

Characterization of Equalized and Repeated Interconnects for NoC Applications

Byungsub Kim and Vladimir Stojanović

Massachusetts Institute of Technology

Editor's note:

As the number of cores increases and on- and off-chip bandwidth demand rises, it is becoming increasingly more difficult to rely on conventional interconnects and remain within the chip power budget. This article explores leveraging equalization for global and semi-global long interconnects to overcome this problem.

—Li-Shiuan Peh, Princeton University

■ **MULTICORE PROCESSOR ARCHITECTURES** with increased parallelism, modularity, and hardware customizations may be the only way to continue scaling processor performance under power constraints.^{1,2} This promise hinges on the assumption that the power overhead of communication between modules is far smaller than the power used for computation, and hence requires an energy-efficient network-on-chip (NoC). Emerging multicore designs will require distributed switching and a relatively large NoC.

One example is the 80-core Intel Terascale NoC,² which Intel designed using traditional repeated interconnects under power and area constraints. This NoC has limited energy efficiency because of the inherent trade-offs in repeated wires and poor wire scaling with technology scaling.³ To improve the bandwidth and latency of long on-chip wires, designers have started focusing on equalized interconnects.^{4,5} Although equalization was originally used to improve speed,⁴ it also offers energy efficiency.^{5,6} In both repeated and equalized interconnects, truly optimizing the interconnect means jointly optimizing the circuits and wire parameters for best performance

under power and metal area constraints.

Estimating the power and performance of equalized interconnects is difficult due to the link's complex nature. In equalized interconnects, the communication channel consists of circuit parasitic and wire components, and the equalized pulse shape deter-

mines both the performance (achievable data rate and latency) and power consumption. To support fast quantitative estimation, we developed a modeling and optimization framework of equalized interconnects and demonstrated their energy efficiency benefits over repeaters in a 90-nm technology node.⁶ Using this tool, we identified trade-offs of repeated and equalized interconnects for various NoC scenarios in a 32-nm technology node. By exploring the interconnect design space, we obtain the optimized interconnect metrics necessary for a higher-level NoC architecture study. Results from our design space exploration indicate that using equalized interconnects instead of repeated interconnects provides significant latency (at least 2×) and energy (up to 10×) savings for 5-mm to 15-mm express lanes in global-routing layers. At lower metal layers, the energy efficiency benefits diminish, with a breakpoint occurring at the M5 and M6 semiglobal layers; however, the latency benefits remain.

Hierarchical-design considerations

Optimizing an NoC's performance under tight power and metal area constraints requires a cross-

layer design approach. Our tool optimizes performance across the circuit and network layers. Using an interconnect model (repeated or equalized), circuit designers can provide a set of trade-off curves and link metrics to enable evaluation of various NoC topologies under power and area constraints. With these metrics, combined with switch and router models, designers can balance resources (power and area) for various NoC topologies. They can then apply NoC benchmarks to extract the network performance for different applications, propagating the decision to lower levels of the design hierarchy.

Interconnection scenarios

A mesh network often serves as a baseline for multicore NoCs because of its simplicity, modularity, and reasonable bandwidth efficiency. Advanced NoC topologies, such as concentrated mesh (Cmesh) and flattened butterfly, improve latency and power consumption for the same total cross-sectional bandwidth by adding long-distance interconnects (or *express lanes*).⁷

By reducing the latency and energy cost of these long-distance interconnects, equalization can make these NoC topologies even more beneficial in a power-constrained setting. Figure 1 presents scenarios in which equalization can play a role in a 64-core processor NoC. Because today's typical processor edge length is around 10 mm to 20 mm,² we assume a 20 mm × 20 mm high-end processor die size. Figure 1a depicts the topology of a mesh network with four cores per router,⁷ connected with 5-mm-long links. Figure 1b shows a Cmesh topology with 10-mm-long express lanes along the processor edges, improving message congestion and latency. The flattened butterfly in Figure 1c uses additional 10-mm and 15-mm links to further improve throughput and latency by fully connecting routers within each row and column. For these die dimensions, the interconnect distances required in other NoC topologies are usually around

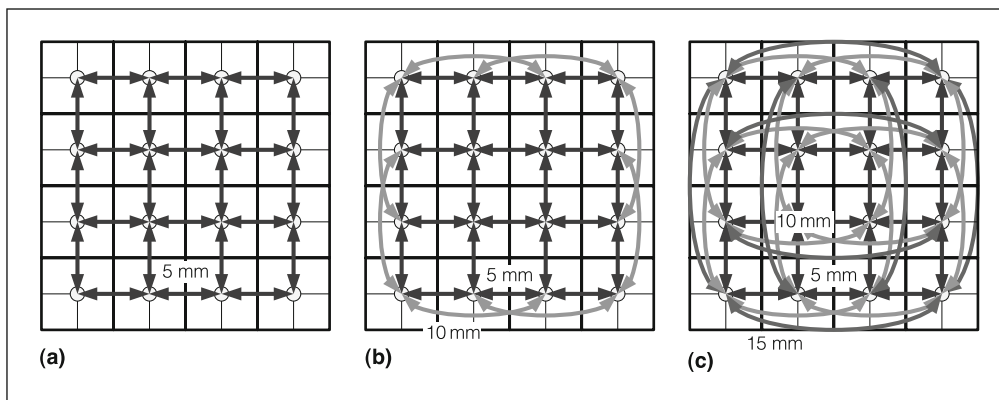


Figure 1. Example link lengths for possible application of equalized interconnects in various network-on-chip (NoC) topologies: 5-mm mesh links (a); 5-mm and 10-mm concentrated mesh (Cmesh) links (b); and 5-mm, 10-mm, and 15-mm flattened butterfly links (c).

