

Rethinking Threshold Voltage Assignment in 3D Multicore Designs

Koushik Chakraborty
Electrical and Computer Engineering
Utah State University
kchak@engineering.usu.edu

Sanghamitra Roy
Electrical and Computer Engineering
Utah State University
sroy@engineering.usu.edu

ABSTRACT

Due to the inherent nature of heat flow in 3D integrated circuits, stacked dies exhibit a wide range of thermal characteristics. The strong dependence of leakage with temperature and process variation plays havoc in achieving system level energy efficiency in such systems, complicating the task of power provisioning in 3D multicores.

In this paper, we address this power provisioning challenge in 3D ICs by advocating a novel microprocessor design paradigm, where the circuit designers are aware of the intended placement of a die in a 3D stack. We present a concrete application of this paradigm through a threshold voltage (V_t) assignment algorithm for a 3D multicore system, where we specifically account for: (a) the change in the role of leakage power, (b) expected operating frequency, and (c) dependency of PV induced leakage variation and V_t levels. Detailed simulation based experiments with our proposed algorithm show 2–15% improvement in energy efficiency for a typical multicore system organized as 3D stacked dies.

1. INTRODUCTION

The recent development in organizing integrated circuits in three dimensions (3D) instead of conventional two dimensional placement is causing major design overhaul in several aspects of computer systems. While 3D circuits promise to solve the long standing problem of rapidly deteriorating interconnect delay in high-performance microprocessors, the severity of thermal challenge in 3D requires both system architects and circuit designers to revisit several design principles. A typical 3D stack is flanked by a heat sink at the lower end and a package at the upper end. This arrangement creates a thermally heterogeneous profile, where the temperature of a die increases rapidly with increasing distance from the heat sink. Consequently, severe restrictions are imposed on the operating voltage and frequency of upper stacks, degrading their energy efficiency.

Much of the energy efficiency problems in dies with high temperature stems from their leakage power, which grows exponentially with temperature [13]. The emergence of process variation (PV) in nanoscale technologies further exacerbates leakage power management. Several recent studies have shown that 30% variation in process parameters can result in 20X increase in leakage [3]. This large variation of leakage power in 3D IC stacks can impair system level power management by creating uncertainty in power estimation. Consequently, designers must either set conservative voltage-frequency in upper layer dies, sacrificing performance, or face the likelihood of extreme operating temper-

atures and possible thermal runaway.

In this paper, we address the power provisioning challenge in 3D ICs by advocating a novel microprocessor design paradigm where circuit designers are aware of the intended placement of a die in a 3D stack. We present a concrete application of this paradigm through a threshold voltage (V_t) assignment algorithm for a 3D multicore system. We show how the design considerations of selecting low V_t devices are fundamentally altered in 3D ICs. While dies close to the heat sink can continue to use 2D design principles [7], upper layer dies must account for: (a) the change in the role of leakage power, (b) expected operating frequency, and (c) dependency of PV induced leakage variation and V_t levels. Detailed simulation based experiments with our algorithm show 2–15% improvement in energy efficiency in a 3D Multicore running multiprogrammed SPEC benchmarks.

Existing literature describes various techniques for dynamic thermal management of 3D ICs, both at the system and circuit levels. Puttaswamy and Loh propose micro-architectural techniques to steer a majority of the processor's switching activity towards the die closest to the heat sink, exploiting its superior cooling efficiency [12]. At a system level, Zhu, et al. explore a combination of software and hardware techniques for runtime thermal management in 3D ICs, targeting the impact of heterogeneous cooling efficiencies across different dies [19]. These works do not explore intrinsic circuit parameters to exploit the operating conditions of 3D stacks. Yu et al. propose a multiple supply voltage arrangement to minimize power in 3D integrated circuits [17]. However, their approach ignores the heterogeneous thermal profiles of 3D stacks. Hua et al. explore the tradeoffs in timing, power, temperature and the optimal number of tiers in 3D ICs, but do not provide any algorithm to improve system level energy efficiency [6]. Goplen et al. and Cong et al. have modified physical design algorithms like floorplanning and placement to optimize the maximum temperature for 3D ICs [5, 4]. However, these techniques consider only temperature and do not explore its complex relationship with leakage power. To the best of our knowledge, none of the previous works have explored multiple threshold voltage assignment techniques specifically tuned for 3D IC stacks.

