Dynamic Voltage and Frequency Scaling (DVFS) Scheme for Multi-Domains Power Management

Jeabin Lee, Byeong-Gyu Nam and Hoi-Jun Yoo
Department of EECS, Korea Advanced Institute of Science and Technology (KAIST)
373-1, Guseong-dong, Yuseong-gu
Daejeon, 305-701, Republic of Korea
Email: jeabin@eeinfo.kaist.ac.kr

Abstract

The power of 3 different power domains is managed by continuous co-locking of voltage and clock, dynamically varying clock frequency and supply voltage level from 90MHz to 200MHz and from 1.0V to 1.8V, respectively. A test 3D-graphics SoC is divided into 3 power domains and their power are managed separately. The workload of each domain is the control parameter to each power management unit (PMU). It takes 0.45mm^2 with 0.18um CMOS process and consumes 5mW. Total SoC takes 17.2 mm^2 and consumes 52.4mW at full operation with triple domains power management.

I. INTRODUCTION

The power reduction is an important design issue for the SoC especially in energy constrained applications. Lowering the supply voltage is one of the attractive approaches to save power of the variable workload system [1, 2]. Furthermore, dynamical scaling the clock frequency and the supply voltage achieves extremely efficient energy saving. Because of this reason, there are many previous researches about DVFS power management for digital system such as RISC, DSP and Video Codec [3-5].

Steady increase of the functional complexity and integration level of SoC separates the chip into multiple the power domains that have different performance requirement. And each power domain operates for different workload with different supply voltage and clock frequency. For example, H.264 Codec, reported in [6], has 2 power domains. However, it manages the power of one domain with only two supply voltages and clock frequencies. And the voltage and frequency of the other domain are fixed. With this passive approach, there is obvious limitation in energy saving.

In this work, we manage the power of 3D graphics engine that has 3 independent domains, and each domain has different performance requirements. Therefore, dynamic frequency and voltage scaling can be applied to each power domain to realize efficient energy reduction. However, there has been no report on multi-domain power management by DVFS. In this paper, multi-domains DVFS power management scheme and its PMU are successfully demonstrated and reported.

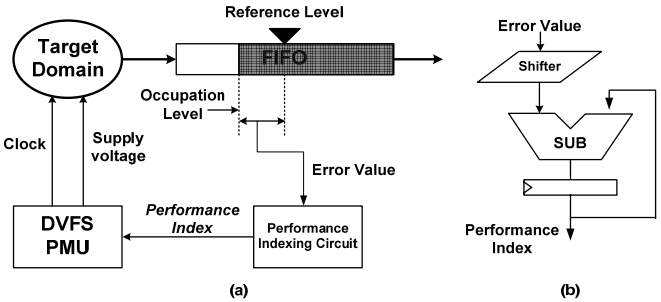

Fig. 1 (a) Power management loop for PMU control
(b)Performance indexing circuit (PI)

II. DVFS SCHEME FOR MULTI-POWER DOMAINS

This work proposes the triple domains DVFS power management scheme for graphics processing unit (GPU). The proposed low-power GPU is composed with RISC, geometry processor (GP) and rendering engine (RE). And each part is independent power domain [7]. The voltage and frequency of each power domain are dynamically scaled according to the performance requirement of each domain, and the performance requirement is generated from occupation level of the FIFO. The output of domain is buffered in the FIFO to transfer to the next pipeline stage. And this performance requirement is used for controlling the power management unit (PMU) that provides the supply voltage and clock to each domain.

In the conventional voltage scaling schemes [3-5], the PMUs feedback control the buck converter digitally to find the voltage level for the minimum energy point (MEP) tracking. However, these schemes are not appropriate for the multi-domain because of their large size. In addition, its digital blocks such as digital filter resulted in area and power overhead to the PMU. In this paper we proposed a new scheme to reduce both of the power and area as well as off-chip inductor and capacitor. The proposed PMU can be reused as an Intellectual Property (IP).

