

Active Power-Gating-Induced Power/Ground Noise Alleviation Using Parasitic Capacitance of On-Chip Memories

Xuan Wang[†], Jiang Xu[†], Wei Zhang[‡], Xiaowen Wu[†],
Yaoyao Ye[†], Zhehui Wang[†], Mahdi Nikdast[†], Zhe Wang[†]

[†]Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology

[‡]School of Computer Engineering, Nanyang Technological University, Singapore

Email: {eexwang,jiang.xu}@ust.hk, zhangwei@ntu.edu.sg

Abstract—By integrating multiple processing units and memories on a single chip, multiprocessor system-on-chip (MPSoC) can provide higher performance per energy and lower cost per function to applications with growing complexity. In order to maintain the power budget, power gating technique is widely used to reduce the leakage power. However, it will introduce significant power/ground (P/G) noises, and threat the reliability of MPSoCs. With significant area, power and performance overheads, traditional methods rely on reinforced circuits or fixed protection strategies to reduce P/G noises caused by power gating. In this paper, we propose a systematic approach to actively alleviating P/G noises using the parasitic capacitance of on-chip memories through sensor network on-chip (SENoC). We utilize the parasitic capacitance of on-chip memories as dynamic decoupling capacitance to suppress P/G noises and develop a detailed Hspice model for related study. SENoC is developed to not only monitor and report P/G noises but also coordinate processing units and memories to alleviate such transient threats at run time. Extensive evaluations show that compared with traditional methods, our approach saves 11.7% to 62.2% energy consumption and achieves 13.3% to 69.3% performance improvement for different applications and MPSoCs with different scales. We implement the circuit details of our approach and show its low area and energy consumption overheads.

I. INTRODUCTION

Multiprocessor system-on-chip (MPSoC) becomes promising to satisfy the growing computation demands by integrating multiple processing units (PUs) and memories on a single chip. In order to maintain the power budget, power gating technique is widely used to reduce the leakage power, by shutting off idle blocks from power supplies. However, the power gated block, e.g. a PU, will inevitably introduce power/ground (P/G) noise [1] [2]. When it is powered on, a large rush current will charge the parasitic capacitance through the power delivery networks, and cause severe P/G noise, e.g. a voltage drop across the parasitic impedance of the networks. The P/G noise will spread and lower the supply voltages of the surrounding PUs. The low supply voltage may lead to failures in the working PUs, e.g. timing violations. All these make power gating induced P/G noise a severe reliability threat in MPSoCs.

This work is partially supported by RGC GRF620911.

Traditional methods rely on reinforced circuits or fixed protection strategies to reduce the P/G noise caused by power gateings. These methods introduce significant area, power and performance overhead. Decoupling capacitance (decap) insertion technique is a prevalent method based on reinforced circuits [3]. The decap is an additional capacitor, which is connected to the power delivery network and placed beside the power gating blocks. When the block is turned on, the decap will discharge and compensate the induced rush current. It reduces the current drawn from the network to decrease the voltage drop. However, the decap will introduce significant area and leakage power overhead in advanced technologies.

In this paper, we propose a systematic approach to actively alleviating the power gating induced P/G noise through sensor network on-chip (SENoC). An active P/G noise alleviation technique is proposed to reduce the P/G noise with negligible area and power overhead, by utilizing the parasitic capacitance of on-chip memories. A control system, SENoC, is developed to coordinate PUs and memories to further improve the system performance by selecting energy efficient power gateings.

II. POWER GATING INDUCED P/G NOISE ANALYSIS AND ACTIVE P/G NOISE ALLEVIATION

Decap insertion method is widely used to reduce the P/G noise, however, the area and leakage power overhead severely limit the size of the capacitance and its effectiveness. It is observed that the charged parasitic capacitance of the idle circuits is able to work as a costless decap and suppress the P/G noise [2], but this feature has not been paid much attention, because most of the previous works evaluate power gateings in a single processor. In this paper, we differentiate

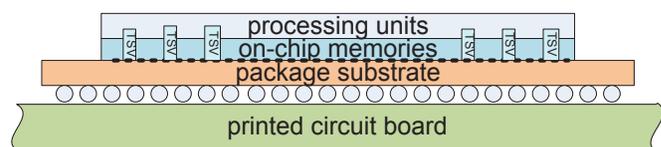


Fig. 1: Cross section of a die with 3D integration technology

between the PUs and on-chip memories because of their different capacitance densities and different magnitudes of induced P/G noise, and try to utilize the large parasitic capacitance of memories to alleviate the P/G noise of PUs with negligible overhead. 3D integration MPSoC provides an opportunity to build an architecture with separated PUs and memories, in which PUs and caches are manufactured in different stacks with the same area [4]. An example of a two-stack 3D integration is shown in Fig. 1.

