EDF-DVS Scheduling on the IBM Embedded PowerPC 405 LP

Aravindh Anantaraman, Ali El-Haj Mahmoud,
Ravi Venkatesan, Yifan Zhu, Frank Mueller

North Carolina State University,
Center for Embedded Systems Research (CESR)
Departments of Computer Science / Electrical & Computer Engineering

Why DVS?

- Energy management
  - vital design constraint especially in battery-constrained embedded systems
- Dynamic Voltage Scaling (DVS)
  - processor core voltage ↑ or ↓ depending on computation demands
  - V ↓ ⟹ f ↓
  - $P \propto V^2 \times f$
- DVS in real time systems
  - calculate safe operation frequency
    - all tasks are guaranteed to meet its deadlines

Why PowerPC 405LP?

- Hardware support for DVS
- Software support to change voltage
  - via. user-defined operating points
    - for e.g., 266 MHz @ 1.8 V to 33 MHz @ 1 V
    - easy to define
  - can define any “stable” frequency/voltage pair
- ability to execute instructions even when frequency/voltage is changed
  - contemporary processors with DVS support enter sleep mode during frequency/voltage transitions

How?

- Real-time earliest deadline first (EDF) scheduling policy as part of user-level threads package under Linux
- 4 hard real-time software DVS techniques
  - static
  - cycle-conserving
  - look-ahead
  - feedback
- Oscilloscope and analog data acquisition board to measure voltage and current and calculate power
- Custom changes to the development board
  - separate measurements of voltage, current of processor, memory and I/O components

PowerPC 405LP

- Diskless Monta Vista Embedded Linux variant
- 50% reduced capacitance
  - DVS switching faster (≤ 200 μs)
- DVS switching
  - synchronous (blocking)
    - conventional DVS-capable processors
  - asynchronous (non-blocking)
    - execution may proceed during switch

PowerPC 405LP(cont.)

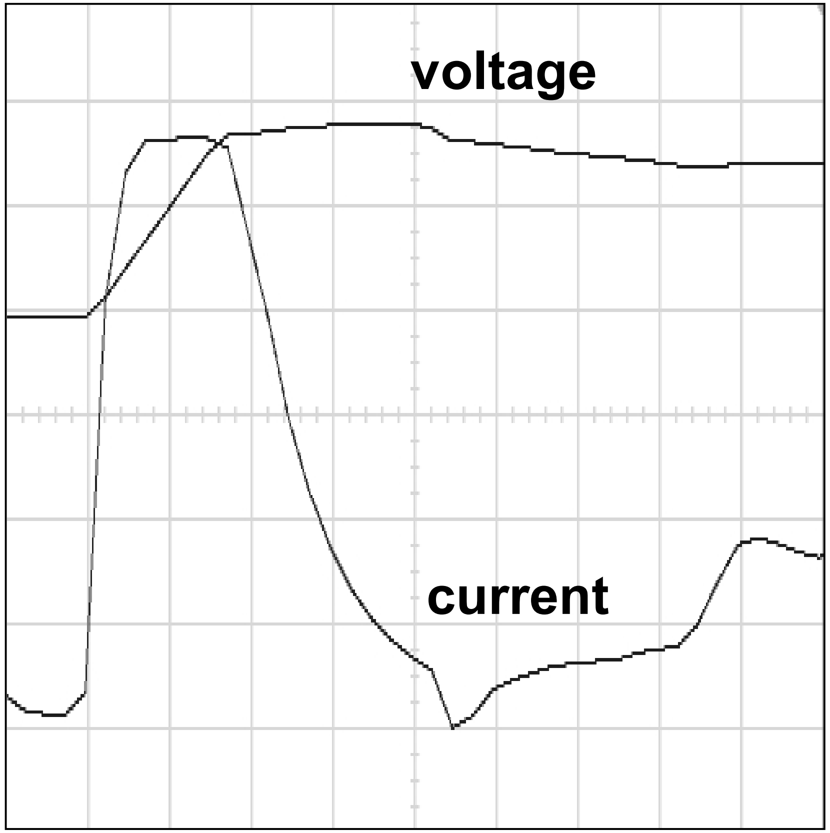

- System call for asynchronous switch
  - voltage ramped up to max
  - time to reach new voltage level is estimated
- High-resolution timer interrupt triggered
  - power management unit reprogrammed
  - new processor frequency activated
  - voltage settles in a controlled manner to new operating point

PowerPC 405LP(cont.)

- Dynamic power management (DPM) facility
  - defines stable operating points (frequency/voltage pairs)
- Operating states
  - system states (idle, task activity, sleep)
- DPM Policy
  - defines set of operating points for each operating state
  - synchronous switching
- DPM task state
  - allows application to select operating points
  - asynchronous switching

Implementation

- EDF Scheduler on 405LP under Linux
- Kernel-level threads
  - complex kernel modifications – error-prone
  - more control over execution
- User-level threads
  - simplicity of design
  - develop generic and portable threads library
  - facilitate easier debugging
  - less control over execution and resources
  - however, OS background activity minimal

