

Title Page

Paper title: Optimal Shielding/Spacing Metrics for Low Power Design

Authors:

Ravishankar Arunachalam, IBM Corporation

Emrah Acar, IBM Corporation

San Nassif, IBM Corporation

Corresponding Author: Ravishankar Arunachalam

Address:

IBM Corporation

11501 Burnet Road M/S 904-6F015

Austin

TX 78758

Phone Number: 512-838-9454

Email: ravia@alumni.cmu.edu

Submission categories:

(1) Low Power Design

(2) Physical Design, Synthesis and Optimization

(3) Advanced VLSI Design

Optimal shielding/spacing metrics for low power design

Author(s)
Affiliation(s)

Abstract

Noise arising from line-to-line coupling is a major problem for deep submicron design, and present technology trends are causing an increase in this type of noise. Common current methods to decrease coupling noise include shielding and buffering, both of which can increase overall power dissipation. An alternative method is spacing, which has the added benefit of improving the manufacturability (i.e. defect insensitivity) of the design. This paper explores the issue of coupling noise reduction, and proposes performance metrics that can be used by the designer to determine which of the alternative methods is best suited for a specific interconnect configuration.

I. Introduction

Capacitive coupling is recognized as one of the most critical problems that designers need to address for deep submicron technologies. A commonly used technique to avoid coupling is to shield signal lines from each other by inserting power/ground lines in between them [3][6]. Although shielding practically eliminates coupling between signal lines, it does result in increased power and area. Other approaches to avoid crosstalk at the physical design stage include net ordering[3] and buffer insertion [5], and active shielding [4].

In this paper, we present a comprehensive analysis approach towards a decision making procedure for effective coupling avoidance. This procedure is primarily developed for design engineers/methodologists of high-performance microprocessors and application specific integrated circuits. Designers often have to optimize for several performance metrics like delay, area and power and these conflicting goals can lead to different coupling avoidance strategies. For example, it is desirable to trade-off noise reduction and power dissipation, a trade-off that has not been well recognized in literature. We present metrics and guidelines for determining the appropriate approach depending on the geometric parameters and the relative driver strengths. Our results demonstrate that spacing is a viable and more effective option than shielding for low power designs, even after budgeting for noise and delay increase due to coupling.

The rest of the paper is organized as follows: in Section 2, we discuss some preliminaries. In Section 3, we describe the important design metrics and our analysis method which forms the proposed evaluation procedure. Section 4 presents an analysis of our results from experiments performed. We conclude in Section 5.

II. Preliminaries

Consider one of the metal interconnect line shown in Fig. 1, with length L , width W , thickness T and height above the ground plane H . The capacitance of the wire is usually obtained by detailed extraction using explicit formula-based

approaches. For a line i , we denote the ground capacitance by C_{ii} and the coupling capacitance between lines i and j by C_{ij} . When two or more wires run parallel to each other, the effect of coupling capacitance depends on the gates driving the wires and the switching activity in the vicinity. For the configuration depicted in Fig. 1, the coupling capacitance between conductors 1 and 2 can be treated as an effective Miller capacitance of value $k \cdot C_{12}$, where k is the Miller factor. The total capacitance that is seen at the victim line is:

$$C_{total} = C_{11} + k_{12} C_{12} + k_{13} C_{13} \quad (1)$$

Due to symmetry, we will assume that $C_{12} = C_{13}$ and only refer to C_{12} in the rest of the paper. For same direction switching, the k value is less than unity, and for opposite direction switching it is greater than unity. The k value also depends on the relative strengths of the drivers. In [2], it is shown that values of -1 and 3 are bounds for the k value.

In order to increase circuit performance, feature sizes are continuously scaled to smaller dimensions. When the length, width and thickness of a wire are scaled down by a factor of α , the value of resistance is scaled up by α . In order to prevent an increase in interconnect delay due to scaling, the thickness of the wires is scaled less than α , thereby increasing the aspect ratio (T/W), and hence the coupling between the lines. With decreasing feature sizes, the coupling capacitance is becoming a larger portion of the total capacitance [7][11].

