

A 2-Slot Time-Division Multiplexing (TDM) Interconnect Network for Gigascale Integration (GSI)

Ajay Joshi

Georgia Institute of Technology
School of Electrical and Computer Engineering
Atlanta, GA 30332-0250
Tel No. 1-404-894-9362
joshi@ece.gatech.edu

Jeff Davis

Georgia Institute of Technology
School of Electrical and Computer Engineering
Atlanta, GA 30332-0250
Tel No. 1-404-894-4770
jeff.davis@ece.gatech.edu

ABSTRACT

In many wire-limited VLSI digital systems the time delay of the longest global interconnect can be a significant percentage of the clock period. However, because semi-global and global wires with shorter wire lengths can have a much smaller time delay, they can remain idle (i.e. not switching) for a significant portion of the clock period even after their data has been transmitted. To capitalize on this phenomenon, this paper proposes that a low-overhead 2-slot Time Division Multiplexing (TDM) network should be incorporated into the multilevel interconnect architectures of future GSI systems to help reduce the escalating manufacturing costs that are primarily due to the projected increase in the number of metal layers. Using system-level interconnect prediction (SLIP) techniques, it is shown that a 20% reduction in the number of metal levels could be achieved for a 100M transistor VLSI system with only a 5-10% increase in power dissipation.

Categories & Subject Descriptors

C.5.4 [Computer System Implementation]: VLSI Systems

General Terms

Design, Theory

Keywords

time-division multiplexing, interconnect area, wire sharing

1. INTRODUCTION

With the continued advances in semiconductor manufacturing and technology scaling, the number of transistors on a single chip is expected to reach one billion before the end of this decade [1]. The resulting digital system on this chip would require a large number of semi-global and global interconnects that could put increasingly restrictive limits on the processor performance [2, 3]. In addition, each added metal layer introduces a non-trivial increase in the manufacturing cost. It is, therefore, imperative to investigate VLSI interconnect design and implementation methodologies that most efficiently utilize wiring tracks in a multilevel wiring network. In this paper, the system-level impact of incorporating a basic 2-slot

TDM VLSI interconnect network is investigated using traditional system-level interconnect prediction techniques.

The premise of this concept is based upon the assumption that a number of local, semi-global and global interconnects in a digital system remain idle (i.e. non-switching) for a significant portion of time either within a clock period after data transmission or across multiple clock periods. This intra-clock period idleness is due to the constraint that the pitch of all wires routed on two adjacent orthogonal metal levels (i.e. a tier) is constant. Shorter wires routed on this tier have a delay that is significantly less than the longest critical-path interconnect on that tier. For a given range of interconnect lengths, this idleness can be utilized by sending two binary transitions on a single wire during the clock period that is triggered by both the positive and negative edge of the clock. The inter-clock period idleness exists because of the low average activity factor, and can be utilized by sending two binary transitions on a single wire during two different clock periods.

This paper will investigate the system-level impact on wire area, performance, and power dissipation of a basic 2-slot TDM network that is incorporated into a VLSI multilevel interconnect architecture. To understand the impact on performance and overhead penalties, a wire sharing and pipelining technique similar to [4] is used to enable efficient wire sharing without a significant penalty in system performance or overhead circuitry. The result of this paper suggests that there is an opportunity to be gained by the use of simpler network schemes to achieve a significant reduction in signal wire routing area (~20%) with only a small increase in power dissipation (~5-10%).

2. TIME DIVISION MULTIPLEXING (TDM) CIRCUIT IMPLEMENTATION

The Network-on-Chip (NoC) concept has been proposed by many researchers, where a circuit-switched or packet-switched network is used to send data between different source-destination pairs. However, most of these complex implementations [5] [6] [7] that are proposed for NoC are focused on sharing a connected network between IP cores in a System-on-Chip (SoC) configuration. Because of the complexity involved in the implementation and operation, it would be difficult to apply these techniques to an arbitrary logic macrocell. The goal of this investigation is to develop a simpler network option that can be employed for both regular routing and irregular routing for intra and inter logic macrocell interconnects.

