

Interconnect-Power Dissipation in a Microprocessor

Nir Magen
Intel Israel (74) Ltd.
Mobile Platform Group
Haifa 31015, Israel
(972) 48656830
nir.magen@intel.com

Avinoam Kolodny
Electrical Engineering Dept.
Technion
Haifa 32000, Israel
(972) 48294764
kolodny@ee.technion.ac.il

Uri Weiser
Intel Israel (74) Ltd.
Corporate Technology Group
Petach-Tikva 49527, Israel
(972) 39207246
uri.weiser@intel.com

Nachum Shamir
Intel Israel (74) Ltd.
Mobile Platform Group
Haifa 31015, Israel
(972) 48655913
nachum.shamir@intel.com

ABSTRACT

Interconnect power is dynamic power dissipation due to switching of interconnection capacitances. This paper describes the characterization of interconnect power in a state-of-the-art high-performance microprocessor designed for power efficiency. The analysis showed that interconnect power is over 50% of the dynamic power. Over 90% of the interconnect power is consumed by only 10% of the interconnections. Relations of interconnect power to wire length distribution and hierarchy level of nets were examined. In light of the results, a router's algorithms were modified, to use larger wire spacing and minimal length routing for the high power consuming interconnects. The power-aware router algorithm was tested on synthesized blocks, demonstrating average saving of 14% in the dynamic power consumption without timing degradation or area increase. The results demonstrate the obtainable benefits of tuning physical design algorithms to save power.

Categories and Subject Descriptors

B.7.1 [Integrated Circuits]: Design Styles – *Microprocessors*;
B.7.2 [Integrated Circuits]: Design Aids – *Placement and Routing*.

General Terms

Performance, Design.

Keywords

Interconnect power, low-power design, routing, wire spacing.

1. INTRODUCTION

Power dissipation of high-performance microprocessors is becoming a limiting factor and hence design for efficient power consumption is becoming a major design consideration. Dynamic power is currently the main component of the power dissipation [1]. Dynamic power consumption due to periodical switching of capacitors is approximated by the well-known expression :

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
SLIP '04, February 14-15, 2004, Paris, France.
Copyright 2004 ACM 1-58113-818-0/04/0002...\$5.00.

$$P = \sum AF_j \cdot C_j \cdot V^2 \cdot f \quad (1)$$

Where AF_j and C_j are the Activity Factor (AF) and capacitance for the j^{th} signal. V is the supply voltage and f is the clock frequency. In this calculation we neglect the short-circuit current, which is later added by using an overall factor of about 10% as shown by [2]. This paper focuses on Interconnect Power, i.e. energy dissipation due to the switching of interconnection capacitances, which are part of the total switched capacitance of each net– C_j .

Previous studies have commented on the growing significance of the interconnect power [3] [4]. Various methods were used to reduce the power consumption. Most efforts were invested in voltage reduction [5] and frequency optimizations, such as [6], gate sizing [7] and clock gating [8]. However no direct design effort to reduce the interconnect switched capacitance is known to the authors.

In this work we study the role of interconnect power in the overall dynamic power consumption. We quantify the magnitude of this significant component and characterize the top power consuming interconnections in order to detect power saving opportunities. We also study methods to reduce the interconnect power consumption by tuning and optimization of routing algorithms. Our approach is based on a detailed case-study of a recent microprocessor designed for power-efficiency.

In Section 2 we present our methodology for estimation and extraction of the interconnect power component. Results of the case study analysis are described in Section 3, along with possible directions for power reduction in interconnect design. As a proof of concept, we developed interconnect-power-aware router algorithms and performed design experiments that are described in Section 4. Ideas for future work are discussed in Section 5 and the conclusions are summarized in Section 6.

