Karlsruhe Institute of Technology

Low Power Design

Volker Wenzel on behalf of Prof. Dr. Jörg Henkel
Summer Term 2016

CES – Chair for Embedded Systems

Lecture Slides

Overview Low Power Design Lecture

- Introduction and Energy/Power Sources (1)
- Energy/Power Sources(2): Solar Energy Harvesting
- Battery Modeling – Part 1
- Battery Modeling – Part 2
- Hardware power optimization and estimation – Part 1
- Hardware power optimization and estimation – Part 2
- Hardware power optimization and estimation – Part 3
- Low Power Software and Compiler
- Thermal Management – Part 1
- Thermal Management – Part 2
- Aging Mechanisms in integrated circuits
- Lab Meeting

Overview today

- „Boring“ Physics: Temperature, Heat, etc.
- Temperature and Reliability
- Thermal Management through DVFS
- Thermal Management through Task Scheduling

Boring Physics: Temperature

- measured with a thermometer
- Units:
  - Kelvin
  - Celsius
  - ~~Fahrenheit~~

$$T = \left(\frac{T_C}{°C} + 273.15\right) K$$

- absolute temperature proportional to avg. kinetic energy of particles

$$E_{kin} = \frac{1}{2} m v^2 = \frac{3}{2} k_B T$$ (ideal gas)

Boring Physics: Heat

- Heat Q is the energy ([J]) transferred between a system and its environment because of a temperature difference that exists between them.
- heat flows from hotter body to colder
- until equilibrium state is reached

Boring Physics: Heat Capacity

- The heat capacity C of an object is the proportionality constant between the heat Q that the object absorbs or loses and the resulting temperature change

$$Q = C\Delta T = C(T_f - T_i)$$

- Specific Heat c

$$Q = cm\Delta T = cm(T_f - T_i)$$

- $c_{Uranium}$ = 0.1 J/g/K
- c_{Fe} = 0.5 J/g/K
- c_{Air} = 0.7 J/g/K
- c_{Si} = 0.7 J/g/K
- c_{Water} = 4.2 J/g/K
- c_{H} = 14.3 J/g/K

src.: en.wikipedia.org

Boring Physics: Heat Transfer

- exchange of thermal energy in different modes
  - Conduction / Diffusion: transfer of energy between objects in physical contact. Thermal conductivity is the property of a material to conduct heat.
  - Radiation: transfer of energy from the movement of charged particles within atoms is converted to electromagnetic radiation.
  - Convection: transfer of energy between an object and its environment, due to fluid motion.
- Heat flux q: rate of heat energy transfer per unit surface.

$$q = \frac{Q}{A \Delta T}$$

Boring Physics: Thermal Conductivity

- Thermal Conductivity λ: property of a material to conduct heat (W/(m K))

$$\vec{q} = -\lambda \nabla T$$

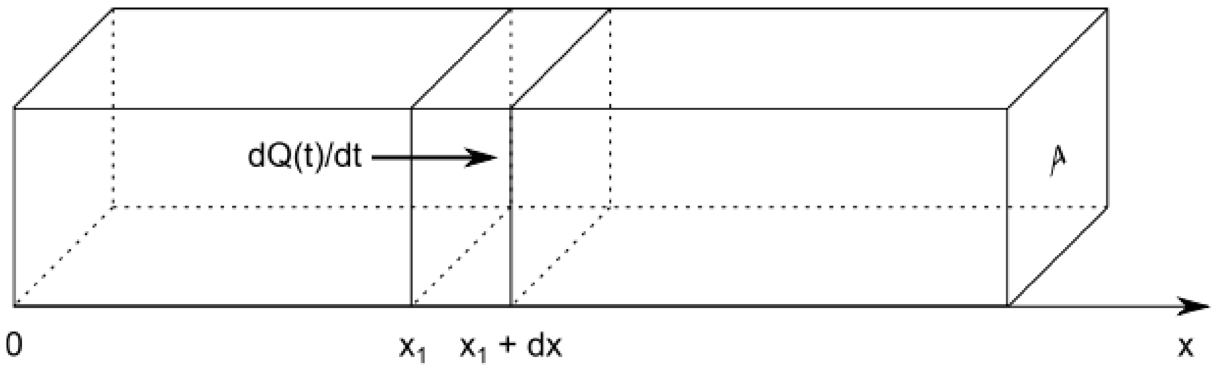

Material	*k*
Silica	1.5
Cu	401
$H_2 0$	0.6
Air	0.025

Boring Physics: Joule Heating

- passage of an current through an electric conductor releases heat
- electrons collide with atoms in conductor and transfer energy

$$Q \propto I^2 R t$$

src.: weller.de

Boring Physics: Power & Temperature

- Basic temperature equation:

$$C\frac{dT}{dt} = -\dot{Q} + P \qquad T(t_1) = T_0 + \frac{1}{C}\int_{t_0}^{t_1} -\dot{Q}(t) + P(t)\,dt$$

where $\dot{Q}$ is the heat dissipation rate.

$$T(t) = T_0 - (T_{SS} - T_0)e^{-\frac{t}{h}}$$

Heating

$$T(t) = T_0 + (T_0 - T_A)e^{-\frac{t}{c}}$$

Cooling

- T_{SS} is the steady state temperature the system will asymptotically reach with current power configuration
- Ambient temperature T_A is minimum reachable temperature

Boring Physics: Power & Temperature

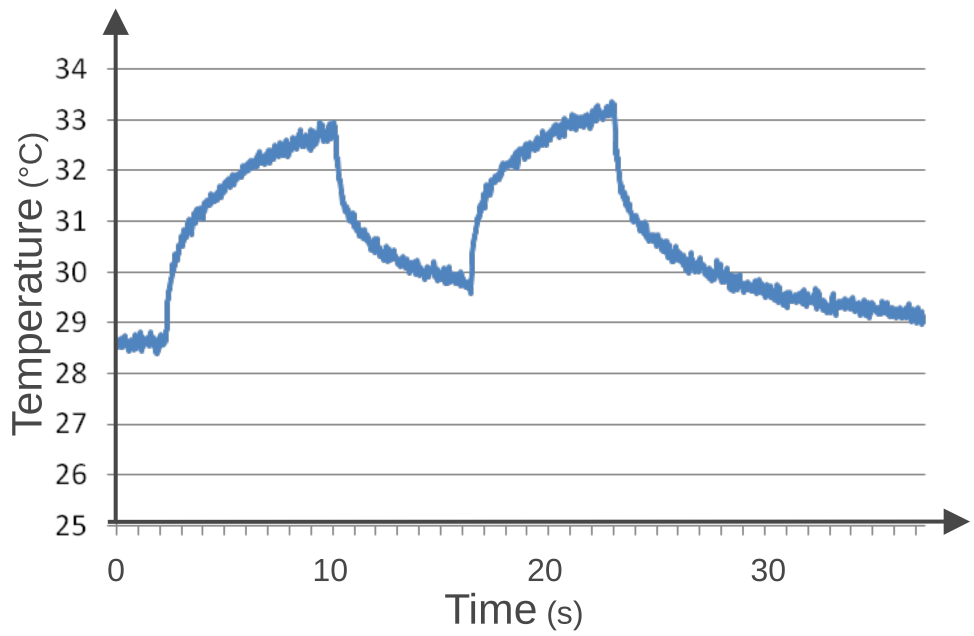

Heat Remains a Problem...