5 mm, 10 mm, or 15 mm. Because these express lanes aren't free, we examine the trade-offs of repeated and equalized interconnects for these example distances to provide metrics for network-level analysis of various NoC topologies under global power and area constraints.

NoC interconnect metrics and interconnect optimization

In metal-area and power-constrained NoC designs, we want to maximize the aggregated data rate (D_a) of interconnect bundles through given cross-sectional metal width (w_c), while satisfying the target latency and given interconnect power constraints. This problem is equivalent to maximizing data rate density (d_d), which we can find by dividing the data rate by each interconnect's cross-sectional width while keeping the target latency and power constraints. We also consider power consumption: Regardless of the number of interconnects in the NoC, the fair comparison of the energy cost is the amount of energy to send 1 bit. Therefore, the interconnect bundle optimization problem is equivalent to the following optimization problem for each interconnect:

$$\begin{aligned}
 & \text{Maximize } d_d \\
 & \text{subject to } T_d \leq \text{target latency constraint} \\
 & \quad E_b \leq \text{target energy cost constraint} \quad (1)
 \end{aligned}$$

where $d_d = D_a/w_c$ (Gbps/ μm), T_d is the actual latency (in ps), and E_b is the actual energy cost (in pJ/bit).

Repeater insertion in RC-dominant wires

Early studies of repeater insertion mostly focused on optimizing delay.⁸ More recent studies have explored delay-energy products,³ with energy efficiency gaining importance. Other studies suggest optimizing data rate density using wave pipelining.⁹ However, none of these design methods handle all three major metrics (data rate density, energy per bit, and latency) simultaneously, which is our goal (Equation 1). Here, we extend previous work by deriving approximate formulas for these metrics, enabling fast computation and intuitive understanding of this multidimensional repeater optimization.

For the best performance and power efficiency, we use wave-pipelined interconnects rather than flip-flop pipelines. In wave pipelining, a transmitter can launch multiple bits, one after another, as long as no interference among bits occurs at the receiver.⁹ In long repeated interconnects, multiple bits can thus be in flight at the same time.

We optimize wire width (w), space (s), repeater size (w_n), and segment length (l) for Equation 1. We express the repeater delay using Bakoglu's formula as

$$T_d = L \left[\frac{(a+b)R_v C_g}{l} + \frac{R_v C_w}{w_n} + \frac{lR_w C_w}{2} + aR_w C_g w_n \right] \quad (2)$$

where L is the wire length, R_v is the resistance of NMOS per unit gate width, C_g is the gate capacitance per unit gate width, and a and b are two technology-dependent constants from the ratio of parasitic capacitance of NMOS and PMOS, respectively.⁸

Ho derived the energy per bit as

$$E_b \approx a_{ct} \left[C_w L V_{DD}^2 + \frac{c_r L}{l} (a_e + b_e) C_g w_n V_{DD}^2 \right] \quad (3)$$

where a_{ct} is the activity factor; V_{DD} is the supply voltage; a_e and b_e are technology-dependent constants computed from the amount of charge stored in inverters' gate and drain parasitic capacitances; and c_r is a technology-dependent constant (typically slightly larger than 1), representing the power consumption due to crowbar currents.³

To address data rate density, we relate the delay and data rate in Equation 4, where N is the number of segments, T_s is the bit time (minimum toggling period), and η is the technology-dependent constant (whose typical value is around 4).

$$T_d = (L/\eta l) T_s = (N/\eta) T_s \quad (4)$$

We can understand Equation 4 from the Elmore delay model's RC time constant by noting that both the delay of one segment (T_d/N) and T_s (for 95% signal swing) are approximately proportional to the segment's time constant:

$$N_d = T_d/T_s = N/\eta \quad (5)$$

By restating Equation 4 for a normalized delay $N_d = T_d/T_s$, we can derive an intuitive formulation in Equation 5, which simply states that the ratio of the delay to the minimum toggling period is approximately proportional to the number of stages. Therefore, even for a given energy budget E_b , we can trade off T_d and T_s . Using Equations 2 through 4 and optimizing parameters w , s , w_n , and l for optimization goals (Equation 1), we can generate 3D trade-offs between energy, latency, and data rate density metrics for repeated interconnects.

Equalization of RC-dominant wires

Rather than inserting repeaters, we can use signal-conditioning techniques such as equalization to overcome the wire bandwidth limitations on a point-to-point link. Figure 2a shows an example equalizer over RC-dominant wires. Figure 2b plots example waveforms of unequalized and equalized pulse responses. Without equalization, a transmitted pulse (dashed bold line in Figure 2b) disperses over a long wire, arriving at the receiver attenuated and with a long tail, causing intersymbol interference (ISI) (solid bold curve in Figure 2b), which limits the data rate. ISI limits the maximum data rate because the transmitter can send the next bit only after the previous bit's ISI tail diminishes.

In an equalized interconnect, a transmitter feed-forward equalizer (FFE) uses a finite impulse response (FIR) filter to shape the transmitted pulse (dashed nonbold line in Figure 2b) so the received response has smaller and shorter ISI (solid nonbold curve in Figure 2b), allowing higher data rate. In addition to the FFE, a one-tap decision feedback equalizer (DFE) at the receiver cancels the first trailing ISI tap, relaxing the FFE. In the frequency domain, the FFE effectively attenuates the transmitted signal's low frequencies to match the wire loss at higher frequencies when the signal arrives at the receiver. Therefore, a highly

sensitive differential regenerative receiver recovers the transmitted bit value.

