

IBM Research Report

Characterization of the Impact of Interconnect Design on the Capacitive Load Driven by a Global Clock Distribution

G. G. Lopez*, Giovanni Fiorenza, Thomas J. Bucelot, Phillip J. Restle,
Mary Yvonne Lanzerotti

IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

*Georgia Tech



Research Division

Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

Characterization of the Impact of Interconnect Design on the Capacitive Load driven by a Global Clock Distribution

ABSTRACT

Power reduction techniques are a critical issue in the design of today's ULSI chips. This paper is concerned with methods to characterize the capacitive load on the CHIP1 (Note to referees: the fake name CHIP1 is used in this paper to ensure blind review) on-chip global clock distribution [1], which is a large contributor to the overall chip power dissipation. A characterization of the capacitive load is needed because the contributions of the on-chip devices and interconnections are typically overestimated and are not well understood for high performance microprocessors. One problem that results from the lack of this information is excessively high power dissipation in the chip global clock distribution; the global clock distribution is over-designed and stronger than necessary to drive the actual (lower) chip load. Information about the capacitive load is difficult to obtain because the data volume is large, and extracting the interconnect data is a complex task. Sophisticated computer software is needed to extract the circuit and physical design data for hundreds of devices and wire segments within the chip design schedule.

This paper presents the first comprehensive characterization of the clock load for ASIC-like control logic designs in the 1.3GHz CHIP1 microprocessor core [1], [2]. This characterization was achieved with the use of sophisticated software written for this study to accomplish the task of extracting the data from these designs. Analysis of the data shows that the wire contribution to the chip capacitive load is significant and can increase the capacitive load of a design by 30% on average and by as much as 130% for some designs. The results also suggest that the wire load contribution on each metal layer can be reduced if an alternate interconnect design style is selected. Two alternate design styles are presented and show that a capacitive load reduction of 8.4% to 20% is expected for each design. Extended to the entire chip, the results show that the load reduction for the core is expected to be as high as 10%. These values are large enough that one alternate design style has been implemented in the design methodology of future chips.

Categories and Subject Descriptors

B.7.2 [Hardware, Integrated Circuits, Design Aids]: Layout, Placement, and Routing

General Terms Design, Optimization

Keywords

Power dissipation, clock, routing, Application Specific Integrated Circuit (ASIC), load capacitance

1. INTRODUCTION

Power reduction techniques are a critical issue in the design of today's ULSI chips. A major contributor to the power dissipation in high-performance chips is the power dissipation in the global clock distribution. This power dissipation results from the power dissipated in the distribution itself and from the capacitive load that is driven by the global clock distribution. In high-performance microprocessor designs, the size of the capacitive loading on the global clock is not well known and is typically overestimated. The use of this load overestimate requires a stronger global clock distribution [1], [2] and thus leads to excess power dissipation in the global clock circuitry and wires. Moreover, the use of the load overestimate also unnecessarily increases the difficulty of the tuning task in which the global clock distribution is carefully tuned to limit clock skew and clock delay [1], [2], [3].

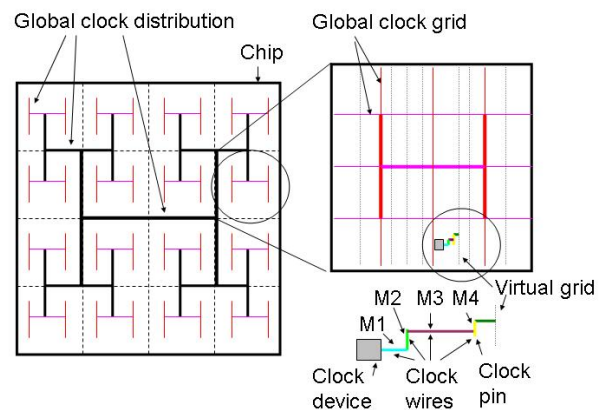


Figure 1. Schematic of the CHIP1 global clock distribution (left-hand side), global clock grid (upper right-hand side), and local clock wiring (lower right-hand side). In the example of local clock wiring shown, the global clock distribution drives a clock device such as a local clock buffer driver (shown in gray) from an design input pin on M4 wiring layer with an interconnection that contains wire segments on M3, M2, and M1 metal wiring layers.

