Analysis of The Enhanced Intel® Speedstep® Technology of the Pentium® M Processor

Efi Rotem, Avi Mendelson, Alon Naveh,
Micha Moffie

June 2004

MOBILE
TECHNOLOGY

Overview

- Mobile computers challenges
  - Energy and Average power
  - Thermal Design Power
- Intel® Pentium® M power management
- The experiments
- Test results
- Conclusions

Overview

- Mobile computers challenges
  - Energy and Average power

  - Thermal Design Power

- Intel Pentium M power management
- The experiments
- Test results
- Conclusions

Power & Density increases

- Think Watts/Cm²
  - Complex architecture at faster and smaller technology
  - Denser power is harder to cool

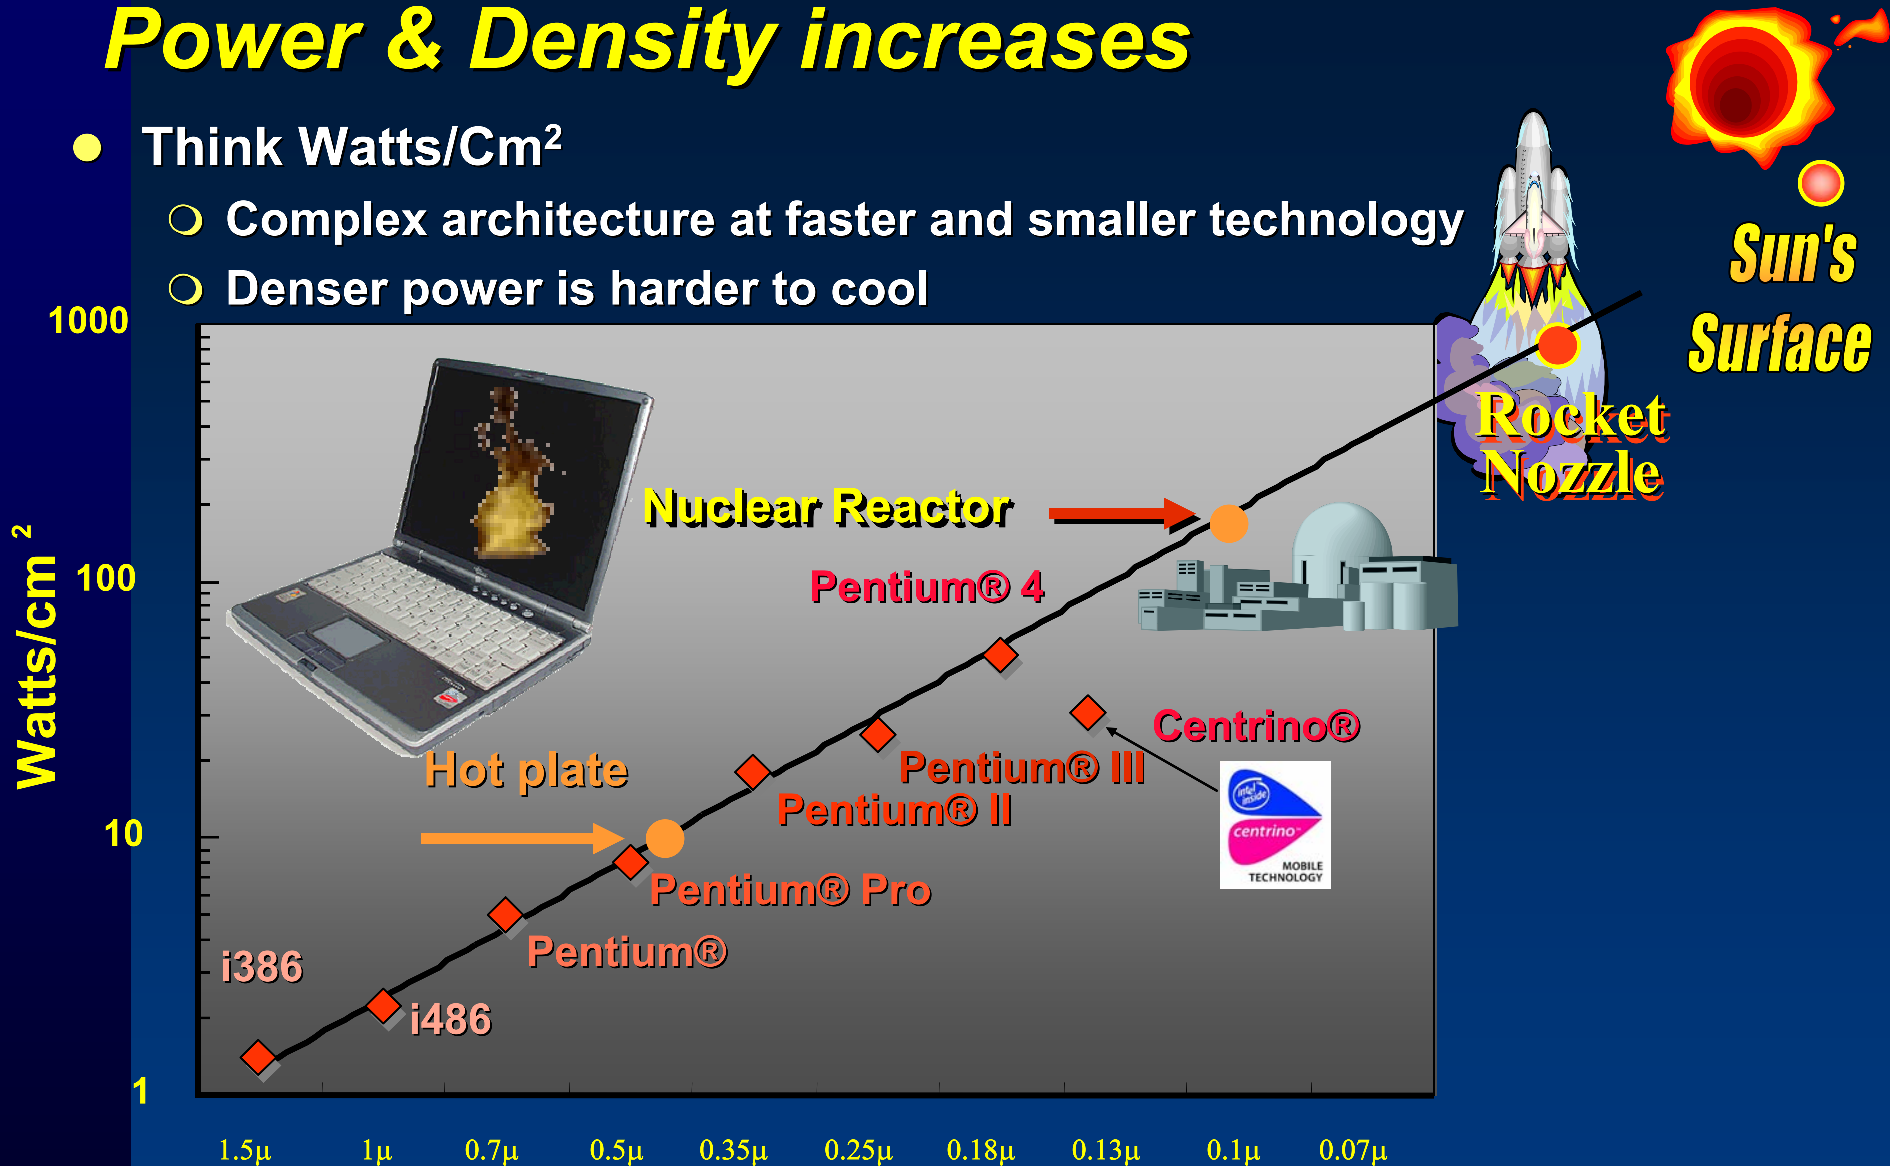

* "New Microarchitecture Challenges in the Coming Generations of CMOS Process Technologies" – Fred Pollack, Intel Corp. Micro32 conference key note - 1999.

The Mobile Environment

- Maximize performance & features within given constraints
  - Power / Thermal
  - Size – form factor
  - Noise
  - Energy / Battery life
  - etc.

- Mobile platforms offer Tradeoff preferences
  - User defined or built-in scheme
    - Compromise performance for longer battery life, lower acoustic noise, cool box etc.

Overview

- Mobile computers challenges
  - Average power
  - Thermal Design Power
- Intel Pentium M power management
- The experiment
- Test results
- Conclusions

Die Temperature measurements

- Temperature based control mechanisms
- A diode connected to an external A/D - reports temperature
- Fixed temperature sensors
  - Max spec junction temperature
  - Critical temperature detector

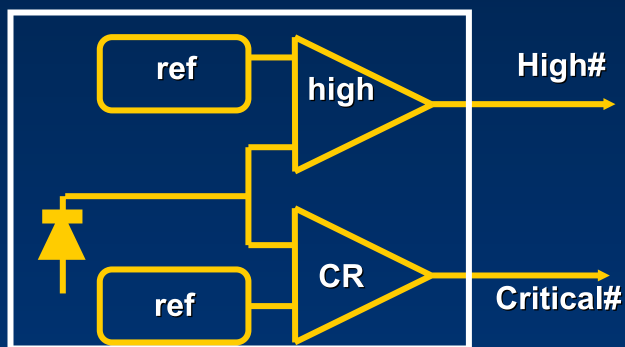

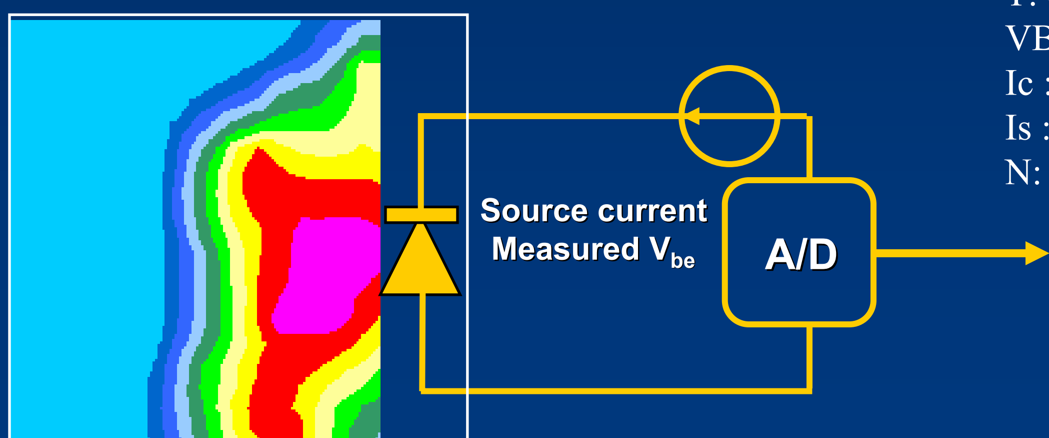

$$T = \frac{\Delta V_{BE}}{n \cdot \ln(N) \cdot \frac{k}{q}}$$

n: diode ideality factor
k: Boltzmann constant
q: electron charge constant
T: diode temperature (Kelvin)
VBE: Voltage from base to emitter
Ic : Collector current
Is : Saturation current
N: Ratio of collector currents

