Adaptive Computing

(… via Timing Error Avoidance)

Gus Uht
(w/ Rick Vaccaro)
Dept. of Electrical and Computer Engineering

Copyright © 2003-5, A. K. Uht, R. J. Vaccaro
Patent applied for.

BARC: January 21, 2005

Background

- Three typical goals (alone or combined):
  1. ~Maximize performance.
  2. ~Minimize power consumption.
  3. **Adapt** to current (or prior) conditions:
     a. Environmental, e.g., temperature. (current)
     b. Operational, e.g., power supply voltage. (current)
     c. Manufacturing, e.g., slower or faster chips. (prior)
- ~All adaptive methods applicable to all.
- "Better-than-worst-case design"

Outline

1. Adaptive Methods
2. Timing Error Toleration
3. Timing Error Avoidance
4. TEAtime – Prototype and Data
5. Adaptive Control Systems
6. TEAPC & TEA42
7. Demo – TEA42
8. Summary

Adaptive System Methods

- Timing Error *Toleration:*
  - Keep changing something (freq., volt.), then:
  - ***Let error occur, then recover.***
  - Change the thing back, repeat.

- Timing Error *Avoidance:*
  - Keep changing something (freq., volt.), then:
  - ***Stop just before error would occur.***
  - Change the thing back, repeat.

Timing Error Toleration

- Operation:
  1. *Speed up clock (resp. reduce voltage) until error occurs.*
  2. *Slow down clock (resp. increase voltage) for no error.*
  3. *Backtrack / repair error.*
  4. *Repeat: GOTO 1.*
- Examples:
  - Uht 2000: *TIMERRTOL*: performance ~maximized; not pursued.
  - Austin *et al* 2003: *Razor:* power ~minimized; chip fabbed.
- Plusses:
  - True (or better than) perf. max., resp. power min. [**within** cycle]
- Minuses:
  - Can be costly (TIMERRTOL: 2x cost; Razor: little extra cost)
  - Is complex and hard to design.
  - Can lead to **increase** in cycles → perf. may NOT be true max.

Timing Error *Avoidance*

- Operation:
  1. *Speed up clock (resp. reduce voltage) 'til **just before** error occurs.*
  2. *Slow down clock (resp. increase voltage) for no error.*
  3. *Repeat: GOTO 1. [No backtracking or repair needed.]*
- Examples:
  - Olivieri, et al, 1999 – used microcontroller.
  - Uht 2003: *TEAtime*: ~performance maximization; prototype built.
    - *One-bit wide slowest-path test logic always errs before real logic does.*
- Plusses:
  - Close to true performance maximization, resp. power minimization.
  - Simple, easy to design and build. (TEAtime)
  - No increase in cycles.
- Minuses:
  - Shorter papers. *(no, wait, that's a plus…)*

