mentor embedded

Power Management
A System Perspective

Stephen Olsen

Embedded Systems Division

April 2010

EDP Symposium

Mentor Graphics

Android · Nucleus · Linux

Mobile & Beyond · 2D/3D User Interfaces · Multi-OS · Networking

Whose fault is it?

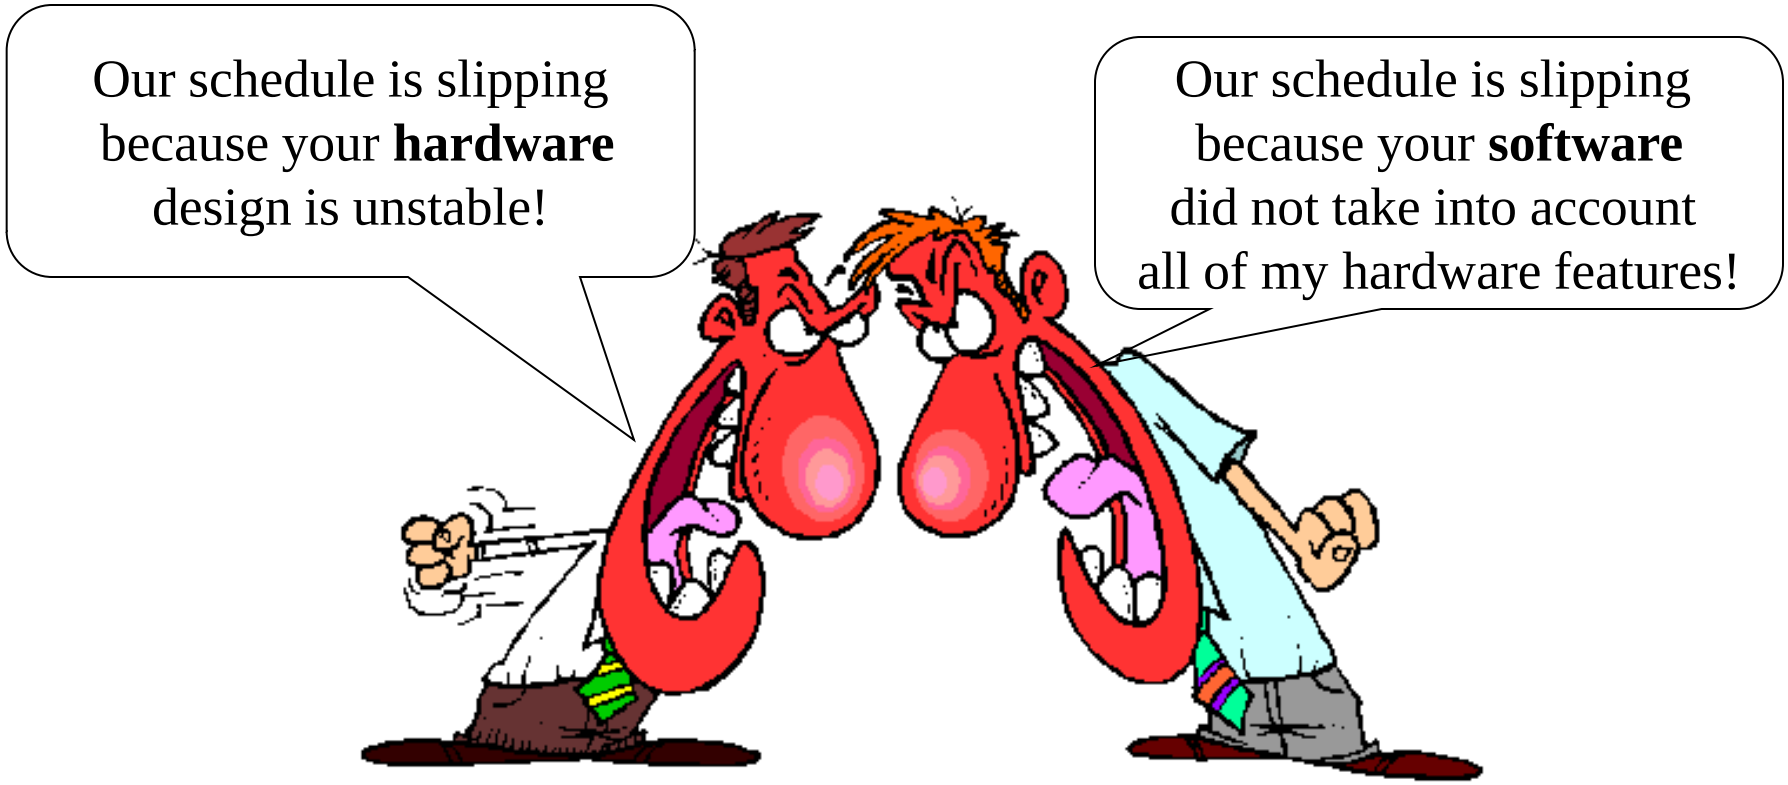

- How do we going to maximize power savings if this is how we work together?

AGENDA

- Why is power management important?
- Introduction to power management
  - A System level issue
  - The business context
  - Power Optimization at every level
- Nucleus Power Manager
  - Power Model
  - Peripheral Power Management
  - Core Power Management
- Next Steps and External loop

Technology Trends – Cell Phones

The first cell phones made phone calls.

Modern 3G Cell phones also make phone calls, and then some… Text, Camera, Email, Webrowser, Email client, GPS, games …

Technology Trends – Batteries

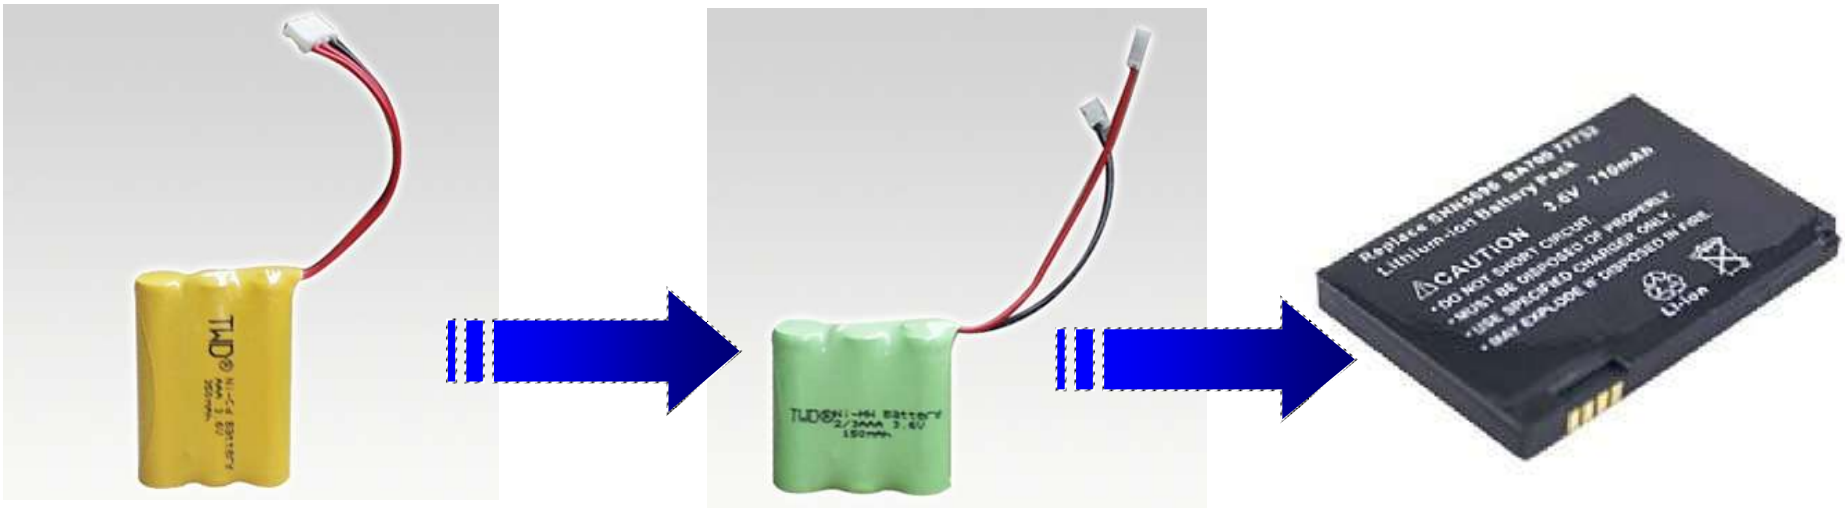

NiCd

- 40-50 wh/kg
- 1.2v
- 500-800 charges

Ni-mh

- 70-80 wh/kg
- 1.2v
- 500-1000 charges

Li-ion

- 130-175 wh/kg
- 3.6v
- 500 charges

Specific Energy: the ratio of energy delivered by the battery to the weight of the battery measured in watt hours per kilogram (wh/kg)

Battery technology is not keeping up with the demand for more functionality

Power Management

For Portable Consumer Electronic Devices

- **Battery Life**
  - **Maximize Battery Life**
  - **Minimize Size, Weight and Capacity**

What is the problem?

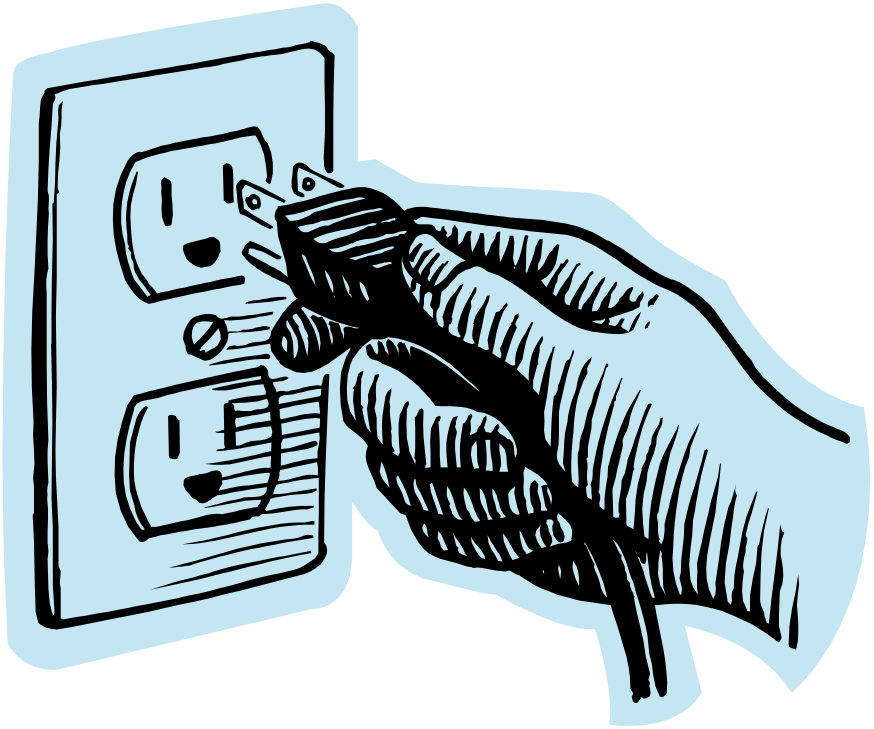

Power Generation – Worldwide

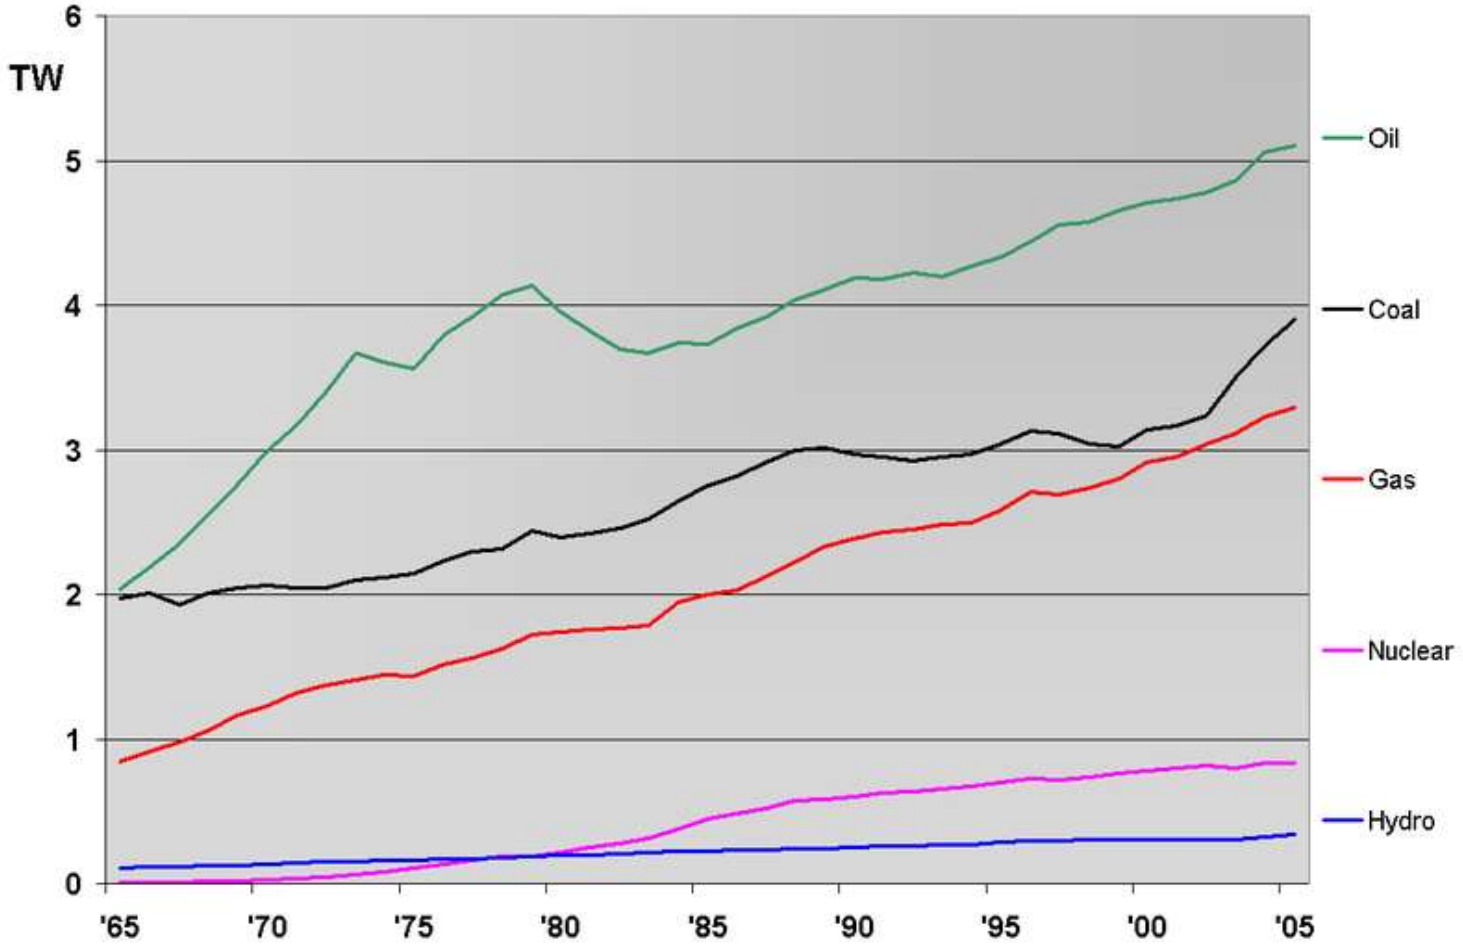

Source: Energy Information Administration, U.S. Department of Energy. July 31, 2006

