Optimal Inductance for On-chip RLC Interconnections

Shidhartha Das Kanak Agarwal David Blaauw Dennis Sylvester

{siddas,agarwalk,blaauw,dennis}@eecs.umich.edu
Department of Electrical Engineering and Computer Science
University of Michigan, Ann Arbor

Abstract

We propose the concept of an optimal inductance value that can substantially reduce delay of global RLC signals while maintaining good signal integrity (low ringing/overshoot). We exploit the fact that inductance results in faster transition times to improve delay of buffers in global signal lines. We observe that voltage overshoot, slew rate, and total line delay all show strong inflection points at the same value of inductance. At this optimal value of inductance significant improvements in signal transition time, and hence in overall signal delay, are obtained with negligible ringing. We propose adjusting the power grid to achieve this optimal inductance. Results show that the delay of a 1cm line with 9 inserted repeaters can be reduced by 8-12% with acceptable ringing by operating at the optimal inductance point.

1. Introduction

On-chip inductance has become significant in designs with GHz clock frequencies. Inductive effects in interconnects manifest themselves as overshoots/undershoots in voltage waveforms that may degrade signal integrity and cause reliability concerns. Hence, interconnect inductance is typically seen as a problem and various approaches have been proposed to minimize inductance [1-4]. However, it is known that inductance provides some useful effects such as improved signal transition times [5,7,8]. This effect of inductance can be used to improve propagation delay in global wires. In this paper we demonstrate that, while the delay of a single interconnect increases with inductance, the delay of global signals containing repeaters can be reduced by exploiting the fact that gate delays improve with faster input transition times. To take advantage of the improved slope effect of inductance, we seek out inductance values that result in improved slew and total delay with acceptable ringing. Our study shows that contrary to the belief that designers should always aim to minimize inductance, there exists an optimal inductance value where better slopes can be harnessed to improve performance of high-speed interconnects while keeping negative inductive effects at acceptable levels.

2. Concept

Inductance results in faster slew rates due to the fact that when a positive voltage is applied to an inductor, it takes some time to build up the charging current. Once the current is established, however, it continues to be supplied for some time, resulting in an overall faster transition time [7]. Figure 1 shows the far-end RC and RLC waveforms of a 2mm long interconnect driven by a 60X inverter in a 0.13µm technology.[1] It is clear from the figure that in this case, the RLC waveform is superior to the RC waveform. This is because the RLC waveform experiences a faster transition than RC with almost the same 50% delay. Also, for this case the RLC waveform does not exhibit significant ringing or overshoot. Increasing inductance further will result in even better transition times but will cause a more significant interconnect delay increase

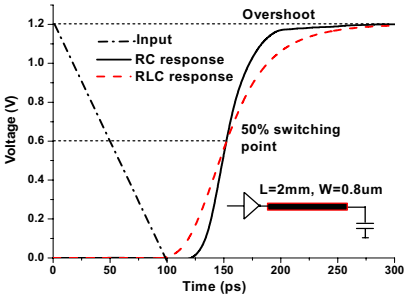

Figure 1. Far-end RC and RLC waveforms of a 2mm line driven by 60X inverter

and signal ringing making the RLC waveforms less desirable than their RC counterparts. However, for the case shown in Figure 1, the RLC waveform has better slope with little negative effects and this useful effect of inductance can be exploited to improve performance.

Typically, repeaters are inserted in long interconnects to improve propagation delay. Gate delays reduce significantly with improvement in input slew and hence, the overall delay of a line containing repeaters can be improved by increasing inductance from its minimum value. In fact, the slew improvement with increased inductance results in a cascading effect where the gain in one stage is propagated to the next stage, resulting in additional improvements with a higher number of stages. This suggests that if we can control the inductance of lines, then the delay of global interconnects containing repeaters can be improved by harnessing the faster transition time effect of inductance. It should be noted, however, that as inductance is increased, interconnect delay and ringing also increase and it can offset the gains due to improved slope and degrade signal integrity. Hence, it is required that the trade-off between the slew improvement and the associated increase in interconnect delay and ringing with inductance is balanced properly. In this paper, we explore this trade-off and propose an *optimal inductance point*. Our results show that this point results in an ideal trade-off between delay and signal integrity. We also propose an analytical expression for obtaining the optimal inductance by controlling the power grid spacing.