The rest of the paper is organized as follows. Section 2 presents an extensive analysis of the interaction between process variation, threshold voltage and leakage power at multiple operating temperatures. In Section 3, we illustrate the challenge of power provisioning in 3D multicores by combining heat flow equations of 3D IC with PV dependent leakage characteristics. In Section 4, we present an algorithm for

threshold voltage assignment in 3D multicore architectures. Section 5 presents experimental results of our algorithm by using a detailed full system simulation of a 3D multicore system. We conclude our discussion in Section 6.

2. PV AND V_T AWARE LEAKAGE ANALYSIS

In this section, we investigate the complex interplay between process variation, threshold voltage, temperature and leakage current.

The major components of the leakage power in a CMOS circuit are the subthreshold leakage (I_{sub}) and the gate leakage (I_{gate}), respectively [13]. Although gate leakage has been steadily increasing for the last few technology generations, the industry wide adoption of high- κ dielectric appears to stem this trend for the next few technology generations [10]. Consequently, the leakage current in a CMOS circuit is dominated by the subthreshold leakage I_{sub} .

Sub-threshold leakage exponentially varies with the threshold voltage V_t , which is sensitive to the transistor effective gate length L_{eff} and the gate oxide thickness T_{ox} . These two physical characteristics of devices are greatly affected by both inter-die and intra-die process variation. Previous work have shown that leakage current variation exhibits a lognormal distribution, under Gaussian distributions of L_{eff} and T_{ox} [3]. For example, a 30% deviation from nominal values can cause 20X variation in leakage.

Interestingly, this effect of process variations on the sub-threshold leakage current is substantially more pronounced at lower V_t levels than at higher V_t levels. Since $I_{sub} \propto \frac{1}{e^{V_t}}$, the rate of change of I_{sub} with V_t accelerates rapidly as we lower the V_t ($\frac{dI_{sub}}{dV_t} \propto e^{-V_t}$). We illustrate this phenomenon using spice simulation next.

We perform Monte Carlo simulations with 10,000 runs on a CMOS inverter circuit in HSPICE using Predictive Technology Model (PTM) [18]. We use the high performance (HP) and low power devices (LP) to capture the characteristics of low V_t and high V_t devices, respectively. The effective gate length L_{eff} and the gate oxide thickness T_{ox} are respectively varied between $\pm 20\%$ of the nominal values using Gaussian distribution functions to account for manufacturing process variations. Circuit simulations are run at two different temperatures ($45^\circ C$ and $110^\circ C$) to show their behavior at a range of operating conditions. We use three different technology nodes spanning between 45nm and 22nm.

Table 1 summarizes the leakage and delay analysis. For each experiment, we present the mean and the σ as a percentage of the mean (i.e., $100 * \sigma / mean$). Using high V_t in the low power devices cuts down the leakage current by more than 100X across these three technology nodes: a fact that has been exploited for last several years [7]. However, the key observation we make in this paper is the dramatic reduction in leakage variation caused by PV. For example, at 45nm node, the ratio of σ and mean is 19X more in low V_t devices, compared to the corresponding ratio in high V_t devices. Most interestingly, *this comparative advantage in leakage variation is growing rapidly with technology scaling*. At the 22nm node, the leakage variation is more than 200X in low V_t devices compared to high V_t devices.

Figure 1 shows a closer look into this phenomenon by presenting the relative spread of leakage current around the mean. We notice a lognormal distribution in leakage varia-

Node	Low V_t		High V_t	
	$45^\circ C$	$110^\circ C$	$45^\circ C$	$110^\circ C$
45nm	I_{off} : 8.7, 204 Delay: 3.8, 7.3	I_{off} : 22.4, 166.4 Delay: 5.01, 12.1	I_{off} : 0.07, 19.3 Delay: 11.6, 5.5	I_{off} : 0.25, 11.2 Delay: 17.63, 7.11
32nm	I_{off} : 10.56, 500.6 Delay: 3.48, 8.8	I_{off} : 29.7, 280.9 Delay: 4.57, 13.75	I_{off} : 0.105, 8.65 Delay: 12.1, 7.1	I_{off} : 0.28, 18.7 Delay: 16.78, 7.9
22nm	I_{off} : 22.01, 1252 Delay: 3.1, 14.5	I_{off} : 58.42, 575 Delay: 6.01, 16.5	I_{off} : 0.54, 6.6 Delay: 11.93, 9.3	I_{off} : 1.2, 8.5 Delay: 15.9, 9.1