[redacted]ime Scheduler Integration

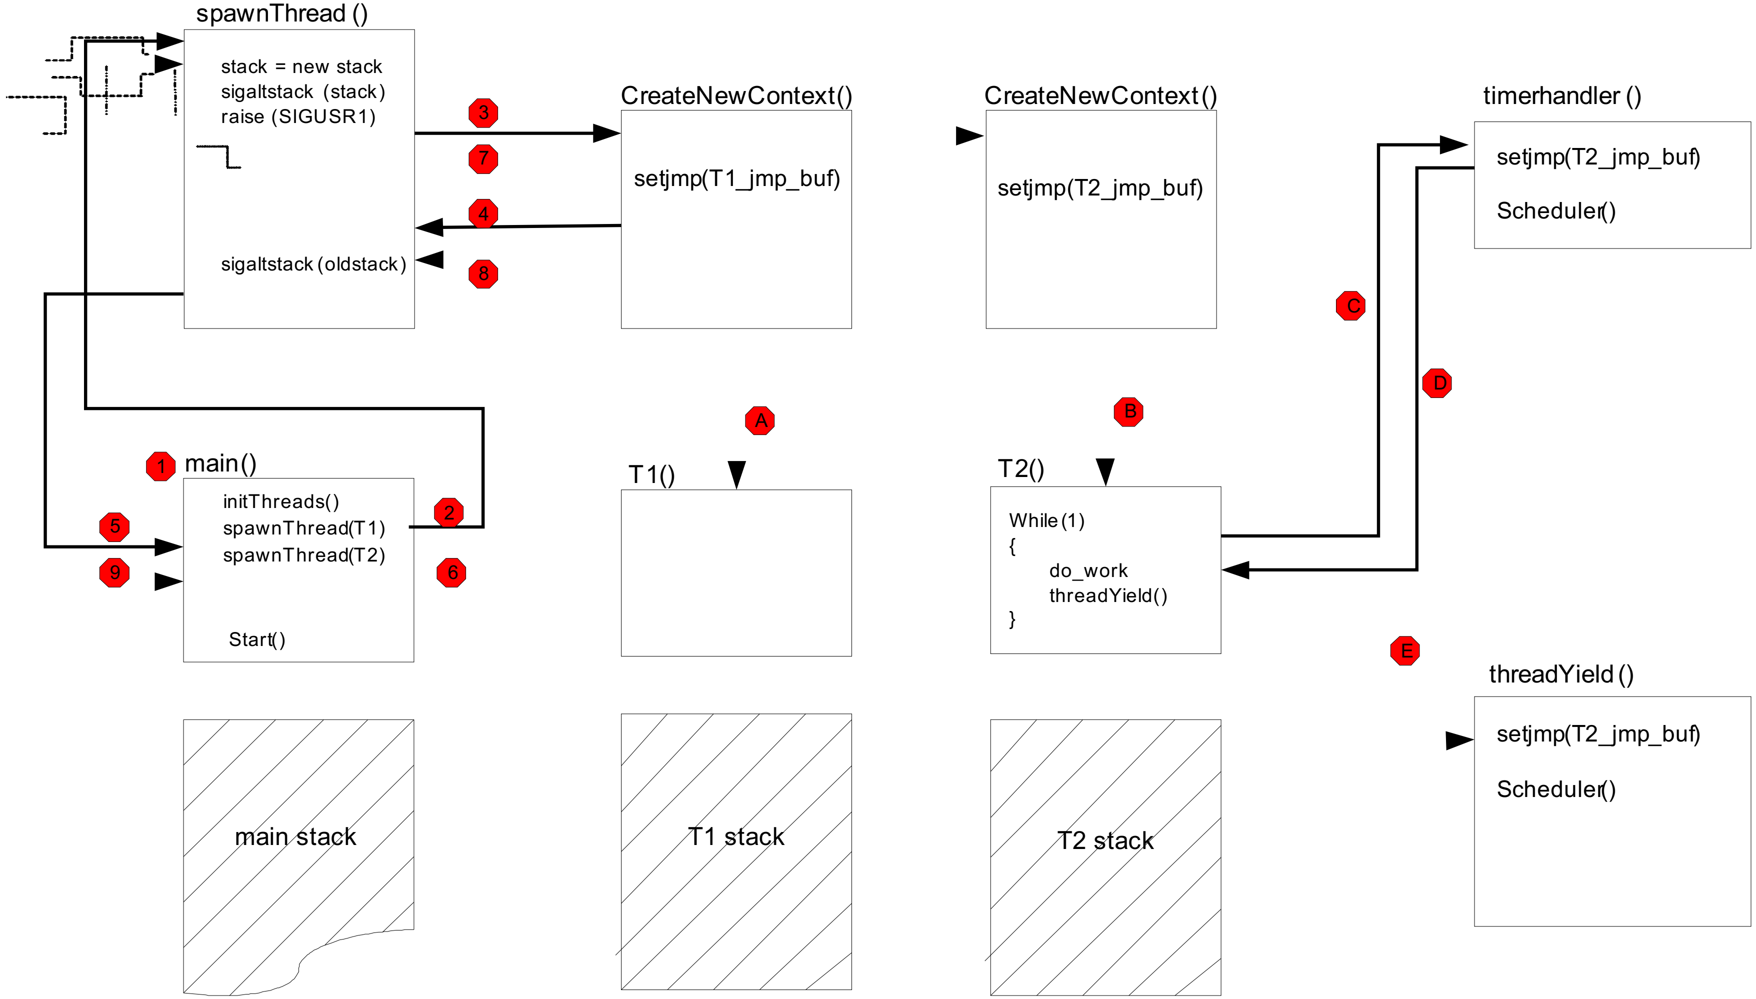

Static RT-DVS EDF

Task	Exec time	Period
1	3	8
2	4	20
3	3	10

$$U = 3/8+4/20+3/10 = 0.875$$

$$f = 0.875 \times f_{max}$$

Cycle-Conserving RT-DVS EDF

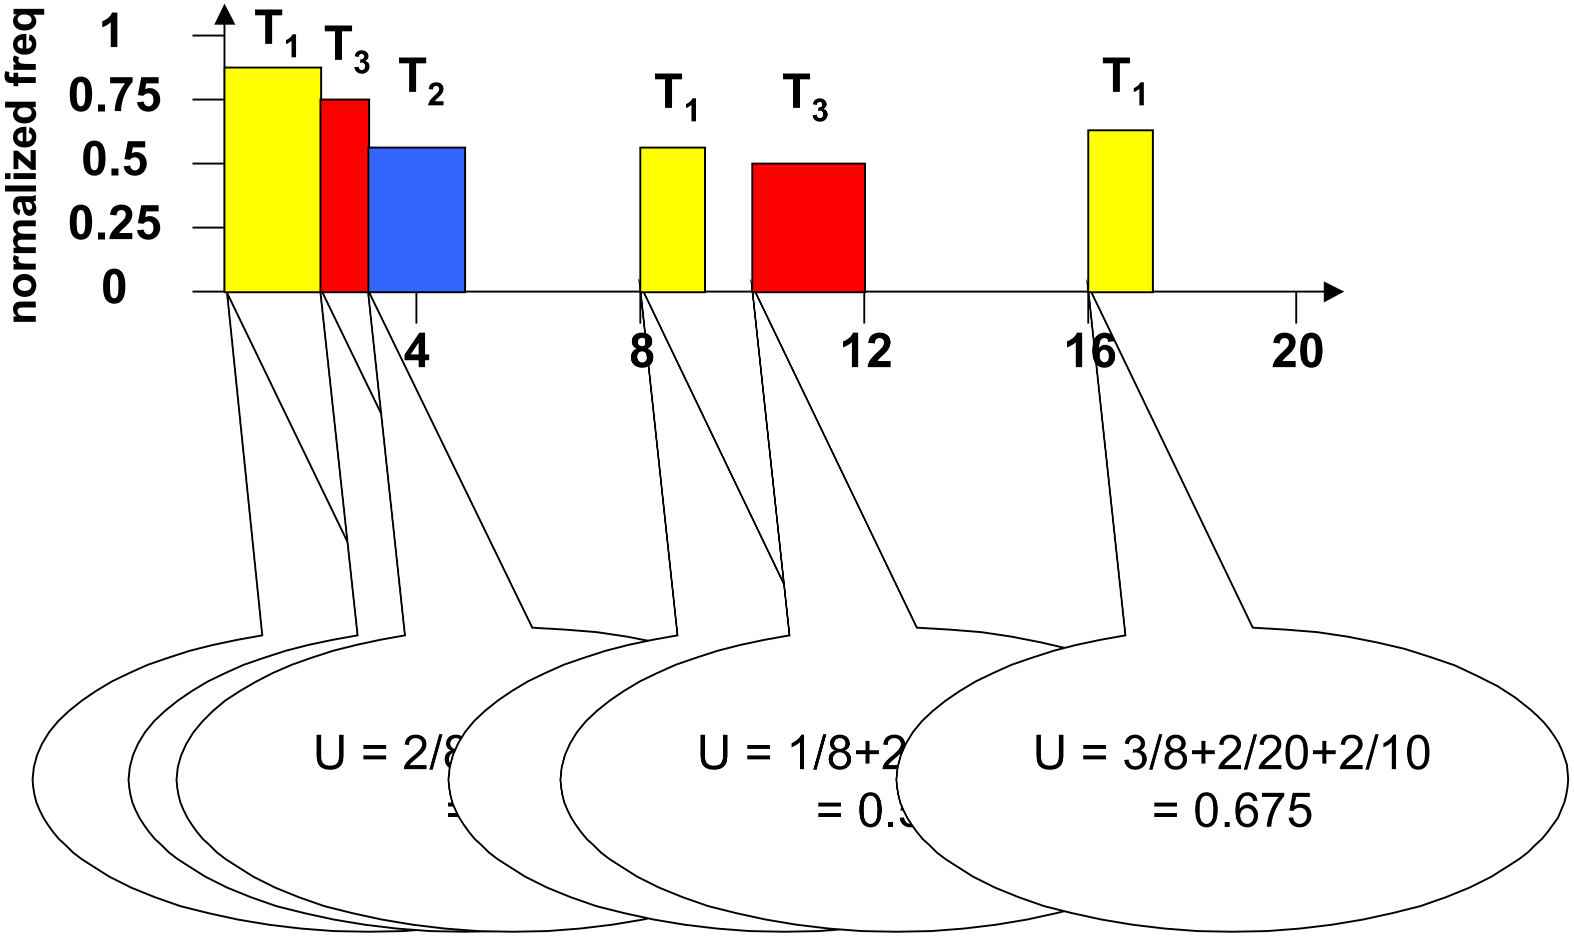

Look-Ahead RT-DVS EDF

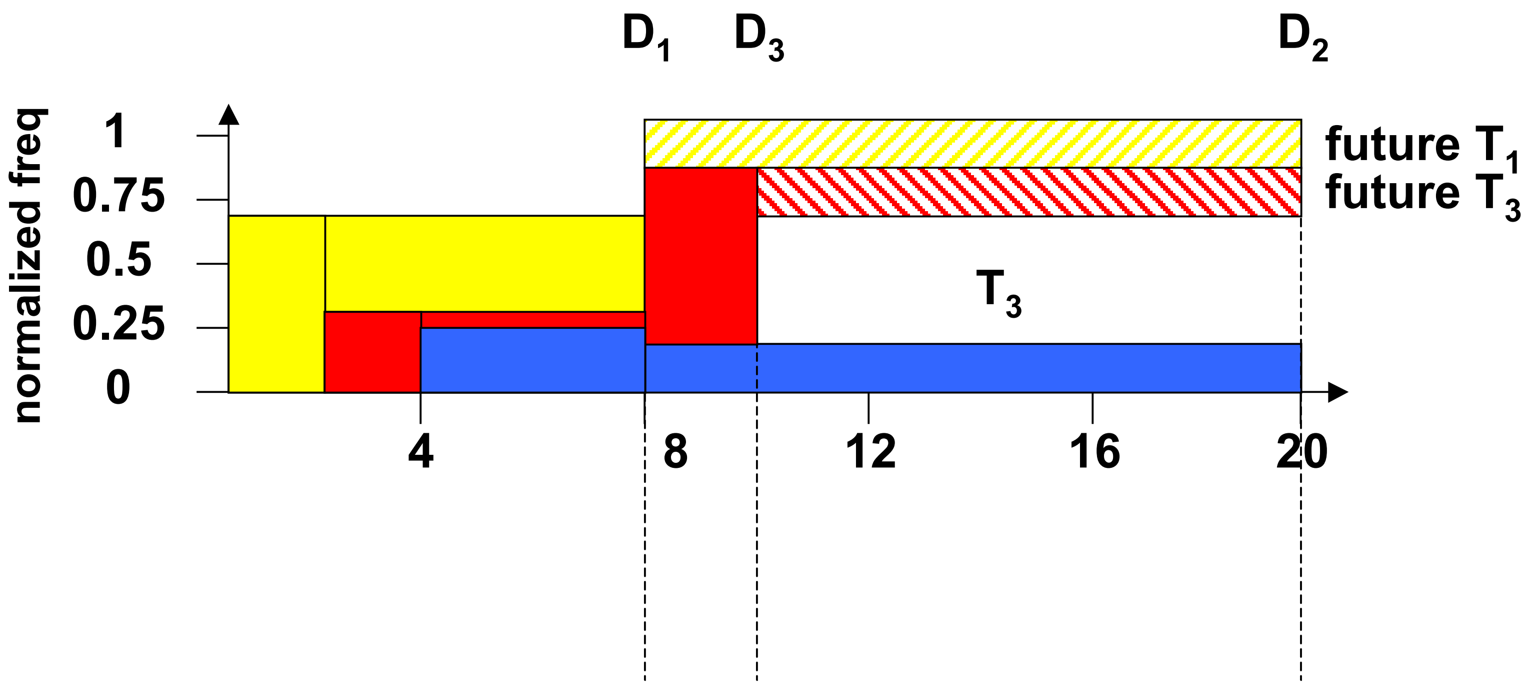

Feedback RT-DVS EDF

- Greedy scheme:
  - assign all idle time/slack to running task
- Assuming all other tasks at the maximum frequency (speed)
- Capitalize on early completion of current task
  - early completion → more slack for other tasks
  - repeat scaling on next task

