

Low-Swing On-Chip Signaling Techniques: Effectiveness and Robustness

Hui Zhang, *Student Member, IEEE*, Varghese George, *Student Member, IEEE*, and Jan M. Rabaey, *Fellow, IEEE*

Abstract—This paper reviews a number of low-swing on-chip interconnect schemes and presents a thorough analysis of their effectiveness and limitations, especially on energy efficiency and signal integrity. In addition, several new interface circuits presenting even more energy savings and better reliability are proposed. Some of these circuits not only reduce the interconnect swing, but also use very low supply voltages so as to obtain quadratic energy savings. The performance of each of the presented circuits is thoroughly examined using simulation on a benchmark interconnect circuit. Significant energy savings up to a factor of six have been observed.

Index Terms—Digital CMOS, low-power design, low-voltage, performance tradeoffs, reliability, special low-power 99.

I. INTRODUCTION

IN THE deep-submicron era, interconnect wires (and the associated driver and receiver circuits) are responsible for an ever increasing fraction of the energy consumption of an integrated circuit. Most of this increase is due to global wires, such as busses, clock, and timing signals. For gate array and cell-library-based designs, D. Liu *et al.* [1] found that the power consumption of wires and clock signals can be up to 40% and 50% of the total on-chip power consumption, respectively. The impact of interconnect is even more significant for reconfigurable circuits. Measured over a wide range of applications, more than 90% of the power dissipation of traditional FPGA devices have been reported to be due to the interconnect [2].

Obviously, techniques that can help to reduce these ratios are desirable. For chip-to-chip interconnects, wires are treated as transmission lines, and many low-power I/O schemes were proposed at both circuit level (e.g., GTL transceiver [3]) and coding level (e.g., work-zone encoding [4] and bus-invert coding [5]). In this paper, the main focus is how to reduce the power consumption of on-chip interconnects. Short of reducing the average length of the wires and their fanout by using advanced processes or improved architectures, reducing the voltage swing of the signal on the wire is one of the best solutions toward getting better energy efficiency. First, we will analyze the effectiveness of a number of reduced-swing interconnect schemes that have been proposed in the literature [6]–[11]. In addition, a number of novel or modified circuits will be introduced, simulated, and critiqued. To present a fair and realistic base for comparison, a single test circuit will be used. Overall, it is found that the proposed schemes present a wide range of potential energy re-

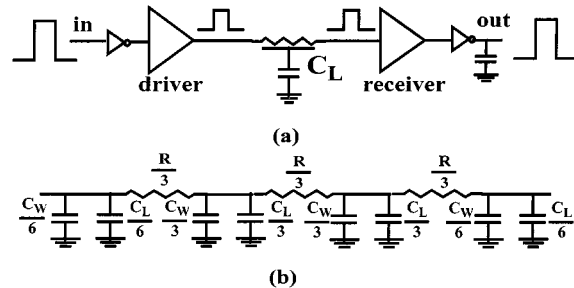


Fig. 1. (a) Benchmark test architecture. (b) Interconnect model.

ductions, yet that other considerations such as complexity, reliability, and performance play an important role as well. We will therefore pay special attention to each of these factors in our analysis.

The paper is organized as follows. First, the benchmark example and the set of quality metrics that will be used in all simulations and comparisons are presented. What follows are a review and comparison of a number of architectures, obtained from the open literature. Several novel or improved low-swing schemes are proposed and analyzed in Section III. Finally, Section IV brings them all together and draws some conclusions. At the end of the paper, an Appendix is attached to provide detailed descriptions for the physical models of important noise sources.

II. TEST ARCHITECTURE AND QUALITY METRICS

Presenting a fair comparison for the various interconnect schemes that are presented in this paper requires a common and fair testbed. Fig. 1(a) illustrates the schematic of our benchmark interconnect circuit. The driver converts a full-swing input into a reduced-swing interconnect signal, which is converted back to a full-swing output by the receiver. The interconnect line is a metal-3 layer wire with a length of 10 mm, modeled by a π 3 distributed RC model with an extra capacitive load C_L distributed along the wire (for fanout), as shown in Fig. 1(b). To fairly compare the delays of the different schemes, we deliberately add an inverter prior to the driver and an inverter after the receiver with 20-fF capacitive load. Both inverters are sized with $W_p = 6 \mu\text{m}$ and $W_n = 3 \mu\text{m}$. All circuit comparisons are based on the MOSIS HP complementary metal-oxide-semiconductor (CMOS) 14TB process parameters and spice models. The minimum drawn channel length for this process is set to $0.6 \mu\text{m}$ with an effective channel length of $0.5 \mu\text{m}$.