2. POWER EVALUATION METHODOLOGY

The interconnect power analysis was performed on a state-of-the-art microprocessor designed for power-efficiency, consisting of 77 million transistors, fabricated in 0.13 μm technology. Dynamic power dissipation was analyzed using a Stochastic Dynamic Power estimation (SDPE) technique [9]. Activity factors for all signals were extracted for 32 high-power and focused stress tests by SDPE simulations. The power data generated by the SDPE showed excellent correlation to silicon measurements of the same tests, showing differences of less than 5% of the total power. In order to analyze the interconnect power

consumption we generated a study database that included structural and power attributes for the entire processor core (excluding the second level cache, since it is very large, highly regular and consumes negligible amount of dynamic power). The stored information included the following data for each signal net:

Interconnect Length summing all metal segments between the drivers and receivers. The net length was extracted from the layout data. Repeater separated segments were considered as part of the original net. The interconnect length based analysis doesn't include the global clock grid, because its unique grid layout makes its total length singular and quite meaningless. Local clock signals, that branch out of the clock grid are included though.

Capacitances summing all types of capacitive loads, including diffusion capacitances of drivers, capacitances of the metal wiring and gate load of the receivers. Repeater gate and diffusion capacitances were added to the original net.

The metal capacitance includes cross-capacitance to neighboring nets, with a Miller factor of 1. When neighboring wires switch simultaneously the energy associated with cross-capacitance may be zero (if both wires switch in the same direction), or double the energy of independent transitions (if they switch in opposite directions). This is similar to the effect of crosstalk on delay. In delay analysis the nets are often decoupled and the cross-capacitance is multiplied by a "Miller Factor" [10]. This is justified since delay analysis seeks a worst-case delay failure. However, unlike delay which is localized in space and time, dissipated power is a cumulative parameter which integrates many transitions by all signals. Therefore, it is reasonable to assume that random simultaneous transitions average-out, using an average Miller factor of 1 (meaning that the extracted cross-capacitance are not multiplied by any factor).

Fan out counting all gate-connected transistors, for modeling simplicity (e.g. an inverter load is considered to be a fan out of 2).

Activity Factor average calculated for all SDPE tests.

Design hierarchy data separating the signals into *local* and *global* sets. Nets inside a functional unit block (FUB) are considered to be *local nets* while those interconnecting the blocks are *global nets*.

Miscellaneous characterization data including: The metal layers information and net classification (clock, signal nets and others).

Using these attributes, we could analyze various aspects of interconnect-related contributions to the total dynamic power expressed in equation (1).

3. MICROPROCESSOR ANALYSIS RESULTS

3.1 Interconnect-Length Based Analysis

The interconnections attributes were examined to reveal dependencies on wire length. The global clock grid is excluded from this analysis. First the wire length distribution in the processor core was analyzed, as in [1]. The number of nets versus the net length is plotted over the figure published in [1], shown here as Figure 1.

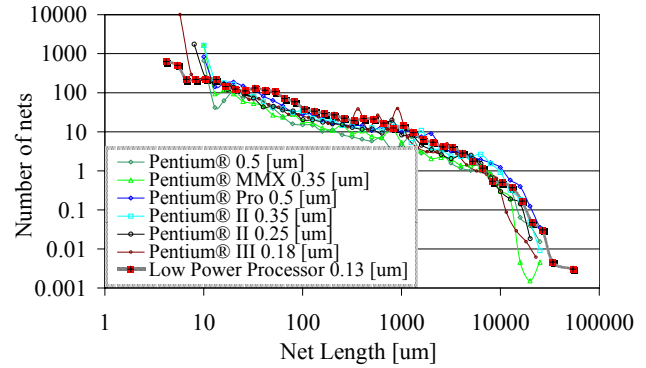


Figure 1. Number of Nets vs. Net Length
Seven processors display the same distribution.
The analyzed processor behaves similarly to previous models

The number of nets decreases exponentially with net length. The same distribution is exhibited by seven different processors, of five different technology generations. The examined processor is plotted by the thickest curve, it has slightly fewer shorter nets – due to the exclusion of the second level cache (L2 cache). The analyzed processor has long nets longer than the other processors, since all the other processors net lengths are approximated by a bounding box metric, as opposed to the sum-of-segment-lengths metric used in our analysis. The data supports the net length distribution predictions of the models derived in [11] and [12].

