

A Self-Consistent Junction Temperature Estimation Methodology for Nanometer Scale ICs with Implications for Performance and Thermal Management

Kaustav Banerjee, Sheng-Chih Lin, Ali Keshavarzi*, Siva Narendra*, and Vivek De*

Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA 93106

*Circuit Research Lab, Intel Corporation, Hillsboro, OR 97124

Abstract

Accurate estimation of the Silicon junction (or die) temperature in high-end microprocessors is crucial for various performance analyses and also for chip-level thermal management. This work introduces for the first time, the notion of *self-consistency* in estimating the die temperature for sub-100 nm technologies by taking into account various electrothermal couplings between supply voltage, operating frequency, power dissipation and die temperature. It also comprehends chip-level reliability constraints and the impact of employing various packaging and cooling solutions in an integrated manner. The self-consistent solutions of die temperature are shown to have significant implications for evaluating various power-performance-reliability-cooling cost tradeoffs and can be used to optimize the performance of nanoscale ICs.

I. Significance of Self-Consistent Junction Temperature in Leakage Dominant Technologies

As CMOS technology scales beyond the 100 nm technology node, power dissipation issues and thermal management of integrated systems are becoming key factors in most high-performance ICs including microprocessors [1,2] (Fig.1). Power dissipation and associated thermal effects have strong impact on the junction temperature, packaging and cooling costs [2,3]. They also impact reliability and circuit performance. For sub-100 nm technologies, leakage power (P_{wl}) (dominated by subthreshold leakage [4]) is increasingly becoming a significant component of total chip power (P_{chip}) due to threshold voltage (V_t) scaling and process variations [5,6]. Also, subthreshold leakage current (or P_{wl}) is exponentially dependent on temperature (Fig.2) and the dependence gets stronger with scaling (Fig.3). Hence, as illustrated in Fig.4, the average junction temperature T_j , which normally varies approximately linearly with the junction-to-ambient thermal impedance θ_j , increases nonlinearly with θ_j due to coupling between P_{chip} and T_j , arising primarily due to the strong dependence of P_{wl} on T_j . Hence, for leakage dominant nanometer scale technologies, as illustrated in Fig.4, in order to maintain a desired value of T_j , a lower value of θ_j will be required, leading to an increase in packaging/cooling cost. Furthermore, as the supply voltage V_{dd} is increased, chip frequency increases, and so does the total power dissipation and hence T_j . Also, frequency itself is dependent on temperature due to the dependence of the transistor on-current (I_{on}) on T_j . The details of the various electrothermal couplings are summarized in Fig.5. It can be observed that supply voltage, power, frequency, reliability and temperature are strongly coupled quantities. Hence, a self-consistent electrothermal analysis method is highly desirable to accurately estimate T_j for any value of V_{dd} (or frequency).

System Level Power/Thermal Constraints

Additionally, at the system level, various packaging and cooling techniques are being developed with different θ_j , cooling efficiencies (η) and costs [7-9] (cooling cost is mirrored by system power, $P_{sys} = P_{chip} + P_{cool}$). The system cooling cost can also be understood using the following equation that relates the thermal

impedance θ_j to the chip junction temperature T_j , and the total power dissipation P_{chip} . T_{amb} is the ambient temperature of the chip.

$$\theta_j = \frac{T_j - T_{amb}}{P_{chip}} \quad (1)$$

From this equation, it can be observed that a larger junction-to-ambient temperature difference allows a smaller heat sink and air flow rate (i.e., larger θ_j and smaller system cooling cost) for dissipating the same power. Reduction in θ_j will increase the packaging and cooling cost rapidly. Hence, in general, maintaining larger T_j relaxes θ_j requirements in an *active power* dominated technology. However, for technologies that are *leakage dominant*, larger T_j will impact P_{wl} (Fig. 3), and hence the total power P_{chip} , and thereby influence θ_j and the cooling cost.