Coupling Avoidance Approaches

Inserting power/ground metal shields is a popular method to avoid the undesirable increases in coupling capacitance. Signal isolation prevents both functional noise and increase in delay due to coupled lines switching. [8] proposes inserting a power or ground shield after every signal wire, resulting in a dense fabric-like structure. It should be mentioned that inserting power/ground shields also reduces inductive effects because of the closer return path to ground for the current flowing through signal wires. However, inserting shield wires between every pair of signal wires is costly in




Fig. 1. Interconnect capacitance

area, increases congestion and may end up requiring more metal layers, leading to an increase in production costs.

Buffer insertion and net ordering have also been proposed as alternative methods to negate the effect of coupling. Buffer insertion requires extra budgeting for space, power and results in increased area. Net ordering involves selection of signals that cannot simultaneously switch in opposite directions and placing them adjacent to each other. This switching orthogonality is achieved by temporal de-correlation[3]. However timing information available before the global routing phase is a crude approximation of the actual timing, leading to inaccurate results.

Alternatively the wires can be simply spaced apart to produce a similar solution. Though spacing does not eliminate coupling noise, it reduces the coupling and at the same time reduces power dissipation since the total capacitance load of the line decreases. This is a significant gain compared to shielding, which eliminates the noise at the cost of extra power dissipation. An effective decision mechanism and detailed discussion of these approaches will help designers establish better methodologies which will better avoid the unwanted problems of coupling. This task requires the formulation of critical performance metrics and objective criteria to compare one approach vs. the other.

III. Proposed Decision Making Approach

Consider the case of three signal wires shown in Fig. 2. Fig. 2(a) shows the case when wires are unprotected, and strong coupling exists between neighboring wires. When shielding wires are inserted as shown in Fig. 2(b), the total capacitance of the victim wire V remains the same, but much of this will be to the neighboring shields. Instead of shielding, we can space the wires such that the same silicon area will be used. The spacing style is directly applicable to the design and does not require use of auxiliary optimization




Fig. 2. Different wire configurations used

algorithms. If S is the planned distance between two wires (assuming it is also the distance between the signal and its shields), the distance between two spaced signal wires will be $2S+W$, as shown in Fig. 2(c). Note also that with the assumption of a $2S+W$ spacing, the number of routing tracks will remain the same as that of the shielding approach. Another important aspect is that the comparison of shielding over $2S+W$ spacing will display the real return of isolating signal lines over simply spacing the signal lines apart. Since the use of silicon area for the shielding and spacing approach would be the same, these configurations are the correct candidates for objective performance comparisons.

III.A. Impact of spacing on coupling capacitance

The C_{11} and C_{12} capacitance values are computed by two methods in this paper: (i) A formula-based approach, which we obtained from the Berkeley interconnect tools evaluation[10], (ii) An in-house capacitance extractor, which uses a boundary element method. The capacitance values from both methods correlate very well and both were used in generating our results.

