Thermal behavior
and
Energy/Frequency Convexity Rule
of
Energy Consumption for Programs

Karel De Vogeleer
Pierre Jouvelot
Gérard Memmi

Motivations for focussing on mobile computing

- **It is NOT the energy saving per se :**
  - A smartphone CPU consumes between 60 to 400mW
  - There are about 7×10^9 smartphones sold in the last 5 years, there will be 50×10^9 ‘smart objects’ in 2020
  - A saving of 30% would provide grossly about 280 MW for the smartphones, about 3 GW for the smart objects
  - This would only save between a tidal and a nuclear power station

- **Focussing on mobile sytems: they are ‘*energy-critical’* : it is being constantly looking for providing more autonomy with a QoS unchanged**
- **It is about a natural-resource-free energy saving**

Thermal Behavior:
Power-temperature rule
Passive Cooling rule

Temperature impacts energy consumption

$$P = a_1 e^{T/a_2} + a_0,$$

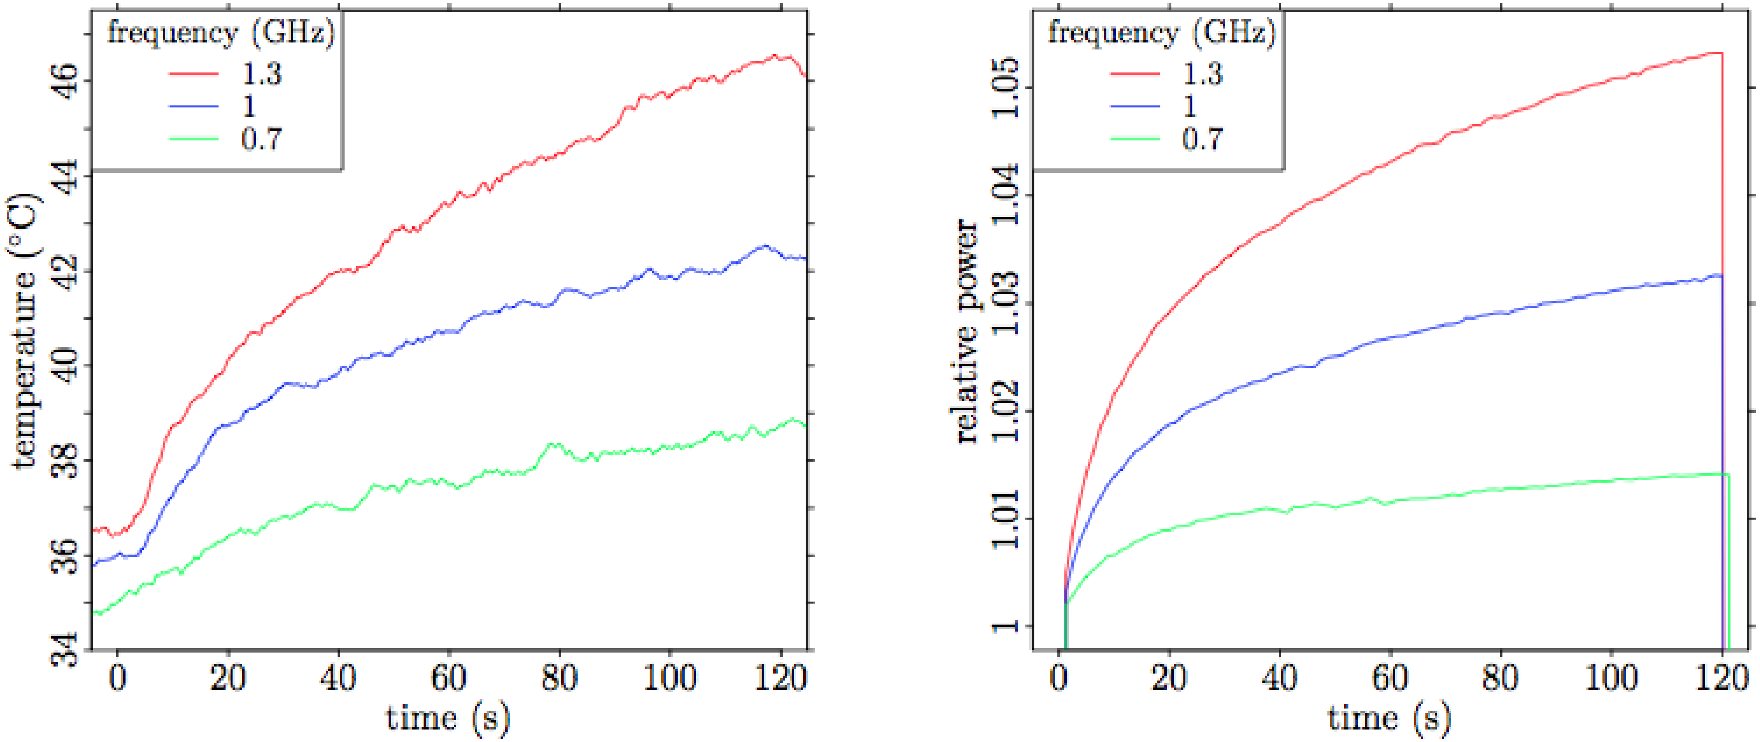

(a) température

(b) puissance

Passive Cooling Rule

The passive lumped system's transient thermal behavior is defined by

$$\text{internal heat generation} + \text{radiation} + \text{convection}$$

$$mC\frac{dT}{dt} = P(T) + \epsilon\sigma(T^4 - T_a^4) + h(T - T_b),$$