The number of nets decreases exponentially along with an exponential increase of the net length, such that between 100 to 1000 μm the total number of nets decreases only by a factor of 5. There are orders of magnitude more short nets than longer nets, however the majority of the wiring length is associated with the long nets. The longer nets are mostly routed on higher metal layers with greater capacitance per unit length. Therefore the long nets are responsible for the majority of the interconnect capacitance. The longer nets have larger average fan out (shown in Figure 4). This associates the longer nets with larger load capacitance, hence increasing the long nets total switched capacitance. The entire net capacitance increases with the net length, and also the relative amount of interconnect capacitance increases with the net length.

Dynamic power was computed by summation according to Equation (1) of nodes capacitances multiplied by the corresponding activity factors. The Supply voltage and frequency are fixed in this analysis. The activity factors seem to be nearly evenly distributed among all the nets lengths. The dynamic power plotted versus the net length is shown in Figure 2. Figure 2 displays two major peaks of power, one at the short nets region and another at the medium and long range. Separation of the power into the capacitance types contributing to it (Figure 2) reveals that over 50% of the dynamic power results from interconnections switching, the gates contribute another 34% and the rest is diffusions. The interconnect power component increases with net length are. The activity factors show no correlation to the interconnect length.




Figure 2. Dynamic Power vs. Net Length
The power has two peaks, correlated with total wiring.
Columns consist of gate, diffusion and interconnect power

Hierarchy Based Interconnect Analysis

The interconnections are separated into two groups based on the design hierarchy (local and global interconnects), as a first order approximation. The results shown in the previous paragraph are refined using the separation of local and global interconnections defined in section 2. The goal of this analysis is to examine the relations between the interconnection properties and design hierarchy, thus further characterizing the interconnect power.

The number of nets versus the net length is redrawn in Figure 3, separating the contributions of local and global nets.




Figure 3. Length distribution of local nets and global nets

Local wires are the great majority of the nets shorter than 1000 μm , while global interconnections dominate the long nets side of the graph. This figure supports the model of [11] where wire length distribution is represented as sum of hierarchy-based Gaussian distributions.

The average fan outs of the two categories are also distinct. Average fan outs vs. net length are plotted in Figure 4. The entire short point-to-point region (up to 100 μm) is local. Among longer nets, the local interconnections appear to have higher fan out than global nets of the same length. Both curves can be fitted by simple power laws with high accuracy. However, as noted for wire load models predicting the average net capacitance based on its fan out, there exists a large scattering of the nets around the average.

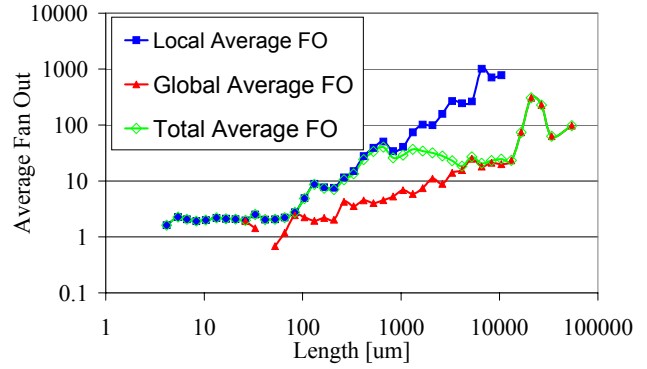


Figure 4. Average Fan Out vs. Net Length
Local nets have higher average fan out for any given net length.

This can cause great errors for single nets, although correct on average. The power laws predicting the fan out are correct on average, which makes them suitable for power predictions using a sum metric.

The capacitance versus net length data is separated into local and global interconnections distributions in Figure 5. The local and global interconnections display totally different characteristics. The local nets have about 25% of interconnect capacitance in the point-to-point region and about 40% above that region. In global nets, the interconnect capacitance component is about 80%.