Increase in T_j will also impact reliability. Hence, for leakage dominant technologies, the cooling cost will also be strongly coupled to the supply voltage, operating frequency and reliability. Therefore, electrothermal models and analysis methodology is required to understand the relationship between these various parameters that are critical metrics for the development of high-performance ICs. Fig.6 shows an equivalent thermal circuit model for active chip cooling. Note that, in order to lower T_j (and thereby increase frequency and lower P_{wl}) packaging/cooling technologies with low θ_j and high η ($0 < \eta < 1$) can be used but this will dramatically increase the cost of the chip. Hence, it is critical to consider the chip packaging and cooling models in the self-consistent electrothermal analysis in order to comprehend various power-performance-cooling cost trade-offs.

II. Self-Consistent Methodology

The methodology is based on an integrated device, circuit, and system level modeling approach. It uses both analytical models and results from circuit simulation. The methodology has been summarized in Fig.7. For a given V_{dd} and initial T_j , the operating frequency of the chip is first estimated using extracted temperature and V_{dd} dependent I_{on} value through a circuit simulator. The estimated frequency is then used in the calculation of the switching (active) power (P_{sw}). The leakage power estimation uses extracted off-current value, I_{off} , (also V_{dd} and T_j dependent) from the circuit simulator. Effect of process (channel length) variation is also taken into account for estimating P_{wl} [10]. The total chip power is then used to calculate the new junction temperature using compact thermal models for the IC packaging and cooling technology. In order to accurately estimate the junction temperature for each iteration in Fig.7, the thermal model shown in Fig.6 is used, where,

$$\begin{aligned} T_j - T_{amb} &= P_{chip} \theta_j \\ T_{amb} - T_{out} &= P_{sys} \theta_{sys} \\ P_{sys} &= P_{chip} + P_{cool} = P_{chip} + (1 - \eta) P_{chip} = (2 - \eta) P_{chip} \end{aligned} \quad (2)$$

here P_{sys} is the system power which includes power dissipation due to the chip as well as power spent due to any dynamic cooling mechanism with efficiency = η (with $\eta < 1$). Note that T_{amb} is the

ambient temperature (typically temperature immediately outside the chip case) and T_{out} is the external room temperature.

The estimated junction temperature is then compared with the initial value of T_j to check for convergence. The process continues till a convergence in the value of T_j is achieved. Although not explicitly shown in Fig.7, chip-level reliability constraints are also applied to check the validity of the estimated value of T_j for the particular value of V_{dd} in the analysis. The above methodology was calibrated against measured power (both active and leakage), frequency and junction temperature data for a 32-bit microprocessor to extract and/or tune certain parameters used in the active and leakage power estimation models. The summary of the power and frequency calculations are outlined below.

Total Chip Power and Frequency Estimation

For the 32-bit microprocessor used in this study, the total chip power was estimated based on the sum of the logic, memory and I/O power. The logic portion of the chip was further sub-divided into various functional blocks/types such as the datapath, register files, RLS, and repeaters. From the total width of devices in each of these logic blocks and the fraction of each of these blocks in terms of the transistor technology, such as low- V_t , high- V_t with short channel-length and high- V_t with long channel-length, P_{logic} was estimated based on both the switching and the leakage power associated with each of the blocks. Since, leakage power can be mostly attributed to subthreshold leakage (P_{swt}) [4] and to a much lesser extent to gate leakage (P_{wgl}), we primarily focused on the subthreshold leakage. Our assumption was also due to the fact that the gate leakage (direct tunneling) is independent of temperature. It is also highly process dependent (high-k gate etc.) and can be tuned to a reasonably low level by suitable process adjustments. Hence, in this work by leakage power we mean power dissipation associated with (off-state) subthreshold leakage only. The chip frequency was estimated based on the low- V_t transistors only since these devices are used in the critical paths that determine effective chip speed. The memory and I/O power was estimated in a manner similar to that described above for the logic. The memory portion of the chip power was attributed entirely to the cache. The activity factor for estimating the switching power of memory was taken as 10^{-2} times the activity factor for logic. Finally, a small fraction of the chip power contributed by the I/O circuitry was estimated as a sum of the switching power (based on the I/O power supply and frequency) and a constant power arising due to the transitioning of the bits at the I/O pads.

