#### Software-level Power-Aware Computing

Lecture 2

## Lecture Organizations

#### Lecture 1:

- Introduction to Low-power systems
- · Low-power binary encoding
- Power-aware compiler techniques
- Lectures 2 & 3
  - Dynamic voltage scaling (DVS) techniques
    - OS-level DVS: Inter-Task DVS
    - Compiler-level DVS: Intra-Task DVS
    - Application-level DVS
  - Dynamic power management
- Lecture 4
  - Software power estimation & optimization
  - Low-power techniques for multiprocessor systems

2

Leakage reduction techniques

Low Power SW.2

J. Kim/SNU

#### Voltage, Frequency & Energy



3

## Basic Idea of DVS



## Key Issues for successful DVS

- Efficient Detection of Slack/Idle Intervals
- Efficient Voltage Scaling Policy for Slack Intervals



### **Commercial DVS Processors**

- Transmeta Crusoe
- AMD K2+ (PowerNow Technology)
- Intel SpeedStep
- XScale



Low Power SW.2

J. Kim/SNU

# Voltage Scaling Processors

	Commercial			Academic		
Processors	Transmeta Crusoe (LongRun)	AMD Mobile K6 (PowerNow)	Intel PXA250	UC Berkely (ARM8)	Ubicom LART(StrongARM)	
Scaling Level	200~700MHz 1.1~1.65V	192~588MHz 0.9~2.0V	100~400MHz 0.85~1.3V	5~80MHz 1.2~3.8V	59~251MHz 0.79~1.65V	
Scaling Time	1.1 ↔ 1.65V < 300μs	0.9 ↔2.0V 200μs	Each step 500μs	1.2 ↔ 3.8V 520μs	59↔251MHz : 140μs 0.79→1.65V : 40μs 0.79←1.65V : 5.5ms	
Scaling Power	??	??	??	130µJ	??	

7

# **DVS Support in PXA250**

- Use Two Registers in PXA250 Xscale Core
  - CCCR (Core Clock Configuration Register):
     Specify memory clock & core clock
  - CCLKCFG (Core Clock Configuration) Register
    - Set FCS (Frequency Change Sequence) bit to change the clock speed



Change if FCS bit = 1

# CCCR Setting



Low Power SW.2

J. Kim/SNU

## Example Voltage Scaling Code



# Voltage Scaling in Linux



# ARM IEM DEMO

#### Successful Low Power S/W Techniques

- 1. Understand workload variations of your target
- 2. Devise efficient ways to detect them
- 3. Devise efficient ways to utilize the detected workload variations using available H/W supports

13

#### Low Power SW.2

J. Kim/SNU

#### Roadmap

- DVS in Non Real-Time Systems
- DVS in Real–Time Systems
  - Compiler-level DVS: Intra-task DVS
  - OS-level DVS: Inter-task DVS
  - Application-level DVS
    - -MPEG-decoder implementation
  - Algorithm-level DVS

     Low-power convolution

Low	Power	SW.2
LOw	1 0 1 0 1	011.2

J. Kim/SNU

### Non Real-Time Jobs

- Non Real–Time Jobs
  - No timing constraints
  - No periodic executions
  - Unknown WCET

#### It is hard to predict the future workload!!

15

# DVS for Non Real-Time Jobs

- Basic Approach:
  - · Predict workload based on history information

- Usually based on *some variations of interval scheduler* 
  - PAST, FLAT
  - LONG\_SHORT, AGED\_AVERAGE
  - CYCLE, PATTERN, PEAK

### **Key Question**

#### How can we predict the future workload?

 Based on long term history: Hard to adapt quickly for the changed workload

17

 Based on short term history: Too many clock/voltage changes

#### PAST

- Looking a fixed window into the past
- · Assume the next window will be like the previous one
- If the past window was
  - mostly busy  $\Rightarrow$  increase speed
  - mostly idle  $\Rightarrow$  decrease speed