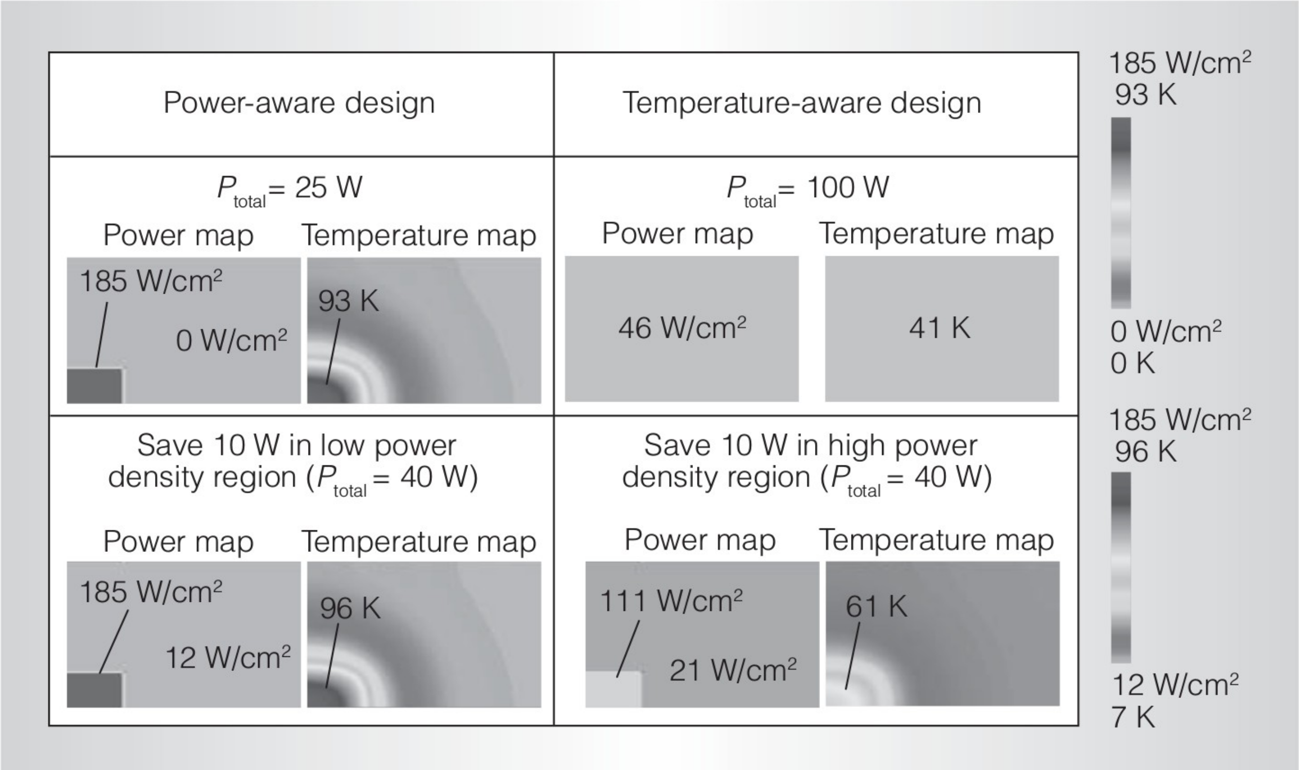

Figure 1. Temperature versus power awareness.[1] The temperatures shown are in terms of degrees Kelvin above room temperature.

(src: K. Skadron: Low-Power Design and Temperature Management; IEEE Micro, Vol. 27, No. 6, 2007)

Heat Remains a Problem… (cont'd)

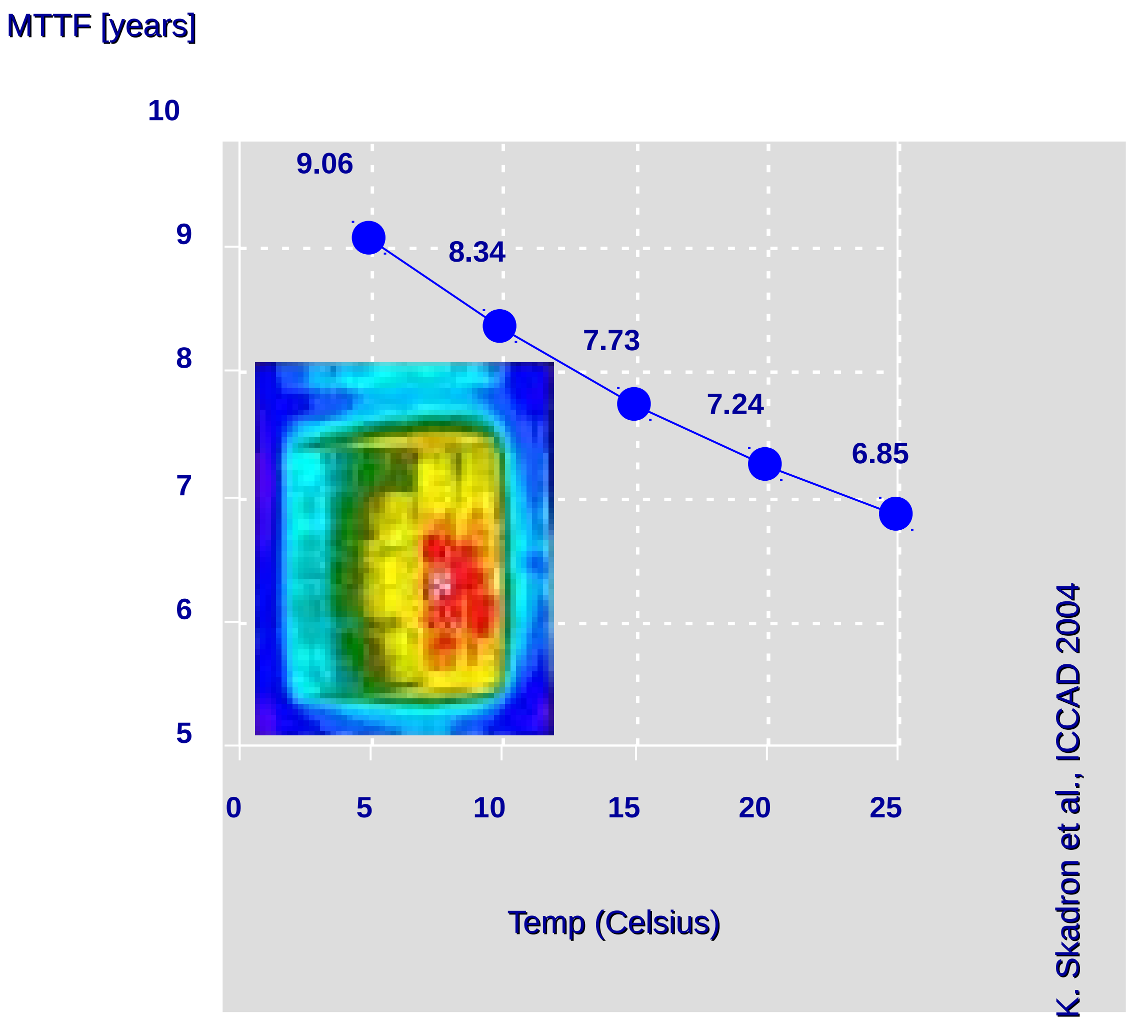

Temperature and Reliabilty

- *MTTF* also affected by thermal gradients

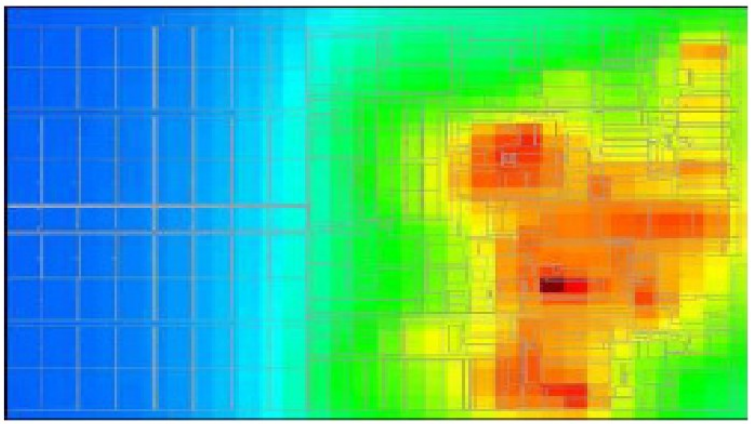

Spatial gradients
Simulated Thermal map Pentium M
[L.Finkelstein, Intel 2005]