III. Implications for Performance and Thermal Management

Although the results shown in Figs. 8-13 are specific to the above mentioned microprocessor, the conclusions are more generic. In Fig.8 the self-consistent method is used to estimate the total chip power along with the active and leakage power components for different values of V_{dd} . It should be noted that the high performance IC product used in this case was fabricated using 130 nm technology with an I_{on}/I_{off} ratio of 660 and nominal $V_{dd}=1.2V$. It can be observed that using a non self-consistent approach will either overestimate or underestimate the total power, depending on the chosen T_j , and will also introduce errors in the estimation of individual power consumption components. In order to highlight the importance of the self-consistent methodology with increasing subthreshold leakage (decreasing I_{on}/I_{off} ratio), the junction temperature is plotted as a function of V_{dd} for increasing values of I_{off} (Fig.9). It can be observed that as V_{dd} increases beyond 0.9 V, T_j begins to increase with increasing I_{off} . This is due to the fact that at higher V_{dd} the chip temperature begins to increase and couples more strongly with the

subthreshold leakage current resulting in higher P_{sw} and higher T_j . This figure also implies that for a leakage dominant technology, lowering of V_{dd} may not only lead to performance degradation as per traditional wisdom, but may also offset the degradation or even improve performance. This is due to the significant lowering of the junction temperature with V_{dd} .

Fig.10 plots the total self-consistent chip power as a function of operating frequency for two different I_{on}/I_{off} ratio. This again shows that at lower frequencies where the active power is low, the junction temperature and hence the subthreshold leakage power remain low. As frequency increases, the active power increases, which results in a corresponding increase in T_j and leakage power, and as expected, the effect is stronger for the leakier technology, resulting in a sharp increase in the total power, P_{chip} . Fig.11, which plots the junction temperature vs. the V_{dd} for $I_{on}/I_{off}=220$, illustrates the impact of using the self-consistent approach on satisfying chip-level iso-reliability constraint ($V_{dd} \leq V_{max} = T_j \cdot R + c$, c is a constant and R represents a chip-level reliability factor with typical value of $-3mV/^\circ C$). It can be observed that for $V_{dd} > 1.1V$, the non self-consistent method predicts lower values of T_j and hence, higher maximum allowable V_{dd} values (determined by the intersection of the curves with the iso-reliability line) and would therefore degrade the reliability of the chip. Fig.12 illustrates the implications of using the self-consistent methodology on some performance-power-cooling cost trade-offs. It can be observed from Fig.12a (for $\theta_j = 0.9^\circ C/W$), that the self-consistent method actually yields a higher frequency for a given P_{sys} . Similarly, as shown in Fig.12b, (for $\eta = 0.1$ and 0.3) a higher performance can be achieved for a given P_{sys} by employing the self-consistent approach. This is due to the fact that for a fixed P_{sys} (or P_{chip}) the self-consistent method always yields a lower T_j as can be observed from Fig.4, resulting in an improved I_{on} and frequency. Finally, Fig.13 provides design guidelines for efficient thermal management of the high-performance IC. This plot can be used for the selection of various packaging and cooling solutions (corresponding to values of θ_j , η and P_{sys}) to get optimum performance and also comprehend the diminishing returns of employing an expensive thermal management system. For example, for the microprocessor used in this study, a maximum frequency of ~ 3.4 GHz will be obtained even with $\eta = 0.8$ and $\theta_j = 0.2^\circ C/W$ along with $P_{sys} \approx 75$ W.

IV. Conclusion

In conclusion, a self-consistent junction temperature estimation methodology has been introduced in this paper for nanometer scale CMOS technologies that takes various electrothermal couplings between supply voltage, frequency, power dissipation, reliability and cooling-cost into account. This methodology will have significant implications for all leakage dominant sub-100 nm technologies and should be employed for various power-performance-reliability-cooling cost trade-offs and their optimization schemes.

Acknowledgement

This work was supported by Intel Corporation, Fujitsu Labs of America and the University of California-MICRO program.