Controlling TDP - Linear

- P states and T states

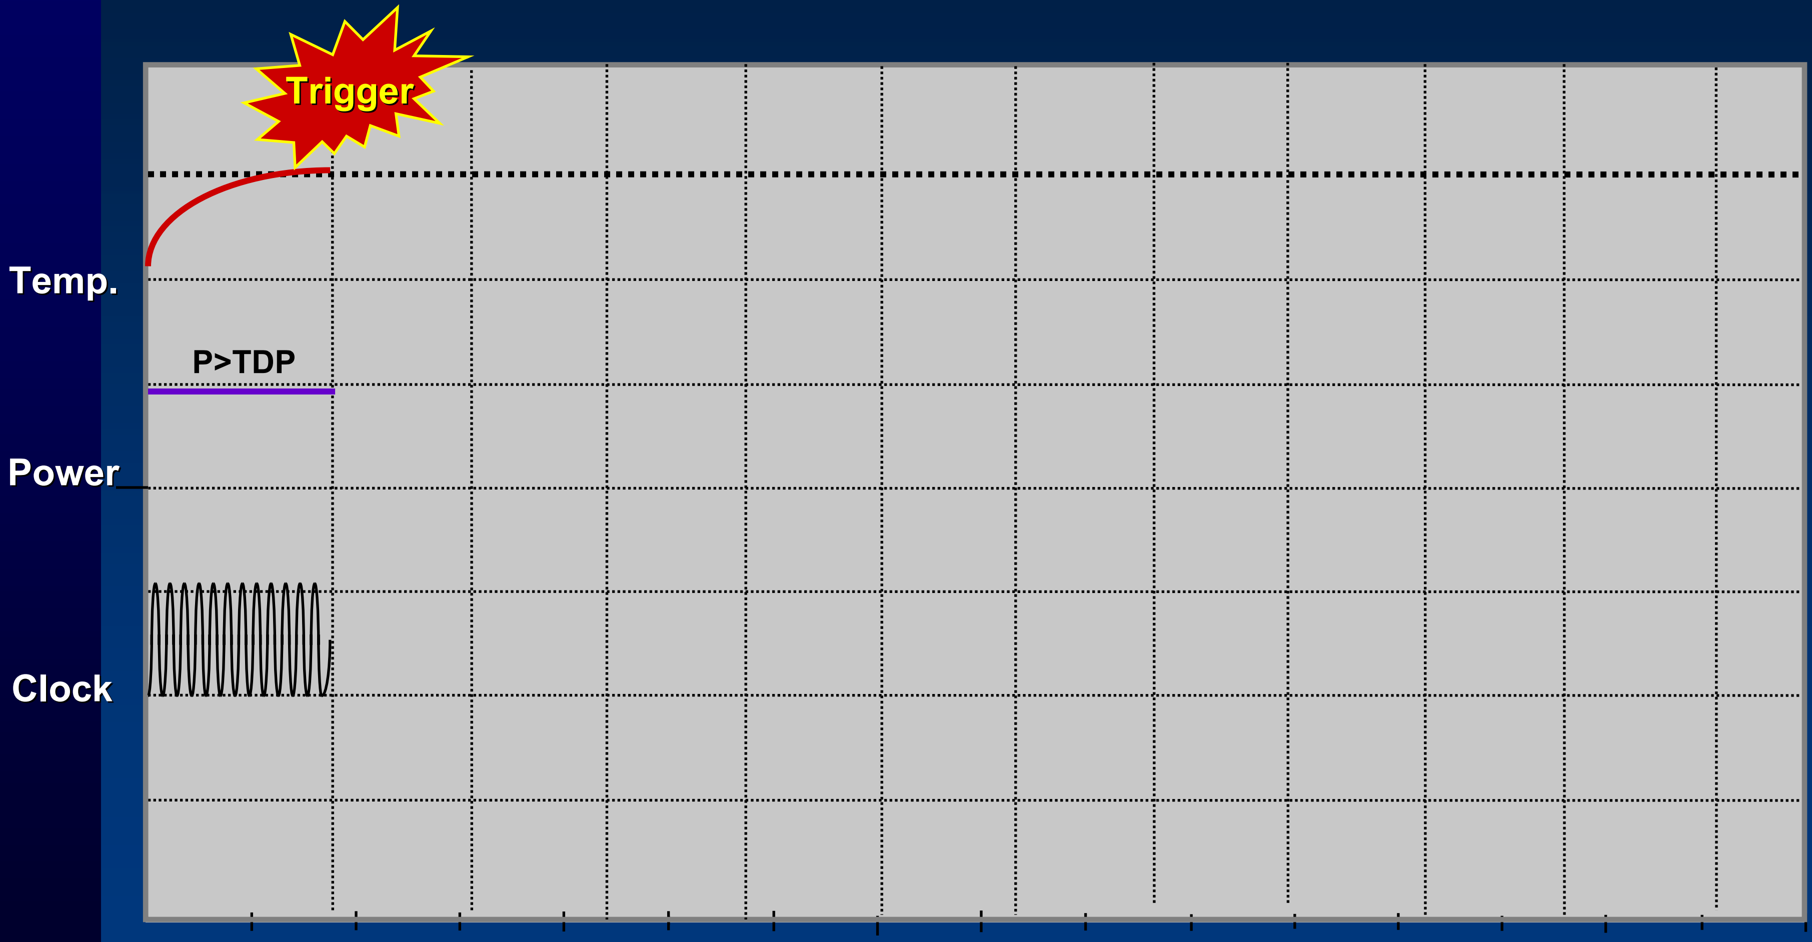

Controlling TDP - Linear

- **Self or S/W trigger – Stop clock asserted, Power and temperature drop**

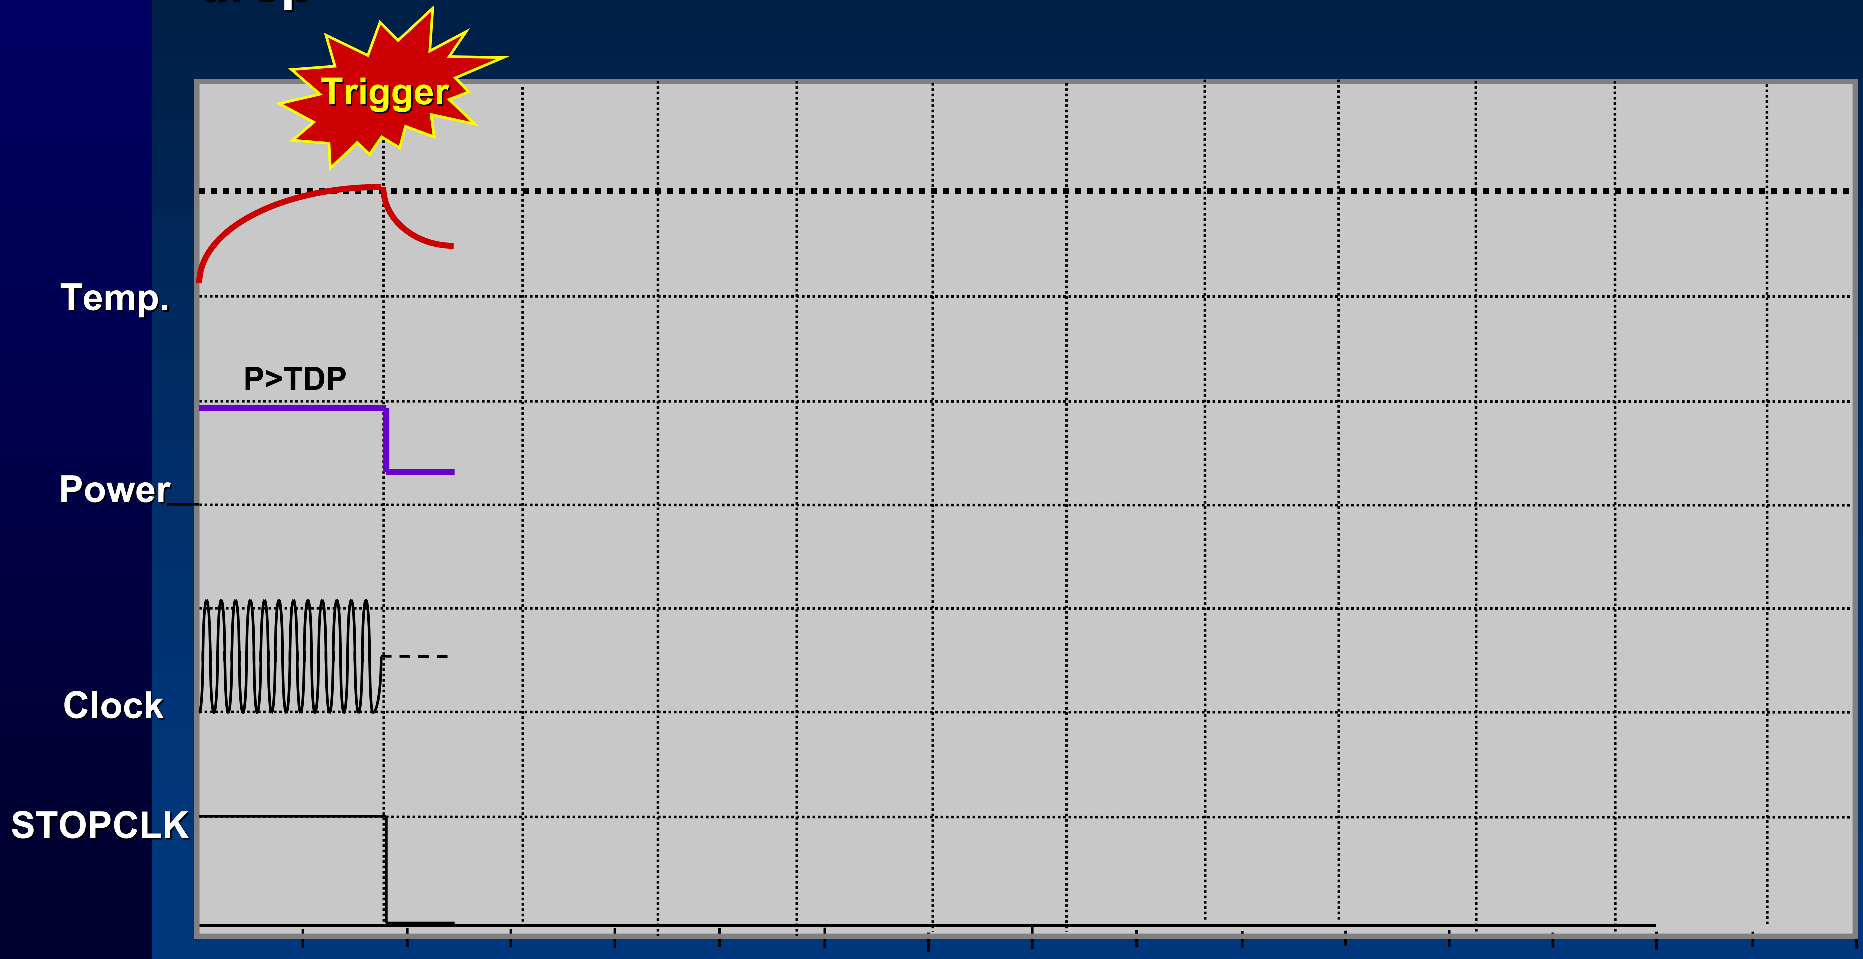

Controlling TDP - Linear

- STOPCLOCK continues toggling for a pre defined time
- Average power and temperature drop – performance drops

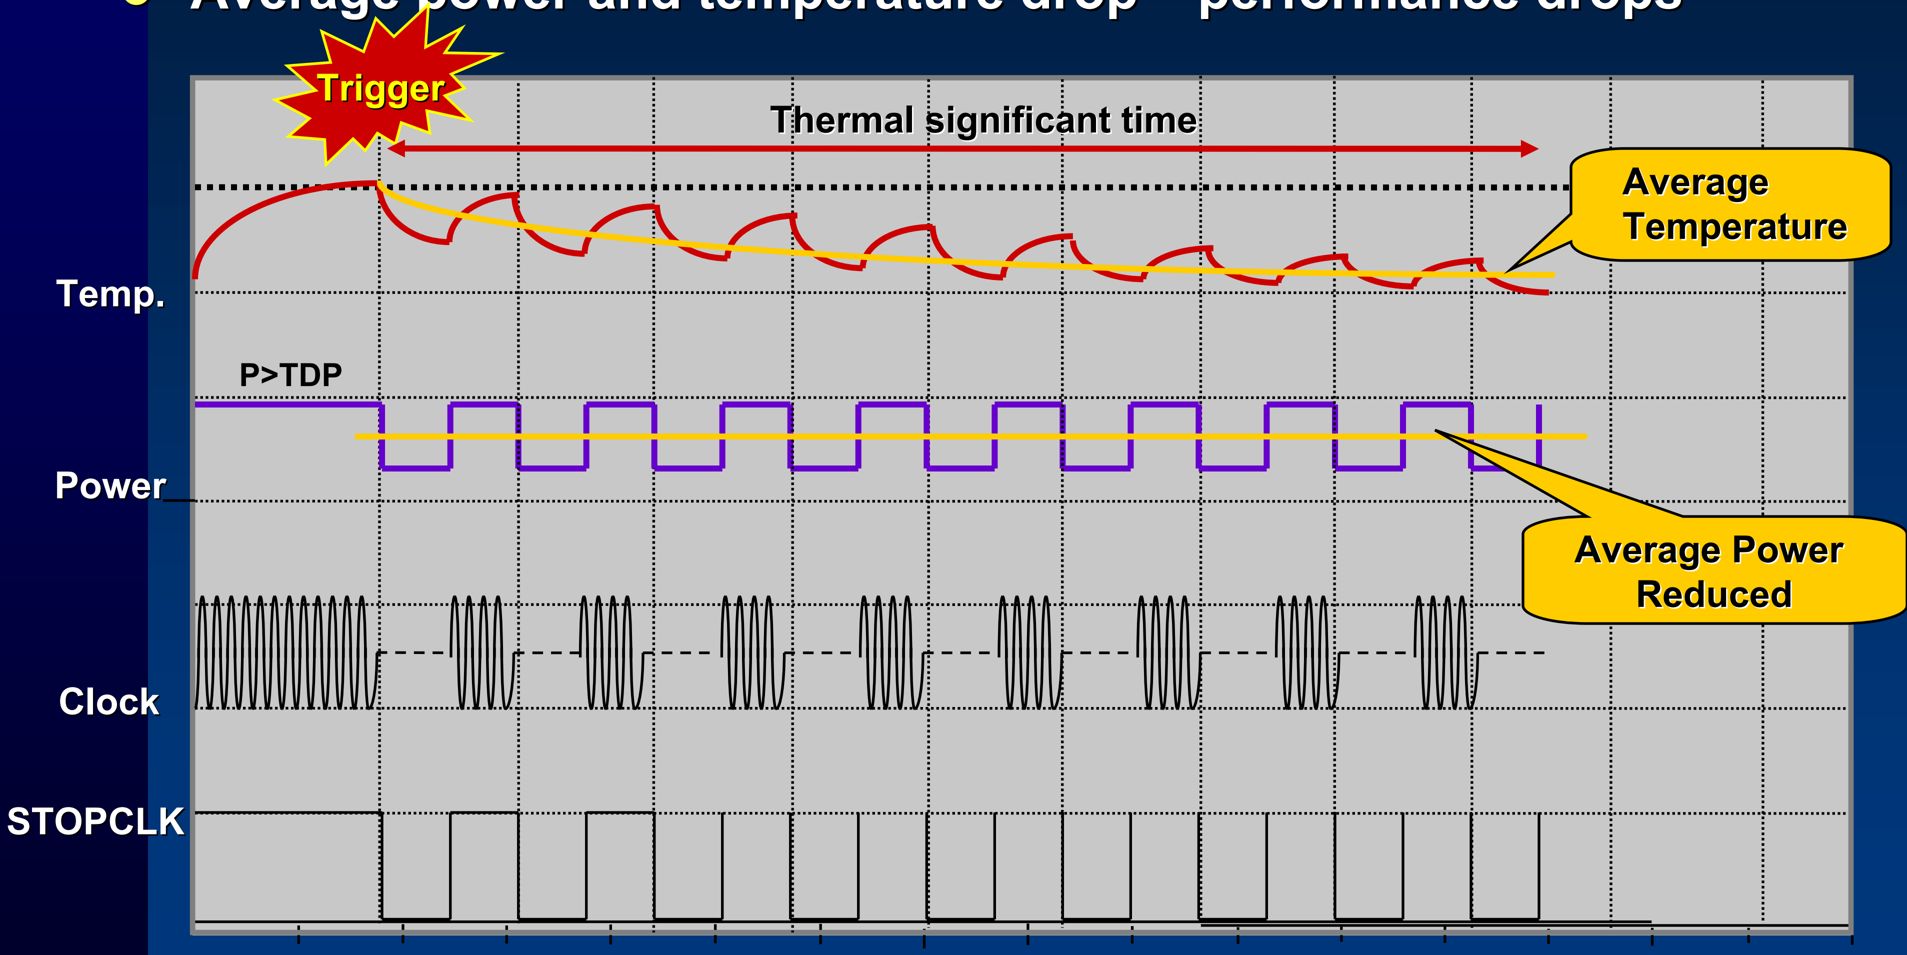

Controlling TDP - Linear

- STOPCLOCK continues toggling for a pre defined time
- Average power and temperature drop – performance drops

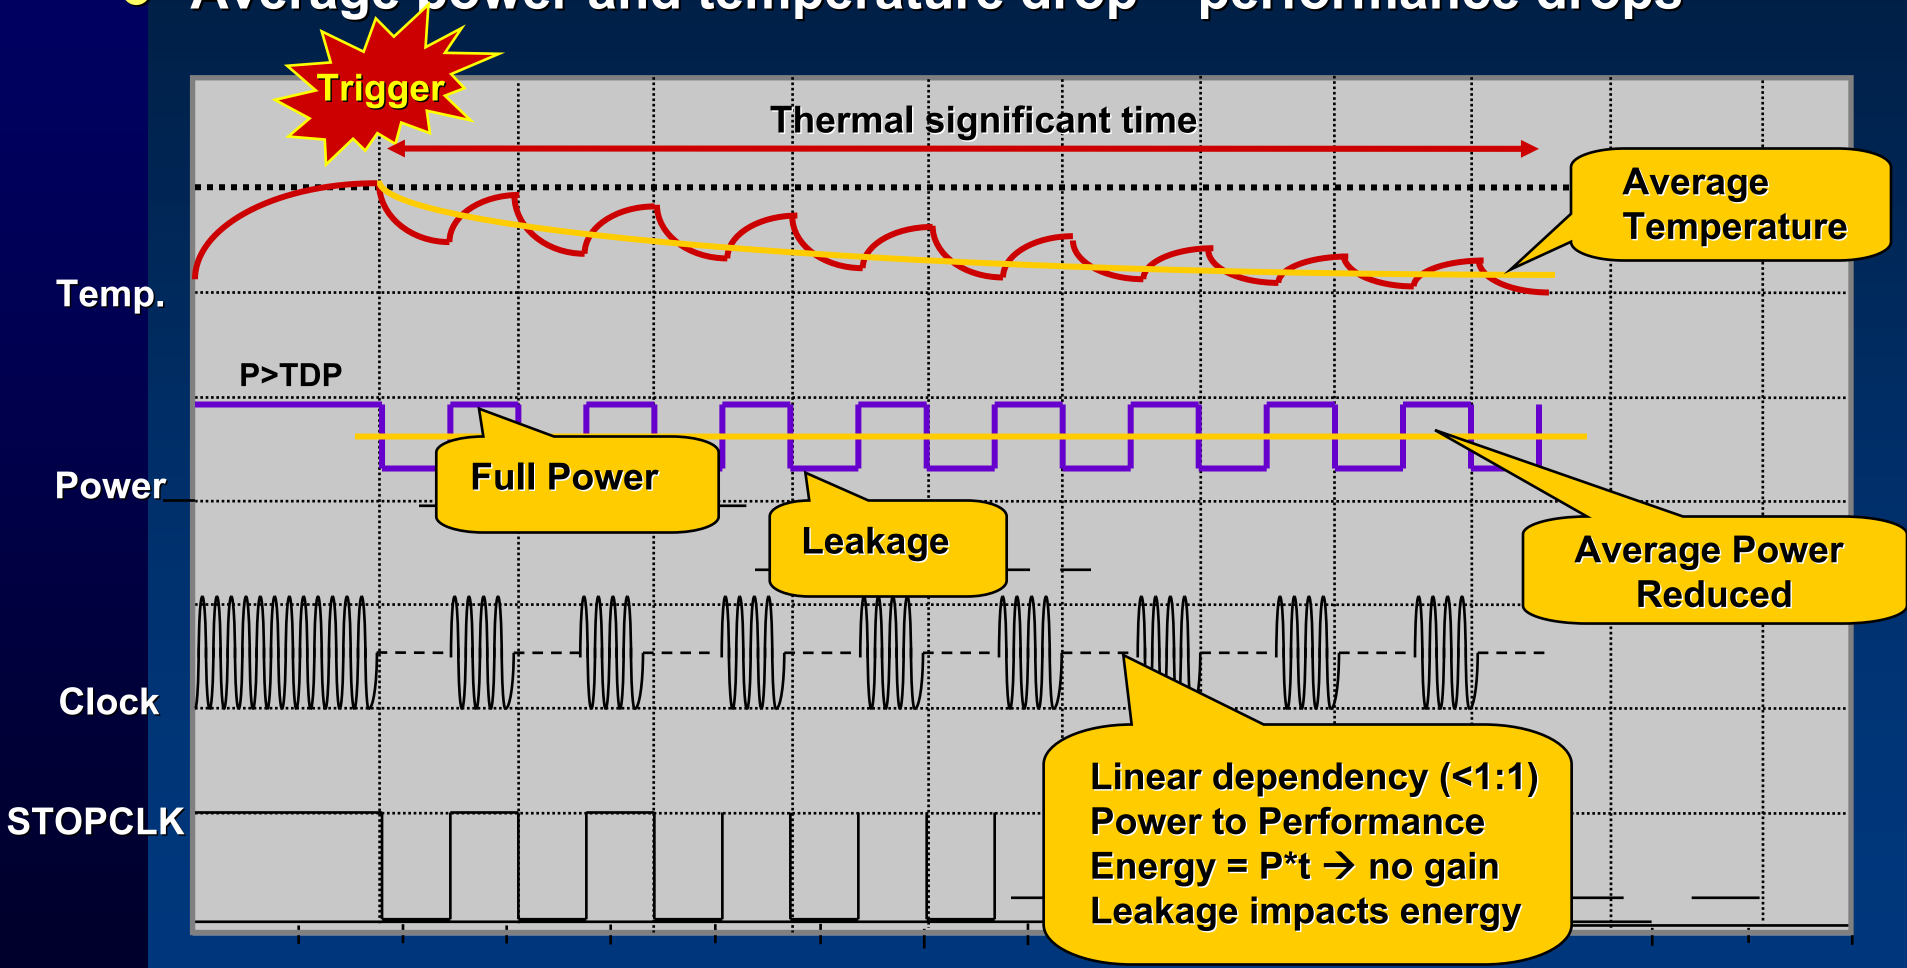

Dynamic Voltage Scaling (DVS)