For each of the circuits under test, we consider the following metrics.

Manuscript received February 18, 1999; revised August 31, 1999. This work was supported by DARPA under the ACS PLEIADES project.

The authors are with the Berkeley Wireless Research Center, EECS Department, University of California, Berkeley, CA 94704 USA (e-mail: hui@eecs.berkeley.edu).

Publisher Item Identifier S 1063-8210(00)04349-3.

TABLE I
 TYPICAL NOISE SOURCES

K_N	K_C	crosstalk coupling coefficient: $K_C = 0.4$ for 10 mm wires with 1pF load and 2 μm spacing.
	Am_C	crosstalk noise attenuation: $Am_C = 1$ for dynamic driver; $Am_C = 0.2$ for static driver.
	K_{PS}	power supply noise due to signal switching: $K_{PS} = 0.05$ for single-ended signaling; $K_{PS} = 0.01$ for differential signaling.
	Worst case: $K_N = Am_C K_C + K_{PS}$	
V_{IN}	Rx_O	receiver input offset: 150 mV for inverter
	Rx_S	receiver sensitivity: 150 mV for inverter
	PS	unrelated power supply noise: 5% of VDD
	Am_{PS}	power supply noise attenuation
	Tx_O	transmitter offset
	Worst case: $V_{IN} = (Rx_O) + (Rx_S) + Am_{PS} PS + (Tx_O)$	

- **Energy:** The dynamic switching energy of the wire for a full switching is given by (1). When comparing schemes with different types of circuit design such as dynamic design versus static design, differences in data activity are taken into account. The short-circuit current and leakage current are relatively less important compared to the dominant switching energy, but will be also under consideration. The total energy shall include the contributions from both the driver and receiver

$$E_{\text{dyn}} = (C_W + C_L) \bullet V_{\text{DD}}(\text{driver}) \bullet V_{\text{swing}}. \quad (1)$$

- **Design complexity.**
- **Delay.**
- **Reliability:** Three main sources of reliability degradation are considered: process variation, voltage supply noise, and interline crosstalk.

We use the worst case noise analysis method presented in [12] to measure the reliability of each circuit. The noise sources are classified into two categories: the proportional noise sources and the independent noise sources

$$V_N = K_N V_S + V_{IN}. \quad (2)$$

$K_N V_S$ represents those noise sources that are proportional to the magnitude of signal swing (V_S), such as crosstalk, and signal-induced power supply noise. V_{IN} includes those noise sources that are independent of V_S such as receiver input offset (due to process variation), receiver sensitivity, and signal-unrelated power supply noise. Table I summarizes the noise sources and their contributions, and detailed descriptions are provided in Section VI (Appendix).

The cross-talk coupling coefficient K_C is derived from the ratio between coupling capacitance and wire load capacitance. The cross-talk noise attenuation for the static driver scenario is achieved by increasing the timing budget for the signal so that the charge loss due to the cross-talk noise can be recovered by the driver. The signal-induced supply noise is estimated to be 5% and 1% of the signal swing for single-ended and differential signaling, respectively. The receiver input offset and sensitivity are dependent on the receiver circuits in question, and will be




Fig. 2. Conventional level converter.

individually assessed for each scheme (e.g., for the CMOS inverter, its input offset and sensitivity are around 150 mV, respectively). The signal-unrelated power supply noise is assumed to be 5% of the magnitude of power supply for a well-designed power distribution network.

The power-supply attenuation coefficient is defined as the change of the switching threshold voltage induced by an unit change of the supply voltage. The transmitter offset results from the parameter mismatch between the transmitter and receiver, such as threshold voltage mismatch and reference voltage variation.

