Incorporating PVT Variations in System-level Power Exploration of On-Chip Communication Architectures

Sudeep Pasricha, Young-Hwan Park, Fadi J Kurdahi, Nikil Dutt
Center for Embedded Computer Systems
University of California, Irvine, CA
{spasrich, younghwp, kurdahi, dutt}@uci.edu

Abstract

With the shift towards deep sub-micron (DSM) technologies, the increase in leakage power and the adoption of power-aware design methodologies have resulted in potentially significant variations in power consumption under different process, voltage and temperature (PVT) corners. In this paper, we first investigate the impact of PVT corners on power consumption at the System-on-Chip (SoC) level, especially for the on-chip communication infrastructure. Given a target technology library, we then show how it is possible to "scale up" and abstract the PVT variability at the system level, allowing characterization of the PVT-aware design space early in the design flow. We conducted several experiments to estimate power for PVT corner cases, at the gate-level, as well as at the higher system-level. Our preliminary results are very interesting and indicate that: (i) there are significant variations in power consumption across PVT corners, and (ii) the PVT-aware power estimation problem may be amenable to a reasonably simple abstraction at the system-level.

1. Introduction

With the advent of the deep submicron (DSM) era, more and more System-on-Chip (SoC) designs are being fabricated in sub-100nm technologies. Unfortunately, *process, voltage and temperature (PVT) variability* makes it hard to achieve 'safe' designs in such nanometer technologies. This is because PVT variability causes fluctuation in timing as well as power for SoC designs [1]. Consequently, timing and power estimates derived early in the design flow are no longer valid, and considerable redesign effort is needed to account for these variability-induced fluctuations. Recently, many research efforts have focused on statistical timing analysis [2-3] to address variability in timing. However, till now very few efforts [4] have looked at addressing the effect of PVT variability on system-level power estimation. Since reducing power consumption is increasingly becoming the most important goal for SoC designs, especially for portable battery-driven embedded systems [5], it becomes essential to address the issue of reliable power estimation for these designs, in the face of PVT variability.

In modern IP-based design, the communication architecture backbone has become a significant factor in influencing overall system power, performance, cost and time-to-market [6]. In particular, it has been shown that for some SoC designs, on-chip communication architectures (wires and bus logic) can consume anywhere between 20-50% of overall system power [7]. The amount of power consumed in the various bus logic components is also steadily increasing with design complexity, and has been shown to be as high as 80% of the total on-chip communication power [7][17]. Furthermore, a significant portion of the on-chip communication architecture power consumption is converted into heat, which has been shown to not only increase interconnect delay (reducing performance), but also increase electro-migration (EM), which significantly increases device failure rate [8]. These observations motivate the need for system-level estimation of on-chip communication architecture power consumption early in the design flow, where design decisions have a much greater impact on power consumption than at lower levels.

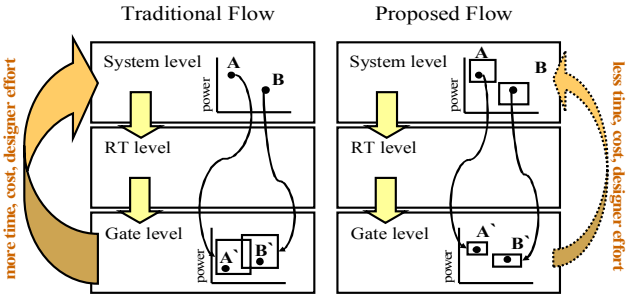

Figure 1. Traditional approach compared with proposed PVT variation-aware system-level exploration approach

While in the past, leakage was negligible and dynamic power did not vary much between technology corners, today the increase in leakage power and the adoption of power-aware design methodologies (such as voltage islands and DVS/DFS) has resulted in considerable variations in power consumption under different process, voltage and temperature (PVT) technology corners. In this paper, we first explore the impact of PVT corners on power consumption at the system level, especially for the on-chip communication architecture. We then show how the variability due to different PVT corners can be abstracted up to the system-level for the on-

1063-9667/08 $25.00 © 2008 IEEE
DOI 10.1109/VLSI.2008.14

chip communication architecture, where the corners can be explored early in the design flow. To the best of our knowledge, this is the first piece of work to incorporate PVT variations at the system-level during power exploration of the on-chip communication architecture.

