

Repeater Insertion and Wire Sizing Optimization for Throughput-Centric VLSI Global Interconnects

Harshit Shah, Pun Shiu, Brian Bell, Mamie Aldredge, Namarata Sopory and Jeff Davis

AIMD Research Group

Georgia Institute of Technology, Atlanta, GA- 30332

{jeff, harshit}@ece.gatech.edu

ABSTRACT

As technology advances towards billion transistor systems, the cost of complex wire networks will require area efficient wiring methodologies. This paper explores the tradeoffs between wire latency, throughput and area for deep submicron (DSM) interconnect technologies. From basic physical models, optimal wiring sizing for repeater networks are rigorously derived and compared to HSPICE simulations. Key case studies from 250nm to 70nm technologies reveal that significant wire area reduction (20-50%) can be achieved with optimal wire sizing to maximize the throughput per unit wire area.

1. INTRODUCTION

As VLSI technology progresses toward integration densities that will allow for more than a billion active devices per chip, the cost of high-speed wire networks will become excessive [1,2]. The economic demands to continue the exponential reduction in price per function will force the use of area efficient wiring methodologies that will require a shift from a low-density latency-centric global wire design to a high-density throughput-centric wire design [3,4]. This is especially important for low-density global interconnect structures that are shielded with co-planar ground lines, which are currently being used in state-of-the-art processors [5,6].

This paper will explore the design tradeoffs of global interconnect architectures that are implemented in 250nm and 70nm technologies. Opportunities for achieving high throughput, low latency, and area efficient designs are revealed through the creation of new physical models for interconnect throughput and HSPICE simulation. Section 2 will explore interconnect design for a long global wires without repeaters, and propose a new wire sizing methodology that maximizes the throughput per unit wire area. In Section 3, a new closed-form expression for the maximum throughput of a high-speed interconnect with repeaters is derived, and the tradeoff with interconnect latency is explored. Finally both these methods are used to find the optimal wire size for a given number of repeaters that maximizes the *throughput per unit wire area*. Maximizing this metric will be essential for the design of low-cost and high-performance digital products in the era of gigascale integration (GSI).

2. SINGLE DRIVER INTERCONNECT

Single driver or cascaded driver circuit design is important in a GSI multilevel interconnect architecture due to constraints on repeater placement and numbers. To obtain physical insight into a throughput-centric design methodology, a simple model for the maximum number of bits per second (bps) that can be sent down a lossy communication channel with a storage latch at either end, can be approximated by the following expression:

$$T = \frac{1}{t_{50\%_wire} + t_{90\%_end_buffer} + t_{hold} + t_{setup} + \delta_{variations}} \quad (1)$$
$$\approx \frac{1}{t_{50\%_wire} + \Delta_{overhead}}$$

where $t_{50\%_wire}$ is the 50% rise time of the signal at the end of the wire, $t_{90\%_end_buffer}$ is the 90% rise time of an end buffer that drives the input to the latch, t_{hold} and t_{setup} are the hold and setup times of the latch, and $\delta_{variations}$ accounts for variations in clock skew, jitter, and manufacturing fluctuations in the wire drivers.

Due to the simultaneous constraints of high-performance and high wire density, a very pertinent metric for VLSI global wire systems is the throughput per unit wire area. Maximizing this metric translates into higher communication performance and lower system cost. Using the above approximation for simplicity and physical insight, the throughput-per-unit-area (T_A) for a single wiring channel is approximately given by:

$$T_A \approx \frac{1}{(t_{50\%} + \Delta_{overhead})\beta WL} \quad (2)$$

where W is the metal wire width, $\beta=(1+S/W)$, S is the metal spacing, and L is the interconnect length.