Power Generation – Worldwide

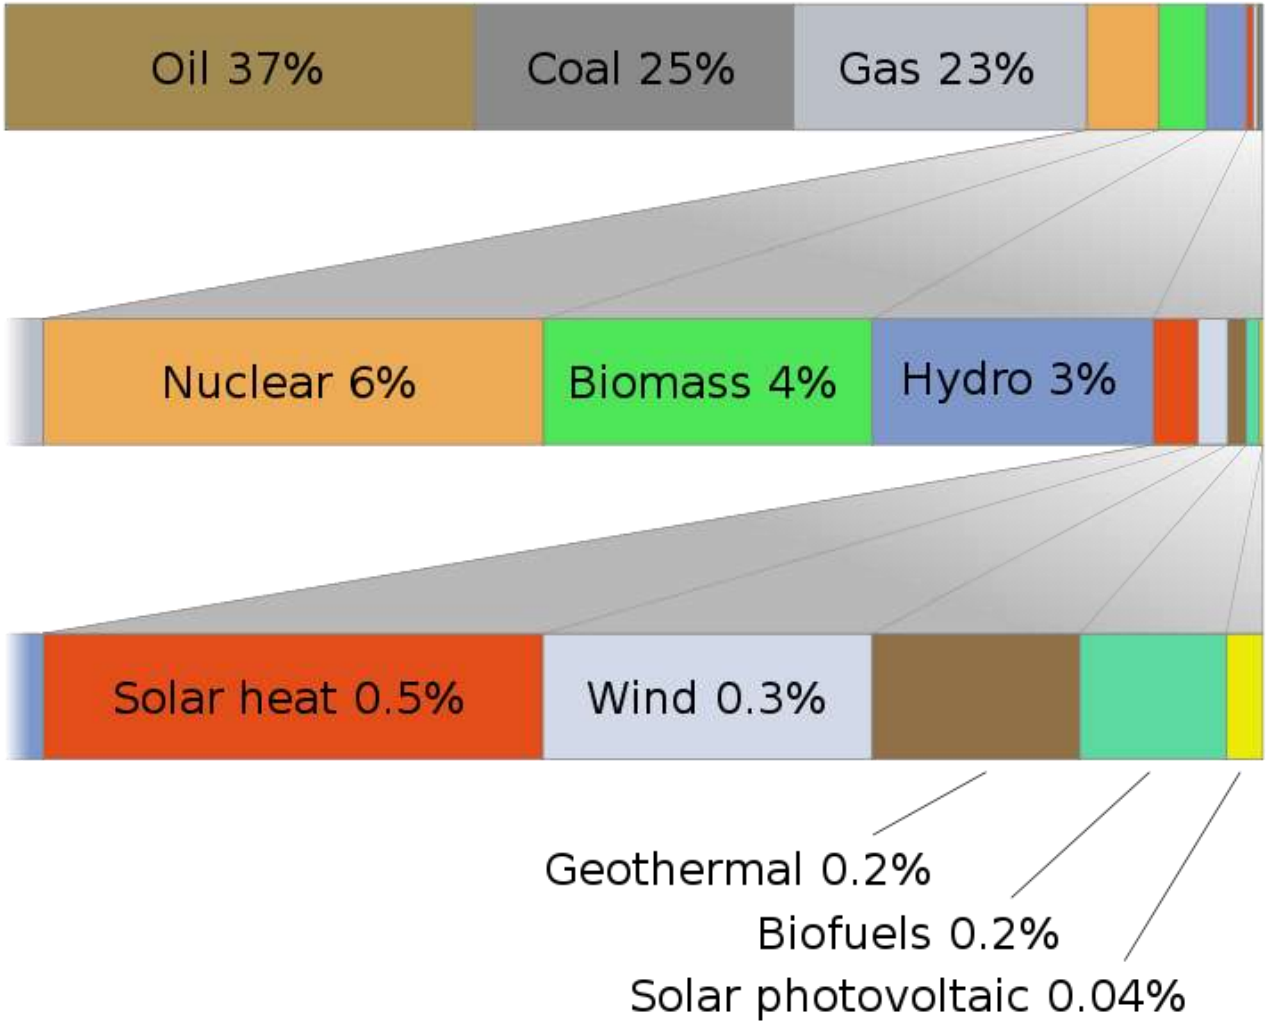

Source: Renewables, Global Status Report, 2006 update; wiki - Omegatron

Temperatures over past 2000 years

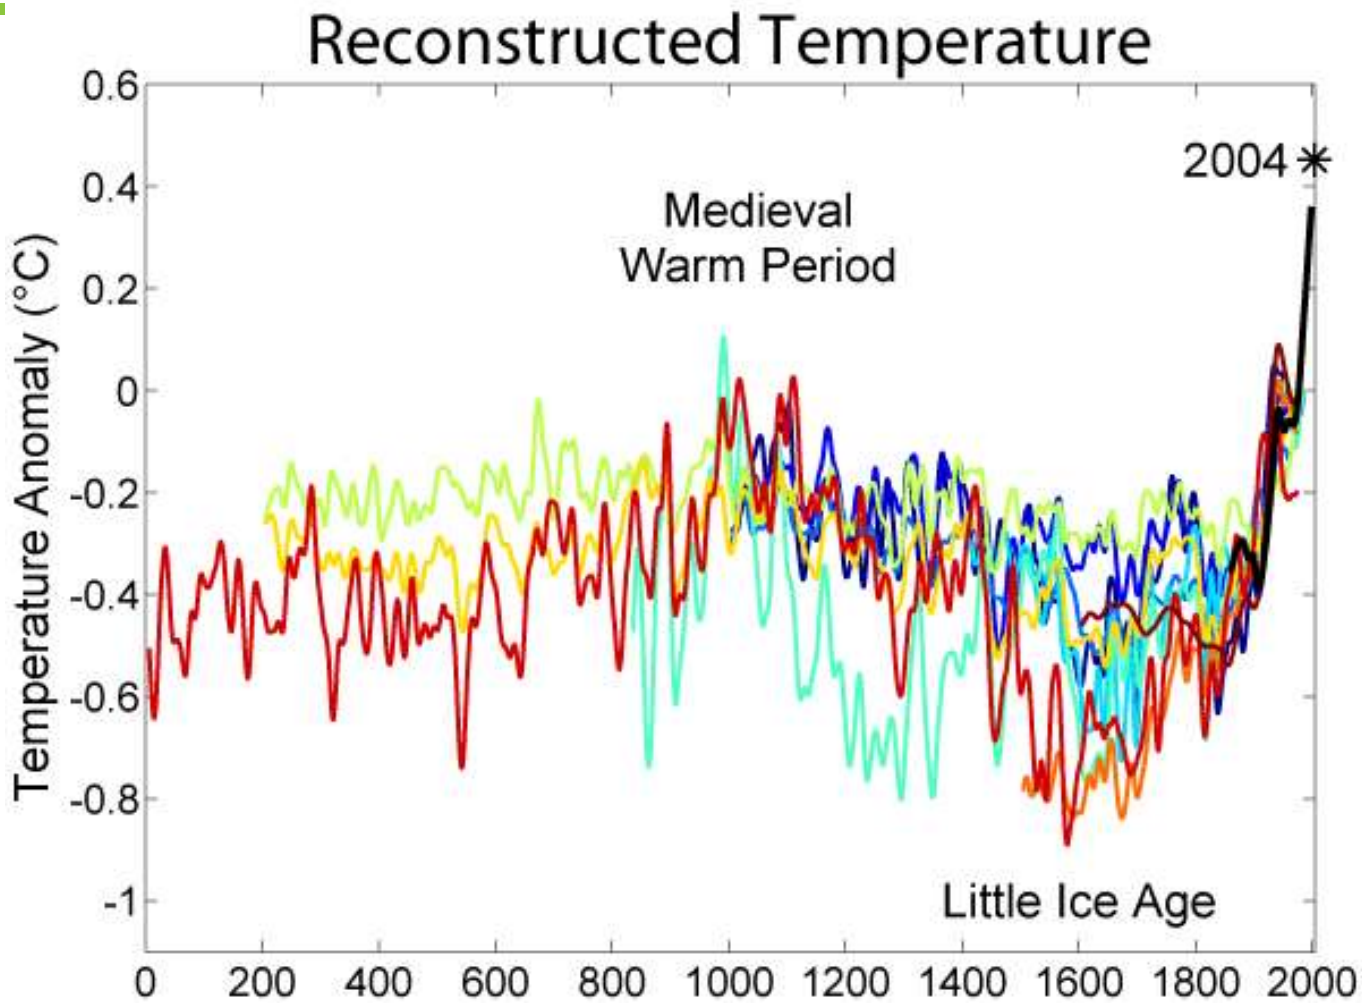

Source: www.globalwarmingart.com; Robert A. Rohde

Power Management

For Portable Consumer Electronic Devices

- **Global Warming**
  - **Not just Set top boxes**
  - **Charging the battery from the Power Grid**

Electricity consumed by US

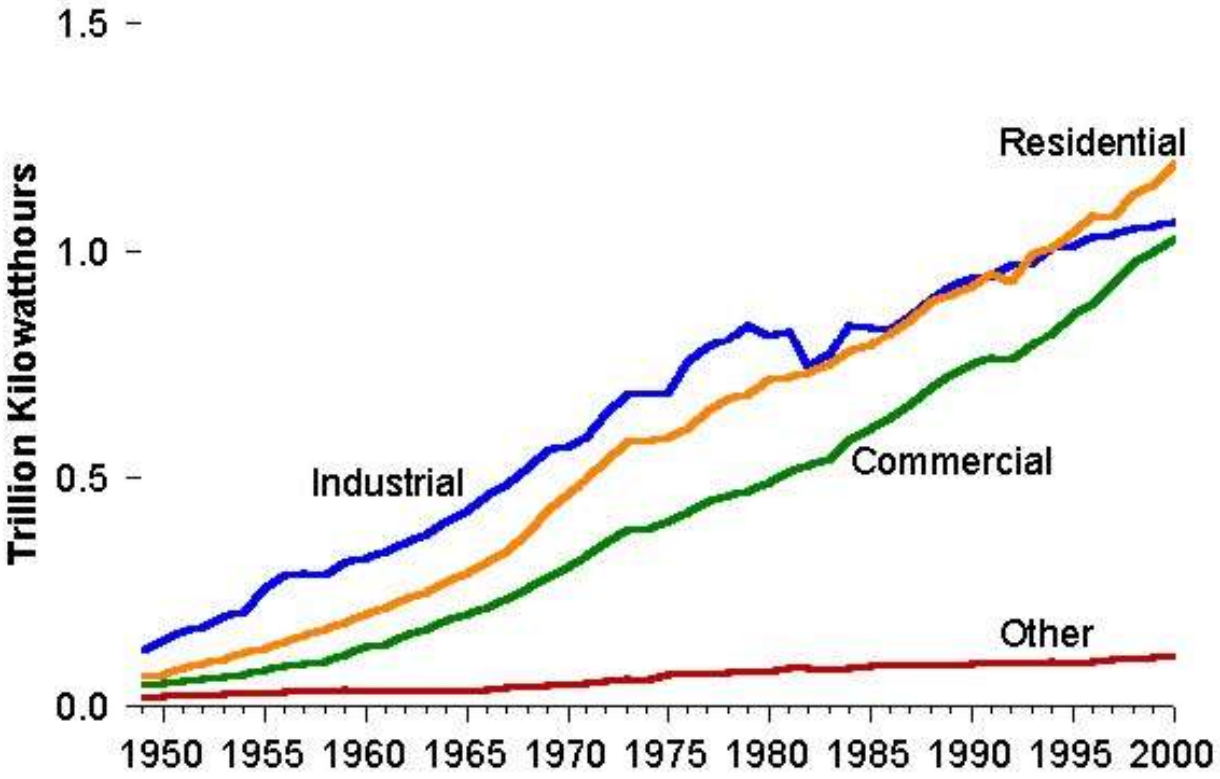

- **11 % of US house hold energy doubled in last 10 years**
  - **~ 100 million kilowatthours**

Government turns up the heat on consumer electronics; Electronic Business; 1/30/2007

http://www.eia.doe.gov/emeu/aer/eh/frame.html

Energy Star Requirements for Charging

- Criteria for Active Mode
  - Output Power (Pno)Minimum Average Efficiency in Active Mode
    - **0 to <= 1 watt >= 0.49 * Pno**
    - **> 1 to <= 49 watts >= [0.09 * Ln (Pno)} + 0.49**
    - **> 49 watts >= 0.84**
- Criteria for No-Load Mode
  - Output Power (Pno)Maximum Power in No-Load
    - **0 to < 10 watts <= 0.5 watts**
    - **>= 10 to <= 250 watts <= 0.75 watts**

Power Optimization at Every Level

- SOC Physical
  - Si Materials, Si Device Design, Device selection (gate choice), Place and Route, Power Gating (island), Back bias schemes (reduce leakage)
- SOC Gate Level:
  - Clock Gating, Clock tree synthesis
- SOC Register Transfer Level:
  - Sequential and combinatorial Clock gating
- HW System Level:
  - Processor choice, Bus choice, Memory Architecture
  - HW / SW choice for implementation
- SW Level
  - Control processor low power modes
  - Algorithmic performance
  - Optimize memory performance
  - Multicore utilization (thread management)

A System Level Issue

Power Management is a system level issue involving:

- Peripherals
  - Off chip peripheral such as LCD backlight, touch panel, camera, memory , SD/MMC controller…but also companion chip such as IPU, Wireless or cellular baseband
  - On-Chip peripherals such as bus controllers, memory controller, MultiMedia acceleration
- Core
  - Frequency (and voltage)
  - Low power modes
- The « system » itself:
  - The operating mode currently active
  - The data currently being processed
- The Tools
  - Power optimized code

Each item in this list represents an opportunity for system optimization and power consumption reduction.

HW Power Optimization

- "Timing Closure" has been the primary focus for SOC designers over the past 2 decades
  - Design tools that specifically target timing closure have generated hundreds of millions, if not billions of dollars of revenue over the history
- "Power Optimization" has become the next holy grail in HW design
  - The market for HW power design tools exceeds $100M today, yet the biggest opportunities for power optimization are in the SW layer
- Low Power design is a leading business driver for all electronic concerns

Power Management

Enabling longer battery life!

