

Modeling and Analysis of Power Supply Noise Tolerance with Fine-grained GALS Adaptive Clocks

Divya Akella Kamakshi*, Matthew Fojtik[†], Brucek Khailany[†], Sudhir Kudva[†], Yaping Zhou[†],
Benton H. Calhoun*

*University of Virginia, Charlottesville [†]NVIDIA Corporation

dka5ns@virginia.edu, mfojtik@nvidia.com, bkkhailany@nvidia.com, skudva@nvidia.com, yapingz@nvidia.com,
bcalhoun@virginia.edu

Abstract— Power supply noise can significantly degrade circuit performance in modern high-performance SoCs. Adaptive clocking schemes have been proposed recently that can tolerate power supply noise by adjusting the clock frequency in response to fast-changing voltage variations. In this paper, we model and quantify power supply noise tolerance with a fine-grained globally asynchronous locally synchronous (GALS) design style together with an adaptive clocking scheme. An experimental setup that includes SPICE and Verilog-A models is used to quantify the effect of clock-tree insertion delay and spatial workload variations on power supply noise tolerance in both traditional synchronous adaptive clocking and a fine-grained GALS adaptive clocking scheme. Compared to the traditional scheme, fine-grained GALS adaptive clocking significantly reduces these effects and the margins required to tolerate power supply noise. The gain is quantified using the uncompensated voltage noise metric, which is defined as the additional voltage margin that is required for failure-free operation of circuits at the frequency dictated by the adaptive clocking scheme. In our experimental setup for a typical high performance SoC, fine-grained GALS adaptive clocking achieves a 78 mV saving in uncompensated voltage noise, which is an equivalent of 15% savings in power.

Keywords—Globally asynchronous locally synchronous (GALS), Power supply noise, Adaptive clocking

I. INTRODUCTION

Power supply noise is a major challenge in modern high performance SoC designs. The power delivery network across multiple stages of regulator, motherboard, package, and on-chip parasitics with the model shown in Fig. 1 causes the supply voltage to fluctuate in response to load current switching. The supply voltage noise is classified into IR drop (static and dynamic) and inductive Ldi/dt droop. Ldi/dt droop is categorized into different types: the third droop with a duration of a few microseconds that is affected by the bulk capacitors, the second droop that is caused by the interaction of the board and package and lasts a few hundred nanoseconds, and the first droop with a duration of a few nanoseconds that is determined by the package inductance and on-die capacitance. The first droop is the shortest and usually the deepest. These power supply noise effects can cause circuit timing errors and severe performance degradation.

Additional margins are allocated in timing closure to overcome these performance failures in the presence of




Figure 1. A typical power delivery network model.

power supply noise. It is also challenging to determine a realistic, yet worst-case scenario for the power supply noise, which subsequently makes estimation of the required margins difficult. Recently, different schemes such as adaptive clocking and reactive clocking have been used in SoCs [1][2][3] to compensate for supply noise by scaling performance in response to noise events. As illustrated in Fig. 2, in a fixed clock scheme, the maximum possible frequency is set to tolerate the worst possible supply noise event. On the other hand, an adaptive clocking scheme improves performance by tracking supply noise and potentially attaining a higher average system clock frequency.

Although adaptive clocking helps reduce margin due to voltage noise, it cannot fully compensate for it due to spatial voltage variation and other local effects. In this paper, we analyze the benefits of combining adaptive clocking with fine-grained GALS to further reduce voltage margin by compensating for local effects. A traditional GALS design




Figure 2. In a fixed clocking scheme, the maximum frequency is set so that it can tolerate worst-case supply noise events. In an adaptive clocking scheme, the system frequency dynamically tracks voltage fluctuations and enables improved performance.

involves synchronous islands operating with local clocks and communicating asynchronously with each other [4]. Recent commercial designs such as [5][6] typically have one clock per core. Despite adopting GALS design methodologies to mitigate design effort, timing closure in these synchronous cores continues to be a challenge with growing chip sizes.

In this paper, we propose the benefits of a fine-grained GALS design methodology in which the SoC is partitioned into a large number of synchronous sub-blocks, each spanning no more than a few mm^2 of physical area and operating with its own local clock. Thereby, the clock generation would include hundreds of locally generated clocks. This is illustrated in Fig. 3. Local clock generation in hundreds of synchronous blocks is a major challenge in adopting the fine-grained GALS methodology. Adaptive clocks generated using replica path circuits [7][8] eliminate the need for PLLs, are of low overhead and are good candidates for local clock generation. For seamless asynchronous boundary crossings, [9] describes a low-latency (1.34 cycles average) bisynchronous FIFO design using pausable clocking techniques and can be integrated easily into standard tool flows.