TEAtime Prototype

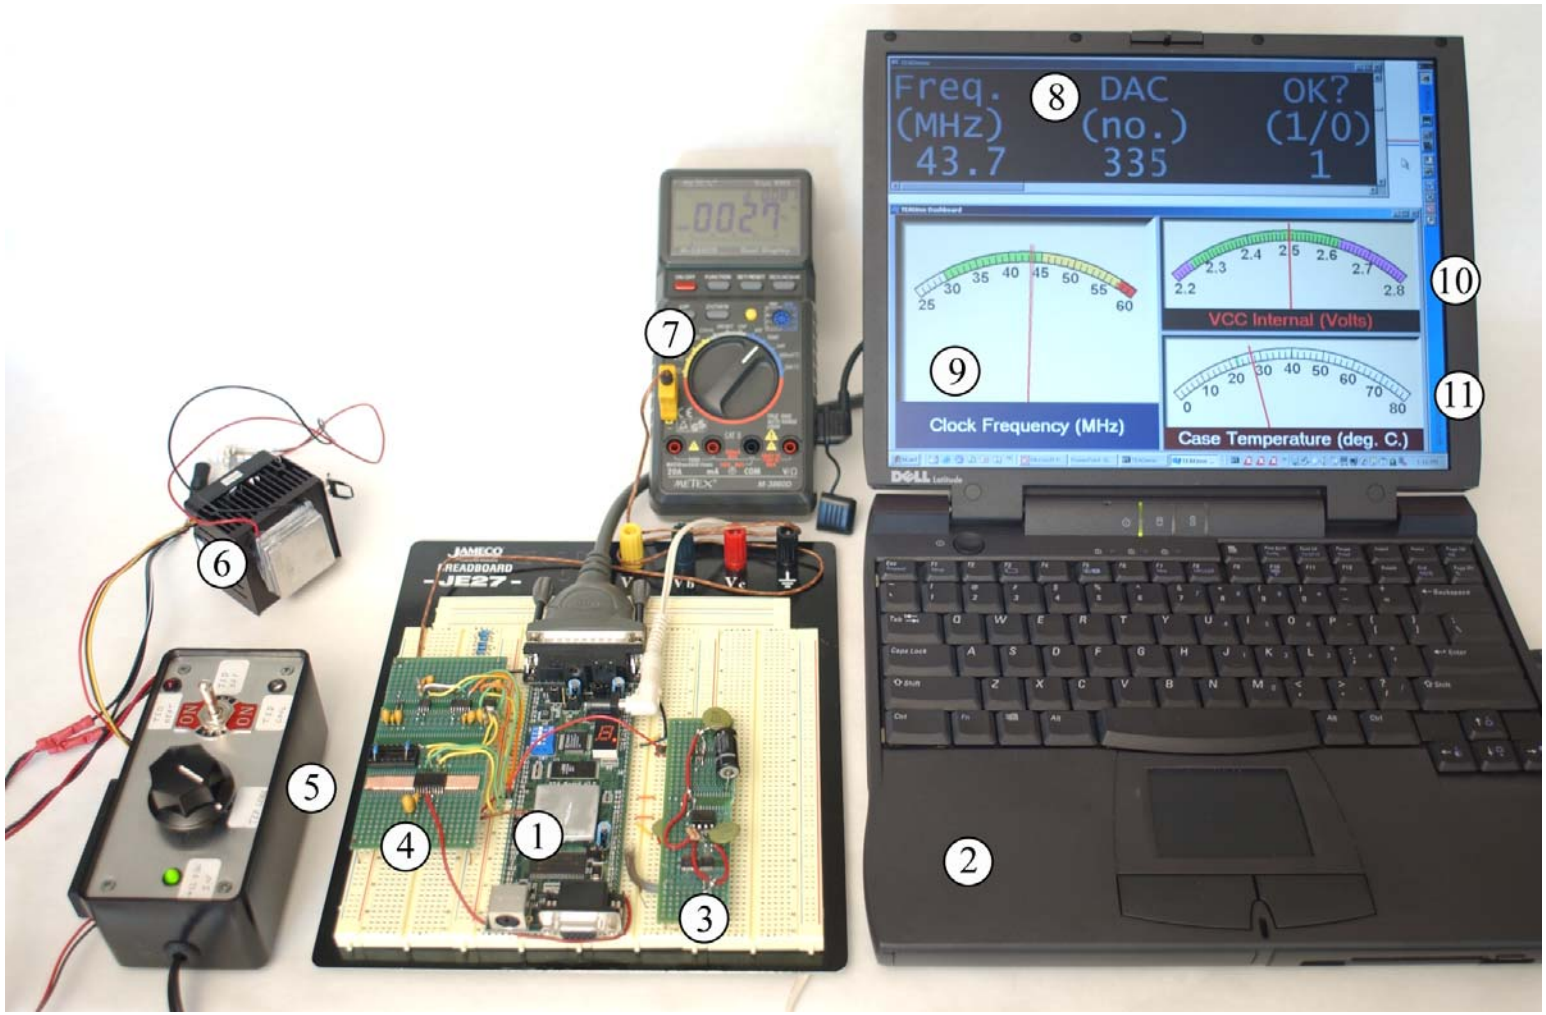

Frequency & DAC setting vs. time (Temp. = 22° C.)