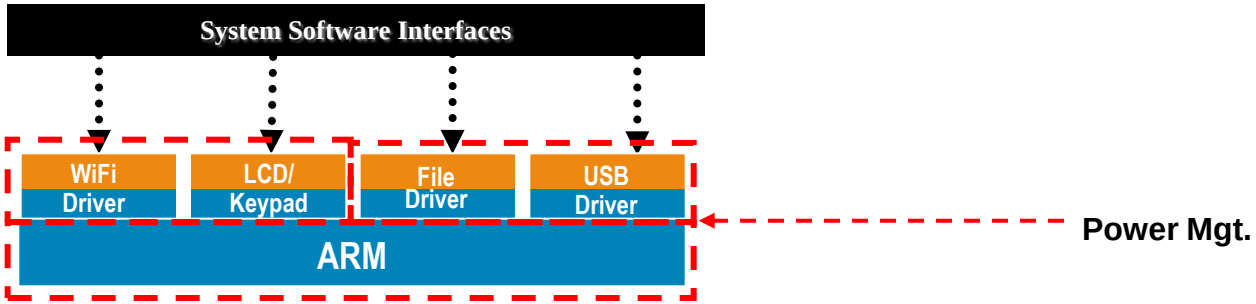

Power Management

Power Manger UI Control Demonstration

- Touch Panel
- LCD
- Backlight
- Keypad

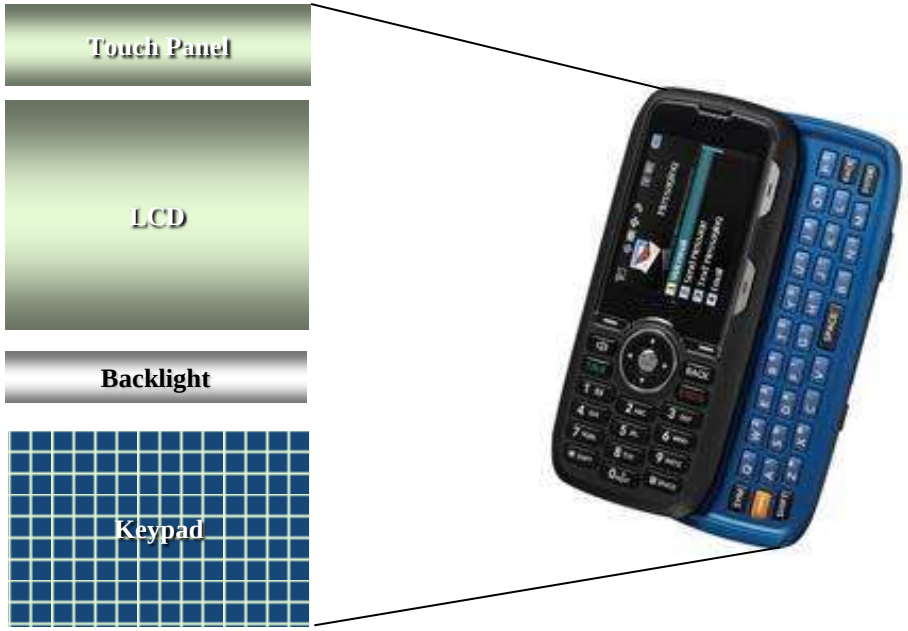

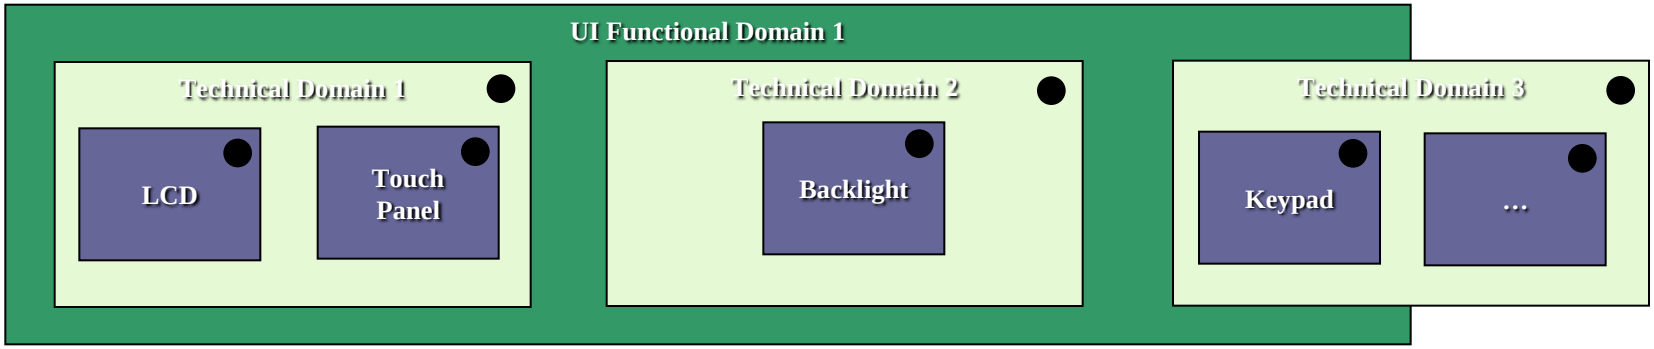

Power Management

Power Manger UI Control Demonstration

Initialization

Initialize technical power domains

Creates associations as functional domains

Initiates the states and respective drivers

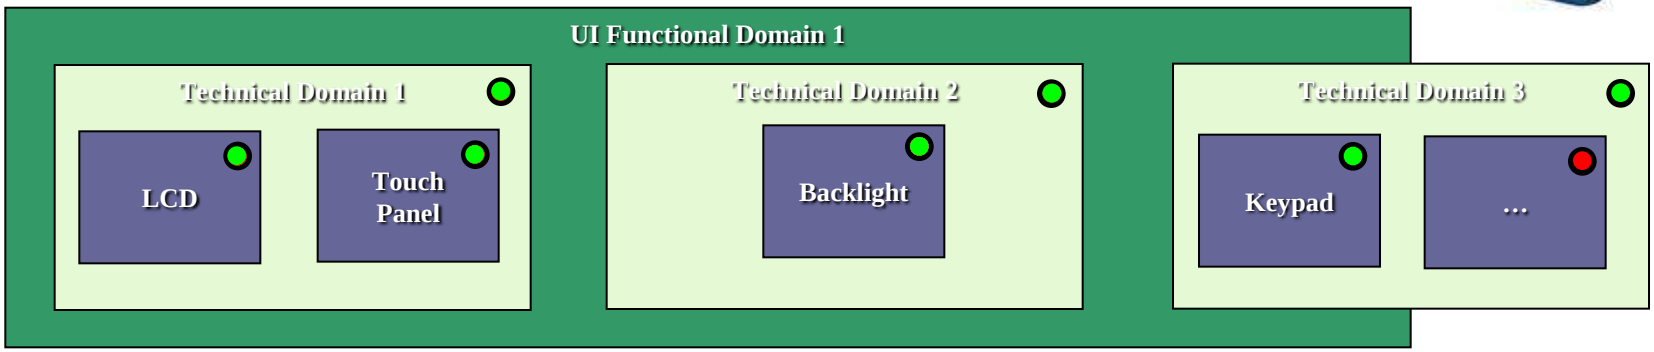

Power Management

Power Manger UI Control Demonstration

Initialize technical power domains

Creates associations as functional domains

Initiates the states and respective drivers

Reset Timer on UI functional domain

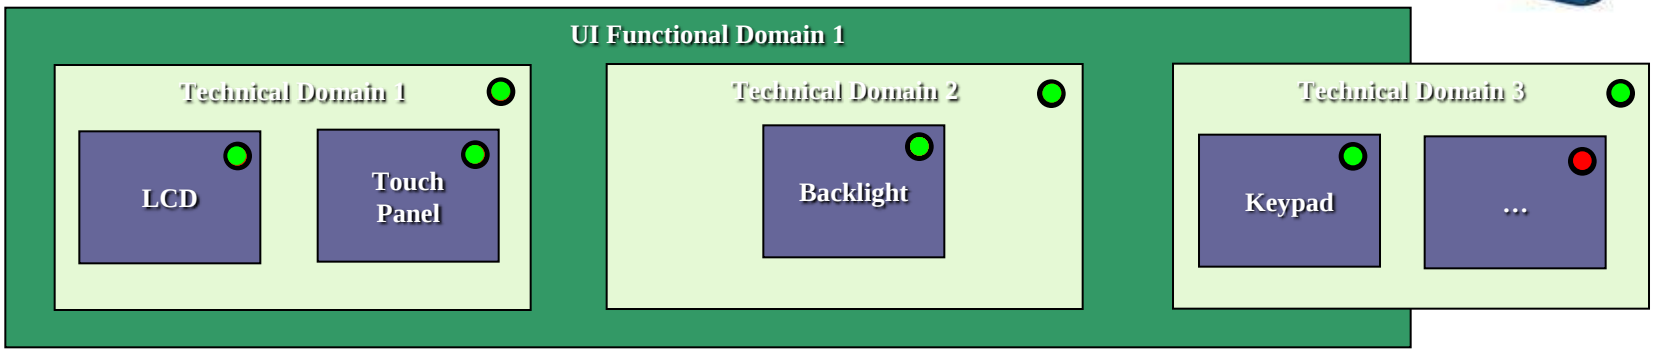

Power Management

Power Manger UI Control Demonstration

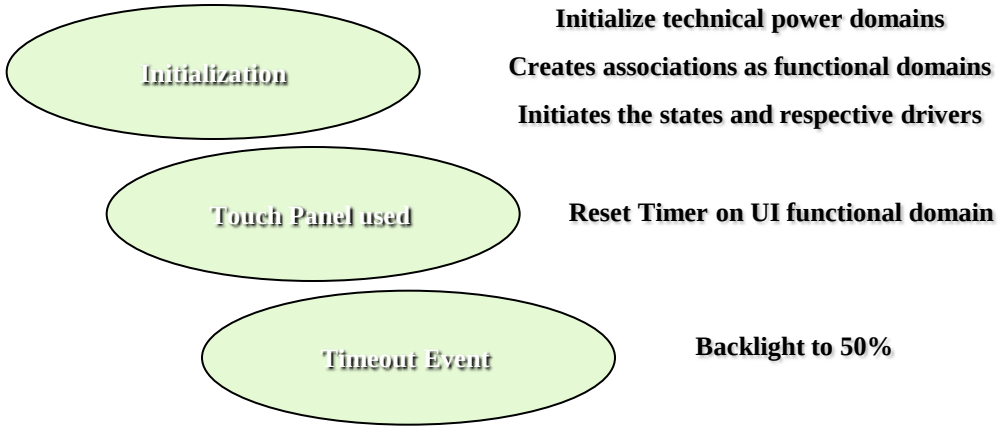

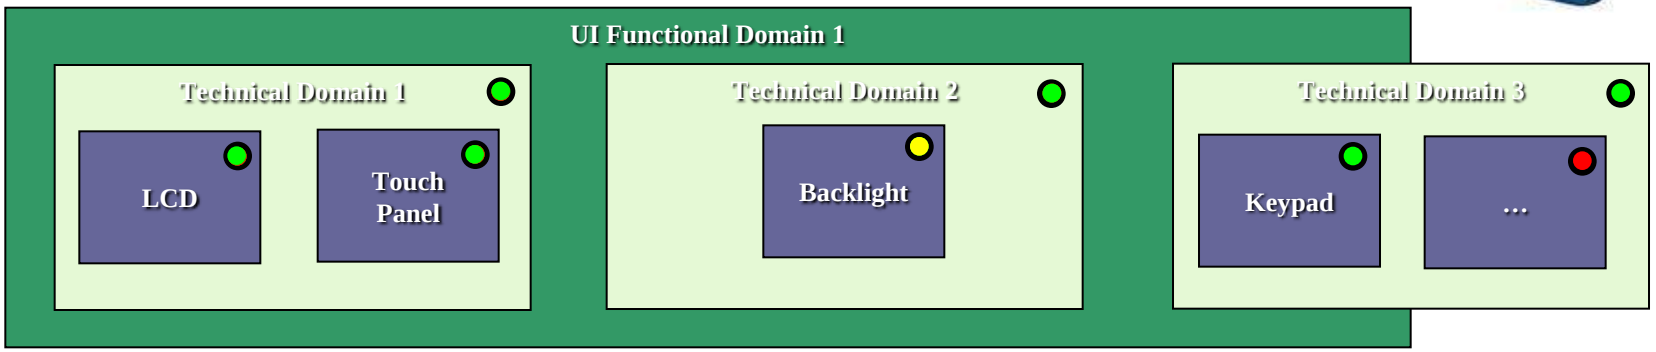

Power Management

Power Manger UI Control Demonstration