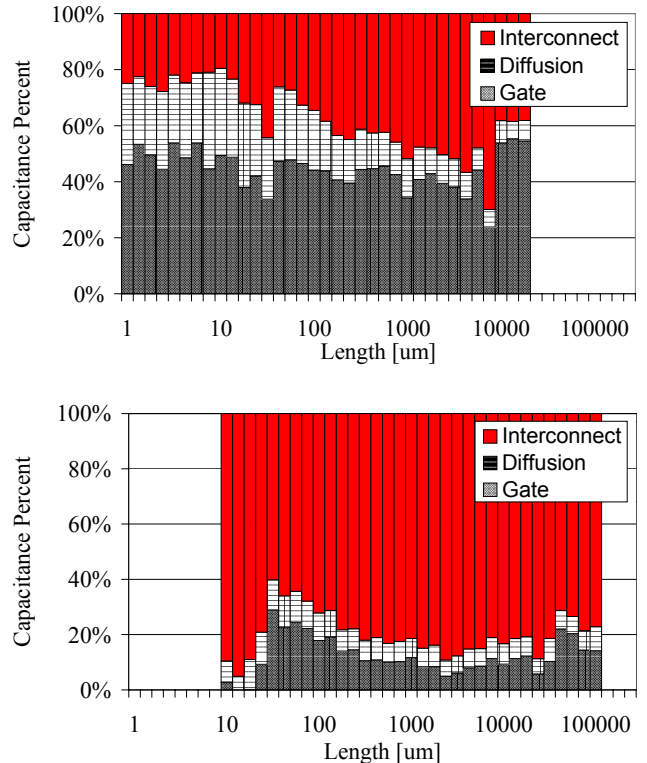


Figure 5. Capacitance Types vs. Net Length
Top) Local nets, Bottom) Global nets.
The types are gate, diffusion and interconnections

The dynamic power vs. net length of Figure 2 is broken into local and global interconnect types, displayed in Figure 6.




Figure 6. Total Dynamic Power vs. Net Length.
The short nets peak is entirely local, while the other peak consists of both local and global nets.

The relation of the two power peaks shown in Figure 2 to the hierarchy level of the interconnection is evident. The power peak due to short nets is entirely in local interconnections, while the other peak consists of both local and global nets, with roughly equal shares. The global nets consume about 34% of the dynamic power and the local nets dissipate the remaining 66%. The interconnect-power is the majority of the long nets peak, where it contributes over 60% of the power, while in the short nets peak the interconnect-power dissipation is only about 30% of the power. The total interconnect-power (that is about 50% of the dynamic power in the processor) is divided into two nearly equal parts of global nets and local nets interconnect-power, each part is approximately 25% of the total dynamic power. The same ratios seen in Figure 5 for the capacitances apply for power as well, since the activity factors distribution is uniform.

The relative portion of global nets power out of the total dynamic power depends on the test benchmark used. The percentage of global nets power is plotted versus the test's total normalized power consumption in Figure 7.

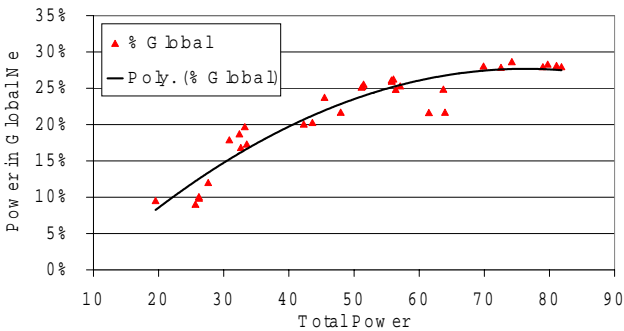


Figure 7. Global Nets Power vs. Total Test Power.
High power tests have high percentage of global nets power.

The tests with total high power consumption also have high percentage of global nets power, compared to tests consuming less power. This means that global communication is important in worst-case high power consumption analysis.

3.2 Clock and Signals Analysis

The previous analyses, based on the interconnect length and the separation into local and global nets excluded the global clock grid. The clock is known to be a major power-consuming net. The local and global clock nets are about 1% of the nets, their length is about 4% of the routing length, however their power consumption is not minor. The total dynamic power breakdown into local and global, signals and clock nets is shown in Figure 8.