With the wire configuration shown in Fig. 2, Fig. 3 shows

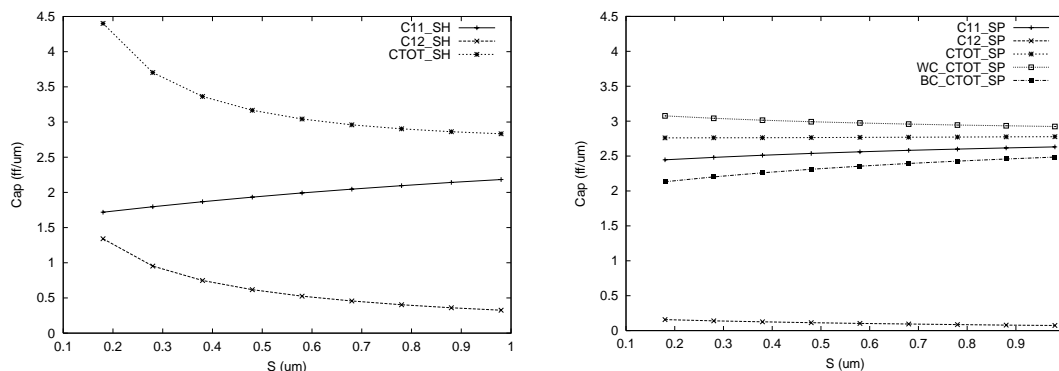


Fig. 3. Ground and Coupling capacitance (C_{11} and C_{12}) as a function of distance between wires. a) Shielding (original) configuration, b) Spacing configuration

plots of the ground capacitance and coupling capacitance to the neighboring line per unit length, as a function of the distance S between the wires. The width, thickness and height are fixed. ($H=T=0.4\mu$, $W=1.6\mu$). In Fig. 3(a) (the shielded configuration), the coupling capacitance is seen to decrease

sharply with increasing distance, while the ground capacitance increases because of more fringing effects as lines are separated farther.

In Fig. 3(b), the self and coupling capacitance values are plotted as a function of S , when the wires are separated by a

distance $2S+W$. In circuit operation, the coupling capacitance is *amplified* due to the Miller effect, as given by (1). Hence the best and worst-case total capacitances are also plotted, obtained using Miller coefficients of -1 and 3. However, it can be observed from the figure that the coupling is almost an order of magnitude smaller than the case in shielding. This can again be explained by the almost-exponential decrease in coupling capacitance with increasing distance. Therefore, the rate of decrease in coupling-capacitance for $2S+W$ spacing is not as aggressive as shielding, but overall the capacitance value is much lower than the shielding (and the original) case. The capacitance plot shown here depends on other geometry parameters like thickness (T), width (W) and height (H) of the wire from the ground plane, though the trends in the curve with respect to spacing remain the same.

III.B. Effect of shielding on inductance

Apart from eliminating noise due to capacitive coupling, an important benefit of shielding is to significantly reduce the inductance of signal wires. Since the shields are tied to Vdd/Gnd, the return path for the current through the signal is significantly closer, hence lowering the inductance. However, the inductive effect is only significant for so called “fat wires” with low resistance and fast transition times. For example, in [12] it is shown that in 0.25μ technology, a 1mm signal wire with a transition time greater than 100ps will exhibit insignificant inductive effect. The target for our approach is a large class of wires in mid-level metal layers in a low-power design which are likely to have more lossy characteristics, experience significant coupling noise but relatively little inductive effects. For the long and wide conductors in the uppermost metal wires, inductive effects become dominant and shielding is required to control noise.

III.C. Experimental setup and procedure

In order to investigate the effects of spacing and shielding on different performance metrics, we used the configurations depicted in Fig. 2 and used clock buffers to drive each line. All experiments were performed for a 0.18μ static CMOS technology. The length of wires was set at 1000μ for all cases because for wires longer than this value, buffer insertion would be applied resulting in a significant reduction in delay and slew. The wire widths were varied over a wide range of values to simulate wide as well as narrow wires. In current practice, the wire width is typically increased in order to reduce the resistance of long wires. Copper was assumed keeping with current trends for interconnect.

With the length of the interconnect fixed, we used different driver sizes to simulate different loads. The effect of coupling on noise as well as on delay is a strong function of driver resistance. Each interconnect line was driven by a clock buffer taken from a high performance commercial processor design. Seven different buffer sizes were used, with the strength/size of the buffers ranging from 4X to 70X of the minimum sized buffer.

The minimum distance between wires is a function of the technology parameters. For our experiments, we considered

the minimum value for S to be 0.18μ , and allowed it to vary up to 0.5μ . The thickness was varied from 0.2μ to 1.2μ to simulate a range of aspect ratios. The height H above the ground plane does not have a significant impact on the nature of results and we used only two different H values.