Table 1: Summary of I_{off} ($nA/\mu m$) and Delay (ps) comparison between low V_t and high V_t devices. Numbers show the mean and σ as a percentage of the mean for each measurement.




Figure 1: Comparison of I_{off} Spread Around Mean.

tion, also reported in [3]. Across all technology generations, leakage is far more concentrated around the mean in high V_t devices (LP). Note that the figure is truncated at the extremes to expand on the distribution around the mean.

Leakage increases rapidly with the rise of temperature, roughly growing to 3X from $45^\circ C$ to $110^\circ C$ in most configurations. The general trend of PV impact remains identical to the characteristics seen at $45^\circ C$. Fundamentally, across different technology nodes and operating conditions, we find that high V_t devices are substantially more tolerant to PV (by one to two orders of magnitude). We present the implication of this characterization in the context of power provisioning in a 3D IC next.

3. IMPACT OF LEAKAGE IN POWER PROVISIONING 3D IC MULTICORE

In this section, we analyze the heat flow in a typical 3D IC stack to investigate the challenge of thermal provisioning in the light of temperature and PV dependent leakage characteristics from Section 2.

A key component of thermal provisioning is to estimate the dynamic power budget, and expected energy wastage through leakage. Given the importance of thermal issues in a 3D IC, we will explore the interaction between die temperature T and leakage power P_{leak} . We use the temperature dependent leakage model proposed by Skadron et al. [16], where leakage power of a die is modeled as a temperature dependent function of the active power P_{dyn} dissipated by the same die. The temperature dependent ratio R_T of the leakage power to the dynamic power is shown in equation 1.




Figure 2: A Typical 3D IC Stack

$$R_T = \frac{P_{leak}}{P_{dyn}} = \frac{R_0}{V_0 T_0^2} e^{\frac{B}{T_0}} \cdot VT^2 \cdot e^{-\frac{B}{T}} \quad (1)$$

where T_0 is the ambient temperature, R_0 is the ratio at T_0 and nominal voltage V_0 , and B is a process technology constant. For a specific process technology and constant supply voltage V , equation 1 can be simplified to

$$R_T = K_0 T^2 \cdot e^{-\frac{B}{T}} \quad (2)$$

where K_0 and B are constants in the absence of PV. Under process variation, the values of K_0 and B change substantially based on the underlying variation of leakage with PV.

3.1 Motivational example

We now integrate this temperature dependent leakage model in the governing heat flow equations of a 3D IC, and illustrate the challenge of leakage aware thermal provisioning in a typical multi-die stack (see Figure 2). The configuration shown has four dies, with the die_1 close to the heat sink and the die_4 close to the package. The design parameters of the chip, shown in Table 2, have been chosen to loosely represent the 45nm Intel Nehalem processor [8]. Each die is only half the area of a quad-core Nehalem processor due to the presence of only two cores (see Figure 3). The 4-die stack operates at 150 Watts (TDP).

To analyze the temperature and power profile of this 3D IC stack, we use an analytical heat flow model proposed by Jain et al. [1]. This model uses a thermal resistor network with multiple heat sources to calculate the temperature rise in each die of a 3D stack. This model captures the first order effects of thermal heat flow in a 3D IC. Using this model, we can calculate the temperature distribution for a given power distribution or vice versa, in an n -die stack.