Equalization improves the latency close to that of the speed of light.⁴ The propagation velocity at high frequency in the RC-dominant channel draws near to the speed of light, whereas at low frequency it's far slower. Unlike the case for a repeated interconnect, equalization uses the fast-propagation velocity at high frequencies and thus achieves small latency.⁴ By shaping the transmitter pulse, the FFE also shapes the received signal's phase, so the receiver observes the constant delay over the baseband frequency range after sampling. This equalized delay is the channel-propagation delay at the Nyquist frequency, as we explain later.

Equalization also improves the transmission's energy efficiency. A transmitting FFE's output voltage swing is attenuated along the wire. For example, although the FFE output swings rail to rail, the received voltage typically swings 50 mV to 100 mV because of high attenuation in the channel. This mechanism effectively reduces the total amount of charge that the FFE driver must inject into the wire, unlike the case in which the repeater drives all wire capacitance rail to rail.

Because the smaller signal amplitude requires less charge for a long capacitive wire, and a narrower pulse response allows a higher data rate, equalized interconnects can have a better performance-power trade-off than repeaters, depending on the equalizer circuit and wire characteristics.

Figure 2a shows a physical implementation of the equalized interconnect with available design parameters (V_s , V_p , w , s , w_{LCM} , \underline{w}^T , and y_1), which we optimize for Equation 1. In this example, we segment the low-common-mode (LCM) driver to form the FFE.¹⁰ An LCM driver consists of a digitally programmable voltage divider. Signal supply voltage V_s and predriver supply voltage V_p directly affect the signal swing and the power consumption. The wire width (w) and the space between wires (s) are

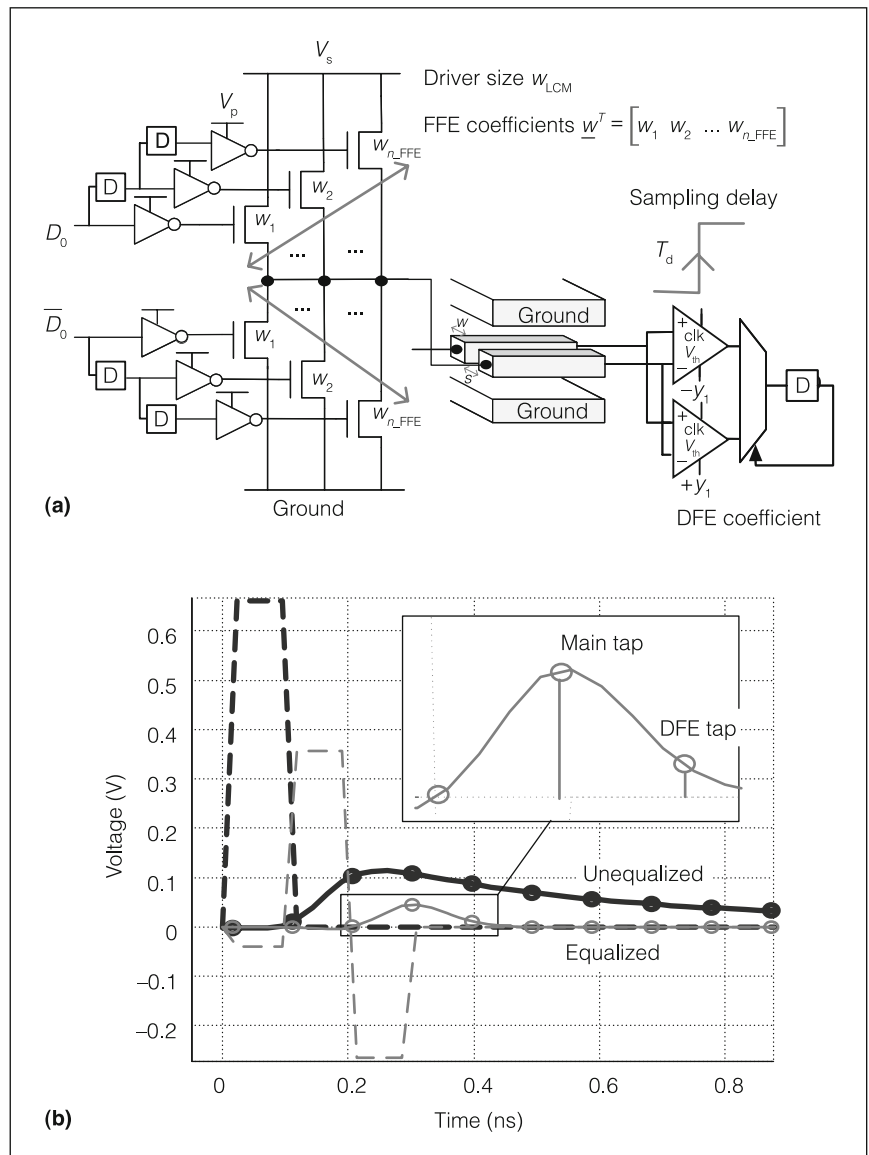


Figure 2. Example equalizer circuit over RC-dominant wires (a) and unequalized and equalized pulse responses (b). (DFE: decision feedback equalizer.) (We generated this figure from our earlier work.⁶)

strongly coupled with the channel's transfer function and the interconnect density. Driver size w_{LCM} affects both the channel transfer function and the driver's power consumption. FFE coefficients \underline{w}^T ($= w_1 w_2 \dots w_{n_FFE}$), DFE tap y_1 , and sampling timing T_d are the key parameters affecting the link eye opening. In addition to Equation 1, we add constraints on eye opening and target latency in Equation 6 to guarantee that the received signal is sampled at the proper time and that the eye is large enough to be sensed.