The best and worst-case delays due to coupling were obtained using the Miller cap approach, and the Miller coefficients were varied between -1.0 and 3.0 to simulate same direction as well as opposite direction switching. To compute the noise on the victim, the victim line was held at a constant value, the aggressor was switched in both directions and the worst-case noise was used. A typical circuit configuration which was used in simulation is illustrated in Fig. 4. The circuit is simulated with a distributed coupled RC network as the load. The delay and noise are measured at the sink node of the victim, marked in the figure. Different characteristics

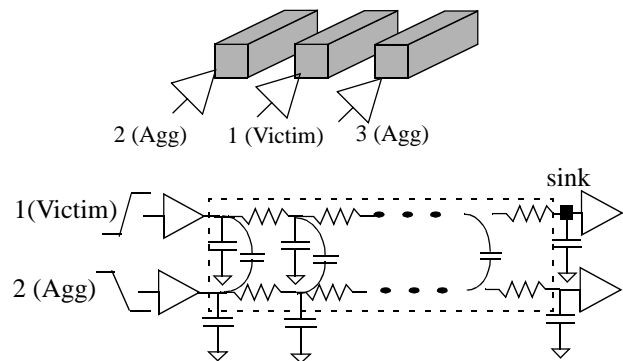


Fig. 4. Typical circuit configuration used in simulation

of the noise pulse - the shape of the waveform, pulse-width or peak value - can be considered. For simplicity, we use the peak of the noise waveform, though any other metric can be used since we have the complete noise waveform from circuit simulation.

IV. Simulation Results

For each configuration in our experimental setup, detailed simulations were carried out using PowerSpice, an in-house SPICE-like simulator. The interconnect was represented by a distributed RC line. The coupling capacitance extracted was also distributed along the lines. The input to all the clock buffers is a ramp waveform with a rise-time of 100ps. The slew at the output of the gate (which determines the current injected into the neighbor) is a strong function of the load, and a weaker function of the input slew. We consider four performance metrics: delay, slew, noise and power which are all critical for present day technologies.

IV.A. Noise

If coupling noise is sufficiently large in magnitude and duration, it may create functional errors by changing the value of the victim net or by creating a wrong state in the circuit. To evaluate the $2S+W$ spacing and shielding alterna-

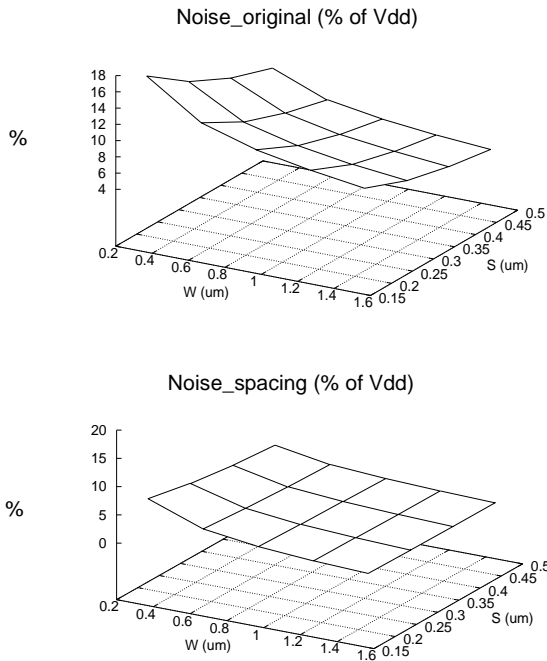


Fig. 5. Relative reduction in peak noise from the 2S+W spacing as a function of W and S.

tives with respect to the noise performance, we followed a similar set of experiments performed for delay. The same interconnect geometry parameters were used along with a medium-size buffer. As mentioned in the previous section, we selected the peak noise value as the metric of interest.