Figure 1 illustrates the difference between the traditional SoC design approach, and the approach proposed in this paper. In the *traditional approach*, designers explore the power space of the design at the system-level, and select the configuration with the least power consumption. In the figure, points A and B represent design configurations, and a designer would select configuration B, with the lower power consumption at the system-level. Later in the design flow, at the gate-level, designers encounter process, voltage and temperature (PVT) variations that alter the power characteristics and behavior of the synthesized design. Each of the configuration points becomes a large region of uncertainty (representing possible power consumption under different PVT conditions), and it is no longer clear whether configuration A or B is the superior one in terms of lower power consumption. It is possible that an instance of the design configuration A (shown as A` in the figure) is found to be superior to the best instance of design configuration B (shown as B` in the figure). As a result, designers end up spending considerable time and effort to explore design configurations at the gate-level. It is also important to ensure that PVT variations do not cause a violation of design constraints for the selected design configuration. Design reiterations might be required if violations are detected (requiring changes in the design at the system-level), which can severely influence design cost and time-to-market. In contrast, in our *proposed approach*, we attempt to "scale up" and abstract the PVT variability at the system-level, to provide a more realistic characterization of the design space, early in the design flow. The designer can then select a design configuration with greater confidence, after analyzing its behavior under PVT variations. This significantly reduces the exploration and redesign effort later in the design flow. We conducted several experiments to estimate power for PVT corners of DSM technology libraries, at the gate-level as well as at the higher system-level, especially for on-chip communication architectures. Our preliminary results are very interesting and indicate that: *(i)* there are significant variations in power consumption across PVT corners, and *(ii)* the PVT-aware power estimation problem may be amenable to a reasonably simple abstraction at the system level.

2. Related Work

System-level power estimation approaches typically create power models for heterogeneous system components (e.g. buses, memories, caches, processors) and integrate them to get overall power estimates [10-12]. Several approaches have proposed power estimation techniques for bus-based communication architectures [13-17]. While early work focused mainly on power estimation for bus wires [13-14], more recent work has shown the importance of considering bus logic components as well [7]. System-level power estimation approaches for communication architectures that consider the contribution of both bus logic and wires, have been proposed for the AMBA hierarchical shared bus [15], STBus interconnection network [16] and the AMBA bus matrix [17]. None of the abovementioned power estimation approaches have studied the effects of PVT variability on power consumption at the system-level. To the best of our knowledge, our work is the first to try and understand how PVT variability affects power consumption, especially for on-chip communication architectures, and then attempt to abstract up this variability to the system-level, for early power exploration of the true design space.

Table 1. PVT Corners in UDSM Technologies

Nominal V_{dd}	Corner	Process	Temp	V_{dd}
1.0V	MaxPerf	F-F	0	1.1
	TypPerf	T-T	25	1
	WorstPerf	S-S	125	0.9
	WorstLeakage	F-F	125	1.1
	TypLeakage	T-T	125	1
0.7V	MaxPerfLowV	F-F	0	0.77
	TypPerfLowV	T-T	25	0.7
	WorstPerfLowV	S-S	125	0.7
1.2V	MaxPerfHighV	F-F	0	1.32
	TypPerfHighV	T-T	25	1.2
	WorstPerfHighV	S-S	125	1.08
	WorstLeakageHighV	F-F	125	1.32
	TypLeakageHighV	T-T	125	1.2

3. PVT Corners in Ultra-Deep Submicron (UDSM) Technologies

Traditionally, the most important means by which a foundry communicates process, voltage and temperature variations to designers is through library characterization at design corners, known as PVT corners, relating cell metrics (timing, power) to Process, Voltage and Temperature variations. Up until the 130nm technology library node, design tools relied on three corners: *Typical*, *Worst*, and *Best* corners. The adjectives associated with these corners relate mainly to timing. The *Worst* corner combines high temperature, low V_{dd} (nominal-10%), and a Slow-Slow (S-S) process that leads to worst case timing. The *Best* corner goes the opposite way, combining low temperature, high V_{dd} (nominal+10%), and Fast-Fast (F-F) process to achieve maximum performance. The *Typical* performance corner lies between these two extreme corners. Synthesis tools currently use the *Worst* and *Best* corners during synthesis, guaranteeing that all the resulting functional chips would meet timing (by using *Worst* case corner to avoid setup time violations at register inputs, and using *Best* case corner to guarantee that no hold time violations occur). The *Typical* corner is usually used to characterize power consumption under nominal conditions. Since leakage was negligible up until the 130nm technology library node, the only factor affecting power consumption (mostly dynamic) was V_{dd}, and variations of about ± 20% were expected between the three corners.

With ultra-deep submicron (UDSM) technologies under 100nm gearing into production, some changes became necessary with corner characterization. This is due to a variety of factors, the most important of which is the drastic increase in the device leakage power. Other factors include IR drop as well as power management strategies such as DVS/DFS and voltage islands. Today, IPs, especially cell libraries, I/O and memories are thus characterized at many more PVT corners, shown in Table 1. These corners become necessary for a variety of reasons: A *TypicalPerf* corner, for example, does not provide a realistic assessment of leakage power under typical conditions because when the application is running, a die would heat up to well above 25ºC. With temperature being an exponential factor in leakage, a more realistic *TypicalLeakage* corner must be considered with Typical-Typical (T-T) process, nominal V_{dd} and 125ºC. A *WorstLeakage* corner is used to assess the absolute maximum leakage under Fast-Fast (F-F) process (i.e. low V_t) and high V_{dd} (Nominal+10%).