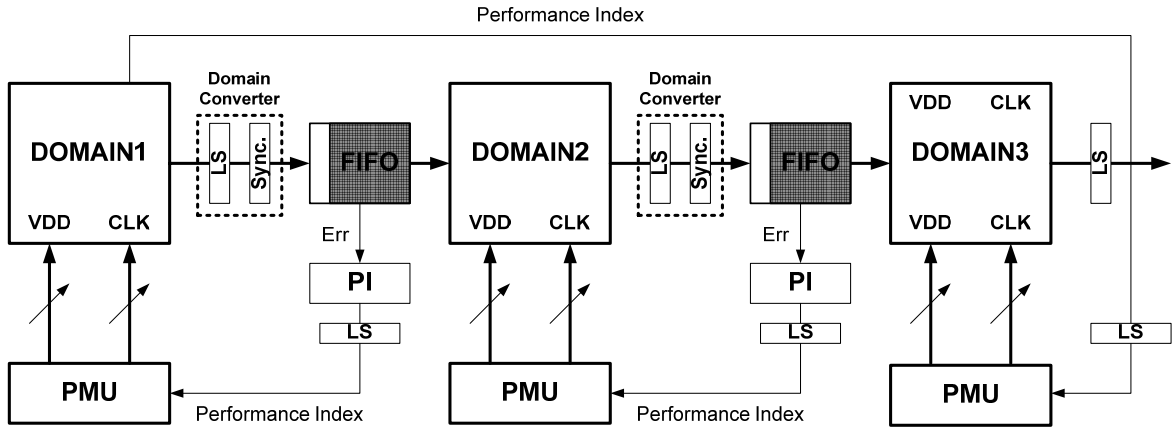

Fig. 2 Triple Domains Power Management

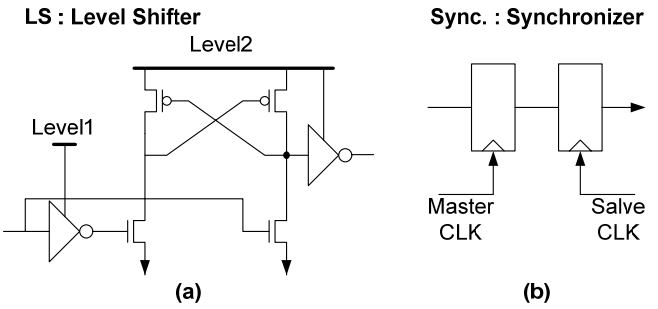

Fig. 3 (a) Level shifter (b) Synchronizer

A. *Power Management with varialble worklaod*

The triple domains which have different workloads should be controlled separately according to their own performance requirements.

Each domain is connected through a FIFO which plays the role of workload monitor. Its occupation level is used as the indicator of the workload of the each domain to control the speed of target domain. Fig. 1 (a) shows how the occupation level is compared with the reference level to get the error value, and how this value is used to determine a new performance requirement. The target performance index is determined by the workload of target domain and it is generated by the performance indexing circuit (PI) in Fig. 1 (b). According to the performance index, the PMU generates the clock and the supply voltage to adjust the speed of target domain.

B. *Triple-domains power management*

This GPU is divided into three pipeline stage in a systemic point of view [7]. And each pipeline stage power can be managed according to the different workload. Fig. 2 shows the triple domains workload monitoring and power management system. The performance requirement of the total system is decided by the RISC in the first domain with directly controlling the PMU3. After that, the FIFOs between each domain monitors the workloads of the first and second domain using difference between occupation level and reference level, and control the PMU with performance index.

In this case, each domain has different clock and supply voltage. System has to convert the voltage level and adjust the phase of data which goes from domain 1 to domain 2 and from domain 2 to domain 3. The level shifter (LS) and synchronizer (Sync.) in Fig 3 change the voltage level and convert the clock frequency of transferred data, respectively. Therefore, the domain converters composed with LS and Sync. are inserted between each domain to adjust voltage level and to avoid metastability of data.