Shielding suppresses almost all of the coupling noise on the victim line. Therefore, we can only compare the noise performance of the 2S+W spacing with the original circuit configuration, where the wires were spaced by S without shields. The comparison of noise performance explains how much of improvement the spacing can provide over the original choice of interconnect.

Fig. 5(a) shows the noise value as a percentage of Vdd, for the original configuration with wires neither shielded nor spaced apart. Fig. 5(b) shows the same noise values for the 2S+W-spacing case. From the plot we see that the original peak noise is reduced considerably, with a reductions of 50-90% in magnitude. Since some amount of noise well within the noise margin can be budgeted for the design process, a quantitative criterion can be chosen by the designer to determine where the noise reduction with 2S+W spacing is acceptable.

IV.B. Power

Since power dissipation is one of the most critical performance bottlenecks for current and future circuits[13], any solution addressing the coupling problem must be considered from a power perspective. This metric is ignored in the shielding papers[8], thus we will discuss the fundamental issues related to the power dissipation in the context of coupling avoidance.

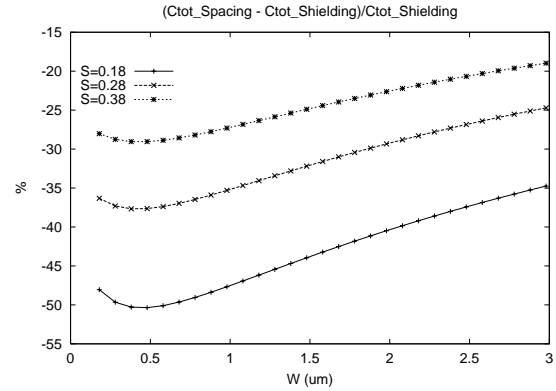


Fig. 6. Percentage reduction in total capacitance between shielding and 2S+W spacing configurations. Power dissipation on victim line varies linearly with total capacitance

Although architectural techniques are the most effective solutions for low-power design, circuit power can be greatly reduced by keeping the load capacitances as low as possible. Here we assume that the total power of a circuit is linearly related with the load capacitance it charges ($P=f_{clk}CV^2$).

Fig. 6 shows the relative difference between the total capacitance for the spacing and shielding approaches, again as a function of the wire geometries. The total capacitance, hence the expected power consumption is seen to be always smaller for the spacing case. From the plot we see that for large S values, the total capacitance value difference between both approaches do not vary significantly, but for smaller S, spacing can result in significantly reduced capacitance and hence power.

Another item that needs to be mentioned here is the extra work shielding requires for power distribution network design and verification. Today's circuits operate at very low voltages to save power and to increase the speed of the operation. Therefore, the reliability of the power supply and ground distribution network is of utmost important than ever. Since the IR-drops and Ldi/dt noise may impact the circuit performance negatively, the power distribution network must be low-resistance and uniform across the chip. With shielding approach, the power distribution network must be connected to the shield lines, using precious routing resources and metal area. Shielding will also increase the number of electrical elements in the power grid model and hence will complicate the analysis. Therefore, 2S+W spacing appears to be more favorable than shielding from a power perspective.

IV.C. Delay

Fig. 7 shows a scatter plot of the worst-case delay with shielding vs. 2S+W spacing for different values of S, W and driver strengths. The worst-case delay is the one computed when the neighboring line is switching in the opposite direction for the non-shielded case. The figure illustrates that for approximately 80% of the cases, the worst-case delay (late-mode) for spacing configuration is generally smaller than

