Staggered Core Activation: A Circuit/Architectural Approach for Mitigating Resonant Supply Noise Issues in Multi-core Multi-power Domain Processors

Ayan Paul, *Matt Amrein, Saket Gupta, Arvind Vinod, Abhishek Arun, Sachin Sapatnekar and Chris H. Kim

*Milwaukee School of Engineering, Milwaukee, WI 53202 USA
University of Minnesota, Minneapolis, MN 55455 USA

Abstract- **In order to reduce the impact of resonant supply noise on processor performance, a simple, fully-digital and scalable technique based on staggering the activation time of the cores sharing the same power domain in a multi-core multi-power domain processor is presented. Measurement data from a 65nm test chip shows an Fmax improvement as large as 20% in a 3-core configuration. This is one of the first approaches to utilize the architecture level behavior for mitigating resonant noise issues in a multi-core multi-power domain processor.**

I. INTRODUCTION

Resonant supply noise which resides in the frequency range 40MHz-300MHz [1-2] and is formed between the package inductance and on-chip capacitance, has been identified as the one of the major performance limiting factors of modern low voltage processors (Fig. 1) [3]. A number of circuit level techniques including active decoupling circuits and clock/data compensation schemes have been recently proposed to mitigate the resonant noise impact on circuit performance [4-7]. However, only few attempts have been made to utilize the architectural level behavior of multi-core, multi-power domain processors which could be more effective, simpler, fully digital, and more scalable than previous brute-force circuit approaches. In this work, we propose a circuit/architectural approach to reduce the resonant supply noise in a multi-power domain processor by staggering the activation of the cores sharing the same power domain.
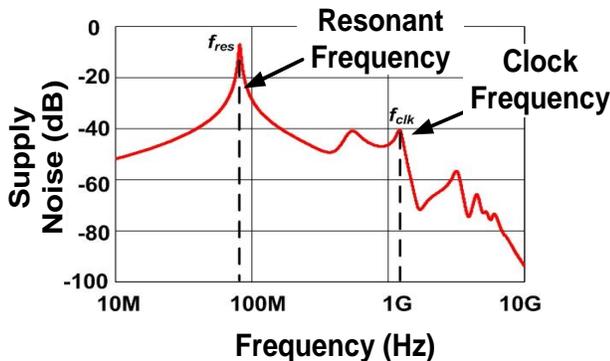

Fig. 1: Supply noise amplitude response showing frequency resonant peak at ~100MHz [3].

II. PROPOSED STAGGERED CORE ACTIVATION SCHEME

Fig. 2, top shows 3 different configurations for multi-core processors, of which in the leftmost configuration all cores are in a single voltage domain; in the middle configuration each core has its own separate power supply; and finally the rightmost configuration has a number of power domains each of which is shared by a number of cores. Our approach is applicable within a single power domain of a multi-core multi-power domain processor, where all the cores share same power supply.
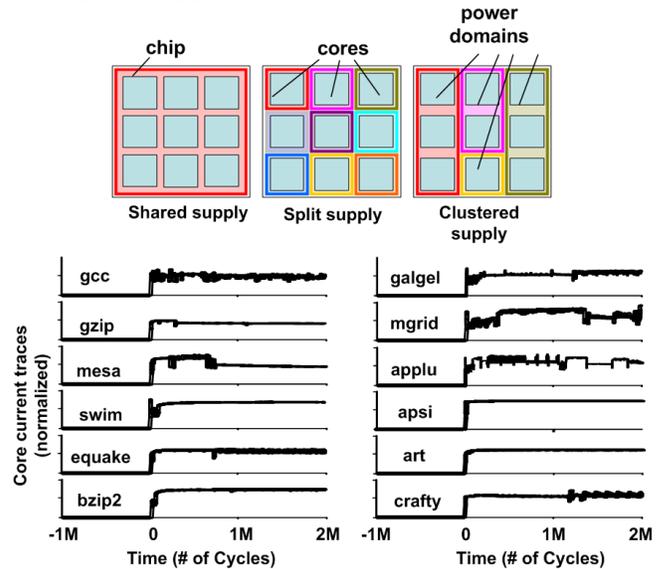

Fig. 2: Representation of multi-core multi-power domain processors. (top). Normalized core current traces for various benchmarks (bottom).

