Continuous Frequency Adjustment Technique Based on Dynamic Workload Prediction

Presenter: Kimish Patel

Authors: Hwisung Jung and Massoud Pedram

University of Southern California
Dept. of Electrical Engineering

Overview

- Introduction
- Prior Work
- Power Management Architecture
- Workload-Driven Frequency Adjustment
- Experimental Results
- Conclusions

Introduction

- Power management becomes a first-order concern
  - More function blocks in SOC are being built with DVFS
  - Different clock and voltage domains on the same chip
- Limitations in dynamic power management
  - Intricate trade-offs between power-saving and performance
  - Power manager becomes a heavy burden due to DVFS scheduling
- DVFS-enabling techniques depend on:
  - Configuration of voltage/frequency control circuits
  - Efficiency of workload prediction mechanism

Prior Works

- M. Najibi et al. (ICCAD 2006)
  - Frequency management based on variable update intervals
- Y. Cho et al. (DAC 2006)
  - Analytical models for selecting an optimal DVFS
- A. Iyer et al. (ICCAD 2002)
  - Online DVFS by utilizing queue in multiple frequency domains
- Q. Wu et al. (HPCA 2005)
  - A voltage island-based power management technique
- J. Liu et al. (ICCAD 2004)
  - Optimization technique for power mode transitions

Shortcomings & Proposal

- Traditional approaches for DPM
  - Highly dependent on the operating system speed
  - Based on software control for DVFS
  - Non-negligible overhead for policy computations

↓

- Propose a continuous frequency adjustment tech.
  - Adjust a frequency value of various functional blocks
  - Minimize the overhead of power-mode transitions
- Key features of the proposed technique
  - Workloads are captured at runtime by using IVP
  - A continuously adjusted frequency is generated
  - Energy penalties due to power-mode transitions are reduced

Power Management Architecture (1/3)

- **Proposed power management architecture**
  - **Power manager**
  - **Performance monitoring unit**
  - **Dynamic frequency adapter**

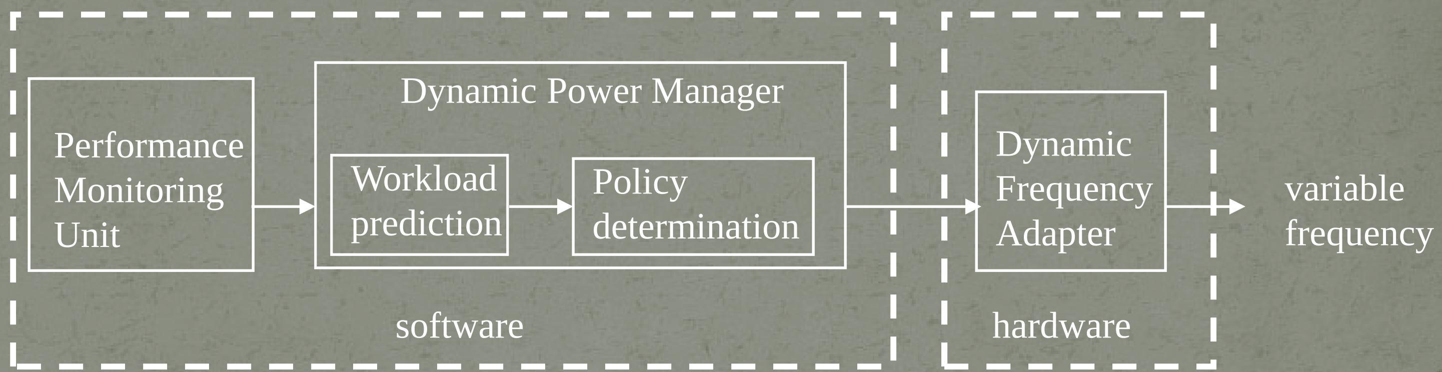

Proposed power management architecture

Power Management Architecture (2/3)

- **Power manager**
  - Determine and execute a power management policy
  - Predict workloads
- **Performance monitoring unit**
  - Profile and analyze workloads by using service queue (SQ)
  - SQ is represented by a G/M/1 queuing model
    - Inter-arrival times are arbitrarily distributed
    - Service times are exponentially distributed
    - G/M/1 mode is more realistic than M/M/1 model
- **Dynamic frequency adapter**
  - Increase (decrease) frequency at slow or fast pace, depending on how slow or fast the workload is changing

Power Management Architecture (3/3)

- Dynamic frequency adapter (cont.)
  - Use a mapping table for selecting an optimal frequency
  - Generate a variable frequency using pulse width modulator