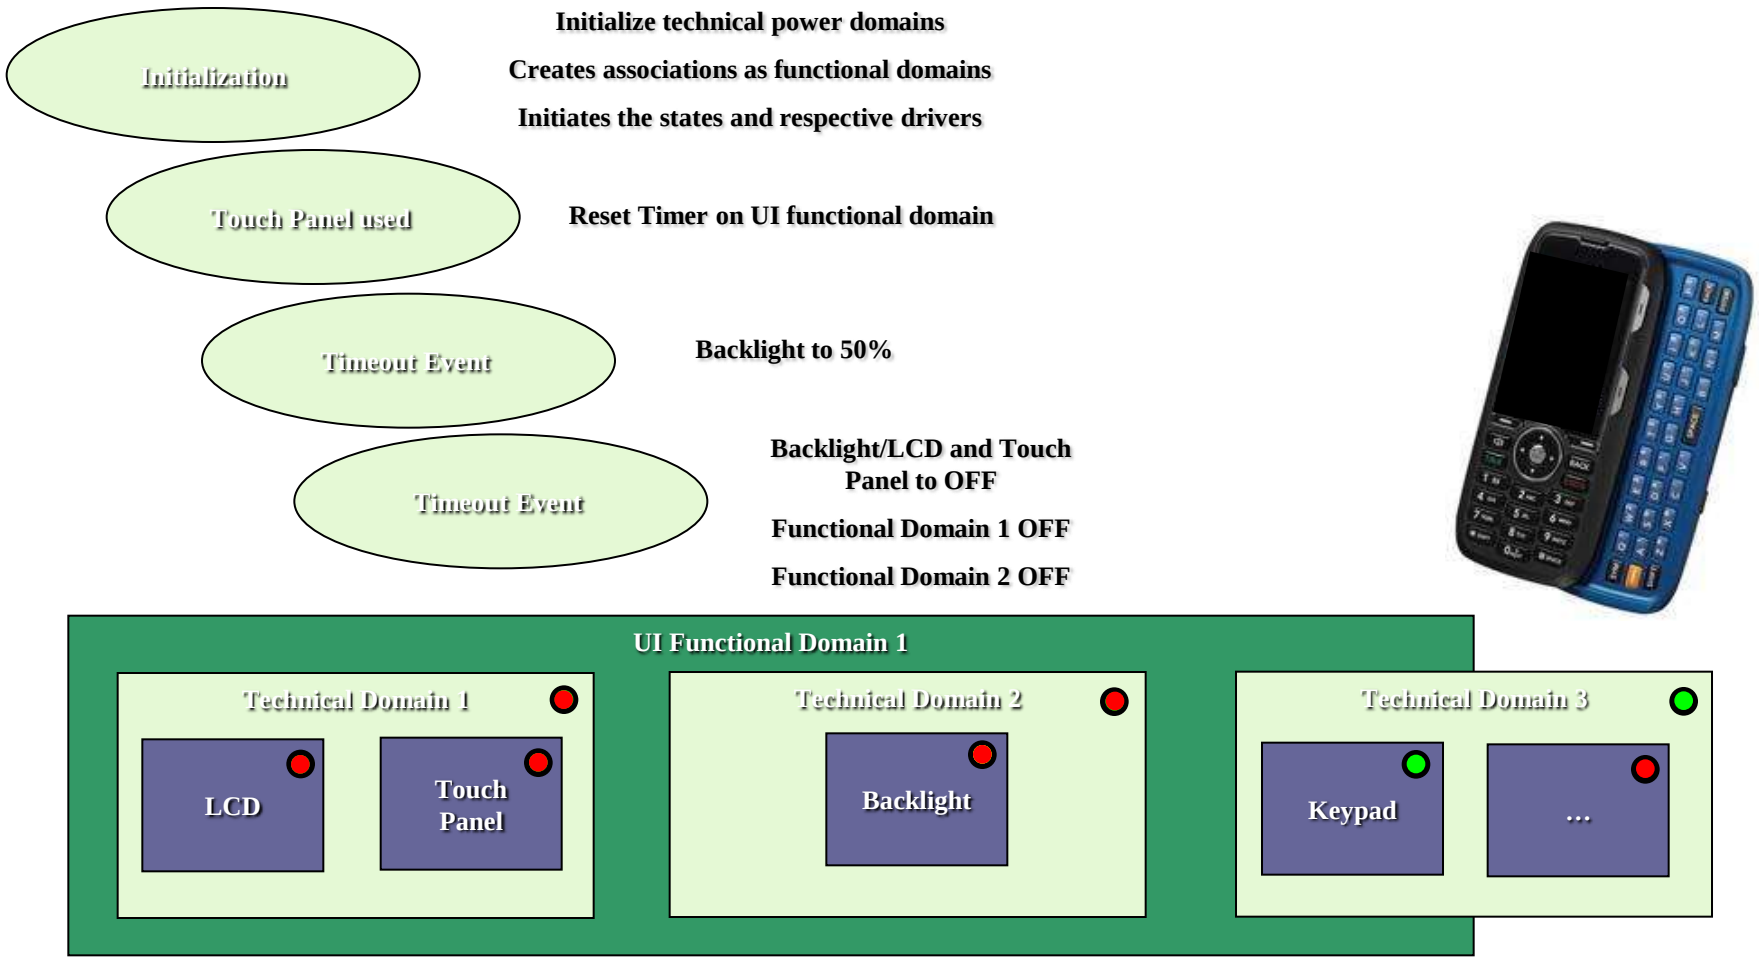

Power Management

Power Manger UI Control Demonstration

Initialization

Initialize technical power domains

Creates associations as functional domains

Initiates the states and respective drivers

Keypad Event

Functional Domain 1 ON

Reinitialize LCD

Reinitialize Touchpanel

Functional Domain 2 ON

Backlight ON 100%

Touch Panel used

Reset Timer on UI functional domain

Timeout Event

Backlight to 50%

Timeout Event

Backlight/LCD and Touch Panel to OFF

Functional Domain 1 OFF

Functional Domain 2 OFF

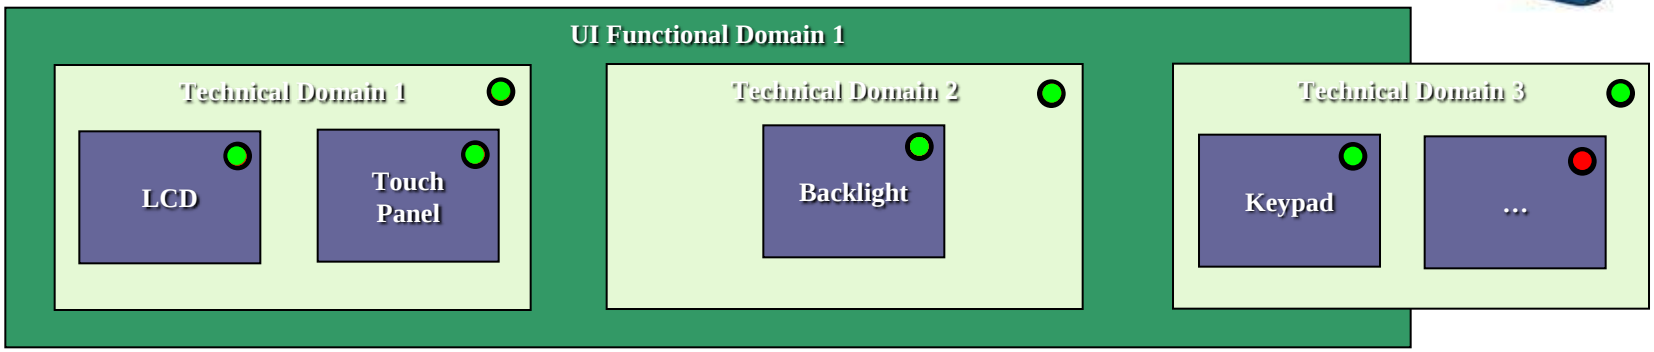

Power Management State Machine

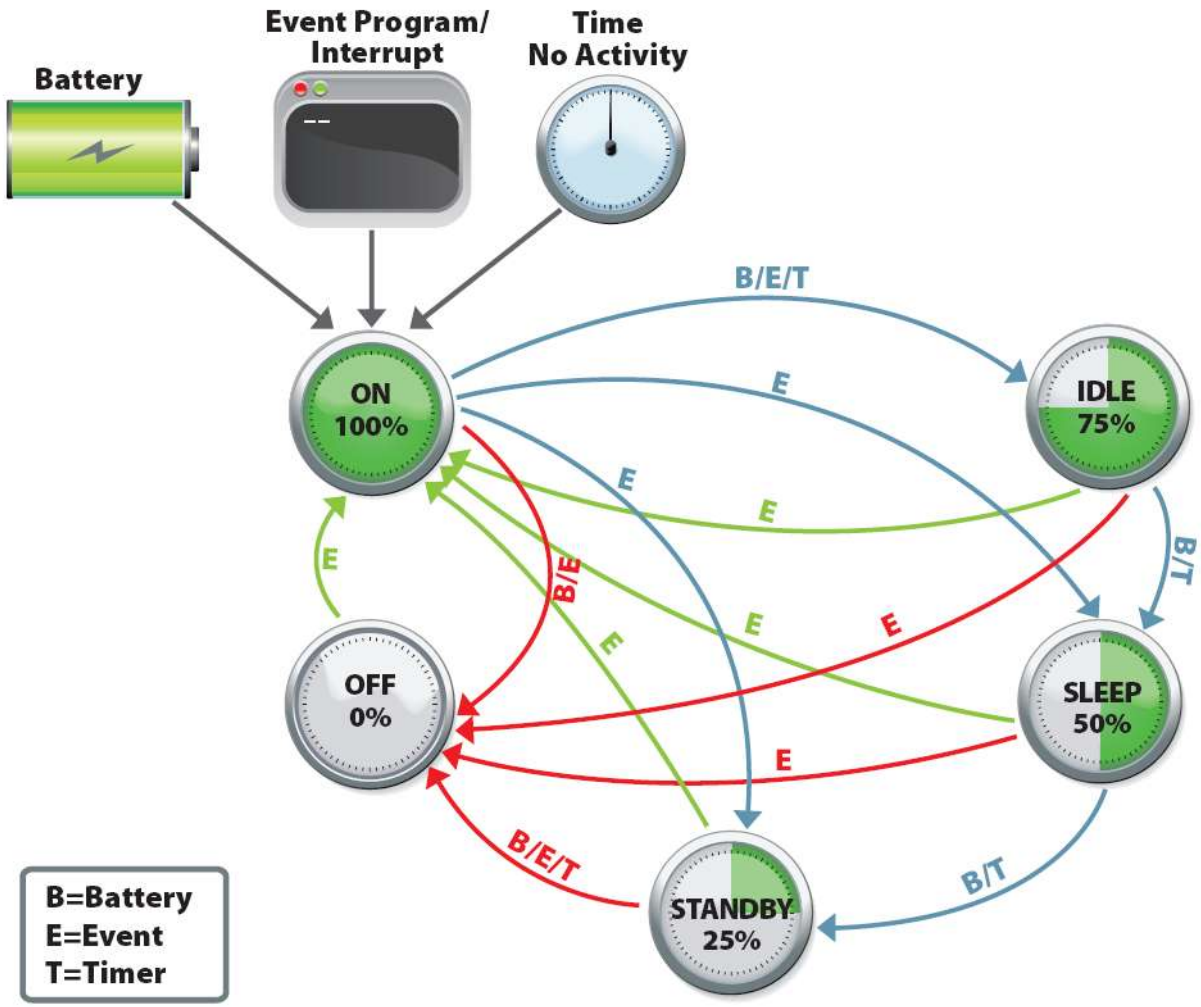

Reactive Power Management

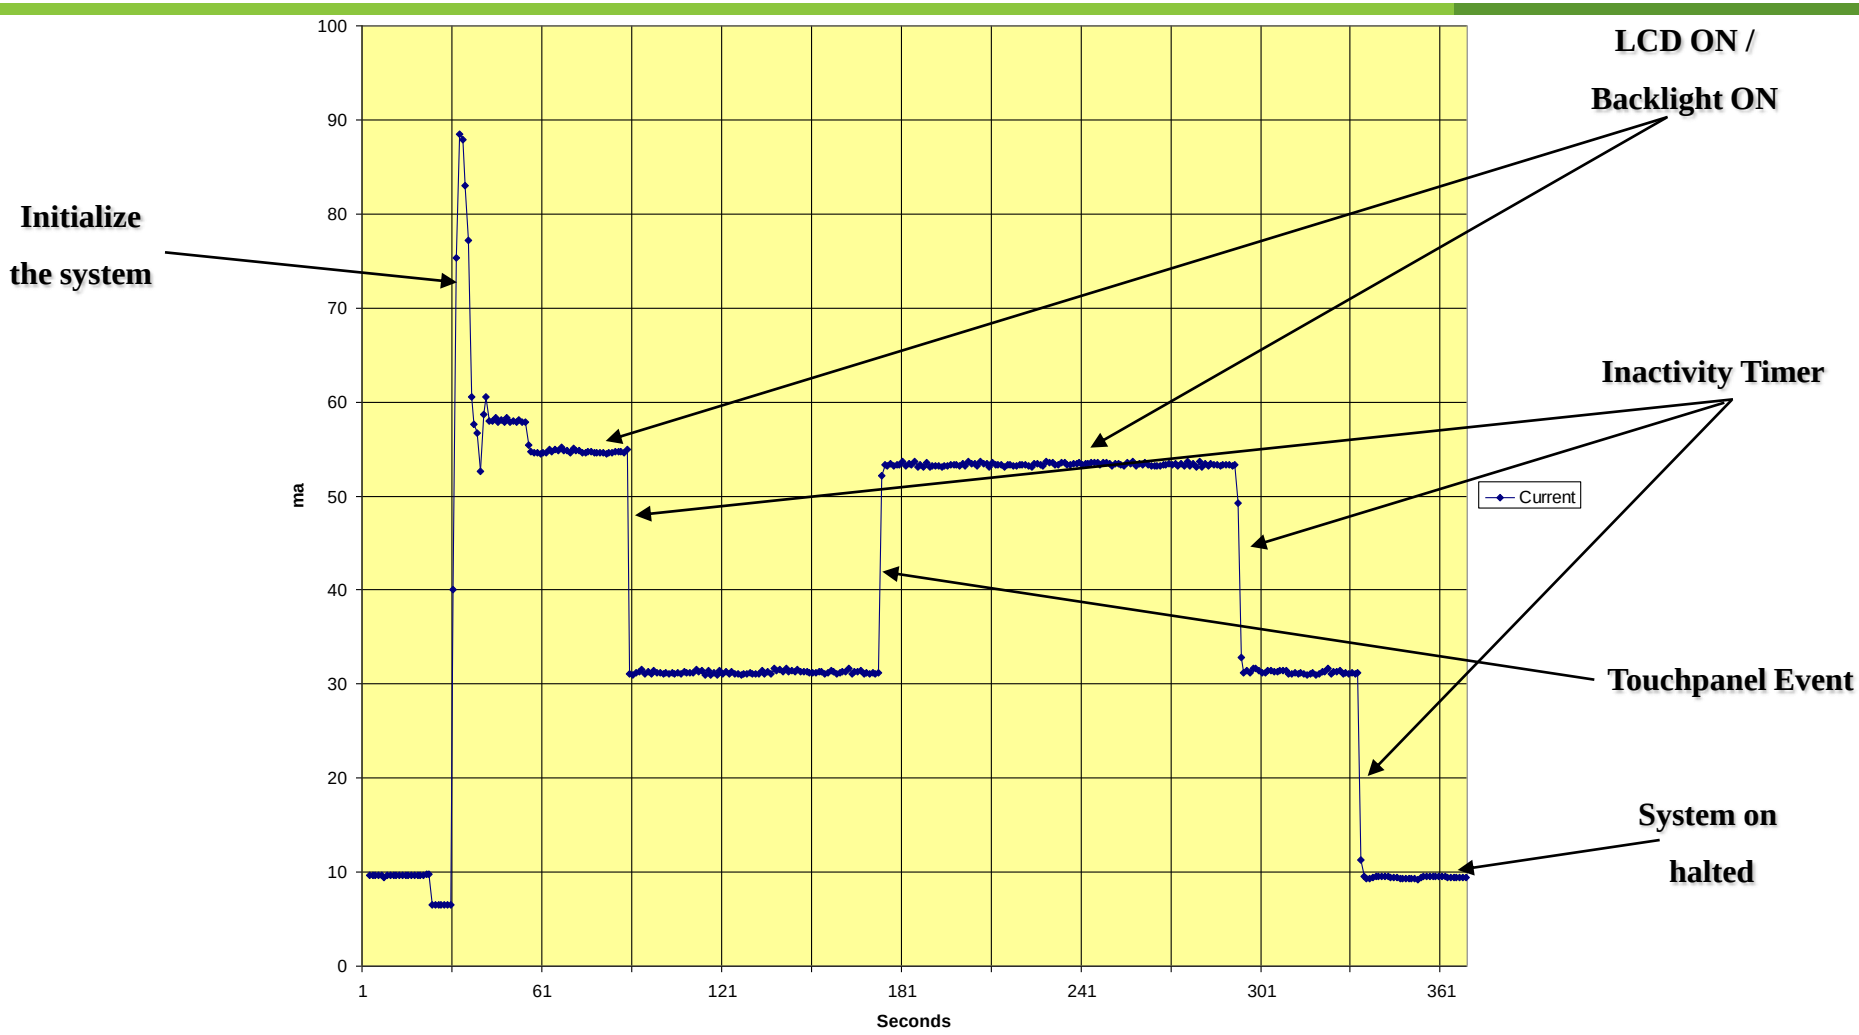

Current draw from User Interface

Power Efficiency

Speed and Peformance variance as process nodes decrease

Source: common platform conference Energy Aware Multi-Processing with ARM Cortex-A9 MPCore