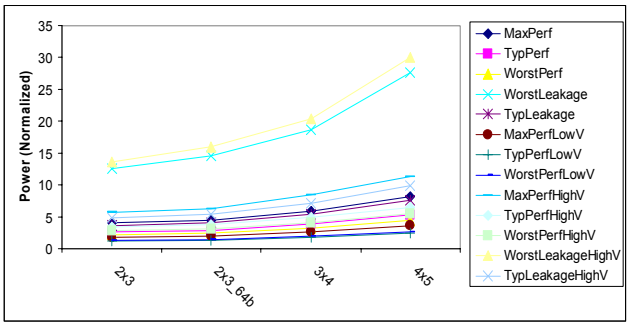

Figure 2. Normalized Power for Bus Matrix Configurations at 90nm

Power management strategies such as voltage islands and discrete voltage scaling (DVS) cannot be validated at the chip level unless IPs are characterized under several low V_{dd} conditions. This requires another set of corners. IPs such as the Metro libraries from Artisan (ARM) [19] are characterized for V_{dd} increments of 100mV for a range of possible V_{dd} values. *MaxPerfLowV*, *TypPerfLowV* and *WorstPerfLowV* are needed for each V_{dd}. For 90nm, the lowest safe V_{dd} is 0.7V. On the other hand, and under certain condition, more performance may be needed. For that case, some cell libraries are characterized for higher than normal operating conditions. In the case of 90nm, IPs can operate up to 1.2V nominal V_{dd}. Thus, another set of corners are needed including *MaxPerfHighV*, *TypPerfHighV*, and *WorstPerfHighV*. Since leakage can be significantly higher under those conditions, additional *TypicalLeakageHighV* and *WorstLeakageHighV* are sometimes available in order to assess typical and worst case leakage under high V_{dd} conditions. Note that the PVT corners shown in Table 1 do not constitute a maximal set. Many more corners can be added during library characterization, to support more elaborate design methodologies and possible operating environments and conditions. Alternatively, some technology libraries may not support *HighV* corners for reliability reasons. The existence of a multitude of these corners motivates the need to mitigate the complexity and achieve more reliable designs by understanding and incorporating PVT effects early in a design flow, at the system-level.

In the case of timing analysis, one may argue that corner characterization is of limited use. However, we note that while timing analysis concentrates on critical path characterization, power characterization introduces significantly more degrees of freedom to the analysis, such as data dependence, power management, etc., all of which are very hard to incorporate into an amenable statistical analysis, especially at the system level. Thus, employing corner based analysis with a larger and more realistic corner set helps reduce the complexity of the designers' task in exploring the design space, albeit at the cost of perhaps slightly more pessimistic assumptions.

4. Impact of PVT Variability on Power Consumption

We will now present some experimental results to show how PVT variability affects power consumption in ultra deep submicron technology nodes, for the on-chip communication architecture (Section 4.1), the entire SoC design (Section 4.2) and a multi-Vt design flow (Section 4.3). Unlike the technology nodes upto the 130nm node, the ultra deep submicron technology nodes such as the 90nm and 65nm have many more PVT corners that need to be explored during the design phase, to ensure that power constraints in a design are satisfied. The experiments in this section explore power consumption characteristics for these different PVT corners, for different ultra deep submicron technology nodes, designs and operating frequencies.

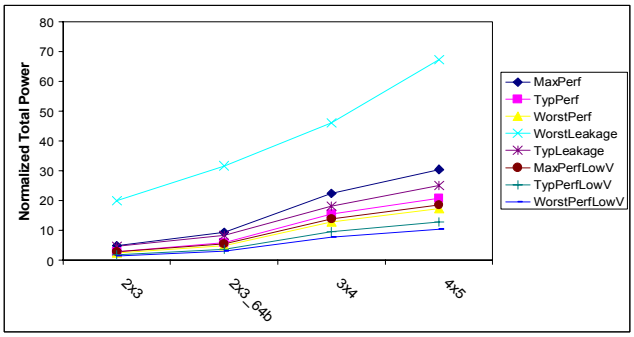

Figure 3. Normalized Power for Bus Matrix Configurations at 65nm

4.1. Impact on the On-Chip Communication Architecture

Our first set of experiments was conducted with the aim of understanding the impact of PVT variation on power consumption of the on-chip communication architecture. In the first experiment, we selected four SoC designs of varying

complexity. Each of the designs had an AMBA AHB bus matrix [18] communication architecture configuration with a different structure and traffic characteristics: *(i)* a 2 master, 3 slave bus matrix with 32 bit data bus width (2x3), *(ii)* a 3 master, 4 slave bus matrix with 32 bit data bus width (3x4), *(iii)* a 4 master, 5 slave bus matrix with 32 bit data bus width (4x5) and *(iv)* a 2 master, 3 slave bus matrix with 64 bit data width (2x3_64b). We targeted the 90nm and 65nm general purpose technology libraries, synthesized these designs for a 100 MHz bus clock frequency and estimated power for the different PVT corners shown in Table 1 using Synopsys PrimePower [20] at the gate-level.