In this paper, we analyze and quantify the benefits of the combination of adaptive clocking with fine-grained GALS to further reduce voltage margin and compare the results to traditional synchronous adaptive clocking approach. Another contribution of the paper is the analysis technique and models used to achieve this result. In Section II, we motivate the advantages of fine-grained GALS adaptive clocks. In Section III, we describe the models and the experimental setup used in our analysis. Section IV presents the simulation results and we conclude the paper in Section V.

II. BACKGROUND

A traditional synchronous clocking scheme distributes a single clock source over a large amount of silicon area, often spanning tens of mm^2 . In such a scheme, an adaptive clock source can dramatically reduce margin due to global voltage noise at low noise frequencies. However, not all the margin can be compensated for. In this section we describe the effects of clock-tree insertion delay and spatial workload




Figure 3. On the left is the baseline, which is a traditional synchronous adaptive clock design. Each core has one adaptive clock generator, represented by a black dot. On the right is the proposed fine-grained GALS adaptive clocking approach, in which synchronous sub-blocks span no more than a few mm^2 and have their own local adaptive clocks.

variations in large synchronous islands that are not effectively compensated for by a traditional synchronous adaptive clocking scheme.

A. Effect of Clock-tree Insertion Delay

Clock-tree insertion delay is the delay between the root of the clock at the clock generator to the leaf of the clock at the flip-flops in load circuits. During a supply noise event, the adaptive clock generator responds with a change in frequency. However, it takes a time equivalent to the clock-tree insertion delay, for the stretched clock pulses to reach the load circuits. This effect is illustrated in Fig. 4.

In prior clocking designs, when the clock generator used is a fixed-frequency PLL, increasing insertion delay is a technique implemented to give enough time for the clock-tree to adapt to supply noise [10]. In this paper, we use a noise-aware or adaptive clocking methodology to tolerate supply noise [11]. An adaptive clock generator itself tracks the critical path delays by design and adjusts the frequency with supply noise. It also does not involve the design effort required to build clock-trees such as in [10].

If the load circuits, clock generator, and the clock-tree sense the same voltage and track similarly, then the adaptive clocking system would perfectly tolerate supply noise. However, clock-trees do not track as well as desired since they are often wire dominated paths and designed for better slew rates unlike digital logic. Moreover, higher the clock-tree insertion delay, the longer it takes for the load circuits to see the stretched clock pulses. On the other hand, the load circuits instantaneously sense any fluctuations in the supply voltage. Hence, there is a time difference between the change in the supply voltage and the change in the operating frequency as seen by the load circuits. In modern SoCs, the clock-tree insertion delay is in the range of one to a few clock cycles (1-2 ns). This effect of clock-tree insertion delay requires additional margin for failure-free circuit operation, which cannot be completely eliminated with the use of traditional synchronous adaptive clocks.

B. Effect of Spatial Workload Variations

In large synchronous islands, spatial variations in the switching current can cause voltage fluctuations of different

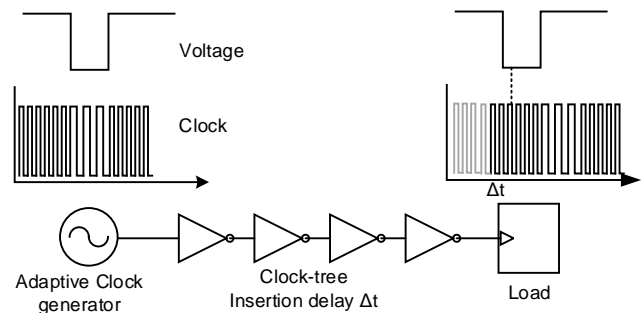


Figure 4. Illustration of the effect of insertion delay: The stretched clock is delayed by Δt , equal to the clock-tree insertion delay. This requires some margin for failure-free operation, which cannot be completely eliminated with the use of traditional synchronous adaptive clocks.

characteristics in different sections of the block. For instance, a local supply droop near the load circuits can cause them to slow down, but the adaptive clock source may be sensing a local supply voltage overshoot that increases the clock frequency. This effect illustrated in Fig. 5 is one such example of the consequence of spatial workload variations that can cause a timing failure. The traditional synchronous adaptive clocking scheme does not compensate for this effect. Additional margin is required to guarantee correct operation due to voltage noise from spatial workload variations.