Die area	135 mm ²
Die thickness	50 μm
Interface layer thermal conductivity	0.3 W/mK
Interface layer thickness	10 μm
Heat sink thermal conductivity	400 W/mK
Thermal conductivity of Silicon	150 W/mK
Package thermal resistance	30X heat sink resistance
T_{amb}	40°C
T_{max}	110°C

Table 2: 3D IC design parameters

To illustrate the impact of leakage in conjunction with PV, we investigate two different scenarios in our example:

1. Case I - Equal power budgeting.
2. Case II - 3D stack aware power budgeting

In Case I, we distribute the power budget of 150W *uniformly* across the 4 dies. We then use the heat flow equations in [1] to calculate the temperature of each die. Subsequently, we calculate the ratio R_T using equation 2, and find the percentage ranges of dynamic and leakage power dissipated by each die. Table 3 summarizes the results of our calculations.

For each case, the first and second column show the total power P_{tot} in Watts and the calculated temperature of each die. The third column shows the percentage of total die power available as dynamic power in each die. After accounting for the leakage variation characteristics from Table 1, we can get a *maximum likelihood range* for the dynamic power honoring P_{tot} . For example in die_4 , the total budgeted power is 37.5 watts, leakage can vary from 70.3% to 98.9%. Such variation leads to a provisioning range of 1.1% to 29.7% of the total power for dynamic power.

To maintain acceptable operating temperatures and avoid extreme temperatures in die_3 and die_4 as in Case I, we must decrease the budgeted power for the dies away from the heat sink. In Case II, we show such a power distribution across the 4 dies. We find that a decreasing power distribution with increasing distance from the heat sink leads to a more acceptable thermal profile across the dies. However, variation in leakage can still severely degrade dynamic power budget by more than 60%, as the dynamic power in die_4 varies from 40% to 15% of the total power.

die	Case I			Case II		
	P_{tot} W	T °C	P_{dyn} %	P_{tot} W	T °C	P_{dyn} %
1	37.5	72.2	46.8-62.2	63.0	72.7	46.3-61.8
2	37.5	97.3	20.5-43.5	45.0	92.2	25.7-47.2
3	37.5	113.1	6.15-33.3	25.0	100.4	17.5-41.4
4	37.5	119.6	1.1-29.7	17.0	102.4	15.7-40.1

Table 3: Thermal and leakage analysis of a 3D stack with constant total power = 150 W

Honoring these dynamic power budget will require significant effort, and we outline the nature of solution obtained through voltage frequency scaling (VFS). We assign the nominal voltage and frequency to die_1 , that has the maximum dynamic power budget. We next scale the voltage and frequency of the other dies to honor their respective dynamic power budgets (Case II in Table 3). Table 4, column II shows the allowable voltage and the range of operating frequencies in the respective dies. Our analysis demonstrates a wide range of operating frequencies in stacked dies: dies away from the heat sink achieving a small fraction of the nominal frequencies seen in contemporary designs.

To exploit the low frequency range in upper two dies, we can use higher threshold devices in the *critical* sections of die_3 and die_4 . (Note that high V_t devices are already used in non-critical sections following traditional 2D design techniques). Such a V_t assignment can drastically reduce their leakage power consumption, allowing them to devote a larger percentage of their limited power budget to dynamic power. For example, using higher threshold devices in die_3 and die_4 allows them to operate at 2GHz and 1.3 GHz frequencies respectively as shown in column III of Table 4. This is because, the leakage power is reduced by *100X*, preserving most of

the budgeted power for useful computation in the form of dynamic power.

die	(Voltage,Freq) at low Vt V,GHz	(Voltage,Freq) at high Vt V,GHz
1	(1,2,2) - (1,3)	Same
2	(0.9,1.1) - (0.9,2)	Same
3	(0.8,0.5) - (0.8,1.2)	(0.95,2)
4	(0.8,0.4) - (0.8,1.1)	(0.95,1.3)

Table 4: Operating voltage frequency in each die

The above case study demonstrates the utility of a fundamental design overhaul in multicore designs for 3D IC, where the intended die placement in a 3D stack becomes a key design consideration. While the example above shows a specific case, we generalize the V_t assignment in a 3D multicore through a comprehensive algorithm next.

4. AN ALGORITHM FOR V_T ASSIGNMENT IN 3D IC

Based on our findings in the previous two sections, we develop an algorithm for threshold voltage assignment in this section. Our algorithm *3DIC_Vt_assign* is shown in table 5. There are two key insights behind this algorithm: (a) The power budget at each die is governed by the heat flow equations in 3D IC, and the allowable operating temperature; (b) Given this total power, the dynamic power and hence allowable frequency at each die is governed by the amount of power wasted as leakage.