# Example: PAST

Low Power SW.2



19

### FLAT

Low Power SW.2

- Try to smooth speed to a global average
- Make the utilization of next window to be <const>

18

 Set speed fast enough to complete the predicted new work being pushed into the coming window

J. Kim/SNU

20

# Example: FLAT



## LONG-SHORT

- Look up the last 12 windows
  - Short-term past : 3 most recent windows
  - Long-term past: the remaining windows
- Workload Prediction
  - the utilization of next window will be a weighted average of these 12 windows' utilizations

Low Power SW.2

22

J. Kim/SNU

# Example: LONG-SHORT



# AGED-AVERAGE

- Employs an exponential-smoothing method
- Workload Prediction
  - The utilization of next window will be a weighted average of all previous windows' utilizations
    - geometrically reduce the weight

### Example: AGED\_AVERAGE



#### CYCLE

Workload Prediction



# Example: CYCLE



27

Predict :

0

0

1 2

3

4

utilization = # cycles of busy interval / window size .4 .8 .1 .3 .5 .7 .0

5

6

error measure =  $\frac{|0-.3|+|.4-.5|+|.8-.7|+|.1-0|}{4} = 0.15$ 

The next utilization will be .3

78

time

current time

# PATTERN

- A generalized version of CYCLE
- Workload Prediction
  - Convert the n-most recent windows' utilizations into a pattern in alphabet {A, B, C, D}.
  - Find the same pattern in the past

Low Power SW.2

29

J. Kim/SNU

# Perspectives-based Algorithm



### Roadmap

Pattern = ABCDD

2 3

.5 1 1

4

Predict : The next utilization will be D

0

.3

Low Power SW.2

DVS in Non Real–Time Systems

Example: PATTERN

 $\begin{array}{c|c} A & B & C & D \\ \hline & & & & & \\ 0 & 0 & 25 & 0 & 5 & 7 & 25 & 1 & 0 \end{array}$ 

8

30

9

. . . . . .

5 .....

*Pattern* = *ABCD* 

.1 .35 .6 .9

10 11

12

current

time

- DVS in Real-Time Systems
  - Compiler-level DVS: Intra-task DVS
  - OS-level DVS: Inter-task DVS
  - Application-level DVS
     MPEG-decoder implementation
  - Algorithm-level DVS
    - -Low-power convolution

32

time

## Two Types of DVS Algorithms

- Inter-task DVS algorithms
  - Determine the supply voltage and clock speed on *task-by-task* basis
- Intra-task DVS algorithms
  - Determine the supply voltage and clock speed within a single task boundary

#### Inter-task DVS

- Inter-Task Voltage Scheduling for Hard Real-Time Systems [Yao95, Hong98, Okuma99, Shin99, Lee99].
  - Problem : Given a set of tasks, how to assign the proper speed to each task dynamically while guaranteeing all their deadlines.
  - Task-by-task Speed Assignment
    - The slack time due to a task used by following tasks, not by the current one.
  - Practical Limitations
    - Requires OS modifications
    - Cannot be applied to a single-task environment
    - Can be ineffective in a multi-task environment

L	Low Power SW.2	33	J. Kim/SNU	Low Power SW.2	34	J. Kim/SNU

# An Example of Ineffective Inter DVS



35

- A dominant task (C)
  - Exploits small slack times from other tasks.
  - · Cannot use its own.

# Earlier Version of Intra-task DVS



#### Run-time voltage hopping [Lee00]

- Each task is partitioned into N timeslots.
- Frequency and voltage determined for each timeslot.
- Voltage scheduling embedded in application programs.
- Can be applied to conventional non-DVS OS.
  - No systematic guideline
  - Manual selection of scaling points
  - Too much burden for average programmers



### Overview of Intra-task DVS



#### Low Power SW.2

WCEP

J. Kim/SNU

Non-WCEP : there is no slack time

But we cannot know the execution path in advance !!