C. Benefits of Fine-grained GALS Adaptive Clocks

In a fine-grained GALS adaptive clocking scheme, the design is composed of myriad synchronous islands of areas as small as a mm^2 . The insertion delay for such small synchronous islands is only a few hundred picoseconds. Moreover, due to the close proximity of the clock generator and load circuits, they also tend to experience similar voltage fluctuations. Therefore, fine-grained GALS adaptive clocks can reduce the effect of insertion delay and spatial workload variations and the additional margins associated with them as illustrated in Fig. 6. In the following section, we describe the

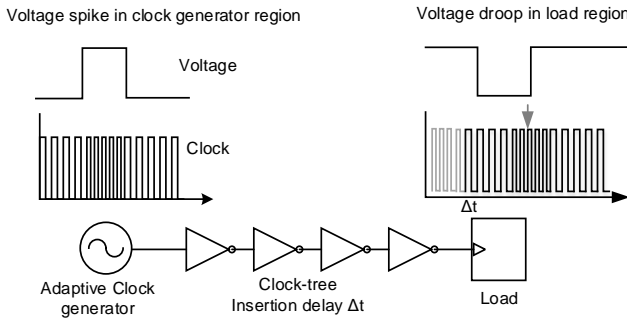


Figure 5. Effect of spatial workload variation: The region of the SoC holding the clock generator experiences a supply overshoot and the load circuits experiences a local voltage droop due to high switching activity. The voltage rail fluctuating in the opposite directions can potentially cause circuit failure. Additional margins are required to compensate for such local noise effects.

Both clock generator and load region experience similar switching activities

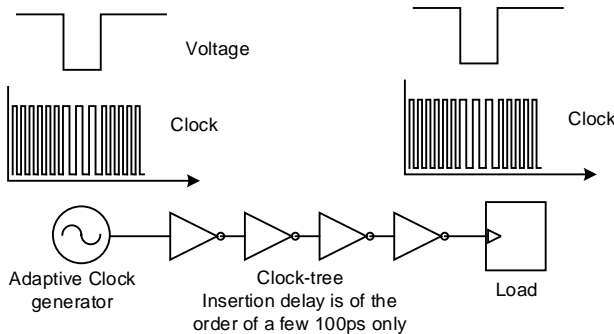


Figure 6. Illustration of benefits of fine-grained GALS: The insertion delay is a few hundred picoseconds which reduces the effect of insertion delay. Due to close proximity of the local clock generators and the rest of the logic, the supply noise experienced by the logic is similar to its clock source, thereby reducing the effect of spatial workload variations.

experimental setup that is used to estimate the margins required in traditional synchronous adaptive clocks and the potential savings using fine-grained GALS adaptive clocks.

III. EXPERIMENTAL SETUP

A system level diagram is shown in Fig. 7. The different components of the experimental setup are described in detail in this section. The power distribution analysis is done using a publicly available power distribution network (PDN) analysis tool, Voltspot [12]. The setup for deriving the adaptive clocking model was obtained from [13]. This experimental setup can be used for the analysis of any microprocessor system. In this paper, we focus on high performance, high-end GPUs that have large areas and high power consumptions [14].

A. Power Distribution Network

A simple power distribution network (PDN) is represented using a lumped model as shown in Fig. 8. The PCB values are similar to the digital processor model developed in [15] and used in GPU supply noise characterizations [16]. We estimated the package and bump resistance and inductance for a die area ($23.5 \text{ mm} \times 23.5 \text{ mm}$) typical of large SoCs or GPUs such as [14]. For our analysis, we assume the presence of decoupling capacitors on the backside of the PCB, but none on the topside of or buried in the package. The on-chip resistance was estimated by running a steady-state IR drop analysis using Voltspot [12] for our nominal supply voltage of 1 V. The PDN impedance curve is shown in Fig. 9. The first droop resonant frequency occurs at 30 MHz.

Although this existing simple lumped PDN model would be sufficient to analyze power-supply noise effects that are common to the whole chip, a distributed power distribution

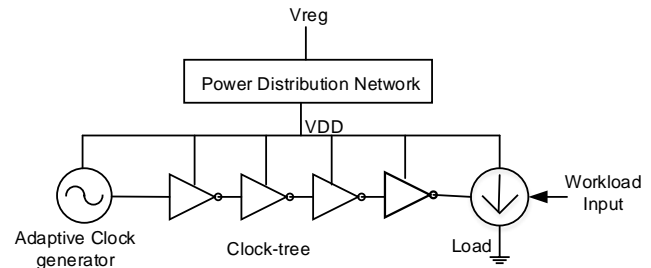


Figure 7. Experimental setup showing the different components.




Figure 8. Lumped power distribution network model.

network model is required to analyze the spatial effect of workload variations and the power-supply noise tolerance of a fine-grained GALS system. For this purpose, we use Voltspot [12] to simulate a distributed PDN with a granularity of the pad pitch. By default, Voltspot captures an on-chip distributed grid (including the dynamic effects of R, L, and C), but assumes a lumped package and ignores PCB parasitics.

