of surface size on the radiative thermal behavior of embedded systems," CoRR, vol. abs/1410.0628, 2014, (submitted to *IEEE TMC* in 2014).
- K. De Vogeleer, G. Memmi, and P. Jouvelot, "Parameter Sensitivity Analysis of the Energy/Frequency Convexity Rule for Nanometer-scale Application Processors," CoRR, vol. abs/1508.07740, 2015, (in submission to *The Elsevier Journal of Parallel and Distributed Computing*, 2015).

References I

Fan, X., Ellis, C. S., and Lebeck, A. R. The synergy between power-aware memory systems and processor voltage scaling. In *Proceedings of the Third international conference on Power - Aware Computer Systems* (Berlin, Heidelberg, 2004), Springer-Verlag, pp. 164–179.

Hager, G., Treibig, J., Habich, J., and Wellein, G. Exploring performance and power properties of modern multi-core chips via simple machine models. *Concurrency and Computation: Practice and Experience* (2013), n/a–n/a.

Le Sueur, E., and Heiser, G. Dynamic voltage and frequency scaling: the laws of diminishing returns. In *Proceedings of the 2010 international conference on Power aware computing and systems* (Berkeley, CA, USA, 2010), HotPower'10, pp. 1–8.

Snowdon, D. C., Ruocco, S., and Heiser, G. Power management and dynamic voltage scaling: Myths and facts. In *2005 WS Power Aware Real-time Comput.* (New Jersey, USA, Sept. 2005).