We use the worst case signal-to-noise ratio (SNR) defined in (3) as a measure of the reliability of each circuit. The noise margin is defined as (SNR 1)

$$\text{SNR} = \frac{0.5V_S}{V_N}. \quad (3)$$

III. REVIEW OF EXISTING LOW-SWING INTERFACE CIRCUITS

In this section, seven low-swing circuit schemes (three static and four dynamic) are reviewed, and the pros and cons of each approach are enumerated. The important design metrics of the circuits are compared based on simulation results.

A. Static Driver with Reduced Supply

The conventional level converter (CLC) shown in Fig. 2 represents the traditional way of converting a low-swing signal back to a full swing one. The driver uses an extra low-voltage supply to drive the interconnect from zero to V_{DD_L} . Although the noise margin is reduced, this circuit is very robust against noise, as the receiver behaves as a differential amplifier, and the internal inverter further attenuates some noise through regeneration. The symmetric driver and level converter (SDLC), proposed in [7], also falls in the same category. It requires two extra power rails to limit the interconnect swing and uses special low- V_t devices (~ 0.1 V) to compensate for the current-drive loss due to the lower supplies.

B. Differential Interconnect (DIFF)

Differential signaling is more immune to noise due to its high common-mode rejection, allowing for a further reduction in the signal swing. Fig. 3 shows a circuit, which is fully analyzed in [14], achieving great energy savings by using a very low voltage supply. The driver uses NMOS transistors for both pull-up and pull-down. The receiver is a clocked unbalanced current-latch sense amplifier, which is discharged and charged at every clock cycle. The receiver overhead may hence be dominant for short




Fig. 3. Differential low-swing interconnect.




Fig. 4. Pulse-controlled driver with sense amplifier.

interconnect wires with small capacitive load. Due to its differential nature, the sense amplifier has a very low power supply noise attenuation coefficient (0.2 from simulation results). Its input offset is determined by the local device mismatch between the two input transistors and is as small as 20 mV. The main disadvantage of the differential approach is the doubling of the number of wires, which certainly presents a major concern in most designs. The extra clock signal further adds to the overhead.

C. Dynamically Enabled Drivers

The idea behind this family of circuits is to control the (dis)charging time of the drivers so that a desired swing on the interconnect is obtained. The pulsed-controlled driver (PCD) shown in Fig. 4 is a typical member of this family. The advantage of this circuit is that the pulse width can be finetuned to realize a very low swing while no extra voltage supply is needed. This concept has been widely applied in memory designs. However, it only works well in the cases when the capacitive loads are well known beforehand. Furthermore, the wire is floating when the driver is disabled, making it susceptible to noise. Another scheme (called RSD_VST, proposed in [10]) also uses a dynamically enabled driver, but with an internally generated EN signal. The driver uses an embedded copy of the receiver circuit (called *voltage-sense translator* or VST) to sense the interconnect swing so as to provide a feedback signal to control the driver. This circuit has a potential problem due to long wire delay—before the input of the receiver reaches the right level to switch the receiver, the driver might already be disabled. Mismatch of the switching voltage threshold between the two VST's, and supply noise can cause similar problems.

D. Low-Swing Bus

The *charge intershared bus* (CISB) [8] and *charge-recycling bus* (CRB) [9] are two schemes that reduce the interconnect swing by utilizing charge sharing between multiple data bit lines

TABLE II
PERFORMANCE COMPARISON OF EXISTING SCHEMES ($V_{\text{dcl}} = 2\text{ V}$, $C_L = 1\text{ PF}$)

Schemes	Energy (PJ)	Delay (ns)	E·D (PJ·ns)	Swing (V)	V_N (V)	SNR	Complexity
CMOS	11.6	2.1	24.5	2.0	0.66	1.52	least
CLC	4.4	3.1	13.6	1.1	0.443	1.24	1 REF
SDLC	3.5	3.1	10.9	0.8	0.373	1.08	low-V _t devices, 2 REFs
DIFF	3.0	2.7	8.1	0.25	0.076	1.64	extra timing, 1 REF, wires doubled
PCD	3.5	2.0	7.0	0.5	0.355	0.70	extra controls
RSD_VST	3.7	2.0	7.4	0.6	0.525	0.57	1 REF, big driver
CISB	3.5	4.4	15.4	0.25	0.19	0.66	extra timing, sense amplifiers
CRB	3.1	3.5	10.9	0.25	0.168	0.74	extra timing, wires doubled