- One μ ARCH implementation – Power and energy control

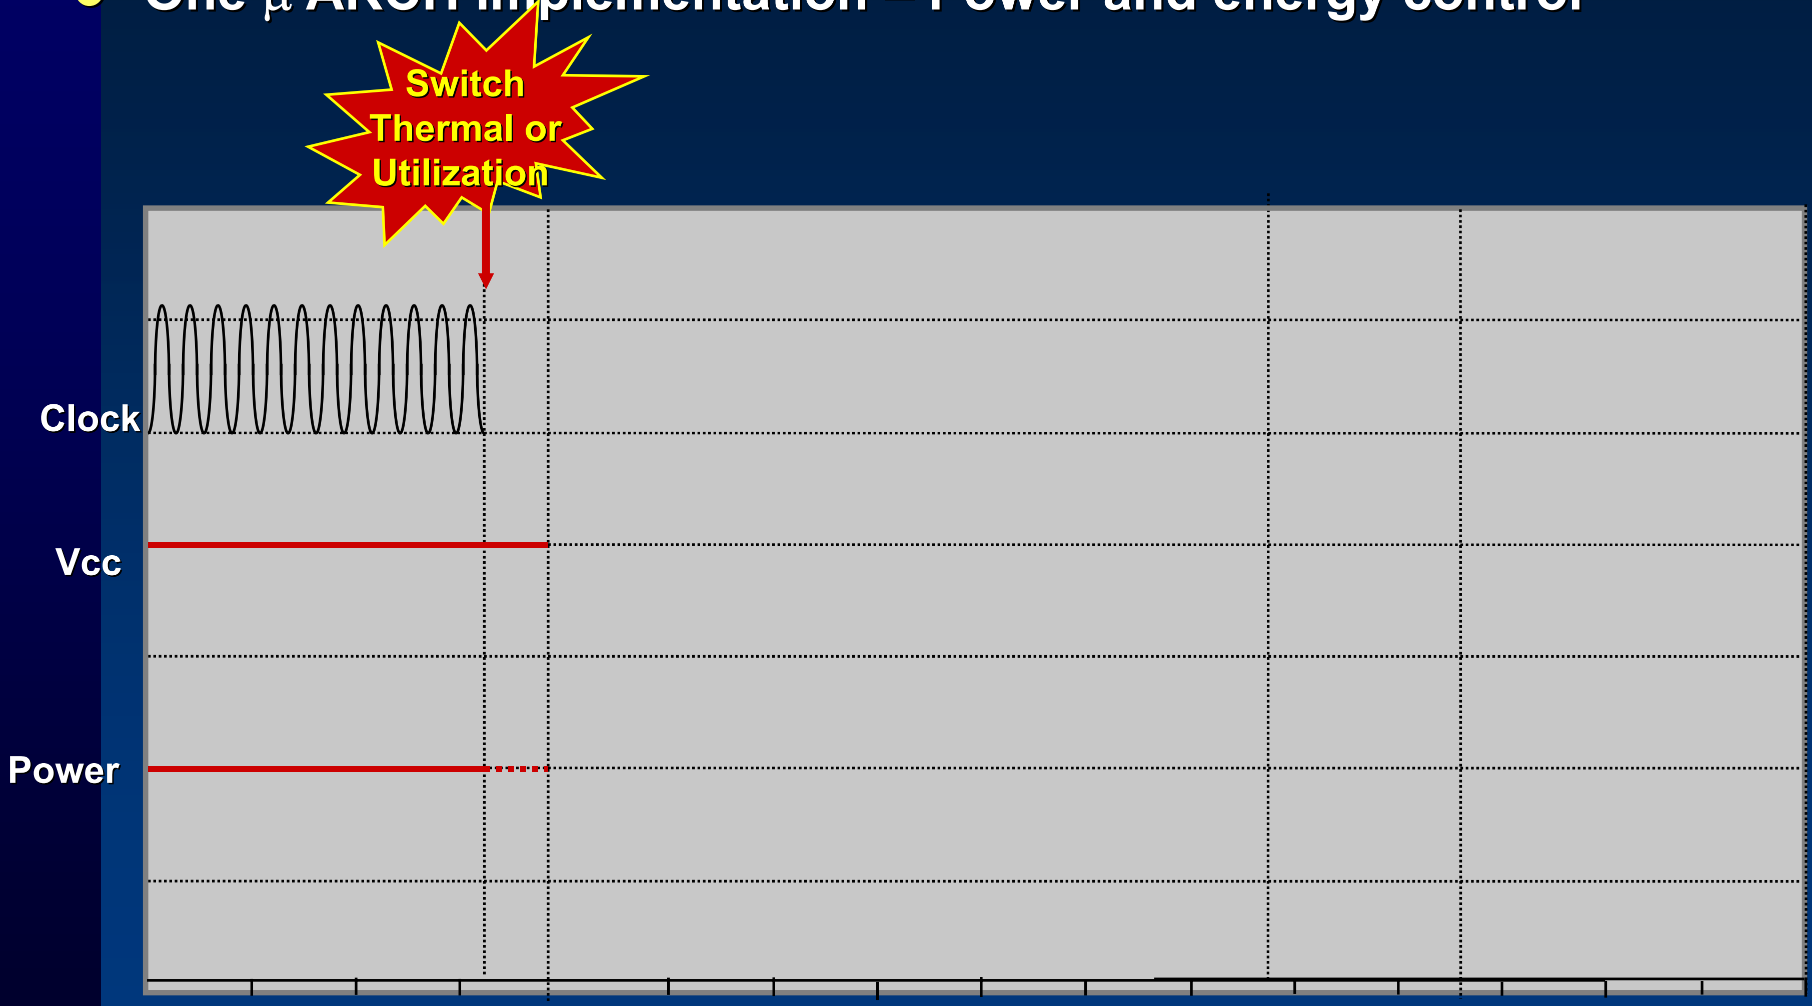

Dynamic Voltage Scaling (DVS)

- PLL relock at lower frequency at same Vcc
- Fast change – no user experience impact

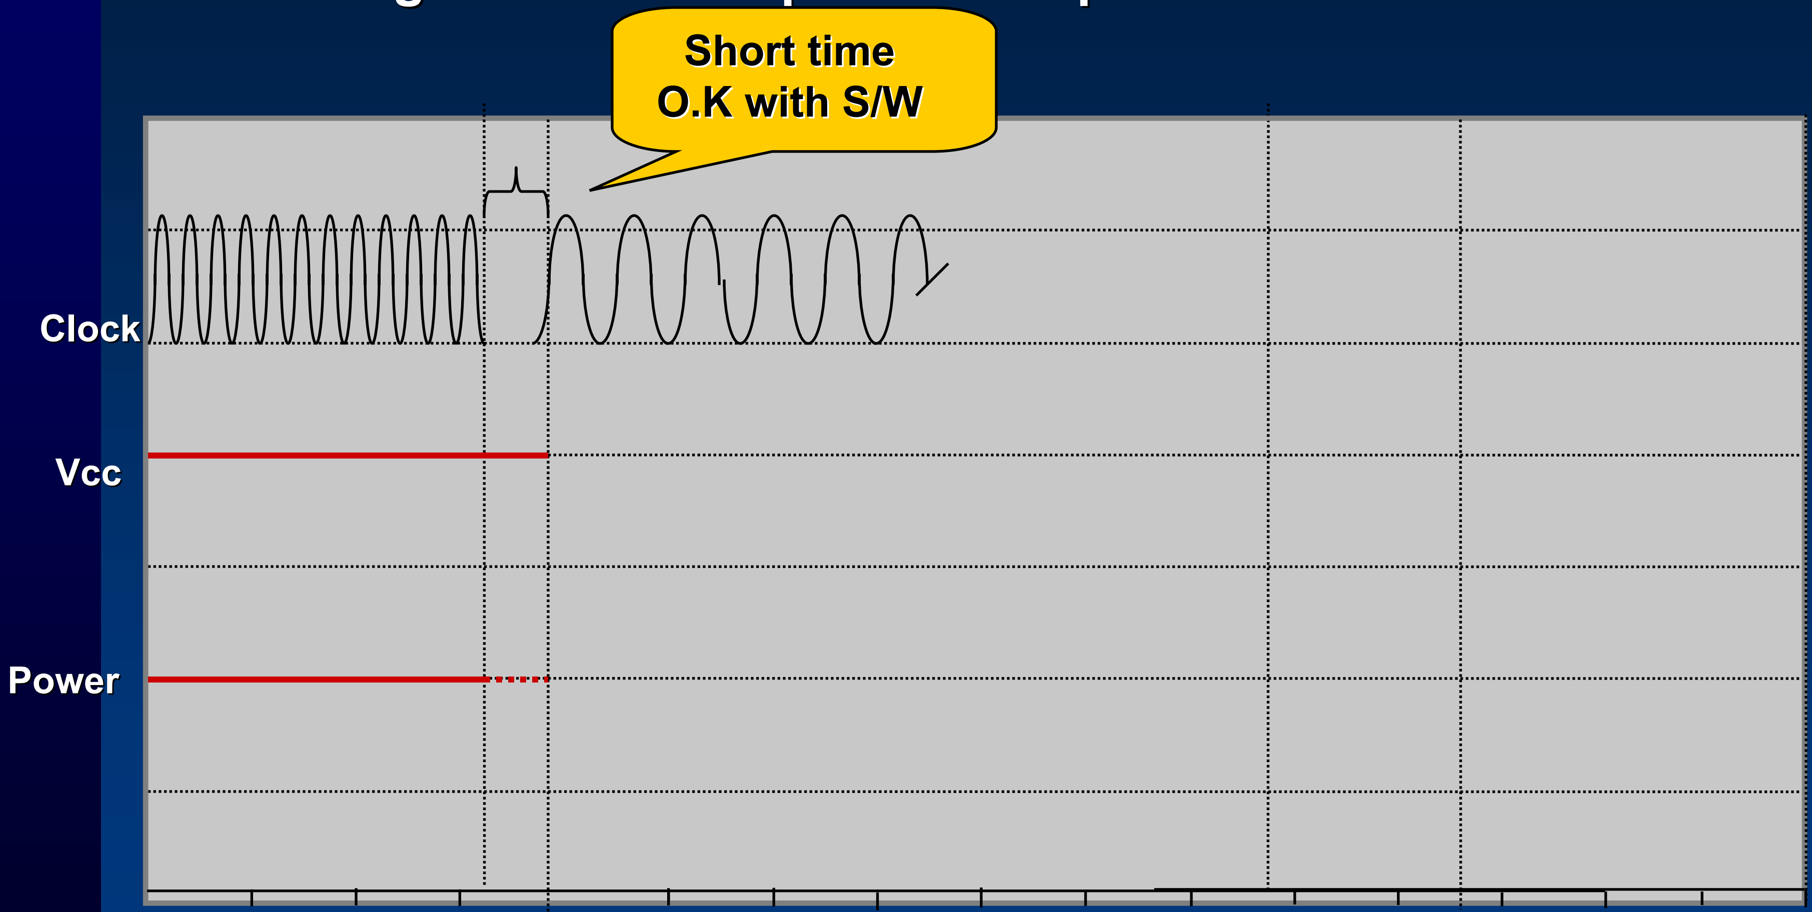

Dynamic Voltage Scaling (DVS)

- Vcc drops gradually while CPU active
- Power savings changes from linear to F^3

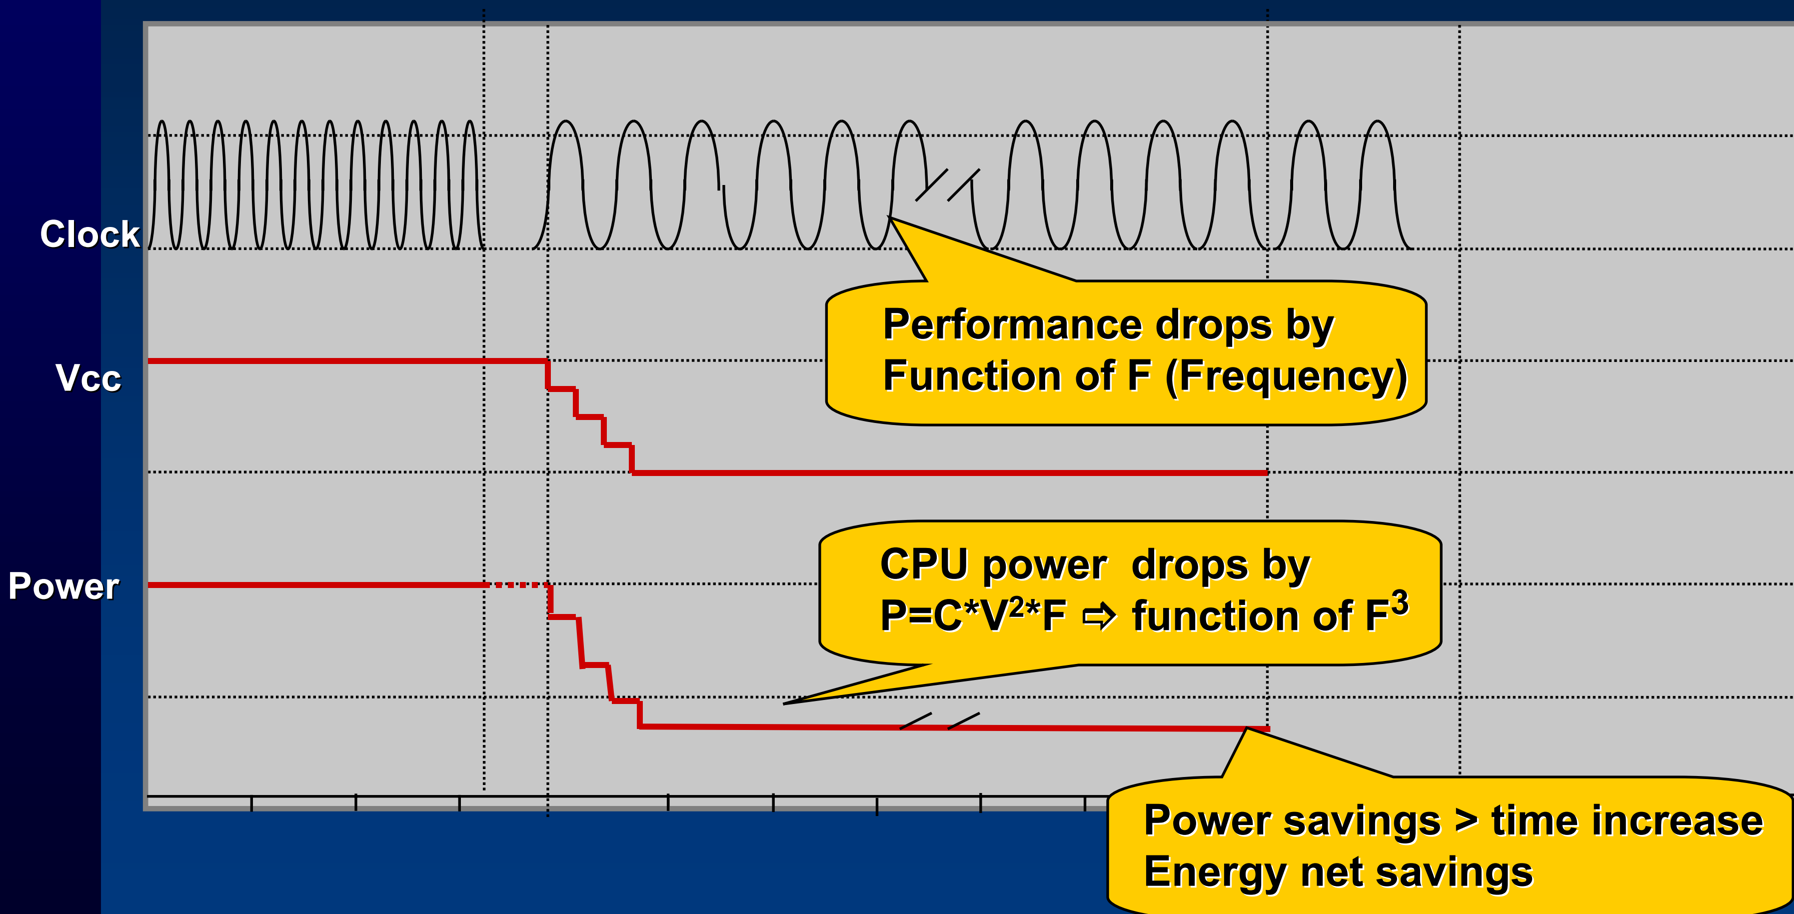

Dynamic Voltage Scaling (DVS)

- Vcc is ramped up increasing power
- Once stable – PLL relock at high frequency

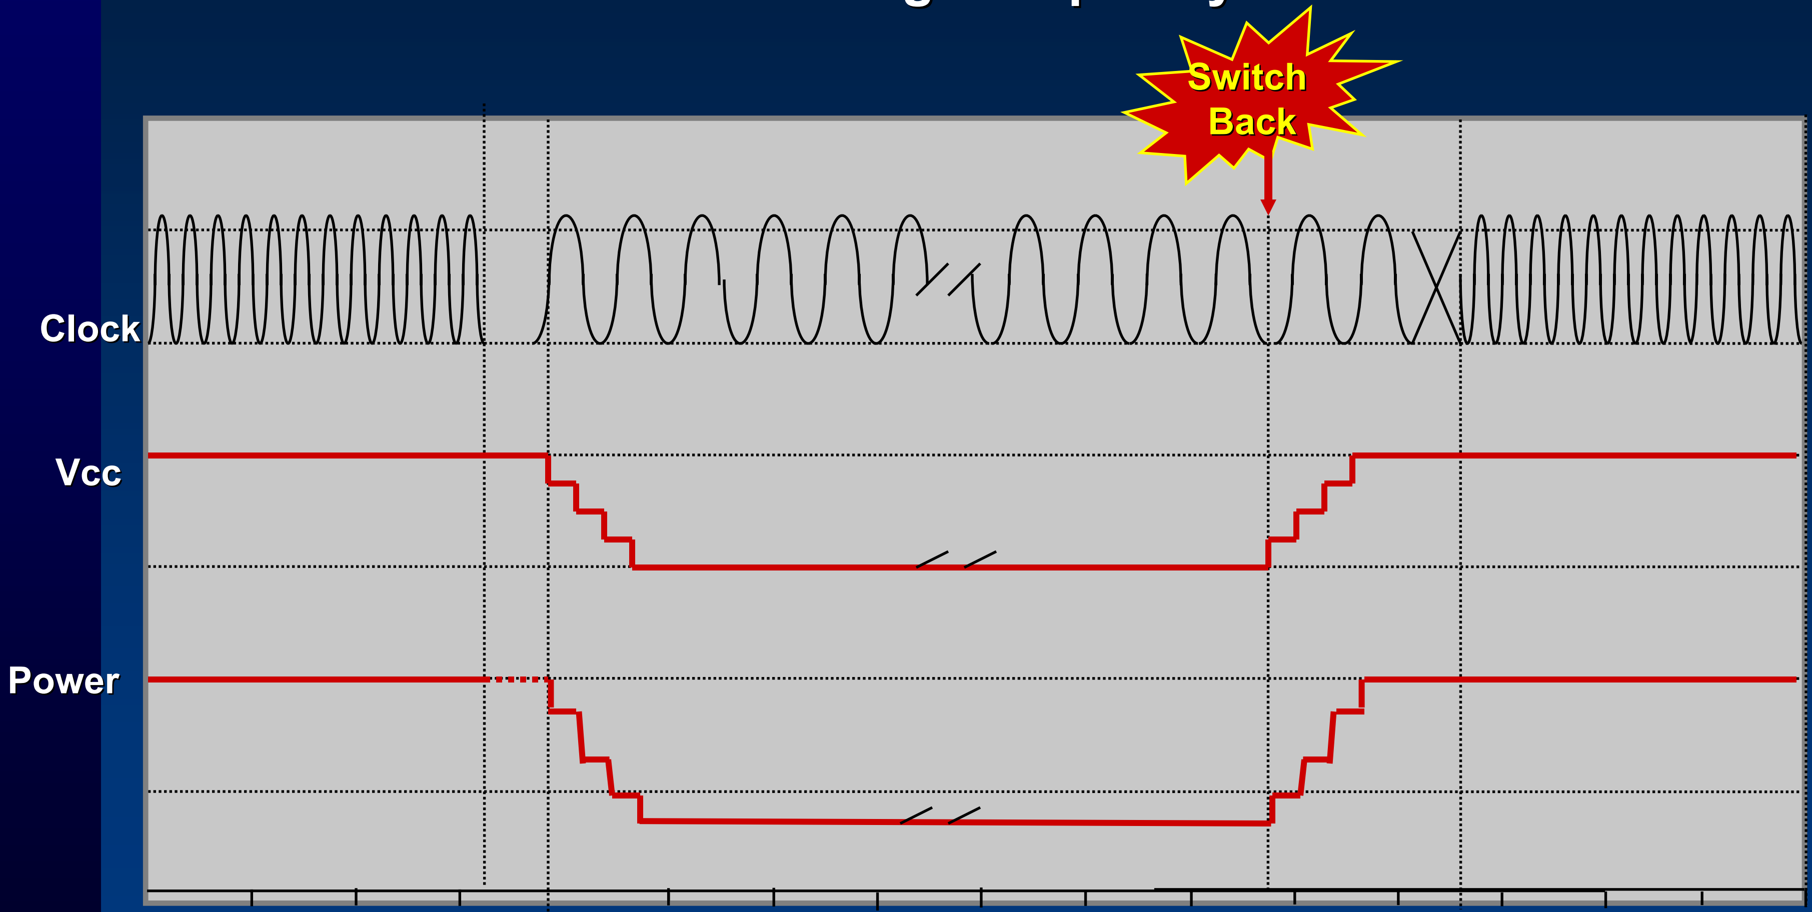

Adaptive Energy Control

- Applications have wide dynamic power range
- Require high power high performance bursts
  - Determine user experience
- Trade power performance as needed
  - Driven by Operating system ACPI
  - Average power control on the fly - **ADAPTIVE**