We propose the use of a variety of 2-slot TDM network channels distributed throughout the design that can be easily applied to both inter and intra logic macrocell communication. This TDM routing

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.

technique requires that interconnects send their data over a shared wire resource that is determined by the placement and routing constraints for a given pair of sources and sinks. In both regular and irregular routing, there could be a variety of routing configurations that could use a shared wiring resource to implement a basic 2-slot TDM circuit. For example, the two wire nets that are routed close to each other in Figure 1 can be replaced by a single net in Figure 2 with a demux/mux circuit and additional latches. This pattern illustrates the most straightforward opportunity for wire sharing; however, this is one of many patterns that demux/mux insertion could be used to reduce wire channels. In fact, a combination of longer and shorter wires could be used to reduce wire demand as well. For instance, consider a long interconnect in parallel with two smaller wires as seen in Figure 3. These two wire tracks could be replaced by a single shared interconnect with appropriate demux/mux insertion as seen in Figure 4. In fact, intelligent insertion of demux/mux pairs and repeaters could be incorporated into the physical design constraints of a GSI design to help reduce signal routing needs or excessive routing congestion.




Figure 1. Two uni-directional Interconnects (signals in same direction)




Figure 2. A single time-shared uni-directional interconnect

The impact on the system-level performance can be understood in part by considering two types of implementation of a 2-slot TDM network. The first implementation option is to have each time slot equal to the clock period of a given digital design. For high-activity interconnects, this could have a negative impact on the overall system performance because it would take twice the time for data to reach its destination; however, for low activity interconnect wires this could be an acceptable option that would not affect the overall throughput. However, this implementation needs to be incorporated at the RTL level of design to make sure that the timing of this network will work in harmony with the existing architecture. Moreover, this technique could require additional overhead of storage buffers to manage data traffic during spurious periods of high activity. Subsequently, for the system-level models discussed in this paper, it is assumed that this technique is used sparingly in the resulting design.




Figure 3. Three uni-directional interconnects (signals in same direction)




Figure 4. A single time-shared uni-directional interconnect

The second implementation option of a 2-slot TDM network that is considered in this paper has each time slot divided between the high and low levels of the global clock period of a given digital system. When comparing to a conventional design, the clock frequency is the same, but the data rate on the shared wire resource is doubled by triggering off the positive and negative edge of the system clock. This is a viable technique as long as there is adequate intra-clock idleness on a given line after it sends a single binary transition. Figure 5 shows the fraction of the clock period (e.g. 1.5 Ghz) that a wire needs to send a single binary transition. It is seen in Figure 5 that for a given range of interconnect lengths on the semi-global and global routing tiers that there could be a significant number of wires that have significant intra-clock period idleness. In fact, the wires that have a time delay that is less than 40% of the clock period can use demux/mux insertion to significantly reduce the wire area without any loss in system performance. The control for these wires is greatly simplified by connecting a global clock with a 50% duty cycle to the select lines on the demux/mux control circuitry.

However, it is also assumed in the system level modeling that an alternate technique of wire pipelining and wire sharing similar to [4] is used. The result of this alternate circuit technique is that a larger number of wires can use the 2-slot TDM routing method and still maintain high-performance. In [4] the minimum sustainable pulse width that can travel down a repeater interconnect circuit is calculated. By using the minimum pulse width for sending the first bit on a shared wire resource, the overall time delay that it would take to send two bits on this line is significantly reduced. In fact the latency for sending two bits is simply the minimum pulse width plus the intrinsic wire time delay. Figure 6 illustrates that this circuit approach significantly increased the range of wire lengths for which the 2-slot TDM routing technique could be used, by more than 50%. The control for this minimum pulse width circuit technique is greatly

simplified by using a global clock that has a duty cycle that is equal to the ratio of the minimum pulse width to the original clock period.