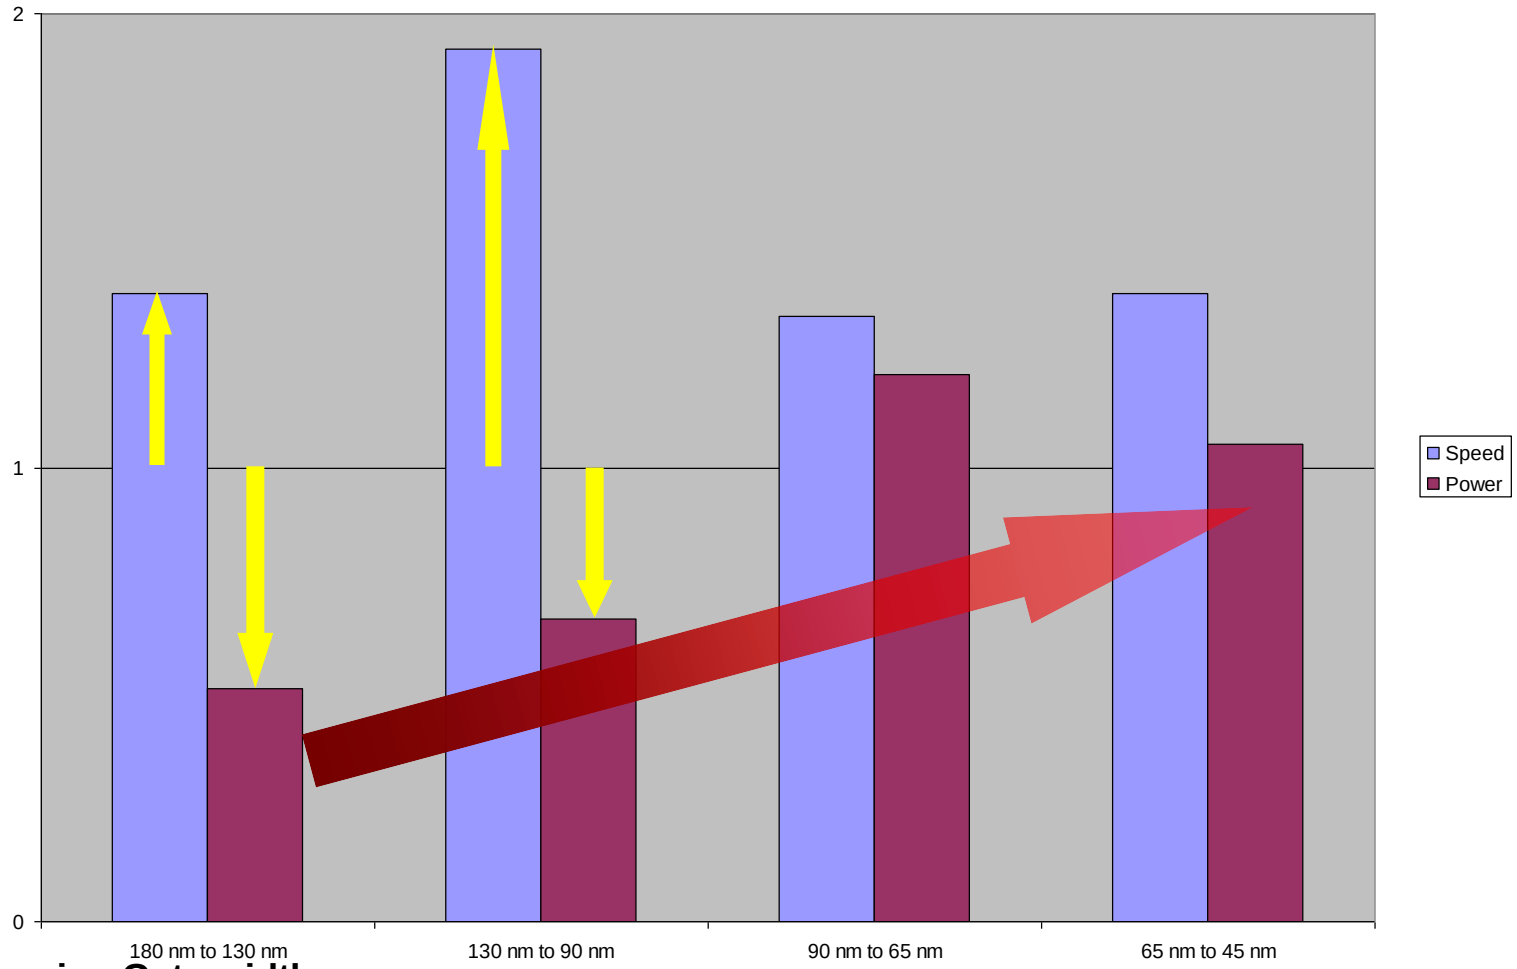

Decreasing Gate width

- Increases performance
- Initially decreased power
- Smaller geometry's hinder power savings due to both static gate leakage and increased frequency

Power Consumption

For Portable Consumer Electronic Devices

$$CV^2 f + V I_{leak}$$

Dynamic Power

Static Power

Two variables that can save power

- **Frequency (only effects the dynamic piece)**
- **Voltage (can be decreased if Frequency is also decreased)**