which has the solution

$$t(T) = -\frac{1}{\kappa_4}\Bigg(A\ln|T-\omega_1| + \frac{C}{2}\ln|(T-\alpha)^2+\beta^2| + B\ln|T-\omega_2| + \frac{\alpha C - D}{\beta}\arctan\left(\frac{T-\alpha}{\beta}\right) + c_o\Bigg).$$

Approximative solutions

- Coefficient approximation:

$$T = \frac{\omega_1 \pm \omega_2 c_o e^{-\frac{\kappa_2}{A} t}}{1 \pm c_o e^{-\frac{\kappa_2}{A} t}};$$

- First-order O'Sullivan:

$$T = \left(T_0 - T_a + \frac{p}{n} \right) e^{-\frac{n}{C} t} - \frac{p}{n} + T_a;$$

- Second-order O'Sullivan:

$$T = \frac{w}{2m} (\mathrm{co}) \tanh \left(\frac{w}{2m} t + c_0 \right) - \frac{n}{2m} + T_a.$$

Isothermal assumption

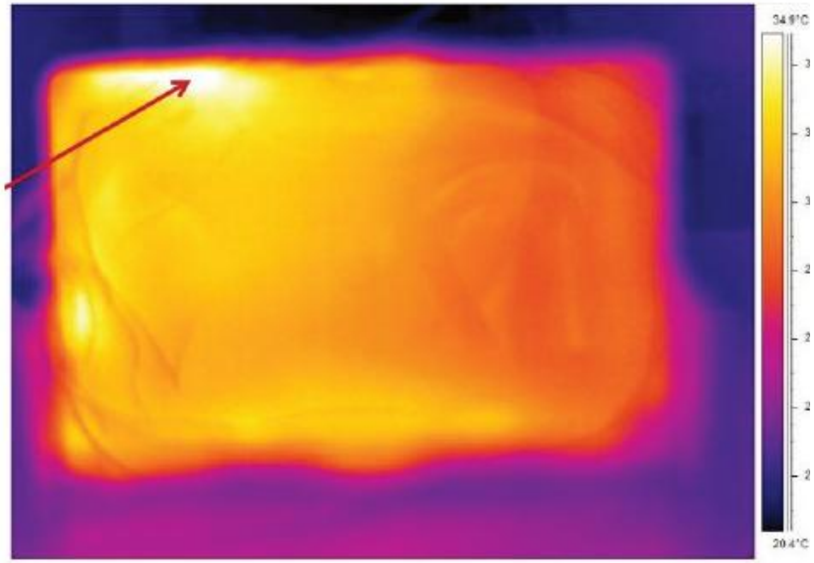

(a) an Apple iPad by Wagner and Maltz [20]

Fig. 1. Experimental thermal imaging of the skin temperature of (a) an Apple iPad [20] and (b) an unnamed thin notebook [21]. The iPad exemplifies the quasi isothermal surface of an embedded system. The surface temperature varies between 30°C and 35°C. On the other hand, the thin notebook shows large temperature variations, between 48°C and 25°C. Here, aggressive active cooling methods extract the heat as fast as possible from the heat sources inside the device.