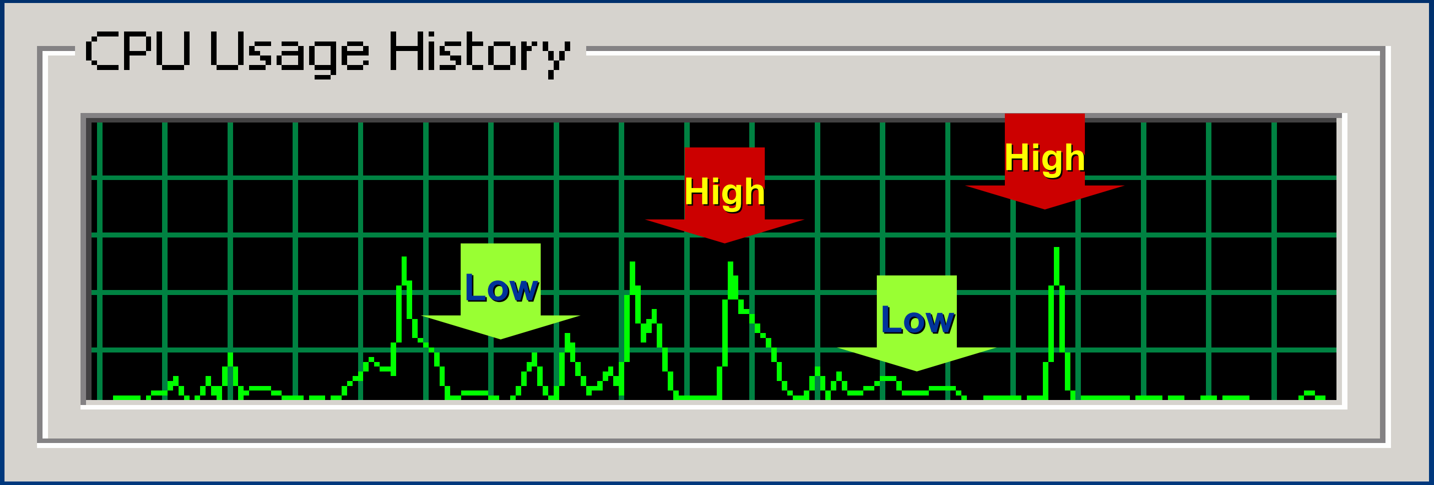

ACPI Control

- Operating system feature – Industry standard
  - Supported by MS and Linux
  - Passive and active policies defined for each zone
  - User preference – Max performance or Max battery
    - Switches _PSV and _ACx
  - DVS or clock throttling used for passive power control
    - Available Power states published by BIOS
    - DVS and once reached min Vcc - linear
    - Implements PD controller algorithm
  - Fan on/off and speed used for active cooling

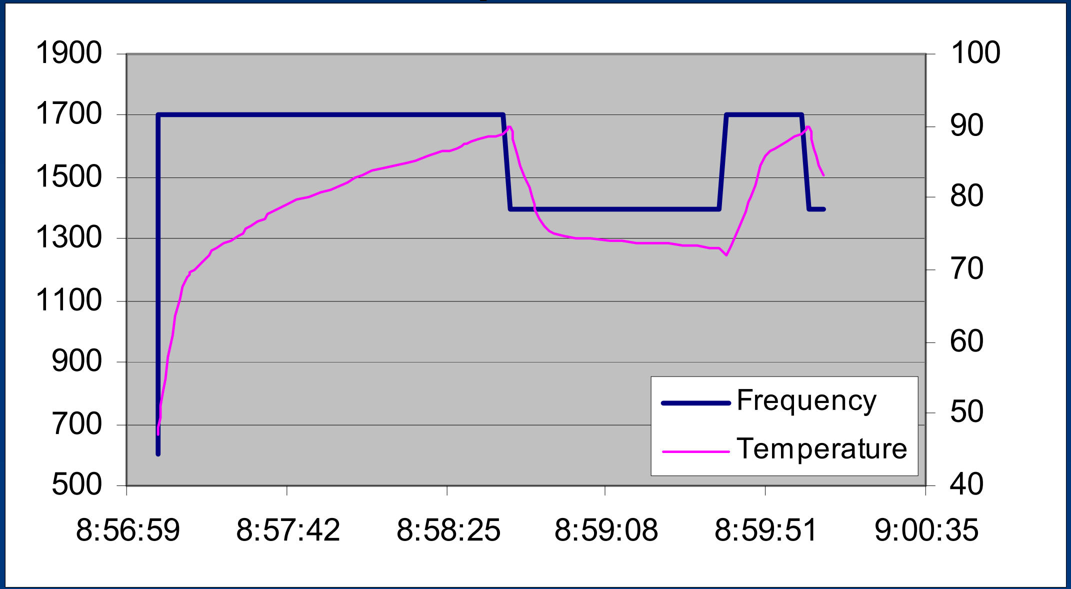

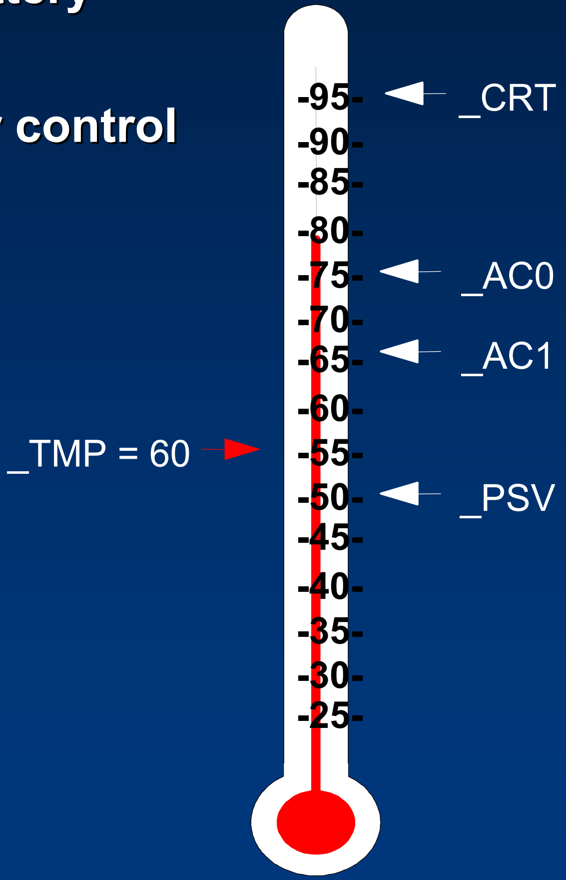

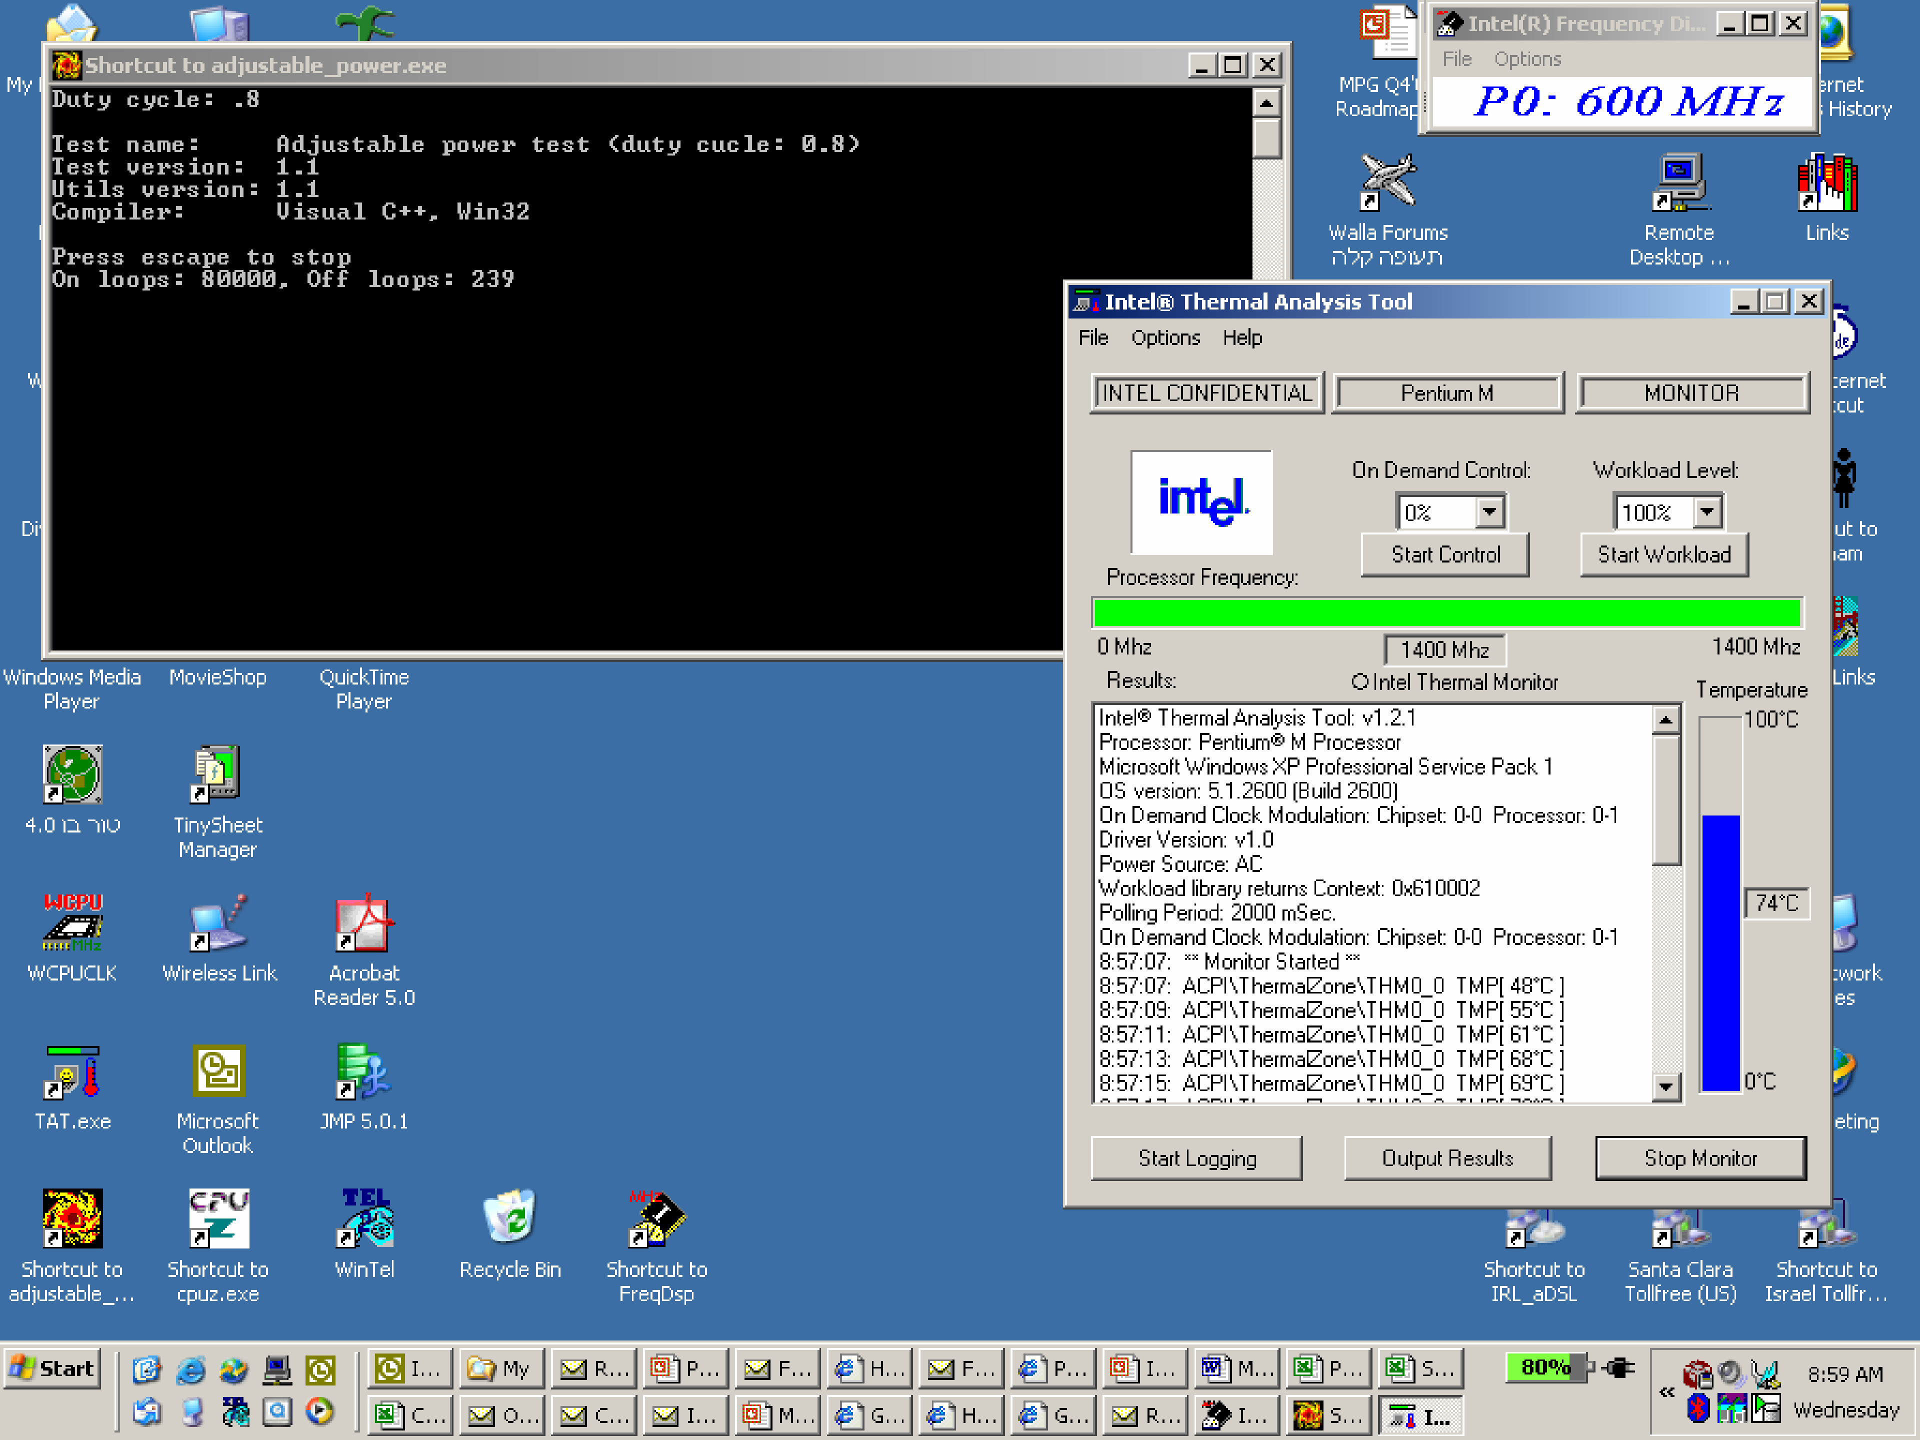
Shortcut to adjustable_power.exe
Duty cycle: .8
Test name: Adjustable power test (duty cucle: 0.8)
Test version: 1.1
Utils version: 1.1
Compiler: Visual C++, Win32
Press escape to stop
On loops: 80000, Off loops: 239
Intel(R) Frequency Di...
File Options
P0: 600 MHz
Intel® Thermal Analysis Tool
File Options Help
INTEL CONFIDENTIAL
Pentium M
MONITOR
On Demand Control:
Workload Level:
0%
100%
Start Control
Start Workload
Processor Frequency:
0 Mhz
1400 Mhz
1400 Mhz
Results:
Intel Thermal Monitor
Temperature
Intel® Thermal Analysis Tool: v1.2.1
Processor: Pentium® M Processor
Microsoft Windows XP Professional Service Pack 1
OS version: 5.1.2600 (Build 2600)
On Demand Clock Modulation: Chipset: 0-0 Processor: 0-1
Driver Version: v1.0
Power Source: AC
Workload library returns Context: 0x610002
Polling Period: 2000 mSec.
On Demand Clock Modulation: Chipset: 0-0 Processor: 0-1
8:57:07: ** Monitor Started **
8:57:07: ACPI\ThermalZone\THM0_0 TMP[48°C]
8:57:09: ACPI\ThermalZone\THM0_0 TMP[55°C]
8:57:11: ACPI\ThermalZone\THM0_0 TMP[61°C]
8:57:13: ACPI\ThermalZone\THM0_0 TMP[68°C]
8:57:15: ACPI\ThermalZone\THM0_0 TMP[69°C]
100°C
74°C
0°C
Start Logging
Output Results
Stop Monitor
Windows Media Player
MovieShop
QuickTime Player
טור בו 4.0
TinySheet Manager
WCPUCLK
Wireless Link
Acrobat Reader 5.0
TAT.exe
Microsoft Outlook
JMP 5.0.1
Shortcut to adjustable_...
Shortcut to cpuz.exe
WinTel
Recycle Bin
Shortcut to FreqDsp
MPG Q4'
Roadmap
Walla Forums
תעופה קלה
Remote Desktop ...
Links
Shortcut to IRL_aDSL
Santa Clara Tollfree (US)
Shortcut to Israel Tollfr...
Start
80%
8:59 AM
Wednesday