A. Modeling of power gating induced P/G noise in MPSoC

The P/G noise generation and propagation vary with different power gatings of the PUs and caches at different locations. In order to evaluate the P/G noise, a fine grain Hspice model of a two-stack MPSoC is built. In this paper, the 3D MPSoC is homogeneous, consisting of identical pairs of PU and cache. The analysis platform targets at the 45nm technology node. The supply voltage is 0.95V according to the International Technology Roadmap for Semiconductors 2009 [5].

The power delivery network (PDN) model is shown in Fig. 2. The Hspice models of the two layers are quite similar with the mesh-based power distribution. The PDN is composed of identical sub-blocks. The sub-block consists of segment resistors and inductors of the wire segments of the VDD and GND network, and the lumped capacitance of the PU or cache per sub-block, which is connected to the PDN through head sleep transistors [6]. The capacitance of the cache is connected to two parallel sleep transistors, which are used to control the power supply, e.g. regular supply voltage, reduced supply voltage for data retention or power gated.

For some of the sub-blocks, there are additional through silicon vias (TSVs), which connect the PDNs of the two layers through some pass transistors. A pair of the pass transistors, e.g. $G1$ and $G2$, is called decap switch. It decides whether the PDNs of the PU and cache pair are directly connected. In order to protect the reliability, we make a conservative assumption that conventionally the switches are turned off, and a PU and

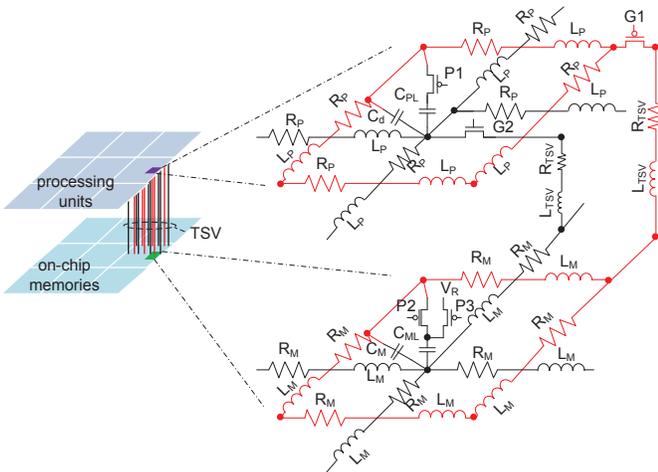


Fig. 2: modeling power delivery network of the MPSoC

cache pair will work with separated PDNs. Besides the PDN model, a conventional ladder package model is also included.

B. Power gating induced P/G noise estimation in MPSoC

After the Hspice model is established, we are able to estimate the P/G noise propagation and distribution. A powering on PU is defined as a PU attacker, and a powering on cache, a waking up cache or a sleeping cache is defined as a cache attacker, because of the noise generation. Particular noise magnitudes are derived to formulate the relationship between the P/G noise and the reliability. The maximum P/G noise magnitude that a PU or cache can tolerate for correct execution is estimated as both 100mV [6] [7]. That of preventing the occurrence of the data upset in the memory cells is 600mV [8]. When an attacker occurs, the unsafe PU or cache is defined as a victim within the impact range of the attacker, and the set of the victims is defined as its impact range.

The distribution of the peak voltage drops during the power gating of different PU attackers is shown in Fig. 3. We assume the side length of a PU is 1, and the average impact range is the PUs within the distance of $4\sqrt{2}$. The impact range of powering on a cache is about the caches within the distance of 5. While, the impact range of waking up or sleeping a cache is that of $\sqrt{2}$. The small voltage change significantly reduces the charge needed to be restored in the parasitic capacitance, and the rush current flowing through the networks.