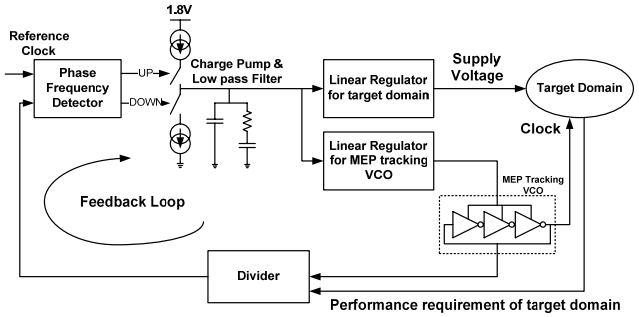

Fig. 4 The DVFS scheme for multi-domains power management

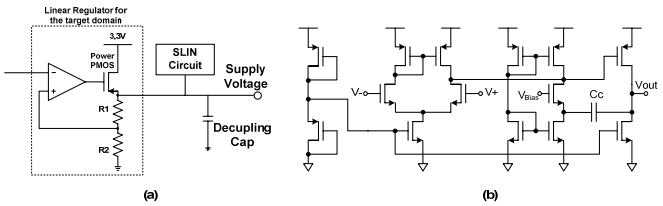

Fig. 5 (a) Dynamically Scaled Supply Voltage Generating Circuit (b) Operational Amplifier in The Linear Regulator [8]

III. CoCoLo (Power Management Unit)

A. Overall Archtecture

The DVFS scheme in Fig. 4 uses the phase locked loop (PLL) structure to find supply voltage level and clock frequency level on the MEP. It merges PLL and linear regulator (LR) as shown in Fig. 5. The LR [9, 10] is controlled by PLL frequency feedback loop and generates supply voltage. The PLL for the system clock generation can be shared with the PMU and only additional LR is required for DVFS power management. Therefore, the proposed scheme can reduce both of the area and power of CoCoLo. The proposed PMU, CoCoLo can remove the off-chip inductor. Another scheme, SLIN, is proposed to reduce on-chip capacitor without degradation of supply voltage variation.

A voltage controlled oscillator (VCO) is inserted in the PLL to model the critical path of target block. And the control voltage and clock frequency of the VCO can track the MEP of target system. The clock frequency and supply voltage are co-locked by the feedback loop of the same PLL. Consequently, the proposed scheme can save area and power for DVFS multi-domain power management without off-chip elements.

C. Programmable VCO

The VCO in Fig. 4 simulate critical path of target block for the MEP tracking. In the multi-domains power management, there are several target blocks. Therefore, the different VCO is needed for each target domain. It means that each target domain needs a specific CoCoLo. The programmable VCO can solve this problem. It digitally controls the length of the critical path up to 50MHz at 1.8V supply voltage. With the programmable VCO, the same PMU can be re-used at each power domain if it exists in coverage range of programmable VCO.

D. Selectively Low Impedance Node (SLIN)

In Fig. 6, the node 1 is biased at slightly bigger than 0.9V. Voltage drop at node X pull down the node 1 voltage under 0.9V. And then, the node 2 voltage goes to 'HIGH', and node 3 drops to 'LOW' to turn on the power PMOS quickly. This process has 0.5nsec response time for stabilize variable supply voltage. The SLIN reduces large amount of capacitor for supply voltage stabilization. With 1.2nF on-chip capacitor and SLIN, the LR can be 160mV load regulation at 100mA load current.