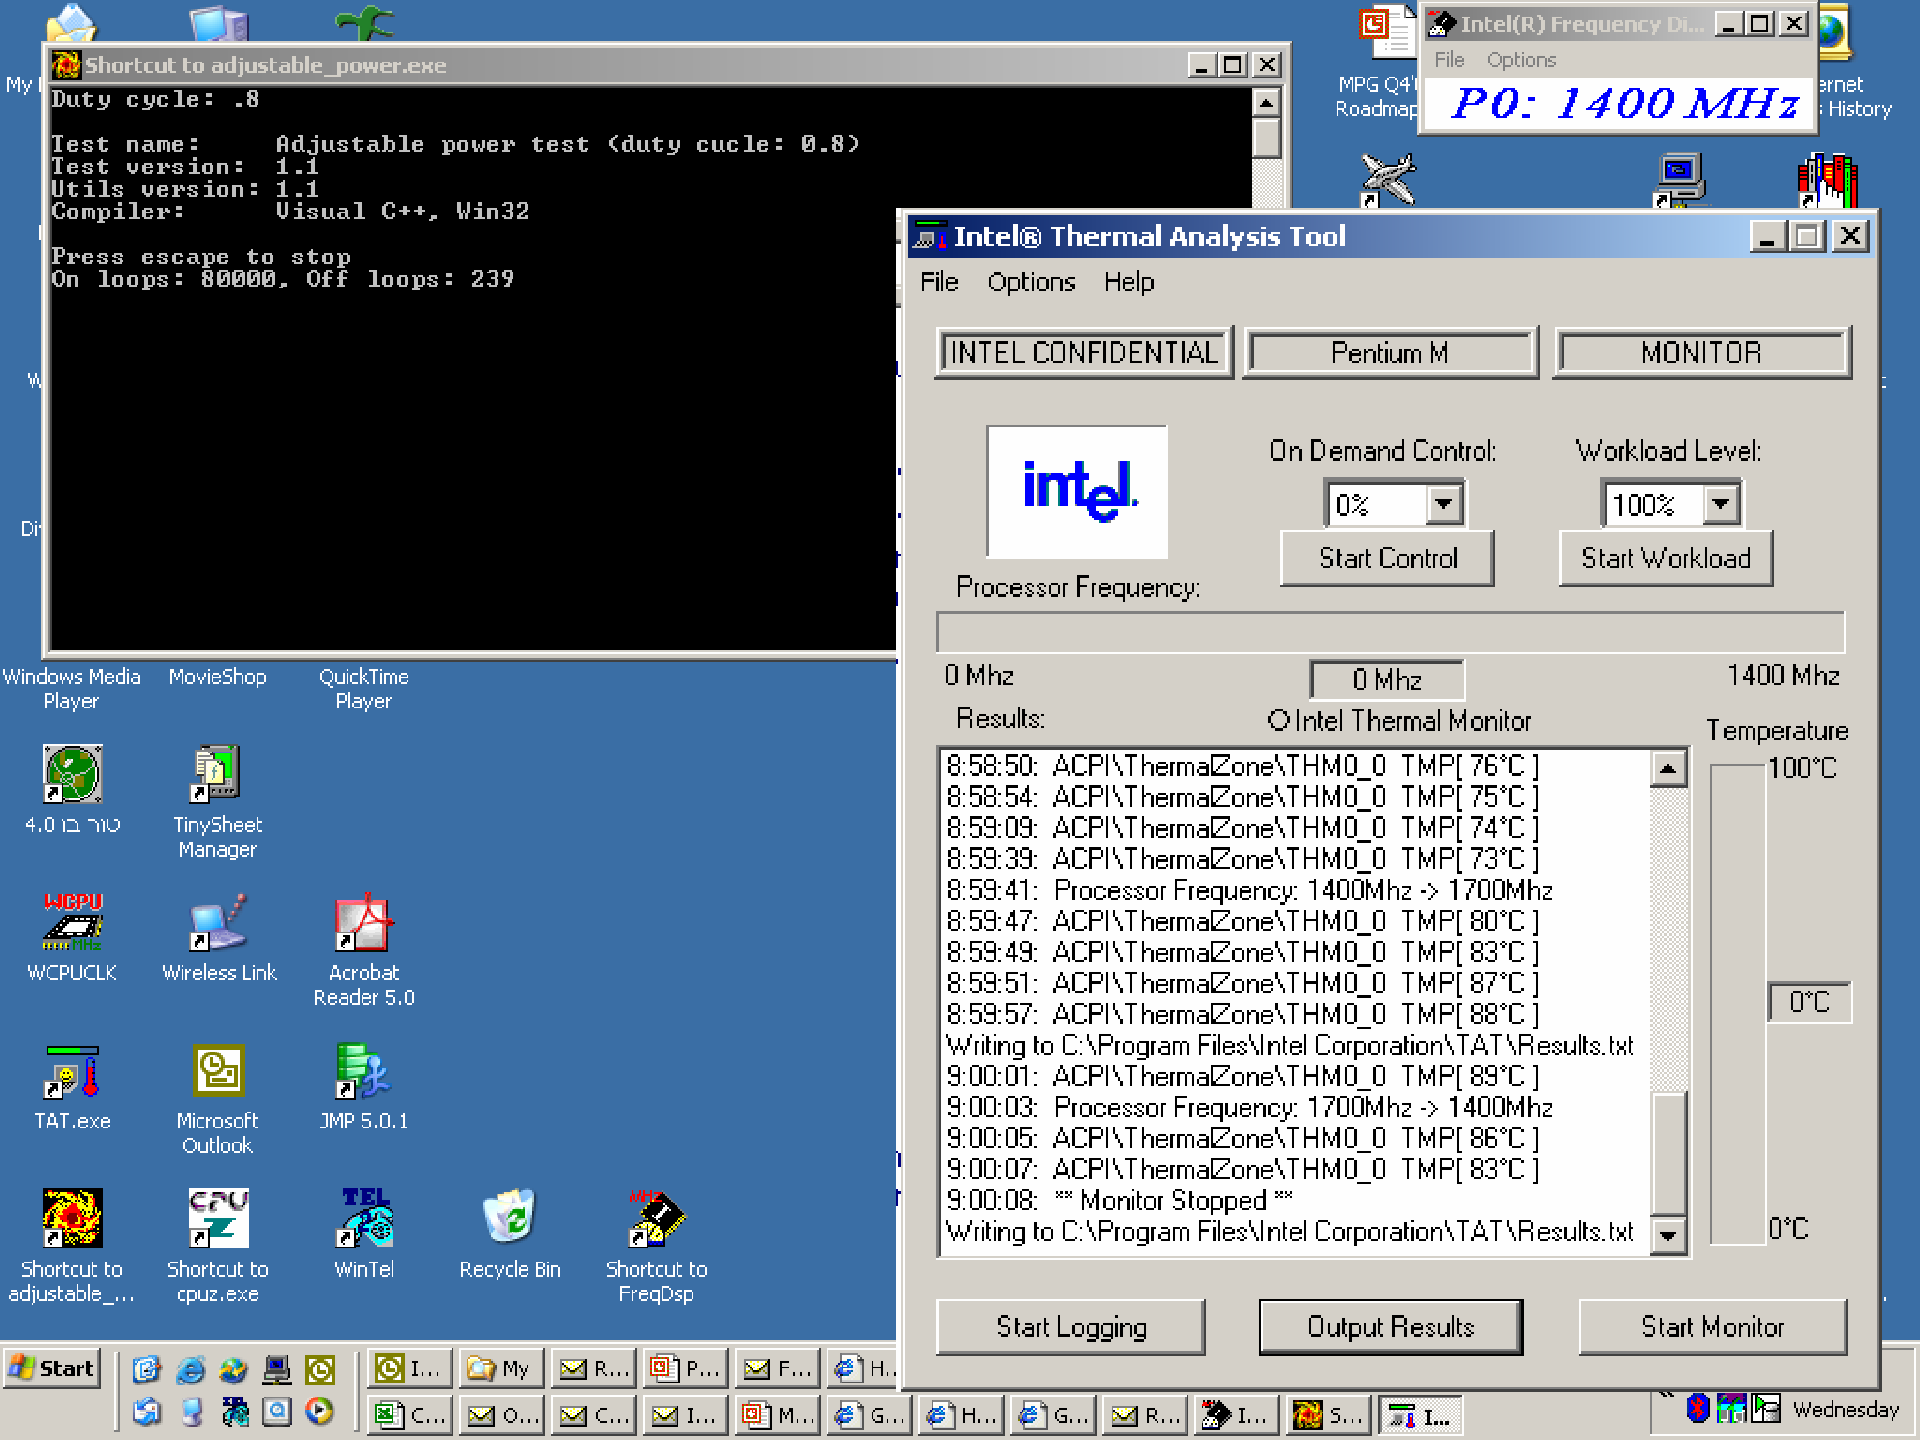
Shortcut to adjustable_power.exe
Duty cycle: .8
Test name: Adjustable power test (duty cucle: 0.8)
Test version: 1.1
Utils version: 1.1
Compiler: Visual C++, Win32
Press escape to stop
On loops: 80000, Off loops: 239
Intel(R) Frequency Di...
File Options
P0: 1400 MHz
Intel® Thermal Analysis Tool
File Options Help
INTEL CONFIDENTIAL
Pentium M
MONITOR
On Demand Control:
Workload Level:
0%
100%
Start Control
Start Workload
Processor Frequency:
0 Mhz
0 Mhz
1400 Mhz
Results:
Intel Thermal Monitor
Temperature
8:58:50: ACPI\ThermalZone\THM0_0 TMP[76°C]
8:58:54: ACPI\ThermalZone\THM0_0 TMP[75°C]
8:59:09: ACPI\ThermalZone\THM0_0 TMP[74°C]
8:59:39: ACPI\ThermalZone\THM0_0 TMP[73°C]
8:59:41: Processor Frequency: 1400Mhz -> 1700Mhz
8:59:47: ACPI\ThermalZone\THM0_0 TMP[80°C]
8:59:49: ACPI\ThermalZone\THM0_0 TMP[83°C]
8:59:51: ACPI\ThermalZone\THM0_0 TMP[87°C]
8:59:57: ACPI\ThermalZone\THM0_0 TMP[88°C]
Writing to C:\Program Files\Intel Corporation\TAT\Results.txt
9:00:01: ACPI\ThermalZone\THM0_0 TMP[89°C]
9:00:03: Processor Frequency: 1700Mhz -> 1400Mhz
9:00:05: ACPI\ThermalZone\THM0_0 TMP[86°C]
9:00:07: ACPI\ThermalZone\THM0_0 TMP[83°C]
9:00:08: ** Monitor Stopped **
Writing to C:\Program Files\Intel Corporation\TAT\Results.txt
100°C
0°C
0°C
Start Logging
Output Results
Start Monitor
Windows Media Player
MovieShop
QuickTime Player
TinySheet Manager
WCPUCLK
Wireless Link
Acrobat Reader 5.0
TAT.exe
Microsoft Outlook
JMP 5.0.1
Shortcut to adjustable_...
Shortcut to cpuz.exe
WinTel
Recycle Bin
Shortcut to FreqDsp
Start
Wednesday

Overview

- Mobile computers challenges
  - Average power
  - Thermal Design Power
- Intel Pentium M power management
- The experiment
- Test results
- Conclusions

Test Setup

Thermal control

- CPU temperature is a function of power and ambient temp.
  - Long time to heat the ambient and heat sink
- Case temperature control formed stable environment
  - Heating plate and fan to keep temperature at fixed temperature
  - Force extreme conditions

Test cases

- SPEC-Int and SPEC-FP components
- Used the self trigger mechanism
  - For repeatable & consistent results
- Collected power temperature and performance

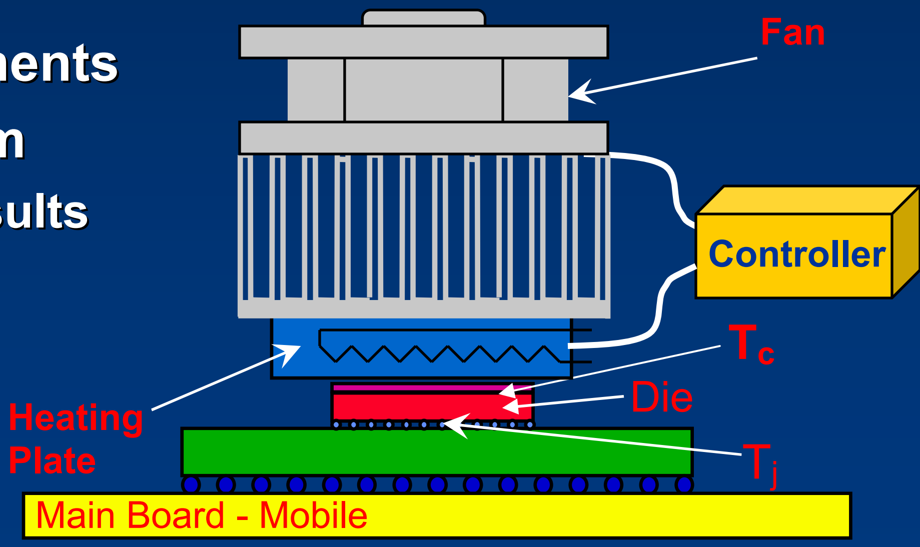

Testing on real silicon the theoretical expected behavior

Overview

- Mobile computers challenges
  - Average power
  - Thermal Design Power
- Intel Pentium M power management
- The experiment
- Test results
- Conclusions