The CHIP1 global clock distribution is designed with a strategy similar to that implemented in previous chips, as described in [4]. The purpose of the global clock distribution is to drive the local clock devices and wires that synchronize the functional circuitry. The global clock distribution consists of global clock interconnections and global clock buffers [1], [2]. Figure 1 shows a schematic of the global clock distribution in the CHIP1 chip. The figure shows the CHIP1 chip on the left-hand side and a schematic of the clock grid on the upper right-hand side. The

lower right-hand side of this figure shows an example of local clock wiring and local clock circuitry in a design. The local clock wiring can be modeled as a capacitance [6]. In each design, the four lowest metal layers are available to wire the local clock interconnections. These layers are: M1, M2, M3, and M4. An internal routing tool wires the global clock signal to M1 input pins on the clock devices from each M4 input pin. The length of each M4 pin is 27.5X the minimum wire pitch and is chosen at project commencement to enable a single M4 pin driving each clock device, for skew constraints [1], [2]. Figure 1 shows an example of a local clock wire that connects the M4 input pin to three wire segments on the metal layers M1, M2, M3. Each clock device is driven by a separate M4 pin.

Local clock circuitry consists of local clock buffers and local clock buffer drivers. In this paper, the term *total device load* refers to the sum of the input loads of the local clock buffers and the local clock buffer drivers. The *total wire load* is the sum of the load contributed by the local clock interconnections; these wires connect the devices to the design input pin that is driven by the global clock signal. The term *total clock load* refers to the sum of the load generated by the local clock circuitry and the load generated by the wires that connect this circuitry to the global clock distribution, and is given by the sum of the *total device load* and the *total wire load*.

Previous work [5] introduced two metrics to describe the interconnect contribution to the capacitive load on the global clock distribution. The first metric is the *Clock Wire Layer Contribution C*, which is the ratio of the metal layer capacitance to the *total clock load*, and can be expressed by the following equation,

$$C = C_{mj} / C_{tot}, \quad [1]$$

where the term C_{mj} represents the capacitance on metal layer m_j , and C_{tot} represents the *total clock load*. The second metric is the *Clock Uplift Factor U*, which is the ratio of the *total clock load* to the *total device load*, and can be expressed by the following equation,

$$U = C_{tot} / C_{dev} = C_{tot} / (C_{LCB} + C_{LCBDRV}), \quad [2]$$

where the term C_{dev} represents the total device load and is the sum of the terms C_{LCB} which represents the capacitive load in the local clock buffers (LCBs) and C_{LCBDRV} which represents the capacitive load in the local clock buffer drivers (LCBDRVs). Note that both these terms contribute to the *total device load* C_{dev} since the LCBs and LCBDRVs in the designs are driven by the global clock distribution. An example of a clock device driven by the global clock distribution is shown as the grey box in the lower right hand side of Fig. 1. Smaller values of the *Clock Uplift Factor* are desirable where possible as long as the total clock wire load and power are not increased, and skew is not increased significantly.

This paper presents a comprehensive characterization of the capacitive load of the CHIP1 core. The load is characterized at three levels of the physical design hierarchy: design (lowest level), unit (next upper level), and core (third level). For the case of the CHIP1, the core contains six functional units, and each unit contains multiple designs; for example, the IFU contains 18 individual ASIC-like control logic designs. The three hierarchy levels are shown in Fig. 2. These levels are the individual ASIC-like designs (at the lowest level of hierarchy; 18 examples of these are outlined in white in Fig. 2), six functional units (the next level of hierarchy; the six units are outlined in black in Fig. 2), and the core itself (the upper-most level of hierarchy; this is the image in Fig. 2). In total, the ASIC-like control logic designs occupy approximately 50% of the core area.