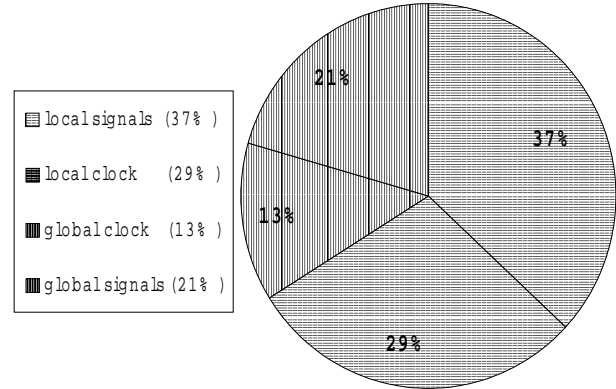


Figure 8. Total Dynamic Power Breakdown.

The breakdown reveals that the local nets consume the majority (over 60%) of the total dynamic power. The clock power is about 40% of the total dynamic power. The majority of the clock power is consumed due to local clock nets (including drivers and loads). The total power breakdown motivates investing power optimization efforts in the local interconnections power first.

3.3 Interconnect Power - Summary

The interconnect power analysis reveals that the interconnections are responsible for over 50% of the dynamic power consumption, as shown in Figure 2. The interconnections attributes, nets distribution, capacitance breakdown, average fan out and power are correlated to the net length. The activity factors of nets are uncorrelated with all these attributes. The local and global interconnections have been characterized showing that interconnections at different hierarchy levels have different attributes and might require different design approaches. For example, power-aware pre-routing of some global power-critical nets can be effective because about 50% of the interconnect-power is consumed by (relatively few) global wires. The important global interconnect-power can also be addressed by choice of architecture, as suggested in [13], [14] and [15]. In local signals, where the power is the sum of great numbers of nets, optimizing the average case seems to be a reasonable approach for interconnect-power saving. However, special local nets such as local clock nets must be treated separately. The separation of the nets into clock and signal nets shows that the clock is the largest consumer, occupying relatively small amount of routing area making it attractive to target for power optimization. All of the trends displayed in this analysis were captured by simple predictive models, however they are beyond the scope of this paper and require further study and validation. Based upon the analysis performed, several interconnect power saving opportunities emerged.

4. POWER-AWARE ROUTING ALGORITHMS EXPERIMENTS

4.1 Interconnect Capacitance Reduction

Experiments were performed in order to demonstrate the potential of power-aware interconnect physical design algorithms for capacitance reduction of the highest power consuming wires. A simple rule of the thumb that was discovered based on the blocks of the processor analyzed, shows that 90% of the total dynamic power is consumed by 10% of the wires. This leads to assuming that applying wire capacitance reduction techniques to a small percentage of the wires, can save the majority of the interconnect power. Great potential appeared to be in interconnect power-aware routing of local nets, affecting the majority of the dynamic power, as shown in Figure 8, by interconnection capacitance reduction. Two possible methods for capacitance reduction suggested are:

1. Interconnect length reduction – For the entire processor, wires lengths are about 30% longer than the lower bound of optimal routing length (length of minimum rectilinear spanning tree).
2. Interconnect spacing – The majority of the wires are routed at minimal spacing, as defined by the process design rules. Increasing the spacing between neighbor wires will reduce capacitance and power dissipation. Using the [16] 0.13 μm technology parameters, assuming that wire width and spacing are equal, and modeling the interconnect capacitance by [17], shows that increasing the wire spacing can reduce total capacitance by up to 40%, as shown in Figure 9, where wire capacitance for a global wire is plotted vs. the wire spacing.




Figure 9. Interconnect Capacitance vs. Wire Spacing, 0.13 μm technology, upper metal wire. Wire spacing is measured in multiples of minimum design rule.