Due to the reducing gap between threshold voltage and supply voltage, a limited number of supply voltages are available for VFS. We assume a given set of m operating voltages including the nominal voltage, and a set of p threshold voltages. For each threshold voltage V_t , we have a set of tuples (V_{dd}, f_{max}) such that the device can operate at a maximum frequency of f_{max} at a supply voltage of V_{dd} . Using the analytical heat flow equations (Section 3.1), we estimate the total power budget P_i for each die, to give an overall acceptable thermal profile in the 3D stack. Furthermore, we assume the total nominal device power P_{nom} dissipated at the nominal voltage, frequency and threshold voltage.

At any given operating configuration $(V_{tj}, V_{ddj}, f_{maxj})$, we calculate the maximum total power $P_{candidate}$ by scaling the nominal power P_{nom} for the respective supply voltage, threshold voltage and frequency. If $P_{candidate}$ for the i^{th} die is within its budgeted power P_i , we add $(V_{tj}, V_{ddj}, f_{maxj})$ to the list of allowable voltage, frequency and threshold voltage for die_i . From the set of all allowable configurations (V_i, V_{dd}, f_{max}) for a particular die, we select the operating point that gives the maximum frequency.

3DIC_Vt_assign calculates the lowest threshold voltage of each die, that allows it to operate at the maximum frequency permissible by its given power budget. As the number of available operating voltages and threshold voltages are small, the computational complexity of our algorithm is approximately linear with the number of dies n .

The fundamental idea behind this selection is to tradeoff leakage with dynamic power via High V_t transistors, when (a) the operating frequencies are substantially lower than the nominal frequency, and (b) saving leakage is paramount to achieving energy efficiency. Both these conditions are precisely applicable for dies away from the heat sink in a 3D

ALGORITHM 3DIC_Vt_assign:

Input:

1. Set of m operating voltages $\{V_{dd1}, V_{dd2}, \dots, V_{ddm}\}$
2. Set of p threshold voltages $\{V_{t1}, V_{t2}, \dots, V_{tp}\}$, each with corresponding max frequencies at different operating voltages $\{(V_{dd1}, f_{max1}), (V_{dd2}, f_{max2}), \dots, (V_{ddm}, f_{maxm})\}$
3. Total power budget of n dies $\{P_1, P_2, \dots, P_n\}$
4. P_{nom} : total device power at nominal (V_i, V_{dd}, f) and max *IPC*

Output: Operating voltage, frequency and V_t of time critical blocks in each die

Begin

1. Using thermal eqns find die temperatures $\{T_1, T_2, \dots, T_n\}$
2. **for** $i := 1$ to n **do** /*for each die*/
3. Set $A := \phi$ /*Allowable set of (V_i, V_{dd}, f_{max}) */
4. **for** $j := 1$ to p **do** /*each V_t */
5. **for** $k := 1$ to m **do** /*each V_{dd}, f_{max} */
6. Compute $P_{candidate}$ for $(V_{tj}, V_{ddk}, f_{maxk})$
7. by scaling P_{nom}
8. **if** $(P_{candidate}) \leq P_i$
9. Add $(V_{tj}, V_{ddk}, f_{maxk})$ to A
10. **else** Reject $(V_{tj}, V_{ddk}, f_{maxk})$
11. **end**
12. **end**
13. Assign $(V_{tj}, V_{ddj}, f_{maxj})$ to die_i
14. s.t. $\forall (V_{tk}, V_{ddk}, f_{maxk}) \in A, f_{maxj} \geq f_{maxk}$
15. **end**

End

Table 5: Algorithm summary for 3DIC_Vt_assign

IC stack. Interestingly, this selection is opposite to nominal and fast design corners where low V_t transistors allow much higher frequency at the cost of higher leakage.

5. EXPERIMENTAL ANALYSIS

In this section, we will describe our 3D Multicore modeling methodology, and present our experimental results.

5.1 Methodology

We use full-system simulation built on top of Virtutech SIMICS [9]. SIMICS provides the functional model of several popular ISAs, in sufficient detail to boot an unmodified operating system. For our experiments, we use the SPARC V9 ISA, and use our own detailed timing model to enforce timing characteristics of an out-of-order microprocessor.