However, a distributed package model in addition to on-chip PDN model is also required for an accurate representation of spatial variation of transient supply noise. We modified Voltspot settings to emulate a distributed package and on-chip PDN distribution as shown in Fig. 10. Further, we also factored in PCB resistance as it affects the magnitude of first droop impedance. The input parameters to Voltspot tool were derived from the lumped model in Fig. 8. Therefore, we expect that the resonant frequency of the Voltspot model also occurs at 30MHz, which we confirm from the power supply noise results in Section IV.

We divide the PDN of the chip of area into a 47 x 47 array of units of ~ 0.5 mm x 0.5 mm. We use Voltspot for




Figure 9. Power distribution network impedance: The resonant frequency of the first droop is 30 MHz.




Figure 10. Distributed power distribution network model.

consistency in both the insertion delay and workload variation analyses. A uniform current distribution is assumed across the 47 x 47 PDN array for the effect of insertion delay analysis and a non-uniform current distribution for workload variations. These details are discussed in Section IV.

B. Adaptive Clock Generator

In [7], an aggressive adaptive clock model is proposed which tracks changes in the supply voltage within each clock cycle. In this paper, we model the adaptive clock generator to respond to voltage variations with single cycle latency. The experimental setup for adaptive clocking modeling is described below and was derived from [13].

The load rail voltage is averaged over the previous cycle to get V_{mean} . The frequency of the current cycle is determined by V_{mean} and a VF curve. The VF curve was obtained by modeling a critical path including low, regular and high-threshold voltage (V_t) devices using the FreePDK 45nm kit. The data and clock path models are shown in Fig. 11. The circuit was simulated for maximum frequency (F_{max}) at different voltages and insertion delays. The global clock path drives long wires and hence includes both device and significant wire delays. In our model for the global clock path, each clock inverter drives a 200 μm long wire. The device and wire parasitics for the data and clock path are extracted from the layout. Fig. 12 shows the clock frequency adapting to the supply voltage.

C. Clock-tree

The clock-tree model includes the global clock distribution and the local clock distribution. The clock-tree delay variation with supply voltage is obtained for different cases of insertion delay. It is assumed that the entire clock-tree sees the same supply voltage. A plot of delay versus supply voltage (V) for a clock tree insertion delay of ~1.2 ns at a nominal supply voltage of 1 V is shown in Fig. 13. Similar trends are obtained for insertion delays 1.5 ns, 0.9 ns, 0.6 ns and 0.3 ns. A curve fit polynomial for each insertion delay is obtained from these plots and used in the Verilog-A model for clock-tree. For instance, the polynomial for 1.2 ns insertion delay is:

$$\begin{aligned} \text{Insertion Delay (ns)} = & -22.9 \times V^6 + 130.3 \times V^5 - 301.7 \times V^4 \\ & + 361.9 \times V^3 - 232.8 \times V^2 + 73.1 \times V - 6.6 \end{aligned}$$

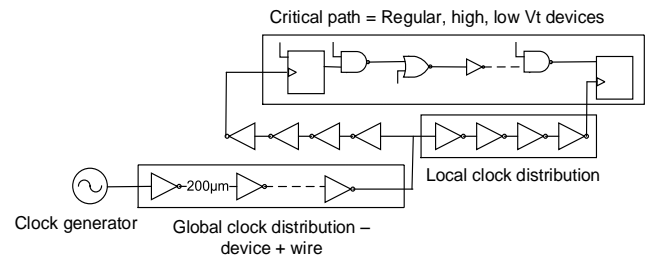


Figure 11. Clock and data path modeling to get Voltage-Frequency relationship.




Figure 12. Clock frequency adapting to supply voltage: V_{mean} is the average of load voltage over the previous cycle, Clock is the adaptive clock source output. The output frequency of the adaptive clock source is also shown.

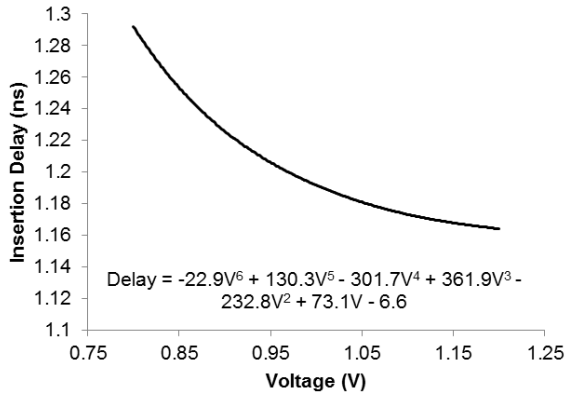


Figure 13. Clock-tree insertion delay versus supply voltage.

