DEMONSTRATION OF POWER ENHANCEMENTS ON AN INDUSTRIAL CIRCUIT THROUGH DELAY MANAGEMENT OF NON-CRITICAL DATA PATHS

Dimitrios Velenis[1], *Kevin T. Tang*[2*], *Ivan S. Kourtev*[3], *Victor Adler*[4*], *Franklin Baez*[5], *and Eby G. Friedman*[1]

[1]Department of Electrical and Computer Engineering, University of Rochester, Rochester, New York 14627
[2]Broadcom Corporation, 2099 Gateway Place, San Jose, California 95110
[3]Department of Electrical Engineering, University of Pittsburgh, Pittsburgh, Pennsylvania 15261
[4]Sun Microsystems, 901 San Antonio Road, Palo Alto, California 94303
[5]Intel Corporation, 2200 Mision College Boulevard, Santa Clara, California 95052

ABSTRACT

On-chip power dissipation has become a fundamental design issue in high performance integrated circuits. A technique to significantly reduce the power dissipated in the non-critical data paths of an industrial circuit is demonstrated. The application of this technique with non-zero clock skew scheduling to the slower data paths is also described. Simulation results demonstrating the application of this technique to certain functional blocks of a high performance microprocessor are presented. A greater than 80% power savings is achieved in specific circuit blocks.

1. INTRODUCTION

The rapid scaling of device geometries in modern VLSI systems supports the system-on-a-chip integration of multiple subsystems, greatly increasing the amount of on-chip power dissipation. High power dissipation penalizes the overall system since more advanced packaging and heat removal technology are necessary. Additionally, wider on-chip and off-chip power busses, larger on-board decoupling capacitors, and more complicated power supplies are required. These factors increase the system size and cost. Furthermore, with the revolution of portable electronic devices, power dissipation has become a system level performance metric since the operation of these devices is limited by the battery life.

Two of the most popular techniques that are used to reduce power dissipation are supply voltage (V_{dd}) scaling and clock gating [1,2]. V_{dd} reduction is an effective way for reducing power since power dissipation is quadratically proportional to the power supply. The disadvantages of supply voltage scaling are effects such as subthreshold and gate oxide leakage and increased sensitivity to noise [3]. Clock gating reduces the capacitance being switched by the clock distribution network [2]. The major disadvantages of clock gating are the increased complexity of the timing analysis process and the greater current transients when large blocks of logic are switched on and off.

Another technique to reduce the dissipated power is the use of smaller size circuit elements from predesigned cell libraries in order to achieve significant power savings. The smaller sized elements introduce smaller load capacitances albeit with a small delay penalty [1]. When this technique is applied to non-critical data paths, the delay penalty has no impact on the overall performance of a synchronous system. A demonstration of the application of this technique to an industrial circuit is presented in this paper. It is shown that significant improvements in power dissipation can be achieved. Additionally, conditions to expand this technique to slower (more critical) data paths are discussed.

The paper is organized as follows. Background on the concept of the technique is presented in Section 2. The necessary conditions to apply this technique to slower data paths are described in Section 3. Simulation results that demonstrate the achievable power savings on an industrial circuit are presented in Section 4. Finally, some conclusions are offered in Section 5.

2. REDUCING POWER IN NON-CRITICAL DATA PATHS

In a large high performance system, such as a microprocessor, the number of critical data paths is small as compared with the total number of data paths in the system. For example, in a specific system described in [4], less than 5% of the total data paths are within 20% of the maximum path delay while more than 65% of the total data paths have path delays less than half of the maximum path delay. Alternatively, more than 65% of the local data paths are at least twice as fast as compared with the slowest local data paths. A similar distribution of path delays is common in the majority of high complexity circuits.

This research was supported by a grant from Intel Corporation.
*Dr. Kevin T. Tang and Dr. Victor Adler contributed to the development of this project during summer internships at Intel Corporation, prior to receiving their Ph.D. degree from the University of Rochester.

0-7803-6741-3/901/$10.00 © 2001 IEEE

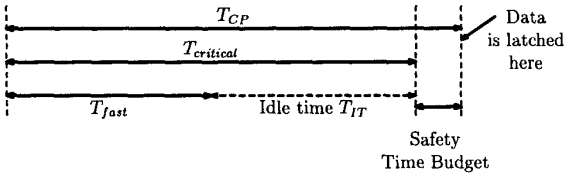

Figure 1: Short data path delay as compared to a critical long data path delay.