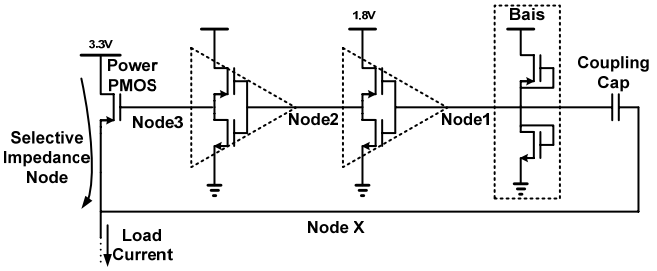

Fig. 6 The Selectively Low Impedance Node (SLIN) circuit

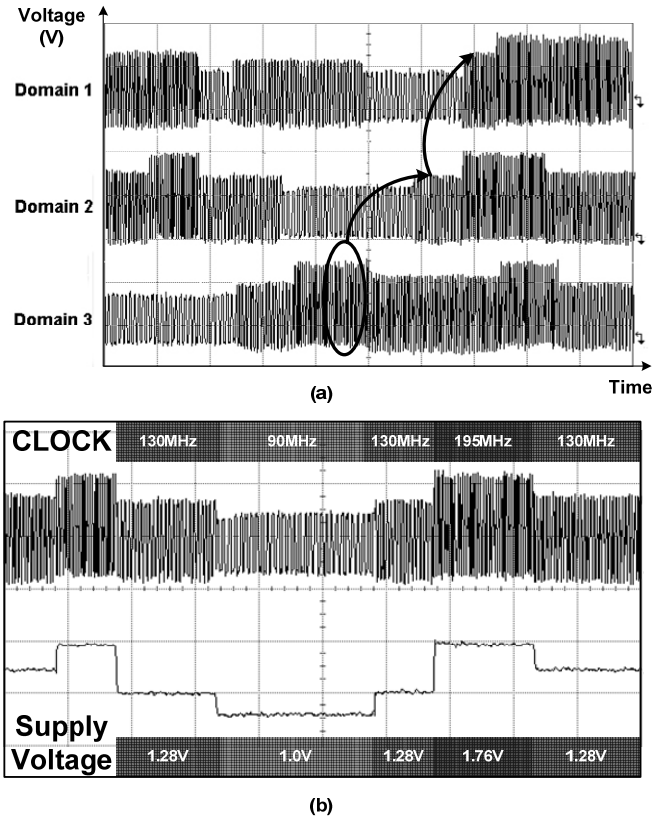

Fig. 7 (a) The clock of 3-domain with DVFS power management
(b) The clock and supply voltage of second- domain

IV. Measurement Results

The PMU and the target system are fabricated by using 0.18um CMOS process, and the CoCoLo is takes 0.45mm^2 area. The target system, which is 3D graphics processing unit (GPU), divided by three parts required different performance each other [7].

A. Multi-domains DVFS power management

Fig. 8 shows chip photo of total system. Three CoCoLos placed on corner of chip provide supply voltage and clock to each power domain for DVFS power management. Fig. 7 (a) shows measurement results of the triple domains DVFS power management, the waveforms of each domain represents clock frequency and supply voltage. The clock frequency and supply voltage of the CoCoLos are dynamically scaled according to runtime work load. As you shown in Fig. 7 (a), the speedup of the third domain increases the workload and the operating speed of the second domain. With the chain reaction, the PMU3 increases the voltage and frequency of the third domain.

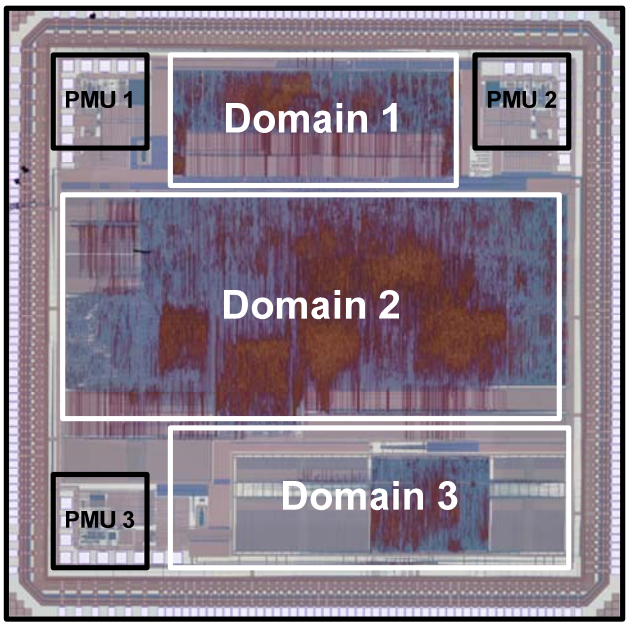

Fig. 8 Chip micro-photograph

B. *Dynamically Scaled Frequency and Voltage*

Fig. 7 (b) shows measured waveform of the output supply voltage and the clock. The peak to peak voltage of clock represents output of the LR for the MEP tracking VCO, and supply voltage is generated by the LR for the target domain, respectively. Because the MEP tracking VOC simulates the critical path of target system, these two voltage level are same, and dynamically vary at the same time as you shown in Fig. 7(b). Also, the clock frequency of CoCoLo is dynamically scaled with the supply voltage for tracking MEP of the target domain.