D. Workload

Switching activity in the workload causes the supply voltage rail to fluctuate. In [15], the current profile characteristics for SPEC benchmarks were broadly classified into three categories: step, pulse and resonating currents. We observed that resonating currents, which are associated with repeating activity patterns, have the worst effect on supply noise as compared to others. When the frequency of the currents is close to the resonant frequency of the PDN, the voltage fluctuations are at their worst.

In [17], it was shown that in a 7.1 billion transistor GK110 GPU implemented in a 28 nm technology, the dynamic current may switch from 10A to 130A with a very fast slew of up to 15 clock periods in computing applications. Our frequency of interest is the resonant frequency of ~ 30 MHz and the operational frequency in our analysis is 850 MHz. Therefore, we choose a 10 cycle slew rate from 10 A to 90 A in our analysis and sweep the

workload frequency up to 40 MHz. The nominal supply voltage is 1 V. A resonating workload waveform is shown in Fig. 14.

E. Metric: Uncompensated Voltage Noise

V_{mean} is defined as the average voltage at the load circuits over a clock cycle. The clock frequency at the load circuits, $freq_{load}$, is impacted by many factors such as the clock-tree insertion delay and spatial voltage variations at the clock generator compared to the load circuits. V_{req} is defined as the actual voltage that is required for operation of the load circuits at frequency, $freq_{load}$. *Uncompensated voltage noise* is the difference between V_{mean} and V_{req} (when $V_{mean} < V_{req}$, the available voltage is less than the required voltage), which is the additional margin that is required to keep the load circuits operational without failure at frequency, $freq_{load}$. A lower uncompensated voltage noise implies power savings.

IV. SIMULATION RESULTS

In our simulation setup, the PDN area is divided into a 47×47 array of units as shown in Fig. 15, each unit of size $0.5 \text{ mm} \times 0.5 \text{ mm}$. To model current distribution for a large die, as shown in Fig. 15, we group the 2209 units into nine partitions: three 12×20 and one 11×20 partitions along the top, three 12×20 and one 11×20 partitions along the bottom, and one 47×7 partition along the middle of the die.

A. Effect of Clock-tree Insertion Delay

To analyze the effect of clock-tree insertion delay, a resonating workload current as shown in Fig. 14 is uniformly applied throughout the PDN area. We expect the worst case supply noise to occur at 30 MHz resonance frequency of the PDN as indicated by Fig. 9. As shown in Fig. 16, uncompensated voltage noise is measured by sweeping workload frequencies near the resonant frequency for several designs, each with different insertion delays. Uncompensated voltage noise is indicative of the additional margin required for failure-free operation of the circuit.

As expected, the worst-case uncompensated voltage noise occurs at higher insertion delays at workload frequency of 30 MHz. With a non-adaptive clocking scheme, the margin required for failure-free operation is $\sim 175 \text{ mV}$. With the traditional synchronous adaptive clocking scheme with a clock-tree insertion delay of 1.5 ns, the uncompensated voltage noise at 30 MHz workload




Figure 14. Resonating workload current profile.




Figure 15. PDN area is divided into a 47 x 47 array of units. Each unit is 0.5 mm x 0.5 mm.

frequency is 49 mV. Therefore, the adaptive clocking scheme considerably reduces the voltage margin as compared to the non-adaptive clocking scheme. At 300 ps insertion delay, which is typical of fine-grained GALS, the uncompensated voltage noise is further reduced to 8 mV. Therefore, fine-grained GALS adaptive clock mitigates the effect of clock-tree insertion delay and further reduces the voltage margin.

B. Effect of Spatial Workload Variations

The fine-grained GALS adaptive clocking scheme with an insertion delay of ~ 0.3 ns is compared to the baseline traditional adaptive clocking scheme with long (1.5 ns) insertion delay. In this section, we analyze the effect of spatial voltage variations in addition to the effect of insertion delays. For the analysis of effect of spatial workload variation on power supply noise, we consider one partition containing 12x20 PDN units as shown in Fig. 17. This is the baseline design that is equivalent to a large synchronous island in systems with a traditional synchronous clocking scheme. Each partition is operated using their own local clock.

Fig. 17(a) represents the traditional clocking scheme in

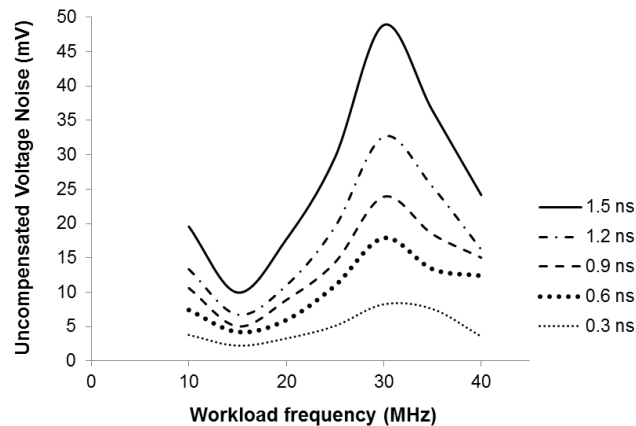


Figure 16. Uncompensated voltage noise vs. workload frequencies for different insertion delays. As insertion delay increases, the uncompensated voltage noise increases.