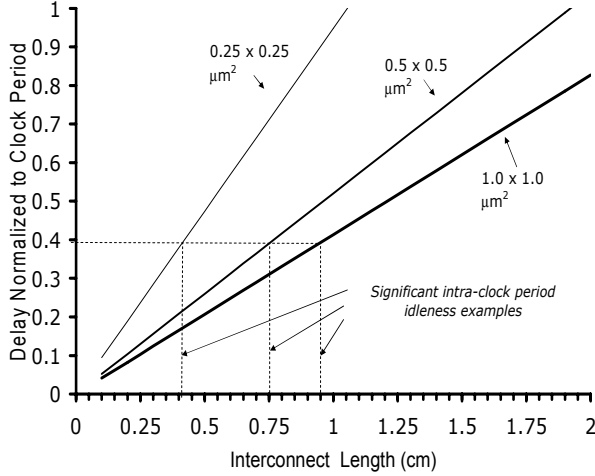


Figure 5. Intra-clock period idleness for wires on different wiring tiers

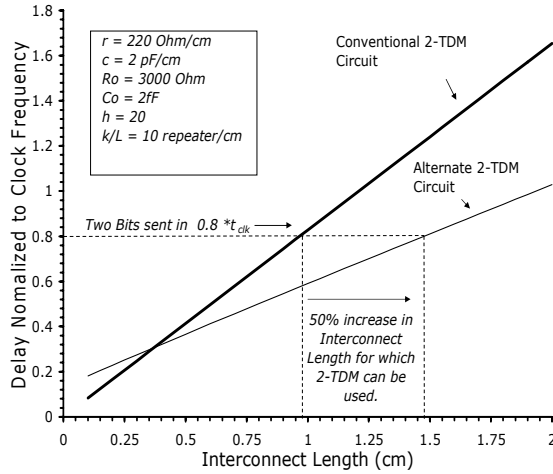


Figure 6. Interconnect wire pipelining technique to increase the number of wires

3. SYSTEM MODELING

3.1 Model Assumptions

In order to quantify the advantages of a 2-slot TDM network, a digital system is modeled as a homogeneous array of logic gates. Using a Rentian interconnect distribution after [10], the implementation of an n-tier multilevel interconnect architecture is modeled after techniques similar to [9]. A multilevel interconnect architecture containing repeaters and no shared wires is first constructed and serves as the base case that is compared to a multilevel interconnect architecture that incorporates a 2-slot TDM network.

The number of interconnects that can be shared is determined by three different factors – the ratio of the interconnect delay to the cycle time, activity factor of the interconnects, or the minimum pulse width of the interconnects with repeaters. The entire range of

interconnect length can be divided into three categories. The first category contains local interconnects which are idle for a significant part of the clock period. This intra-time period idleness can be used to send a second signal over the shared wire. These interconnects lie on the lowermost tier of the multilevel architecture. The second category consists of interconnects that are idle for smaller part of clock period (<20%), however they remain idle for a significant number of clock cycles due to a low average activity factor of all interconnects. The time-division multiplexing technique utilizes this inter-time period idleness for sending signals over the shared wire. While designing these interconnects, care should be taken that the sources that are chosen to send data over the shared wire do not contend for the shared wire in the same clock cycle. The third category encompasses the global interconnects that have repeaters inserted upon them. The enhanced throughput of these interconnects with repeater insertion is utilized to send two signals in a single clock period [4].

The reduction in the number of metal layers has been modeled for a large logic macrocell using system level interconnect prediction techniques. An n-tier multilevel interconnect architecture is constructed to incorporate this wire sharing technique for various pairs of interconnects. Given the physical layout constraints for interconnects and transistors, it is impossible to apply the 2-slot TDM routing to all interconnects; therefore, a sharing efficiency factor e_{share} is introduced to determine the number of interconnects for which the techniques in section 2 can be utilized. In addition, to explore the limitations on the usage of 2-slot TDM routing, the range of interconnect lengths on which this technique is applied is varied as a part of the case study in the next subsection. Various "cut-off" lengths will be assumed such that the TDM routing is applied to all interconnects larger than this cut-off length value.

3.2 Case Study Description

In this case study, a digital system that contains 100 million transistors (assuming the use of three input six transistor NAND gates as a standard gate, which corresponds to approximately 1.66×10^7 gates) is examined. The logic gates are assumed to be uniformly distributed across the system. A 100 nm technology generation is used and the system is operated at a clock frequency of 1.5 GHz. The die area is varied from 3 sq.cm to 6 sq.cm. The values for other parameters are presented in Table I.