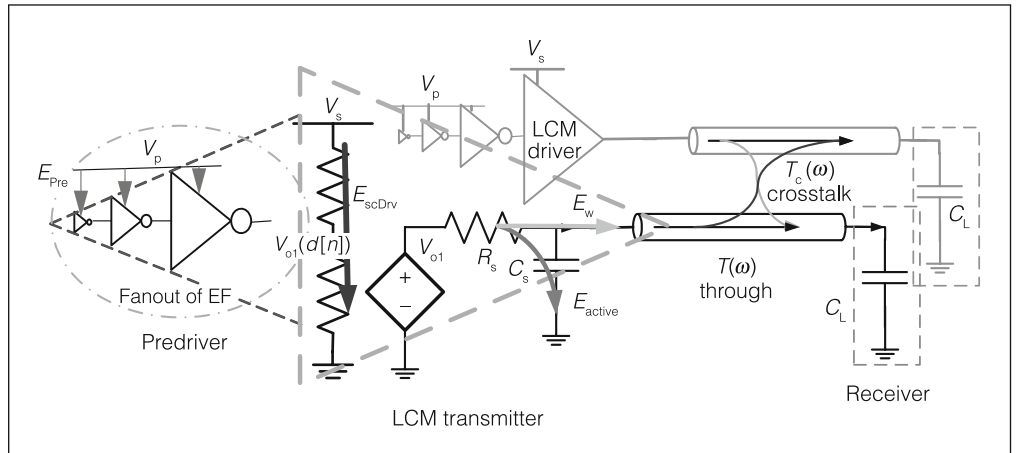


Figure 3. Thevenin equivalent model of equalized interconnect with power-consumption breakdown. (LCM: low common mode.) (We generated this figure from our earlier work.⁶)

Maximize d_d

subject to $T_d = \text{target latency}$

$$\begin{aligned} &= T_s/2 + (n_{\text{main}} - 1)T_s \\ &\quad - \frac{\mathcal{L}T(f_N)}{2\pi f_N} \end{aligned} \quad (6)$$

$E_b \leq \text{target energy cost}$

eye ≥ 50 mV

tion, similar to the logical effort model. We model the wire as a lossy transmission line:

$$\begin{bmatrix} z_o & z_c \\ z_c & z_o \end{bmatrix} = \mathbf{Z} = \mathbf{R} + j\omega\mathbf{L} \quad (7)$$

$$\begin{bmatrix} y_o & y_c \\ y_c & y_o \end{bmatrix} = \mathbf{Y} = \mathbf{G} + j\omega\mathbf{C} \quad (8)$$

The greatest challenge in equalized interconnects is modeling the link transmission quality (eye opening or bit error rate), performance, and power metrics, while jointly tuning the circuit and wire parameters. Our earlier work suggested an efficient modeling approach and verified the modeling quality (channel transfer function, eye diagram, and power model accuracy) using Spice simulation.⁶ The model provides fast design space exploration over millions of design points, letting the tool choose the best design within a practical design space.

Figure 3 shows a linearized circuit model of the links in Figure 2a. To model the crosstalk, we analyze two driver-receiver pairs simultaneously. The Thevenin equivalent voltage sources controlled by the programmable voltage divider model the LCM driver, whereas the far-end load capacitors model the receiver input capacitance. We derived driver parasitic resistance and capacitance from Spice models using lineariza-

where z_o and y_o are the wire's through impedance and admittance per wire length, z_c and y_c are the wire's crosstalk impedance and admittance per wire length; \mathbf{R} , \mathbf{L} , \mathbf{G} , and \mathbf{C} are 2×2 RLGC matrices (where \mathbf{G} is conductance) of wires;⁶ j is a complex number (square root of -1), and ω is the angular frequency. Equation 9 shows the link's through- and crosstalk-transfer functions.

$$T(\omega) \approx T_{\text{com}}(\omega) + T_{\text{diff}}(\omega) \quad (9)$$

$$T_c(\omega) \approx T_{\text{com}}(\omega) - T_{\text{diff}}(\omega)$$

where $T_{\text{com}}(\omega)$ and $T_{\text{diff}}(\omega)$ are the common and differential mode transfer functions between adjacent channels. With good accuracy, we derive the transfer function's closed-form solution in Equation 9 in terms of $T_{\text{com}}(\omega)$ and $T_{\text{diff}}(\omega)$, as defined in Equations 10 and 11:

$$T_{\text{com}}(\omega) = \frac{e^{-d\sqrt{(z_o + z_c)(y_o + y_c)}}}{\left(j\omega C_L \sqrt{\frac{(z_o + z_c)}{(y_o + y_c)}} + 1\right) \left(1 + R_s \left(\sqrt{\frac{(y_o + y_c)}{(z_o + z_c)}} + j\omega C_s\right)\right)} \quad (10)$$

$$T_{\text{diff}}(\omega) = \frac{e^{-d\sqrt{(z_o - z_c)(y_o - y_c)}}}{\left(j\omega C_L \sqrt{\frac{(z_o - z_c)}{(y_o - y_c)}} + 1\right) \left(1 + R_s \left(\sqrt{\frac{(y_o - y_c)}{(z_o - z_c)}} + j\omega C_s\right)\right)} \quad (11)$$

where d is the wire length, C_L is the receiver input capacitance, R_s is the driver parasitic resistance, and C_s is the driver parasitic capacitance.