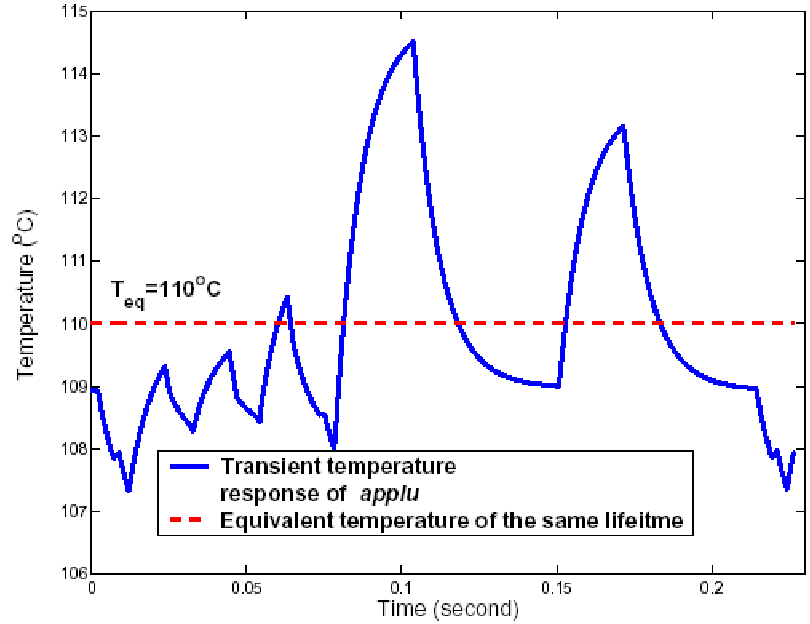

Temporal gradients
[K. Skadron, 2005]

  - → Goal: balance temperatures

Thermal Gradients

- Same scenario, this time measured on an FPGA

Scr: Amrouch, Ebi, Henkel

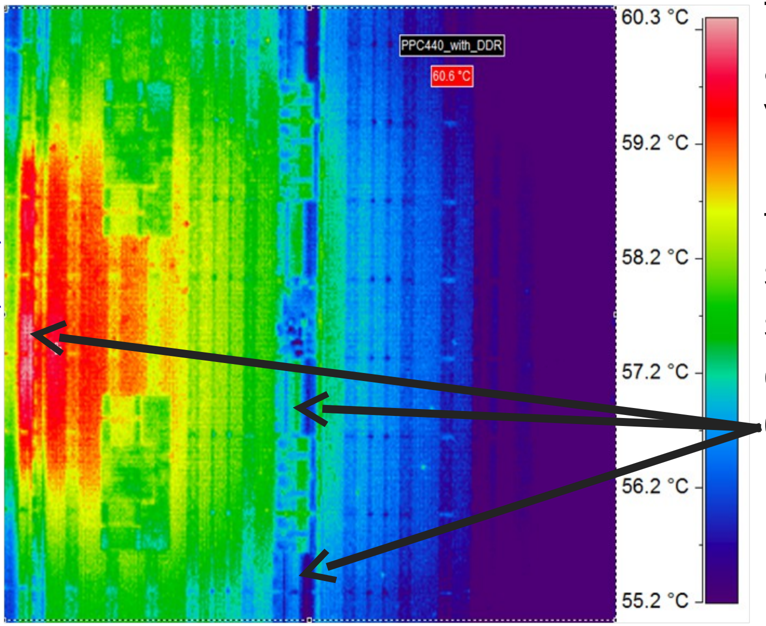

Temperature analysis of Virtex-5 FPGA using infrared thermal camera showing peak spatial thermal gradients of 0.12°C/µm

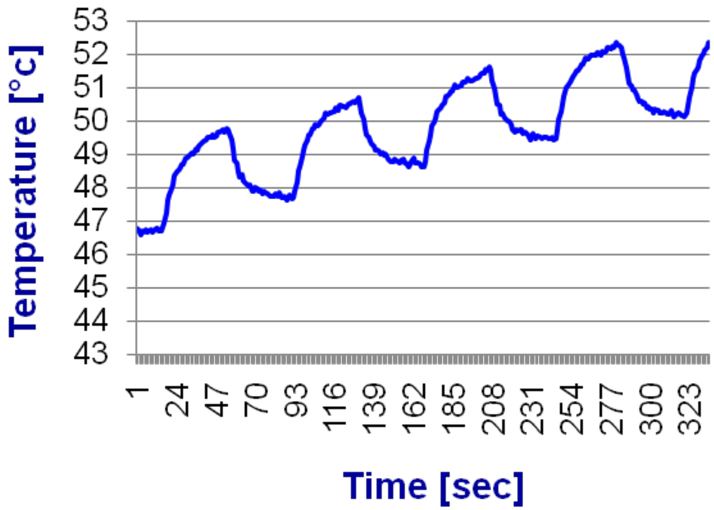

Virtex-5 with two PowerPC CPUs

Spatial gradients

Temporal gradients

Temperature and Reliability

- Transient errors may result due to timing errors
  - Approx. 5% decrease in delay every 10°C temperature increase [Xie 2006]
  - Timing errors result from spatial temperature variations

  → localized hotspots need to be avoided
  - Clock trees are particularly vulnerable
    - Span across multiple thermal areas
    - Additional buffers can be inserted to cope with thermal clock skew

[Chakraborty, 2008]

Temperature and Reliability (cont'd)

TMU:
Temperature
Management
Unit

TDB

sensor

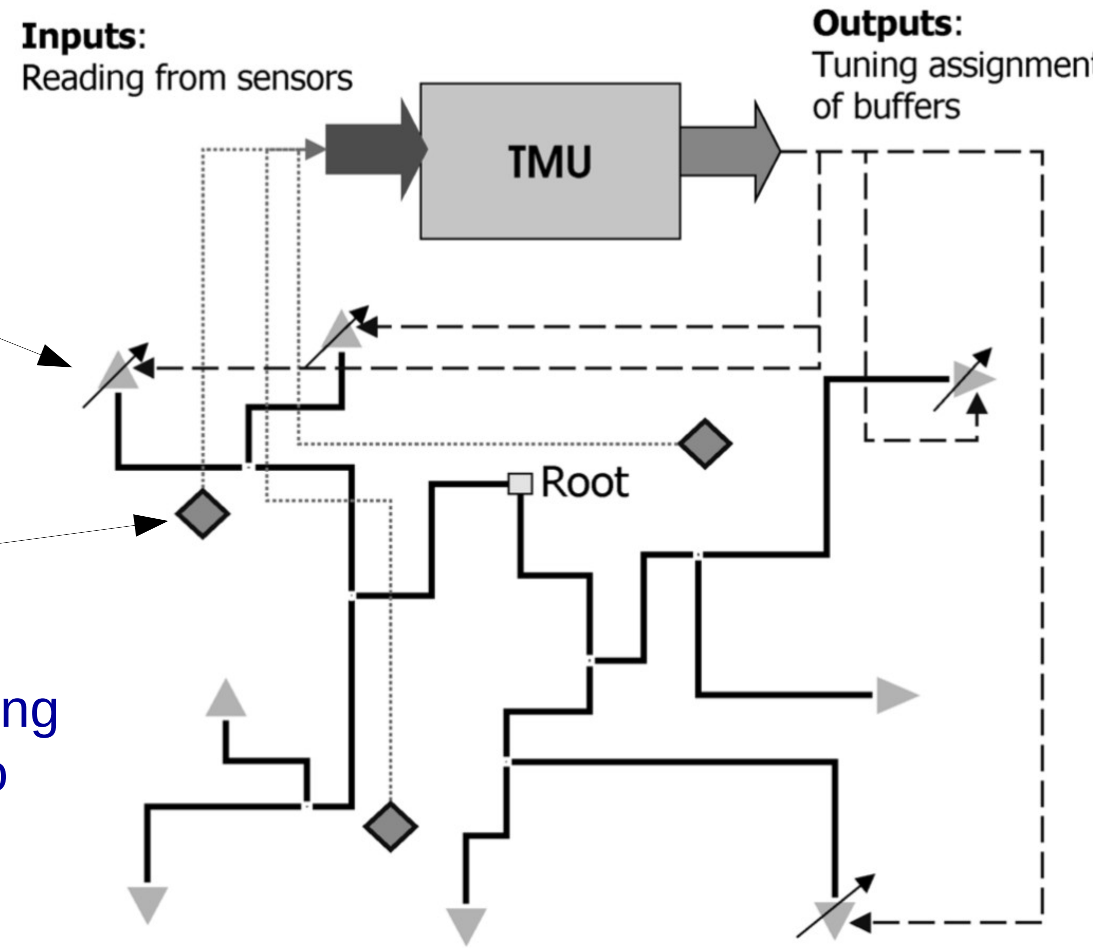

Clock skew compensation using a thermal management unit to control tunable delay buffers inserted into clock tree

(src.:[Chakraborty, 2008])

Temperature and Leakage

- Thermal "runaway" problem:
  - Increase in temperature leads to increase in leakage power → feedback loop possible!
  - Sub-threshold leakage approximated by

$$I_{sub} \approx A \times e^{-\frac{B}{T}}$$

where *A* and *B* are constants → exponential growth!