The first goal of this paper is to provide a characterization of the capacitive load of the ASIC-like control logic chip designs that are driven by the global clock distribution on the CHIP1 chip. The characterization consists of a quantification of the total wire load, the total device load, *Clock Uplift Factor*, and percentage of the wire load on each of the four available metal wiring layers. A second goal of this paper is to evaluate the impact of interconnect on the total capacitive load that is driven by the global clock distribution. A third goal of this paper is to provide alternate

interconnect design styles that reduce the interconnection contribution to the capacitive load.

2. MOTIVATION

A characterization of the interconnections in ULSI circuit designs is needed because the contributions of the on-chip interconnections to the size of the capacitive loading on the global clock are not well defined during the majority of the duration of the chip design project. Specifically, well-understood values of the load in these designs are not typically available until as late as near the time of tape-out. As a result, the global clock distribution is over-designed and, in fact, at tape-out is strong enough to drive a load that is significantly larger than the actual load that exists at tape-out. This over-design of the global clock distribution leads to unnecessary excess power dissipation. An understanding of the components of capacitive load will also enable alternative physical design methodologies that reduce the overall capacitive load on the global clock distribution and thus reduce the total power dissipation of the clock network.

3. EXPERIMENTAL DESIGN

One of the main obstacles to the existence of comprehensive data analysis studies of chip physical design data is the lack of sophisticated computer-aided analysis software, such as the software developed in this study. For this study, sophisticated computer software programs were written to extract the circuit and interconnect information from the chip designs. In this study, the CHIP1 designs are contained in a Cadence database format; thus the data analysis software was written in the Cadence SKILL programming language.




Figure 2. Image of the CHIP1 core. The locations of the six CHIP1 functional units (FPU, ISU, FXU, IDU, IFU, LSU) are shown with black outlines. The locations of 18 ASIC-like control logic designs in the IFU are shown with white outlines.

3.1. DATA ANALYSIS SOFTWARE

The development of this software eliminated the need for a time-intensive and costly manual analysis of the numerous individual designs. The development of this software also provided the results in time for application to the next chip project; without the use of this software, the results described in this paper would not have been available. Note that while the quantities (such as device count and wire segments) that are being measured in this paper appear to be straightforward, the data

analysis problem is complex. The complexity of the problem is in reality exacerbated by the data volume, including several hundreds of devices and wire segments. Specifically, the use of the data analysis software permitted the measurements of the number of clock devices and the interconnect wire load to be obtained in a few months and in time to provide results for decision-making in future microprocessors.

3.2. ALGORITHM

The data analysis software utilizes a hierarchical approach to the data acquisition by using the chip design hierarchy and design nomenclature to organize the results. Figure 3 shows a schematic of the algorithm for data extraction. The dark box annotated 'Unit Analyzer Report Generator' in the upper left represents the main engine of the data analysis software. The connections in the figure show that the algorithm actually extracts several types of data from each design. If design data exists (where the term MACRO represents a design), the algorithm extracts the area, clock data, and net data including wire-length information, as indicated by the shaded box in the middle of the middle row in the figure. This data is collected for all designs in each unit, as shown on the left in the middle row. After all data is collected, ASCII text files are generated according to the information desired; these files are referred to as *Reports*. As shown in the figure, separate reports are generated for each unit, the control logic macros (RLMs, or random logic macros), SRAM, or custom macros. The wire-length data generated by these reports is used to obtain the capacitance information described in this study.

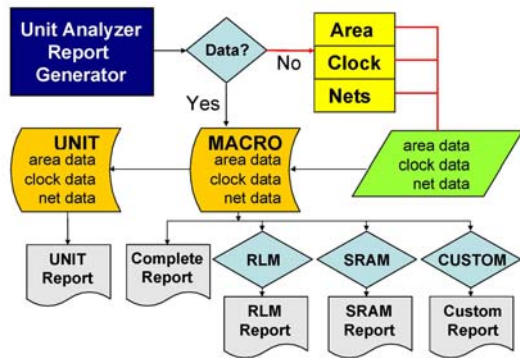


Figure 3. Schematic of the algorithm for data extraction.