DVFS – Triangle Utilization Pattern

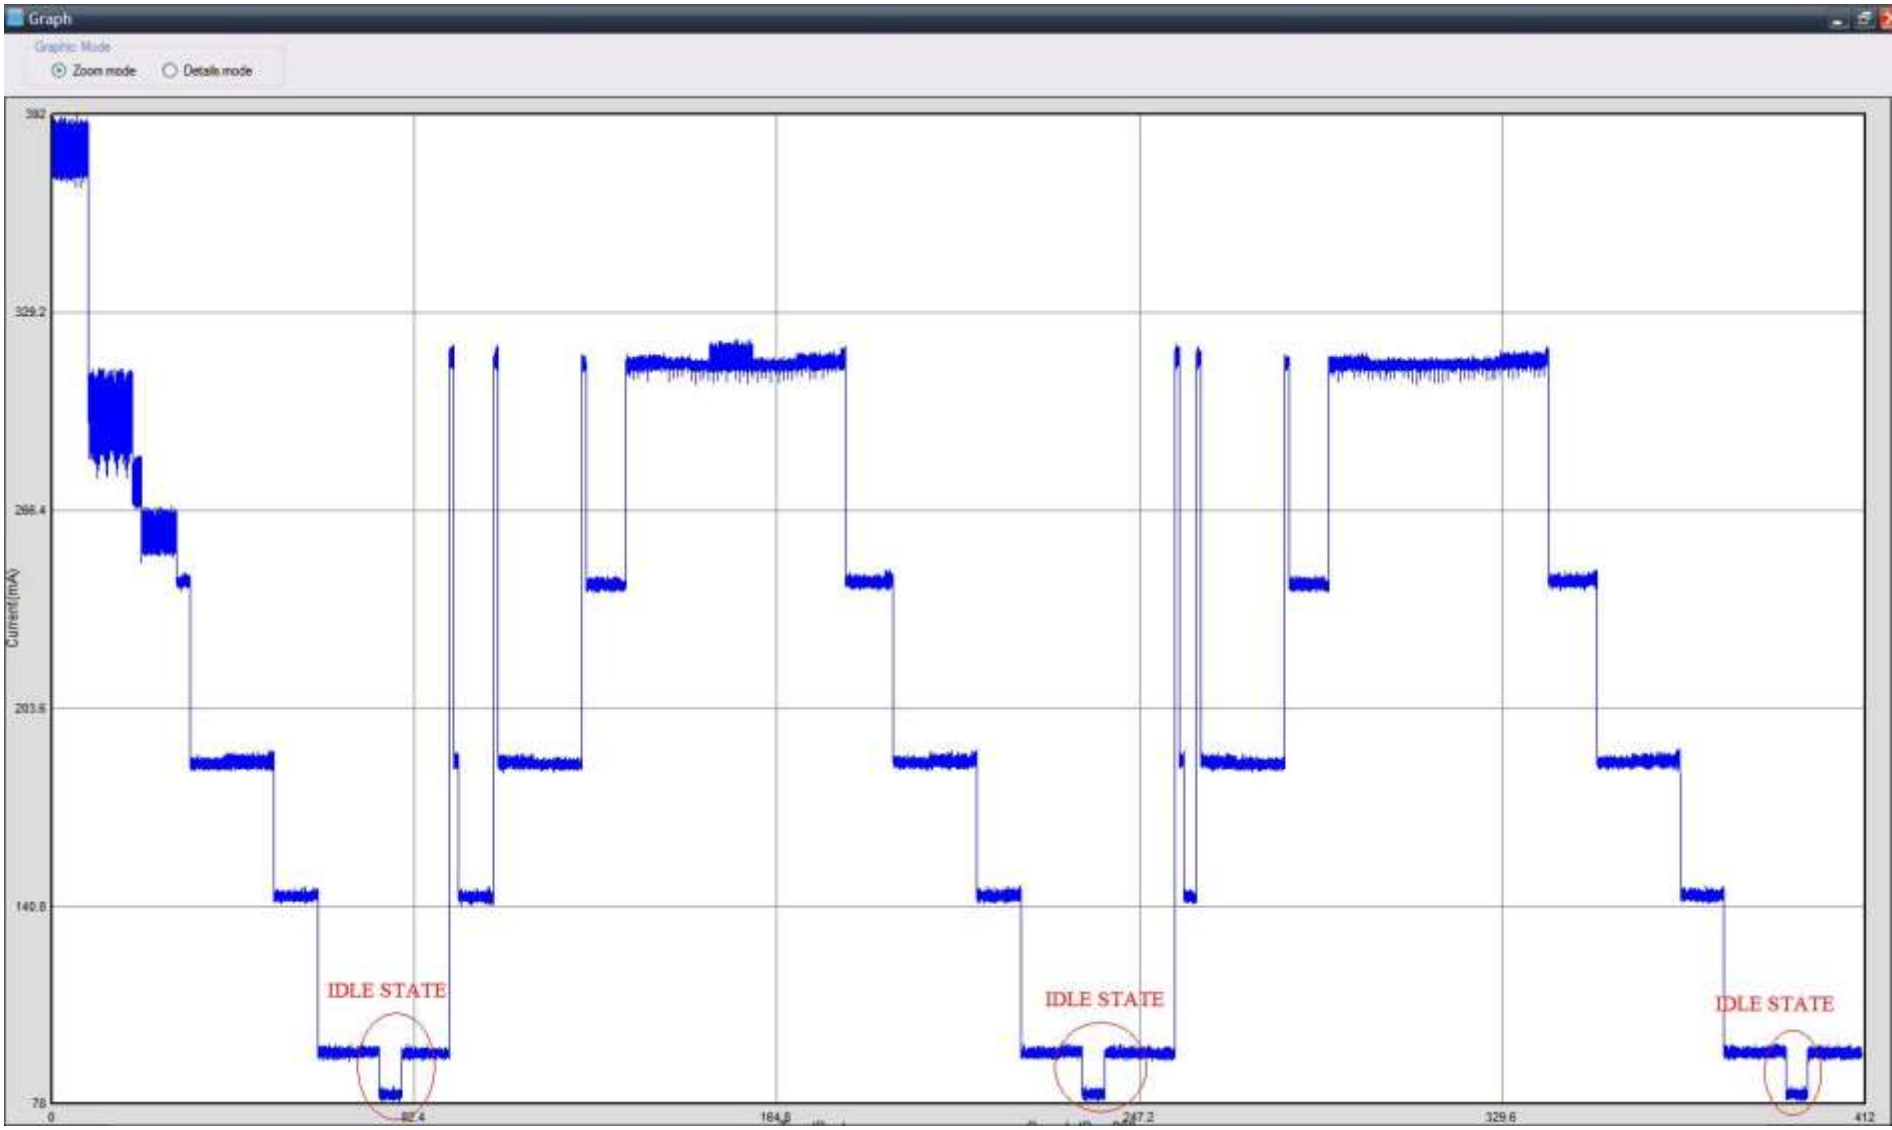

PMP Application with and without Power Management

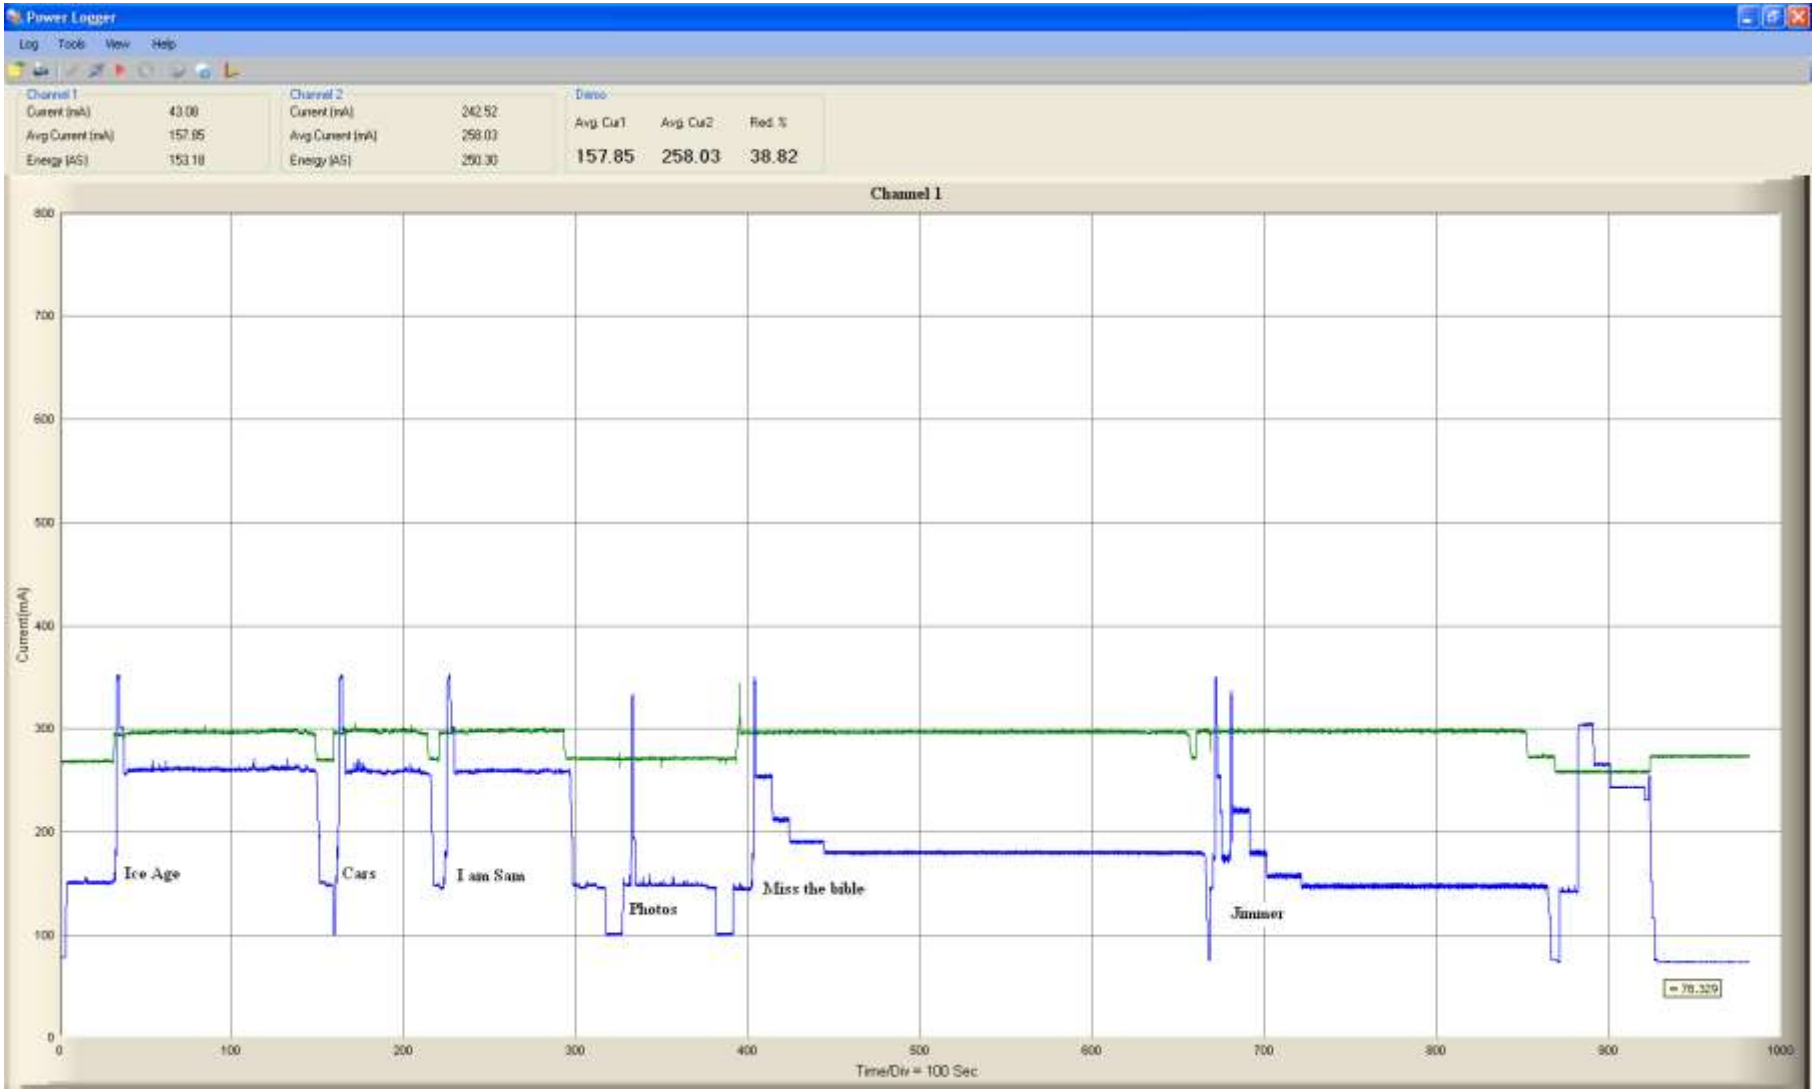

Results PMP Sinewave 220Hz

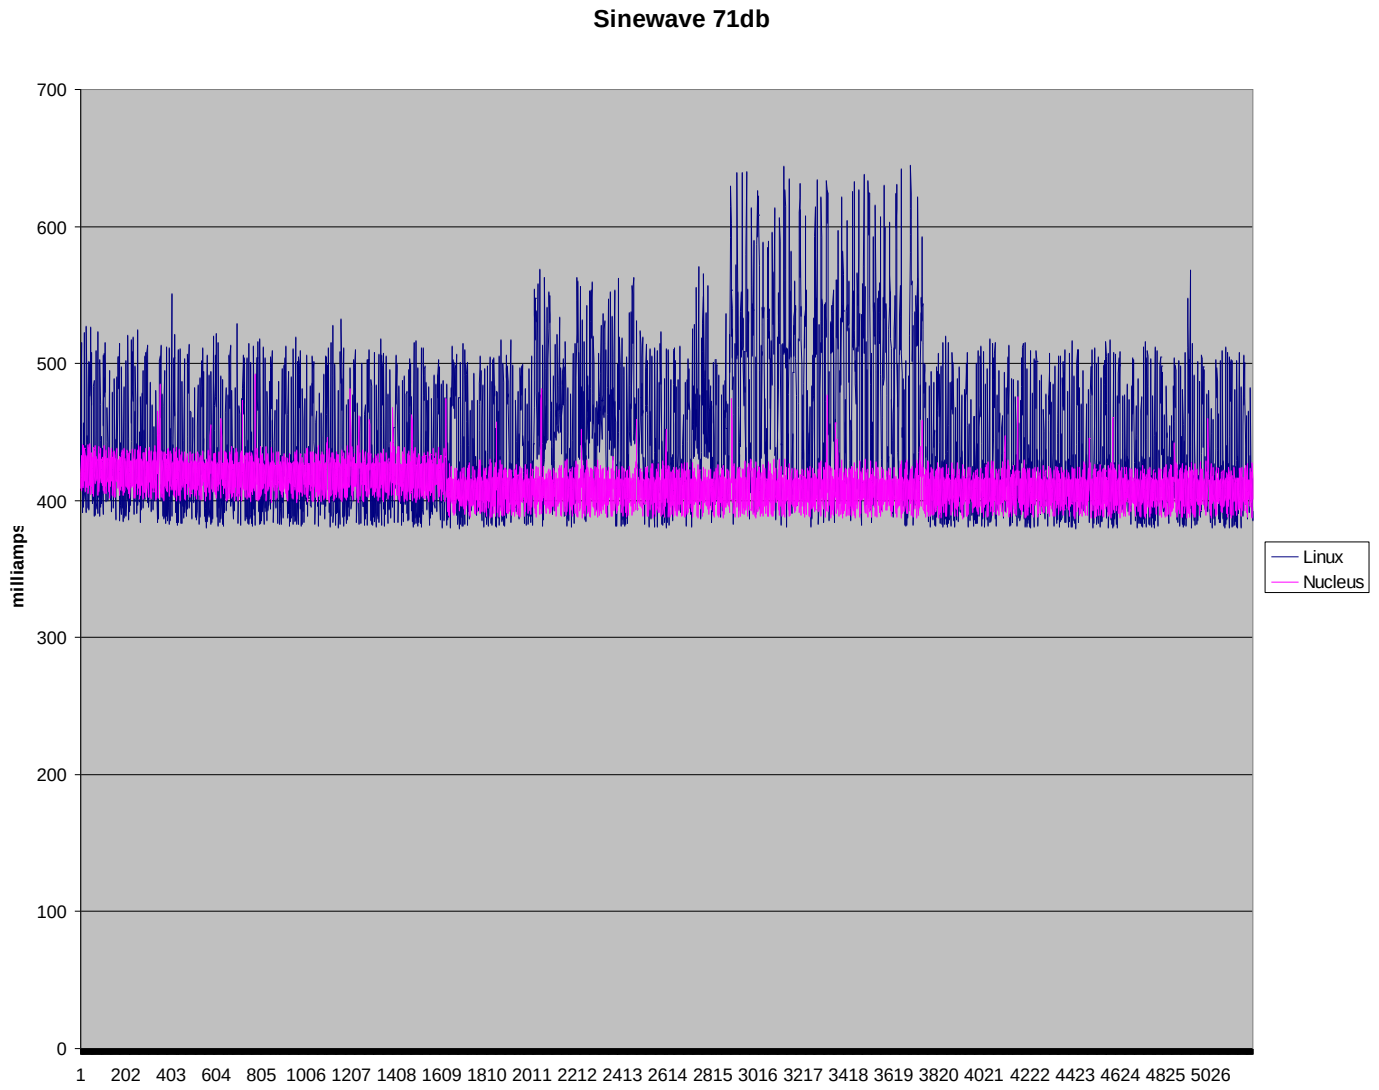

Next Generation Embedded Systems

System-Wide

Hardware

- CPU Control of sleep modes
- Dynamic Frequency and Voltage Scaling (DVFS)
- Dynamic Process Temperature Compensation (DPTC)
- Independent Peripheral Clock Gating
- Peripheral Low Power Modes

Software

- Power Aware Drivers
- DVFS enabled control of CPU
- Battery Charging Control
- Reactive Power Management
  - Use it or lose it.
  - Control power aware hardware.
- Proactive Power Management
  - Predict what the power use needs are before you need them.

Conclusion: Power Management

A System Perspective

- Are not just for battery operated devices
- Must consider the whole life cycle of how the device is utilized
- Must consider the whole infrastructure surrounding the system
- The HW must provide power saving IPs
- The OS must enable power saving techniques
- The system integrator must use them

This is our challenge for the next generation of embedded systems. The next generation depends on us.

Questions

Thanks!

www.mentor.com/embedded

System Example

System Attributes

- **Set Top Box**
- **Hardrive to store movies**
- **10 new HD movies per week**
- **Uses broadcast subcarrier**
- **Always on**
- **Has USB and Ethernet**
- **Runs COTS OS**
- **User can select any movie of the 100 stored on STB.**
- **Oldest movies are deleted**

System Example

System Attributes

- **Set Top Box**
- **Hardrive to store movies**
- **10 new HD movies per week**
- **Uses Internet multi-cast**
- **Has back-channel**
- **User can select any movie of the 100 stored on STB.**
- **Oldest movies are deleted**
- **Only on for downloads and on demand for viewing movies**

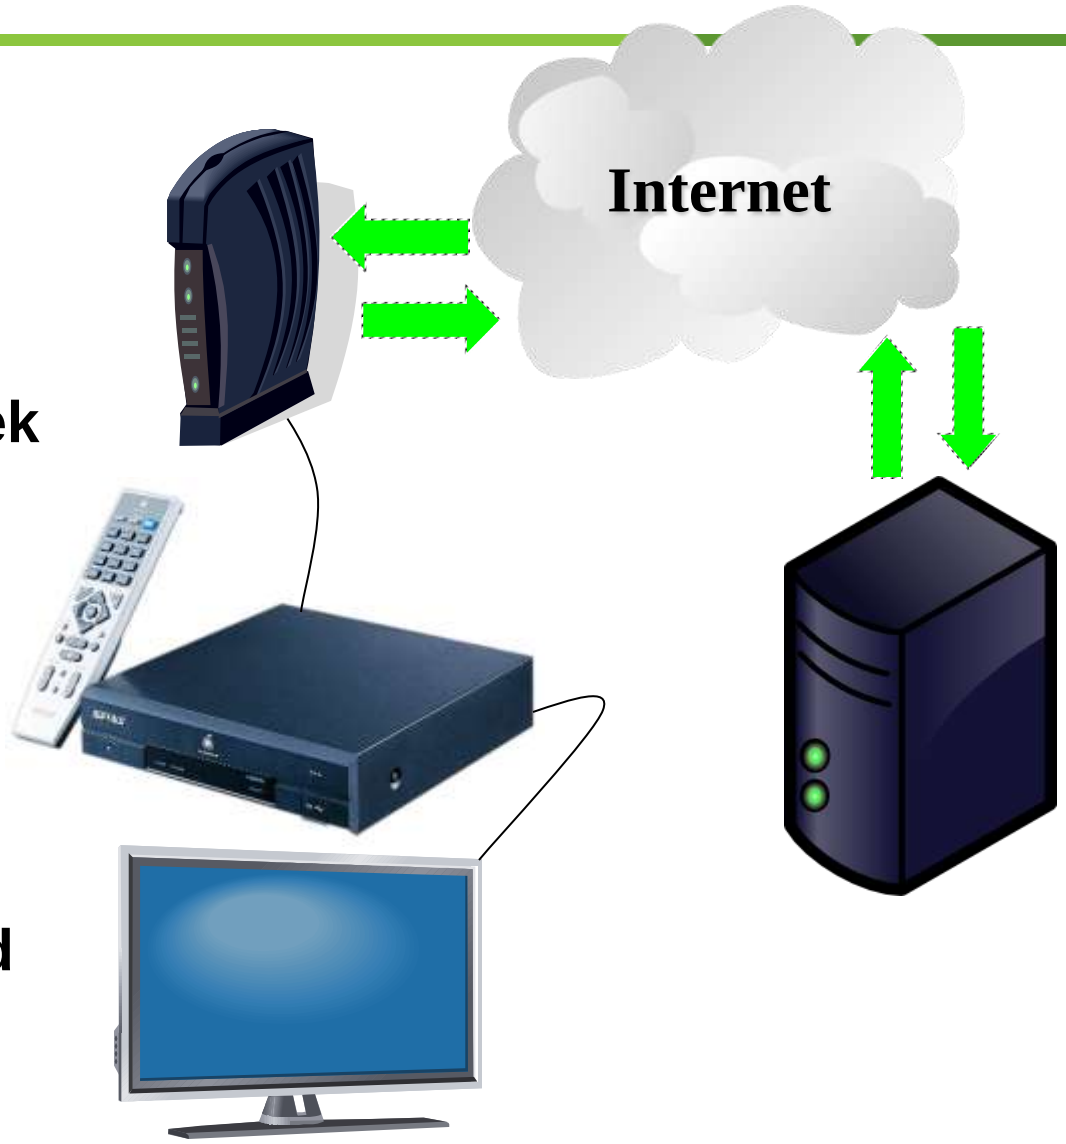

Thanks!

www.mentor.com/embedded

Power Optimization at Every Level

- SOC Physical
  - Si Materials, Si Device Design, Device selection (gate choice), Place and Route, Power Gating (island), Back bias schemes (reduce leakage)
- SOC Gate Level:
  - Clock Gating, Clock tree synthesis
- SOC Register Transfer Level:
  - Sequential and combinatorial Clock gating
- HW System Level:
  - Processor choice, Bus choice, Memory Architecture
  - HW / SW choice for implementation
- SW Level
  - Control processor low power modes
  - Algorithmic performance
  - Optimize memory performance
  - Multicore utilization (thread management)

AGENDA

- Introduction to power management
  - A System level issue
  - The business context
  - Power Optimization at every level
- Nucleus Power Manager Phase 1
  - Power Model
  - Peripheral Power Management
  - Core Power Management
- Next Steps and External loop

A System Level Issue

Power Management is a system level issue involving:

- Peripherals
  - Off chip peripheral such as LCD backlight, touch panel, camera, memory , SD/MMC controller…but also companion chip such as IPU, Wireless or cellular baseband
  - On-Chip peripherals such as bus controllers, memory controller, MultiMedia acceleration
- Core
  - Frequency (and voltage)
  - Low power modes
- The « system » itself:
  - The operating mode currently active (the use case in progress)
  - The data currently being processed
- The Tools
  - Power optimized code

Each item in this list represents an opportunity for system optimization and power consumption reduction.