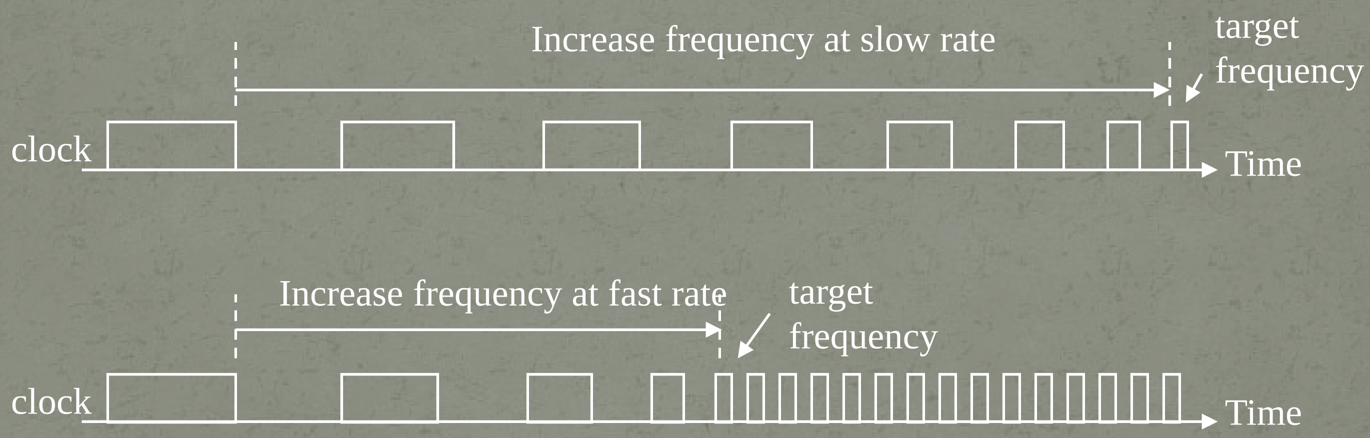

Continuously changing the frequency at a slow or fast pace

Workload-Driven Freq. Adjustment (1/3)

- **IVP-based workload prediction**
  - **Initial Value Problem is defined to predict workload *w*(*t*)**

$$\partial w / \partial t = f(t, w), \quad w(t_i) = w_i$$

  - ***w*(*t*): workload (i.e., arrival rate of tasks) to functional block**
  - ***f*: function of providing an operating frequency**
  - **w_i: workload at the beginning of current interval [t_i, t_{i+1}]**

- **A method to solve the IVP is to approximate the solution of ordinary differential equation**
  - **Predict the next value of *w***

$$w(t+h) = w(t) + hw'(t)$$

Workload-Driven Freq. Adjustment (2/3)

- Evaluate different methods for solving the IVP
  - Set $w(0) = 0.3$ as the initial value
  - Euler method (the simplest approach) shows low accuracy
  - 4th-order Runge-Kutta method exhibits low error, but higher computation

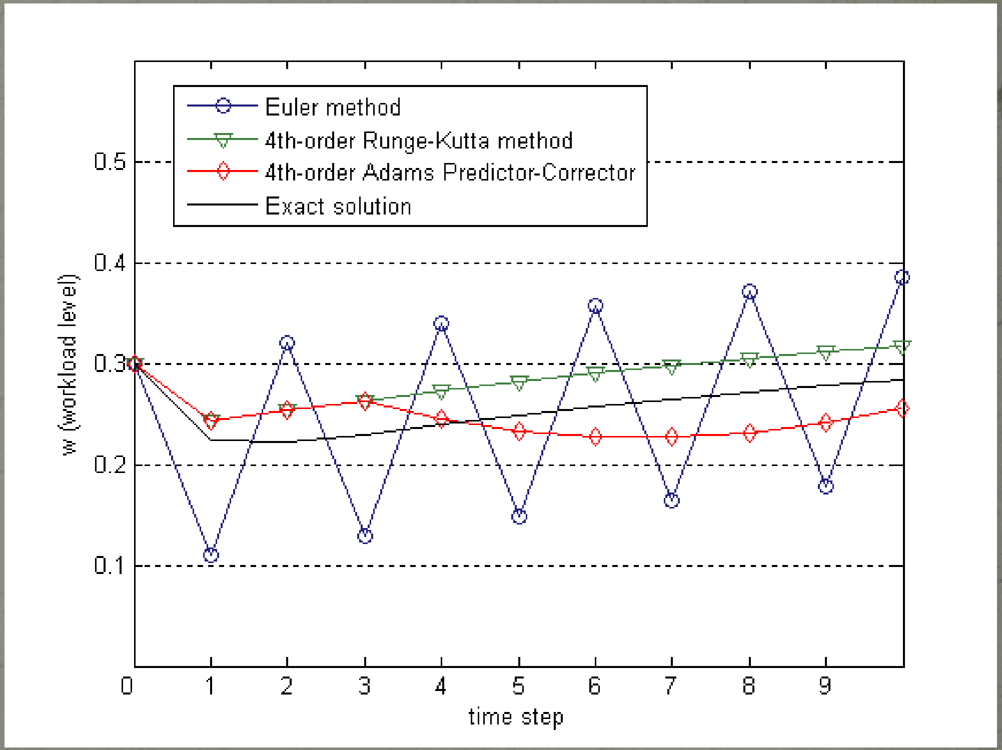

Evaluation of various IVP solutions

Workload-Driven Freq. Adjustment (3/3)