3. Analysis and Discussion

In this section, we analyze the above concept by simulations. For our experimental setup, we consider a 10mm long global interconnect using repeaters as shown in Figure 2. The total number of repeaters inserted is nine with a 1mm interconnect segment between two repeaters. Line width and thickness are

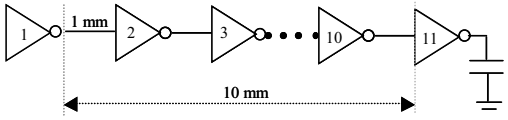

Figure 2. Line configuration

[1] Driver size 60X implies an NMOS width in the inverter of 60 times minimum channel length (2*0.13µ). PMOS is twice as wide as NMOS.

This work is supported by grants from the NSF and Intel Corporation.

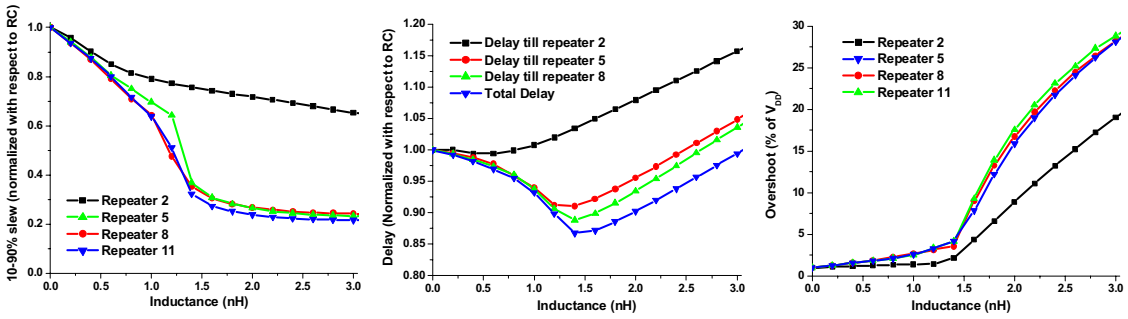

(a) (b) (c)

Figure 3. (a) 10-90% slew vs. inductance, (b) 50% delay vs. inductance and (c) overshoot vs. inductance at different stages

0.8μm and 1μm respectively. The size of the repeaters is 45X in a 0.13μm technology. The resistance and capacitance are extracted using the commercial extraction tool Raphael. To find an optimal inductance value for this configuration, we keep resistance and capacitance fixed and sweep the inductance value of each 1mm segment from 0 to 3nH.

Figure 3a shows the improvement in slopes with inductance. The figure shows 10-90% slew (normalized to the RC case) measured at the input of repeaters 2, 5, 8, and 11. It is clear from the figure that significant improvements (up to nearly 5X in this case) in slew can be obtained by increasing the inductance. The figure also shows that slope improves rapidly when small inductances are included, but saturates beyond some point, showing only small improvements with further increases in inductance. The improvement in slope at the input of Repeater 2 is less than that at the input of other repeaters. This occurs because Repeaters 3 through 11 benefit from a cascading effect as improved slopes are propagated through the inverter chain, while the slope at the input of Driver 1 is fixed in the simulation setup.

Next, we plot delay at various stages with inductance as shown in Figure 3b. As inductance is increased, the delay first decreases, becomes minimum at an optimal point, and then starts to increase. The initial delay reduction is due to the significant improvement in repeater delay as slew rates decrease, while the interconnect delay itself stays roughly constant in this region. When the inductance is increased beyond a certain value, however, the gain in repeater delay is overcome by the increase in interconnect delay and hence total line delay starts to rise. Figure 3b also shows the cascading effect where the gain in delay improves with the number of repeaters (for instance, the relative gain in delay at Repeater 8 is more than the gain at Repeater 5). We note that the delay until Repeater 2 increases monotonically with inductance. This is because the input of Driver 1 has a fixed input slope and hence no improvement in its delay is obtainable while the larger inductance yields an increase in the delay of the first wire segment.