Worst-Case Timing Analysis

```
t1 = gettimeofday();
for ( i = 0; I < N; i++);
t2 = gettimeofday();
work();
t3 = gettimeofday();
for ( i = 0; I < N; i++)
     work();
t4 = gettimeofday();
```

- No tool support to derive WCETs of tasks on PPC405LP
- Timing loop computes average execution of N instances of the task taking into account loop overhead and cold cache misses

$$t_{work} = \frac{(t4 - t3) - (t2 - t1)}{N}$$

Experimentation Methodology

- ➢ Need to measure voltage/current at a high rate
- ✗ Multimeters – coarse-grained and low precision
- ✓ High frequency analog data acquisition board
  - current measured as voltage level over a resistor with 1V drop per 360 mA
- ➢ Oscilloscope
  - visualizing voltages and currents
  - high precision

Valid Frequency/Voltage Pairs

Setting	0	1	2	3	4
Frequency (MHz)	33	44	66	133	266
Voltage (volts)	1.0	1.0	1.1	1.3	1.7

DVS Overhead

Activity	Sync. DVS	Async. DVS	Signal handling
Overhead	117-162 μs	8-20 μs	0.07-0.6 μs

Task Sets

Task	Task Set 1 (msec)		Task Set 2 (msec)		Task Set 3 (msec)	
	Period (P_i)	WCET (C_i)	Period (P_i)	WCET (C_i)	Period (P_i)	WCET (C_i)
1	2400	400	600	80	90	12
2	2400	600	320	120	48	18
3	1200	200	400	40	60	6