C. Active P/G noise alleviation using on-chip memories

During the application execution, a PU will be power gated to reduce the leakage power, however, the caches will be just slept because of the stored data. Therefore, powering on a PU will introduce significant performance penalty, due to its large impact range. Active P/G noise alleviation technique is effective to suppress the P/G noise of a PU attacker using the parasitic capacitance of the caches. The decap switches of the attacker and victims are turned on before the power gating to attenuate the coming P/G noise. After the PU and cache become safe again, the switches are turned off. Assuming PU (3, 3) is a PU attacker, by turning on the decap switches of the attacker, the impact range is significantly reduced to the PUs

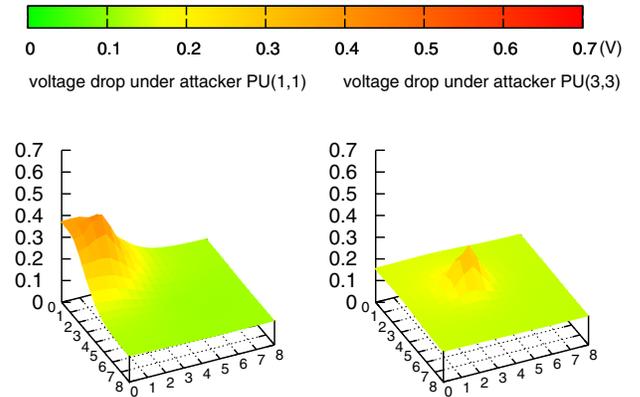


Fig. 3: Spatial distribution and impact range of the power gating induced P/G noise in a 8x8 MPSoC

within the distance of $2\sqrt{2}$. With more caches used, voltage drops can be further reduced. Moreover, the active P/G noise alleviation technique is also efficient to mitigate the noise.

First, the P/G noise and the victims are significantly reduced by using the parasitic capacitance of caches, without additional decaps. Second, we only utilize the caches of the attacker and victims. It will not degrade the performance, because originally those caches are unable to work due to the unsafe PUs. Moreover, the safe PUs faraway from the attacker will have safe caches, and will not be disturbed, because the voltage drop of the cache is much smaller than that of the PU at the same location. Third, the victims are prone to be the closer PUs, whose caches are more effective to reduce the noise.

III. SENoC BASED POWER GATING INDUCED P/G NOISE ALLEVIATION

The active P/G noise alleviation technique is able to suppress the P/G noise effectively and efficiently for PU attackers. However, some system-level decisions, e.g. power gating selection, cannot be determined without run-time information. Therefore, a system-level strategy is needed to collect the information, analyze it and make correct decisions. SENoC is a control system prototype based on MPSoCs, which is composed of on-chip sensors, node agents, and a central controller (CC). It is integrated with network-on-chip (NoC) and uses NoC as the communication medium. Fig. 4 illustrates the SENoC for a 4×4 MPSoC with a mesh-based NoC. A node agent is an interface between the CC and the local PU and cache. It is in charge of interpreting the information and the requirement. The CC is used to maintain run-time information of PUs, caches and tasks, and make system-level decisions, e.g. task scheduling, power gating selection and etc.

We propose a two-stage scheduling strategy to improve the performance and energy consumption under the run-time variations, e.g. power gatings. The strategy is a composition of an offline task mapping and scheduling and a light-weight online dynamic adjustment. At design time, we decide the task mapping and scheduling to fully utilize the parallelism of MPSoCs and optimize the performance. We use the load balancing strategy, which is one of the most popular heuristic algorithms to deal with large scale task scheduling problems.

During run time, an online dynamic adjustment is proposed

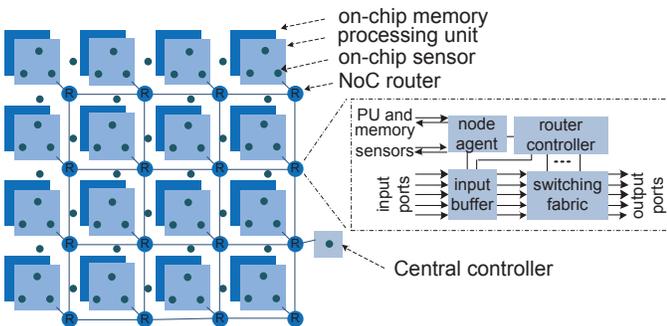


Fig. 4: Overview of SENoC architecture on a 4×4 MPSoC

to deal with the power gating penalty. In this paper, we alleviate the penalty, focusing on two issues. One is to decide whether a PU and cache pair should be power gated after a task is finished, if the next scheduled task is not ready. The other is how to efficiently utilize the active P/G noise alleviation by selecting the optimal combination to minimize victims.