On-chip inductance is very significant for noise considerations in high-speed wires; however, if the time-of-flight condition is not violated, then it will be assumed that inductance slightly perturbs the maximum throughput. Making this assumption allows for the use of simple distributed rc models after [7,8]. An expression for the optimal interconnect width can be calculated by setting the derivative of (2) with respect to the wire width, W , equal to zero. The optimal wire size for maximum throughput per unit area is:

$$W_{opt} = \sqrt{\frac{\rho L}{0.7\alpha R_o} \frac{(0.4c_w L + 0.7C_o h)}{\left(\frac{c_w L}{h} + C_o + \frac{\Delta_{overhead}}{0.7R_o}\right)}} \quad (3)$$

where ρ is the resistivity of the metal; R_o and C_o are the driver output resistance and driver input capacitance of a minimum size driver, respectively; α is the metal thickness to width ratio; c_w is the capacitance per unit length, which is dependent only on the ratio of the wire dimensions; and h is a factor that indicates the size of the drivers being used (with $h=1$ being defined as a minimum size driver). Substituting (3) into (2) gives the optimal throughput per unit area, T_A , for a single driver.

$$T_A = \frac{\sqrt{\alpha}}{2\beta L \sqrt{0.7\rho L R_o (0.4c_w L + 0.7C_o h) \left(C_o + \frac{c_w L}{h} + \frac{\Delta}{0.7R_o}\right)}} \quad (4)$$

A plot of the throughput per unit area of a one centimeter global interconnect as a function of the interconnect width appears in Figure 1. Using interconnect circuit models after [7,8], this design case study assumes a 70 nm technology node corresponding to ITRS projections for 2006. The clock frequency for this design is 5.631Ghz with a clock period of 177ps.

Table 1 includes the data for a conventional “latency-centric” design that uses the largest global interconnect dimension that is available for this technology, which is assumed to be a $1\mu\text{m} \times 2\mu\text{m}$ line. The driver size is approximately 20x greater than the minimum size driver and is chosen to achieve a latency of approximately 177ps. This is marked on Figure 1 as “latency-centric design.” However this wire width is far from optimal throughput per unit area. In fact, the optimal width is approximately 525nm, which results in a throughput per unit area that is 1.24x greater than the conventional latency centric design point.

Moreover, two other throughput-centric design points are marked on Figure 1 and appear in Table 1. The “minimum clock frequency throughput-centric design” in Figure 1, for example, has the same wire area and throughput of the conventional latency design, but the synchronizing global clock frequency has been cut by more than 70%, which could significantly relieve the burden of distributing a high-speed clock across the entire chip without a loss in throughput performance.

Table 1 also illustrates the area needed to send 720Gbps, which corresponds to a 128 parallel lines times the ITRS projected clock frequency for 2006. For this particular examples, a throughput-centric design strategy indicates that the wire pitch for a one-centimeter global interconnect bus should be approximately *half* the dimension of a latency-centric design strategy. Reducing

the wire dimensions results in slower global wires, but the number of interconnect channels is increased until the aggregate throughput of 720Gbps is again achieved. As seen in Table 1, the area needed to send this data is reduced by almost 20% and the global clock frequency for this area efficient design has been reduced by 36%.

Additional latches are needed to achieve the throughput per unit area for a given technology; however, if ULSI systems are severely wire-limited then this could be an acceptable tradeoff. Additional HSPICE simulations of these circuits are being presented in Section 4.

	Optimal Latency	Max T_A	Optimal T_A	Min. Clock F.
T_A [Tbps/cm ²]	27.76	34.5	33.86	29.07
Wire width	1000nm	525nm	420nm	280nm
Latency	177ps	276ps	351ps	614ps
T/line [Gbps]	5.64	3.63	2.84	1.63
Clock F. [Ghz]	5.64	3.63	2.84	1.63
#lines (720Gbps)	128	199	253	442
Wire Area (720Gbps)	0.0256 cm ²	0.0208 cm ²	0.0212 cm ²	0.0248 cm ²

Table 1. Single Driver Wire Sizing Design Points

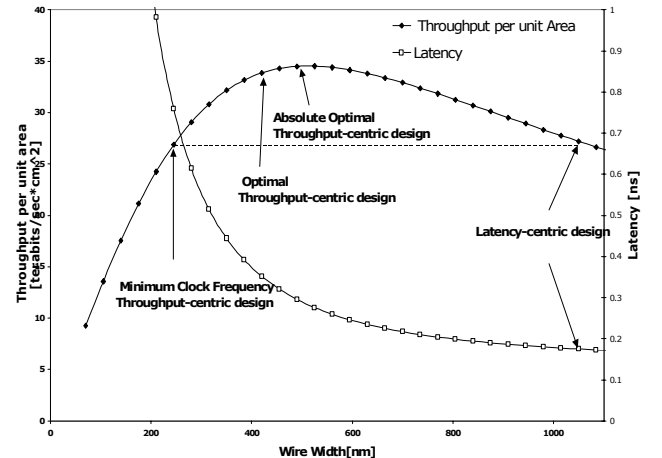


Fig. 1. Comparison of a throughput–centric design to a latency centric design for 70nm technology