5.1.1 Multicore System

We model an eight core multicore system spread over 4 dies in a 3D IC, where each stacked layer consists of two cores with a few shared L2 cache banks (see Figure 3). Each processing core has a private instruction and data cache: both of them are 4-way, 32KB with 2-cycle latency. Spreading the L2 banks across multiple dies improves latency and performance with marginal complexity overhead [11]. The shared L2 cache modeled is a 32-way associative 16MB with uniform access latency of 50 cycles. We use a 4-wide superscalar microprocessor core with 192 entry instruction window, 11 cycle branch mis-prediction loop, aggressive 2-level branch predictors, and a 64-entry load-store disambiguation predictor. The entire stack uses 32nm technology node, operating at 1.0V and 3GHz (nominal voltage and frequency).




Figure 3: 4-Die Stacked Configuration

Name	Composition
MP1	astar-lakes(both), xalancbmk (2), sjeng(2), h264 (both), libquantum(2), hmmer(1)
MP2	mcf(both), hmmer(both), h264(1), xalancbmk(1), gobmk(both)
MP3	hmmer(2), gcc(both), omnetpp(both), libquantum(1), perlbench(2)
MP4	gobmk(both), perlbench(1), sjeng(1), h264(both), mcf(2), libquantum(2)
MP5	mcf(2), sjeng(1), astar-lakes(both), omnetpp(2), h264(1), gcc(both)
MP6	libquantum(both), xalancbmk(2), gcc(both), mcf(both), bzip2(1)
MP7	mcf(2), xalancbmk(2), h264(both), omnetpp(both), astar-lakes(2), bzip2(2)
MP8	xalancbmk(both), perlbench(both), bzip2(1), mcf(2), sjeng(1), omnetpp(1)

Table 6: Multi-programmed Workload Composition. Specific phases included from each benchmark is shown in parentheses.

5.1.2 Power and Thermal Analysis

Power and thermal analysis is performed by combining HotSpot and Wattach toolset [2, 16] in our simulation infrastructure. Wattach provides power estimation based on pipeline activity, and we calibrate this estimation using HotSpot, validating acceptable temperature range. Using the layered configuration of HotSpot, we model a 4-die stacked IC with their grid model, and update the temperature at $10\mu\text{sec}$ intervals to obtain temperature dependent leakage estimation.

5.1.3 Workloads

We use several multiprogrammed workloads composed of representative phases of SPEC CPU2006 benchmark on Solaris 9. We use two of the most representative phases from these SPEC benchmarks in our study, extracted using SimPoint toolset [15]. Table 6 shows the composition of our multiprogrammed workloads. We run each simulation for 200 million cycles, and count the total instructions committed as the aggregate throughput. representative portions of each toolset

5.2 Results

We show the results of $3DIC_{V_t}\text{assign}$ for two different libraries shown in Table 7. *Library 1* has two different threshold voltage levels, while *Library 2* has five. *Library 2* is more representative of the current and forthcoming technology generations.

The baseline system replicates the same design in all the

Name	Library Characterization
Library 1	$Low V_t = (1, 1, 1)$
	$High V_t = (0.9, 0.4, 0.06)$
Library 2	$Low V_t = (1, 1, 1)$
	$Low - Mid V_t = (0.95, 0.6, 0.1)$
	$Mid V_t = (0.9, 0.4, 0.06)$
	$Mid - High V_t = (0.9, 0.3, 0.03)$
	$High V_t = (0.8, 0.2, 0.005)$

Table 7: Device Libraries. For each device we show the $(V_{dd}, f_{max}, P_{leak})$, normalized to the $Low V_t$ device in Library 1.