Lineal vs. DVS control

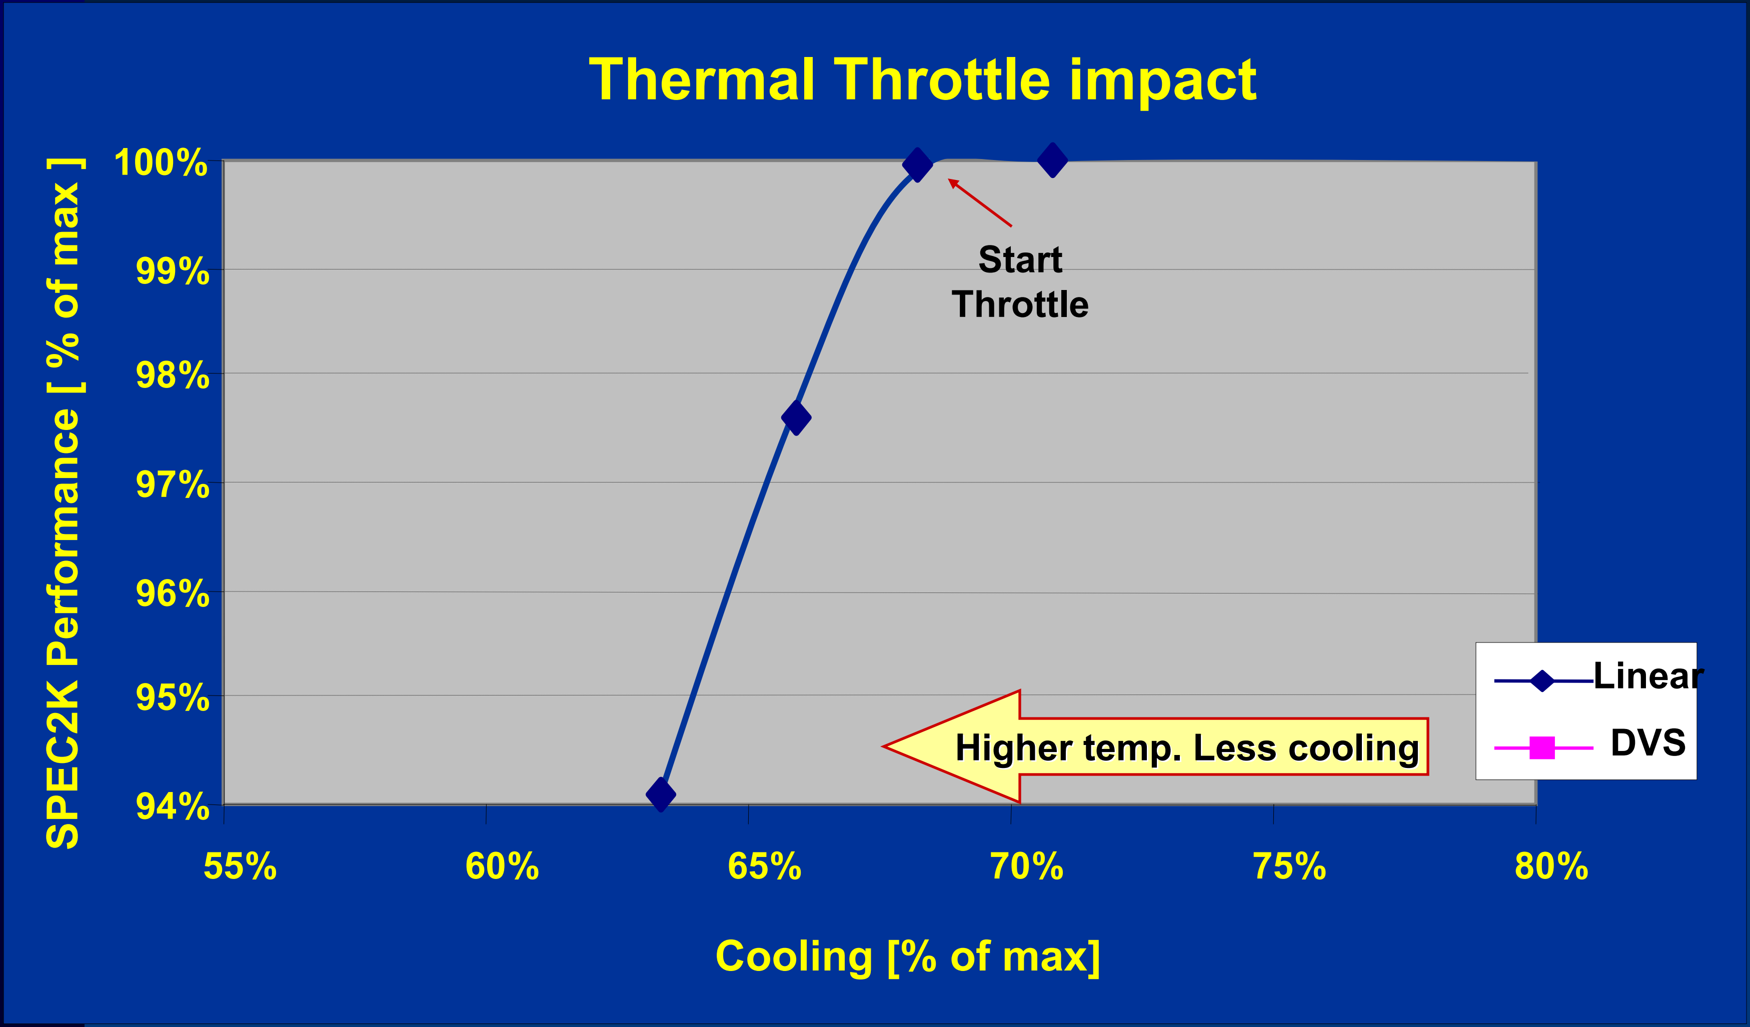

Lineal vs. DVS control

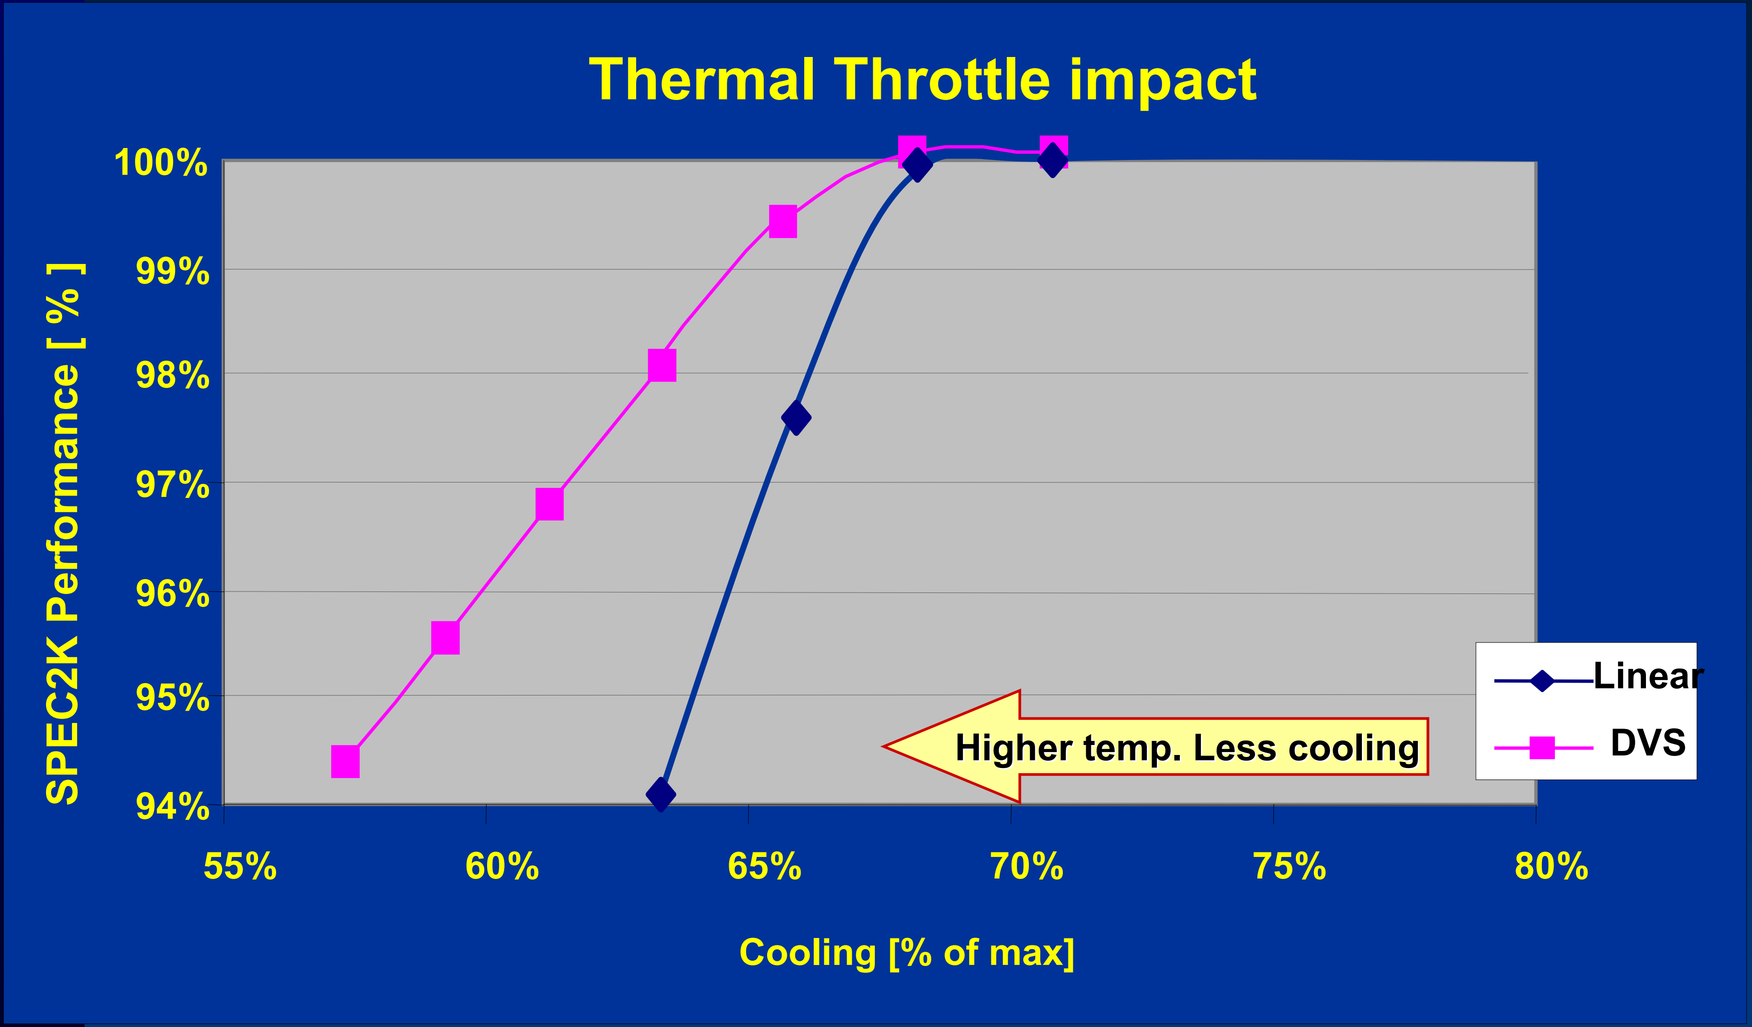

Lineal vs. DVS control

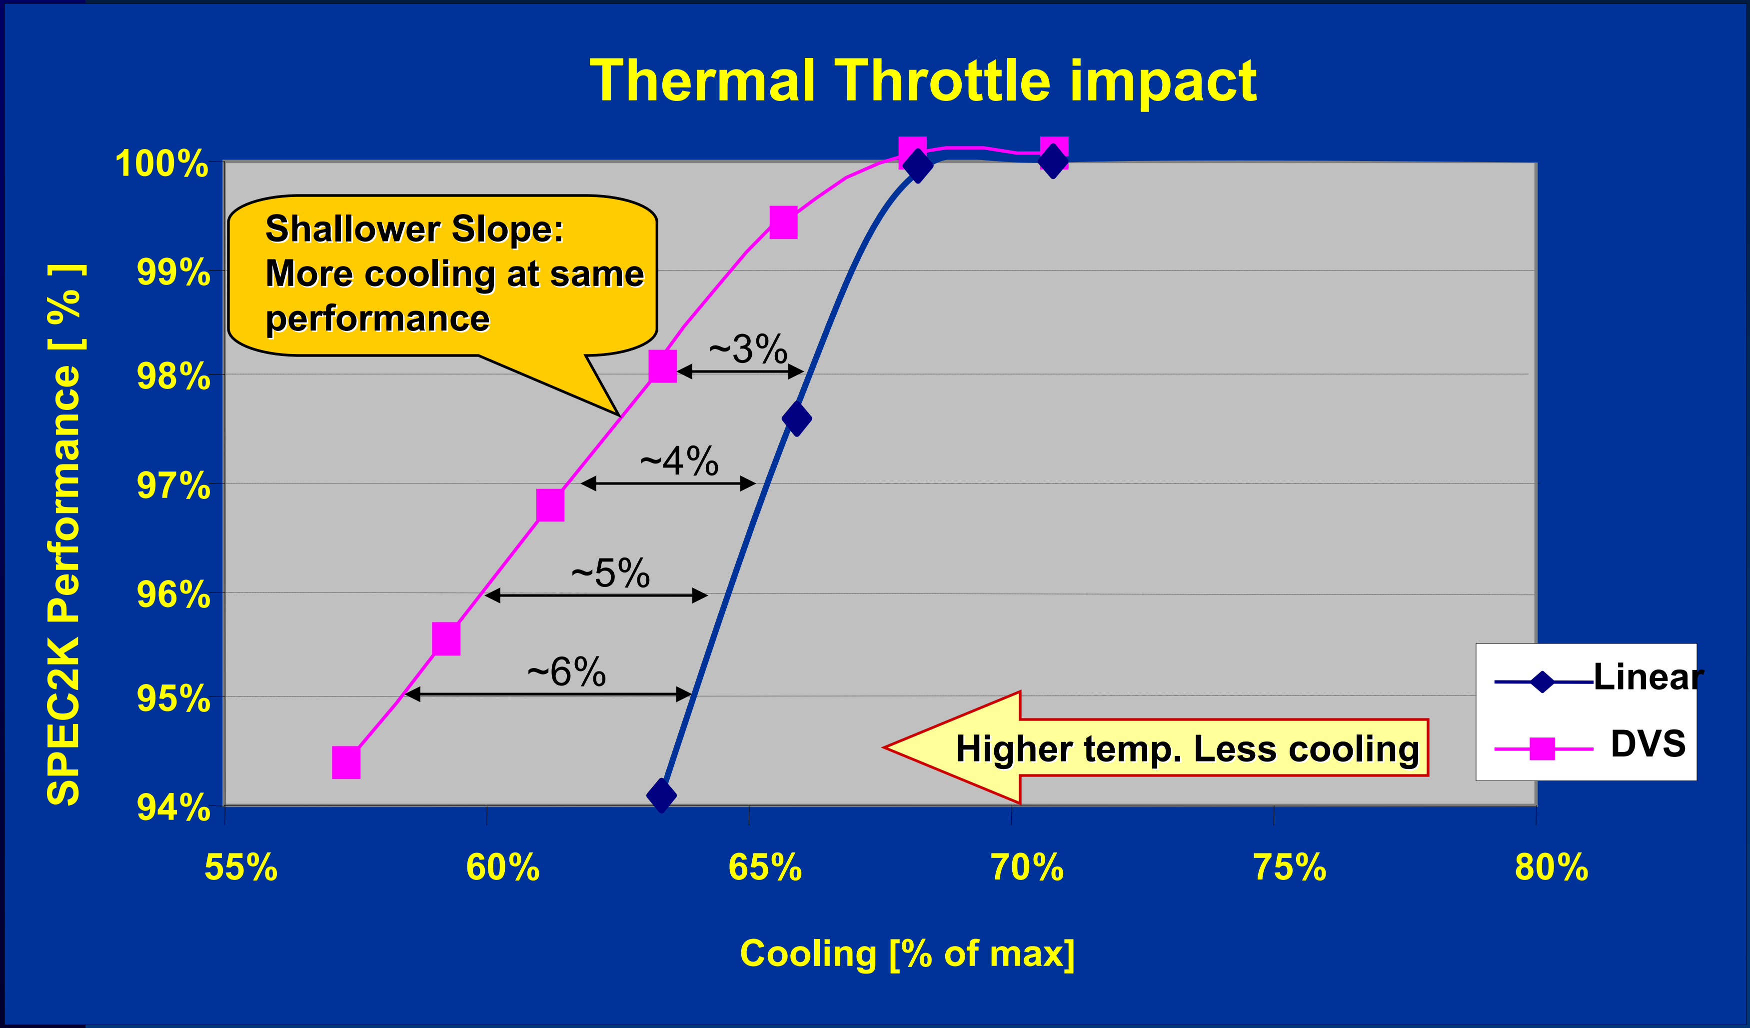

DVS set point

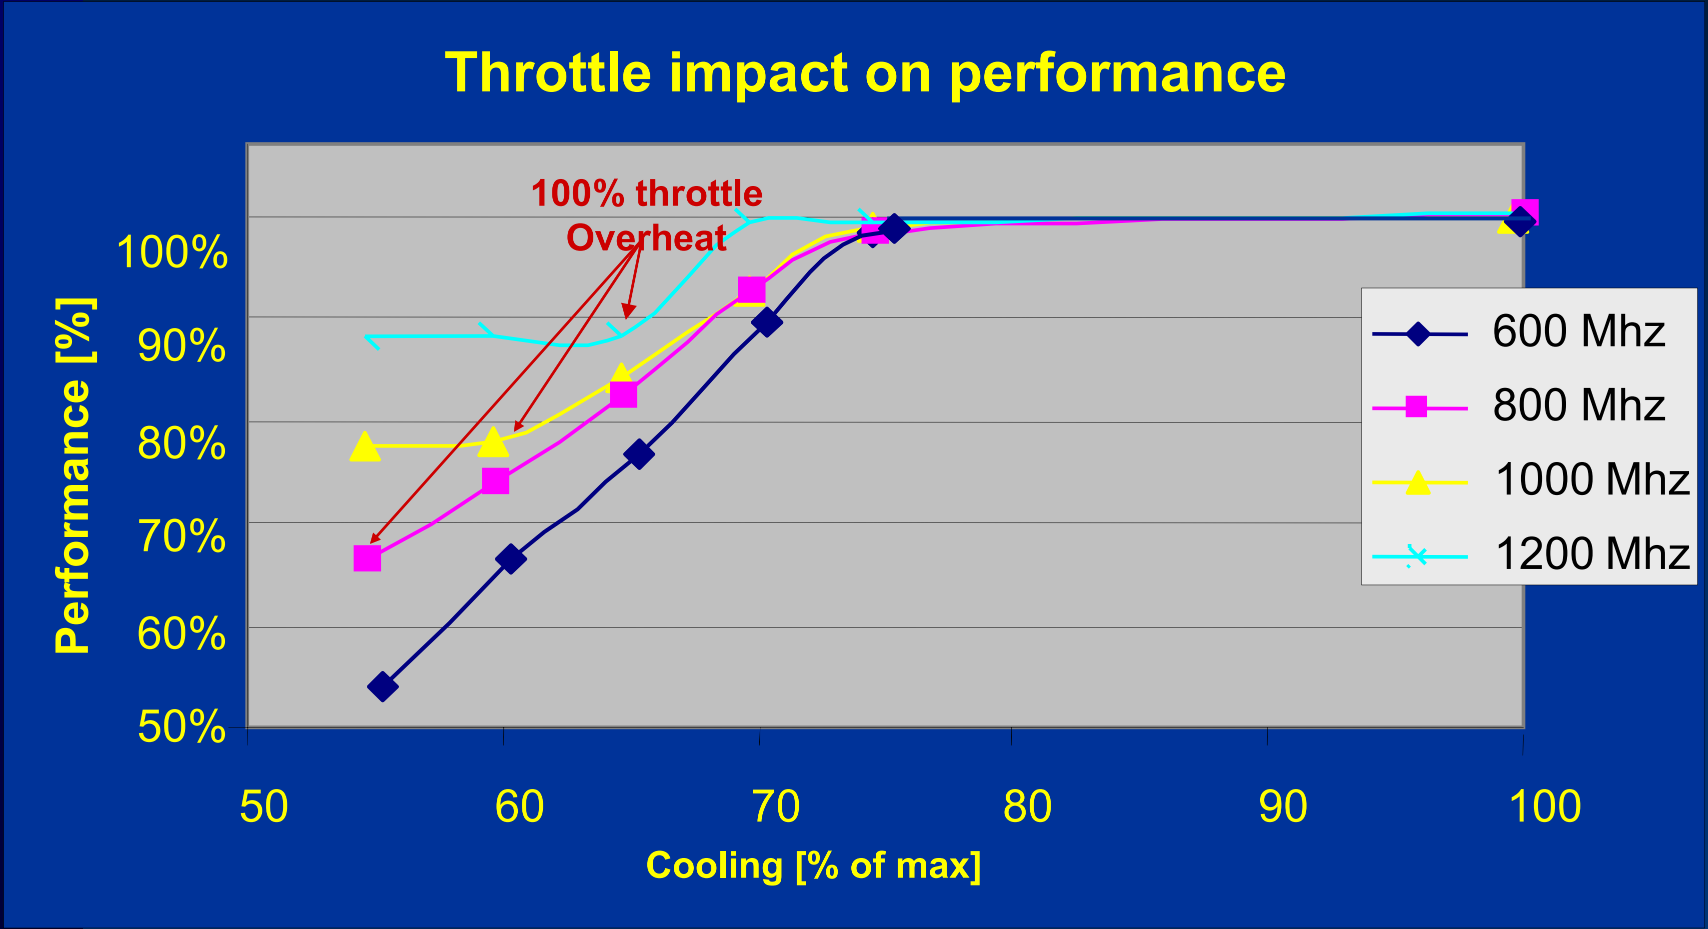

Optimal point (performance) is at the highest frequency that dose not accede Tj_max with no transitions up and down