Typical workloads of a number of SPEC2000 benchmarks [8-10] are shown in Fig. 2, bottom. All benchmark programs show a step function-like jump at the onset of running an application, and similarly an abrupt step down will occur upon their completion. Resonant noise excited by the step increase in the supply current at the instant of core wake-up creates damping sinusoid like ripples in the VDD and ground network. The amplitude of the step depends on the amount of current drawn by a single core and on the number of cores going active at the same time. It determines amplitude of the damping sinusoid noise. Turn-on delay dependent resonant noise interactions of two cores are graphically represented in Fig. 3.

In case of a 2-core processor, if the turn-on times of the two cores are apart by several hundreds of cycle (Fig. 3, top-left),

the resultant noise is identical to the noise generated from a single core. However, as both the cores' turn-on times come closer, the noise waveforms start interacting with each other (Fig. 3, mid-left), which reinforces or diminishes noise depending on the time difference between the turn-on times of the two cores. As one would expect, worst-case supply noise corresponds to the case when both the cores turn on simultaneously (Fig. 3, bottom-left). As we sweep the turn-on delay of two cores and plot the minimum VDD (VDD_MIN), we find a number of valleys and plateaus, corresponding to large noise and small noise, respectively. Our proposed Staggered Core Activation (SCA) scheme to reduce resonant noise is based on staggering the turn-on time of the latter core(s). A core requesting activation first will turn on immediately. However, latter core's activation request and its actual turn-on event may differ in time. If the latter core's activation request time corresponds to the valleys of the VDD_MIN vs. delay waveform, we stagger the turn-on event of the latter core until the next plateau. This ensures minimum supply noise in a 2-core processor.
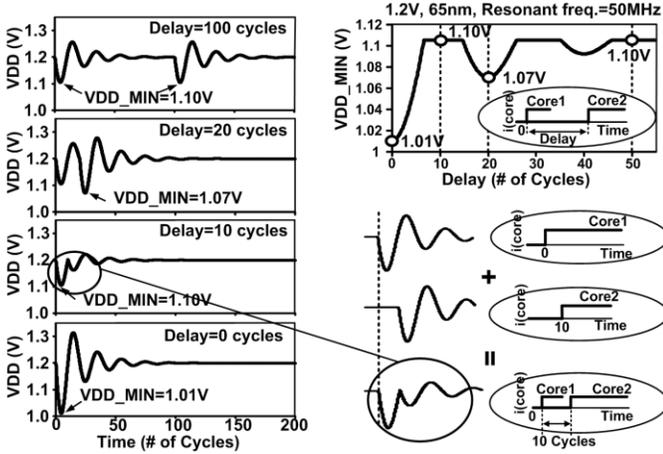

Fig. 3: Graphical representation of stagger core activation scheme in a 2-core configuration.

SCA can be extended to more than 2-core processor scenarios without significant modification. For an arbitrary N-core processor, core requesting activation first will turn on immediately. The core whose request comes second should wait until the plateau due to first core's turn-on, as discussed in 2-core scenario. The core coming third in requesting activation should wait until the plateau due to second core's turn-on, and so on. Although for the absolute lowest possible supply noise, a core's stagger time may depend on all previous cores' turn-on times, its implementation will be too complicated while the effectiveness will diminish quickly.