3. INTERCONNECT REPEATERS

In the previous section, single driver circuits had a throughput that was approximately the reciprocal latency (i.e. time delay) because high resistance prevents wave propagation. However, with the insertion of repeaters on high-speed interconnects a type of wave propagation can be achieved in VLSI interconnects. This section illustrates the derivation of a physical model that captures this phenomenon for repeater circuits.

3.1 Mathematical Model for Throughput

Because repeaters provide shorter wire segments, the models developed in this section assumes that inductance will only slightly perturb the solution; therefore, to derive the equation for interconnect throughput, consider a single pole approximation for the transient voltage at the far end of the line

$$\frac{V(l, t)}{V_{dd}} = 1 - \sum_{k=1}^{\infty} K_k e^{-t/\sigma_k} \approx 1 - K_1 e^{-t/\sigma_1} \quad (5)$$

where expressions for σ_1 and K_1 are given in [8].

To derive a throughput model based upon (5), first consider an interconnect that is divided into n equal size segments as seen in Figure 2. Assuming there is a constant 50% delay on each repeater segment and a constant delay in each buffer that is denoted by Δ_{buf} , then the 90% rise time for the n^{th} segment is given by

$$t_{n90\%} = \tau_{90\%seg} + (n-1)\tau_{50\%seg} + (n-1)\Delta_{buf} \quad (6)$$

Using (5) and (6) gives the time it takes for the n^{th} segment to change to v_n is

$$t_n = \sigma_{RCseg} \ln \left(\frac{K_1}{1-v_n} \right) + (n-1)\sigma_{RCseg} \ln(2K_1) + (n-1)\Delta_{buf} \quad (7)$$

As seen in Figure 3 when the voltage of the final segment (w_n) reaches $0.9V_{dd}$, the voltage of the previous segment (w_{n-1}) will have $0.5V_{dd}$, which is the threshold voltage of a symmetric CMOS inverter. If w_n reaches $0.9V_{dd}$ at time $t=t_n$, the rising and falling transient voltage of w_{n-1} can be found by solving for t_{n-1} and v_{n-1} in the following equations

$$V_{dd}(1 - K_1 e^{\frac{-(t_{n-1} - (n-2)\tau_{50\%seg} - (n-2)\Delta_{buf})}{\sigma_{RCseg}}}) = V_{n-1} \quad (8)$$

$$V_{n-1} K_1 e^{\frac{-(t_n - t_{n-1})}{\sigma_{RCseg}}} = 0.5V_{dd} \quad (9)$$

where V_{n-1} is the peak voltage of w_{n-1} and t_{n-1} is the time at which this peak occurs. Solving for t_{n-1} in (8) and (9) as a function of t_n gives

$$t_{n-1} = t_n + \sigma_{RCseg} \ln \left(\frac{0.5}{K_1} + K_1 \exp \left(-\frac{t_n - (n-2)(\tau_{50\%seg} + \Delta_{buf})}{\sigma_{RCseg}} \right) \right) \quad (10)$$

where t_n is given by (7). Generalizing the index n to k gives a complete recursive expression for the waveform at the k^{th} segment.




Figure 2: Physical repeater model




Fig. 3. Example waveform for n^{th} , $(n-1)^{\text{th}}$ and $(n-2)^{\text{th}}$ segments

Letting $K_1=1$ and assuming that the time interval ($t_k - t_{k-1}$) is approximately constant, then the minimum pulse width for a binary signal is given (11). This is the minimum pulse width needed to achieve a 90% voltage at the end of the last repeater segment, and the reciprocal of this pulse width is define as the maximum channel throughput (i.e. $T_{max} = 1/t_1$)

$$t_1 = \tau_{90\%seg} + (n-1)\tau_{50\%seg} + (n-1)\Delta_{buf} - (n-1)\sigma_{RCseg} \ln \left(\frac{1}{0.5 + \exp \left(-\frac{(\tau_{90\%seg} + \tau_{50\%seg} + \Delta_{buf})}{\sigma_{RCseg}} \right)} \right) \quad (11)$$