In Figure 3c, we plot the overshoot as a function of inductance. Two important observations can be made from this plot. First, the overshoot is almost negligible for small values of inductance and second, there exists a clear inflection point beyond which the overshoot starts to increase rapidly. The key point is that the inductance value at which the delay is minimized is the same value after which the overshoot starts increasing rapidly. This is also the same inductance value after which the improvement in the slope saturates (Figure 3a). For example, in the above experiment, this value is about 1.4nH. We refer to this point as the *optimal inductance point*, as it represents an operating point with minimum total delay while ringing is still negligible.

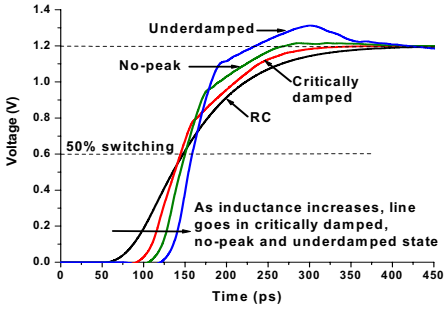

Figure 4. Comparing waveforms for RC, critically damped, no-peak, and underdamped RLC cases.

It was observed in [6] that the inductive overshoot is initially a weak function of inductance and remains negligible until the so-called *no-peak* condition, beyond which it starts to increase rapidly. Reference [6] shows that the no-peak condition can be analytically calculated using a second-order approximation of the distributed RLC response and that at the no-peak condition, overshoot is limited to only ~4% of the voltage supply. For a distributed line of length l with a driver resistance R_D, load capacitance C_L and per unit length line parameters R, L, C, the no-peak condition satisfies following equation [6]:

$$rl = \frac{-br_D + \sqrt{d - fr_D^2}}{g} \quad (1)$$

$$b = 2(1 + 3c_L + 3c_L^2) \qquad f = 2(1 + 6c_L + 15c_L^2 + 12c_L^3)$$

$$d = 6(1 + 2c_L)(1 + 4c_L + 6c_L^2) \qquad g = 1 + 4c_L + 6c_L^2$$

where, $rl = \frac{Rl}{\sqrt{L/C}}$, $r_D = \frac{R_D}{\sqrt{L/C}}$ and $c_L = \frac{C_L}{Cl}$

In this paper, we make the key observation that the inductance yielding an optimal delay also occurs at the no-peak condition, as shown in Figure 3. Our experiments show that the no-peak point not only yields acceptable overshoot, but also provides optimal delay and very good slew as well. This is illustrated in Figure 4, which shows the far-end waveforms of a 4mm line driven by a 100X driver. Inductance of the line is increased and the waveforms for RC, critically damped, no-peak, and underdamped states are plotted. The figure shows that at the no-peak point, slope is improved with controlled overshoot and negligible delay overhead. The figure also shows that the no-peak point is different from the critically damped point and additional improvement in slope (without significant ringing) can be obtained by operating at this point instead of the critically damped state.

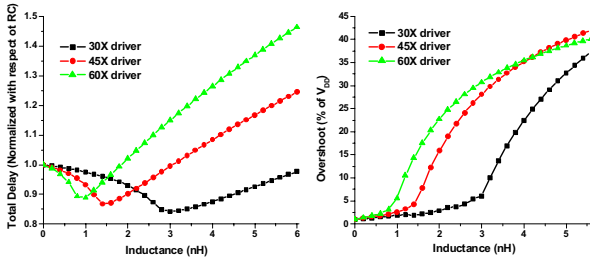

(a) (b)

Figure 5. (a) 50% delay vs. inductance and (b) overshoot vs. inductance for different driver sizes