Table I. Values of various parameters

Paramater	Value
$\alpha = \text{fanout} / (\text{fanout} + 1)$	0.75
Wiring efficiency e_w	0.4 [12][13]
$\chi = \text{factor for converting point to point wire length to wiring net length}$	0.667 [11]
$\beta = \text{wire delay expressed as a fraction of the cycle time}$	0.25 for small wires and 0.9 for large wires

In addition, the Rent's exponent is $p=0.66$ and the Rent's coefficient is $k=4.0$. Local interconnects with no repeaters require a small fraction of the available time period to send the data. For the system having the parameters mentioned above, these local interconnects use less than 25% of the clock period, thus creating the opportunity for utilizing the intra-clock period idleness. The semi-global and global interconnects can utilize up to 80% of the clock period. Hence, intra-clock period idleness cannot be used for

all of these wires; however, for those few that operate in close to 80% of the clock it is assumed that the inter-clock period 2-slot TDM technique will be used. However, as mentioned earlier, it is not possible to implement interconnect sharing on all interconnects greater than the cut-off length due to physical layout constraints. Hence, a sharing efficiency factor has been introduced and its value is chosen to be 0.6. Given the introduction of a wiring efficiency factor of 0.4, effectively only 0.24 of the available routing channels provide opportunities to have shared interconnects. As mentioned earlier, each shared line uses one or more pairs of multiplexer and demultiplexer and latches. The sizing of multiplexer-demultiplexer pairs and latches is done based on the sizing of the driver of the interconnect and any inserted repeaters on the interconnect.

The technique described in [9] starts inserting repeaters from the topmost tier and then goes on inserting repeaters for successive lower tiers based on the availability of the silicon area. The multiplexers and demultiplexers are inserted along with the repeaters while checking for availability of silicon area. In addition to the multiplexer-demultiplexer pair, latches are inserted at the terminals of the multiplexers and demultiplexers to ensure that the data is not lost before the beginning of the next clock cycle. In this technique though interconnects are eliminated, the overall dynamic power of the interconnects does not decrease due to the increased activity of the shared interconnects. The overhead circuitry consisting of multiplexer-demultiplexer pair and latches result in an increase in the dynamic power of the system. The transistors of this overhead circuitry have the same activity factor as that of the shared interconnects. As an additional note, a multiphase clock source is required for the system due to the need for varying duty cycles of the inputs to the multiplexer-demultiplexer pairs.

3.3 Case Study Results

Figure 7 shows the percent reduction in the number of metal layers obtained as a result of 2-Slot TDM interconnect network for different cut-off lengths. It can be observed from Figure 7 that the percent reduction in the number of metal layers decreases as the value of the cutoff increases and the percent reduction obtained saturates as the die-size is increased. Table III shows the number of metal layers for the conventional case and the various models.

Figure 8 shows the percent increase in the power dissipation with the increase in the wire-limited area for various cut-off lengths. As the cut-off length decreases there is less opportunity for wire sharing and hence a smaller number of interconnects that can be eliminated. However, this reduction in the number of shared interconnects does not contribute to the power calculation as the power dissipated by a single shared wire is the same as that dissipated by the two wires due to the increase in the activity factor of the shared wire. It can be seen from Figure 8 that the total power dissipated increases with a decrease cut-off length. This increase in the power comes from switching of the demux/mux pair and latches that are required to implement wire sharing. The dissipated power increases with larger die area because of the requirement of thicker interconnects to send signals over longer distances. For cut-off lengths below 400 gate pitches, the increase in power due to the overhead circuitry is extremely high and the resulting transistors limited area (logic area + repeater area + overhead area) becomes almost equal to the wire-limited area, thereby making it impossible to implement the technique for lower cut-off lengths. In case of cut-off lengths above

1100 gate pitches, there is no significant decrease in the number of metal layers.