40

*b*7 10

Need to know the remaining

worst case execution cycles

for a new speed.

Low Power SW.2

Basic Idea : Inter-task DVS

*b*<sub>5</sub> 10





### The Change of C<sub>RWEC</sub>



# B-type VSE



# **VSE** Selection



# **Code Generation for VSEs**



## Automatic Voltage Scaler



#### **Simulation Results**



- Less than 25% and 7% of the original program
- There is a large difference between WCET and ACET of the MPEG-4 decoder

50

Low Power SW.2

J. Kim/SNU

## Simulation Results



- How many times voltage scaling code were executed
- When C<sub>VTO</sub> < 30μsec in MPEG-4 encoder, the number of voltage transitions decreases sharply, and energy consumption does not increase rapidly.
- How many copies of voltage scaling code ?
  - 20 VSEs are inserted when  $C_{VTO} > 50 \ \mu sec.$

#### A Profile-Based Intra-Task Voltage Scheduling

- IntraVS algorithm based on average-case execution information
- Average-case execution paths (ACEPs) are the most frequently executed paths
- More effective than the original intraVS algorithm
- The timing constraints of a hard real-time program is still satisfied, even if the ACEPs are used for voltage scaling decisions.

# General IntraVS Algorithm



### **RAEP-based IntraVS**

Motivations

- To make the common case more energy-efficient
- If we use one of hot paths as a reference path for intraVS, the speed change graph for the hot paths will be a near flat curve with little changes in clock speed.
- Even for the paths that are not the hot paths, they are more energy-efficient because they can start with a lower clock speed that RWEP-based IntraVS.
- RAEP is the best representative of the hot paths.

54

Low Power SW.2

J. Kim/SNU

## RAEP-based IntraVS



# **RAEP-based IntraVS**



## **Reference Path Modification**



### **Experimental Results**



## Conclusion

- Presented a novel intra-task DVS algorithm using static timing analysis on RWECs
- Provides a framework for automatic DVS-aware lowpower program generation
- The RAEP-based IntraVS algorithm exploits the fact that the average-case execution paths are more likely to be followed at run time than the WCEP.

59

 Demonstrated the effectiveness of the approach using MPEG-4 encoder/decoder programs

## Experiments on Itsy



### **Experimental Environment**

- Itsy Pocket Computer V2.6
  - CPU : Intel StrongARM SA1110
  - Frequency scaling: 11 levels (59.0 MHz ~ 206.4 MHz)
  - Voltage scaling: 30 levels (1.00 V ~ 2.00 V)
  - Default setting: 1.55 V/206.4 MHz
- Linux operating system (ver. 2.0.30)



61

Low Power SW.2	Low	Power	SW.2	
----------------	-----	-------	------	--

J. Kim/SNU

Experimental Results



#### **Experimental Results**

#### DVS EXPERIMENTS ON ITSY.

		MPEG-4	Decoder	MPEG-4 encoder	
Factors		DVS-aware	normal	DVS-aware	normal
Energy (mJ)		0.11	0.22(0.18)	0.28	0.81(0.62)
Normalized Energy		0.51(0.62)	1	0.35(0.46)	1
Execution Time (sec)		1.18	0.46	5.34	1.54
WCET (sec)		3.2		21.0	
Selected	B-VSE	2		1	
VSEs	L-VSE	1		2	
Management	Function	2		3	
Code	Loop	5		8	

MANAGEMENT CODE OVERHEAD.