Equation 1 shows that the no-peak point is a strong function of driver resistance. For very strong drivers (small driver resistance), the optimal inductance point is reached at a small value of inductance and vice versa. In order to clarify this point, we consider the configuration from Figure 2 (a 10mm global line with 9 inserted repeaters) and plot delay and overshoot for varied repeater sizes. Figure 5a shows the total line delay and Figure 5b shows the far-end overshoot for 30X, 45X, and 60X driver sizes. These plots demonstrate that as driver size is decreased, the optimal inductance value increases because the no-peak state is reached at higher values of inductance. The plots further confirm that the optimal inductance point for minimum delay corresponds to the no-peak point after which the overshoot starts increasing.

4. Designing for Optimal Inductance

In order to apply the concept of optimal inductance to on-chip interconnects, it is necessary to control their inductance value. The optimal inductance for an interconnect signal in a real design can be controlled in a number of ways, such as by adjusting driver size, wire width, inserting shields, and power/ground grid spacing. In this paper, we propose the use of power grid spacing to obtain the optimal inductance for high-speed interconnects. While power grid spacing at lower metal layers is typically dictated by physical design and IR-drop constraints, significant flexibility in the power grid spacing exists for higher layer metals that are used for high-speed global interconnects. IR-drop requirements for the power supply lines can be met by compensating for increased power grid pitches with increased power grid line widths, thereby allowing the pitch to be determined by optimal inductance.

We now show how the optimal inductance can be determined by controlling the power grid pitch. Inductance is a function of the current return paths. In on-chip signals, the return current primarily flows through the power grid and shields inserted between signals. Therefore, one approach to obtain optimal interconnect inductance is by controlling either the position of the shields, the pitch of the power grid, or both. We propose that while designing the power grid and performing shield insertion, inductance should be considered along with IR-drop constraints. One issue with controlling shield insertion for optimal inductance is that shields are important for reducing cross coupling effects and the optimal shield placement for inductance may not be ideal for minimizing coupling. However, once the shield positions are fixed for optimal inductance, the remaining coupling noise violations can be minimized by increasing metal-to-metal spacing as well. In fact, spacing has an additional advantage over shielding in that it reduces total capacitance and further improves performance [12].

We consider the experimental setup shown in Figure 6. We have a 10mm long signal line containing repeaters similar to the one shown in Figure 2. There are power and ground lines on either side of the signal as shown in the figure. To control inductance, the pitch of the coplanar power lines is varied. As the pitch is increased, the inductance of the signal line also increases since the area of the loop formed by the signal and its return current becomes larger. For each position of the power grid line, we extract inductance using FastHenry [10]. Figure 7 shows the 10-90% slew, 50% delay, and overshoot as a function of the power grid pitch. These plots are similar to the ones in Figure 3 except that now the inductance is varied by changing the position of the current return paths. Inductance has a logarithmic dependence on spacing and hence a log scale is used for the x-axis. The power grid pitch was varied from a minimum possible of 4.8μm to 1000μm and extracted inductance values for these endpoints were found to be 0.4nH and 1.8nH respectively.

Figure 7 shows that slew, delay, and overshoot exhibit the same pattern as obtained by directly sweeping the inductance. The figure also shows that the optimal operating point corresponds to a power grid pitch of 200μm. By operating at this optimal power grid pitch, a considerable improvement in delay (12% in this case) is obtained. We point out, however, that not all lines can operate at

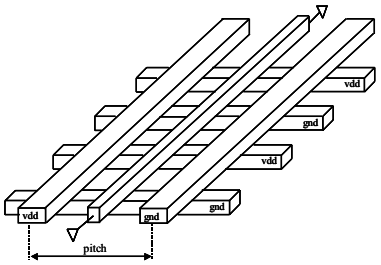

Figure 6. Experimental setup for controlling inductance by varying power grid pitch