Table II. Number of metal layers for various models

Cutoff Length (Gatepitches)	Wire area			
	3 sq.cm	4 sq.cm	5 sq.cm	6 sq.cm
Conventional case	8.34	8.31	8.31	8.31
400	6.45	6.4	6.4	6.4
500	6.47	6.41	6.41	6.41
600	6.49	6.43	6.43	6.43
700	6.54	6.45	6.45	6.45
800	6.59	6.49	6.49	6.49
900	6.64	6.53	6.53	6.53
1000	6.69	6.58	6.58	6.58
1100	6.73	6.61	6.61	6.61

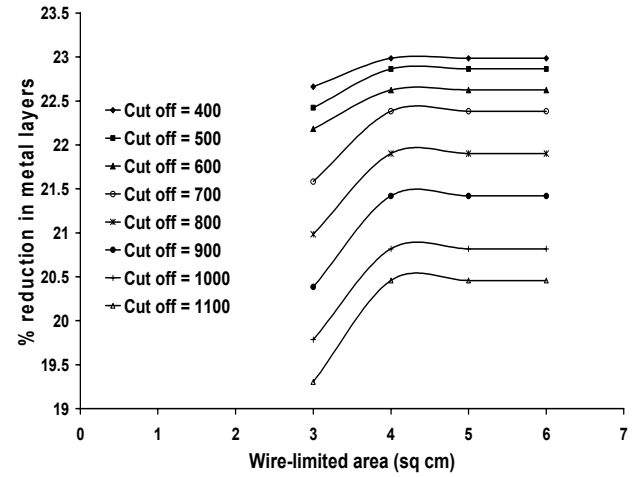


Figure 7. Percent reduction in the number of metal layers using uni-directional shared interconnect technique

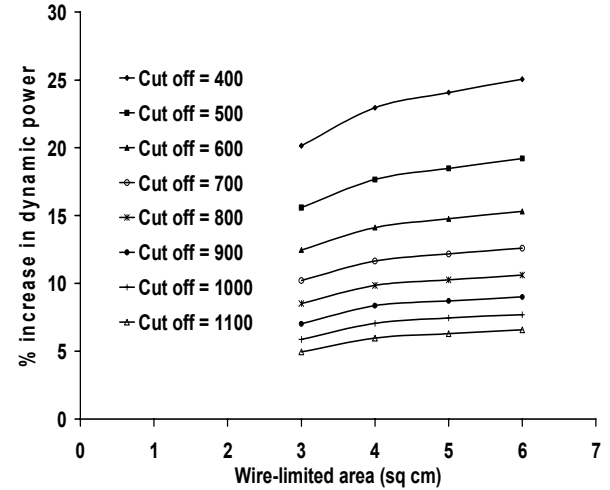


Figure 8. Percent increase in dynamic power for uni-directional shared interconnect technique

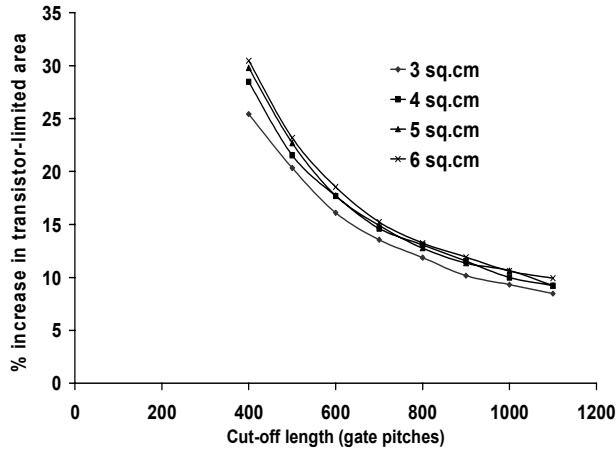


Figure 9. Percent increase in the transistor-limited area

Table III. Transistor limited area for various models

Cut off Length (Gatepitches)	Wire area			
	3 sq.cm	4 sq.cm	5 sq.cm	6 sq.cm
Conventional case	1.18	1.3	1.41	1.51
400	1.48	1.67	1.83	1.97
500	1.42	1.58	1.73	1.86
600	1.37	1.53	1.66	1.79
700	1.34	1.49	1.62	1.74
800	1.32	1.47	1.59	1.71
900	1.3	1.45	1.57	1.69
1000	1.29	1.43	1.56	1.67
1100	1.28	1.42	1.54	1.66