This paper describes the use of this algorithm to extract data from the ASIC-like control logic designs in the six CHIP1 functional units. These six units are the Floating Point Unit (FPU), Fixed Point Unit (FXU), Instruction Decode Unit (IDU), Instruction Fetch Unit (IFU), Instruction Store Unit (ISU), and Load Store Unit (LSU).

3.3. CAPACITANCES OBTAINED FROM EXTRACTED DATA

A measurement of the *total device load* in each design is obtained by counting the number of devices driven by the global clock distribution. This number is then multiplied by the capacitive

load of each device. The value of the capacitive load per device is obtained from hardware measurements provided separately [7].

A measurement of the *total wire load* in each design is obtained by locating each of the wire segments of the local clock wiring for the global clock signal and by identifying the metal wiring layer for each segment. The load on each metal segment is obtained by multiplying the wire-length of each segment by the capacitive load per wire-length for that metal layer; the value of the capacitive load per layer is provided separately from hardware measurements. The *total wire load* is obtained by summing the load contributions for each of the metal wiring layers. The values of the *total clock load* and the *Clock Uplift Factor* are then obtained from the values of the *total device load* and *total wire load*. Values for the *Clock Wire Layer Contribution* can also be obtained from these quantities.

4. EXPERIMENTAL RESULTS

This section describes a comprehensive characterization of the clock load of the ASIC-like control logic designs in the CHIP1 core. Specifically, this section presents the *total wire load*, *total device load*, *Clock Uplift Factor*, and the portion of the wire load on each of the four available metal wiring layers for individual ASIC-like control logic designs, the six functional units, and the core itself. Values for these characteristics for the individual designs, units, and core itself are shown in Tables I-III.

4.1 Characterization of Individual Designs

Table I shows values of the *total wire load*, *total device load*, *Clock Uplift Factor*, and the portion of the wire load on each of the four available metal wiring layers for 18 IFU designs. This results in this table show that the contribution of the wire load to the total load averages 31% and ranges from 22% to 43%. The average *Clock Uplift Factor* is 1.5 and ranges from 1.3 to 1.7. The results also show that the main contribution to the wire load is provided by the metal segments on the M3 and M4 wiring layers; on average, 62% of the wire load is contributed by M3, for which the load contribution ranges from 43% to 81%. Also, on average, 36% of the wire load is contributed by M4, for which the load contribution ranges from 19% to 47%. Similar results are observed for individual designs in the other units.

The distribution of values of the *Clock Uplift Factor* for the individual designs in the CHIP1 core is shown in Fig. 4. This figure shows that the average value of the *Clock Uplift Factor* is 1.5, and can take on values from 1.1 to as large as 2.3.

Figure 5 shows the distribution in wire load for all the designs on each of the four metal layers. This figure shows that the wire load is contributed primarily from interconnect segments on M3 and M4. The results in this figure show that on average, 52% of the wire load is contributed by M3 and that 40% of the wire load is contributed by M4. These results show that the load contribution on M3 ranges from 19% to 81%, and the load contribution on M4 ranges from 16% to 79%.

4.2 Characterization of Functional Units

Table II shows values for the *total wire load*, *total device load*, *Clock Uplift Factor*, and the portion of the wire load on each of the four available metal wiring layers for the six CHIP1 functional units. The results in the table show that the *Clock Uplift Factor* takes on values between 1.3 and 1.5. Three of the units – the FPU, ISU, and LSU - have the largest value of *Clock Uplift Factor* of 1.5. The results in the table also show that the contribution of the wire load to the total unit load is 30% on average and that this contribution is provided mainly from the wire capacitance of the M3 and M4 metal wiring layers.