The fast data paths of a system are synchronized by the same clock signal that synchronizes the critical long data paths. Therefore, *idle time* (T_{IT}) exists in these short data paths since the data signal arrives at the final register well before the clock signal arrives at the same register, as shown in Fig. 1. This idle time can be exploited to slow down these short data paths in order to save power. One way to accomplish this technique is by downsizing (*i.e.*, decreasing the geometric width) of the latch R_i that drives the data path, as shown in Fig. 2(a), using smaller sized circuit cells from a predesigned cell library. By downsizing the latch the effective capacitance of the latch is decreased and the power required to drive the latch is reduced. Also, the geometric width of the output driver within the latch is decreased, thereby reducing the output current of the latch [5]. This procedure results in a decrease in power consumption, albeit with an increase in the data path delay.

There are constraints, however, that limit the minimum size of an output driver and thereby the additional delay that can be introduced. One constraint is that the additional delay should not exceed the maximum permissible path delay constraint as shown in Fig. 2(b). The summation of the initial data path delay $T_{initial}$, the additional delay T_{add}, and the safety time budget T_{safe} should be less (or, in the worst case, equal) to the clock period T_{CP}.

Another constraint is that the introduction of smaller sized output drivers should not degrade the signal rise and fall times below some target level. Due to the reduced size of the output driver, the output signal transition time of the latch is slower, increasing the short-circuit power dissipation within the gates that are driven by the latch. The short-circuit power dissipation is due to the current that flows directly from the power supply to the ground of a CMOS gate when the input voltage is within the range V_{tn} and $V_{dd}+V_{tp}$ (when both the PMOS and NMOS transistors are on). When the transition time of the input voltage is longer, the time during which both transistors are on is also longer, increasing the short-circuit power dissipation. A close approximation of the short-circuit power dissipation is given by [6]

$$P_{SC} = \frac{1}{2} I_{peak} t_{base} V_{dd} f, \qquad (1)$$

where I_{peak} depends upon the size of the transistors of the

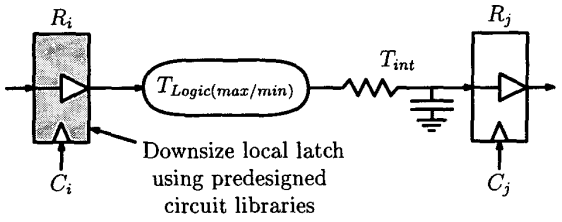

(a) A data path with a downsized latch to decrease the power of the fast data paths

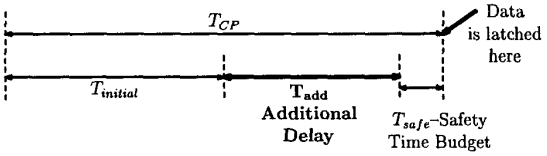

(b) The added delay of the fast data path does not violate the long path timing constraint

Figure 2: Increasing the delay of the fast data paths by downsizing the local latches that drive these paths.

driven gate, t_{base} is the input signal transition time, and f is the switching frequency of the input signal.

As shown by (1), as the size of the output buffer of the latch is decreased, the input signal transition time t_{base} increases, increasing the short-circuit power in the load gates. Therefore, there is a lower limit on decreasing the size of the output driver to achieve less power.

3. CLOCK SKEW SCHEDULING APPLIED TO CRITICAL DATA PATHS

The concept of slowing down fast data paths in order to save power can be further applied to slower, more critical data paths with the aid of *non-zero clock skew scheduling* [4, 7]. Given two sequentially-adjacent registers, R_i and R_j, the clock skew between these two registers is defined as $T_{skew} = T_{C_i} - T_{C_j}$, where T_{C_i} and T_{C_j} are the clock delays from the clock source to the registers, R_i and R_j, respectively. If the clock delay to the initial register T_{C_i} is greater than the clock delay to the final register T_{C_j}, the clock skew is described as positive. Similarly, if the clock delay to the initial register T_{C_i} is less than the clock delay to the final register T_{C_j}, the clock skew is described as negative. As described in [4, 7, 8], the individual clock skew for each data path should satisfy the following two constraints:

$$T_{CP} \geq T_{skew_{pos}} + T_{PD_{max}}, \qquad (2)$$

$$T_{PD_{min}} \geq T_{skew_{neg}} + T_{hold}. \qquad (3)$$

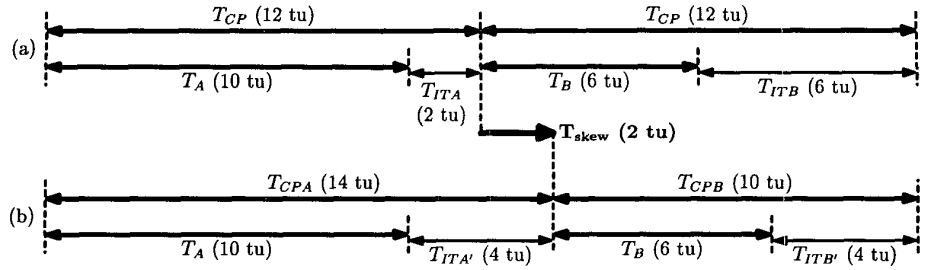

Figure 3: Application of local clock skew to equalize the available idle time between the long and short delay data paths. (a) Initial timing of the data paths. (b) Timing of the data paths after the application of local clock skew