The normalized power of the bus matrix communication architecture, for the different PVT corners is shown for all four bus matrix configurations, in Figure 2 (90nm) and Figure 3 (65nm). It can be seen from the figures that there is significant variability in estimated power for 90nm and 65nm libraries, especially between the *WorstLeakageHighV* and *TypPerfLowV* corners (more than a **10**× difference). As mentioned earlier, just considering the traditionally used corners (e.g. *TypPerf* vs. *TypLeakage*) is not realistic because there is a large variation in power consumption for sub-100nm libraries due to DSM effects (Section 3) which can only be captured by additional corners. In order to meet power goals, designers thus need to consider multiple PVT corners to understand the power characteristics of a design. ***It can be concluded from our experimental results that there is a significant (as much as 10×) variation in power consumption across PVT corners, for on-chip communication architectures.***

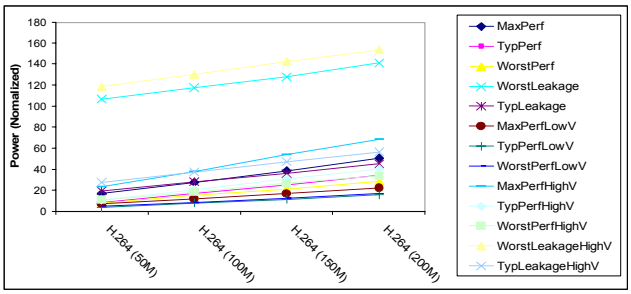

Figure 4. Normalized Power for H.264 SoC Subsystem

4.2. Impact on SoC Design

In our next set of experiment, we were interested in investigating the impact of PVT corners on power consumption for not just the communication architecture, but for an entire SoC design. We were also interested in estimating the impact of clock frequency on power consumption for the corners. For the purpose of this experiment, we implemented a complex SoC subsystem for the H.264/AVC codec [21] at the RTL level, consisting of the chroma inter-prediction IP, buffers and several memory blocks interconnected together using the AMBA AHB bus matrix [18]. Figure 4 illustrates normalized total power of the synthesized H.264 subsystem, obtained after detailed gate-level power analysis [20], for a 90nm technology library implementation, with clock frequency ranging from 50 MHz to 200 MHz. A near linear increase in power is observed with frequency, which implies that power increases by an approximately constant ratio due to clock scaling. The impact of leakage power can be observed by considering the difference in power (almost **2**×) between *TypPerf* and *TypLeakage* corners under the same V_{dd}. ***From our experimental result, it can be concluded that there is significant variation (as much as 60×) in power consumption across PVT corners, for SoC designs.***

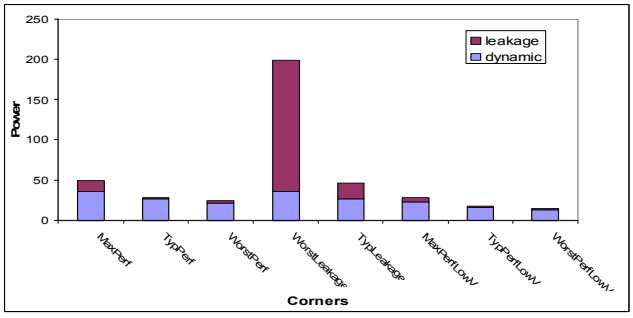

Figure 5. Decomposed Power for Bus Matrix Configurations at 65nm

4.3. Impact on Multi-Vt Design Flow

The multi-Vt technique is an effective way to reduce sub-threshold leakage current without sacrificing performance. High-Vt libraries can be used to reduce leakage current while low-Vt libraries can be used to get high performance on critical paths. We conducted several experiments for the 90nm and 65nm libraries, where we performed multi-Vt synthesis for different configurations of the AMBA AHB bus matrix communication architecture, running at 100 MHz bus clock frequency. The bus matrix power consumption across corners for the Multi-Vt case was found to be similar to the results shown in Figures 2 and 3, for the single-Vt case. Synthesis of the bus matrix configurations across different frequencies ranging from 100-300 MHz showed some interesting results. Unlike the near linear increase in power with frequency that was noticed in Figure 4, some distortions in the power consumption were noticeable at high frequencies. This was because additional low-Vt libraries were used to meet the more stringent timing requirements at higher frequencies in multi-Vt synthesis, unlike for the single-Vt synthesis case. ***Our experimental results indicated significant variations in power consumption across PVT corners (as much as 6×) for multi-Vt implementations of the on-chip communication architecture.*** Detailed experimental results obtained with multi-Vt synthesis can be found in our technical report [26].

5. Scaling Relation for Power Estimation Across PVT Corners

From the results of the experiments presented in the

previous section, we found that the power consumption for a PVT corner scales almost linearly with frequency. However, a much more interesting and important observation is that the power consumption numbers obtained for the different PVT corners show an almost constant ratio relative to each other. Thus if the power consumption of the bus matrix on-chip communication architecture for an implementation with PVT corner *C1* is expressed as:

$$P_{C1} = P_{L1} + P_{D1} \times f \quad \text{... (1)}$$

where P_{C1} gives the base level total power for a corner *C1* that has base level leakage power P_{L1} and base level dynamic power P_{D1}, at frequency f; then the power consumption for an implementation under any other PVT corner *C2* can be expressed as:

$$P_{C2} = \alpha_{1\text{-}2} \times P_{L1} + \beta_{1\text{-}2} \times P_{D1} \times f \quad \text{... (2)}$$

where the total power P_{C2} for another corner *C2* can be linearly scaled from the base level power relation in Eq. (1) by using scaling factors $\alpha_{1\text{-}2}$ and $\beta_{1\text{-}2}$ for leakage and dynamic power respectively. The scaling factors can be easily obtained by decomposing the total power into dynamic and leakage components (as shown in Figure 5, for the 65nm case shown in Figure 3), and averaging the ratio values. ***Thus, knowing the leakage and dynamic power for one PVT corner can enable us to obtain the power for other PVT corners using a simple linear model with good accuracy, which speeds up power exploration across corners considerably.***