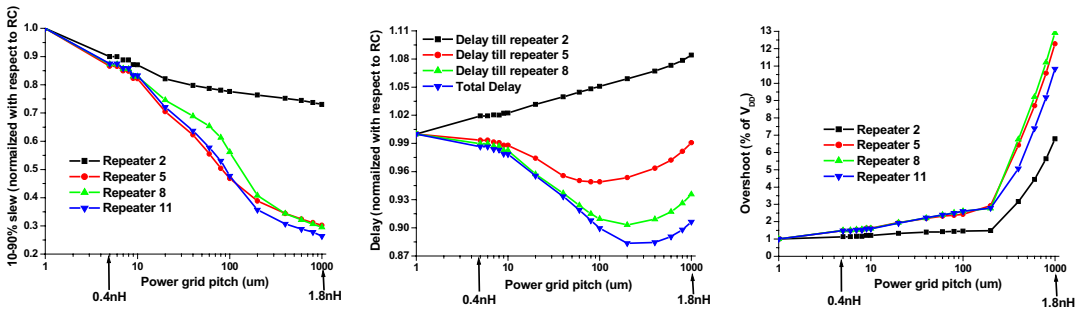

(a) (b) (c)

Figure 7. power grid pitch vs (a) 10-90% slew (b) 50% delay and (c) overshoot

their optimal inductance as some must be routed near the grid returns. Hence, priority should be given to timing critical signals.

We now develop an analytical expression for optimal power grid pitch. This expression will help designers to predict the optimal pitch for high-speed interconnect, eliminating the need to find the optimal pitch through laborious extraction and circuit simulation runs. To find the optimal power grid pitch, we first need to find the optimal inductance value using Equation 1. Solving the no-peak expression of Equation 1 for optimal inductance gives

$$L_{no\_peak} = \frac{(Rl + bR_D)^2 + fR_D^2}{d} Cl \quad (2)$$

Here, R and C are line resistance and capacitance per unit length. R_D is driver resistance, l is line length, and b, f, and d are coefficients as defined in Equation 1.

The loop inductance of a line can be approximated by

$$L_l = L_{self\_sig} + L_{self\_power} - 2M_{sig\_power} \quad (3)$$

L_{self_sig} and L_{self_power} are the self-inductances of signal and power line respectively and M_{sig_power} is the mutual inductance between them. Self and mutual inductances can be expressed as [9]:

$$L_s = \frac{\mu_0 l}{2\pi}[\ln(\frac{2l}{w+t}) + 0.5 + 0.2235(\frac{w+t}{l})$$
$$M = \frac{\mu_0 l}{2\pi}\left[\ln(\frac{2l}{s}) - 1.0 + (\frac{s}{l})\right] \quad (4)$$

Here, w, t, and l correspond to the width, thickness, and length of the lines, and s denotes the center-to-center spacing between power and signal lines. By combining Equations 3 and 4, the loop inductance can be expressed as:

$$L_l = \frac{\mu_0 l}{2\pi}\left[3 + \frac{w_{sig}}{l}\{0.2235(1 + k_p + 2k_t) - 2k_s\} + \ln\left\{\frac{k_s^2}{(1+k_t)(k_p+k_t)}\right\}\right]$$
$$\text{where } k_t = t / w_{sig}\ , k_p = w_p / w_{sig}\ , k_s = s / w_{sig} \quad (5)$$

To operate at the optimal inductance point, loop inductance should be equal to the no-peak inductance value. Hence, optimal power grid pitch can be obtained by equating L_l from Equation 5 with L_{no_peak} from Equation 2 and solving iteratively for k_s. Once the value of k_s is obtained, then the optimal spacing s and hence the optimal power grid pitch ($2*s$) can be easily calculated.

We verified the above analytical approach by comparing experimental results with analytical calculations. It was shown in Section 3 that optimal inductance is a function of driver resistance. We use the same setup as in Figure 6 and consider different driver sizes. For each driver size, we sweep the power grid pitch and find the optimal pitch value (i.e., the pitch at which the delay is minimum) and measure gain in delay corresponding to this optimal pitch. This experiment is similar to the one shown in Figure 6 where the optimal pitch was measured to be 200µ and the gain in delay at that pitch was 11.7%. Next, we calculate the optimal pitch value using our analytical approach and find the delay gain using SPICE for that optimal pitch. The results are shown in Table I and demonstrate that the analytical approach matches very well with the experimental results. The table also shows that as driver size is reduced, the optimal pitch increases. This is due to the fact that, for smaller driver sizes the optimal point occurs at higher inductance values.