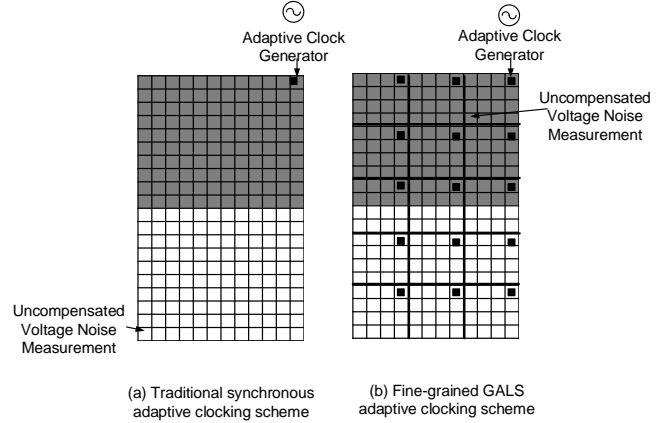


Figure 17. Effect of spatial workload variations: (a) In traditional clocking scheme setup, the uncompensated voltage noise is measured at a unit farthest away from the clock unit. (b) In fine-grained GALS adaptive clocking scheme, a clock generator is present in every 2 mm x 2 mm partition.

which the unit containing the clock generator (clock unit) is assumed to be in the top right corner. The uncompensated voltage noise is measured in the lower-left corner unit that is farthest away from the clock unit for worst-case measurements. This is called the measurement unit. Fig. 17(b) represents the fine-grained GALS clocking scheme. We assume a 2 mm x 2 mm size for each GALS block with its own local clock.

In this paper, we present two examples of workload variations to demonstrate the mitigation of uncompensated voltage noise by fine-grained GALS adaptive clocks.

1) *Case 1:* All units of the block consume equal current. A workload switching frequency of 30 MHz is chosen to generate a worst-case supply noise scenario. However, the workload current switches at a phase difference of 180° between the lower and the upper half of the block as shown in Fig. 18(a). The remaining seven blocks of the SoC are idle. Fig. 18(b) shows supply noise variation between the clock and the measurement units. Idle units in close proximity to the clock unit cause it to have lower supply noise than the measurement unit. For a design with an insertion delay of 1.5 ns, an uncompensated voltage noise of 59 mV is measured. Fig. 18(c) shows that there is very little supply noise difference between clock and measurement units in a fine-grained GALS clocking design. In this case, uncompensated voltage noise is 5 mV. Therefore, the voltage margin savings provided by fine-grained GALS scheme in this case is ~ 54 mV.

2) *Case 2:* For a higher supply noise variation, the lower half units of the partition is assumed to consume 80% of the total block power and units in the top half consume 20% of the total block power. The current profile of the block is shown in Fig. 19(a). The remaining seven blocks of the SoC are idle. Workload switching frequency of 30 MHz is chosen for the worst-case supply noise scenario.

As expected, the supply noise is higher in the measurement unit (lower half) than the clock unit for the traditional clock scheme as shown in a Fig. 19(b). The higher difference in supply noise gives an uncompensated

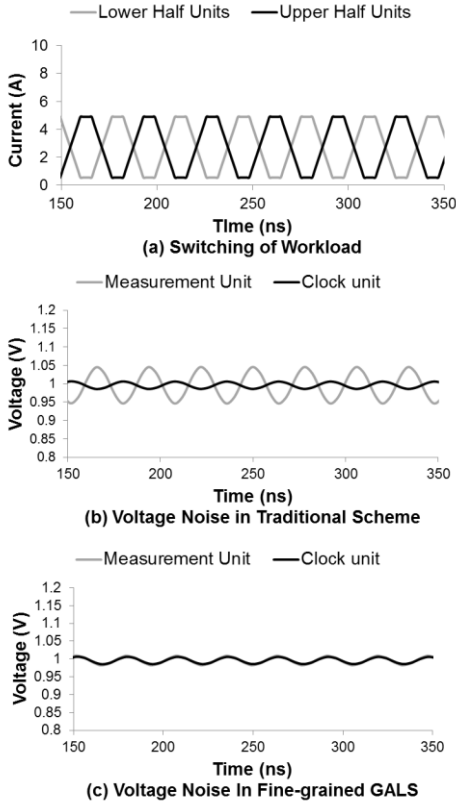


Figure 18. Case 1: Effect of spatial workload variations. (a) Current profile (b) Supply noise variation in clock unit and measurement unit in traditional clocking scheme (c) Supply noise in both clock and measurement units is similar in fine-grained GALS clocking scheme.