By ignoring crosstalk terms ($z_c = y_c = 0$), we can further simplify the transfer function:

$$T(\omega) \approx \frac{2e^{-d\sqrt{j\omega R_o C_o}}}{\left(j\omega C_L \sqrt{\frac{R_o}{j\omega C_o}} + 1\right) \left(1 + R_s \left(\sqrt{\frac{j\omega C_o}{R_o}} + j\omega C_s\right)\right)} \quad (12)$$

Equation 12 tells us that the overall transfer function has exponential dependencies on wire length and operating frequency. In a typical link design, the signal amplitude at the receiver is far smaller than the supply rail because of the high loss in the channel. We therefore choose the sampling phase for the maximal eye opening. In a 1-tap DFE over an RC-dominant link, we can approximate the latency that maximizes the vertical eye opening as

$$T_d \approx T_s/2 + (n_{\text{main}} - 1)T_s - \frac{\angle T(f_N)}{2\pi f_N} \quad (13)$$

where T_s is the symbol period, and n_{main} is the index of the FFE's main tap (typically 0 or 1). $T_s/2$ is the time duration from the pulse edge to its center. The term

$$(n_{\text{main}} - 1)T_s$$

is the precursor flip-flop delay in the FFE, and

$$-\frac{\angle T(f_N)}{2\pi f_N}$$

is the propagation delay over the wire at the Nyquist frequency.⁶ The exponential term

$$e^{-d\sqrt{j\omega R_o C_o}}$$

in Equation 12 comes from the RC-dominant transmission line, and gives us imprecise but intuitive propagation velocity:

$$\frac{\sqrt{2\omega}}{\sqrt{R_o C_o}}$$

This propagation velocity strongly depends on the

frequency, and the value draws near to the speed of light at the link's Nyquist frequencies in the multi-GHz range, leading to low latency in Equation 13.

After we compute the optimal latency, we use the rail-constrained least-mean-square-error (LMSE) algorithm to get a fast closed-form solution for the FFE coefficients⁶:

$$\begin{aligned} \underline{w}_{\text{lmse}} &= \frac{(\mathbf{H}_{\text{isi}}^T \mathbf{H}_{\text{isi}} + \mathbf{H}_c^T \mathbf{H}_c)^{-1} \underline{h}_{\text{sig}}}{\underline{h}_{\text{sig}}^T (\mathbf{H}_{\text{isi}}^T \mathbf{H}_{\text{isi}} + \mathbf{H}_c^T \mathbf{H}_c)^{-1} \underline{h}_{\text{sig}}} \\ \tilde{\underline{w}}_{\text{lmse}} &= \frac{\underline{w}_{\text{lmse}}}{\|\underline{w}_{\text{lmse}}\|_1} \end{aligned} \quad (14)$$

where $\underline{w}_{\text{lmse}}$ is the vector of the FFE coefficients, obtained via the LMSE method; $\tilde{\underline{w}}_{\text{lmse}}$ is the normalized vector of the FFE coefficients, obtained by normalizing $\underline{w}_{\text{lmse}}$; \mathbf{H}_{isi} and \mathbf{H}_c are the ISI and crosstalk matrices; and $\underline{h}_{\text{sig}}$ is a vector computed from the sampled raw-pulse response.⁶ Under the assumption that ISI typically reduces the eye opening from 20% to 70% of the signal amplitude, we can derive the following equation to predict the eye opening from the channel's frequency response:

$$\begin{aligned} \text{eye} &\approx (1 - a_{\text{isi}}) \left(\frac{Y(0)}{T_s} + \frac{2}{T_s} |Y(f_N)| \right) \\ Y(f) &= T_s \text{sinc}(f T_s) e^{-j\pi f T_s} T(f) \underline{L}(f) \underline{w} \\ \underline{L}(f) &= \left[1 \quad e^{-j2\pi f T_s} \quad \dots \quad e^{-j\pi f T_s (n_{\text{FFE}} - 1)} \right] \end{aligned} \quad (15)$$

where a_{isi} is the nonnegative coefficient less than 1 (typically around 0.5), n_{FFE} is the number of FFE taps, f_N is the Nyquist frequency, $Y(f)$ is the frequency spectrum of the pulse at the receiver input, and $\underline{L}(f)$ is a vector of the transfer-function representation of flip-

flop delay in the FFE. Equation 15 indicates that the eye opening strongly depends on the voltage level at DC ($f = 0$) and at the Nyquist frequency of the transmitted pulse equalized by the FFE.

In the equalized interconnect, the driver draws most of the energy because a long wire is a large capacitive load. The receiver typically burns 20% to 30% of the link energy, according to our implementation estimates. We therefore focus mostly on modeling driver power dissipation. Figure 3 also shows the energy dissipation model of the equalized interconnect. We estimate the predriver energy E_{Pre} from the total inverter capacitance sized to drive the LCM driver at a desired data rate (represented by the effective fan-out factor EF):

$$E_{\text{Pre}} = \alpha \left(\frac{1/EF}{1 - 1/EF} \frac{C_{\text{gpre}} + C_{\text{dpre}}}{C_{\text{gpre}}} \right) \times C_{\text{gLCM}} w_{\text{LCM}} V_p^2 \quad (16)$$

where α is the activity factor, C_{gpre} and C_{dpre} are the gate and drain capacitances of the predriver per unit width, C_{gLCM} is the gate capacitance per unit width of the LCM driver, w_{LCM} is the driver size, and V_p is the predriver's supply voltage.