III. TEST CHIP IMPLEMENTATION

A 65nm test chip was designed to verify the effectiveness of the proposed SCA scheme. The block diagram is shown in Fig. 4. The test block consists of 16 simple cores that can be activated at different times. Since the wake-up transition

current of a core exhibits a step-like function, we use simple replica circuits such as individual devices and random logic blocks whose current profiles resemble that of typical processor cores as shown in Fig. 2.
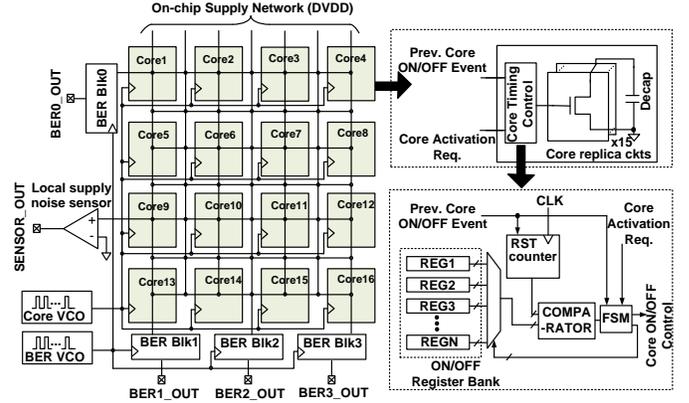

Fig. 4: Block diagram of the 65nm test chip.

Each core has a precise timing control circuit and several noise injection blocks. The register bank inside the control unit stores the timing information corresponding to the beginning and the end points of the plateaus of the VDD_MIN vs. delay waveform. Other core's ON/OFF event resets the counter of control unit and with current core's activation/deactivation request, counter output is compared with the proper register value and current core's ON/OFF control signal is sent out. A number of noise injection blocks inside the core give us the flexibility in terms of injecting noise of different amplitudes. These noise injection blocks are clocked by an on-chip VCO.
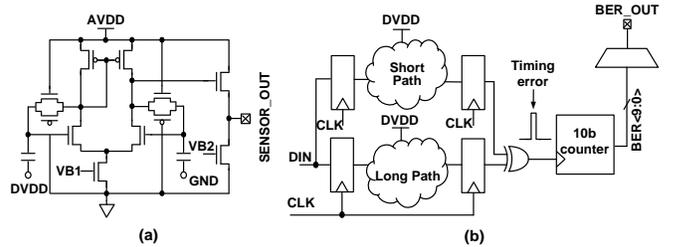

Fig. 5: (a) Local supply noise sensor and (b) BER measurement block.

There is a local analog noise sensor for direct measurement of the local supply noise as shown in Fig. 5(a). It takes the noisy supply and ground signals as the differential inputs, and its output indicates the on-chip supply noise frequency and amplitude [4]. The frequency response of the noise sensor can be characterized using external high frequency inputs and a spectrum analyzer.

To measure the Fmax improvement using SCA, four on-chip Bit Error Rate (BER) measurement blocks (Fig. 5(b)) with a dedicated VCO were added. The BER blocks count the number of errors produced as a result of supply noise by comparing the outputs of a very short logic and interconnect path and a long logic and interconnect path. By measuring the average period of the 10-bit ripple counter output and the

VCO frequency, BER can be readily calculated. Changing the frequency that the BER block runs at will allow us to determine the Fmax for different time delay schemes. Here, without loss of generality we assume Fmax to be the frequency at which BER=10^{-6} as shown in Fig. 6.
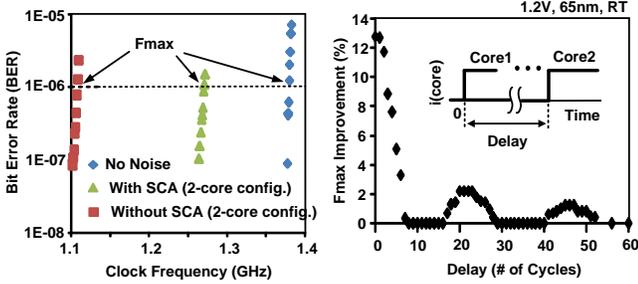

Fig. 6: Measured BER versus clock frequency (left). Measured percentage Fmax improvement with SCA versus delay in a 2-core configuration (right).