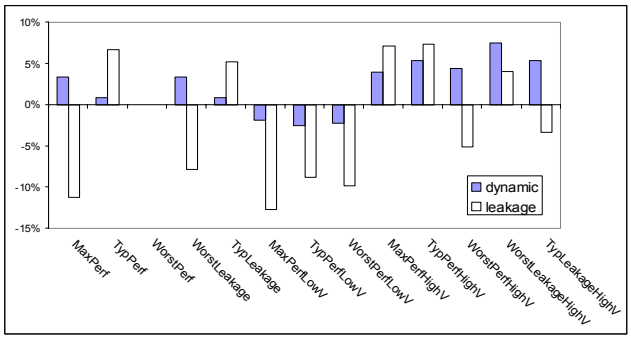

Figure 6. Normalized Power Estimation Error (90nm)

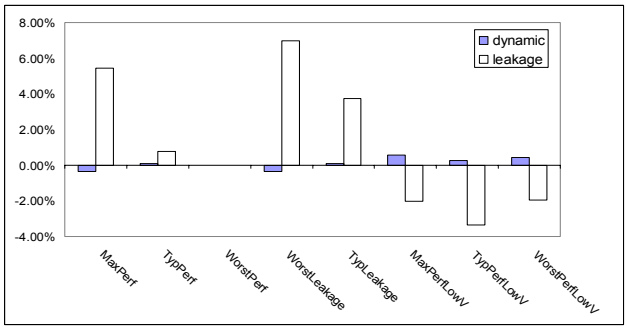

Figure 7. Normalized Power Estimation Error (65nm)

It is important to differentiate the proposed linear corner-to-corner scaling proposed in (2) from the power (specifically leakage) dependence on P, V, and T. While it is well known that a super-linear relation exists between leakage and each of those P, V, and T parameters [27], what (2) reflects is an observation that, outside of frequency, which affects dynamic power, the ratio of power (dynamic+leakage) shows little variations *across design instances*. In other words, once a design implementation is characterized for power across the PVT corners (say through detailed gate level simulation), the corner ratios obtained can be re-used to "scale" another design instance without having to do another full characterization run for that instance.

Figures 6 and 7 show the maximum estimation error for the leakage and dynamic power for the 90nm and 65nm libraries respectively, when the scaling factors are used to estimate power consumption for different PVT corners, for different design configurations: the AMBA AHB bus matrix communication architecture, the H.264 SoC subsystem and a register file, operating at different clock frequencies. It can be seen from the figures that using the scaling technique we propose, it is possible to estimate power consumption for different PVT corners with extremely good accuracy for dynamic power (< 5% in most cases), and fairly good accuracy for leakage power (< 10% in most cases). ***This amenability to scaling for PVT corners is an extremely important result, and has been obtained for cell libraries characterized with industrial strength numbers.***

6. PVT-Aware System Level Power Exploration

In this section, we will use the results of our observations from the previous sections, and show how to abstract up the power exploration across PVT corners early in the design flow, up to the system level, specifically for the AMBA AHB bus matrix communication architecture. The results of all the power exploration experiments conducted at the system-level have been verified by detailed gate-level simulation.

6.1. Incorporating PVT Corners in a System-level Power Estimation Methodology

To estimate power for the bus matrix communication architecture early in the design flow, at the system-level, we make use of the power estimation methodology for the bus matrix proposed by us in [17]. The approach makes use of energy macro-models to determine power consumption for the bus matrix logic components such as the input buffer stages, decoders, arbiters and output stages. A high level simulated annealing floorplanner [22] is used for early core placement and Manhattan routing is used to determine wire lengths. The wire lengths are subsequently used to determine wire energy, in formulations proposed in [23], which we extended to incorporate power for delay-optimally inserted repeaters. The power estimates for the bus matrix communication architecture were shown to be within 5% accuracy of gate-level power estimates in [17]. Creating the energy macro-

models for the bus matrix however requires a one-time effort to identify macro-model variables and coefficients using multiple linear regression analysis [25], which can take several hours. Since a different energy macro-model is required for every PVT corner, the task can take a long time.