Energy Savings of DVS Techniques

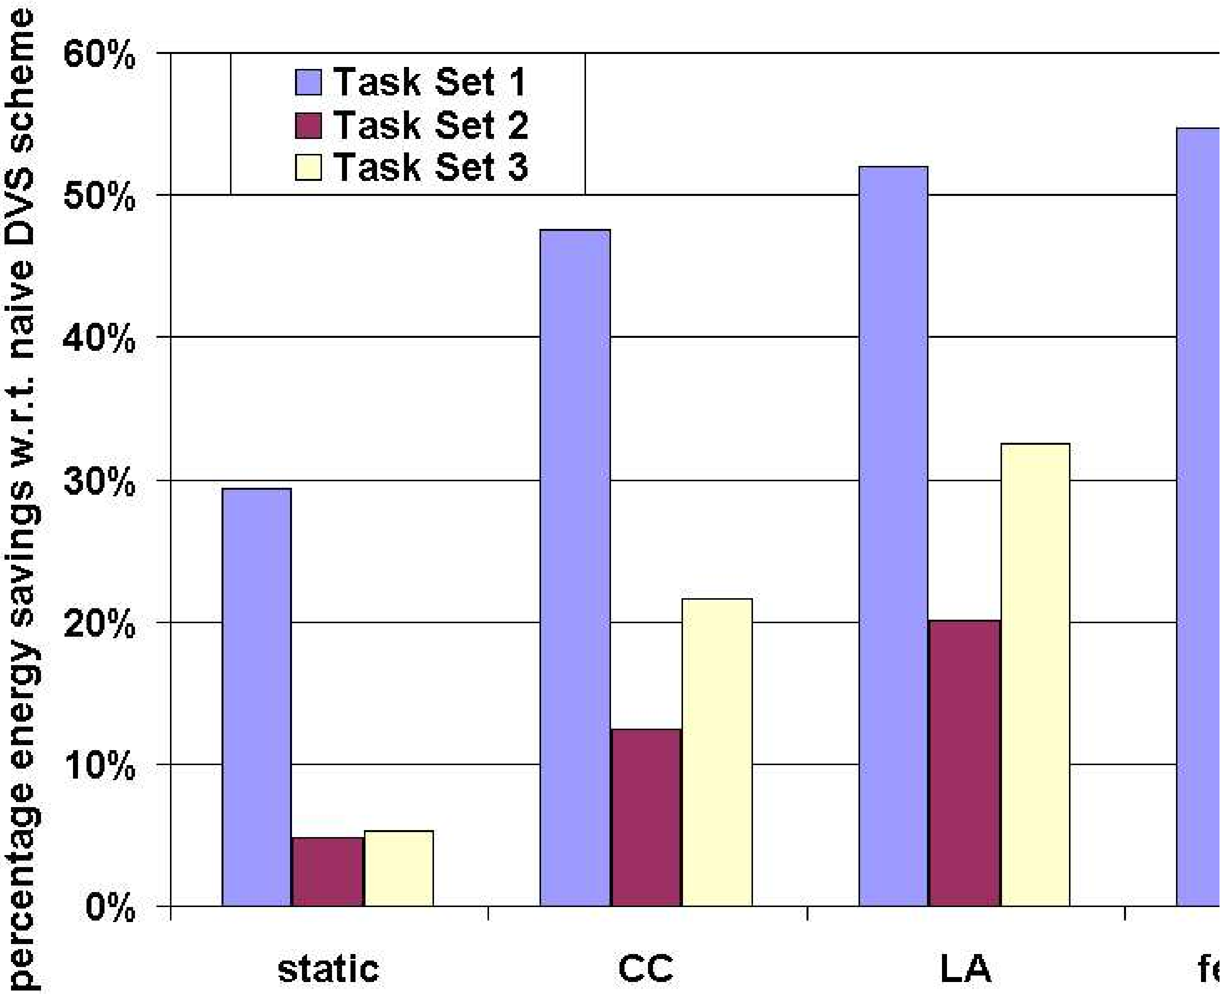

Energy Savings of fast DVS modulation

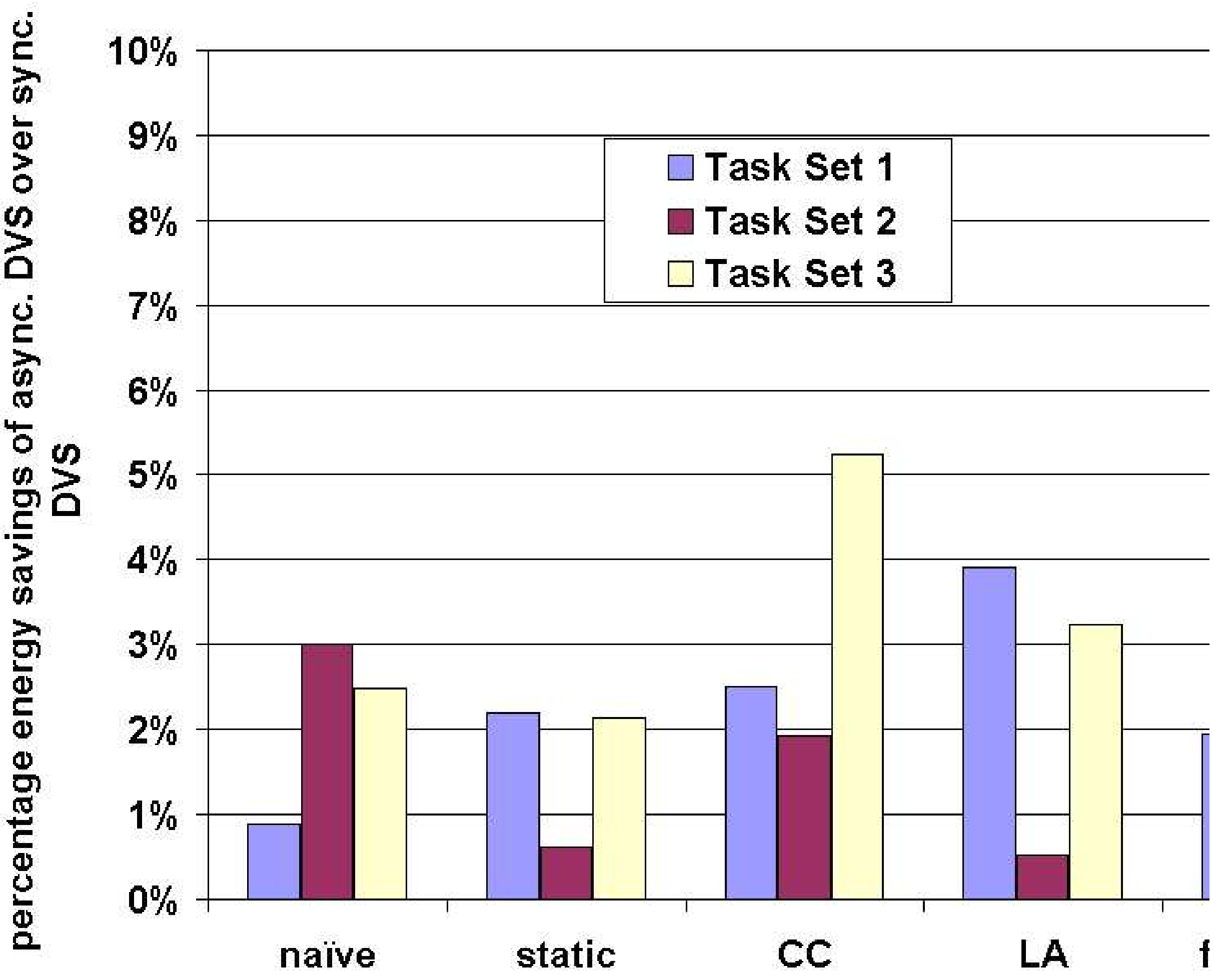

Oscilloscope – Loose Task Set