The percent increase in the transistor-limited area as a result of the application of the 2-slot TDM routing is shown in the Figure 9. The transistor-limited area includes the total silicon area occupied by the logic, repeaters and network overhead circuitry. It can be observed from Figure 9 that as the cutoff length increases the required overhead circuitry decreases and results in less overhead transistor area. Given the constraints of the physical layout, it is not possible to utilize the entire available silicon area for the transistors. This puts an upper limit on the cut-off length that can be used for different die sizes. Table III shows the transistor-limited areas for the various die size options

4. CONCLUSION

A 2-slot TDM routing technique is proposed in this paper which makes use of the existing idleness in many semi-global and global interconnects. The control circuitry for this network sends data on a single wire channel during both the positive and negative edge of the clock cycle. Using system-level interconnect prediction techniques, it is seen that incorporation of a 2-slot TDM network into the multilevel interconnect architecture of a 100M transistor case study reveals an average reduction of 20% in the number of metal layers. For current VLSI and future GSI systems this basic technique could contribute to significant savings in the interconnect manufacturing costs. It is shown in this paper that the reduction in number of metal

layers could be obtained with a minimal loss in performance and only a 5-10% increase in power dissipation.

5. ACKNOWLEDGEMENTS

This material is based upon work supported by the National Science Foundation under grant no. 0092450. We would also like to thank Raguraman Venkatesan and Vinita Deodhar for the various interesting discussions we had while investigating this problem.

6. REFERENCES

- [1] <http://public.itrs.net/>
- [2] J. D. Meindl, "Low-power microelectronics: Retrospect and prospect," *Proc. IEEE*, vol.83, pp. 619-635, Apr.1995.
- [3] J. A. Davis, R. Venkatesan, A. Kaloyeros, M. Bylansky, S. J. Souri, K. Banerjee, K. C. Saraswat, A. Rahman, A. Reif, and J. D. Meindl, "Interconnect limits on gigascale integration (GSI) in the 21st century," *Proc. IEEE*, vol.89, pp. 305-324, March 2001
- [4] V. V. Deodhar and J. A. Davis, "Voltage scaling and repeater insertion for high throughput low power interconnects," *Proc. ISCAS 2003*, vol.5, pp. V-349 – V-352, May 2003
- [5] S. Kumar, A. Jantsch, J. Soininen, M. Forsell, M. Millberg, J. Oberg, K. Tiensyrja and A. Hemani, "A Network on Chip Architecture and Design Methodology," *Proc. IEEE Comp Soc*, pp. 105-112, April 2002
- [6] J. Liu, L-R. Zeng, D. Pamunuwa and H. Tenhunen, "A Global Wire Planning Scheme for Network-on-Chip," *Proc. ISCAS 2003*, vol.4, pp IV-892 – IV-895, May 2003
- [7] P. Bhojwani and R. Mahapatra, "Interfacing cores with on-chip packet switched networks," *Proc. VLSI design*, pp. 382-387, Jan. 2003
- [8] P. Zarkesh-Ha, K. Doniger, W. Loh and P. Wright, "Prediction of Interconnect Pattern Density Distribution: Derivation, Validation and Applications," *Proc. SLIP*, pp. 85- 91, Apr. 2003
- [9] R. Venkatesan, J. A. Davis, K. A. Bowman and J. D. Meindl, "Optimal n-tier Multilevel Interconnect Architectures for Gigascale Integration (GSI)," *IEEE Trans. VLSI systems*, vol. 9, pp. 899-912, Dec. 2001
- [10] J. A. Davis, V. K. De and J. D. Meindl, "A stochastic wire-length distribution for gigascale integration (GSI)—Parts I and II," *IEEE Trans. Electron Dev.*, vol.45, pp. 580-597, Mar. 1998
- [11] G. A. Sai-Halasz, "Performance trends in high end processors," *Proc. IEEE*, vol.83, pp. 18-34, Jan. 1995
- [12] H. B. Bakoglu, *Circuits, Interconnections and Packaging for VLSI*. Reading, M.A: Addison-Wesley, 1990