We compute the energy dissipated to charge the driver's parasitic capacitances E_{active} and wire load E_w from the spectrum of current and voltage in the frequency domain using Parseval's theorem and the transfer function from Equations 9 through 12:

$$E_{\text{active}} = \int_{-\infty}^{\infty} |V_o(f)|^2 \text{Re} \left\{ 1 / \left(R_s + \frac{1}{j2\pi f C_s + Y_w(f)} \right)^* \right\} df \quad (17)$$

$$E_w = \int_{-\infty}^{\infty} |V_o(f)|^2 \text{Re} \left\{ \left(\frac{Y_w(f)}{1 + (j2\pi f C_s + Y_w(f)) R_s} \right)^* \right\} df \quad (18)$$

where $\text{Re}\{ \}$ represents the real part of each respective complex number. The short-circuit energy E_{scDrv} dissipated by the current through the resistive divider in the equalizing driver is

$$\langle E_{\text{scDrv}} \rangle = \frac{V_s^2}{4R_s} \left(1 - \|w\|_2^2 \right) T_s \quad (19)$$

In power- and area-constrained design, the optimizer

tries to keep the eye opening constant (for example, at 50 mV) for signaling robustness, while adjusting other parameters to meet the optimization goal (Equation 6) over various data rate densities and thus extract the trade-off curves for link performance metrics (data rate density, energy per bit, and latency).

Unlike repeater interconnects, in which designers can trade latency for throughput and power by inserting more stages, the equalized interconnect's latency is tied to the signal phase delay at a given Nyquist rate (to maximize the eye opening). Hence, the equalized interconnects have a single latency value for a given energy per bit and data rate density.

Comparison of equalized and repeated interconnects

Using our modeling framework, we examine the trade-offs of equalized and repeated interconnects projected to the 32-nm technology node under the NoC scenarios we discussed earlier. Table 1 summarizes the interconnect scenarios and parameter setup for the 32-nm technology node.^{11,12}

Figure 4a shows the energy cost versus data rate density of equalized global and semiglobal wires (M9, M6). As wire length increases, the energy cost per bit increases dramatically, as Equation 12 predicted. On the other hand, the propagation delay decreases as data rate density increases, as Figure 4b shows. This is clear from Equation 13: First, as the data rate density increases, the link data rate generally also increases, so the delay term proportional to T_s decreases. Second, as we explained earlier, the channel propagation term

$$- \frac{\angle T(f_N)}{2\pi f_N}$$

decreases because of the phase delay's frequency dependency, as the data rate (and thus the Nyquist frequency) increases.⁶

Figure 5 compares the trade-off contours of repeated interconnects to the equalized interconnect trade-off curve with the energy cost label. Each contour of the repeater trade-off surface represents the equivalent energy level. The energy contours indicate the sensitivity of the trade-off between data rate density and latency for a given energy budget. In the M9 metal layer, the plots show that repeated interconnects cost $3\times$ to $10\times$ more energy than equalized interconnects for the same range of data rate densities but have at least $2\times$ higher latency.

Table 1. Simulation setup for interconnect trade-off generation.

Feature	Description	
Die size	20 mm × 20 mm	
NoC topologies	Mesh, Cmesh, flattened butterfly	
Wire length	5 mm, 10 mm, 15 mm	
Metal level	M6 semiglobal wire, M9 global wire	
Technology	32 nm, copper, low- <i>k</i> interconnect	
Parameters	M6 semiglobal wire	M9 global wire
Effective dielectric constant	2.3	
Interlayer dielectric thickness	400 nm	800 nm
Metal thickness	403 nm	576 nm

In the M6 metal layer, the equalized interconnect's energy efficiency gain is smaller than in the M9 metal layer because of poorer wire resistance and higher capacitance per unit length. Still, the equalized interconnect's latency benefit is about 2× better than that of the repeated interconnect in both layers. The high attenuation in the M6 layer doesn't significantly affect phase delay in the equalized interconnect but requires many repeater stages, resulting in poor delay, as Equation 4 predicted.

Table 2 shows the optimal wire pitch and width, and the total driver width of the equalized interconnect. It also shows the wire pitch and width of the repeated interconnect whose latency is closest to the equalized interconnect in the M6 and M9 metal layers. As the data rate density increases, the wire pitch decreases to increase the density, and the wire width increases to reduce the wire resistance. The table also summarizes the total driver width, which is the sum of each repeater width for the repeated interconnects. The repeater is significantly larger than the equalizer to overcome the inherent repeater delay.

**OUR DESIGN SPACE EXPLO-
RATION** of equalized and repeated interconnects projected to the 32-nm process node shows that equalized interconnects offer 2× to 10× better energy efficiency and at least 2× better latency for the same data

rate densities than repeated interconnects. The advantage of equalized interconnects over repeaters is larger in global layers than semiglobal layers, implying that reverse wire scaling will make equalized interconnects even more attractive in the future. The improved latency and energy efficiency of long (10 mm and 15 mm) equalized wires makes the long express lanes more affordable, further improving the efficiency of advanced NoC topologies such as Cmesh or flattened butterfly over a standard mesh topology. ■

Acknowledgments

We thank the Nanoscale Integration and Modeling Group at Arizona State University for the use of the Predictive Technology Model (<http://www.eas.asu.edu/~ptm>) to generate the 32-nm transistor models.