$T_{skew_{pos}}$ ($T_{skew_{neg}}$) is the positive (negative) clock skew [8] between the registers of a data path, and $T_{PD_{max}}$($T_{PD_{min}}$) is the maximum (minimum) propagation delay of a data path. T_{CP} is the minimum clock period and T_{hold} is the amount of time that the input data signal must be stable at the input of register R_j once the clock signal changes state.

By applying negative clock skew to the slower, more critical data paths, the idle time in these data paths can be increased, permitting these paths to be further slowed down. However, there is one condition that must be satisfied for this concept to be feasible. This condition is that the data path that follows the slow data path should be sufficiently fast to satisfy the constraint set by (2). An example of the application of this concept to long data path delays is shown in Fig. 3. As shown in Fig. 3(a), data path A has a long delay of T_A = 10 time units (tu) and data path B has a short delay of T_B = 6 tu. The clock period of the system is T_{CP} = 12 tu. Because the delay of data path B is short as compared to the target clock period, the clock signal that controls the latching operation of the register located between data paths A and B can be delayed by 2 tu, as shown in Fig. 3(b). This strategy delays the data signal propagating into data path B without creating any timing hazards, satisfying $T_{CPB} = T_{CP} - T_{skew} \geq T_B + T_{ITB\prime}$. Alternatively, delaying the arrival of the clock signal at the register delays the latching of the data signal that propagates into data path A, adding more idle time to data path A. Therefore, both data paths have sufficient idle time, permitting the drivers to be downsized so as to reduce the power dissipation of the overall circuit. If the slow data path A is not further slowed down, the application of non-zero clock skew can increase the safety margin of data path A which can be used to relax the strict timing constraints, making the circuit less sensitive to process parameter variations [9].

The approach presented above and illustrated in Figs. 3(a) and 3(b) provides an additional technique for saving power through the application of negative clock skew. The negative clock skew across data path A can be produced either by inserting a delay element along the clock line that distributes the clock signal to the final register of data path A, or by decreasing the transistor size of the clock buffer that drives this clock line. In the latter case, decreasing the size of the clock buffer results in less output current, providing an additional savings in power.

4. APPLICATION TO AN INDUSTRIAL CIRCUIT

In a joint research project between the University of Rochester and Intel Corporation, the process of decreasing the power dissipation of an industrial circuit [10] by changing the timing of the non-critical data paths by applying non-zero clock skew scheduling has been investigated. The technique described in this paper has been applied to specific functional unit blocks (FUBs) within a high performance microprocessor. A graph representing one of these FUBs is shown in Fig. 4 with normalized maximum and minimum local data path delays. All of the timing information in the analysis that follows is described in terms of these normalized path delays.

The concept of slowing down a data path in order to save power has been applied to the fast data paths, B, C, and D, of the FUB shown in Fig. 4. Each of these data paths is slowed down by downsizing the driving latch R_2 by using a different latch from an existing circuit library. The maximum and minimum delay of these data paths prior to and after decreasing the size of the data path driver is listed in Table 1. As listed in Table 1, the delay of these data paths is increased on average by 21.6%.

The effect of downsizing the local latches that drive the data paths is to substantially reduce the power dissipated within the circuit block that contains these latches. As shown

Table 1: Comparison between the original and the increased delay of data paths B, C, and D within the FUB illustrated in Fig. 4

Data path	Original max/min data path delay (tu)	Increased max/min data path delay (tu)	Increased Delay (%)
B	(21/19)	(25/21)	14.7
C	(20/16)	(25/20)	22.5
D	(19/17)	(25/21)	27.5
	Average Increase in Delay (%)		21.6