It is visible that wire spacing quickly enters the region of diminishing returns, the exact trade-off between area and power reduction should be established. Previous work of [18] applied wire spacing and sizing for delay optimizations, and [19] utilized wire spacing of clock networks for skew reduction. Both commented that wire spacing reduced the

interconnect power consumption, it was however a by-product and not a design optimization goal.

4.2 Interconnect Power-Aware Routing Algorithms

Interconnect power reduction algorithms were implemented as extensions to an existing industrial router. Interconnect length reduction for the top power consuming nets was achieved by giving routing priority to selected nets. Some routing area was allocated for spacing of certain high power consuming nets. Both methods have some negative effects, increasing the congestion of other wires, thus increasing the cross capacitance between wires and the interconnect power consumption. In extreme cases, it becomes impossible to route with the given resources. Due to the negative effects, the wire capacitance reduction techniques were applied to a limited number of power-critical nets. The clock network is a major power consumer as seen in Figures 8. Clock nets have a fixed topology in this router because of clock skew considerations. Therefore wire spacing alone was applied to the clock network. The wire spacing chosen is twice the minimum design rule, which appears to be an acceptable trade off between routing area consumption and power saving. The percentage of power critical nets for routing priority was determined by empirical testing per block (under 5% of the nets).

The simplified flow of the interconnect power-aware router is described by figure 10:

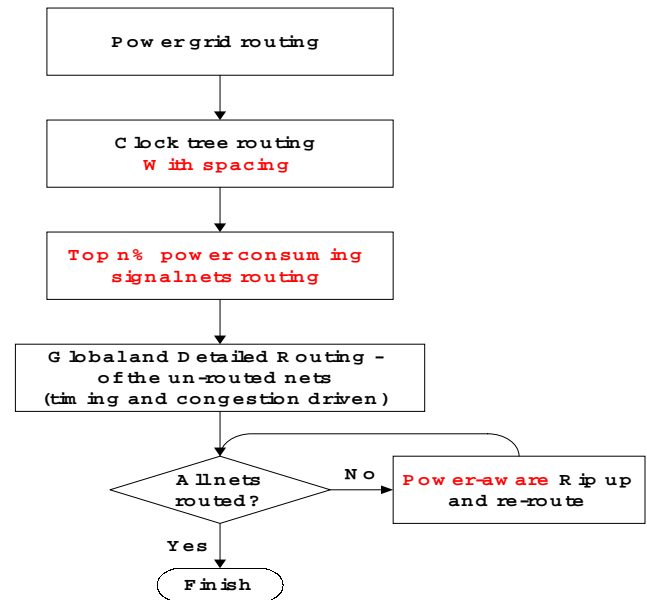


Figure 10. Interconnect Power-Aware Router flow

The router applies the common technique of rip-up and reroute [20] when encountering wires blocking each other. It turns out that the rip-up can be harmful to the power consumption if it rips up power-critical nets. To avoid this drawback, the rip-up mechanism was enhanced, allowing it to select a power-critical net only if no other blocking candidate exists. The rip-up and reroute loop is bounded, if it fails to improve the routing after several iterations the router ends the flow with some open nets. The power-aware router flow is followed by a driver resizing stage based on post-routing parasitic extraction and timing analysis. For every net whose capacitance was reduced by the

router, the drivers of the net are downsized, without degrading the timing performance. The downsizing saves both area and power consumption at the drivers.

4.3 Experimental Results

The power-aware router flow was used to route blocks of the analyzed processor. Five blocks of various sizes, chosen out of different sections of the chip, were selected for the experiments. The original placement of the cells was used. The dynamic power analysis performed on the original routed blocks was used to set priorities for the power-aware router algorithm. The same timing specifications (required and arrival times) were applied to the power-aware designs, too. The power-aware designed blocks didn't violate timing requirements. The power-aware router flow resulted in considerable dynamic power reduction for all the blocks analyzed. The experiment results are shown in figure 11. For an average block, the flow should reduce about 14% of the dynamic power without any performance penalty. This amount is projected based on the average interconnect capacitance, power and clock-to-signals ratio. These experiments show that significant power-saving potential exists in the interconnections design.