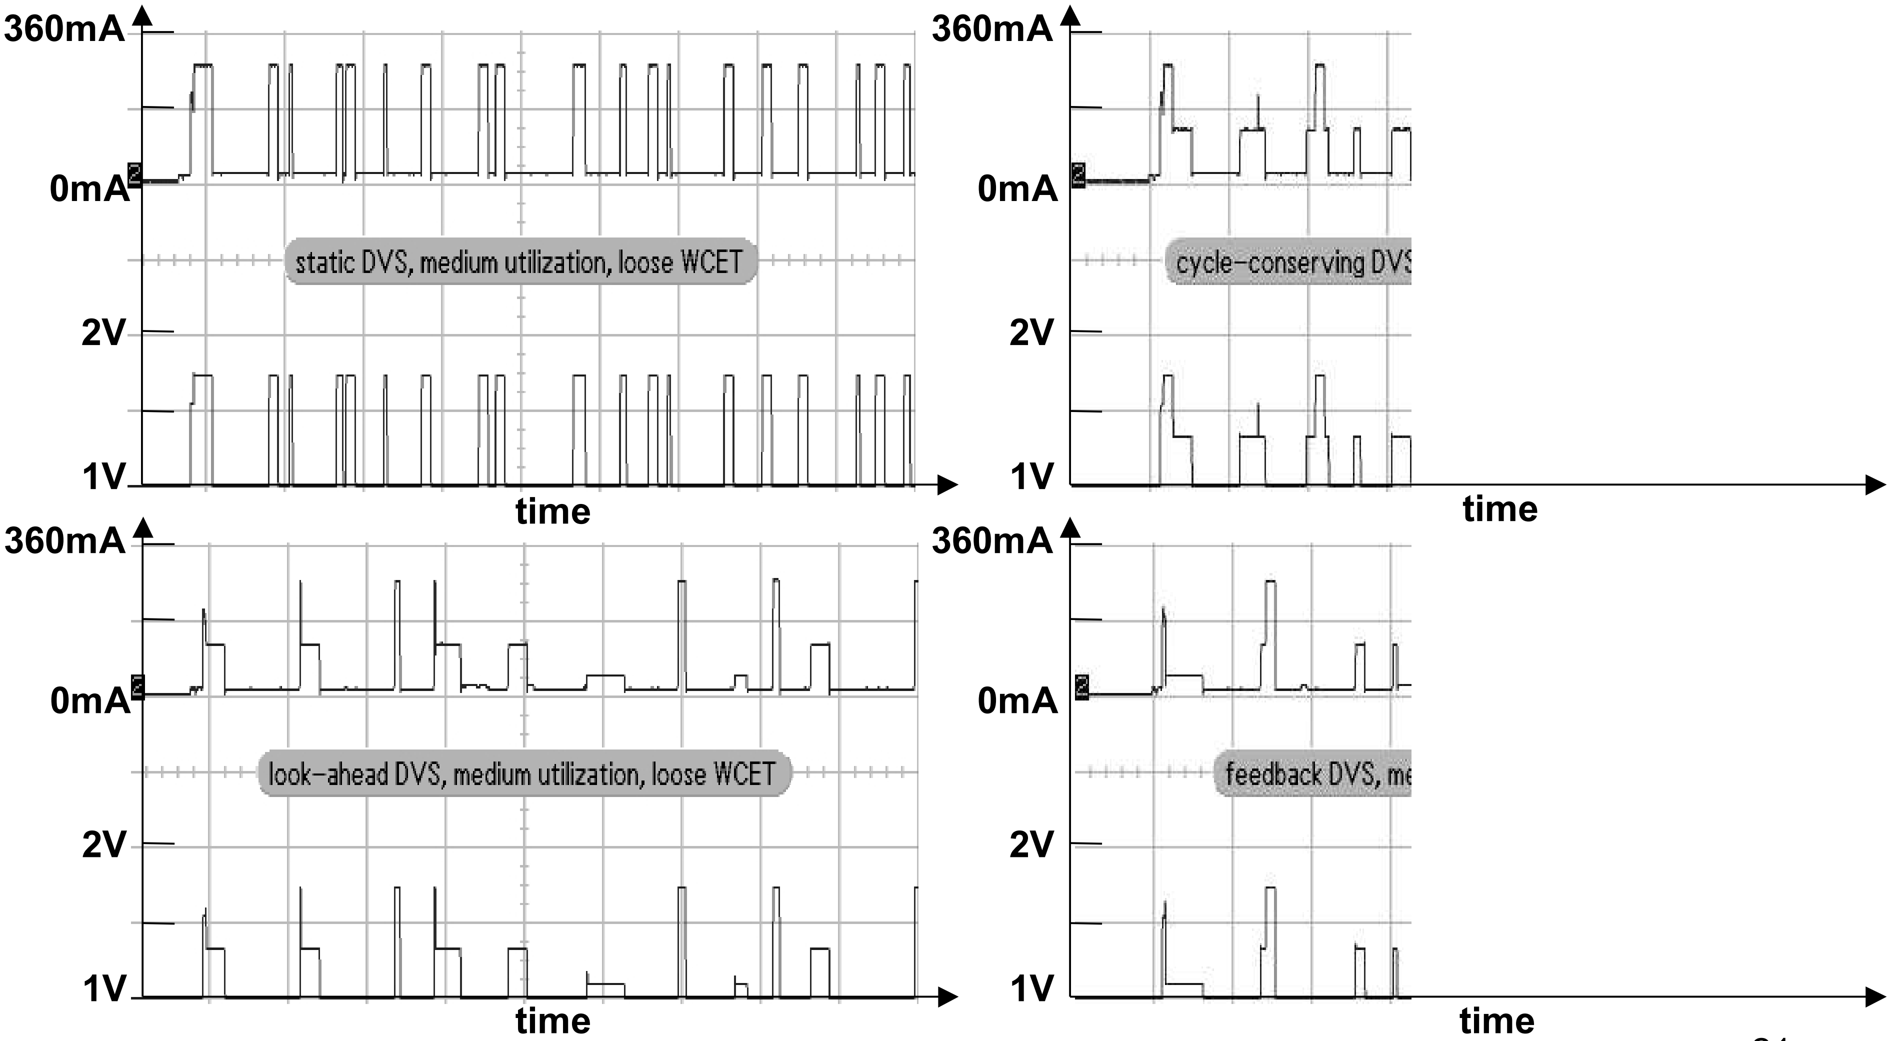

Oscilloscope – Tight Task Set

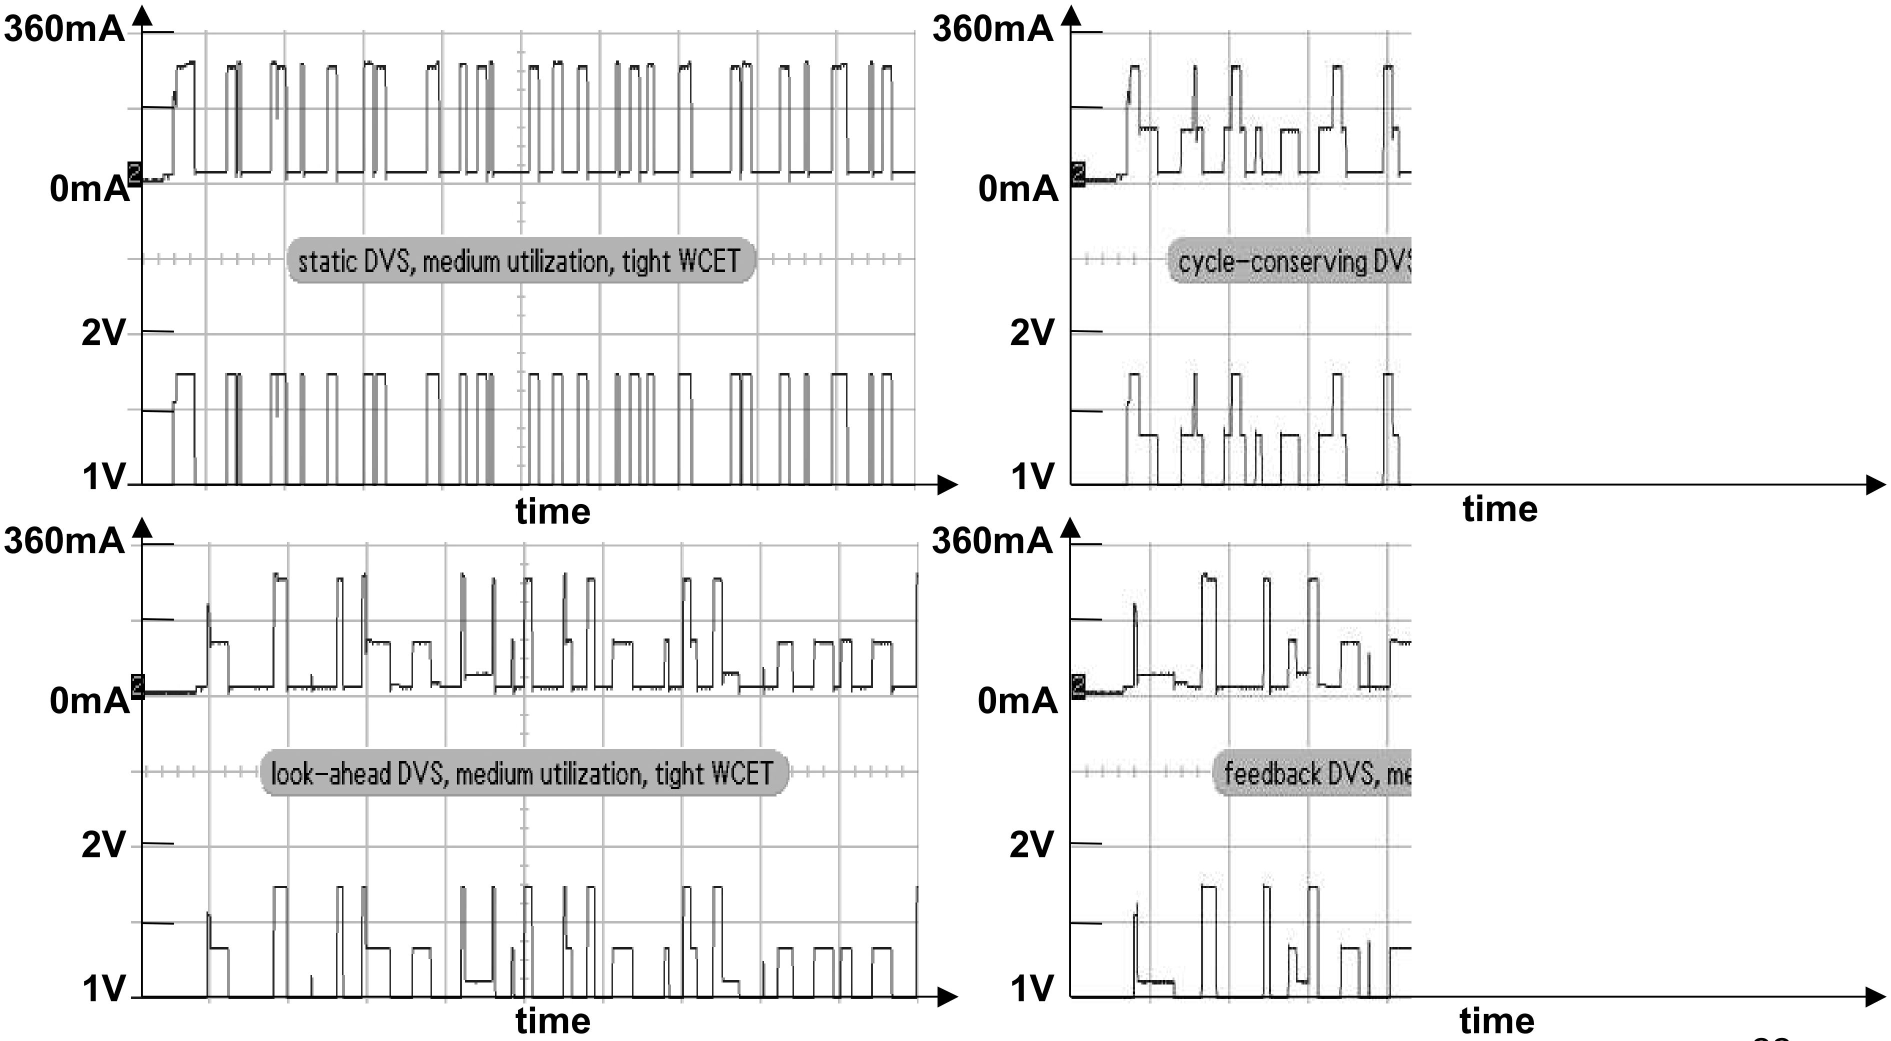

Conclusions

- Successfully created an infrastructure for investigating EDF-DVS schemes on IBM PowerPC 405LP
- Benefits in energy reduction of up to 5% for fast DVS modulation
  - continue execution while switching voltage/frequency
- Aggressive real-time DVS scheduling algorithms can achieve up to 54% reduction in energy consumption over naive DVS scheme
  - energy over hyperperiod of real-time tasks
  - modulate both processor voltage and frequency