- **Workload prediction technique**
  - **Compute $w^{exact}(t)$ by using performance monitor**
  - **Estimate $w^{pred}(t+h/2)$ by using the midpoint method**

$$w^{pred}(t+h/2) = \frac{w^{exact}(t) + w^{exact}(t-h)}{2}$$

  - **Predict $w^{pred}(t+h)$ by solving the IVP**

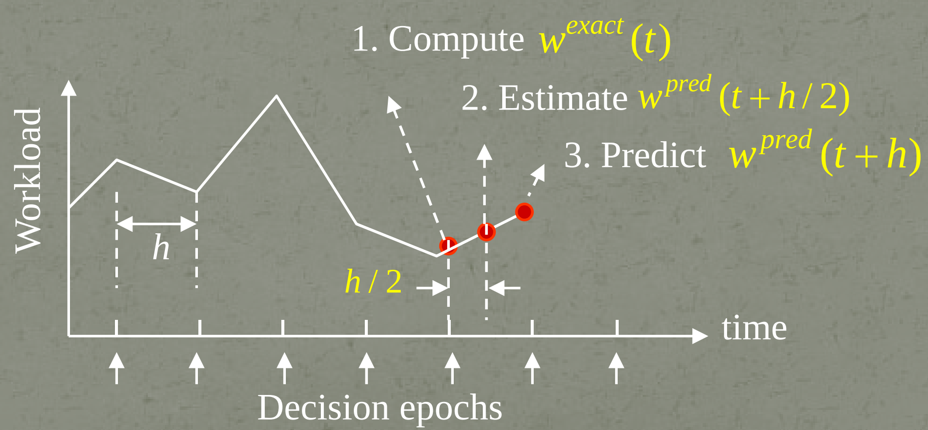

Workload prediction technique

Mapping of Workloads to Optimal Freq.

- **Build a pre-characterized mapping table**
  - **Through extensive offline simulations during design time**
  - **Consider performance characteristics of functional block**

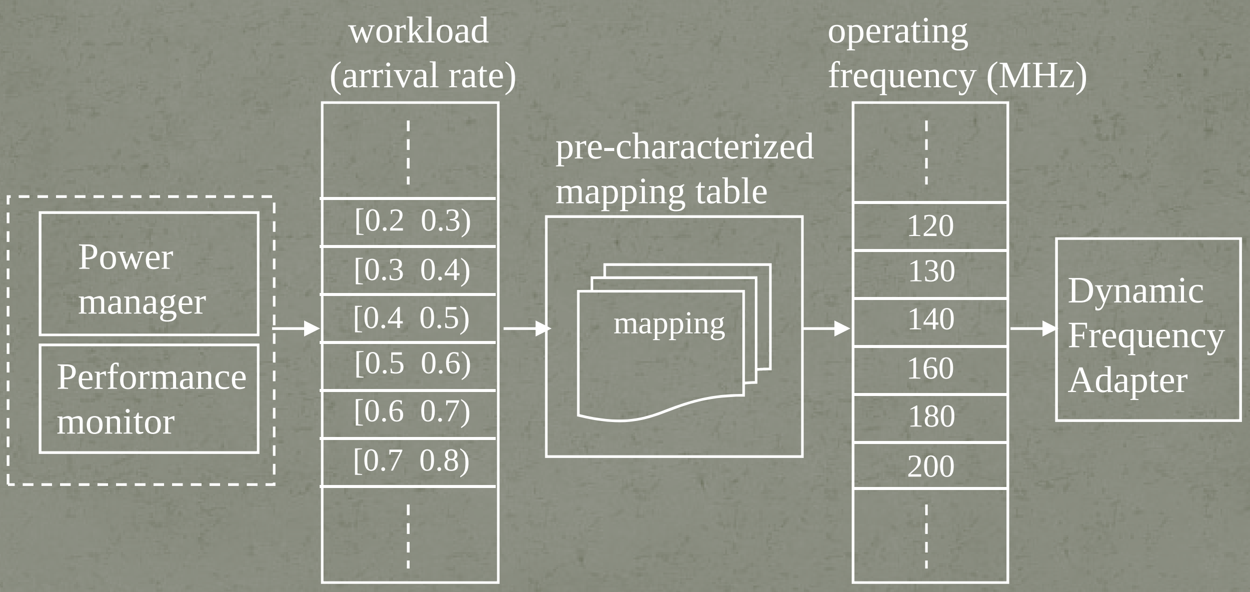

Mapping of workloads to optimal operating frequency values for each FB

Experimental Setup

- Applied the technique to a gigabit Ethernet controller
  - Implemented with a TSMC 65nm cell library