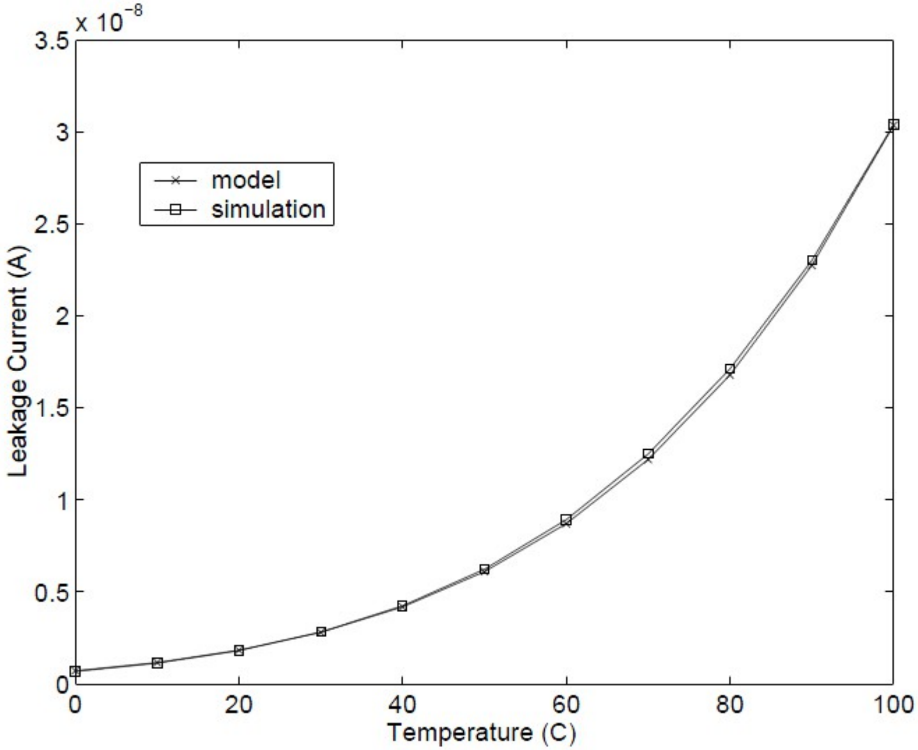

(src.: [Zhang 2003])

Thermal Design Power

- The thermal design power (TDP): maximum amount of heat generated by the CPU that the cooling system in a computer is required to dissipate in typical operation.
- controversial
  - TDP is not maximum power
  - no agreement on measurement setup
- used as an estimate to provison cooling

- Typical TDPs
  - 486DX2 : ~6W (old)
  - 3205U : 15W (mobile)
  - Core i7-5960X : 140W
  - Geforce GTX 1080 : 180W

Cooling Methods

- Heat sink
  - Passive cooling element designed to maximize surface area of heat dissipation
- Air cooling
- Liquid cooling
- etc…

Actively increase convection away from heat source.
Commonly used together with heat sink

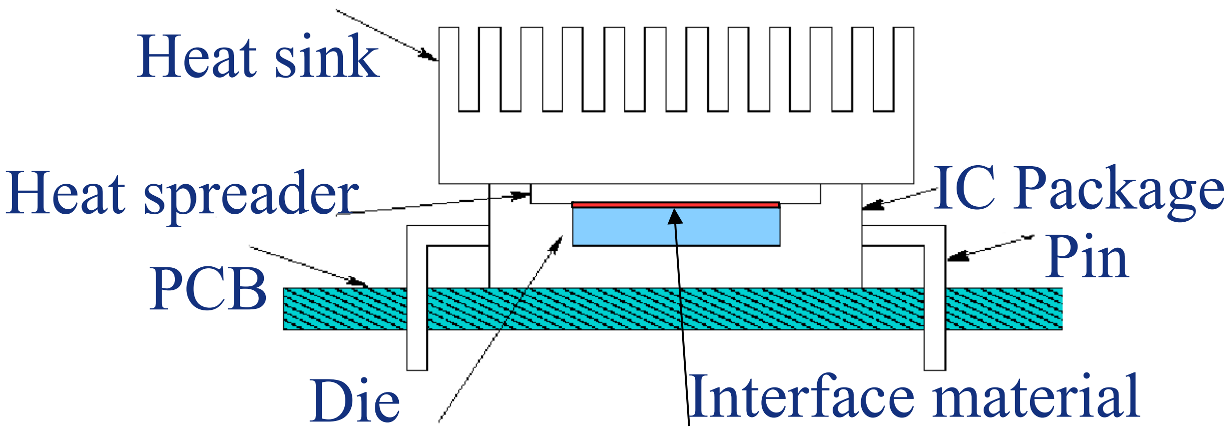

[K. Skadron, 2005]

An exemplaric cooling system

Be Quiet! Dark Rock Pro 3 Silentwings CPU Cooler

(src.:www.bequiet.com)

An exemplaric cooling system (cont'd)

An exemplaric cooling system (cont'd)

Karlsruhe Institute of Technology

- Dimensions (LxWxH in mm): 150x137x163
- Weight: 1.197kg
- Power: 250W
- Noise: 26.1dB
- Price: ~105$ (amazon.com)
- Price of Intel Core i7-6700K: ~350€ (amazon.de)

Heatpipe

- heat-transfer device
- combines thermal conductivity and phase transition to efficiently manage the transfer of heat
- At hot interface: liquid in contact with a thermally conductive solid surface turns into a vapor
- Vapor travels along heat pipe to the cold interface and condenses back into a liquid – releasing the latent heat.
- Liquid returns to hot interface through either capillary action, centrifugal force, or gravity
- Very high heat transfer coefficients for boiling and condensation → heat pipes are highly effective thermal conductors