Contribution on thermal behavior

- Necessary for reproducible measurement and for accurate energy consumption models
- Power – temperature relationship
- Approxiations for practical use

EFCR: the energy – frequency convexity rule

Fragmenting energy consumption per system module

- System's energy consumption E_{sys} definition

$$\begin{aligned} E_{\text{sys}} &= \int_0^{\Delta t} P_{\text{total}}\, dt \\ &= \int_0^{\Delta t} \left(P_{\text{cpu}} + P_{\text{back}} \right) dt; \end{aligned}$$

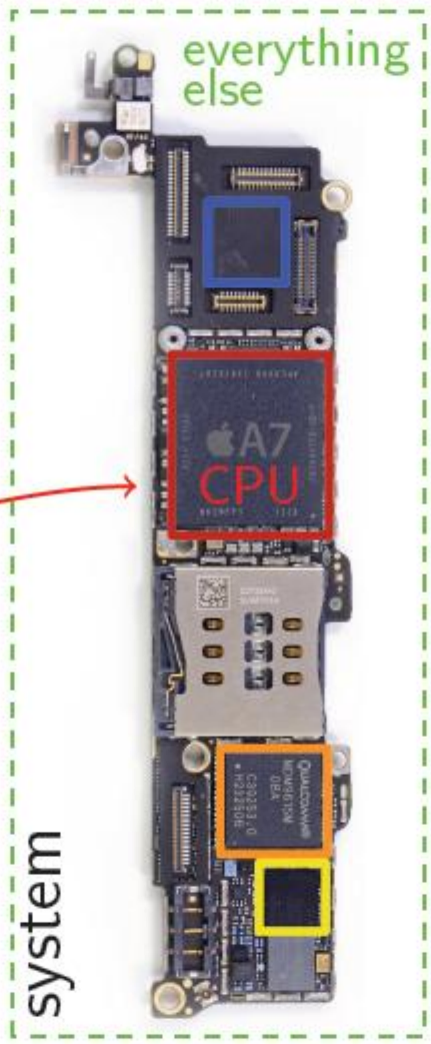

- Examples of P_{back} include:
  - LCD screen,
  - radio interface,
  - power supply;
- If P_{cpu} and P_{back} are constant over Δt:

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

Power and time model

Microprocessor Power Model

CPU power P_{cpu} consists of:

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

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

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;

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

Optimal frequency and Convexity

- System's energy consumption model (EFCR)

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

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

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

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

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

3 classes of processors

- ***There exists a domain of frequencies where the processor is delivered by the manufacturer:*** *f_k must be smaller than f since f_k a fraction of f*
- ***Save for overclocking or underclocking***