HW Power Optimization

- "Timing Closure" has been the primary focus for SOC designers over the past 2 decades
  - Design tools that specifically target timing closure have generated hundreds of millions, if not billions of dollars of revenue over the history
- "Power Optimization" has become the next holy grail in HW design
  - The market for HW power design tools exceeds $100M today, yet the biggest opportunities for power optimization are in the SW layer
- Low Power design is a leading business driver for all electronic concerns

Power Model

The power model represents the system power states

The Power manager moves from one mode to the other based on events such as battery level, timer, user event, peripheral notification, interrupt etc.. Each transition is attached a dedicated processing. Such processing can be simple (register operation) or complex (frequency scaling, system shut down)

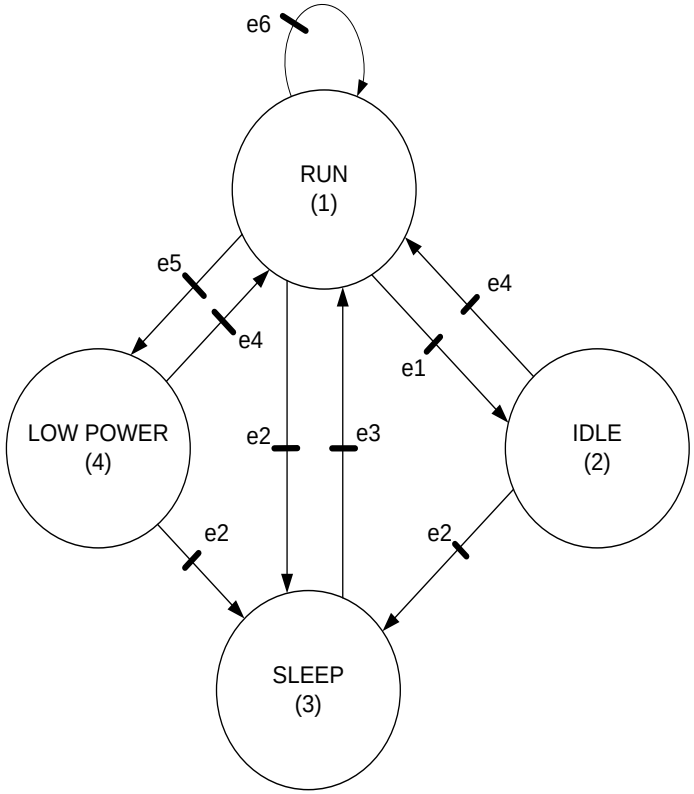

	e1	e2	e3	e4	e5	e6
s1	S2, f1	S3, f8			S4, f2	S1, f0
s2		S3, f9		S1, f7		
s3			S1, f7			
s4		S3, f6		S1, f7		

Battery powered systems

Peripheral Power Management

Peripherals have a large contribution to the overall power consomption. They deserve a special attention.

In most cases, it is possible to manage the peripheral power modes independently

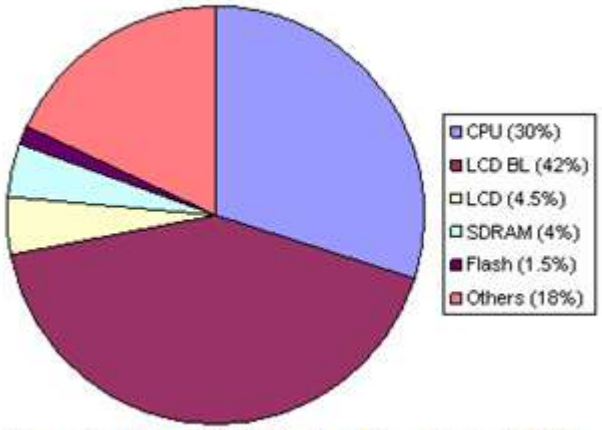

Figure 1. Power Distribution for a Typical PDA

Peripheral Power Management (cont.)

Each peripheral is associated with its own power model, implemented as state machine just like the system power model.

All the power-enabled peripherals then report their status to the system power model where system level decisions can be made.

Not all the peripherals are independant from the logical or functional view point.

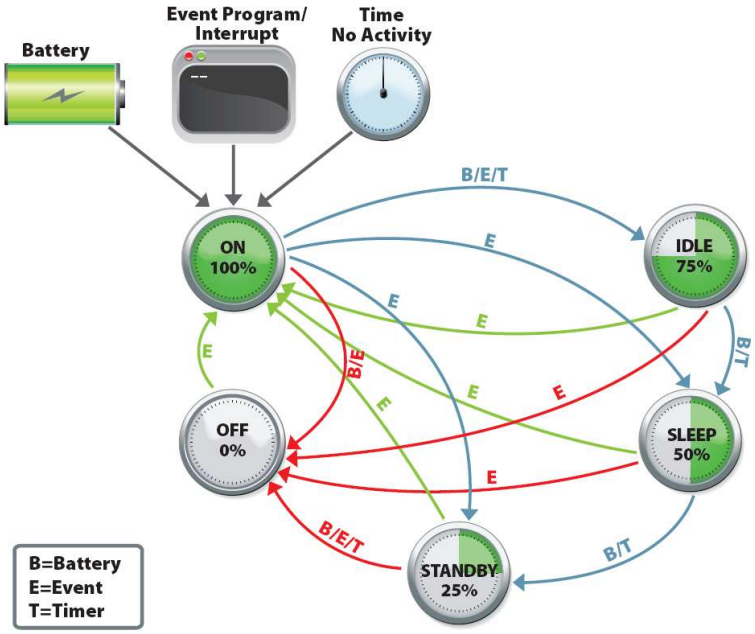

Core Power Management

The power consumption of CMOS digital circuits is proportional to the frequency and to the square of the core supply voltage.

$$CV^2 f + V I_{leak}$$

Dynamic **Static**

The static contribution is very small (<figure needed>), so the focus is on the dynamic part: CV^2f

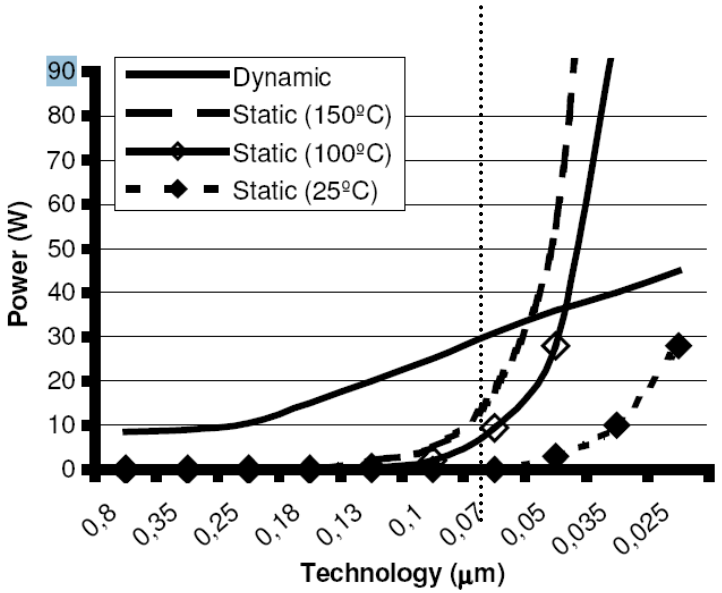

The system does not require 100% of the processing power all the time. So, modern processors offer means to adapt the processing capacity to the system demand:

- DVFS (during processing)
- Low power modes (during inactivity)

Core Power Management (Cont)
DVFS

- DVFS management is about setting the most appropriate CPU function point, based on the system need.
- The challenge is to be able to evaluate and/or predict the system need at a particular point in time. Multiple policies can be defined with different levels of complexity.
- We will start with a « minimum idle time policy » that implements a negative feedback loop.
- Note: Voltage and Frequency are linked: not all frequency are possible at a particular voltage. The processor has actually a discrete number of function point available: e.g. for PXA270: (0.9,104), (1.15, 208)

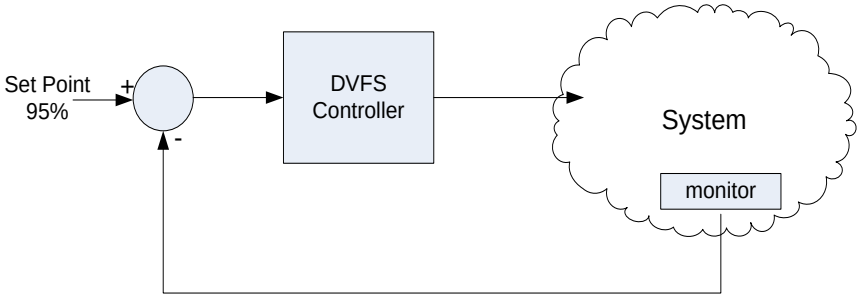

A Negative Feedback for Power Optimization

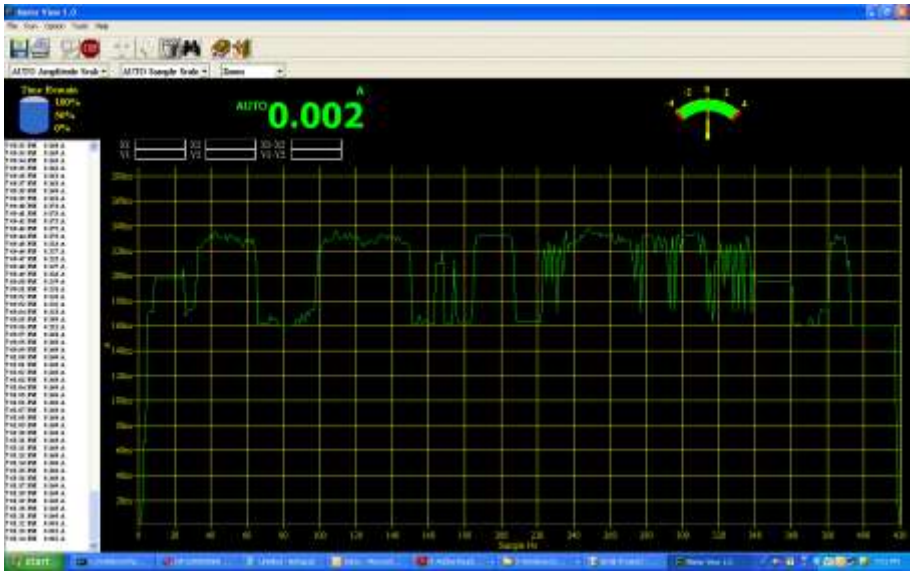

Core Power Management (Cont)

Low Power Modes

- Low power mode management is about placing the core processor into the lowest possible power mode when there is no processing demand. Modern processors have multiple low power modes: Idle, Stand-by, Sleep, Deep-Sleep

- The challenge is to be able to set the ‘alarm clock ’ so that no event is missed. Bearing in mind that wake up time can be long.

PXA270 low power modes

- Running mode:
  - 925mW @ 625MHz
  - 44mW at 13MHz.
- IDLE mode
  - 260mW @ 624MHz
  - 8.5mW at 13MHz.
  - It takes 1usec to enter / exit this mode
- Stand By:
  - 1.7 mW.
  - It takes up to 11msec to exit this mode.
- Sleep Mode
  - 0.16mW
  - it takes up to 136msec to exit this mode.
- Deep Sleep Mode
  - It consumes 0.1 mW.
  - It takes up to 261msec to exit this mode.

Next Steps and External Loop

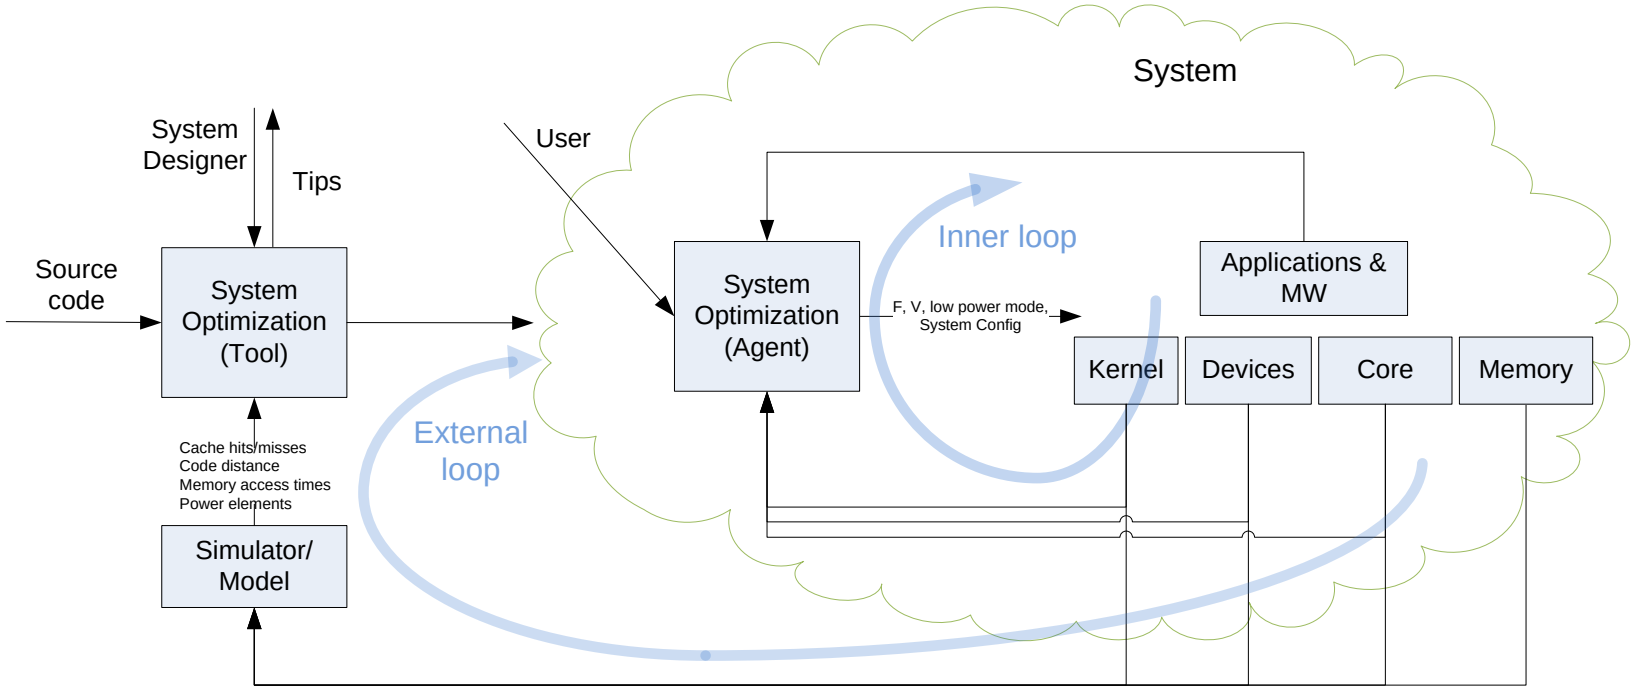

- **Power management agent takes more inputs and output to more component**
- **The external loop optimizes the memory mapping and code efficiency**

SUMMARY

- We are committed to power saving.
- The first implementation of our power framework will have
  - Peripheral management
  - DVFS with a ‘minimal idle thread’ policy
  - Simple low power modes management
- Moving forward, the framework will have
  - more policies, fine grain control over the peripherals, more kernel awareness, smart low power modes management.
  - We also are looking at providing the user with power profile integrated into our tool offering

Thank you!

Questions?