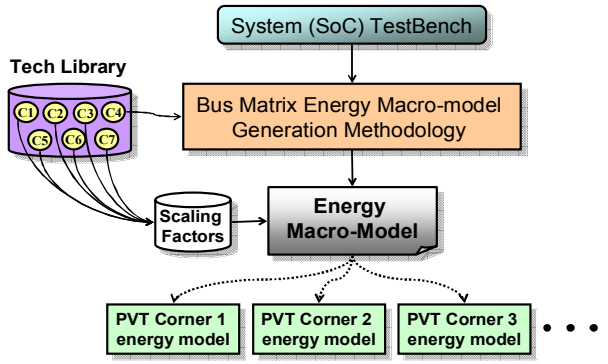

Figure 8. System-level Energy Macro-model Generation for PVT Corners

Figure 8 shows our proposed methodology to speed up bus matrix energy macro-model creation for different PVT corners of a technology library. This approach extends the energy macro-model based power estimation methodology from [17]. Initially, a system testbench consisting of a diverse set of bus matrix-based SoC designs is used to generate an energy macro-model for one of the PVT corners of a technology library, according to [17]. Subsequently, scaling relations presented in Section 5 are used for the other PVT corners to modify the base energy macro-model and create macro-models for each of the other PVT corners. This enables a considerable saving in time because only one macro-model generation iteration needs to be performed in order to obtain the energy models for all the corners in the selected technology library. *Since a single macro-model iteration can take in the order of hours, this approach can save us in the order of days to estimate power for the various PVT corners in ultra deep submicron technology library nodes.* In the next section, we present experiments to show how accurately these PVT corner macro-models can estimate power at the system-level for the bus matrix communication architecture.

6.2. System-level PVT-Aware Power Estimation for Bus Matrix Communication Architectures

To verify if the scaling factor-based PVT corner power estimation approach can be used to accurately explore power consumption across different corners for the bus matrix communication architecture, we performed an experimental study at the system-level. We selected an industrial strength multi-processor networking SoC application used for data packet processing and forwarding, with 4 ARM processors and more than 25 master, slave and memory IP blocks interconnected using the AMBA AHB bus matrix. We modeled the SoC application in SystemC [9] at the Transaction-based Bus Cycle Accurate (T-BCA) [24][28] abstraction. The goal of our experiment was to estimate power for the PVT corners at the system-level using two methods: first, with the traditional approach of creating energy macro-models for each PVT corner separately, and second, with the proposed scaling based approach.

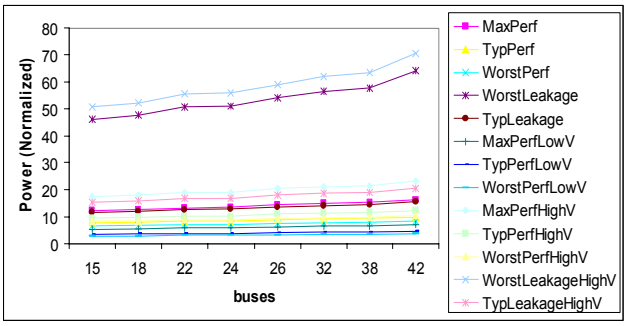

Figure 9. Normalized Power for networking SoC

In the first experiment, we used our communication architecture synthesis framework from [17] to generate a set of bus matrix solutions (each solution having a different number of buses) that satisfies the performance constraints of the networking SoC application. Next, we created bus matrix energy macro-models for each of the different PVT corners of the 90nm general purpose technology library. We plugged the energy macro-models into a T-BCA simulation model of the networking SoC, and simulated the design for each of the solutions in the solution set to get power consumption for the application. Figure 9 and 10 show the normalized power and energy obtained after simulating the design for all the PVT corners, for each of the solutions. As mentioned earlier, these numbers are within 5% accuracy of gate-level power estimates. The X-axis shows different bus matrix solutions, each having a different number of buses (and consequently different number of logic components such as arbiters, decoders etc.).

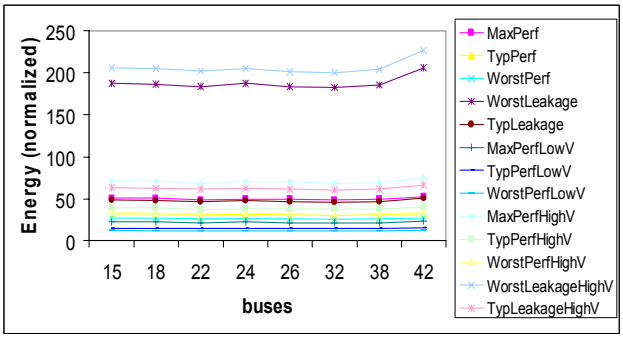

Figure 10. Normalized Energy for networking SoC

From the figures, it can again be seen that there is a significant variation in power and energy consumption across different PVT corners, and the power and energy ratios for different corners are fairly constant. Table 2 shows the percentage change in performance for the different solutions

having fewer buses, compared to the solution with 42 buses. It can be seen that solutions having fewer number of buses have lower performance. This is because there are more delays due to traffic conflict when data streams originating from different masters must share fewer buses. *Note that the performance numbers in Table 2 represent a lower bound on performance that can be achieved for all the PVT corners.*

Table 2. % performance variation for bus matrix configs

No. of buses	15	18	22	24	26	32	38	42
% perf. variation	-26.5	-21.9	-12.6	-14	-6.1	-0.7	-0.2	0

Creating energy macro-models for each PVT corner case turns out to be very time consuming, requiring a few days in designer effort. There is a need to speed this process up. Clearly, one solution is to use the scaling relations from Section 5, and the methodology presented in Figure 8. The question is, *can this approach be used at the system-level to accurately estimate power for PVT corners?* To answer this, we selected one of the PVT corners (*WorstPerf*) as a base reference and scaled its power results to create energy-macro models for the other PVT corners. We then plugged these energy macro-models into our T-BCA simulation model, and simulated the design for each solution in the bus matrix solution set, to obtain power numbers.