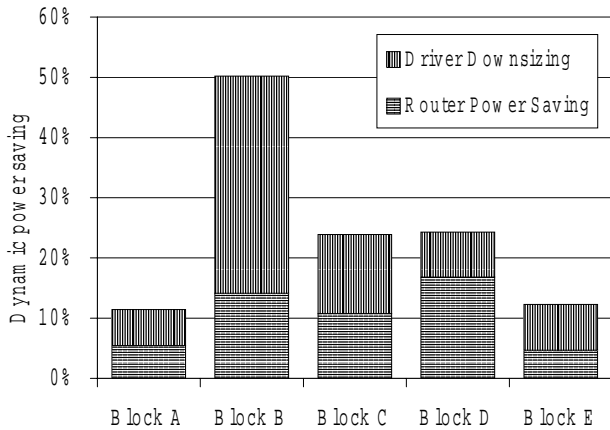


Figure 11. Dynamic Power Savings in Power-Aware Design Experiments.

Power saving (lower) portion is due to the router, the (upper) portion is due to the driver downsizing.

5. FUTURE WORK

The interconnect power analysis presented is based on a single case study. Further study of other processors should take place to verify the results. ITRS [16] projects an increase in the interconnect capacitance (relative to gate and diffusion capacitances) in future technology generations, relative to the total switched capacitance. Future generation processors should be evaluated to determine if the importance of interconnect power increases with technology scaling. Modeling and prediction of the interconnect power are required to evaluate these phenomena.

The interconnect power-aware flow shows that physical design can save considerable amounts of dynamic power by interconnect optimizations. The highest potential for further interconnect-power saving is probably found in interconnect-power-aware placement. Both placement and routing algorithms should be enhanced for better interconnect power saving

techniques. The power-awareness should not be added as an extension (as in our study), the entire physical design optimization should consider simultaneously timing, routing congestion and interconnect power to achieve better overall results.

6. CONCLUSION

In this paper we analyzed the dynamic power consumption in a state-of-the-art microprocessor, designed in 0.13 μm CMOS technology for power-efficient operation. The results indicate that interconnect switching consume over 50% of the overall dynamic power. The local and global net topologies were analyzed showing major difference in length distribution, average fan out and capacitance components associated with gates, diffusions and interconnections. For local and global nets of the same length, the local net has typically higher fan out and significantly lower interconnect capacitance component (30% vs. 80%). However the interconnect power is divided fairly equally between local and global nets, because there are much more local nets than global nets. No differences were found in activity factors between local and global nets, for averaged tests. Clock nets consume about 40% of the total power, two thirds of which is consumed by local nets.

90% of the dynamic power is consumed by about 10% of the nets, however these nets show great variance in the attributes of wire length, fan out, activity factor and design hierarchy level. Power reduction for these nets is a major design challenge.

Many interconnect-power saving opportunities exist at all levels of the design from architecture definitions minimizing the global communication to physical design starting at partitioning, floor planning, placement and routing and post-layout optimizations. As a proof-of-concept an existing industrial router was modified, making its routing algorithms interconnect-power-aware. This router was used to redesign five blocks of the analyzed processor, resulting in an average saving of 14% of the dynamic power, demonstrating that Considerable dynamic power dissipation reduction is achievable by interconnect-power-aware physical design.

7. ACKNOWLEDGMENTS

The authors acknowledge the support and information sharing by Shekhar Y. Borkar.

8. REFERENCES

- [1] Borkar, S. Y. Design Challenges of Technology Scaling, IEEE MICRO, Jul./Aug. 1999, Vol. 19, No. 4, 23-29.
- [2] Chatterjee, A., Nandakumar, M., and Chen, I.C. An Investigation of the Impact of Technology Scaling on Power Wasted as Short-Circuit Current in Low Voltage Static CMOS Circuits. in International Symposium on Low Power Electronic Design, 1996, 145-150.
- [3] Meindl, J. D., Davis, J. A., Zarkesh-Ha, P., Patel, C. S., Martin, K. P., and Kohl, P. A. Interconnect Opportunities for Gigascale Integration. IBM J. Res. & Dev., vol. 46, Mar/May 2002, 245-263.