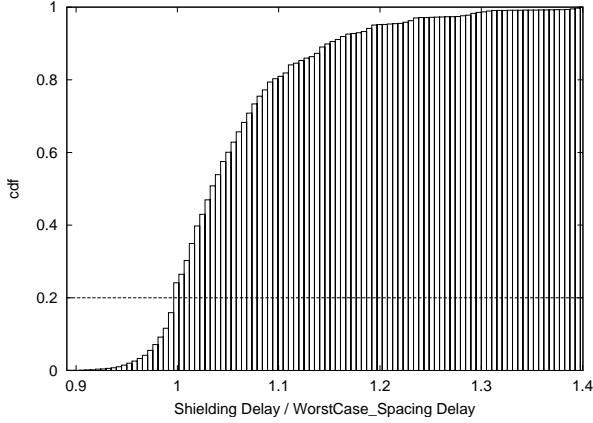


Fig. 7. Cumulative distribution function of the ratio of shielding delay to worst-case 2S+W delays for 2500 samples with different buffer sizes and S,W values

that obtained from shielding. Hence, with lesser metal area, power dissipation and less complex routing procedures, designers can achieve the same objectives of minimizing delay impact. This result confirms our previous results for capacitance distribution shown in Fig. 3. In many of the cases (roughly 80%) shown in Fig. 7, coupling capacitance of shielding is far bigger than that of the 2S+W case, and dominates the total effective capacitance load on victim. Although the uncertainty in delay is eliminated, this results in a larger delay for the shielding configuration for these cases. We also noticed that shielding actually eliminated the coupling impact and reduced the delay of the victim line for cases where wires are narrow and considerably spaced apart (low W and high S). For such cases, the approach of 2S+W spacing is not effective for delay performance.

IV.D. Slew

The slew on the victim line signal is also affected by the switching activities in the vicinity. Slew is an important metric since it affects noise and downstream delay and designers have constraints on the range of slew values possible for a signal wire. In our experiments, we found that the victim line slew performance closely tracks the delay. Hence the same decision criteria that are used to minimize delay can be applied to control slew degradation.

IV.E. Other considerations

The approach taken for coupling avoidance approach may impact several other parameters important for current design technologies. An important consideration for present-day technologies is the manufacturability of the circuit. In the 2S+W-spacing configuration, the area between metal lines is sufficiently large leading to a drastic reduction in the density of defects that may cause shorts and bridging faults. So spacing can improve the yield of the design. One of the most critical items is that of inductance, due to the manner in which it affects signal delays and may cause ringing. In order to control inductance values on global nets, shielding must be


applied at a periodic basis to provide a shorter current return path.

V. Optimal Decision Criterion

Our experimental results show that depending on the metric of consideration, shielding or 2S+W spacing approaches may yield different advantages and disadvantages. For better results, a comprehensive analysis and decision procedure must be employed by the designer for particular design goals. We expect this analysis procedure to be performed in advance and the values stored. In the light of the results obtained, we propose two simple but effective decision criteria applicable in determining the proper coupling avoidance approach.

Fig. 8(a) shows a contour plot of the difference between

$$\frac{(\text{WorstDelay_spacing} - \text{Delay_shielding})}{\text{Delay_shielding}}$$



$$\frac{(\text{Noise_spacing} - \text{Noise_original})}{\text{Noise_original}}$$




Fig. 8. Proposed criteria to perform decision making between shielding and 2S+W spacing approaches to overcome coupling effects

shielding case delay and worst-case 2S+W spacing delay for the victim line. The case for $H=T=0.8\mu$ is considered, and the contour levels are plotted as functions of W and S. For a simple delay-centric criterion, one can choose the use of a separator line on the (W,S) plane to prefer one approach over

the other. Depending on the interconnect size, a quick comparison between the lines can determine the advisable avoidance method. Another important decision criterion between the methods discussed is the magnitude of noise. Fig. 8(b) shows the relative difference in magnitude between the original noise and the noise after $2S+W$ spacing as a contour plot. Depending on the budget on the delay uncertainty, the designer may choose to use $2S+W$ spacing approach over shielding which seems to be increasing the delays, total capacitance and power for most of the cases. Therefore, a threshold can be also used with a simple constraint on delay uncertainty which can be easily measured by the help of Miller factors.