Inductance has a logarithmic dependence on spacing (or pitch). For smaller drivers, the value of optimal power grid pitch is very high. At such high values, the analytically calculated optimal pitch may differ more significantly from experimentally obtained result. However, at such high values of pitch, the sensitivity of inductance to power grid pitch is very small. Hence, although the calculated optimal pitch may differ significantly from the measured optimal pitch, the delays in both cases match well.

Table I. Comparison between analytical and experimental optimal power grid pitch

Driver Size	Optimal pitch (µm)		Delay at optimal pitch (ps)		Delay Gain	
	Exp.	Anal.	Exp.	Anal.	Exp.	Anal.
75x	12	15	303	304	8.1%	7.7%
60x	24	20	323	324	9.6%	9.3%
45x	200	190	359	359	11.7%	11.7%
35x	1200	900	407	416	12.6%	10.7%

5. Conclusion

We proposed an approach where the delay of global interconnects can be improved by exploiting the faster transition time effect of inductance. An analytical expression for this optimal inductance is given. We propose the use of varying power grid pitch to obtain optimal inductance and develop analytical expressions for optimal power grid pitch. Our results show that by designing power grid with this optimal pitch, the delay of a repeated 1 cm line can be improved by 8-12%.

References

[1] Y. Massoud, J. Kawa, D. Macmillen and J. White, "Modeling and analysis of differential signaling for minimizing inductive crosstalk", *Design Automation Conference*, pp. 804-809, 2001.

[2] G. Zhong, C. K. Koh and K. Roy, "A Twisted-Bundle layout structure for minimizing inductive coupling noise", *Intl. Conf. On Computer Aided Design*, pp. 406-411, 2000.

[3] Y. Massoud, S. Majors, T. Bustami, and J. White, "Layout techniques for minimizing on-chip interconnect self inductance", *Design Automation Conference,* pp. 566-571, 1998.

[4] L. He and K. M. Lepak, "Simultaneous shield insertion and net ordering for capacitive and inductive coupling minimization", *Intl. Symp. on Physical Design,* pp. 55-60, 2000.

[5] Y. Ismail, E. Friedman and J. Neves, "Exploiting the on-chip inductance in high-speed clock distribution networks", *IEEE Trans. On VLSI*, pp. 963-973, Dec. 2001.

[6] J. Brews, "Overshoot-controlled RLC interconnections", *IEEE Trans. On Electron Devices*, pp. 76-87, Jan. 1991.

[7] M. Shoji, "*High-Speed Digital Circuits",* Addison-Wesley-Longman, 1996.

[8] K. Gala, D. Blaauw, J. Wang, V. Zolotov and M. Zhao, "Inductance 101: Analysis and Design Issues*," Design Automation Conference*, pp. 329-334, 2001.

[9] Y. Lu, M. Celik, T. Young and L. T. Pileggi, "Min/max On-chip Inductance Models and Delay Metrics," *Design Automation Conference*, pp. 341-346, 2001.

[10] M. Kamon, M. Tsuk and J. White, "FastHenry: A Multipole Accelerated 3D Inductance Extraction Program", *IEEE Trans. On Microwave Theory & Techniques,* pp. 1750-1758, Sept. 1994.

[11] G. Zhong, C. K. Koh and K. Roy, "A Twisted-Bundle layout structure for minimizing inductive coupling noise", *Intl. Conf. On Computer Aided Design*, pp. 406-411, 2000.

[12] R. Arunachalam, E. Acar and S. Nassif, "Optimal Shielding/Spacing Metrics for Low Power Design", *IEEE Annual Symp. on VLSI*, pp. 167-172, 2003.