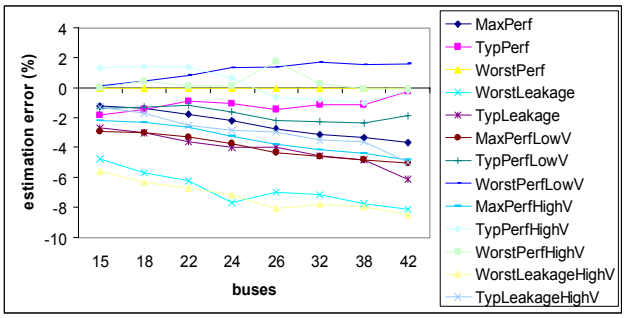

Figure 11. % power estimation error for corner cases when Eq, (2) is used to obtain power at PVT corners

Figure 11 shows the error in power estimation for the solutions, when we used the scaling technique to estimate power at PVT corners, compared to the traditional technique of creating energy macro-models separately for the corners (as was done in the previous experiment). It can be seen that the maximum absolute error compared to macro-model estimates is less than 9%. *The maximum absolute error compared to gate-level estimates is 14%* (since the macro-models are within 5% accuracy of gate-level estimates [17]), *which is an extremely good accuracy for PVT corner power estimation at the system-level.* The results imply that we only need to create an energy macro-model for one of the PVT corners, and use scaling factors to quickly obtain power for other PVT corners at the system-level. The overall time taken for creating energy macro-models for all the PVT corners from the base reference macro-model in this case is in the order of a few seconds, and *several orders of magnitude less* than the case where macro-models have to be separately created for every PVT corner.

Figures 12 and 13 show the percentage change in power and energy for the solutions in the bus matrix solution set, compared to the solution with the most number of buses (42). It can be seen that solutions with fewer number of buses have lower power and energy dissipation, at the cost of performance (Table 2). There is a large variation in energy and power consumption across the different PVT corners, for the solutions.

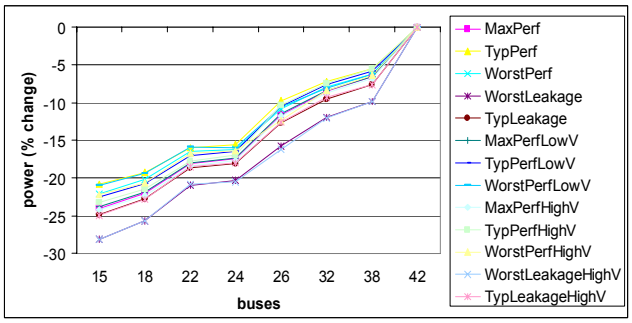

Figure 12. % change in power for networking SoC

6.3. Discussion of Experimental Results

From the experiments, it is clear that it is extremely important for designers to consider different PVT corners during early power and energy exploration at the system-level. Consider, for instance, the case where the average power dissipation constraint for the design is 80 mW. Assume that average power dissipation for the solution with 42 buses is 100 mW. In such a case, from Figure 12, it can be seen that the solution with 15 buses meets the constraint for all corners (i.e. % power reduction compared to 42 bus case > 20% for all corners). The solution with 18 buses meets the constraint for all corners except the *TypPerf* and *WorstPerfLowV* corners. The solutions with 22 and 24 buses meet the constraint only for the *WorstLeakage* and *WorstLeakageHighV* corners. If the designer used traditional techniques, he would have only a single power number for a technology library to rely upon during early exploration; now the designer can select a solution based on a more comprehensive characterization of the technology library across various PVT corners. The designer in this case can select the 15 bus solution conservatively, or go for the better performing solution with 18 buses (refer to Table 2), with a reasonable degree of confidence that the solution will not violate constraints under most PVT conditions.

Consider another case where a solution with minimum energy dissipation needs to be selected. From Figure 13, it can be seen that under maximum leakage conditions, the lowest energy is obtained for the solution with 15 buses. However, if such maximum leakage conditions will not be encountered, then the lowest energy dissipation (for the other corners) is obtained for the solution with 22 buses. Such PVT-aware

exploration information can aid the designer in selecting the appropriate solution with greater confidence and more accuracy, than the traditionally used approach of considering only a single corner for a technology library, during early system-level power exploration.