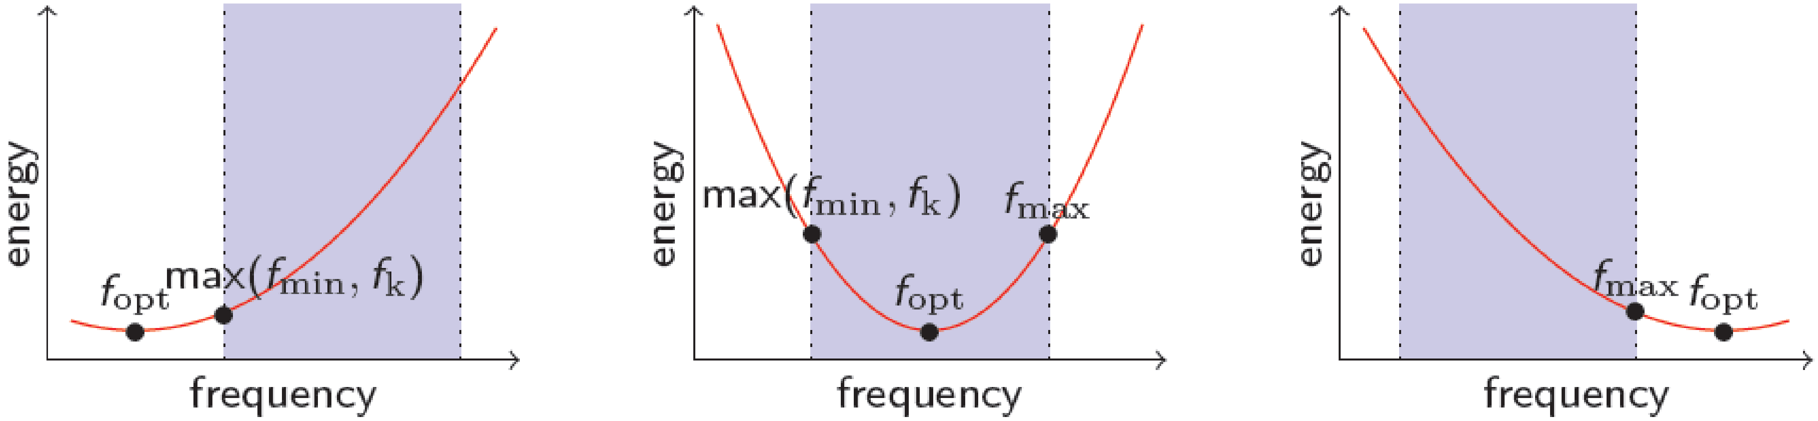

1			f_{opt}	$<$	$\max(f_{min}, f_k)$	the slower the better
2	$\max(f_{min}, f_k)$	$\leq$	f_{opt}	$\leq$	f_{max}	chase f_{opt}
3	f_{max}	$<$	f_{opt}			race-to-halt

State of the art

- **Convexity was already observable, however no analytical studies were performed**