3.2 HSPICE verification

In this section, the physical models derived in the previous sections have been validated using HSPICE simulations. For 250nm technology, it is assumed that the distributed capacitance of interconnect is equal to $2.7612[pF/cm]$ and the distributed resistance of interconnect is equal to be $2672 [\Omega/cm]$. Level 49 HSPICE models are used for the repeater drivers that are 56x larger than a minimum size repeater. HSPICE simulations closely follow the results calculated using the physical models as seen in Figure 4 and Figure 5.

The throughput eventually saturates because of the limit on transistor switching speed for 250nm technology. This view is validated by a recent experiment in 800nm technology [4] in which the maximum bandwidth of a repeater circuit was measured.

3.3 Throughput and Latency Tradeoff

Latency-centric repeater insertion attempts to achieve low latency with the smallest possible number of repeaters.

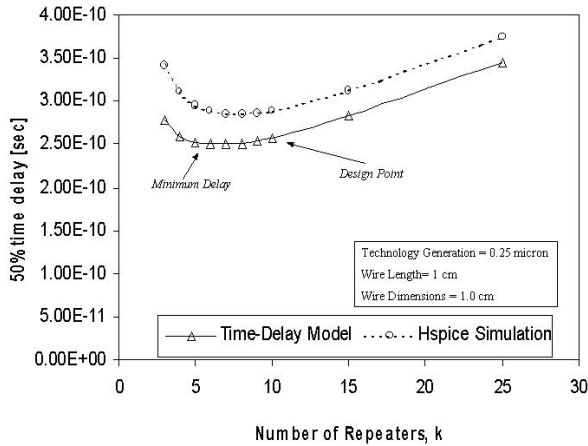


Fig. 4. Optimal interconnect delay for an ideal repeater circuit

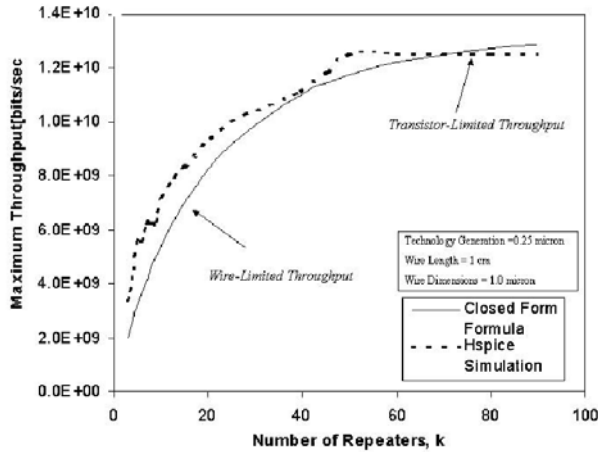


Fig. 5. Maximum throughput for an ideal repeater circuit

For example, Figure 4 illustrates the existence of an absolute minimum point for interconnect delay. However, it is well known that the region around this optimal latency is relatively flat. In a throughput-centric design, however, one could choose the design point in Figure 4 with more repeaters that would significantly increase the communication throughput (as seen in Figure 5) while maintaining low latency (as seen in Figure 4).

As seen in Figure 6, according to HSPICE simulations, a throughput oriented design gives 2.5x increase in throughput with a small increase in the interconnect latency. The number of repeaters in this case is greater, but this increase in throughput per interconnect could partially offset this penalty because of the significant decrease in the number of routing channels. For example, in this case the number of repeaters is increased by a factor of 4x, but because the throughput is increased by a factor of 2.5x, the number of repeaters would increase by only a factor of 1.6x.