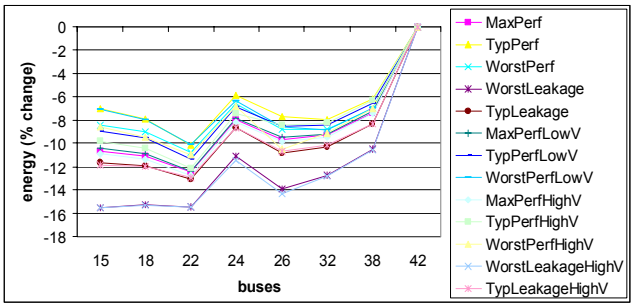

Figure 13. % change in energy for networking SoC

7. Conclusion

In this paper, we investigated the impact of PVT corners on power consumption at the system level, especially for the bus matrix on-chip communication architecture. We first conducted several experiments to show how there are significant variations in power consumption across different corners for a given technology library. Next we showed how it is possible to "scale up" and abstract the PVT variability to characterize the true design space early on in the design flow, at the system level. We used scaling relations to quickly create power models for the different PVT corners, to estimate the power consumption of the bus matrix at the system level. The scaled power models took several orders of magnitude less time to create than the traditional macro-modeling approach, with a maximum absolute estimation error of 14%, which is extremely good for early power estimation at the system-level. Finally, we experimentally established the importance of considering PVT corners during system-level power exploration. While we currently do not consider intra-die PVT variations, we believe that the simplicity of the abstractions described in this paper will make it feasible to incorporate such variability into future work.

Acknowledgement

This research was partially supported by grants from SRC (2005-HJ-1330).

References

[1] S. Borkar et al., "Design and reliability challenges in nanometer technologies", *DAC 2004*

[2] V. Khandelwal, A. Srivastava, "A General Framework for Accurate Statistical Timing Analysis Considering Correlations", *DAC 2005*

[3] J. Le, X. Li, L. Pileggi. "STAC: Statistical Timing Analysis with Correlation", *DAC 2004*

[4] A. Papanikolaou et al. "A System Level Methodology for Fully Compensating Process Variability Impact of Memory Organizations in Periodic Applications", *CODES+ISSS 2005*

[5] K. Lahiri et al., "Battery driven system design: A new frontier in low power design", *ASPDAC 2002*

[6] R. Ho, K. W. Mai, M. A. Horowitz, "The Future of Wires", *IEEE, vol. 89, April 2001*

[7] K. Lahiri, A. Raghunathan, "Power Analysis of system-level on-chip communication architectures", *CODES+ISSS 2004*

[8] K. Banerjee et al., "Coupled Analysis of Electromigration Reliability and Performance in ULSI Signal Nets", *ICCAD 2001*

[9] SystemC initiative, www.systemc.org

[10] M. Lajolo et al, "Cosimulation-based power estimation for system-on-chip design," *IEEE TVLSI, vol. 10, no. 3, 2002*

[11] N. Dhanwada, I-C. Lin, V Narayanan, "A Power Estimation Methodology for SystemC Transaction Level Models", *CODES+ISSS 2005*

[12] I. Lee et al. "PowerViP: SoC Power Estimation Framework at Transaction Level", *ASPDAC 2006*

[13] P. P. Sotiriadis, A. P. Chandrakasan, "A Bus Energy Model for Deep Submicron Technology," *IEEE TVLSI June 2002*

[14] L. Benini et al. "Architectures and Synthesis Algorithms for Power-efficient Bus Interfaces," *IEEE TCAD Sept. 2000*

[15] M. Caldari et al. "System-level power analysis methodology applied to the AMBA AHB bus", *DATE 2003*

[16] A. Bona et al., "System level power modeling and simulation of high-end industrial network-on-chip", *DATE 2004*

[17] S. Pasricha, Y. Park, F. Kurdahi, N. Dutt, "System-Level Power-Performance Trade-Offs in Bus Matrix Communication Architecture Synthesis", *CODES+ISSS 2006*

[18] AMBA AHB Interconnection Matrix, www.synopsys.com/products/designware/amba_solutions.html

[19] http://www.arm.com/products/physicalip/metro.html

[20] Synopsys CoreTools, PrimePower www.synopsys.com

[21] Join Video Team (JVT) ISO/IEC MPEG, ITU-T VCEG: Final draft international standard of joint video specification (ITU-T Rec. H.264/ISO/IEC 14496-10 AVC). *JVT-G050, 2003*

[22] S. N. Adya, I. L. Markov, "Fixed-outline Floorplanning: Enabling Hierarchical Design", *IEEE TVLSI, Dec. 2003*

[23] C. Kretzschmar, et al., "Why transition coding for power minimization of on-chip buses does not work", *DATE 2004*

[24] S. Pasricha, N. Dutt, M. Ben-Romdhane, "Extending the Transaction Level Modeling Approach for Fast Communication Architecture Exploration", *DAC 2004*

[25] J.J. Faraway, "Linear Models with R", *CRC Press, 2004*

[26] S. Pasricha, Y. Park, F. Kurdahi, N. Dutt, "PVT Variation Aware On-Chip Communication Architecture Synthesis", CECS Technical Report, *Jan 2008*

[27] J. Rabaey, A. Chandrakasan, B. Nikolic, "Digital Integrated Circuits: A Design Perscpective", *2nd Ed. Prentice-Hall, 2003*.

[28] S. Pasricha, "Transaction Level Modeling of SoC with SystemC 2.0", *SNUG 2002*