A power gating selection policy is well developed based on the fact that the slack time after a task should be long enough to compensate the penalty of power gatings. First we will decide the threshold of the slack for efficient power gating by comparing the energy overhead, e.g. leakage energy of the halted victims, and energy saving, e.g. the reduced leakage energy, to estimate the break even point. Then we will predict the slack time by estimating the finish time of preceding tasks. This process will search the application graph in an adverse direction like a depth-first search of a tree. For the real time constraint, the depth is limited to 2 in this paper.

For selecting an optimized cache combination, potential victims will be evaluated one by one, according to the impact range under certain cache combinations. If the PU is unsafe, it will be protected, and its cache will be used. The combination will be updated by adding this cache for the next victim. After all the potential victims have been evaluated, the optimized combination is determined. The impact range checking for all the combinations is achieved by using a lookup table.

IV. IMPLEMENTATION AND PERFORMANCE EVALUATION

In order to evaluate our approach, at the circuit level, a well-developed Hspice model is proposed to analyze the P/G noise and the penalty. At the system level, we develop a cycle-accurate SENoC simulator using SystemC, and conduct performance evaluations on a real application benchmark.

The PDN parameters are extracted from the interconnect model of Predictive Technology Model [9]. The parameters of a PU comes from ARM11 by scaling down to 45nm, e.g. 30mW average dynamic power, 10mW leakage power and etc. That of a cache are estimated according to Synopsys 90nm Library by scaling down to 45nm, e.g. 42.6mW average dynamic power, 115.7mW leakage power and etc. The data retention voltage is assumed as 0.41V [8], and the leakage power is 39.8mW. We use 72 additional TSVs per PU with a uniform location distribution. The settling time of powering on a PU is 10ns, and that of the cache are 100ns and 87ns, according to the Hspice simulation. The time penalty of clocking on or off a PU are both 10ns.

The SENoC simulator is built as a mesh-based MPSoC architecture with XY routing and wormhole switching protocols. The clock frequency is set to 1GHz. Extensive performance evaluations are conducted with real application traffic patterns [10]. During the simulation, task executions are simulated with random execution times following Gaussian distribution.

In order to demonstrate the efficiency, we compare our approach, labeled as *proposed*, to 4 different strategies with the same task mapping and scheduling result. *w/o power gating* doesn't carry out power gatings to save the leakage power. *stopgo* is a traditional power gating protection strategy to

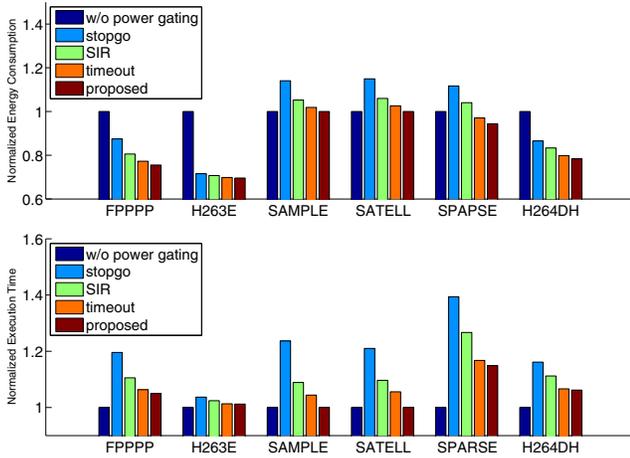


Fig. 5: Energy consumption and execution time comparison for different applications on 4x4 MPSoC

compare with, and the other two, *SIR* and *timeout*, are used to evaluate our approach in detail. For the victim protection of PUs or caches, *stopgo* method will halt all the PUs for protection. *SIR* evaluates a static impact range for both PU and cache attacker. *timeout* utilizes the proposed active P/G noise alleviation technique. For the power gating selection, those strategies uses a timeout technique [11]. A PU and cache pair will be power gated, if there is no task ready for a certain period, e.g. 1000 clock cycles in this paper.

The normalized energy consumption is defined as the energy consumption for MPSoCs with different scales divided by the energy consumption of *w/o power gating* on 4x4 MPSoCs. The definition of the normalized execution time is similar. Fig. 5 shows the simulation results for each application on 4x4 MPSoCs. Our approach has smaller energy consumption and execute time compared with *stopgo* and *SIR*, because the power gating penalty is slight due to the active P/G noise alleviation technique. It also outperforms *timeout* for all the applications, because it is able to adapt with different applications. For some applications with few energy efficient power gatings, our approach cancels the power gatings.