Management Code	Code number in Fig. 7	Number of assembly instructions
Loop Enter	code 1 and 2	30
Loop Header	code 3	16
Loop Exit	code 4	16
Function Enter	code 5	14
Function Return	code 6	11
B-type VSE	code B	$\approx 200$
L-type VSE	code L	$\approx 200$

62

Low Power SW.2

J. Kim/SNU

## Roadmap

- DVS in Non Real-Time Systems
- DVS in Real-Time Systems
  - Compiler-level DVS: Intra-task DVS
  - OS-level DVS: Inter-task DVS
  - Application-level DVS
    - -MPEG-decoder implementation

64

Algorithm-level DVS

 Low-power convolution

#### Inter-task DVS Overview

- Inter-task DVS algorithms
  - Determine the supply voltage and clock speed on *task-by-task* basis
- Inter-task DVS
  - Is similar to that of imprecise computation in conventional real-time systems
    - -Imprecise computation
      - · Use the slack time to increase the values of results
      - · While guaranteeing the feasible schedule of tasks
    - Dynamic voltage scaling
      - Use the slack time to lower the voltage/clock speed
      - · While guaranteeing the feasible schedule of tasks

65

Low	Power	SW/ 2
LOW	r uwei	3VV.Z

#### Preliminaries

- Computing model
  - Non-real-time
    - -tasks have no timing constraints
  - Real-Time
    - -Timing constraints
    - -Periodic and(or) aperiodic tasks
    - -Scheduling policy : EDF, RM, and etc.
- Different DVS algorithms are necessary depending on different computing models.

```
Low Power SW.2
```

J. Kim/SNU

# Inter-Task DVS

- "Run-Calculate-Assign-Run" strategy for the supply voltage determination
  - Running the current task
  - Calculating the maximum allowable execution time for the next task
    - -WCET plus slack time
  - Assigning the supply voltage for the next task

67

Running the next task

## Generic Inter-DVS Algorithms

- Consist of two parts
  - Slack estimation
    - -Identify as much slack times as possible
    - -Slack times
      - Static slack times

 $^{-}$  Extra times available for the next task that can be identified statically

- Dynamic slack times
  - Ones caused from run-time variations of the task executions
- Slack distribution
  - Adjust the speed so that the resultant speed schedule is as flat as possible

# Static and Dynamic Approaches

- Off-line (Static) voltage scheduling approaches
  - The execution times are assumed to be known a priori
  - There are several optimal solutions for EDF, RM, and etc.
- On-line (Dynamic) voltage scheduling approaches
  - The execution times are assumed to be not known

69

There cannot be an optimal solution

## **Slack Estimation Methods**



# Maximum Constant Speed

- The lowest possible clock speed that guarantees the feasible schedule of a task set
  - EDF scheduling
    - If the worst case processor utilization U of a given task set is lower than 1.0 under the maximum speed  $f_{\rm max}$ , the task set can be scheduled with a new maximum speed

$$U = \sum_{i=1}^{n} \frac{c_i}{p_i} \qquad f_{MSC} = U \cdot f_{max}$$

Rate Monotonic scheduling

$$L_{i}(t) = \frac{\sum_{j=1}^{i} \left[ \frac{t}{T_{j}} \right] C_{j}}{t} \qquad f_{MCS} = f_{\max} \times \max_{i} \{ L_{i}(t) \mid 1 < i \le n, 0 < t < d_{i} \}$$

71

## Example



Low Power SW.2

Maximum constant speed

# Stretching to NTA

- Even though a given task set is scheduled with the maximum constant speed, since the actual execution times of tasks are usually much less than their WCETs, the tasks usually have dynamic slack times
- For the task  $\tau$  which is scheduled at time t
  - If the next task is later than  $t + WCET(\tau)$
  - We can slow down the execution of  $\tau$  so that its execution completes exactly at this next task arrival time (NTA)

73

#### Example



# Priority-Based Slack Stealing

- Exploits basic properties of priority-driven scheduling such as EDF and RM
  - When a higher-priority task completes its execution earlier than its WCET, the following lower-priority tasks can use the slack time from the completed higher-priority task
- Advantage

Low Power SW.2

- Most task instances in a hyper-period may have chances to utilize dynamic slack times
  - -Because most task executions complete earlier than WCETs

75

-Therefore, many task instances can be scheduled with lowered voltages and speeds