The relative contributions of the load in each functional unit to the total core load are also obtained. These results show that of the total capacitance load in the CHIP1 core, the LSU contributes 52%; the IFU contributes 20.7%, the ISU contributes 14.1%, the IDU contributes 9.5%, the FXU contributes 3.4%, and the FPU contributes 0.3%.

4.3 Characterization of the CHIP1 Core

Table III shows values of the *total wire load*, *total device load*, *Clock Uplift Factor*, and the portion of the wire load on each of the four available metal wiring layers for the designs in the CHIP1 core. The results in this table show that overall the local clock interconnections contribute nearly one-third (32%) of the total clock load and that the total device load is 68% of the total clock load. The results also show that the *Clock Uplift Factor* for the core is 1.5, from which it follows that the contribution of the wires to the total load is 50%.

Table III also shows the contribution of the wire load in the core from each of the four metal wiring layers: the table shows that 5% of the wire load is contributed by M1, 2% of the wire load is contributed by M2, 56.4% of the wire load is contributed by M3, and 36.3% of the wire load is contributed by M4.

5. DISCUSSION

The previous section provides the first characterization of the interconnection contributions to the capacitive load on a global clock distribution in a high-performance microprocessor. This characterization provides values for the increase in load provided by the local clock interconnections. Prior to this study, the information presented in the previous section was not well understood. For example, it was not known that the typical *Clock Uplift Factor* provided by the wires was on the order of 50% for the entire CHIP1 core. It was also not known that the *Clock Uplift Factor* for an individual design could be as high as 130%. Knowledge of this information enables future global clock distributions to be designed to drive a load capacitance for which an accurate estimate exists. It is therefore expected that the power dissipation in the global clock distribution will be reduced.

The results presented in the previous section show that the wire load in ASIC-like designs is a significant contribution to the total clock load. Specifically, the results show that the contribution of the wires can increase the capacitive load of ASIC-like control logic designs by as much as 30%.

The previous section also shows that the main contribution to the wire load in ASIC-like designs arises from load contributed by wires on two of the four available metal wiring layers: M3 and M4 metal layers. The load contribution from the M3 layer results from the use of M3 to connect the clock devices with a horizontal wire segment to the M4 input pins. The load contribution from the M4 layer results from the 27.5X M4 pin wire-length chosen during the project for all M4 pins.

5.1 Proposed Alternate Physical Design Methodologies

The results provided in the previous sections suggest two alternate physical design styles that are expected to reduce the total wire capacitance in the local clock wiring. The first alternate style is a reduction in the M4 pin length from 27.5X to a minimum length of 0.5X the minimum wire pitch. Implementation of this methodology in the CHIP1 core is expected to decrease the capacitive load by up to 8.9% in the LSU, up to 9.7% in the IFU, up to 8.9% in the ISU, up to 8.4% in the IDU, up to 9.7% in the FXU, and up to 9.7% in the FPU. These improvements and the expected reduction in the clock power are large enough that this style has been implemented in future microprocessors including

the POWER5 chip. In these microprocessors, additional M4 wire segments stitch together the 0.5X M4 pins at the next upper level (unit) of the physical design hierarchy. These additional wire segments increase the load contribution at the unit level.

Table IV compares the maximum load reduction and actual load reduction that are expected for five IFU designs. The results presented in this table show that for these designs, the value of the maximum reduction ranges from 8.1% to 13.2% and that the value of the expected reduction ranges from 3.4% to 6.2%. The range in these values exists since each design requires different lengths of M4 segments to connect the M4 pins. Extended to the chip, the reduction is expected to be as high as 10% with the implementation of this methodology.

The results described in the previous section also suggest a second alternate physical design style. In this style, each M4 pin is relocated to the same location as the clock device to which it is connected. Alternately, each clock device can be relocated to the position of the corresponding M4 pin. In these cases, the M3 contribution is expected to be eliminated, and the load is expected to be reduced by 12% to 20%; these results can be obtained by subtracting the M3 load contribution from the total load in Tables I-III.