voltage noise of 111 mV for a design with 1.5 ns insertion delay. For a fine-grained GALS adaptive clock system, the difference in supply noise fluctuation between clock and measurement is lower as shown in Fig. 19(c). The uncompensated voltage noise in this case is 33 mV. Worse uncompensated voltage noise may be possible when the remaining blocks in the SoC are also switching. In this case, the voltage margin savings provided by the GALS clocking scheme is ~ 78 mV.

C. Savings versus Overhead

Significant savings in uncompensated voltage noise of ~ 78 mV is possible with the use of fine-grained GALS adaptive clocks. At a nominal operating voltage of 1 V, this is equivalent to a power savings of $\sim 15\%$ for the same performance of the SoC. However this involves a design with myriad local clocks. In the floorplan used in our analysis, the number of local clock generators is 138. The ideal candidates for local clocks generators are digitally controlled oscillators that generate clock frequency using simple critical-path replica circuits driven off the core power supply [7][8]. Such circuits consume only a few milliwatts of power [18] and we expect the area and power overheads to be small ($<1\%$).

The overall savings are dependent on the GALS partition size. The lower the dimensions of the GALS block, the

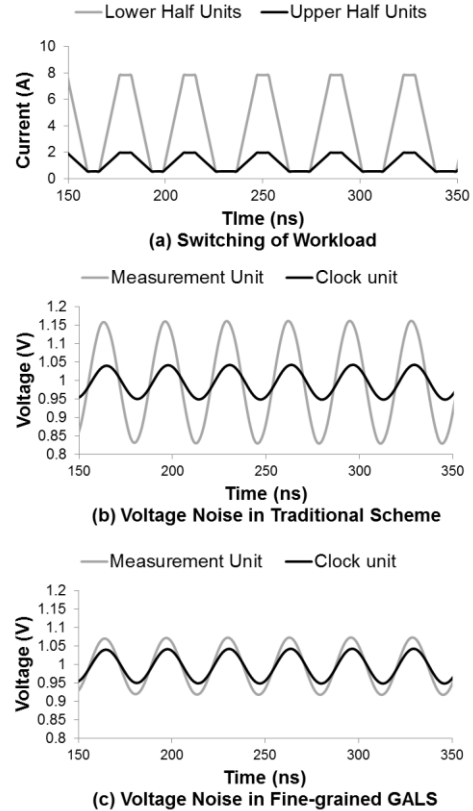


Figure 19. Case 2: Spatial workload variation. (a) Current profile (b) Supply noise variation in clock unit and measurement unit in traditional clocking scheme (c) Supply noise in fine-grained GALS clocking scheme

lower the uncompensated voltage noise. However, the power and area overheads of the clock generators, the latency of the clock domain crossing circuits and implementation constraints of the logic functionality also play a role in the actual partitioning. As a part of future work, the proposed modeling methodology can be further utilized to quantify the overheads of the clock generators and clock domain crossing circuits and to derive the optimal partition size.

V. CONCLUSION

With increasing area of SoCs, traditional synchronous adaptive clocking schemes alone cannot fully compensate for power supply noise. In this paper, we analyzed power supply noise tolerance using the fine-grained GALS adaptive clocking scheme. We observed that this scheme mitigates local effects of supply noise due to clock-tree insertion delay and spatial workload variations. We developed models and an analysis technique to quantify these benefits of fine-grained GALS adaptive clocks as compared to the traditional synchronous adaptive clocking scheme. From our experimental setup for high-performance processors, we obtained a 78 mV savings in uncompensated voltage noise at 1 V supply voltage, with the use of fine-grained GALS adaptive clocking scheme compared to the traditional synchronous adaptive clocking scheme. This is

an equivalent of 15% savings in power consumption for the same processor performance.

ACKNOWLEDGMENT

This work was funded in part by DARPA through a subcontract from NVIDIA and in part by the NSF NERC ASSIST Center (EEC-1160483). This research was, in part, funded by the U.S. Government. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government. The authors would also like to thank Dr. Ke Wang, research scientist with the Department of Computer Science, University of Virginia for his valuable guidance on Voltspot.