of a bus. The CRB scheme uses differential signaling while the CISB scheme is single ended with references. Both schemes reduce the interconnect swing by a factor of n (where n is the number of bits). The CRB scheme presents quadratic power savings (by a factor of n^2) due to its charge-recycling mechanism, although the potential savings are offset by the fact that the bus is discharged and charged for every cycle (i.e., 100% switching activity). Both of their receivers use clocked current-latch sense amplifiers and require multiple timing signals. One stringent requirement for these bus schemes to work reliably is that all the wire capacitances must be matched very well, which is certainly nontrivial in real system designs. In both schemes, but especially in CRB, noise immunity is compromised by the floating nature of the interconnects between different evaluation cycles.

E. Simulation Results and Comparison

Each of the above presented circuits is optimized individually against the benchmark test architecture. Their important metrics and simulation results are tabulated in Table II. The CMOS scheme in the first row represents the full swing case (assuming a 2-V supply). Most of the low-swing schemes can achieve energy savings with a factor of around three, but only few of them have good reliability. The schemes with static drivers have SNR's larger than one, while the dynamic ones have SNR's less than one, which implies negative noise margin. Differential interconnect has the best SNR even with a very small swing of 0.25 V. It achieves energy savings with a factor of close to four, but requires a dual-wire structure. CLC is robust but can only reduce energy by 60% with respect to the original circuit at the expense of a bigger delay and an extra lower voltage supply. The SDLC scheme can reduce the energy by 70%, with low- V_t devices and two reference voltages. The CISB and CRB schemes are only suitable for multiple-bit bus units with large capacitive load. Simulation results predict energy savings of up to 3.5 times. Both of them are slow compared to the other schemes due to the charge-sharing mechanism. Their SNR's are much lower than one due to the floating interconnect. The RSD_VST scheme is susceptible to device mismatch and has the worst SNR. To improve SNR's of dynamic schemes, the cross-talk noise should be minimized (e.g., by wider wire spacing). Overall, existing schemes either

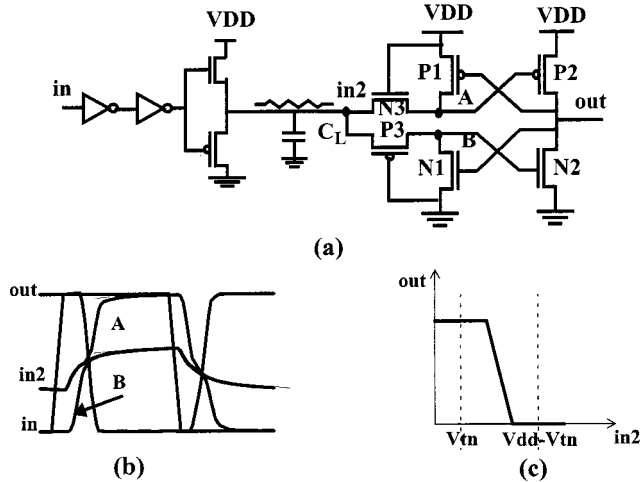


Fig. 5. (a) Symmetric source-follower driver with level converter. (b) Simulated waveforms. (c) Voltage transform curve.

are short of significant energy savings with good reliability, or introduce lots of overhead (e.g., dual wires per bit).

IV. PROPOSED INTERFACE CIRCUITS

We now present several improved or novel low-swing interconnect interface circuits to address some of the problems encountered in the earlier schemes.

- **Reliability:** Only static drivers should be used to avoid floating interconnect, especially for long wires. To reduce the independent noise sources, the receiver must have small input offset, good sensitivity, as well as high common-mode noise rejection.
- **Energy:** Static drivers are also preferred because they will result in lower signal switching activity. The supply voltage of the driver should be as low as possible (while still ensuring reasonable noise margin). The key challenge is how to detect a “one” signal at the receiver end.
- **Complexity:** Although the extra power supplies can be realized on-chip with power efficiencies around 90% [13], it is desirable to keep their number to a minimum. Since wire area is also a major concern in most chip designs, only single-ended signaling schemes will be considered.