Background

ST Low Power Solutions Roadmap

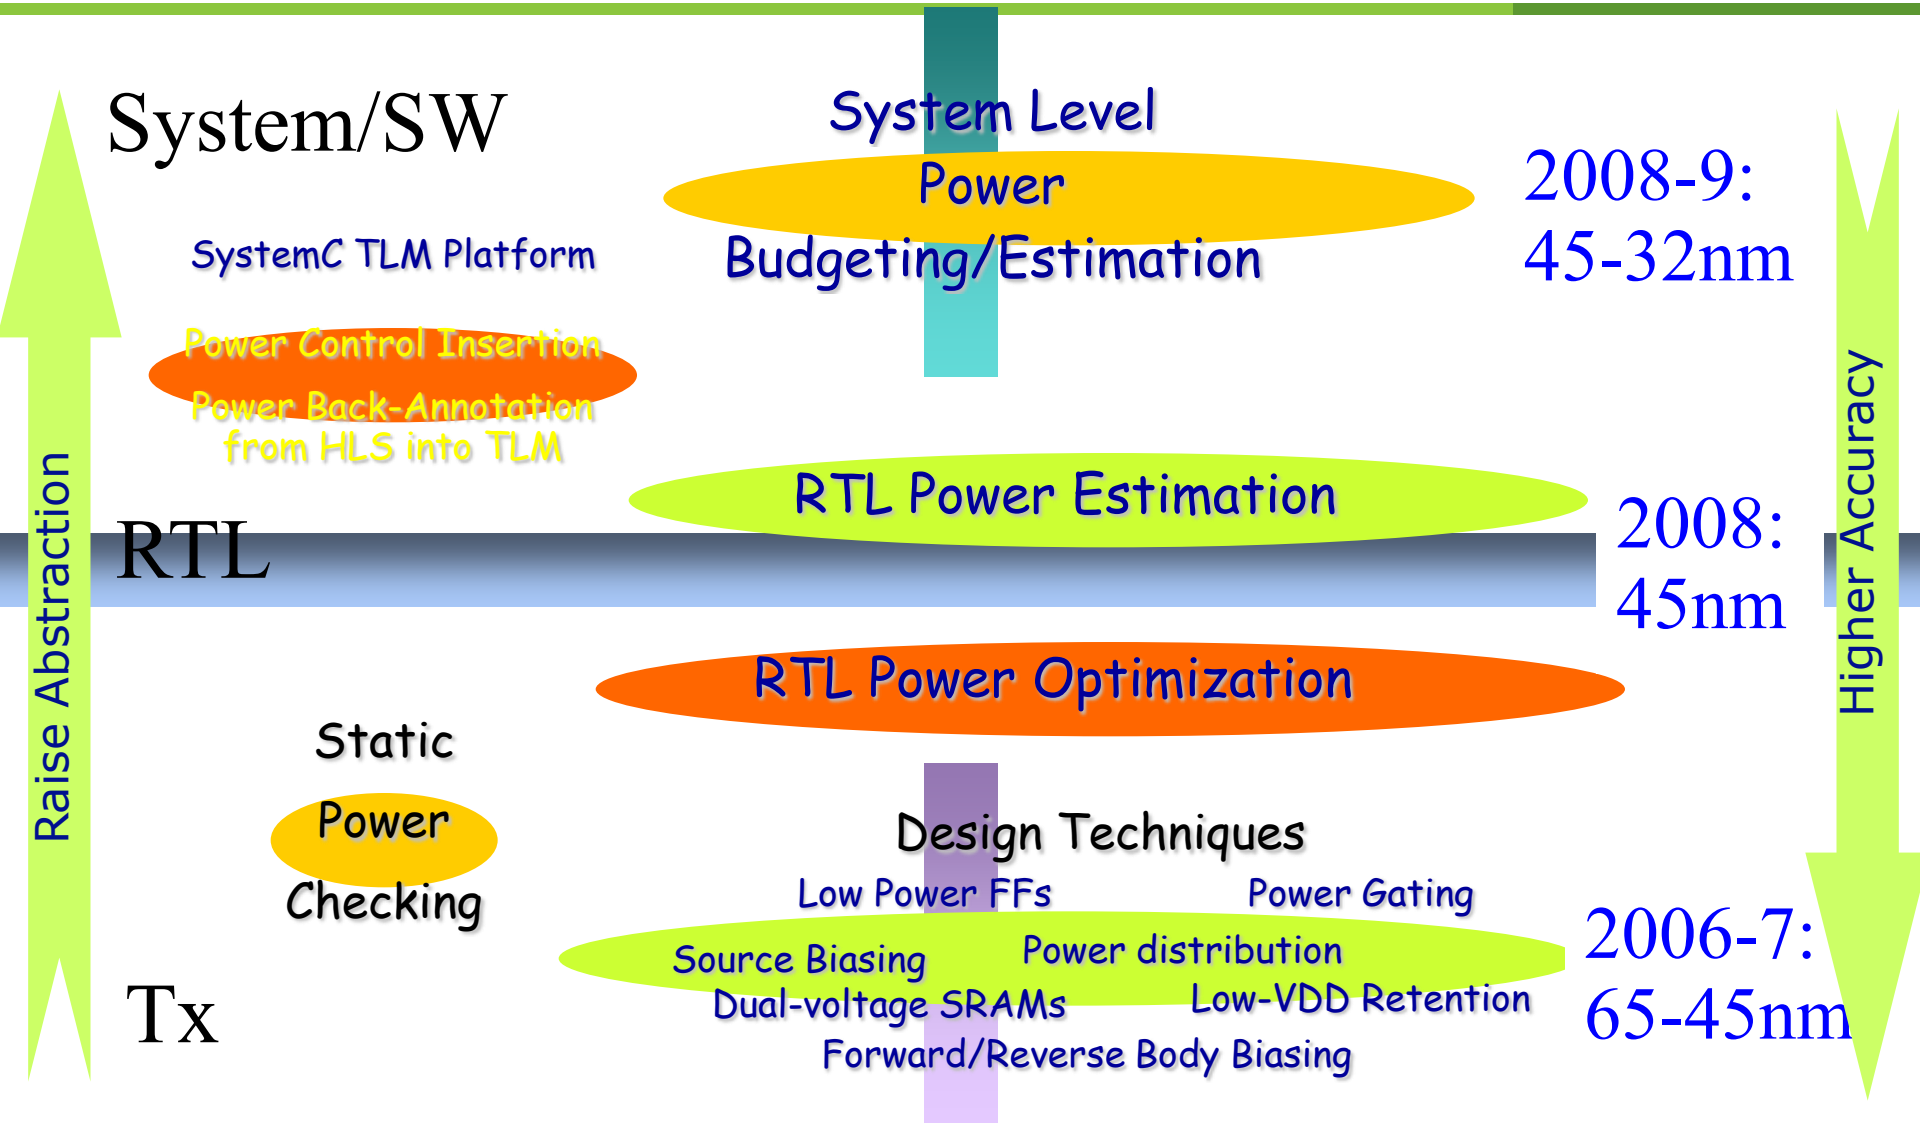

Environmental Effects Regulations

In 2004-2006 a Cisco business unit shipped 5 Million Systems With a total of 500 Megawatts Usage

If these 5 Million Systems had shipped with 1 Watt LESS each…

5 Megawatts of Power would not have shipped

16 Million KWatt Hours saved per Year

$3.8-$7 Million Electrical Expense

31 K Metric Tons CO_2 eliminated from Atmosphere

Double for Cooling & Auxiliary (C & A)

Equivalent to taking 12,000 Cars off the Road…

Big Volume Amplifies Small Gains!

Source: Cisco **Lynelle McKay**, Senior VP and General Manager Networking & Computing Systems Group , a Cisco Business Unit

Cost of Power Cooling

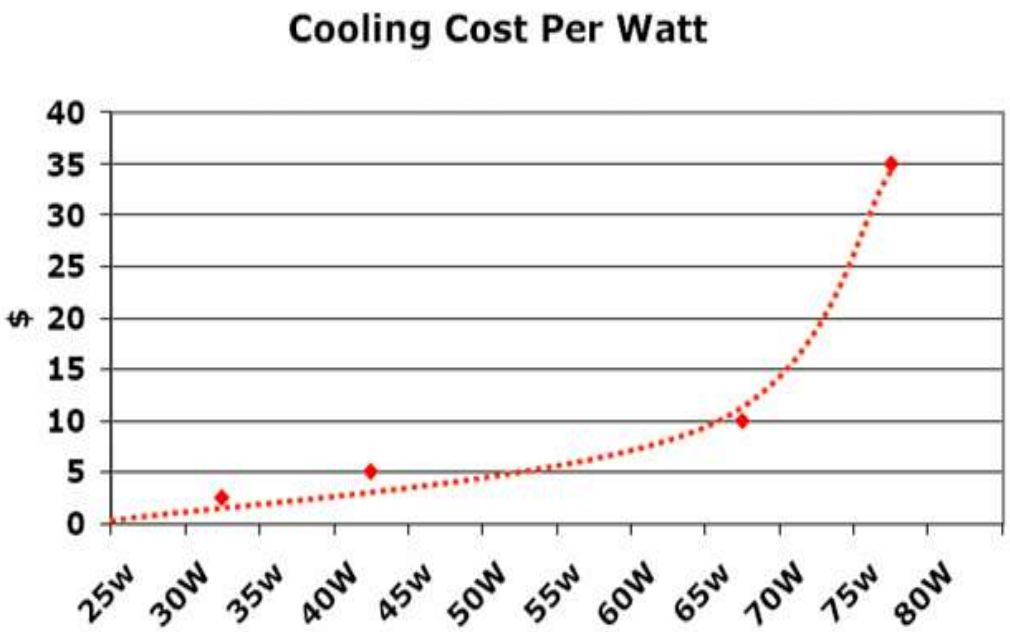

Source: Power Architect development Conference, Steve Carlson

50% Cost

50 percent of the cost of a data center is associated with expensive power and cooling equipment.

The hot and cold of power and cooling / HP June 2006

Power is a Key Competitive Advantage

"Earlier this month the giant chip maker [Intel] pushed up the date to June 26 in an attempt to take back momentum from rival AMD, which has gained market share over the past couple of years based on the strength of the performance-per-watt capabilities of its Opteron

Source: What Intel Gets to Prove Back with 'Woodcrest'?
BASELINE, June, 2006 by Jeffrey Burt

Intel Xeon 5160
CPU Dual-Core
3.0GHz, 4 MB L2
shared,

AMD Opteron 2216HE
CPU, Dual-Core, 2.4GHz,
2MB L2 Cache


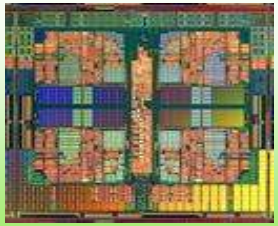

Source: AMD Processor/Memory Comparison of Intel and AMD product. Energy includes power input & cooling, Power Utility cost: $0.10/KW-hr, Publicly available processor & chipset specifications. The examples contained herein are intended for informational purposes only. Other factors will affect real-world power consumption and cost.

Power is a Key Competitive Advantage

Dual-Core Intel Xeon MP ('Paxville MP')

- 692 watts for processors **(173w MAX POWER)**
- 92.4 watts for chipset, 140.8W for memory

$1,297 per/year (1 server)

$648,380 per/year (500 servers)

Intel Xeon 5160 CPU Dual-Core 3.0GHz, 4 MB L2 shared,

Dual-Core AMD Opteron™ 8000 series

- 380 watts for processors **(95w MAX POWER)**
- 24 watts for chipset, 140.8W for memory

$764 per/year (1 server)

$381,796 per/year (500 servers)

AMD Opteron 2216HE CPU, Dual-Core, 2.4GHz, 2MB L2 Cache


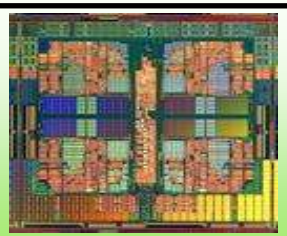
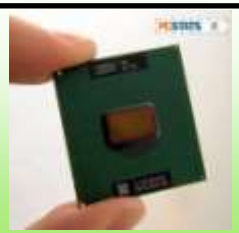


Source: AMD Processor/Memory Comparison of Intel and AMD product. Energy includes power input & cooling, Power Utility cost: $0.10/KW-hr, Publicly available processor & chipset specifications. The examples contained herein are intended for informational purposes only. Other factors will affect real-world power consumption and cost.

Delivered as Electronic Design Webcast February 13, 2008

Power Impacts Business!

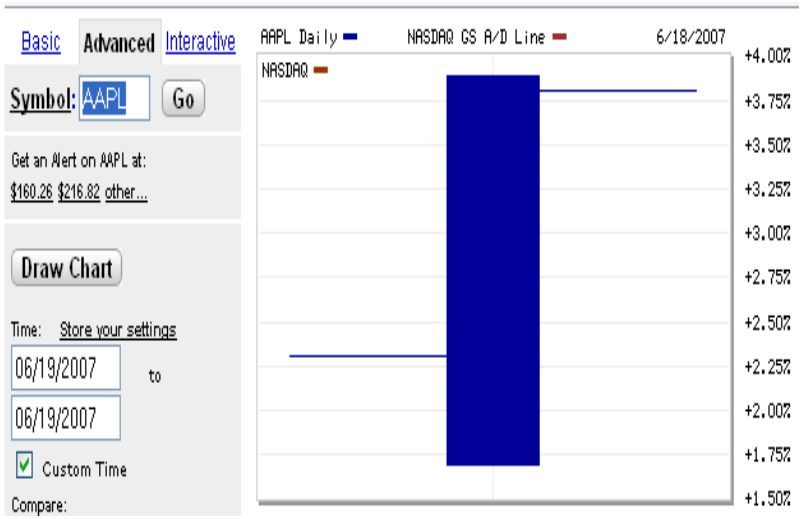

iPhone Battery Improved; Apple Stock Up

There were announcements this morning that Apple had addressed two issues with the iPhone - the non-user-replaceable battery had its battery life runtimes increased and the covering for the touchscreen has been changed from a plastic material to glass. Apple's stock was up almost 4% on the day due to that news.

Apple Inc. (F

125.09

+4.59 (3.81%)

Jun 18 - Close

= ~$3B

In one day

How to Reduce Power?

- Physical Optimization Methods
  - Limited Optimization Margins

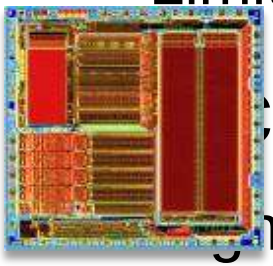

PHYSICAL

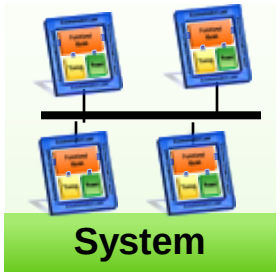

FUNCTIONAL

	Static Power	Dynamic Power
Area Optimizations	10%	10%
Clock Gating		20%
Multiple Supply Voltages (MSV)	20%	20%
Shut Down Power Modes	~50%	0%
A/DVFS (Adaptive/Dynamic Voltage Frequency scaling)	~10%	>40%
SW [User Experience]		>60%

Source: Institute of Electronic Design Automation, Technical University of Munich
Nikkei Electronics Asia -- September 2005
Robert Aitken, George Kuo, and Ed Wan EE Times

Source: Department of Electrical Engineering and Computer Science, Northwestern University, Evanston, Illinois, USA