REFERENCES

- [1] A. Grenat, S. Pant, R. Rachala, and S. Naffziger, "5.6 Adaptive clocking system for improved power efficiency in a 28nm x86-64 microprocessor," in Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2014 IEEE International, vol., no., pp.106-107, 9-13 Feb. 2014
- [2] N. Kurd, P. Mosalikanti, M. Neidengard, J. Douglas, and R. Kumar, "Next Generation Intel Core™ Micro-Architecture (Nehalem) Clocking," in Solid-State Circuits, IEEE Journal of , vol.44, no.4, pp.1121-1129, April 2009
- [3] J. Cortadella et al., "Reactive clocks with variability-tracking jitter," Computer Design (ICCD), 2015 33rd IEEE International Conference on, New York, NY, 2015, pp. 511-518.
- [4] D. M. Chapiro, "Globally-asynchronous locally-synchronous systems," Ph. D. Thesis, vol. 1, p. 50, 1984.
- [5] E. Fluhr et al., "Power8: A 12-core server-class processor in 22nm soi with 7.6tb/s off-chip bandwidth," in Proc. IEEE International Solid State Circuits Conference, 2014, pp. 96-97.
- [6] S. Rusu et al., "Ivytown: A 22nm 15-core enterprise Xeon processor family," in Proc. IEEE International Solid-State Circuits Conference, 2014, pp. 102-103.
- [7] R. Jevtic, Hanh-Phuc Le, M. Blagojevic, S. Bailey, K. Asanovic, E. Alon, and B. Nikolic, "Per-Core DVFS With Switched-Capacitor Converters for Energy Efficiency in Manycore Processors," in Very Large Scale Integration (VLSI) Systems, IEEE Transactions on , vol.23, no.4, pp.723-730, April 2015
- [8] R. Mullins and S. Moore, "Demystifying Data-Driven and Pausible Clocking Schemes," in Proc. IEEE Symposium on Asynchronous Circuits and Systems, 2007, pp. 175-185.
- [9] B. Keller, M. Fojtik, and B. Khailany, "A Pausible Bisynchronous FIFO for GALS Systems," in Asynchronous Circuits and Systems (ASYNC), 2015 21st IEEE International Symposium on , vol., no., pp.1-8, 4-6 May 2015.
- [10] K. A. Bowman, C. Tokunaga, T. Karnik, V.K. De, and J. W. Tschanz, "A 22 nm All-Digital Dynamically Adaptive Clock Distribution for Supply Voltage Droop Tolerance," in Solid-State Circuits, IEEE Journal of , vol.48, no.4, pp.907-916, April 2013
- [11] M. S. Floyd, A. J. Drake, R. W. Berry, H. Chase, R. Willaman, and J. Pena, "Voltage droop reduction using throttling controlled by timing margin feedback," in VLSI Circuits (VLSIC), 2012 Symposium on , vol., no., pp.96-97, 13-15 June 2012
- [12] R. Zhang, K. Wang, B. H. Meyer, M. R. Stan, and K. Skadron, "Architecture implications of pads as a scarce resource," in Computer Architecture (ISCA), 2014 ACM/IEEE 41st International Symposium on , vol., no., pp.373-384, 14-18 June 2014
- [13] Y. Zhou, S. Sudhakaran, A. Naik, X. Chang, D. Lin, T. Raja, and S. Idgunji, "Modeling and measurement of noise aware clocking in power supply noise analysis," in Electrical Performance of Electronic Packaging and Systems (EPEPS), 2014 IEEE 23rd Conference on , vol., no., pp.7-10, 26-29 Oct. 2014
- [14] Whitepaper, NVIDIA's Next Generation CUDA Compute Architecture: Kepler GK110; <http://www.nvidia.com/content/PDF/kepler/NVIDIAKepler-GK110-Architecture-Whitepaper.pdf>
- [15] M. S. Gupta, J. L. Oatley, R. Joseph, Gu-Yeon Wei, and D. M. Brooks, "Understanding Voltage Variations in Chip Multiprocessors using a Distributed Power-Delivery Network," in Design, Automation & Test in Europe Conference & Exhibition, 2007. DATE '07 , vol., no., pp.1-6, 16-20 April 2007
- [16] J. Leng, Y. Zu, M. Rhu, M.S. Gupta, and V. J. Reddi, "GPUVolt: Modeling and characterizing voltage noise in GPU architectures," in Low Power Electronics and Design (ISLPED), 2014 IEEE/ACM International Symposium on , vol., no., pp.141-146, 11-13 Aug. 2014
- [17] W. Mao, Y. Zhou, A. Naik, and S. Sudhakaran, "Study of BGA package cap for high-performance computing GPU," in Electromagnetic Compatibility (EMC), 2013 IEEE International Symposium on , vol., no., pp.858-862, 5-9 Aug. 2013
- [18] J. Zhao and Y.-Bin Kim, "A Low-Power Digitally Controlled Oscillator for All Digital Phase-Locked Loops," VLSI Design, vol. 2010, Article ID 946710, 11 pages, 2010