IV. TEST CHIP MEASUREMENTS

Fig. 7 shows the measurement results of the SCA for a 2-core processor. Timing sequences for core turn-on events are shown in Fig. 7(a). Without SCA, delay vs. Fmax shows certain valleys and plateaus (Fig. 7(c)). Clearly, Fmax is worst (1.10GHz) when both cores turn on simultaneously. However, in the plateaus, Fmax increases to 1.24GHz. SCA staggers second core activation until the plateau regions starts, thereby ensuring 1.24GHz of Fmax independent of the activation request time of the second core (Fig. 7(d)).

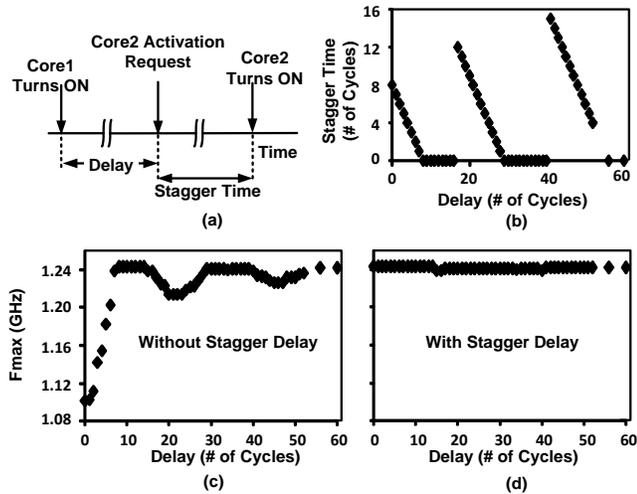Fig. 7: (a) Core activation time sequences in a 2-Core Processor. (b) Measured stagger time versus delay. Measured Fmax versus delay in a 2-core configuration (c) without SCA and (d) with SCA.

Measured data shows that for certain delays stagger time is zero indicating that those delays themselves correspond to the plateaus and additional staggering is unnecessary for those delays (Fig. 7(b)). Stagger time overhead for our measurements was found to be in the order of tens of cycles which is negligible compared to the typical number of cycles ($>10^9$) of a benchmark program.
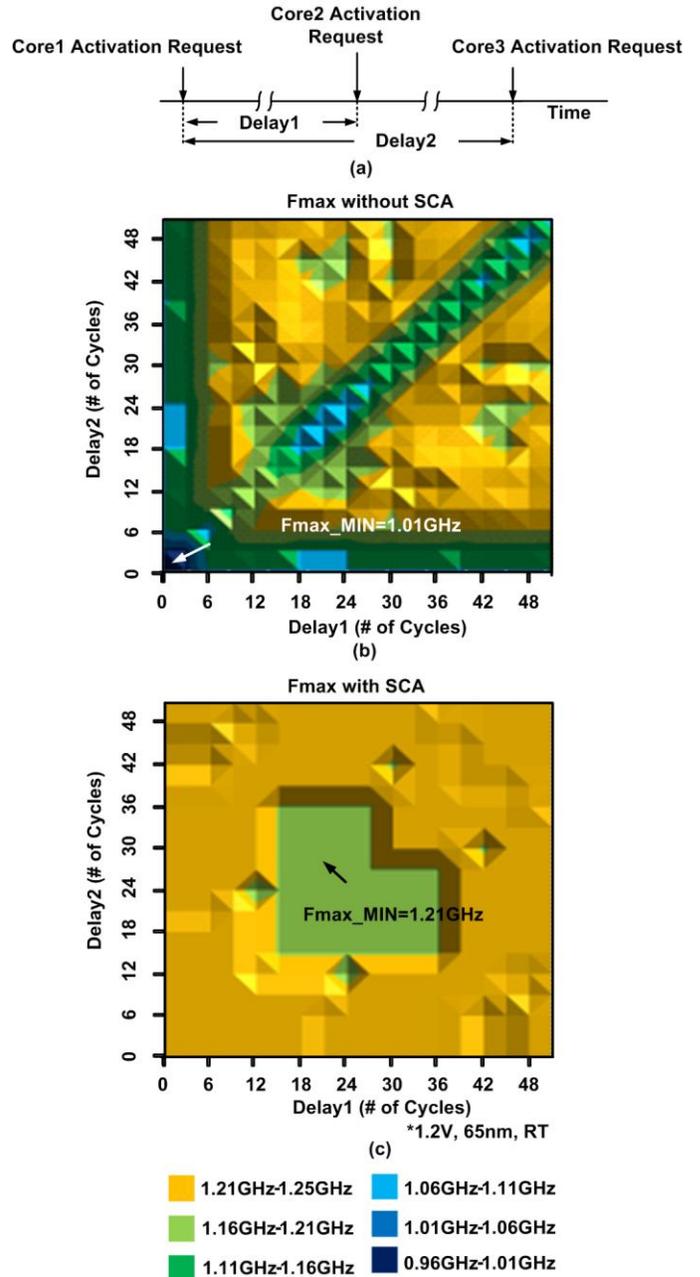