# Example



# **Utilization Updating**

- The actual processor utilization during run time is usually lower than the worst case processor utilization
- This method is to estimate the required processor performance at the current scheduling point
  - By recalculating the expected worst case processor utilization

77

-Using the actual execution times of completed task instances

#### Example



# Slack Distribution Method

Greedy approach

Low Power SW.2

- All the slack times are given to the next activated task
- Most inter-task DVS algorithms have adopted it
- Clearly, this approach is not an optimal solution, but is widely used because of its simplicity

79

### Existing Inter-Task DVS Algorithms

Category	Scheduling Policy	DVS Policy	Used Method
		IppsEDF	(3)+(4)
		ccEDF	(6)
	FDF	laEDF	(6)*
Inter-task	k EDF	DRA	(3)+(4)+(5)
DVS		AGR	(4)*+(5)
		IpSEH	(3)+(4)+(5)*
	DM	IppsRM	
	RIM	ccRM	(4)*
Intra-task	Path-based method	intraShin	(1)
DVS	Stochastic method	intraGruian	(2)

J. Kim/SNU

#### SimDVS: A Unified DVS Evaluation Environment



## Experimental Results



# **Experimental Results**



# **Experimental Results**



### **Experimental Results**



## Experimental Results (IntraDVS)



#### Performance Evaluation DVS Algorithms for Hard Real-Time Systems Using DEW

#### **DEW – DVS Evaluation Workbench**

- XScale-based DVS evaluation environment
  - Pros
    - Allows to monitor real system behaviors under DVS
  - Cons
    - Slower than software simulation
      - Because DEW runs actual applications
    - Less flexible for experimental studies
      - Because DEW represents a single machine specification





#### Evaluation Results Using SimDVS and DEW



# Sources of Differences

- Impacts of
  - System overhead
    - Basic : context switching overhead and tick scheduler overhead
    - DVS : slack computation and clock/voltage scaling
  - System timing resolution
    - Simulator : continuous time model
    - Real system : discrete time model
  - · Memory behavior
    - Changes in cache and memory access behavior

91

- Data/Instruction fetch latency

# Example of System Overheads on a Real Platform



## System Overhead



#### DVS H/W

The ratio of time delay caused by the clock/voltage scaling hardware

DVS S/W

The ratio of time delay caused by the slack computation in a DVS algorithm

#### SYS rest

The ratio of the rest of the system overhead such as context switching and timer service

J. Kim/SNU

#### System Overhead Variations

- The system overhead increases very quickly as the task execution frequency increases
  - · In particular, DVS parts increase quickly



## **Energy Efficiency Variations**

 In DRA, AGR, and IpSHE, the increased system overhead (due to the increased execution frequency) significantly affect the energy efficiency



#### Changes in Memory System Behaviors (1)



- Under a DVS-enabled RTOS, Task's execution time increases due to the lowered clock speed
  - Desirable for reducing energy consumption
  - But, it can introduce negative side effects as well
    - An increase in the number of task preemptions which increases the number of memory accesses
- In aggressive algorithms, the number of preemptions increases more rapidly than the others

#### Changes in Memory System Behaviors (2)



#### PXA250

- Performance Monitoring Unit
- 32-way set-associative cache of Inst/Data cache
- Each application
- 16–KB program code

The increases in memory accesses can be attributed to two sources

- The increase in the number of preemptions
- The increase in memory accesses from the algorithm itself

#### J. Kim/SNU

#### References

- Transmeta Corporation. Crusoe Processor. http://www.transmeta.com, June 2000.
- AMD Corporation. PowerNow! Technology. http://www.amd.com, December 2000.
- Intel Corporation. Intel XScale Technology. http://developer.intel.com/ design/ intelxscale/, November 2001.
- I. Hong, G. Qu, M. Potkonjak, and M. B. Srivastava. Synthesis Techniques for Low-Power Hard Real-Time Systems on Variable Voltage Processor. In Proceedings of the IEEE Real-Time Systems Symposium, pages 178-187, December 1998.
- Y. Shin, K. Choi, and T. Sakurai. Power Optimization of Real– Time Embedded Systems on Variable Speed Processors. In Proceedings of the International Conference on Computer– Aided Design, pages 365–368, November 2000.
- H. Aydin, R. Melhem, D. Mosse, and P. M. Alvarez. Dynamic and Aggressive Scheduling Techniques for Power-Aware Real-Time Systems. In Proceedings of IEEE Real-Time Systems Symposium, December 2001.