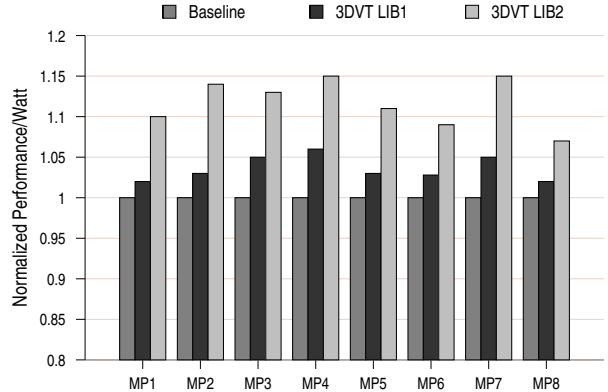


Figure 4: Power-Performance Comparison. From left to right, three bars in each benchmark show the results for the baseline, $3DIC_{V_t}$ using Library 1 (3DVT LIB1), and $3DIC_{V_t}$ using Library 2 (3DVT LIB2), respectively.

stacks, using $Low V_t$ devices in the critical sections. To remain within the allowed power budget at each die, baseline scales down the frequency, with marginal voltage scaling. For each die, designs guided by our algorithm $3DIC_{V_t}\text{assign}$ choose different sets of V_t devices along with supply voltage and frequency combination. In all cases, choice of V_t in die_1 remains identical to the baseline. However, the restriction on operating frequency as outlined in Table 4 allows different choices for other dies. Subsequently, we scale 67% of the die leakage using leakage factors mentioned in Table 7, modeling typical leakage contribution from low V_t devices [14].

5.2.1 Energy Efficiency

Figure 4 presents the improvement in energy efficiency measured as Performance per Watt. Using *Library 1*, our algorithm shows a modest improvement in energy efficiency (2–6%), while *Library 2* yields a substantially better improvement (7–15%). Much of this improvement stems from better use of power budget by reducing leakage, and exploiting higher available dynamic power. Using *Library 2* leads to a relatively higher operating frequency in top stacks due to the higher number of available V_t s. Although using *Library 1*, leads to savings in leakage, it is unable to effectively improve performance as the operating frequency is limited for $High V_t$. These improvements reported are for the entire 3D stack, and are typically lower than individual improve-




Figure 5: Power Consumption Comparison.

ments seen in die_3 and die_4 (Table 4).

Note that these improvements are reported without conservatively scaling the operating frequency in the baseline to allow for leakage PV variation as discussed in Section 3. Conservative scaling of frequency in the baseline can further improve the energy efficiency from our scheme, as high V_t devices are more tolerant to leakage fluctuation from PV.

5.2.2 Power Comparison

Figure 5 presents a comparison of power consumption across the three schemes. From left to right, three bars in each benchmark show the results for the baseline, 3DIC_ V_t using Library 1, and 3DIC_ V_t using Library 2, respectively. Each bar is broken down to show the components of dynamic power (lower) and leakage power (upper). Using *Library 2* sees a substantial reduction in leakage power across all benchmarks as almost all dies except die_1 benefit from comparatively higher V_t assignment in their critical devices. The dynamic power budget is increased in this case, boosting performance and yielding better system level energy efficiency. Using *Library 1* only yields savings in leakage power in die_3 and die_4 , but the system is unable to exploit more available dynamic power due to the delay constraint of a single *High V_t* device.

6. CONCLUSION

This paper presented a novel technique for designing high performance microprocessor dies specifically targeting 3D ICs. 3D integrated circuits pose an intriguing design challenge due to the extreme variation in operating conditions. While dies placed close to the heat sink operate similar to conventional 2D designs, dies away from the heat sink typically operate at high temperature and substantially lower power budget. A large percentage of this limited power budget is wasted through leakage. Furthermore, uncertainties in leakage current estimation due to process variation complicates power provisioning in 3D IC stacks. Solutions to these problems seek a fundamentally new paradigm of microprocessor design where designers must be aware of the die placement in a 3D stack. We demonstrate one application of such a design philosophy by developing an algorithm for 3D IC aware threshold voltage assignment. Using detailed simulation analysis for a 3D multicore, we observe 2–15% energy efficiency improvement through our technique.

Acknowledgement

This work was partially supported by the Women and Gender Research Institute at Utah State University. We thank our anonymous reviewers for their comments on this paper.