Fig. 8: (a) Core activation time sequences in a 3-core processor. Measured Fmax versus delay1 and delay2 in a 3-core configuration (b) without SCA and (c) with SCA.

Fig. 8 shows measurement results of SCA in a 3-core processor configuration. Timing sequences for core turn-on events are illustrated in Fig. 8(a). Fmax improves from 1.01GHz when all three cores turn on at the same time (Fig. 8(b)), to 1.21GHz with SCA (Fig. 8(c)). Comparison of the contour plots of with and without SCA indicates significant increase in Fmax in the former for different activation request time of core2 and core3. For a processor with more number of cores (N>3), Fmax without SCA is worse because of even larger current during simultaneous all-core activation, and with SCA we get more percentage improvement in Fmax.
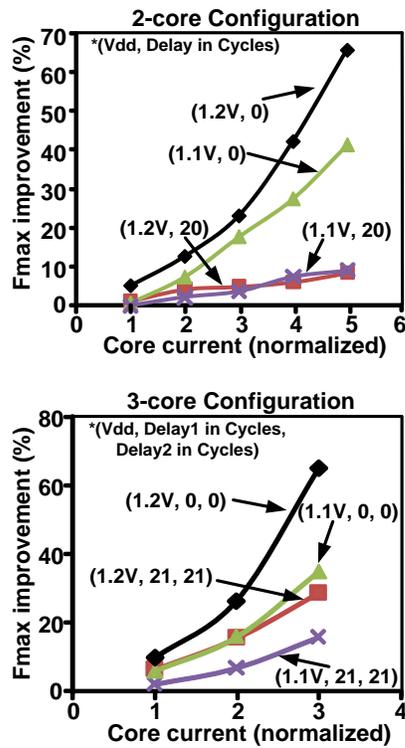
Fig. 9: Measured percentage Fmax improvement with SCA versus normalized core current as a function of VDD.

Fig. 9 shows Fmax improvement with core current as a function of VDD, for a 2-core and a 3-core configuration. For a fixed VDD, percentage improvement in Fmax increases with increasing current as one would expect. For a fixed core current, as VDD goes low, core frequency reduces, and percentage improvement in Fmax reduces.

Test chip microphotograph and sample measurement waveforms are shown in Fig. 10 and Fig. 11, respectively.
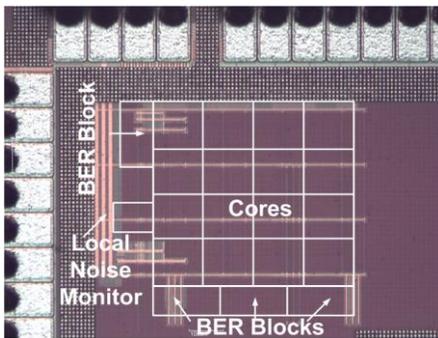

Fig. 10: Test chip microphotograph

V. CONCLUSION

We have presented a novel circuit/architectural approach based on staggering activation of cores in a multi-core multi-power domain processor that reduces the resonant supply noise. A 1.2V, 65nm test chip demonstrates approximately