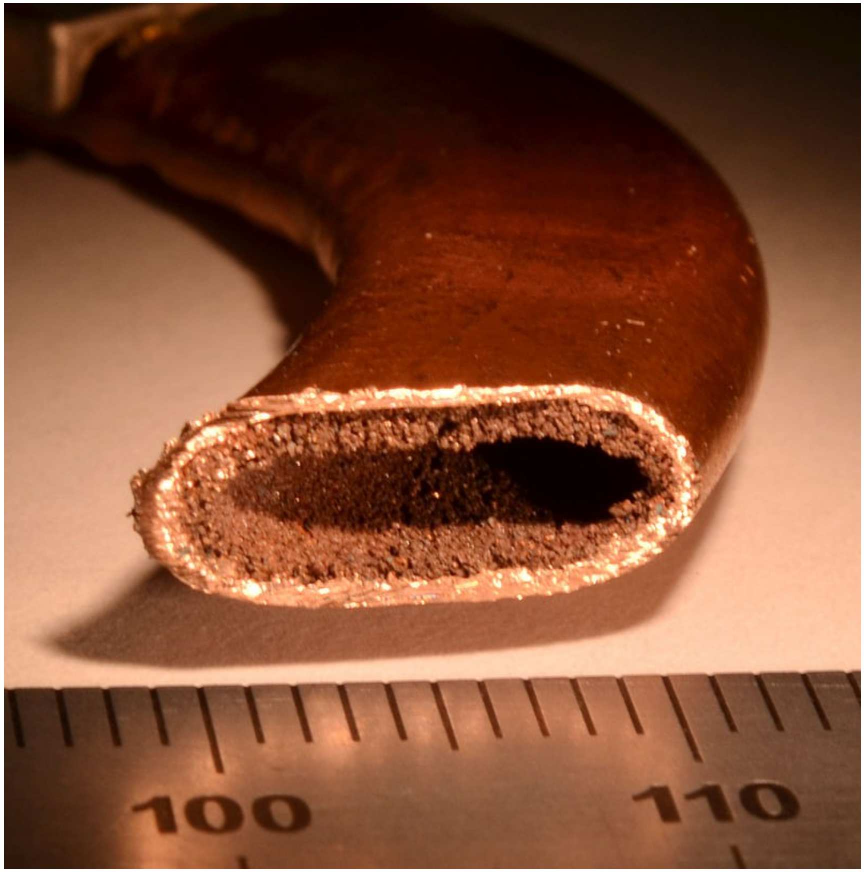

src.:en.wikipedia.org

Heatpipe (cont'd)

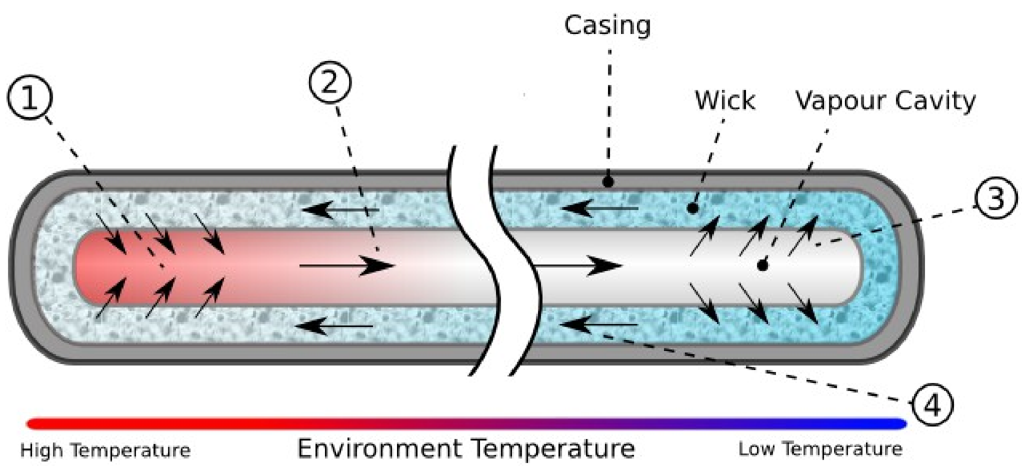

Heat pipe thermal cycle

1) Working fluid evaporates to vapour absorbing thermal energy.
2) Vapour migrates along cavity to lower temperature end.
3) Vapour condenses back to fluid and is absorbed by the wick, releasing thermal energy
4) Working fluid flows back to higher temperature end.

src.:en.wikipedia.org

DVFS and temperature

Frequency (MHz)	Voltage (V)
2000	1.340
1800	1.292
1600	1.244
1400	1.196
1200	1.148
1000	1.100
800	1.052
600	0.988

Table 1: P-States

Figure 1: Pentium M (left) and data acquisition (right)

(src.: [Hanson 2007])

DVFS and temperature (cont'd)

Most straightforward technique since power reduction decreases temperatures

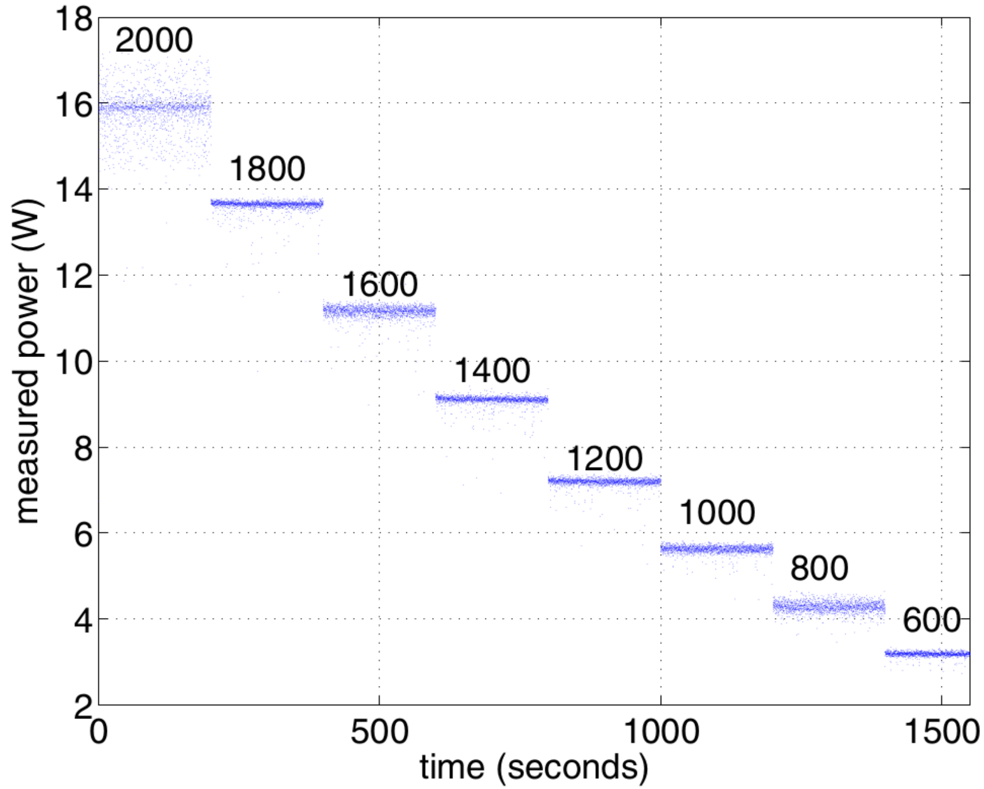

Figure 2: CPU power for daxpy, 200 seconds per p-state (denoted in MHz).

(src.: [Hanson 2007])

DVFS and temperature (cont'd)