Average Power management

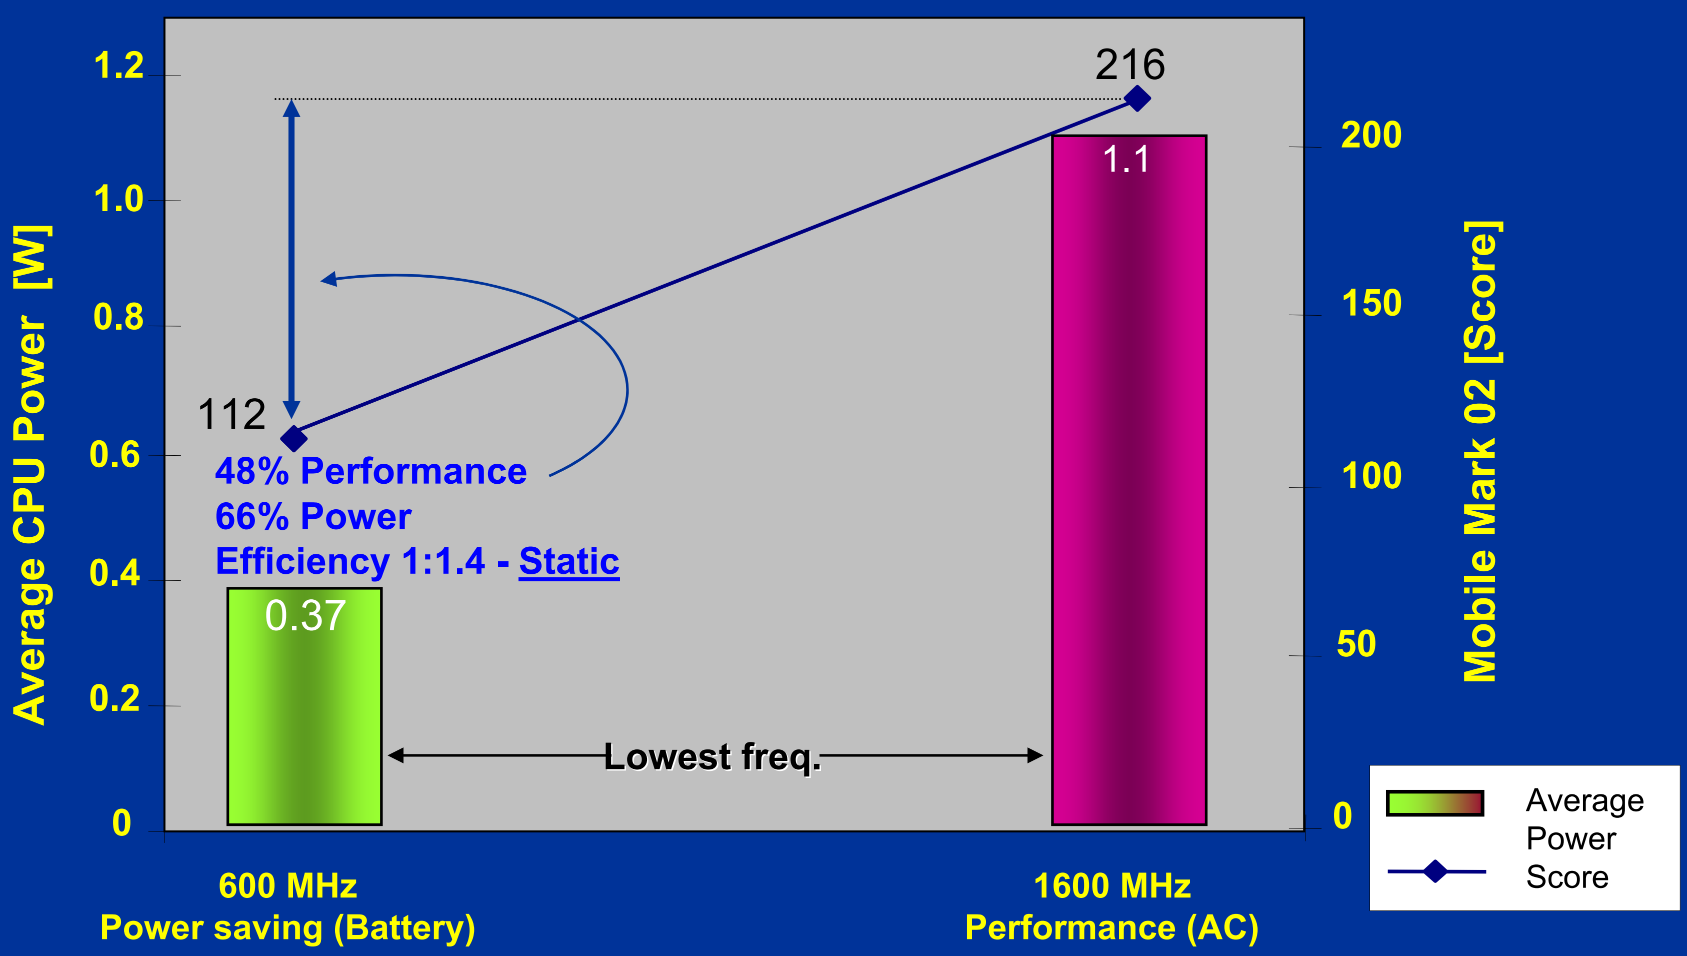

Average Power management

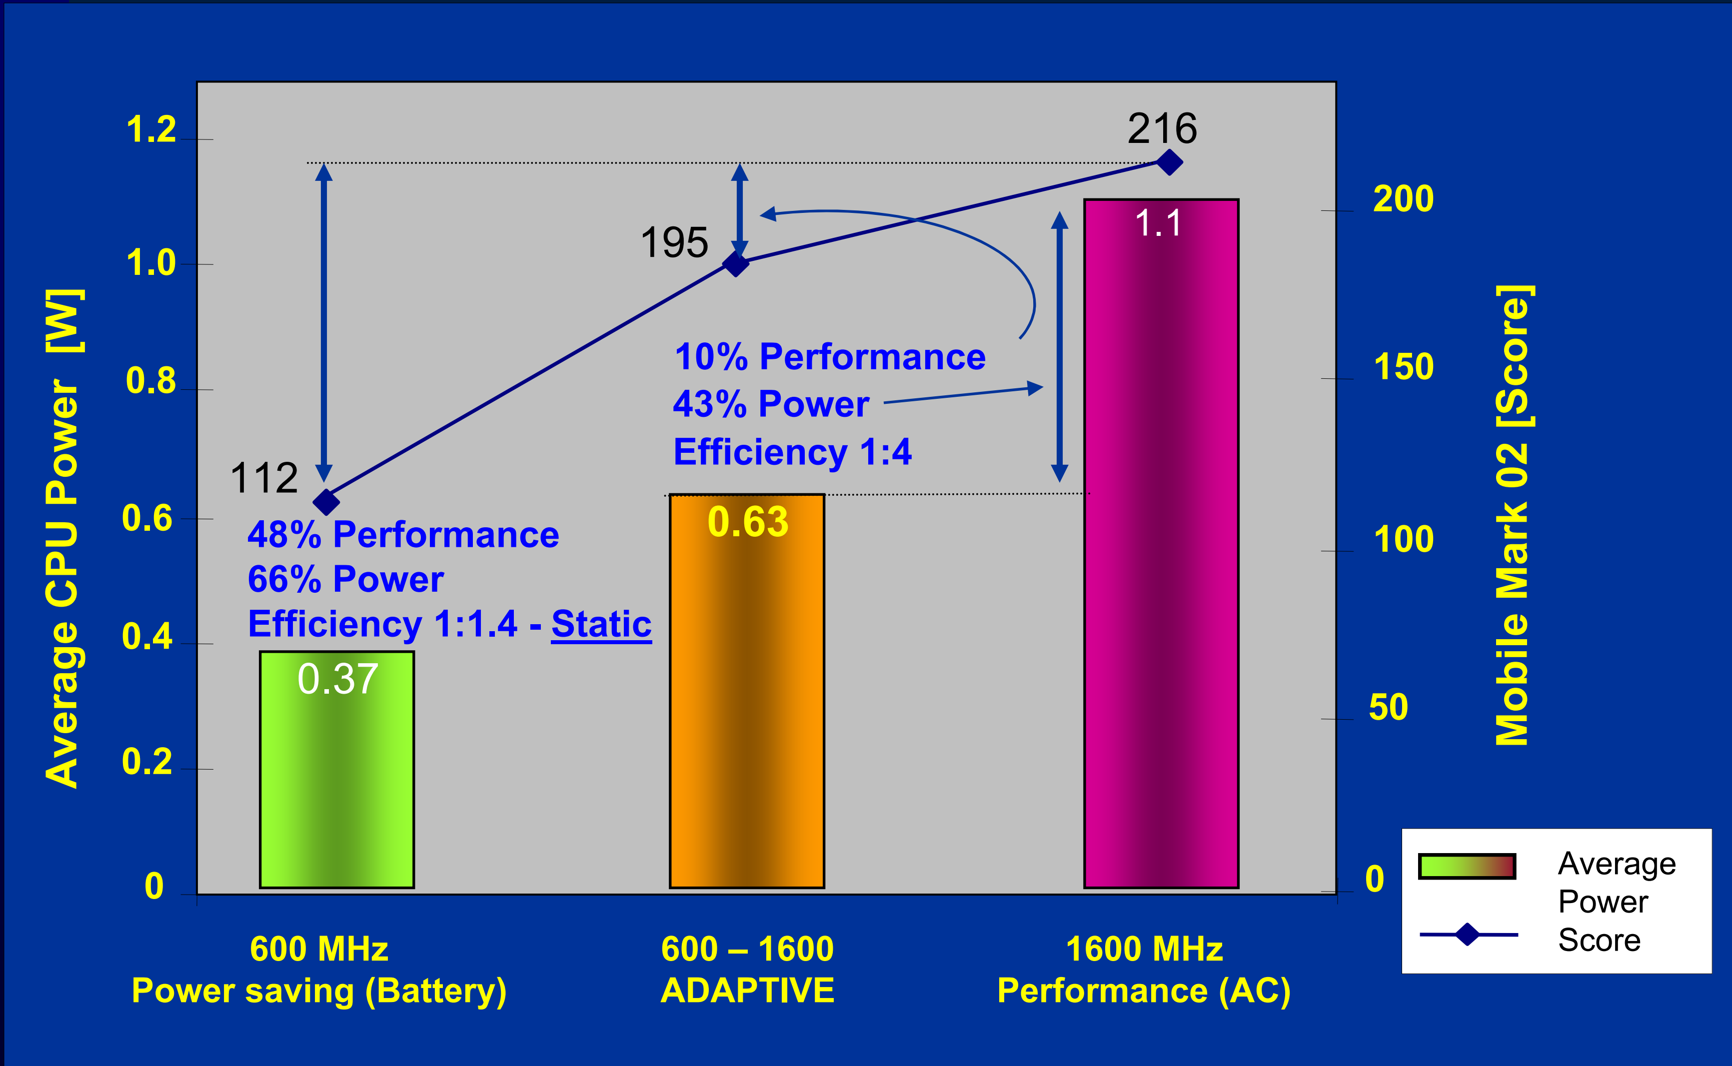

Energy efficiency

- Energy consumed for the SPEC-Int and SPEC-FP
  - We measured energy as $E=\int_t P(t)$
  - Energy consumed at 1600 MHz defined as 100%

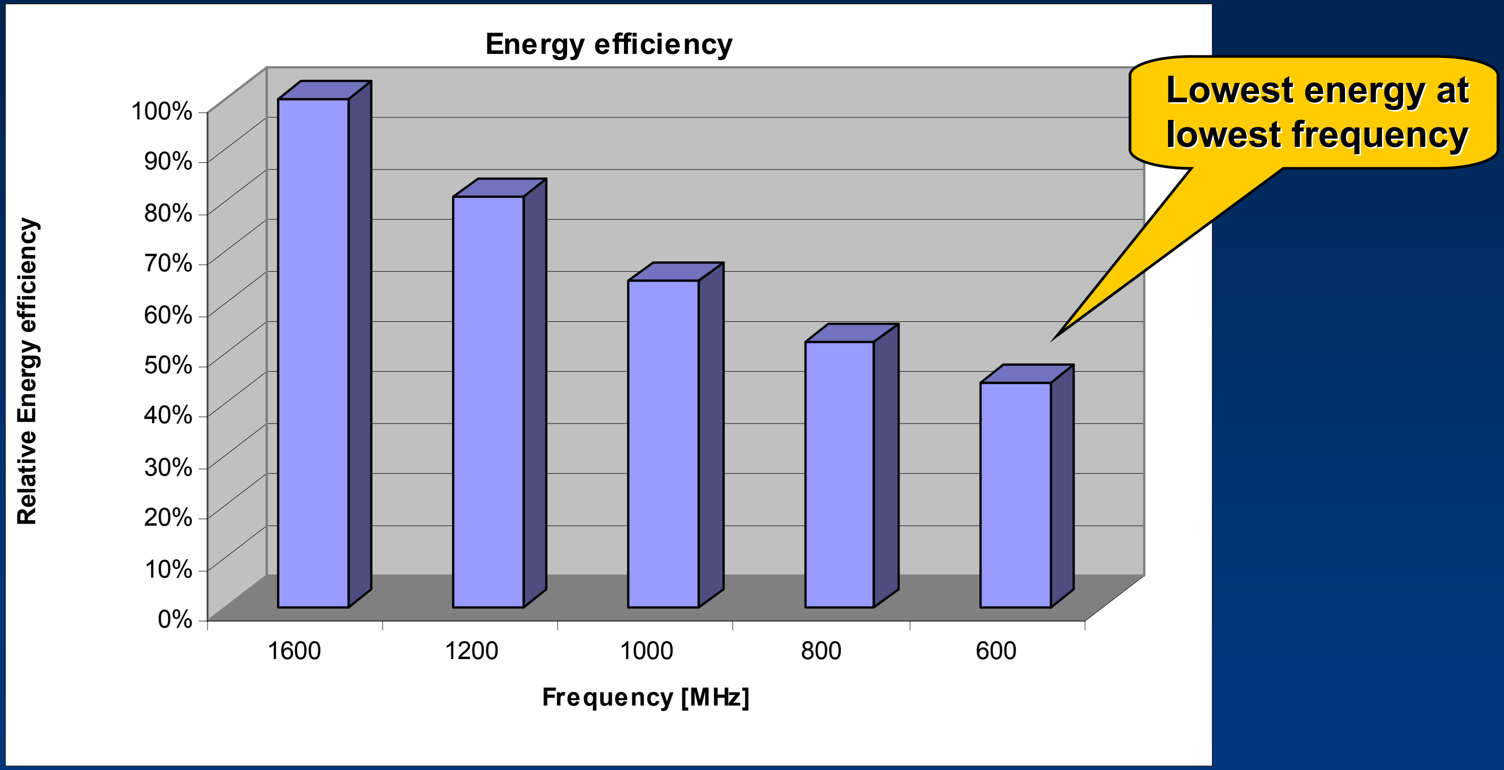

Overview

- Mobile computers challenges
  - Average power
  - Thermal Design Power
- Intel Pentium M power management
- The experiment
- Test results
- Conclusions

Summary and Conclusions

- Intel Pentium M Specifically designed for mobile
  - Targeting energy and power efficiency
- Mobile systems trade performance for power
  - To meet user preference
- Enhanced Intel SpeedStep Technology provides significantly improved power to performance and energy control scheme
  - Silicon measurement confirm the theoretical work
  - Optimal point for performance to power is at the highest frequency that dose not accede Tj_max
    - Policy implemented into ACPI algorithm
  - Optimal point for energy and battery life is at the minimum frequency possible with DVS

Thanks

The authors would like to thank Nachum shamir, Irina Ilatov and Riad Durr for the thermal clamping measurements, to Jessie Garcia, Ali Saeed, Marco Wirasinghe for the evaluation and data collection for this article and to Cohen Aviad and Lev Finkelstein for results verification using simulation.