13% and 20% performance improvement in Fmax for a 2-core processor and a 3-core processor, respectively. We have also shown that Fmax improvement using the proposed staggered core activation scheme is larger as the amplitude of the current drawn from the supply increases. This approach can be further extended to the case of many-core processors, where we expect to see a further improvement in terms of Fmax.
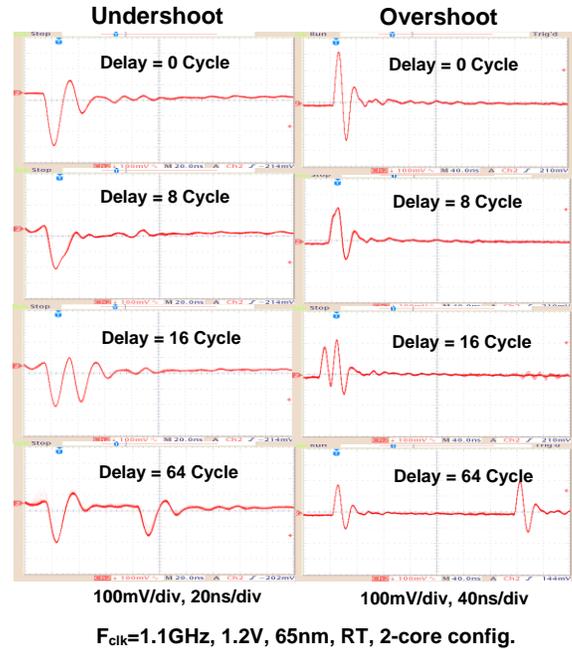

Fig. 11: Example supply noise undershoots (left column) and overshoots (right column) as a function of delay in a 2-core configuration measured using the on-chip supply sensor.

REFERENCES

[1] S. Pant and E. Chiprout, "Power Grid Physics and Implications for CAD," in *Proc. Design Automation Conf.*, pp. 199-204, Jul. 2007.

[2] E. Hailu, D. Boerstler, K. Miki, *et al.*, "A Circuit for Reducing Large Transient Current Effects on Processor Power Grids," in *Proc. Int. Solid-State Circuits Conf.*, pp. 2238-2245, Feb. 2006.

[3] J. Gu, H. Eom, and C.H. Kim, "A Switched Decoupling Capacitor Circuit for On-Chip Supply Resonance Damping," presented at *VLSI Circuits Symp.*, Jun. 2007.

[4] J. Gu, H. Eom, and C.H. Kim, "On-chip Supply Noise Regulation Using a Low Power Digital Switched Decoupling Capacitor Circuit," *J. Solid-State Circuits*, vol. 44, no. 6, pp. 1765-1775, Jun. 2009.

[5] J. Xu, P. Hazucha, M. Huang, P. Aseron, et al., "On-Die Supply-Resonance Suppression Using Band-Limited Active Damping," in *Proc. Int. Solid-State Circuits Conf.*, pp. 286-603, Feb. 2007.

[6] M. Ang, R. Salem, and A. Taylor, "An on-chip voltage regulator using switched decoupling capacitors," in *Proc. Int. Solid-State Circuits Conf.*, pp. 438–439, Feb. 2000.

[7] K. L. Wong, T. Rahal-Arabi, M. Ma, and T. Greg, "Enhancing microprocessor immunity to power supply noise with clock/data compensation," *J. Solid-State Circuits*, vol. 41, no. 4, pp. 749–758, Apr. 2006.

[8] D. Brooks, V. Tiwari, and M. Martonosi, "Wattch: a framework for architectural-level power analysis and optimizations," in *Proc. Int. Symp. on Computer Architecture*, pp. 83-94, Jun. 2000.

[9] AJ KleinOsowski, and D. J. Lilja, "MinneSPEC: A new SPEC benchmark workload for simulation-based computer architecture research," *Computer Architecture Letters*, vol. 1, pp. 7–7, 2002.

[10] SPEC 2000: http://www.spec.org/cpu2000/