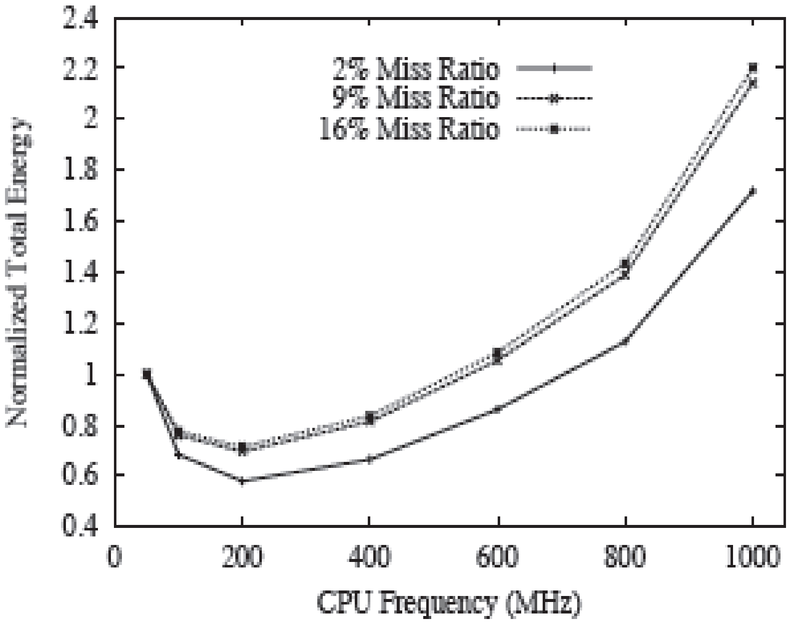

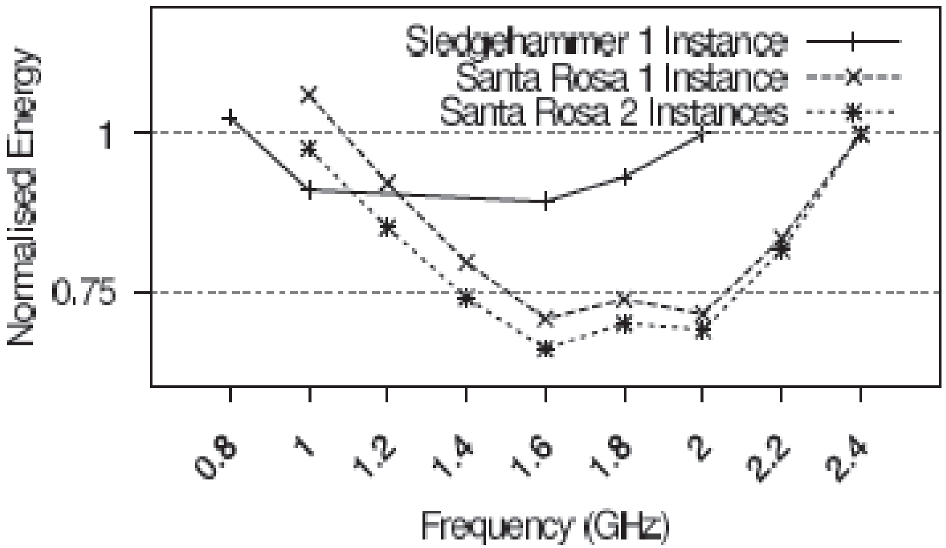

Fan, X., Ellis, C. S., and Lebeck, A. R.
The synergy between power-aware memory systems and processor voltage scaling. In PACS'04
Le Sueur, E., and Heiser, G.
Dynamic voltage and frequency scaling: the laws of diminishing returns. In *PACS'10*