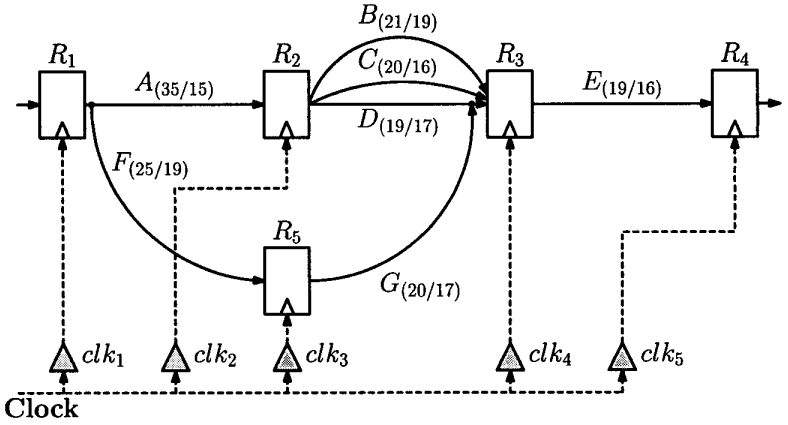

Figure 4: Circuit graph of an industrial FUB with normalized data path delays

Table 2: Normalized power dissipation within the circuit block containing the latches

No optimization	Downsize latches w/o clock scheduling	Downsize latches w/ clock scheduling
100	18	17.3

in Table 2, the total power dissipation of the circuit block is reduced by 82% by downsizing a total of 69 latches.

The remaining data paths within the FUB are unchanged. The effect of changing the latch on the data signal rise and fall times in data paths B, C, and D is negligible. Also, no maximum data path delay constraint is violated since the larger maximum delay of the affected data paths (25 tu) is less than the maximum delay of the most critical data path (35 tu). Since the difference between the delay of data path A and the delays of the data paths, B, C, and D, is significant, the circuit performance can be improved with the application of non-zero clock skew scheduling, as described in [4, 7]. In this case, the clock period can be reduced to $T_{CP} = 25 + \frac{35-25}{2} = 30$ tu. The performance of the circuit can therefore be further enhanced by approximately 14%. Furthermore, the application of negative clock skew using downsized clock buffers results in an additional 4% decrease in power dissipation as listed in Table 2. This decrease is due to the reduced capacitance of the clock buffer clk_2 that drives the downsized latches.

5. CONCLUSIONS

A strategy for decreasing the power dissipation by reducing the size of the driving latches and increasing the delay of the non-critical data paths has been demonstrated. The constraints, advantages, and disadvantages have been discussed. The application of non-zero clock skew scheduling to increase the idle time of slower data paths has also been presented. Simulations of specific FUBs within a high performance industrial microprocessor demonstrate that a substantial reduction in power of up to 82% can be achieved by applying this strategy.

6. REFERENCES

[1] V. Tiwari, D. Singh, S. Rajgopal, G. Mehta, R. Patel, and F. Baez, "Reducing Power in High-Performance Microprocessors," *Proceedings of the IEEE/ACM Design Automation Conference*, pp. 732-737, June 1998.

[2] L. Benini and G. De Micheli, *Dynamic Power Management of Circuits and Systems: Design Techniques and CAD Tools*, Norwell, Massachusetts, Kluwer Academic Publishers, 1998.

[3] K. Chen and C. Hu, "Performance and V_{dd} Scaling in Deep Submicrometer CMOS," *IEEE Journal of Solid-State Circuits*, Vol. SC-33, No. 10, pp. 1586-1589, October 1998.

[4] I. S. Kourtev and E. G. Friedman, *Timing Optimization Through Clock Skew Scheduling*, Norwell, Massachusetts: Kluwer Academic Publishers, 2000.

[5] V. Adler and E. G. Friedman, "Repeater Design to Reduce Delay and Power in Resistive Interconnect," *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, Vol. CAS II-45, No. 5, pp. 607-616, May 1998.

[6] V. Adler and E. G. Friedman, "Delay and Power Expressions for a CMOS Inverter Driving a Resistive-Capacitive Load," *Analog Integrated Circuits and Signal Processing*, Volume 14, No. 1/2, pp. 29-39, September 1997.

[7] J. P. Fishburn, "Clock Skew Optimization," *IEEE Transactions on Computers*, Vol. 39, No. 7, pp. 945-951, July 1990.

[8] E. G. Friedman, *Clock Distribution Networks in VLSI Circuits and Systems*, Piscataway, New Jersey: IEEE Press, 1995.

[9] J. L. Neves and E.G. Friedman, "Optimal Clock Skew Scheduling Tolerant to Process Variations," *Proceedings of the ACM/IEEE Design Automation Conference*, pp. 623-628, June 1996.

[10] U. Desai, S. Tam, R. Kim, and J. Zhang, "Itanium Processor Clock Design," *Proceedings of the ACM/SIGDA International Symposium on Physical Design*, pp. 94-98, April 2000.