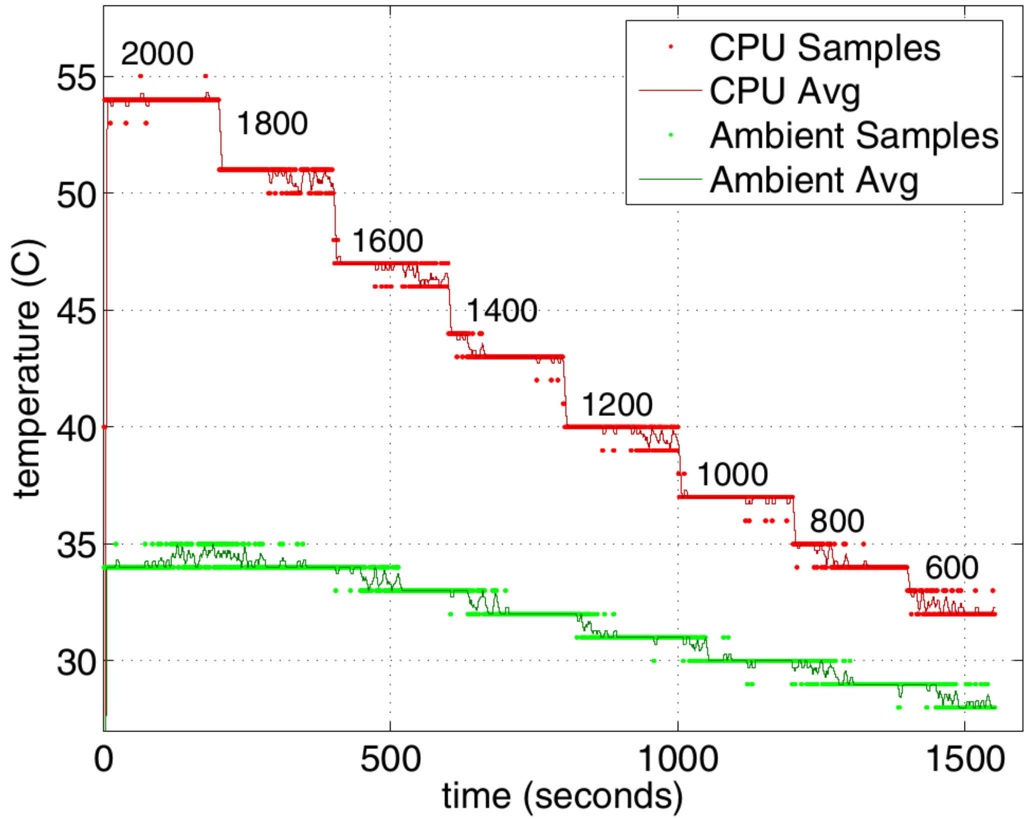

Figure 3: CPU and ambient temperatures for daxpy, 200 seconds per p-state (denoted in MHz).

(src.: [Hanson 2007])

DVFS and temperature (cont'd)

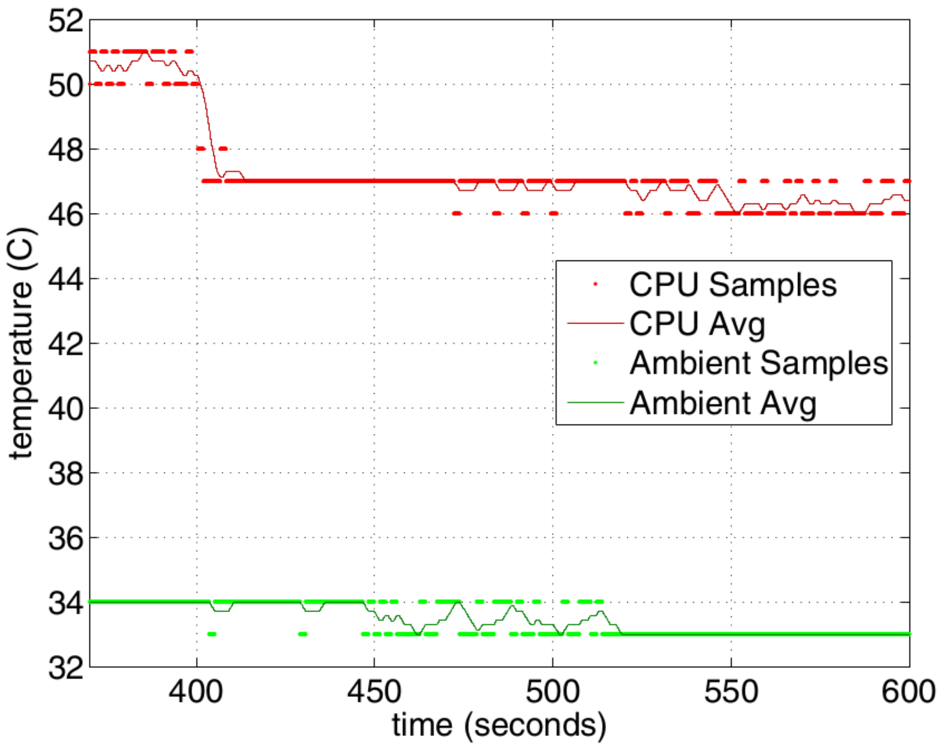

Figure 4: CPU temperature detail for daxpy p-state transition. Temperature adjusts in two stages: initial drop, then additional drift as ambient settles.

(src.: [Hanson 2007])

DVFS and temperature (cont'd)

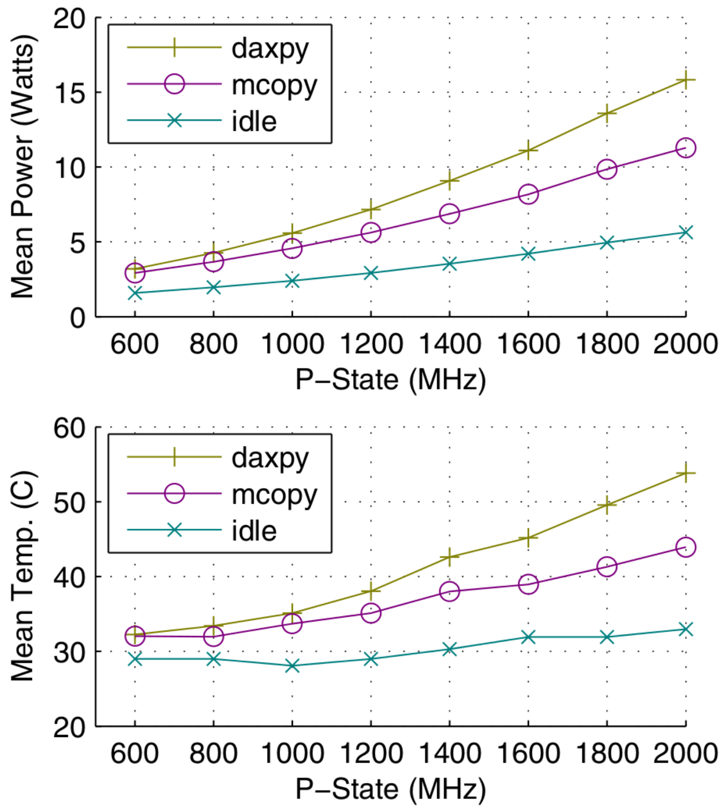

Figure 5: CPU power and temperature vary by benchmark. Power closely tracks p-state; CPU temperature loosely tracks p-state for given benchmark.

(src.: [Hanson 2007])

DVFS and temperature (cont'd)

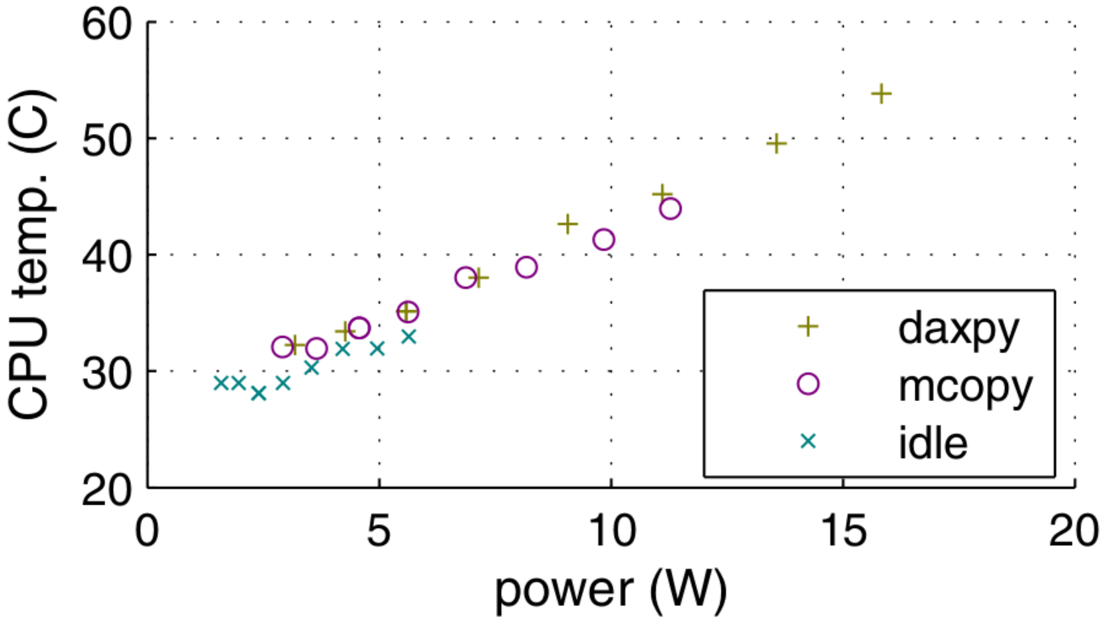

Figure 6: Linear power-thermal relationship under steady-state conditions for single instance of each benchmark and frequency.