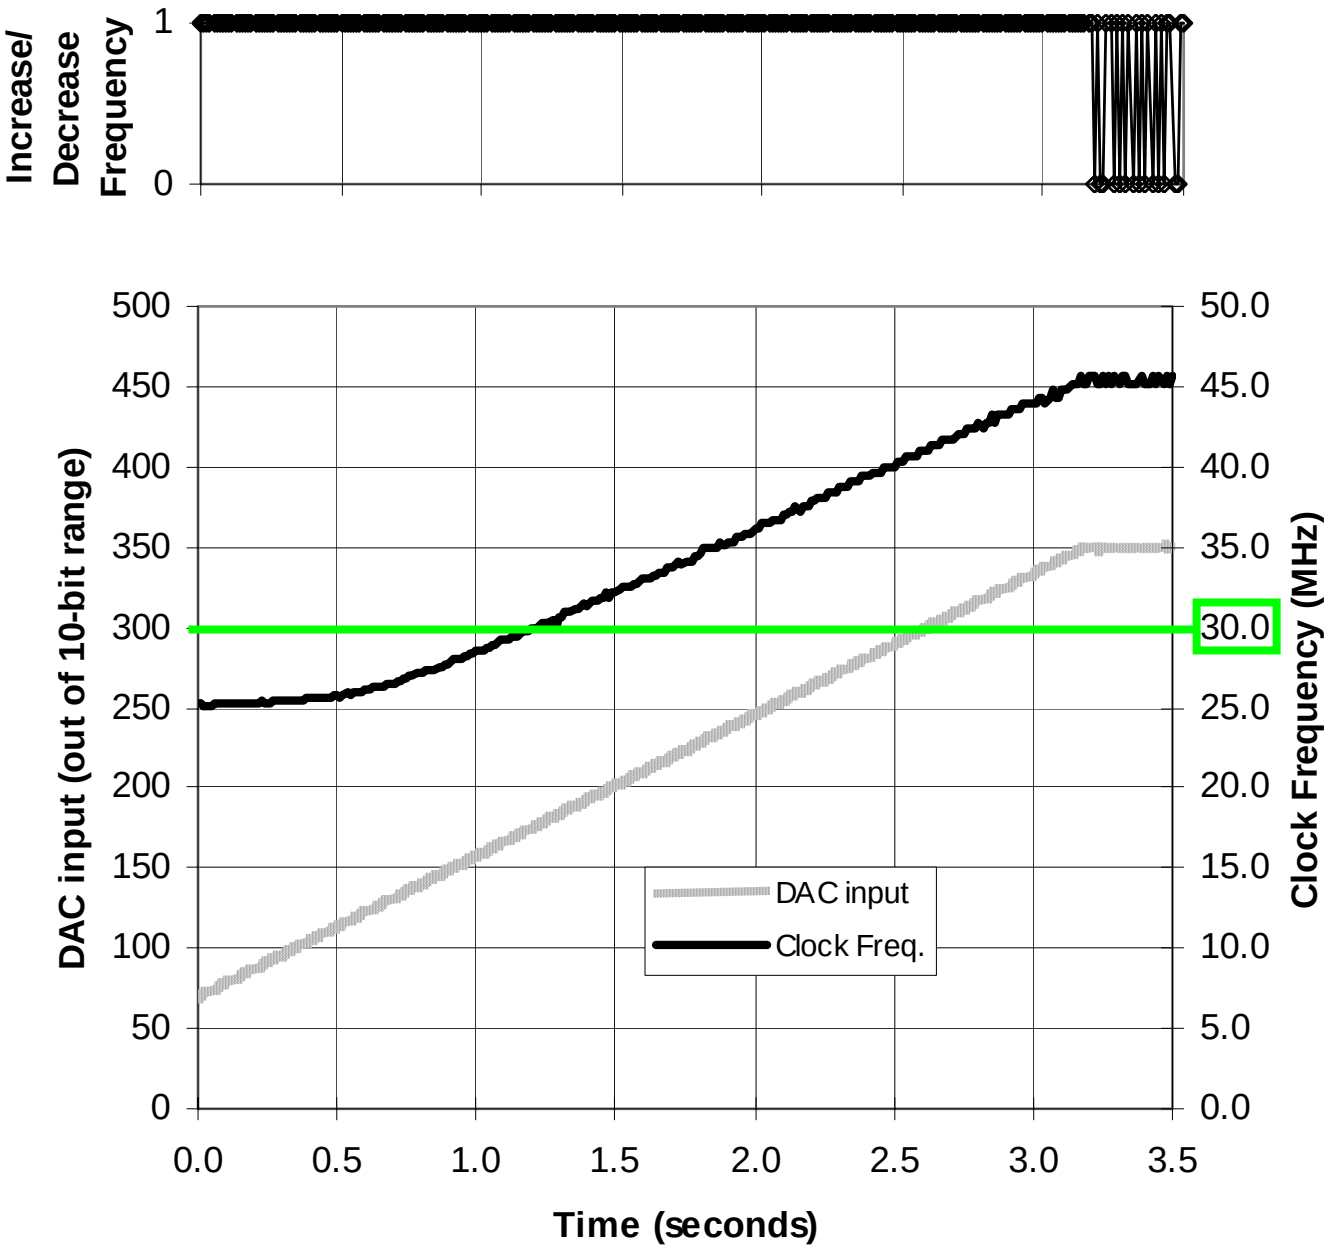

TEAtime Summary

- Performance of prototype improves by >= 34%
- TEAtime provides adaptable frequency control of any synchronous digital system
- Always ~maximizes performance
- Very cheap
- Very easy to add to existing or future designs
- ~Can be adapted to physically existing systems
- ***It Works!!***

Adaptive Control Systems

- TEAtime:
  - Simple up/down control system – 0 delay.
- Skadron, Bahar:
  - Classic feedback control system.
  - In hardware.
- TEAPC, TEA42:
  - Big delays; used as opportunity, not problem.
  - State-space feedback control system.
  - In software; little overhead.