Testbed

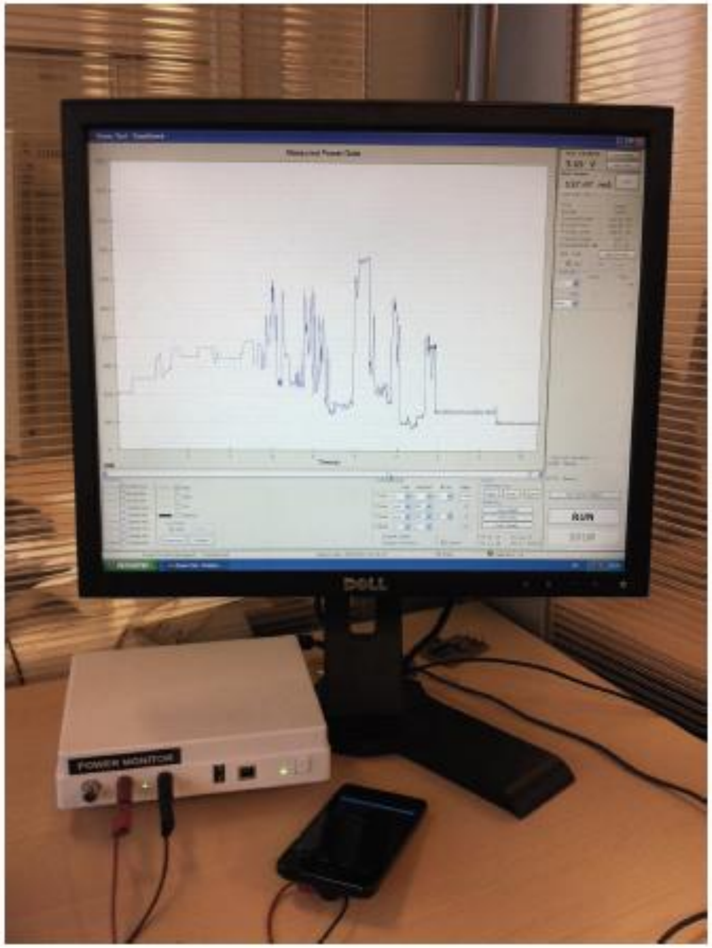

- **Benchmark**: bit-reverse algorithm, part of the FFT algorithm:

```
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.

Experimental validation

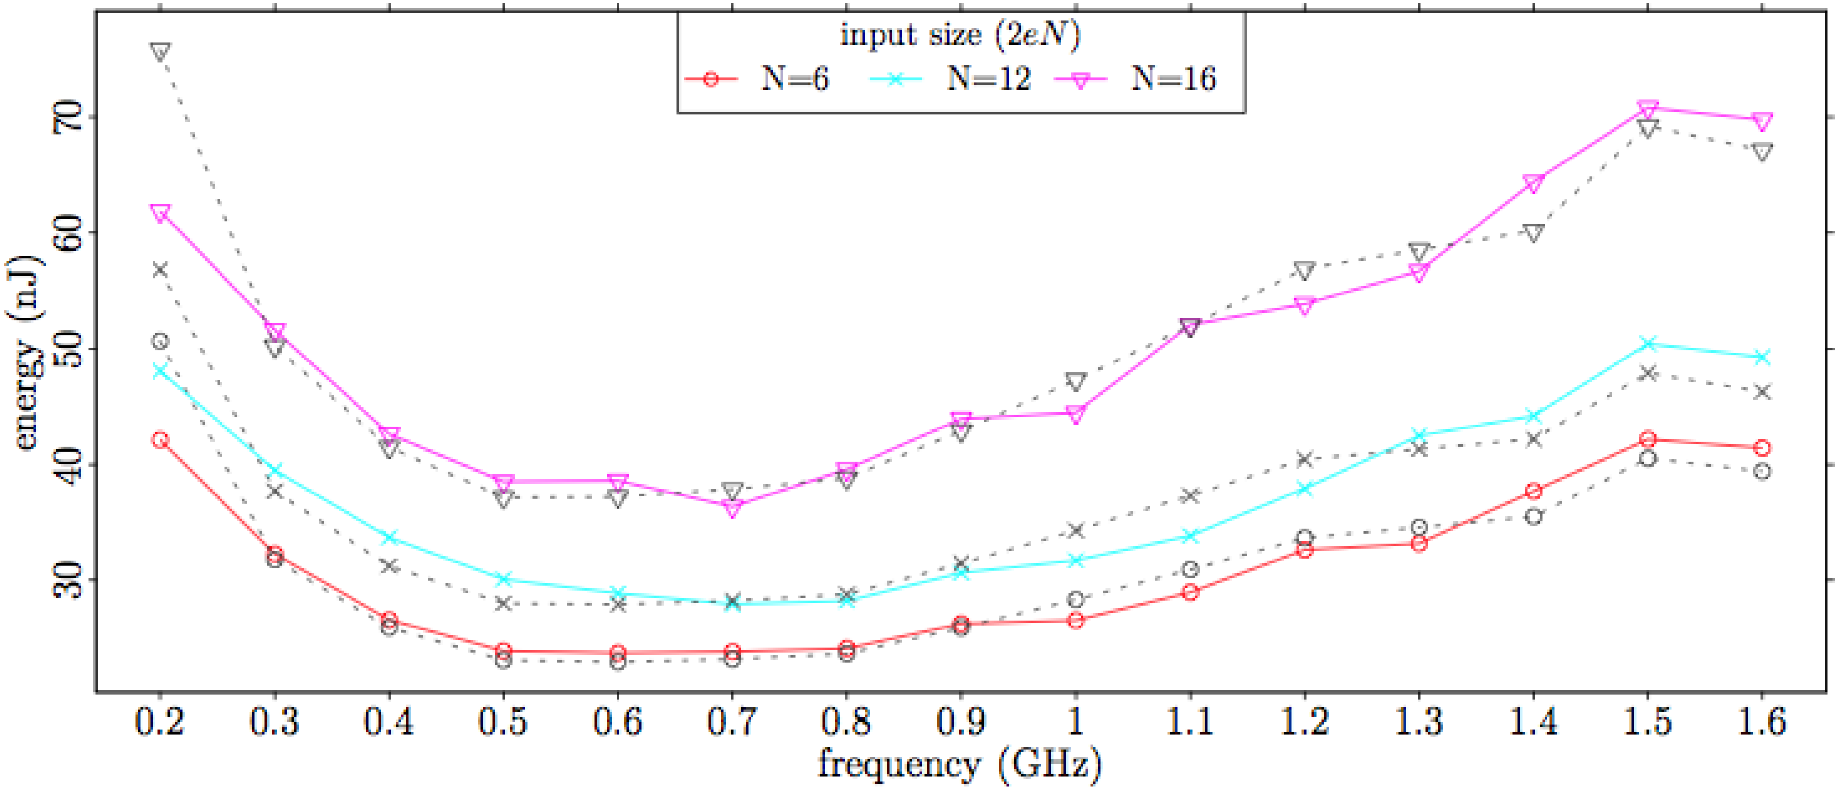

In color, the measurements

In doted lines the theoretical EFCR calculation

When N increases, f_{opt} stays stable

Conclusion

Towards energy program profiling

- Created by tuning clk frequency and performing standard program transformation

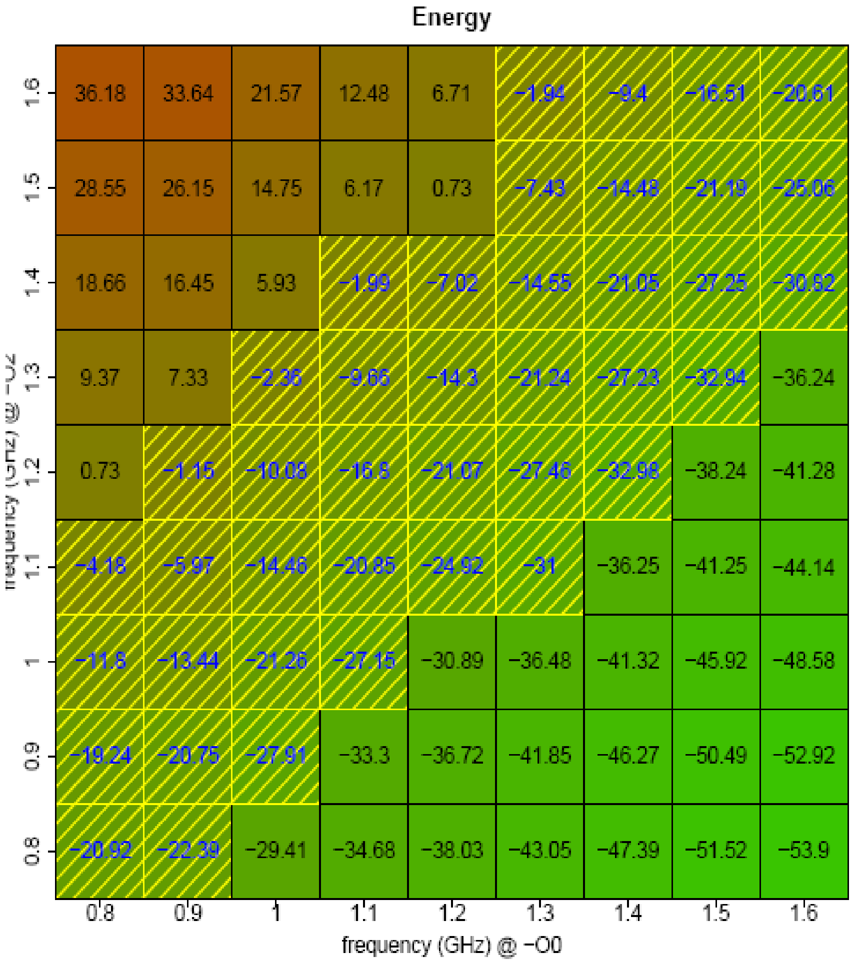

	0.8	0.9	1	1.1	1.2	1.3	1.4	1.5	1.6
1.6	36.18	33.64	21.57	12.48	6.71	−1.94	−9.4	−16.51	−20.61
1.5	28.55	26.15	14.75	6.17	0.73	−7.43	−14.48	−21.19	−25.06
1.4	18.66	16.45	5.93	−1.99	−7.02	−14.55	−21.05	−27.25	−30.82
1.3	9.37	7.33	−2.36	−9.66	−14.3	−21.24	−27.23	−32.94	−36.24
1.2	0.73	−1.15	−10.08	−16.8	−21.07	−27.46	−32.98	−38.24	−41.28