Low Power SW.2

J. Kim/SNU

#### References

- P. Pillai and K. G. Shin. Real-Time Dynamic Voltage Scaling for Low-Power Embedded Operating Systems. In Proceedings of 18th ACM Symposium on Operating Systems Principles (SOSP'01), October 2001.
- D. Shin, J. Kim, and S. Lee. Intra-Task Voltage Scheduling for Low-Energy Hard Real-Time Applications. IEEE Design and Test of Computers, 18(2):20-30, March 2001.
- F. Gruian. Hard Real-Time Scheduling Using Stochastic Data and DVS Processors. In Proceedings of the International Symposium on Low Power Electronics and Design, pages 46– 51, August 2001.
- W. Kim, J. Kim, and S. L. Min. A Dynamic Voltage Scaling Algorithm for Dynamic-Priority Hard Real-Time Systems Using Slack Time Analysis. To appear in Proceedings of Design, Automation and Test in Europe (DATE'02), March 2002.

#### References

- D. Grunwald, P. Levis, and K. I. Farkas. Policies for Dynamic Clock Scheduling. In Proceedings of the 4th Symposium on Operating Systems Design and Implementation, pages 73–86, October 2000.
- S. Lee and T. Sakurai. Run-time Voltage Hopping for Lowpower Real-Time Systems. In Proceedings of the 37th Design Automation Conference, pages 806-809, June 2000.
- T. Burd and R. Brodersen. Design Issues for Dynamic voltage scaling. In Proceedings of the International Symposium on Low Power Electronics and Design, pages 9–14, July 2000.
- D. Burger and T. M. Austin. The SimpleScalar Tool Set, version 2.0. Technical Report 1342, University of Wisconsin-Madison, CS Department, June 1997.
- F. Yao, A. Demers, and A. Shenker. A Scheduling Model for Reduced CPU Energy. In Proceedings of the IEEE Foundations of Computer Science, pages 374–382, 1995.



#### Direct Implementation: Constant-Workload Algorithm



### Low-Power Implementation



#### Variable Workload Algorithm Based on Kernel Characteristics

104



#### Modified Convolution Algorithm: SDMK

- For 1 or –1, no multiplication.
- For 0, no addition & no multiplication.
- For the same absolute values, a single multiplication.

0	1	1
-1	0	1
-1	-1	0

	Direct	SDMK
Number of	9 additions &	6 additions &
Operations/pixel	9 multiplications	0 multiplications

# Direct vs. SDMK



107

#### **Exec Time Prediction & Speed Setting**

#### By a static method

 Based on the number of required arithmetic operations

#### By a dynamic-method

- Based on actual measurements of execution times
- In the direct algorithm, by pre-constructed speed table

Low Power SW.2	109	J. Kim/SNU		Low Power SW.2	110

# **Results (Energy Dissipation)**



Average 67.6% energy saving in the core processor, and 62.8% in the whole Itsy system.