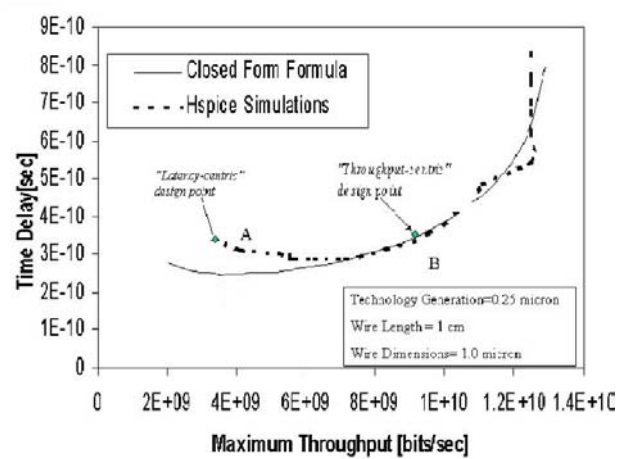


Fig. 6. Tradeoff of interconnect latency and throughput for a one centimeter global wire

3.4 Throughput per unit wire area

These physical models for the throughput in a repeater circuit can be used to understand the implications of repeater design on optimal wiring sizing to achieve the maximum possible throughput per unit wire area. Figure 7 reveals that as the number of repeaters inserted on an interconnect increases, the throughput/area also increases for given interconnect width. This is because, as shown in Figure 5, the throughput increases with greater repeater insertion. However, as repeater design rules become more aggressive (i.e. smaller repeater segment length) the optimal wire width for throughput/area metric will also become smaller.

4. CASE STUDY (L=1 cm, F=250nm)

This case study compares a latency-centric and throughput-centric global interconnect design philosophy for 1 cm global line. Using level 49 transistor HSPICE models for a 250nm MOSIS process [13], the advantages of a throughput centric design are illustrated to corroborate the physical analysis in the previous sections.

In a typical 250nm technology, the target clock period is approximately 1.5ns (i.e. 667Mhz clock frequency) [2]. It is the assumption of this case study that the overhead delay due to setup and hold times, clock skew and jitter totals to approximately 250ps [6,10,12].

First, consider the single driver interconnect circuit ($k=1$). After subtracting the overhead delay from the clock period, the maximum delay of a global wire would be approximately 1.25ns in a latency-centric design approach and is marked as d_1 in Figure 8. At a clock rate of 667MHz, the throughput per unit area of this wire is 4 Terabits per second (Tbps) per cm^2 and is labeled as d_1 in Figure 9. As indicated in Figure 9, a clear optimal wire width of 400nm gives the maximum throughput per unit area of 6 Tbps/ cm^2 . This design point is marked d_2 , and represents a 50% increase in the throughput per unit area of this global interconnect structure. This optimal point

can be translated into either a 50% reduction in area for constant throughput or a 50% increase in throughput for constant area.

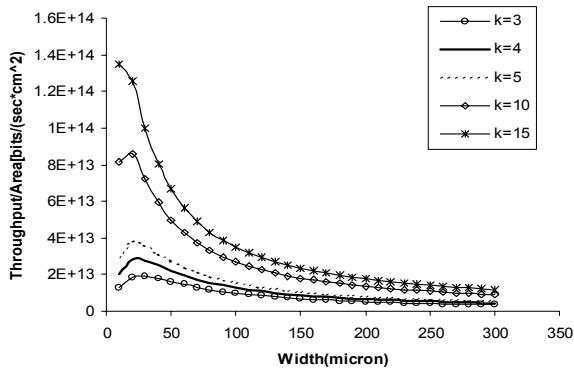


Fig. 7. Throughput/area for global interconnect with

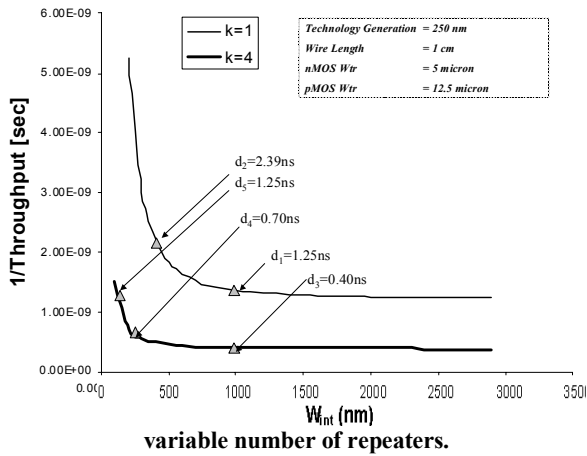


Fig. 8. HSPICE results for delay v/s wire width of a single driver and with 4 repeaters.