The scalability is also important for MPSOC design, and the scalability of the strategies is shown in Fig. 6. Compared with *w/o power gating*, our approach shows a considerable energy consumption improvement with slight performance degradation. Moreover, it provides increased improvement of energy consumption and performance compared with *timeout*, because it can dynamically adjust to the new task mapping and scheduling results for different scales.

Additional hardware is implemented to achieve the proposed techniques. The area overhead of the TSVs and decap switches is 1.97%, and the power consumption is about 0.13% for different applications and scales. The CC is synthesized and validated by Synopsys and Cadence tools in 45nm technology [12]. The area and power consumption are 2.05% and 2.21% for 4x4 MPSoC, and 1.02% and 1.05% for 8x8 MPSoC.

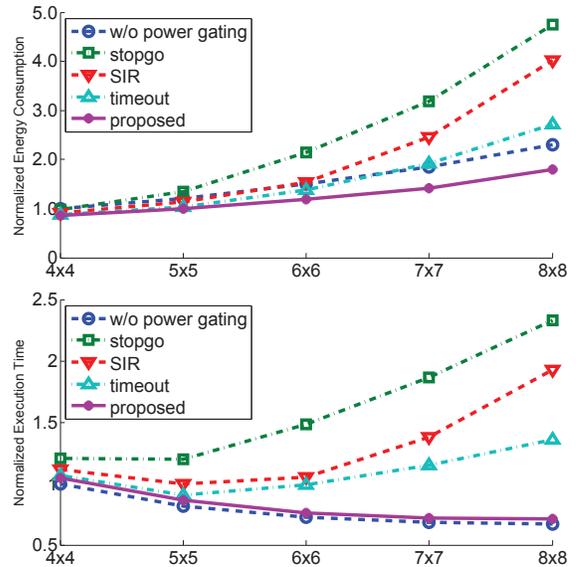


Fig. 6: Average performance improvement trend for MPSoCs with different scales

V. CONCLUSIONS

In this paper, we have proposed a systematic strategy to protect the MPSoC from the reliability threat of power gating induced P/G noise. In addition to maintaining the reliability, the power gating penalty is alleviated in both micro-architecture level and system level by using the active P/G noise alleviation technique through SENoC. It effectively reduces the leakage power with slight performance degradation compared with the case without power gatings. It also outperforms other power gating aware strategies in terms of both energy consumption and performance for MPSoCs with different scales.

REFERENCES

- [1] K. Shi and D. Howard, "Challenges in sleep transistor design and implementation in low-power designs," in *DAC, 2006 43rd ACM/IEEE*, 0-0 2006, pp. 113–116.
- [2] M. Popovich, M. Sotman *et al.*, "Effective radii of on-chip decoupling capacitors," *TVLSI*, vol. 16, no. 7, pp. 894–907, jul. 2008.
- [3] C.-Y. Yeh and M. Merek-Sadowska, "Timing-aware power-noise reduction in placement," *TCAD*, vol. 26, no. 3, pp. 527–541, mar. 2007.
- [4] D. H. Kim, K. Athikulwongse *et al.*, "3d-maps: 3d massively parallel processor with stacked memory," in *ISSCC*, feb. 2012, pp. 188–190.
- [5] "International technology roadmap for semiconductors 2009 edition," Semiconductor Industry Association, <http://public.itrs.net>.
- [6] W. Liu, Y. Wang *et al.*, "On-chip sensor network for efficient management of power gating induced power/ground noise in multiprocessor system-on-chip," *TPDS*, vol. 99, no. PrePrints, pp. 1–1, 2012.
- [7] A. Kumar, J. Rabaey *et al.*, "Sram supply voltage scaling: A reliability perspective," in *ISQED*, march 2009, pp. 782–787.
- [8] A. Nourivand, A. J. Al-Khalili *et al.*, "Postsilicon tuning of standby supply voltage in srams to reduce yield losses due to parametric data-retention failures," *TVLSI*, vol. PP, no. 99, p. 1, 2011.
- [9] "Nanoscale integration and modeling (nimo) group. predictive technology modle (ptm)," <http://ptm.asu.edu/>.
- [10] W. Liu, J. Xu *et al.*, "A noc traffic suite based on real applications," in *ISVLSI*, july 2011, pp. 66–71.
- [11] N. Madan, A. Buyuktosunoglu *et al.*, "A case for guarded power gating for multi-core processors," in *HPCA*, feb. 2011, pp. 291–300.
- [12] "Nangate 45nm open cell library," <http://www.nangate.com>.