1.1	−4.18	−5.97	−14.46	−20.85	−24.92	−31	−36.25	−41.25	−44.14
1	−11.6	−13.44	−21.26	−27.15	−30.89	−36.48	−41.32	−45.92	−48.58
0.9	−19.24	−20.75	−27.91	−33.3	−36.72	−41.85	−46.27	−50.49	−52.92
0.8	−20.92	−22.39	−29.41	−34.68	−38.03	−43.05	−47.39	−51.52	−53.9

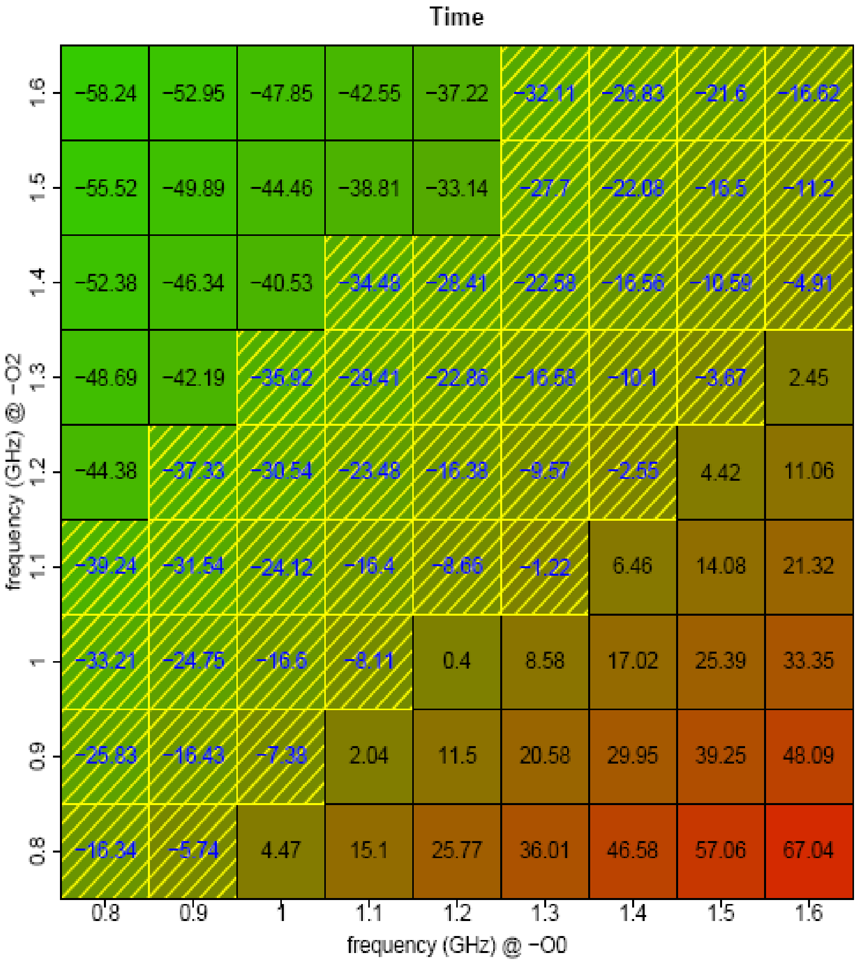

	0.8	0.9	1	1.1	1.2	1.3	1.4	1.5	1.6
1.6	−58.24	−52.95	−47.85	−42.55	−37.22	−32.11	−26.83	−21.6	−16.62
1.5	−55.52	−49.89	−44.46	−38.81	−33.14	−27.7	−22.08	−16.5	−11.2
1.4	−52.38	−46.34	−40.53	−34.48	−28.41	−22.58	−16.56	−10.59	−4.91
1.3	−48.69	−42.19	−35.92	−29.41	−22.86	−16.58	−10.1	−3.67	2.45
1.2	−44.38	−37.33	−30.54	−23.48	−16.38	−9.57	−2.55	4.42	11.06
1.1	−39.24	−31.54	−24.12	−16.4	−8.66	−1.22	6.46	14.08	21.32
1	−33.21	−24.75	−16.5	−8.11	0.4	8.58	17.02	25.39	33.35
0.9	−25.83	−16.43	−7.38	2.04	11.5	20.58	29.95	39.25	48.09
0.8	−16.34	−5.74	4.47	15.1	25.77	36.01	46.58	57.06	67.04

Energy-Oriented Environment

First measurement and results are very encouraging, setting expectations in the 10-40% saving range rather than in the 2-5% initially anticipated saving range

- ✓ Exploiting energy-frequency convexity
- ✓ Eliminating temperature impact in our models

- **Wider array of experimentation, more accurate measurement and mathematical models, wider temperature range**
- **More research on energy program profiling**
  - Handling various architectures (eg cache)
  - Understanding how and when to play with clock frequency changes
  - Temperature on line monitoring

Towards an energy-oriented compiler middle-end, and operating system technology

Bibliography

- 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).

TELECOM ParisTech

Institut
Mines-Télécom

Thank you