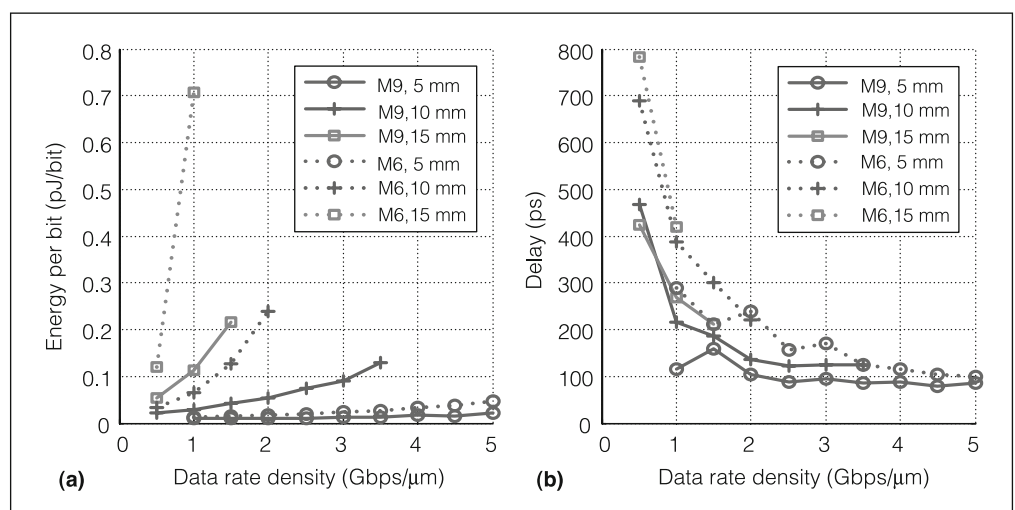


Figure 4. Trade-offs of equalized interconnect over M9 and M6 metal wire: energy versus data rate density (a) and latency versus data rate density (b). (M9: metal-layer level for global wire; M6: metal-layer level for semiglobal wire.)

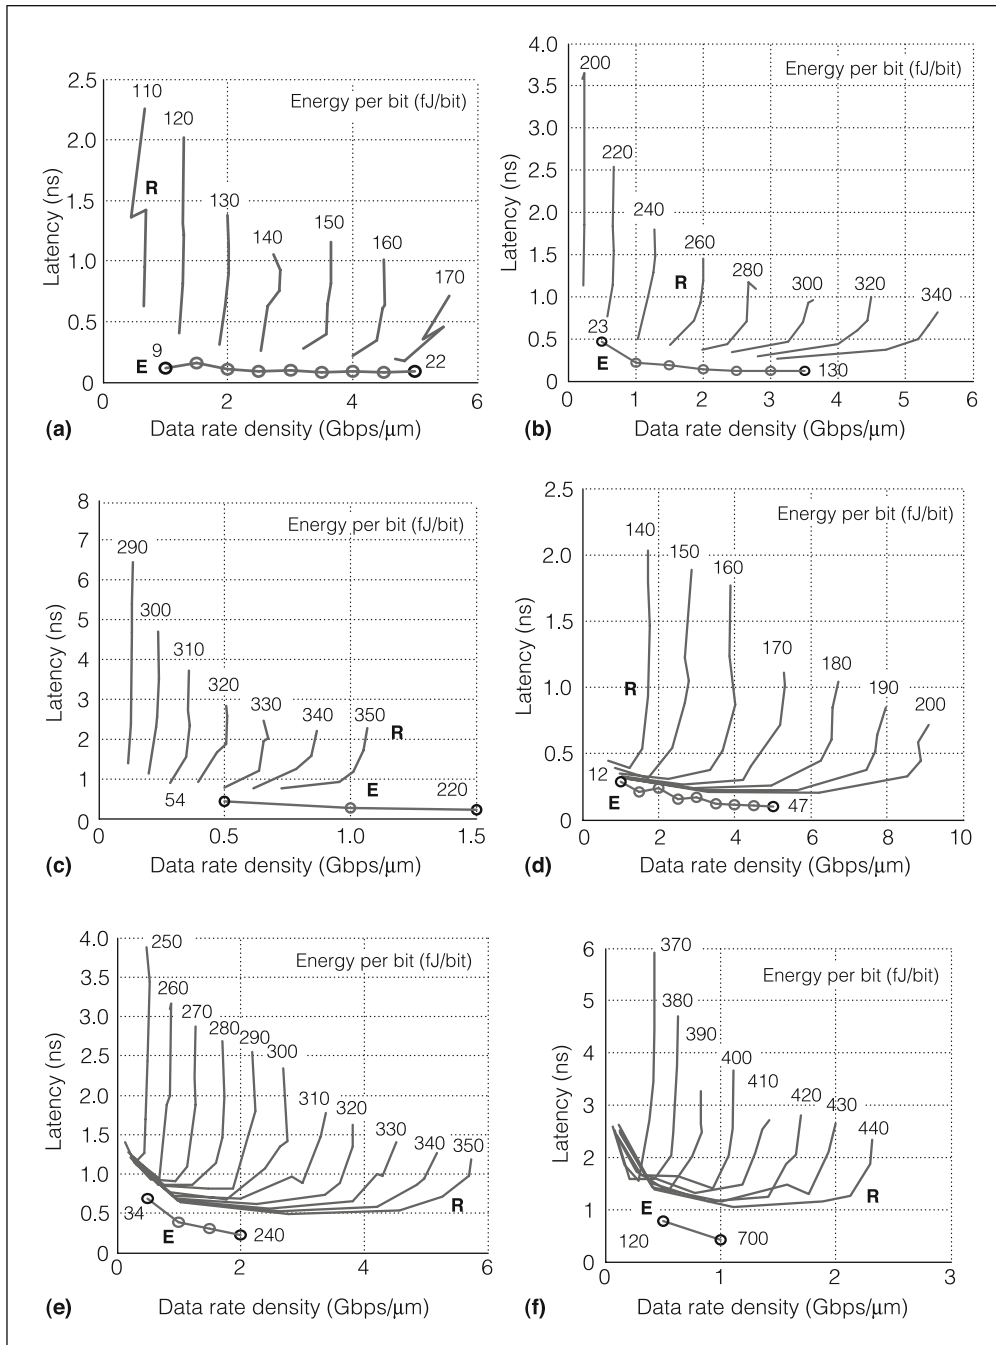


Figure 5. Trade-off contours of repeated and equalized interconnects over M9, 5 mm (a); M9, 10 mm (b); M9, 15 mm (c); M6, 5 mm (d); M6, 10 mm (e); and M6, 15 mm (f) wires. (E: equalized; R: repeated.)