In case of second-domain, the range of the supply voltage covers from 1.0V to 1.8V and clock frequency 90MHz to 200MHz, respectively. The CoCoLo has 21.6mV/usec slew rate according to the varying time of second-domain full range transition of supply voltage.

TABLE 1
Overall Characteristic of CoCoLo

Measurement factors	Value
Voltage Range[V]	1.0 – 1.8
Programmable Range[MHz]	150 – 200 (@ 1.8V)
Slew Rate of transition[mV/us]	21.6 (1.0V – 1.8V)
Load Regulation[mV]	160 (I_L = 100mA)
Phase Noise[dBc/Hz]	- 100 (@ 500kHz offset f)
Power Consumption[mW]	5.0
AREA[mm^2]	0.45

Total GPU system consumes only 52.4mW with triple domains power management and about 153mW without power management and every domain operates at their full speed. In these two cases, the GPU produces almost same output results. Finally, this triple domains power management scheme achieves about 65% power reduction.

V. Conclusion

We propose dynamic voltage and frequency scaling scheme for multi-domains power management. The three power management units used this scheme are fully integrated on-chip to apply DVFS for 3D graphics processing unit (GPU) which has triple independent power domains. With this triple domains DVFS power management, the GPU consumes only 52.4mW at running benchmark test. In PMU, the LR removes off-chip inductor and the SLIN circuit reduces size of decoupling capacitor migrated into on-chip. It generate continuous co-locking the supply voltage and clock frequency, the supply voltage can be dynamically varies 1.0V to 1.8V within 37usec. TABLE 1 shows the overall characteristics of the CoCoLo. It is used as the Intellectual Property (IP) with the programmable VCO. The PMU successfully applies DVFS for triple domains power management of GPU, with 5mW power consumption and 0.45mm^2 areas.

References

[1] A. P. Chandrakasan, S. Sheng, and R. W. Brodersen, "Low-power CMOS digital design," *IEEE J. Solid-State Circuits*, vol. 27, pp. 473-484, Apr. 1992.

[2] D. Liu and C. Svensson, "Trading speed for low power by choice of supply and threshold voltages," *IEEE J. Solid-State Circuits*, vol. 28, pp. 10-17, Jan. 1993.

[3] T. Kuroda and K. Suzuki et al, "Variable Supply-Voltage Scheme for Low-Power High-Speed CMOS Digital Design," *IEEE J. Solid-State Circuits*, vol. 33, pp. 454-462, Mar. 1998.

[4] Thomas D. Burd and Trevor A. Pering, et al, "A Dynamic Voltage Scaled Microprocessor System," *IEEE J. Solid-State Circuits*, vol. 35, pp. 1571-1580, Nov. 2000.

[5] T. Kuroda and Mototsugn Hamada, "Low-Power CMOS Digital Design with Dual Embedded Adaptive Power Supplies," *IEEE J. Solid-State Circuits*, vol. 35, Apr. 2000.

[6] T. Fujiyoshi, et al., "An H.264/MPEG-4 Audio/Visual Codec LSI with Module-Wise Dynamic Voltage/Frequency Scaling," *IEEE International Solid-State Circuits Conf. Fig. Tech. Papers*, Feb. 2005

[7] Byeong-Gyu Nam, et al, "A 52.4mW 3D Graphics Processor with 141Mvertices/s Vertex Shader and 3 Power Domains of Dynamic Voltage and Frequency Scaling", *ISSCC 2007*, session. 15, 15.5 pp. 278-279, Feb. 2007.

[8] J. Kim and M. Horowitz, "An efficient digital sliding controller for adaptive power-supply regulation," *IEEE J. Solid-State Circuits*, vol. 37, pp. 639-647, May. 2002.

[9] Peter Hazucha et al, "Area-Efficient Linear Regulator With Ultra-Fast Load Regulation," *IEEE J. Solid-State Circuits*, vol. 40, pp. 933-939, Apr. 2005.

[10] G. A Ricon-Mora and P.A Allen, "A low-voltage, low-quiescent current, low drop-out regulator," *IEEE J. Solid-State Circuite*, vol. 33, pp 36-44, Jan 1988.