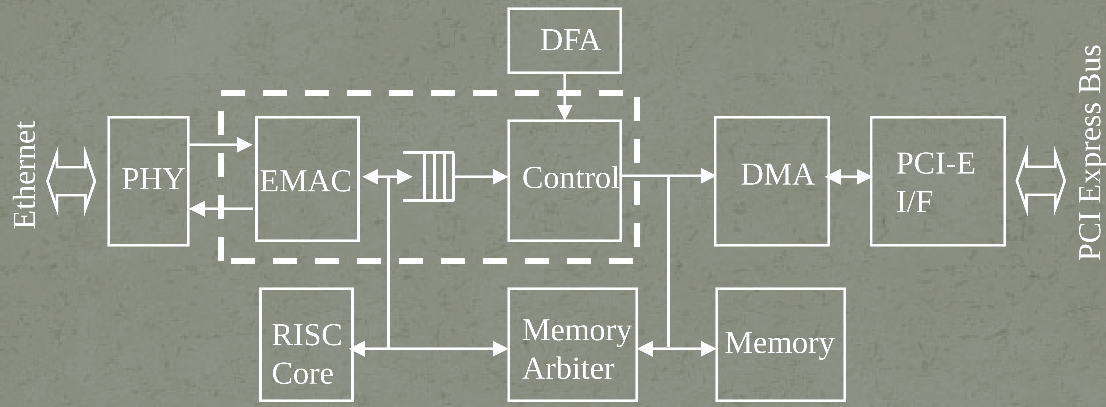

Block diagram of simplified Ethernet controller

- Consider EMAC as a service requestor (SR)
- Consider Control block as a service provider (SP), where frequency is controlled by DFA

Experimental Results (1/3)

- **Analyze energy dissipation of the service provider**
  - **Use Synopsys's Power Compiler**
  - **Use different workload (i.e., arrival rate of tasks), where 0.1 ≤ arrival rate ≤ 0.9**
  - **Achieve full-duplex maximum bandwidth**

1000Base-T	Arrival rate of tasks								
	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Energy (normalized)	1.96	2.62	4.37	5.04	5.85	6.67	8.89	9.86	10.83
Utilization (%)	9.5	14.2	20.6	32.1	41.8	54.0	63.2	67.5	73.5

Energy dissipation (normalized) and utilization of the service provider

- **Overhead of designing the DFA is negligible**
  - **150 standard cells with 2uW power consumption**

Experimental Results (2/3)

- **Evaluate the effectiveness of proposed CFA**
  - **Assume workload changes between 0.1 and 0.9**
  - **Use three sets of frequency values ($F_1 < F_2 < F_3$)**
  - **Implement conventional PM policies for comparison**
    - **PM1:**

      - utilize dynamic frequency scaling technique account for 100us power-mode transition overhead
      - Use the lowest F_1 when $0.1 \leq$ workload ≤ 0.3
      - Use the highest F_3 when $0.6 \leq$ workload ≤ 0.9
      - Use F_2, otherwise
  - **PM2:**

      - Same as PM1 except that frequency change is avoided when the same frequency change occurs

Experimental Results (3/3)

- **Generate dynamic workload randomly**
  - **With 100, 500, 1000, and 5000 decision epochs**
  - **Workload distribution, e.g., Low = 0.1 ≤ arrival rate ≤ 0.3**

No. of decision epoch	Workload distribution			Average Power (mW)			Power saving over		Energy saving over	
	Low	Mid	High	PM1	PM2	CFA	PM1	PM2	PM1	PM2
100	26	43	31	13.2	11.8	11.4	12.9%	2.3%	11.4%	8.8%
500	143	196	161	13.3	11.7	11.5	13.3%	2.2%	12.0%	9.4%
1000	277	417	306	13.4	11.8	11.5	13.6%	2.1%	12.3%	9.5%
5000	1494	1954	1552	13.4	11.9	11.6	13.1%	2.1%	12.1%	9.6%

Power and energy savings of the CFA

- **Our approach achieves power and energy savings up to 13.6% and 12.3%, compared to PM1 policy**

Conclusions

- Addressed the problem of power management technique in the context of handling dynamic frequency management
- Proposed a continuous frequency adjustment technique based on a workload prediction method
- Workload prediction is formulated with IVP
- Experimental results shows that the proposed technique ensures robust energy savings under dynamic workloads.