6. CONCLUSION

This paper presents a comprehensive characterization of the capacitive load on the global clock distribution in the ASIC-like designs in the CHIP1 chip. For this study, which involved a large data volume of chip design information, sophisticated computer programs were written to extract the circuit and wire data from these designs. The paper shows that the wiring contributes nearly one-third of the total load driven by the global clock distribution; the extent of this contribution was unknown prior to this study. For these designs, the main contributions to the wiring load arise from the M3 and M4 metal wiring layers. The results presented in this paper suggest two alternate physical design styles to reduce load on the global clock distribution. The first proposed style is to reduce the M4 contribution to the wire load by reducing the M4 input pin length to 0.5X the minimum wire pitch, from the initial 27.5X length. The second proposed style is to eliminate the M3 contribution to the wire load by either relocating the clock devices to positions near the M4 pins or by relocating the M4 pins to positions near the clock devices. The results presented in this paper show that the values of the capacitive load reductions provided by these styles are expected to be 8.4% to 20%. The paper shows that these improvements and the expected reduction in the clock power are large enough that the changes suggested in this study have been incorporated in future chips, including the POWER5 microprocessor.

7. ACKNOWLEDGMENT

We thank Jeff Davis, of the Georgia Institute of Technology, for discussions. We also thank Izzy Bendrihem and Kelvin Lewis at our lab for their superior IT support for this project.

8. REFERENCES

- [1] P. J. Restle *et al.*, "The clock distribution of the CHIP1 microprocessor," Proc. ISSCC, Feb. 2002, pp. 144-145.
- [2] J. D. Warnock *et al.*, "The circuit and physical design of the CHIP1 microprocessor," *IBM J. Res. Dev.*, vol. 46, pp. 27-51, Jan. 2002.
- [3] V. Mehrotra, D. Boning, "Technology scaling impact of variation on clock skew and interconnect delay," Proc. IITC, 2001, pp. 122-124.
- [4] P. J. Restle, T. G. McNamara, D. A. Webber, P. J. Camporese, K. F. Eng, K. A. Jenkins, D. H. Allen, M. J. Rohn, M.

P. Quaranta, D. W. Boerstler, C. J. Alpert, C. A. Carter, R. N. Bailey, J. G. Petrovick, B. L. Krauter, and B. D. McCredie, "A clock distribution network for microprocessors," *IEEE Jnl. Solid-State Circuits*, Vol. 36, May 2001, pp. 792-799.

[5]. Previous work by the authors, Mar. 2004.

[6] H. B. Bakoglu, *Circuits, Interconnections, and Packaging for VLSI*. New York: Addison-Wesley, 1990.
 [7] G. Plumb, private communication, 2002.

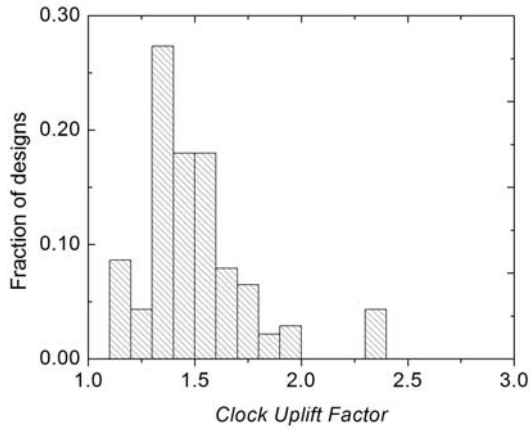


Figure 4. Distribution of Clock Uplift Factor for CHIP1 ASIC-like control logic designs.