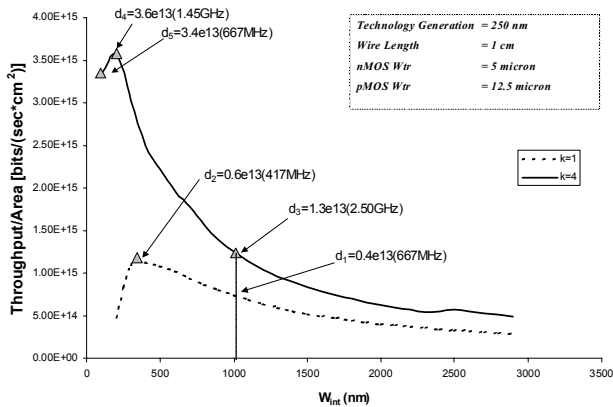


Fig. 9. HSPICE results for throughput /area of a single driver and with 4 repeaters.

The next simulation assumes that 4 repeaters are inserted into this 1cm interconnect [11]. In a latency centric

approach, the insertion of these repeaters with a 1 micron wire width reduces the delay to 400ps and is marked by d₃ in Figure 8. However, Figure 9 indicates that to maximize the throughput per unit area the interconnect width needs to be substantially reduced to 200nm (which of course is unrealizable in this technology but will be used for illustrative purposes only). The resulting optimal design requires a 1.45Ghz global clock.

Though with optimal wire sizing, the throughput per unit area can be increased by 3x to 36 Tbps/cm². Moreover, if the width is decreased further to 125nm, the clock returns to 667MHz and still has a throughput per unit area value that is only 6% off the absolute optimal as indicated by d₅ in the Figure 9.

5. SUMMARY

Interconnects are rapidly becoming a bottleneck for the performance and cost in high-speed VLSI circuits. This paper has explored the implications of DSM technology to quantify the benefits of shifting from a latency-centric to a throughput-centric design strategy. A new physical model has been derived in this paper that approximates the throughput of an interconnect repeater circuit, and is utilized to explore optimal wiring sizing for a throughput-centric design. Key case studies from 250nm to 70nm have indicated that optimal wiring sizing for a throughput-centric methodology can reduce wire area from 20-50%. Moreover, a throughput centric repeater circuit design could increase the throughput of a wire (e.g. 2.5x) with only marginal loss in wire latency.

6. ACKNOWLEDGEMENTS

The authors would like to thank the support of the National Science Foundation (NSF#0092450), the SRC Education Alliance, and the Georgia Tech SURE program for their support of this research.

7. REFERENCES

- [1] J.A. Davis, et al., *Proc. IEEE*, vol. 89, No. 3, pp. 305-324, March 2001.
- [2] ITRS 2001 edition. (<http://public.itrs.net>)
- [3] W.J. Dally, *Proc. IITC*, June 1999, pp.15-17.
- [4] P. Larsson-Edefors, *IEEE Trans. on Circuits and Systems-I*, vol.47, no.4, pp.602-606, Apr 2000.
- [5] A. Kowalczyk, et al., *ISSCC Tech Digest*, 2001, pp. 236-237.
- [6] S. Rusu, et al., *IEEE JSSC*, vol. 35, no. 11, pp.1539-1544, Nov. 2000.
- [7] H.B. Bakoglu and J.D. Meindl, *IEEE TED*, vol. ED-32, no.5, pp.903-909, May 1985.
- [8] T. Sakurai, *IEEE TED*, vol. 40, pp.118-124, Jan. 1993.
- [9] Bakoglu H. B. , "Circuits, Interconnections, and Packaging for VLSI", Addison Wesley 1990.
- [10] B. Geuskens and K.Rose, "Modeling Microprocessor Performance," Kluwer, Boston, 1998.
- [11] J. Culetu, et al., *Proc. DAC*, June 1998, pp.392-395.
- [12] G. Sai-Halasz, *Proc. of IEEE*, vol. 83, no. 1, pp.20-36, Jan. 1995.
- [13] <http://www.mosis.org/Technical/Testdata/> (MOS models)