# **Results (Execution Time)**

**Experimental Environments** 

Voltage scaling: 30 levels (1.00 V ~ 2.00 V)

Multimeter

1

Multimeter

2

Frequency scaling: 11 levels (59.0 MHz ~ 226.4 MHz)

ltsy

V2.6

Rest of

Itsy HW

Recoding 번 Computer

J. Kim/SNU

• Itsy Pocket Computer V2.6

•

8

6

• CPU : Intel StrongARM 1110

Linux operating system (ver. 2.0.30)

0

0 0 0

> -S



## Conclusions

- Presented a low-power implementation of image convolution algorithm for variable voltage processors.
- The energy efficiency of the proposed implementation comes from:

 $\mathbf{E} \propto \mathbf{N}_{\text{cycle}} \cdot \mathbf{V}_{\text{DD}}^2$ 

- Smaller N<sub>cycle</sub>
- Lower V<sub>dd</sub>

Low Power SW.2

- Fewer memory references
  - i.e., less energy consumed in non-CPU components

113

## **MPEG Decoder Implementation**

#### 버퍼를 활용한 DVS

- Workload-variation 슬랙 시간을 활용하는 것이 가능
- 추가적인 메모리 자원을 소모
  - one-buffer-size = image\_width \* image\_height \* byte\_per\_pixel



## **Measurement Results**

- 버퍼 방식 DVS 기법 사용 (WCET사용)
  - Bitrate of sample video : 163Kbps



• Energy saving : up to 53%

# Demo (1)



115

# Demo (2)

각 DVS policy 들 사이의 총 전력 소모를 비교



#### Energy-Optimal Off-Line Voltage Scheduling

# Off-Line Volt. Sched. Problem

- Voltage schedule (speed schedule) : S(t)
  - · the processor speed as a function of time
  - The energy consumption under S(t) is given by  $E(S) = \begin{bmatrix} P(S(t)) & dt \end{bmatrix}$ 
    - $E(S) = \int_{interval} P(S(t)) dt$
    - P is a convex function from speed to power
- Given N jobs  $J_1$ ,  $J_2$ , ... ,  $J_N$  where
  - $r_i$ : the release time of  $J_i$
  - $d_i$ : the deadline of  $J_i$
  - $c_i$ : the workload (# of execution cycles) of  $J_i$ 
    - assumed to be known a priori
  - p<sub>i</sub>: the priority of J<sub>i</sub>
  - compute a feasible voltage schedule S(t) that minimizes E(S)
- S(t) is feasible iff S(t) gives J<sub>i</sub> its workload c<sub>i</sub> between r<sub>i</sub> and d<sub>i</sub> for all J<sub>1</sub>, J<sub>2</sub>, ... , J<sub>N</sub>

#### Low Power SW.2

J. Kim/SNU

## Existing Works for the Problem

- Note that the system model covers
  - Fixed-priority (RM, DM) periodic/aperiodic task set
  - EDF periodic/aperiodic task set
    - $p_i < p_j$  iff  $d_i < d_j$
- For EDF job sets (a special case), the problem can be solved in poly. time by Yao's algo.[FOCS'95]
  - solution space = convex , obj. func. = convex
- For general job sets, the problem becomes much difficult
  - main source of difficulty : feasibility condition
  - Quan & Hu [TCAD'03]: exhaustive optimal algo.
  - Yun & Kim [TECS'03]: NP-hardness & FPTAS

### References

- Optimal algorithm for EDF job sets
  - F. Yao, A. Demers, S. Shenker, "A Scheduling Model for Reduced CPU Energy", In Proc. Foundations of Computer Sciences (FOCS'95), 1995
- Heuristic for FP job sets
  - G. Quan and X. Hu, "Energy Efficient Scheduling for Real-Time Systems On Variable Voltage Processor", In Proc. Design Automation Conference (DAC'01), 2001.
- Exhaustive optimal algorithm for FP job sets
  - G. Quan and X. Hu, "Minimum Energy Fixed Priority Scheduling for Variable Voltage Processors", IEEE Transactions on Computer Aided Design and Systems, 2003.
- NP-hardness proof & FPTAS for FP job sets
  - H.-S. Yun and J. Kim, "On Energy-Optimal Voltage Scheduling for Fixed-Priority Hard Real-Time Systems", ACM Transactions on Embedded Computing Systems, 2003.

121

Low Power SW.2