REFERENCES

- [1] ANKUR JAIN, ROBERT E JONES, R. C., AND POZDER, S. Analytical and Numerical Modeling of the Thermal Performance of Three-Dimensional Integrated Circuits. *IEEE Transactions on Components and Packaging Technologies* (in press, 2009).
- [2] BROOKS, D., TIWARI, V., AND MARTONOSI, M. Watch: a framework for architectural-level power analysis and optimizations. In *Proc. of 27th ISCA* (2000).
- [3] CHANG, H., AND SAPATNEKAR, S. S. Prediction of leakage power under process uncertainties. *ACM Transactions on Design Automation of Electronic Systems* 12, 2 (2007).
- [4] CONG, J., AND ZHANG, Y. Thermal Via Planning for 3-D ICs. In *Proc. of ICCAD* (2005).
- [5] GOPLEN, B., AND SAPATNEKAR, S. S. Placement of Thermal Vias in 3-D ICs Using Various Thermal Objectives. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 25, 4 (April 2006), 692–709.
- [6] HUA, H., MINEO, C., SCHOENFLIESS, K., SULE, A., MELAMED, S., JENKAL, R., AND DAVIS, W. R. Exploring Compromises among Timing, Power and Temperature in Three-Dimensional Integrated Circuits. In *Proceedings of the 43rd annual conference on Design automation* (2006).
- [7] KETKAR, M., AND SAPATNEKAR, S. Standby power optimization via transistor sizing and dual threshold voltage assignment. In *Proceedings of the IEEE/ACM international conference on Computer-aided design* (2002).
- [8] KUMAR, R., AND HINTON, G. A Family of 45nm IA Processors. In *Proceedings of the IEEE International Solid-State Circuits Conference* (2009).
- [9] MAGNUSSON, P., CHRISTENSSON, M., ESKILSON, J., FORSGREN, D., HÅLLBERG, G., HÖGBERG, J., LARSSON, F., MOESTEDT, A., AND WERNER, B. Simics: A Full System Simulation Platform. *IEEE Computer* 35, 2 (Feb 2002), 50–58.
- [10] OHKURA, Y., SUZUKI, C., AMAKAWA, H., AND K.NISHI. Analysis of gate currents through High-K dielectrics using a Monte Carlo Device Simulator. In *International Conference on Simulation of Semiconductor Processes and Devices* (2003).
- [11] PUTTASWAMY, K., AND LOH, G. Implementing caches in a 3D Technology for High Performance Processors. In *IEEE International Conference on Computer Design* (2005).
- [12] PUTTASWAMY, K., AND LOH, G. H. Thermal Herding: Microarchitecture Techniques for Controlling Hotspots in High Performance 3D-Integrated Processors. In *Proc. of 13th HPCA* (2007).
- [13] ROY, K., MUKHOPADHYAY, S., AND MAHMOODI-MEIMAND, H. Leakage Current Mechanisms and Leakage Reduction Techniques in Deep-Submicrometer CMOS Circuits. *Proceedings of the IEEE* 91, 2 (2003).
- [14] RUSU, S., ET AL. A 65-nm Dual-Core Multithreaded Xeon Processor With 16-MB L3 Cache. *J. of Solid-State Circ.* (2007).
- [15] SHERWOOD, T., PERELMAN, E., AND CALDER, B. Basic Block Distribution Analysis to Find Periodic Behavior and Simulation Points in Applications. In *Proc. of 10th PACT* (2001).
- [16] SKADRON, K., STAN, M. R., SANKARANARAYANAN, K., HUANG, W., VELUSAMY, S., AND TARJAN, D. Temperature-Aware Microarchitecture: Modeling and Implementation. *ACM Transactions on Architecture and Code Optimization* 1, 1 (2004).
- [17] YU, S.-A., HUANG, P.-Y., AND LEE, Y.-M. A multiple supply voltage based power reduction method in 3-D ICs considering process variations and thermal effects.
- [18] ZHAO, W., AND CAO, Y. New Generation of Predictive Technology Model for sub-45nm Early Design Exploration. *IEEE Transactions on Electron Devices* 53, 11 (2006).
- [19] ZHU, C., GU, Z., SHANG, L., DICK, R., AND JOSEPH, R. Three-Dimensional Chip-Multiprocessor Run-Time Thermal Management. *Trans. on Comp.-Aided Design of Integrated Circuits and Sys.* 27, 8 (2008).