It is important to note that these contour plots can be generated from multiple circuit simulation runs which are performed in advance. During the design process, the designer can just look up the S , W and T values and determine if the resultant noise reduction satisfies the noise constraints. Other criteria can be easily developed with the mixtures of power and noise performances, and inequality type constraints can be deployed to make the coupling avoidance decision. The primary contribution of this paper is the objective evaluation method of the coupling avoidance approaches with considerations of the impacts on power.

VI. Conclusions and Future Work

In this paper, we have presented a comprehensive analysis and decision making methodology to determine a coupling avoidance strategy. The impact of shielding and spacing approaches result in critical differences in a large design when applied on a general basis to overcome the coupling problem. Excessive (unnecessary) shielding may significantly increase the total capacitance of the signal line, which dissipates more dynamic power in operation. On the basis of results obtained, we predict that spacing is a more viable option for low-power designs. However, it is unavoidable to insert some shield wires for fat top-level metal lines because of effects of inductance.

We have shown how to determine the optimal noise and coupling avoidance strategy for a specific interconnect configuration. We have explained how the designer can achieve the optimal trade-off between lower power and noise immunity by using a combination of spacing/shielding depending on the extent of coupling and the noise budget. As part of future work, we plan to make a more thorough study of the power-grid distribution problem as it impacts noise avoidance approaches and also consider top-level metal lines with extracted inductance in formulating the optimal criteria.

VII. References

- [1] K.L.Shepard, V.Narayanan and R.Rose, "Harmony: Static Noise Analysis of Deep Submicron Digital Integrated Circuits," *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, Vol.18, pp.1132-1150, August, 1999.
- [2] P.Chen, D.A.Kirkpatrick and K.Keutzer, "Miller Factor for Gate-level Coupling Delay Calculation," *Proc. IEEE/ACM Intl. Conf. on Computer-Aided Design*, 2000.
- [3] J.Ma and L.He, "Formulae and Applications of Interconnect Estimation Considering Shield Insertion and Net Ordering," *Proc. IEEE/ACM Intl. Conf. on Computer-Aided Design*, 2001.
- [4] H.Kaul, D.Sylvester and D.Blaauw, "Active Shields: A New Approach to Shielding Global Wires," *Proc. Great Lakes Symposium on VLSI*, pp. 112-117, April, 2002.
- [5] C.J.Alpert, A.Devgan and S.T.Quay, "Buffer Insertion for Noise and Delay Optimization," *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, pp.1633-1645, Nov., 1999.
- [6] H.P.Tseng, L.Scheffer, C.Sechen, "Timing and Crosstalk Driven Area Routing," *Proc. Design Automation Conference*, pp.378-381, 1998.
- [7] International Technology Roadmap for Semiconductors, <http://public.itrs.net/Files/2001ITRS/Interconnect.pdf>
- [8] S.Khatri, A.Mehrotra, R.K.Brayton, A.S.Vincentelli and R.H.J.M.Otten, "A novel VLSI Layout Fabric for Deep Submicron Applications," *Proc. Design Automation Conference*, pp. 491-496, 1999.
- [9] T.Xue and E.Kuh, "Post global-routing crosstalk synthesis," *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, Vol.16, pp. 1418-1430, Nov., 1997.
- [10] Handy Interconnect Tools, www-device.eecs.berkeley.edu/~xuejue/interconnect.htm
- [11] H.B.Bakoglu, "Circuits, Interconnections and Packaging for VLSI," Addison-Wesley, 1990.
- [12] Y.I. Ismail, E.G. Friedman and J.L. Neves, "Figures of Merit to Characterize the Importance of On-Chip Inductance," *Proc. Design Automation Conference*, pp.560-565, 1998.
- [13] www.intel.com/research/silicon/TeraHertzshort.pdf
- [14] PowerSPICE User's Guide, *IBM Corp.*, Hopewell Junction, NY