(src.: [Hanson 2007])

DVFS and temperature (cont'd)

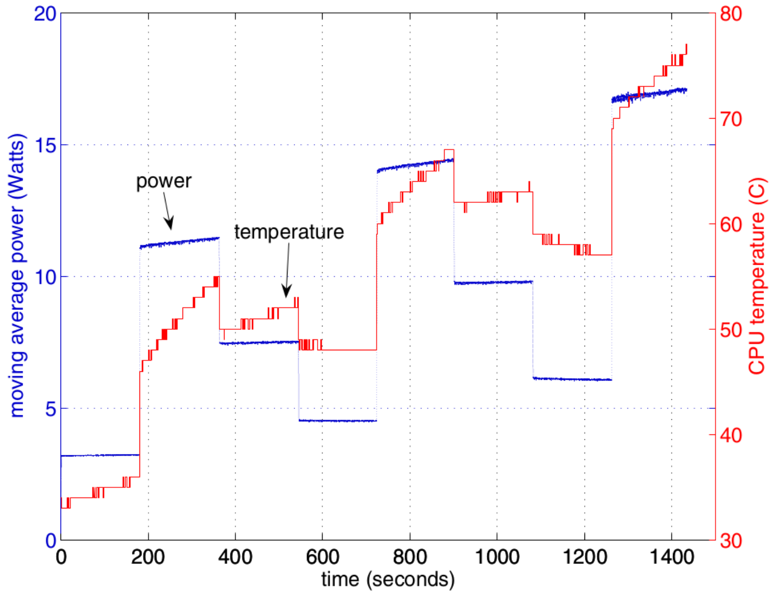

Figure 7: Power and temperature for daxpy benchmark with under-cooled conditions, with 200 seconds of each p-state in order: 600, 1600, 1200, 800, 1800, 1400, 1000, 2000 MHz.

- Effect of running application in multiple power states
- Power changes almost Instantaneous Temperature changes more gradually
- Thermal effect on leakage visible in higher power states

(src.: [Hanson 2007])

DVFS and temperature (cont'd)

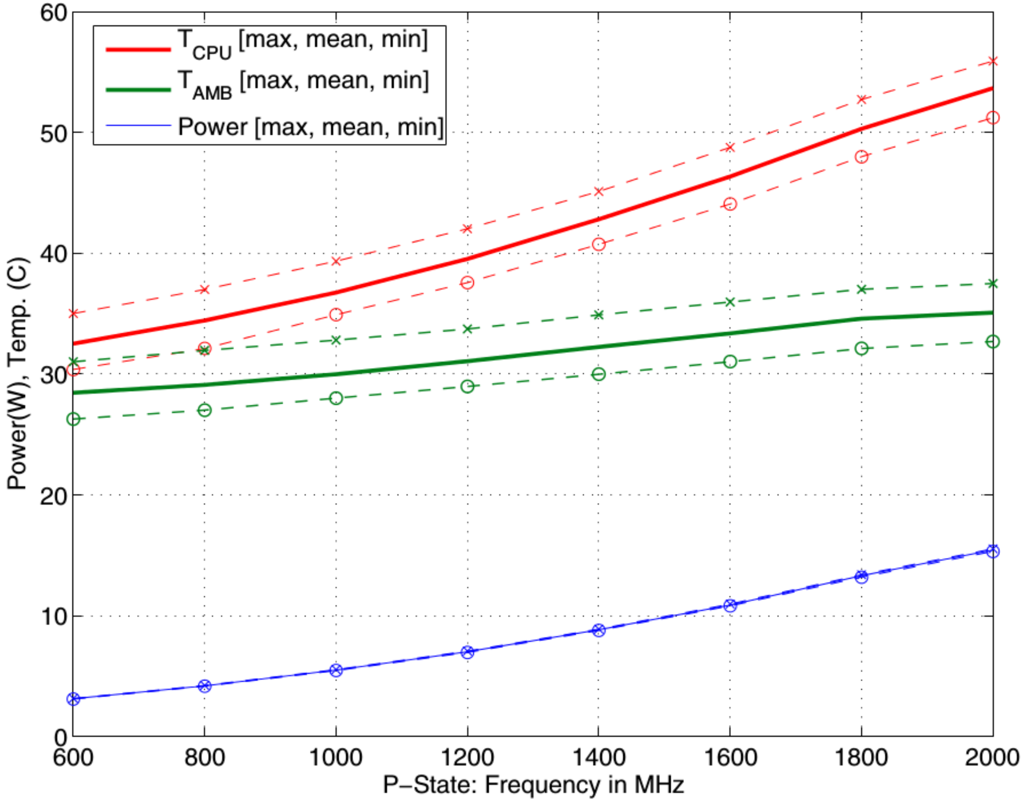

Figure 8: Steady-state power and temperature measurements for multiple invocations of daxpy.

(src.: [Hanson 2007])

DVFS and temperature (cont'd)

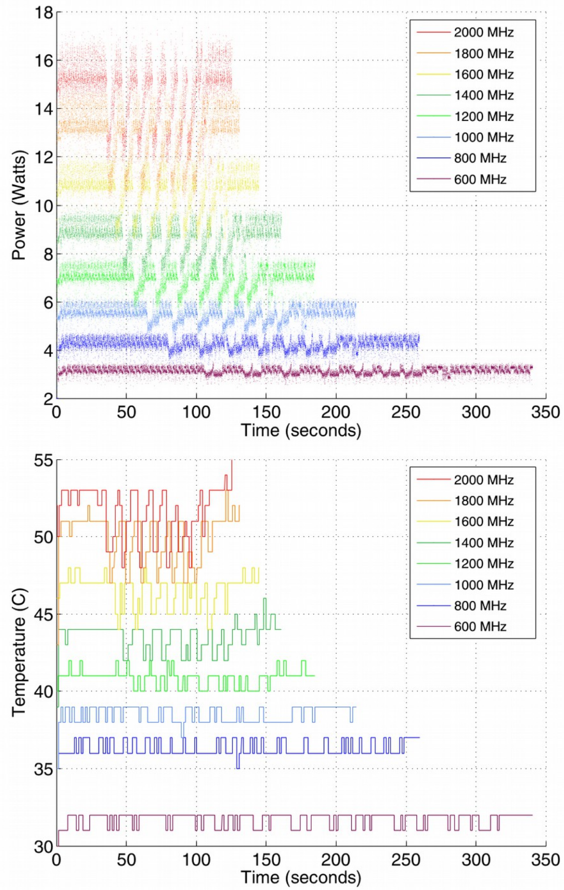

Figure 9: CPU power and temperature for galgel at each DVFS p-state.

(src.: [Hanson 2007])

DVFS and temperature (cont'd)

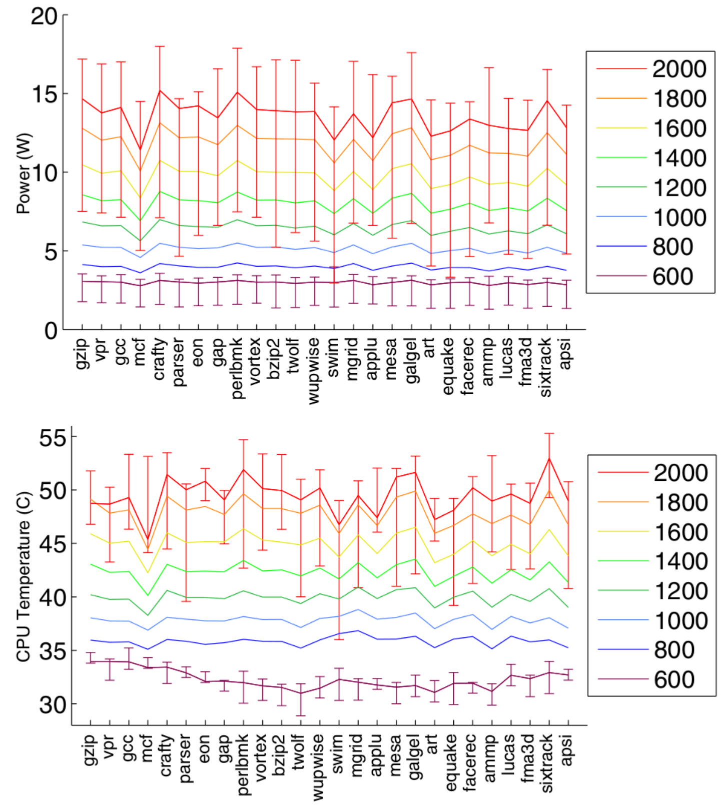

Figure 10: Mean and range of CPU power and temperature for each SPEC CPU2000 benchmark, at each DVFS p-state.