References

1. P. P. Gelsinger, *ISSCC, Dig. Tech. Papers*, 2001, pp. 22-25.
2. R. Viswanath, et al., *Intel Technology Journal*, Q3, 2000.
3. S. Im and K. Banerjee, *IEDM*, 2000, pp. 727-730.
4. Y.-S. Lin et al., *IEEE TED*, Vol. 49, No. 6, pp. 1034-1041, 2002.
5. V. De and S. Borkar, *Proc. ISLPED*, 1999, pp. 163-168.
6. K. Banerjee and A. Mehrotra, *IEEE TED*, Vol. 49, No. 11, pp. 2001-2007, 2002.
7. A. Bar-Cohen, *Int. Work. Thermal Investigations of ICs & Systems* 1999, pp. 1-14.
8. U. Ghoshal and R. Schmidt, *ISSCC*, 2000, pp. 216-217.
9. A. Heydari, *Inter. Society Conf. on Thermal Phenomena*, 2002, pp. 371-378.
10. S. Narendra et al., *Proc. ISLPED*, 2002, pp. 19-23.




Fig.1. Power trends in high-performance microprocessors.




Fig.2. Simulated 100 nm NMOS subthreshold leakage current at different temperatures using BSIM3 models.




Fig.3. Leakage power scaling of an NMOS device based on SPICE simulations using BSIM3 models showing the effect of temperature. The normalization has been performed w.r.t. the leakage power at 130 nm node at 25°C.




Fig.4. Schematic diagram illustrating the dependence of junction temperature on the chip-to-ambient thermal impedance for typical and self-consistent power estimation methods.




Fig.5. Models for various metrics are expressed in functional format. Electrothermal couplings are indicated using broken lines.




Fig.6. An equivalent thermal circuit illustrating relationship between junction, ambient, and external temperatures in terms of chip and system level power dissipation and thermal impedances. Typically, $T_{amb} \approx 45^\circ\text{C}$, $T_{out} \approx 25^\circ\text{C}$. The value of θ_{sys} was estimated to be 0.285 °C/W for the IC.

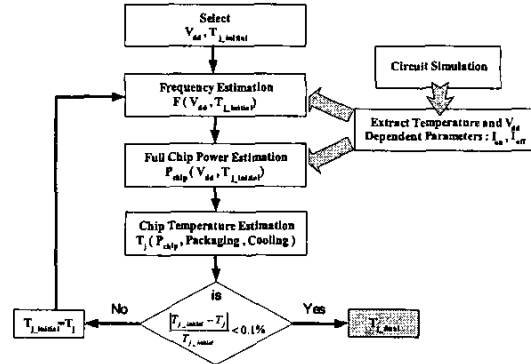


Fig.7. An overview of the self-consistent junction temperature estimation technique. The methodology has been implemented as an automated computer program.




Fig.8. Estimation of total power dissipation for different supply voltages for a high-performance chip using self-consistent and non self-consistent methods.






Fig.9. Self-consistent junction temperature vs. V_{dd} for different I_{off} .



a)



b)

Fig.10. Self-consistent power consumption as a function of frequency using a) $I_{on}/I_{off}=660$ and b) $I_{on}/I_{off}=220$.




Fig.11. Implications of self-consistent approach on satisfying chip-level reliability constraints.




Fig.12a. Frequency vs. system power consumption at different junction-to-ambient thermal impedance θ_{ja} . The efficiency η of dynamic cooling mechanism is 0.7. For $\theta_{ja} = 0.9$, the non self-consistent frequency is also plotted showing the advantage of employing the self-consistent method, i.e., for a given system power, the self-consistent approach yields a higher frequency.




Fig.12b. Frequency vs. system power consumption at different cooling efficiency (η). The junction-to-ambient thermal impedance, $\theta_{ja} = 0.85$. For $\eta = 0.1$ and 0.3, the non self-consistent frequency is also plotted showing the advantage of employing the self-consistent method.

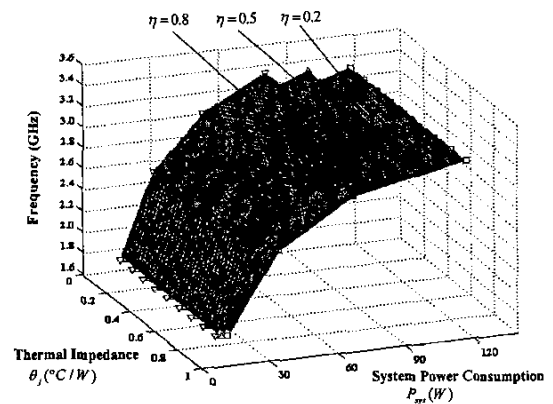


Fig.13. Design guidelines for integrated packaging and cooling solutions.