TEAPC, TEA42 Goals

- Motivating Goals:
  - Realize TEAtime characteristics in a real computer
    - Adaptive computing
    - Improved performance – "Better-than-Worst-Case"
- Additional Goals:
  1. ***Workload adaptation***
  2. ***Reduced power consumption***
  3. ***Improved reliability***
  4. ***Disaster tolerance (always enabled)***
  5. ***…and all in a real machine***
- BUT: can't redesign or build Pentium 4's
- SO: use real IBM/Intel-standard PC

Current Control System

- Only input is CPU temperature (feedback line).
- Primary output is CPU frequency (N).
  [sometimes: Vcore = *f* (N)]

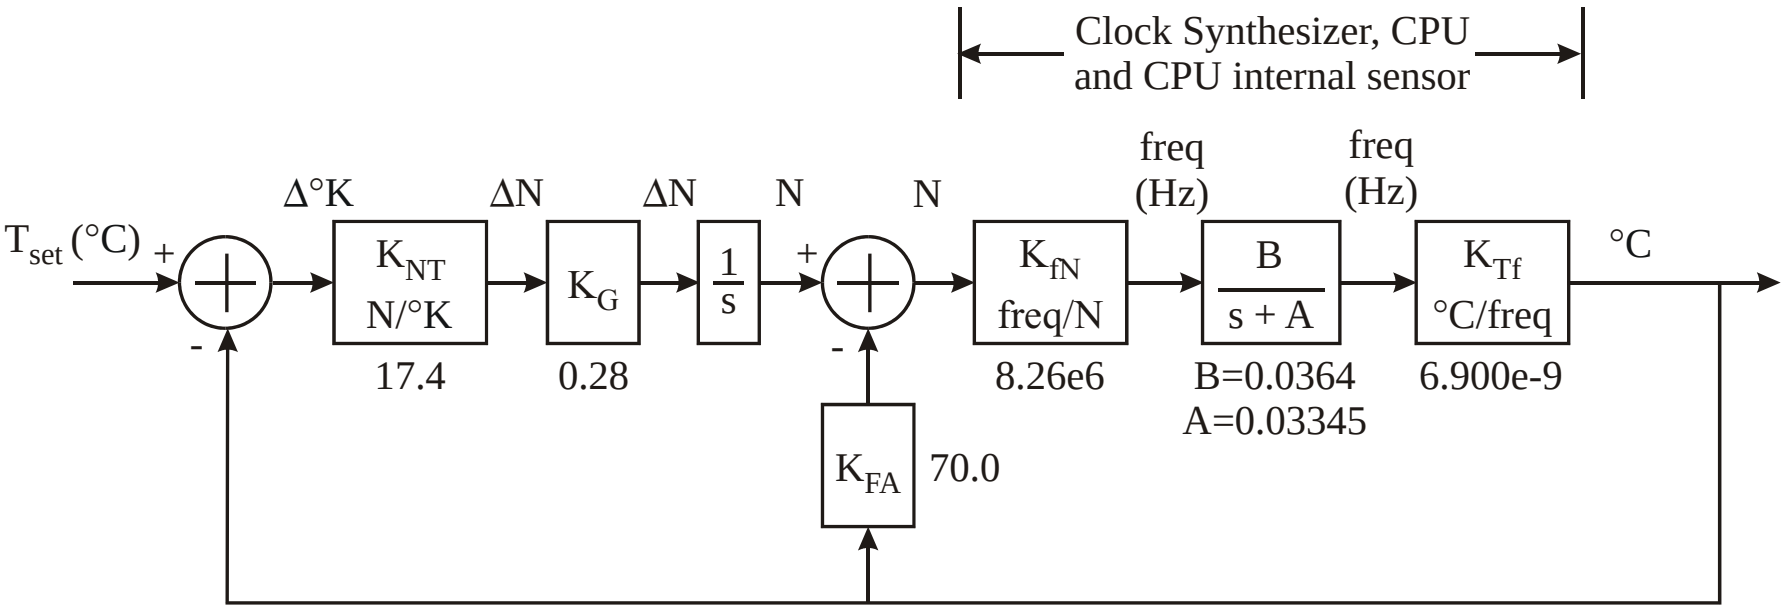

TEAxx System Features

- Hardware:
  - No modifications → all COTS parts.
  - System: ~3.0 GHz Pentium 4; 800/533 MHz FSB; Intel chipset; 1 GB RAM.
- Software (`teapcwin` program):
  - Realizes feedback control system.
  - Standard MS Windows application.
  - (Standard OS: W2K SP4; no modifications [of course].)
  - Small: 1.1 megabytes.
  - Fast: < 1% CPU utilization (without display).

freq., Vcore **un**linked < Load Adaptation > freq., Vcore linked

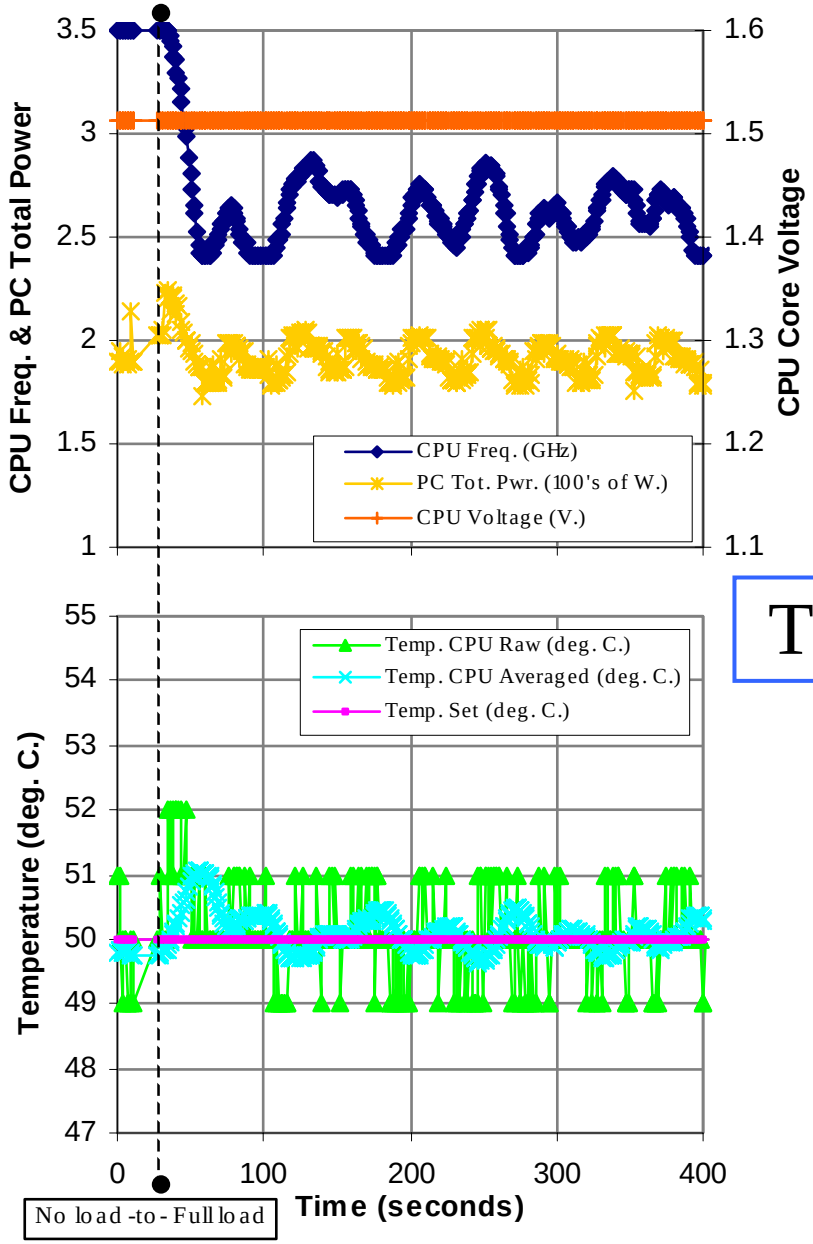

Freq. (GHz)

Vcore (V)

Power (100 W)

TEAPC

Tset

Traw

Tavg

(all deg. C)