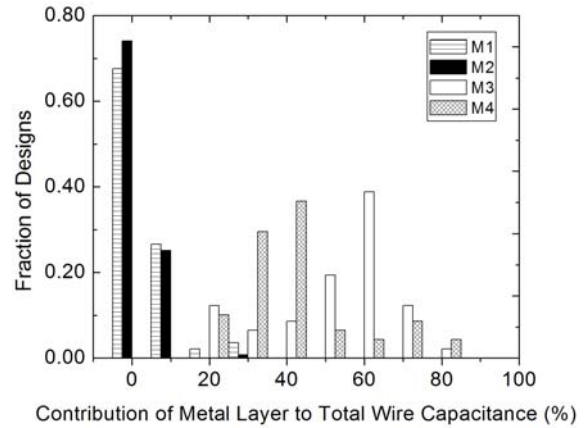


Figure 5. Capacitance distribution for the four metal wiring layers (M1, M2, M3, M4) in ASIC-like control logic designs in the CHIP1 core.

Table I: Characterization of the Load Capacitance of the 18 ASIC-like designs in the CHIP1 IFU.

Design	Number of Copies in IFU	Wire Load (%)	Device Load (%)	Clock Uplift Factor	M1 Wire Load (%)	M2 Wire Load (%)	M3 Wire Load (%)	M4 Wire Load (%)
<i>i1</i>	3	34	66	1.5	0	0	64.3	35.7
<i>i2</i>	1	24	76	1.3	0	0	81.1	18.9
<i>i3</i>	1	28	72	1.4	0	0.5	66.8	32.6
<i>i4</i>	1	23	77	1.3	0	0.5	52.0	47.5
<i>i5</i>	1	29	71	1.4	0	1.5	75.9	22.5
<i>i6</i>	1	31	69	1.4	6.6	3.9	52.8	36.7
<i>i7</i>	1	42	58	1.7	0	0.7	75.8	23.5
<i>i8</i>	1	26	74	1.3	0	1.8	64.1	34.1
<i>i9</i>	1	32	68	1.5	0	1.1	57.2	41.7
<i>i10</i>	1	27	73	1.4	0	1.5	66.5	32.1
<i>i11</i>	1	33	67	1.5	0	1.2	69.3	29.6
<i>i12</i>	1	36	64	1.6	16.0	2.6	42.7	38.6
<i>i13</i>	1	32	68	1.5	0	0.8	71.1	28.1
<i>i14</i>	1	33	67	1.5	0	1.3	52.9	45.9
<i>i15</i>	1	35	65	1.5	7.0	4.0	57.3	31.7
<i>i16</i>	8	30	70	1.4	0	1.1	58.0	40.9
<i>i17</i>	1	35	65	1.5	0	1.1	66.6	32.3
<i>i18</i>	1	36	64	1.6	0	1.2	58.4	40.4

Table II: Characterization of the Load Capacitance of the ASIC-like control logic designs in the six functional units in the CHIP1 core.

Unit	Wire Load (%)	Device Load (%)	Clock Uplift Factor	M1 Wire Load (%)	M2 Wire Load (%)	M3 Wire Load (%)	M4 Wire Load (%)
<i>FPU</i>	35	65	1.5	1.0	14.7	49.9	34.3
<i>FXU</i>	24	76	1.3	0	1.0	51.6	47.4
<i>IDU</i>	29	71	1.4	3.1	4.7	54.9	37.3
<i>ISU</i>	26	74	1.4	1.4	2.6	55.0	41.0
<i>IFU</i>	32	68	1.5	1.9	1.5	60.9	35.7
<i>LSU</i>	34	66	1.5	7.5	2.2	55.5	34.8

Table III: Characterization of the Load Capacitance of the ASIC-like control logic designs in the CHIP1 core.

Unit	Wire Load (%)	Device Load (%)	Clock Uplift Factor	M1 Wire Load (%)	M2 Wire Load (%)	M3 Wire Load (%)	M4 Wire Load (%)
<i>CORE</i>	32	68	1.5	5.0	2.3	56.4	36.3

Table IV: Impact of a Proposed Physical Design Change on the Expected Reduction in Load for five IFU designs.

Design	Maximum Expected Reduction (%)	Actual Expected Reduction (%)
<i>i9</i>	11.8	3.4
<i>i12</i>	11.7	5.4
<i>i13</i>	8.1	3.5
<i>i16</i>	11.3	3.8
<i>i18</i>	13.2	6.2