■ References

1. M. Horowitz and W. Dally, "How Scaling Will Change Processor Architecture," *Proc. IEEE Int'l Solid-State Circuits Conf. (ISSCC 04)*, IEEE Press, 2004, pp. 132-133.

Computer Architecture Letters, vol. 6, no. 2, Feb. 2007, pp. 37-40.

8. H.B. Bakoglu and J.D. Meindl, "Optimal Interconnection Circuits for VLSI," *IEEE Trans. Electron Devices*, vol. 32, no. 5, May 1985, pp. 903-909.

2. S. Vangal et al., "An 80-Tile 1.28TFLOPS Network-on-Chip in 65 nm CMOS," *Proc. IEEE Int'l Solid-State Circuit Conf. (ISSCC 07)*, IEEE Press, 2007, pp. 98-99, 589.
3. R. Ho, "On-Chip Wires: Scaling and Efficiency," doctoral dissertation, Dept. of Electrical Engineering, Stanford Univ., 2003.
4. A.P. Jose, G. Patounakis, and K.L. Shepard, "Near Speed-of-Light On-Chip Interconnects Using Pulsed Current-Mode Signaling," *Proc. IEEE Symp. VLSI Circuits*, IEEE Press, 2005, pp. 108-111.
5. E. Mensink et al., "A 0.28pJ/b 2Gb/s/ch Transceiver in 90 nm CMOS for 10 mm On-Chip Interconnects," *Proc. IEEE Int'l Solid-State Circuits Conf. (ISSCC 07)*, IEEE Press, 2007, pp. 414-415, 612.
6. B. Kim and V. Stojanović, "Equalized Interconnects for On-Chip Networks: Modeling and Optimization Framework," *Proc. IEEE/ACM Int'l Conf. Computer-Aided Design (ICCAD 07)*, IEEE Press, 2007, pp. 552-559.

7. J. Kim, J. Balfour, and W.J. Dally, "Flattened Butterfly Topology for On-Chip Networks," *IEEE*

Table 2. Optimal wire pitch and width, and driver width of equalized (E) and repeated (R) interconnects over 5-mm, 10-mm, and 15-mm wire on M6 and M9 metal layers.

Length (mm)	Data rate density (Gbps/ μ m)	Wire pitch (μ m)				Wire space (μ m)				Total driver width (μ m)			
		M6		M9		M6		M9		M6		M9	
		R	E	R	E	R	E	R	E	R	E	R	E
5	1.0	0.73	2.20	1.15	6.0	0.23	0.34	0.32	0.48	3.72	0.15	2.40	0.3
	3.0	0.71	1.40	1.01	2.6	0.23	0.67	0.32	0.64	5.79	0.60	4.95	0.5
	5.0	0.78	1.80	1.37	1.8	0.23	0.67	0.46	0.48	5.79	2.00	31.40	1.2
10	1.0	0.76	1.96	1.25	3.5	0.23	0.70	0.32	1.00	8.71	0.90	5.31	0.6
	2.0	0.72	2.28	1.36	3.5	0.23	1.00	0.32	1.60	5.09	3.00	7.07	2.4
15	0.5	0.77	1.96	2.56	3.5	0.23	0.70	0.32	1.00	19.1	0.60	4.22	0.6
	1.0	0.87	2.28	2.49	3.5	0.23	0.70	0.32	1.30	5.32	5.00	6.39	0.9

9. V. Deodhar and J.A. Davis, "Designing for Signal Integrity in Wave-Pipelined SoC Global Interconnects," *Proc. IEEE Int'l SOC Conf.*, IEEE Press, 2005, pp. 207-210.
10. H. Hatamkhani et al., "A 10-mW 3.6-Gbps I/O Transmitter," *Proc. IEEE Symp. VLSI Circuits*, IEEE Press, 2003, pp. 97-99.
11. W. Zhao and Y. Cao, "New Generation of Predictive Technology Model for Sub-45nm Design Exploration," *Proc. IEEE Int'l Symp. Quality Electron Design (ISQED 06)*, IEEE CS Press, 2006, pp. 585-590.
12. K. Mistry et al., "A 45 nm Logic Technology with High-k+ Metal Gate Transistors, Strained Silicon, 9 Cu Interconnect Layers, 193 nm Dry Patterning, and 100% Pb-Free Packaging," *Proc. IEEE Int'l Electronic Devices Meeting (IEDM 07)*, IEEE Press, 2007, pp. 247-250.

Byungsub Kim is a PhD student in the Electrical Engineering and Computer Science Department at the Massachusetts Institute of Technology. His research interests include modeling and design of on-chip signaling. He has a BS in electronic and electrical engineering from Pohang University of Science and Technology, Pohang, Korea, and an MS in electrical engineering and computer science

from the Massachusetts Institute of Technology. He is a student member of the IEEE.

Vladimir Stojanović is an assistant professor in the Electrical Engineering and Computer Science Department at the Massachusetts Institute of Technology. His research interests include optimization of analog and VLSI circuits, modeling of noise and dynamics in circuits and systems, communications and signal-processing architectures, high-speed electrical and optical links, on-chip signaling, and high-speed digital and mixed-signal IC design. He has a Dipl Ing in electrical engineering from the University of Belgrade, Serbia, and a PhD in electrical engineering from Stanford University.

■ Direct questions and comments about this article to Byungsub Kim, Electrical Engineering and Computer Science Dept., Rm. 38-266c, Massachusetts Institute of Technology, 77 Massachusetts Ave., Cambridge, MA 02139; byungsub@mit.edu.

For further information about this or any other computing topic, please visit our Digital Library at <http://www.computer.org/csdl>.