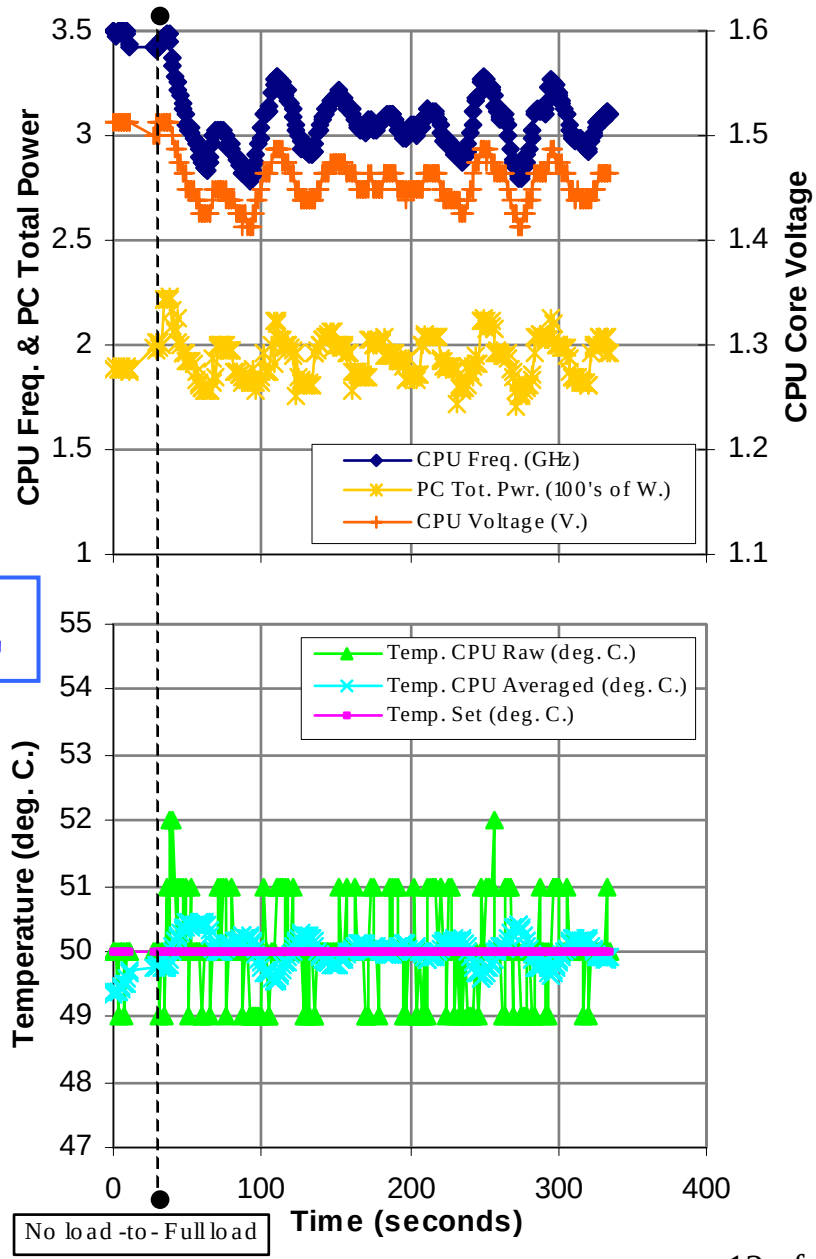

TEA42 Disaster Tolerance

- Example: CPU Fan dies….
- Changes (automatic, via feedback system):
  - Freq: 3.15 GHz → 1.5 GHz
  - Vcore: ~1.5 V. → ~1.1 V.

  → Power: ~150 W. → ~90 W. (40% savings)
- CPU temperature stabilizes at safe value (with this CPU).
- System still works.

…and now it's time for the:

DEMO

Overall Summary

- TEAtime, TEAPC & TEA42 realize:
  1. Better-than-worst-case performance.
  2. ***Adaptive*** operation to both environment and/or loading.
  3. Low-power, high-reliability operation.
  4. Disaster tolerance.
- Feedback-control great for a system, too.
- ***Adaptive systems*** are the way to go.
- ***They Work!***

Pointers…

- *Computer*, March, 2004 Special Issue.
- My website: www.ele.uri.edu/~uht
- Or µRI website: www.ele.uri.edu/muri

Adaptive Computing

(… via Timing Error Avoidance)

Gus Uht
(w/ Rick Vaccaro)
Dept. of Electrical and Computer Engineering

Copyright © 2003-5, A. K. Uht, R. J. Vaccaro
Patent applied for.