(src.: [Hanson 2007])

DVFS and temperature (cont'd)

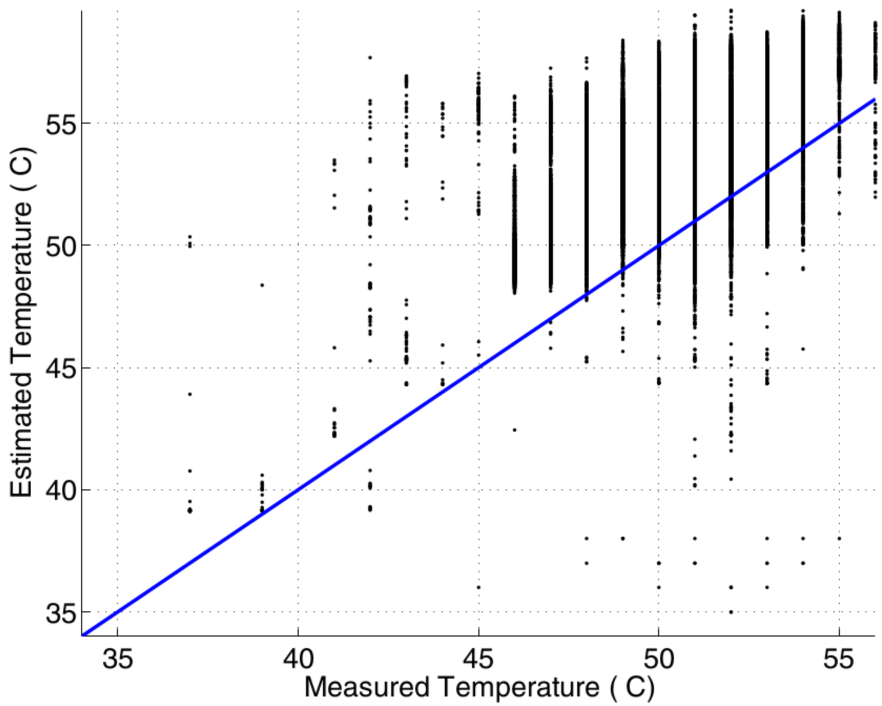

Figure 11: Comparison of estimated and measured CPU temperature.

(src.: [Hanson 2007])

DTM through Task scheduling

- order of task execution influences peak temperature
- Idea: determine a threshold temperature and run “hottest” task for which (predicted) threshold is not reached
- Achieves reduction in power state transitions

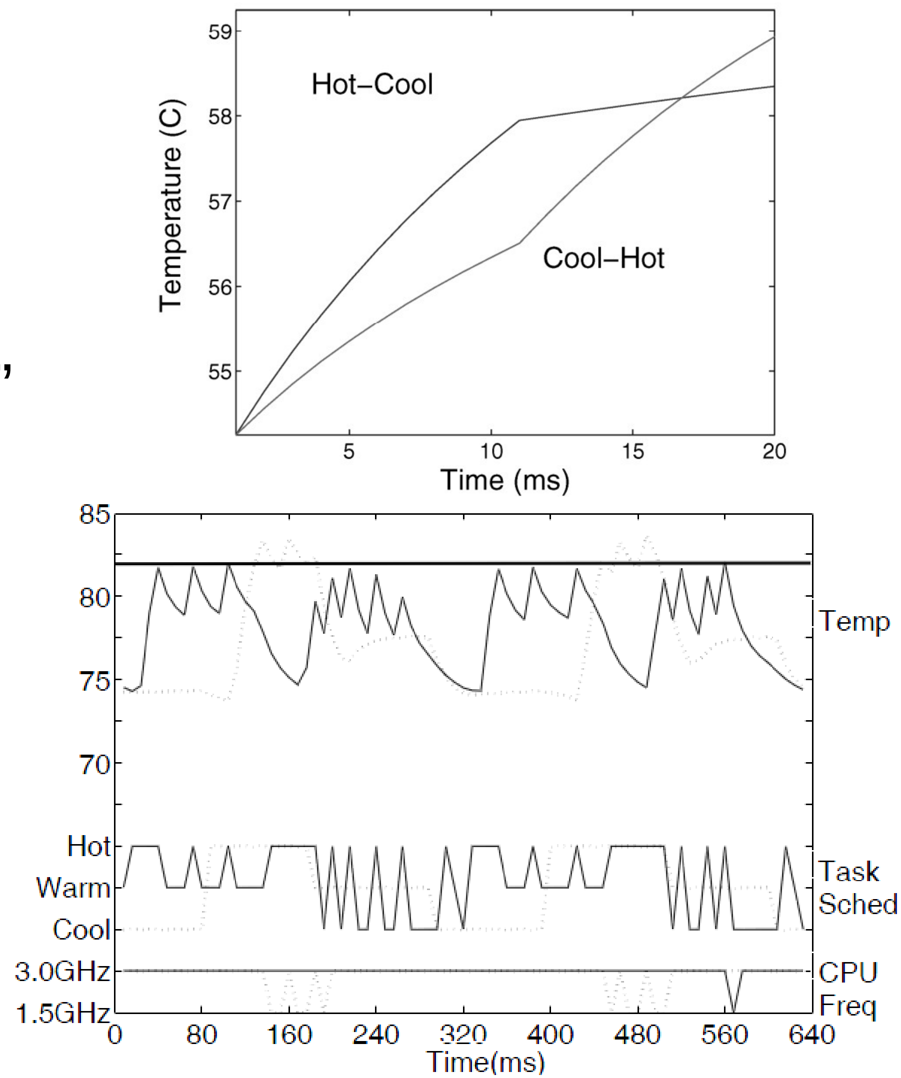

Linux Tools for Temperature Analysis/Control

- sensors (lm-sensors): view temperatures from thermal diodes
- cpufreq-info: view cpu frequencies
- ksysguard: plot frequencies, temperatures, load, and many more over time
- cpu-burn, openssl speed: exemplaric CPU benchmarks
- taskset: pin processes to specific cores
- Linux Scaling Governors
  - Ondemand
  - Powersave
  - Conservative
  - Performance

Sources

David Halliday, Robert Resnick, Jearl Walker Fundamentals of Physics John Wiley & Sons, 2010

Paul A. Tipler, Gene Mosca Physics for Scientists and Engineers W. H. Freeman, 2007

[Hanson 2007] H. Hanson et al, Thermal Response to DVFS: Analysis with an Intel Pentium M, ISLPED 2007

[Xie 2006] V. Narayanan et al, Reliability Concerns in Embedded System Designs. Computer 39, 1 (Jan. 2006).

[Yang 2008] J.Yang et al, Dynamic Thermal Management through Task Scheduling, ISPASS 2008

[Zhang 2003] Y.Zhang et al, HotLeakage: A Temperature-Aware Model of Subthreshold and Gate Leakage for Architects, Tech Report 2003

[Chakraborty 2008] A.Chakraborty et al, Dynamic Thermal Clock Skew Compensation Usint Tunable Delay Buffers, TVLSI 2008

[Skadron 2005] K. Skadron, A Quick Thermal Tutorial, Invited talk.