- [4] Rabaey, J., Chandrakasan, A., and Nikolic, B. *Digital Integrated Circuits: Second Edition*. Prentice Hall, 2002.
- [5] Liu, D., Svensson, C. Trading Speed for Low Power by Choice of Supply and Threshold Voltages. *IEEE Journal of Solid-State Circuits*, Vol. 28, No. 1, Jan. 1993.
- [6] Pering, T., Burd, T., and Brodersen, R. The Simulation and Evaluation of Dynamic Voltage Scaling Algorithms. in *Proceedings of the 1998 International Symposium on Low Power Electronics and Design*, Aug. 1998, 76-81.
- [7] Berkelaar, M.R.C.M., and Jess, J.A.G. Gate Sizing in MOS Digital Circuits with Linear Programming. in *Proceedings of the European Design Automation Conference*, Mar. 1990, 12-15.
- [8] Farrahi, A.H., Chen, C., Srivastava, A., Tellez, G., and Sarrafzadeh, M. Activity-Driven Clock Design. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, Vol. 20, No. 6, Jun. 2001.
- [9] Genossar, D., and Shamir, N., Intel® Pentium® M Processor Power Estimation, Budgeting, Optimization, and Validation. *Intel Technology Journal*, Vol. 07, Iss. 02, May 2003, 43-50.
- [10] Arunachalam, R., Rajagopal, K., and Pileggi, L.T. TACO: Timing Analysis with Coupling. in *Proceedings of the Design Automation Conference*, Jun. 2000, 266-269.
- [11] Stroobandt, D., Van Campenhout, J. Accurate Interconnection Length Estimations for Predictions Early in the Design Cycle. *VLSI Design*, Vol. 10 (1), 1999, 1-20.
- [12] Davis, J., De, V.K., and Meindl, J. A Stochastic Wire-Length Distribution for Gigascale Integration (GSI) – Part I: Derivation and Validation. *IEEE Transaction on Electron Devices*, Vol. 45, No. 3, Mar. 1998, 580- 589.
- [13] Dally, W. J., and Lacy, S., *VLSI Architecture: Past, Present, and Future*. in *Proceedings of the Advanced Research in VLSI conference*, Atlanta, GA, 1999.
- [14] Horowitz, M., Ho, R., and Mai, K. The future of wires. in *Proceedings of the IEEE*, Vol. 89, no. 4, Apr. 2001.
- [15] Ronen, R., Mendelson, A., Lai, K., Lu, S-L., Pollack, F., Shen, J.P. Coming Challenges in Microarchitecture and Architecture. *Proceedings of the IEEE*, Vol. 89, No. 3, Mar. 2001.
- [16] *International Technology Roadmap for Semiconductors (2001 Edition)*. Available: <http://public.itrs.net/Files/2001ITRS/Home.htm>
- [17] Chern, J.H., Huang, J., Arledge, L., Li, P.C., and Yang, P. Multilevel Metal Capacitance Models for CAD Design Synthesis Systems. *IEEE Electron Device Letters*, vol. 13, Jan. 1992, 32-34.
- [18] Cong, J., He, L., Koh C.K., and Madden, P. Performance Optimization of VLSI Interconnect Layout. *Integration, the VLSI Journal*, Vol. 21, Nov. 1996, 1-94.
- [19] Lee, Y.-M., Lai, H.Y., Chen, C.C.-P. Optimal Spacing and Capacitance Padding for General Clock Structures. in *Proceedings of the conference on Asia South Pacific Design Automation Conference*, Jan. 2001, 115-119.
- [20] Dees, W.A., and Smith, R.J. Performance of Interconnection Rip-Up and Reroute Strategies. in *Proceedings of the 18th Design Automation Conference*, Jun. 1981, 382-390.