Die Temperature measurements

- A diode connected to an external A/D reports temperature
- Fixed temperature

$$V_{BE} = \frac{nkT}{q} \cdot \ln\left(\frac{I_C}{I_S}\right)$$

Appling 2 Currents

$$T = \frac{\Delta V_{BE}}{n \cdot \ln(N) \cdot \frac{k}{q}}$$

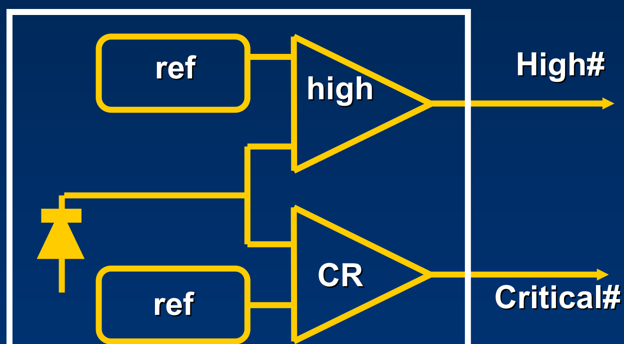

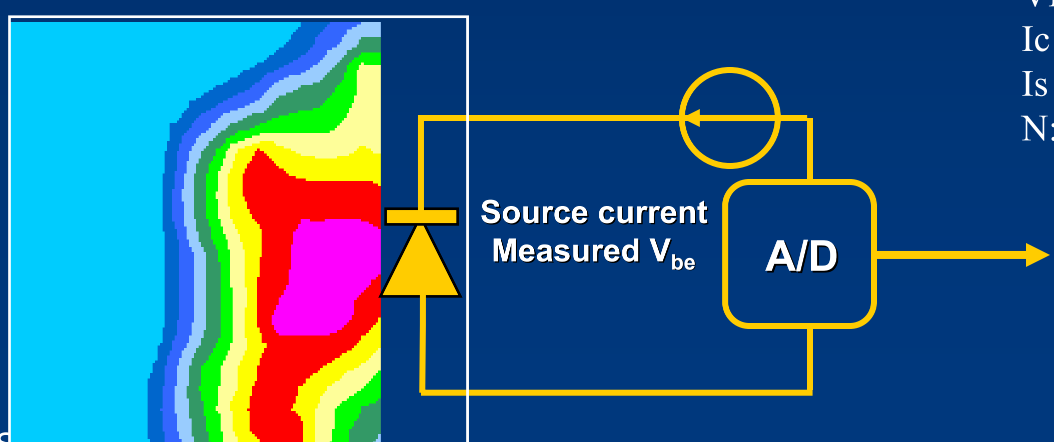

n: diode ideality factor
k: Boltzmann constant
q: electron charge constant
T: diode temperature (Kelvin)
VBE: Voltage from base to emitter
Ic : Collector current
Is : Saturation current
N: Ratio of collector currents

Controlling TDP - Linear

- STOPCLOCK continues toggling for a pre defined time
- Average power and temperature drop – performance drops

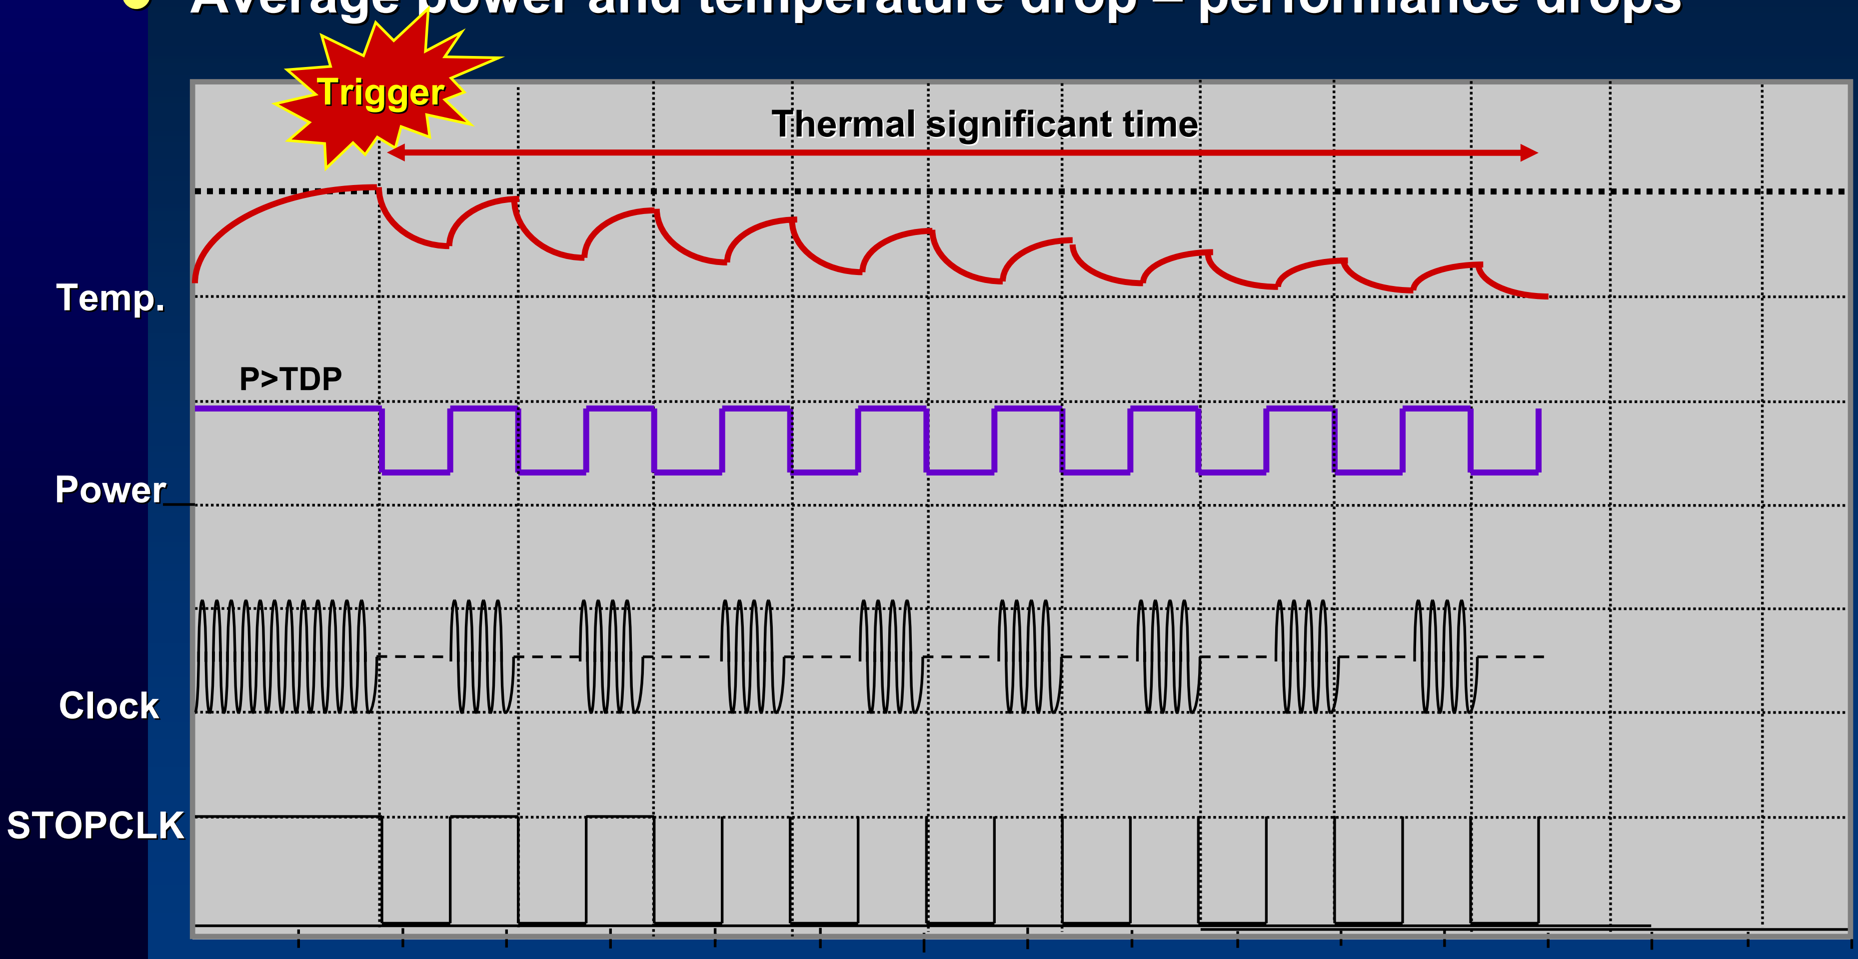

Dynamic Voltage Scaling (DVS)

- Vcc drops gradually while CPU active
- Power savings changes from linear to F^3

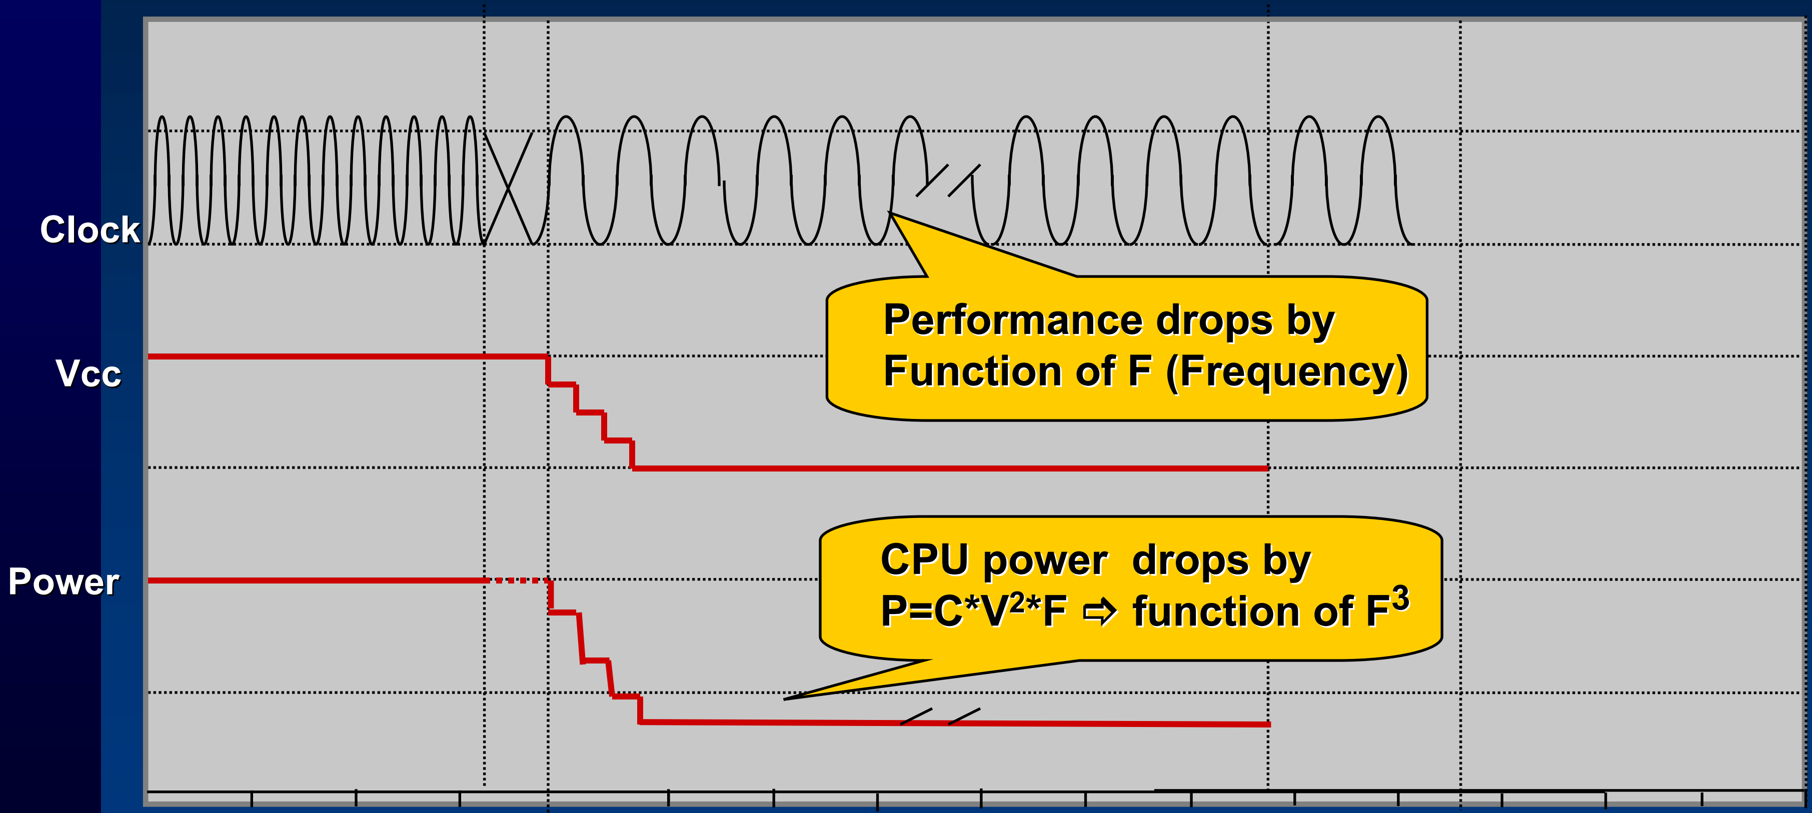

The Pentium M Power Control Schemes

Linear power scaling

- Change CPU frequency only
  - reduces both power and performance linearly with frequency
  - Saves power, No energy savings

New Dynamic Voltage Scaling (DVS)

- Reduces Voltage and frequency on the fly
  - Frequency is dependent ~ linearly on Voltage
  - Power is a function of $C*F*V^2$
- Reduction of both voltage and frequency provides F^3 power reduction for linear performance reduction
  - Also results with energy savings

ACPI Interface

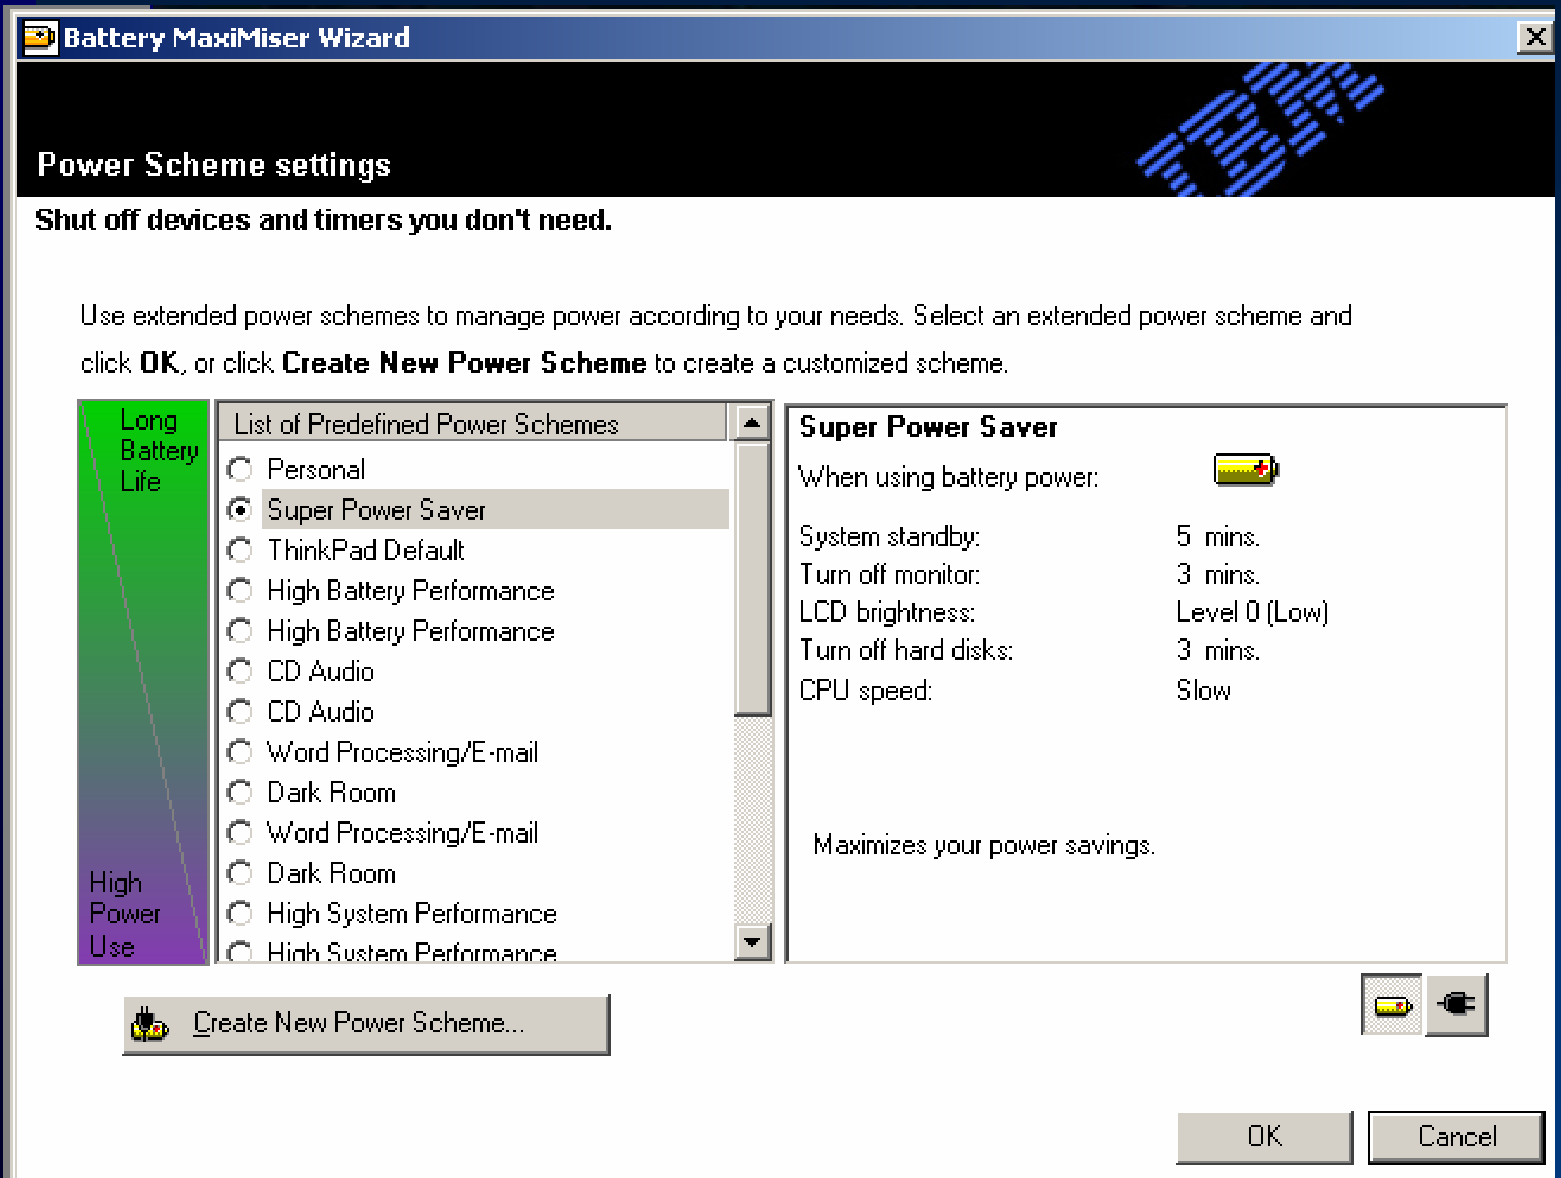

The Pentium M Power Control Schemes

Linear power scaling

- Change CPU frequency only
  - reduces both power and performance linearly with frequency

New Dynamic Voltage Scaling (DVS)

- Reduces Voltage and frequency on the fly
- Reduction of both voltage and frequency provides F^3 power reduction for linear performance reduction

- See details…