BARC: January 21, 2005

Appendix

21. TEAtime Block Diagram.
22. TEAPC Block Diagram.
23. TEAPC Components.
24. [TEAPC] Experiment Setup.

TEAtime Block Diagram

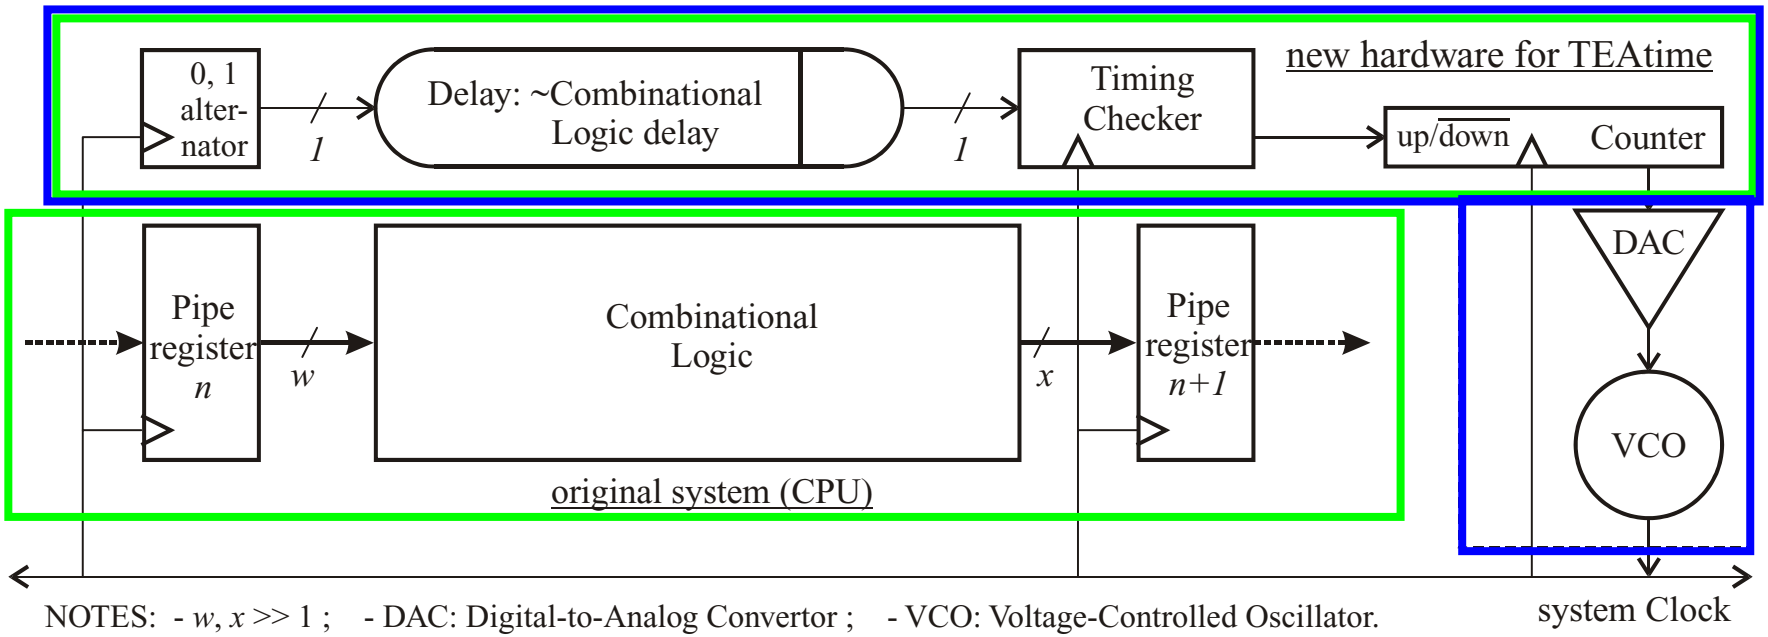

- Timing Error Avoidance system
  - Blue: TEAtime hardware
  - Green: on FPGA

TEAPC Block Diagram

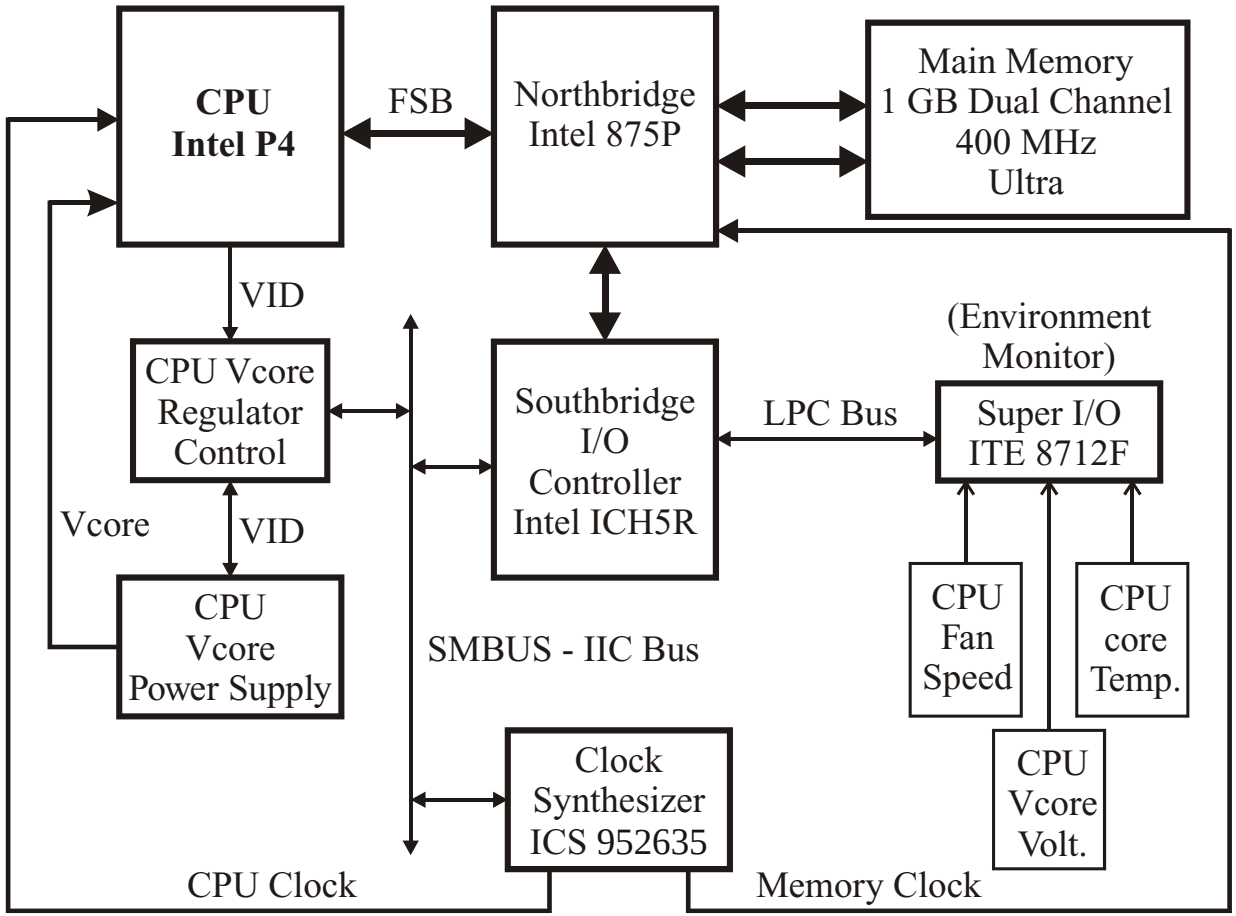

(FSB - Front Side Bus) Only directly relevant components and connections are shown. (LPC - Low Pin Count)

TEAPC Components

PC Component	Manufacturer	Part Number/Description
Motherboard	Gigabyte	GA-8KNXP (Rev. 2); w/DPS regulator
CPU	Intel	P4 3.0 GHz, 800 MHz bus
Chipset	Intel	875P, ICH5R
Clock Synthesizer	ICS	ICS952635
Super I/O (Environment Mon.)	ITE	IT8712F V0.6
CPU Volt. Regulator Control	ITE	IT8206R V0.1
Main Memory	Ultra	U10-5903R; 2 x 512 MB; 400 MHz DDR, Dual Channel (Operated at 320 MHz.)
Operating System	Microsoft	Windows 2000 SP4, HT disabled
Disk System – RAID 0+1	ITE	GigaRAID IT8212F
Disks	Maxtor	4 x 6E040L0, 40 GB, 133MHz IDE
Equipment for experiments only		
Fan Controller & Temp. Mon.	Thermaltake	Hardcano 12; for 4 fans, 4 thermocouples
Power Meter	Electronic Educational Devices	watts up? PRO (Note: this is the unit's model name.)
CPU Fan Controller	custom	On/Off, control sel. (MOBO or Hardcano)

Microarchitecture Research Institute
University of Rhode Island

Experiment Setup