In this section, six schemes are presented. The first two try to avoid any extra reference supplies to minimize the complexity, while still getting a decent amount of energy savings. The rest four schemes use very low supply voltages to further reduce the signal swing. The last two schemes also need additional timing signals.

A. Static Source-Follower Driver

Without extra reference supplies, a natural way to limit the interconnect signal swing is to utilize the threshold voltage drop of source followers. Two circuits based on this concept are introduced.

1) *Symmetric Source-Follower Driver with Level Converter (SSDLC):* The SSDLC scheme is shown in Fig. 5(a). The driver limits the interconnect swing from $|V_{tp}|$ to $V_{dd}-V_{tn}$, as shown in Fig. 5(b). The symmetric-level converter/receiver is similar




Fig. 6. Asymmetric source-follower driver with level converter.

to the one in *SDLC* circuit, except that the gates of the two pass transistors N3 and P3 are biased at V_{dd} and Ground, respectively. Moreover, no special low V_t devices are needed. Assume that node *in2* goes from low to high: V_{tn} to $V_{dd}-V_{tn}$. Initially, node A and B sit at V_{tn} and Ground, respectively. During the transition period, with both N3 and P3 conducting, A and B rise to $V_{dd}-V_{tn}$ as shown in Fig. 5(b). Consequently, N2 is turned on, and *out* goes to low. The feedback transistor P1 pulls A further up to V_{dd} to cut off P2 completely. *in2* and B stay at $V_{dd}-V_{tn}$. Note that there is no standby current path from V_{dd} to Ground through N3 although the gate-source voltage of N3 is nearly V_{tn} . Since the circuit is symmetric, the same explanation can be applied for the high-to-low transition. Ignoring the feedback transistors P1 and N1, the dc voltage transform curve of the level converter (Fig. 5c) is virtually a “compressed” version of the one of the P2-N2 pair. Since transistors P1 and N1 are mainly to provide positive feedback to completely cut off P2 or N2, they can be very weak to minimize their fight against the driver. The sensing delay of the receiver is as small as two inverter delays. The predicted interconnect energy-savings ratio is given in


$$\frac{E_{\text{new}}}{E_{\text{full}}} = \frac{V_{dd} - 2V_{tn}(\text{body})}{V_{dd}} \quad (4)$$

where the threshold voltage is subject to the body effect (and equals 1 V for the targeted technology). To have a reasonable swing on the interconnect, this scheme requires a relatively large V_{dd} (>2.8 V in this case).

2) *Asymmetric Source-Follower Driver with Level Converter (ASDLC):* An asymmetric version of the *SSDLC* scheme is shown in Fig. 6, enabling operation for a V_{dd} around 2 V. The driver swings the wire from REF_L to $V_{dd}-V_{tn}$. The internal voltage supply REF_L is set below V_{tn} of N2. The receiver is a variation of the voltage sense translator and is actually an asymmetric version of the level converter in the *SSDLC* scheme. Their operation is similar for the low-to-high transition. In case of the high-to-low transition, N2 turns off after A and B are discharged to a voltage level below V_t of transistor N2, and P2 pulls out up to V_{dd} . Transistors P2 and N2 are sized wide enough to have large transconductances in order to quickly sense the small V_{gs} applied on them. The feedback transistor N3 provides extra current drive to discharge the output. The following energy-savings ratio is obtained:

$$\frac{E_{\text{new}}}{E_{\text{full}}} = \frac{V_{dd} - V_{tn} - REF_L}{V_{dd}} \quad (5)$$

REF_L can be between zero to V_t (0.7 V for the targeted technology), and it is set at 0.2 V in our simulations to make sure the leakage current of N2 is negligible. Compared to *RSD_VST*

Fig. 7. Level converter with low- V_t devices.

scheme, ASDLC is more robust because of the static nature of the driver.

B. NMOS-Only Push-Pull Driver with Low-Power Supply

The previous two schemes only get linear energy reductions, as their drivers still use the regular power supply. To further reduce the interconnect energy consumption, NMOS-only push-pull drivers (as shown Fig. 7) with very low power supply are used. In the following, four different receiver techniques are proposed to effectively detect the low-swing signal. The expected ratio of the interconnect energy savings is given by

$$\frac{E_{\text{new}}}{E_{\text{full}}} = \left(\frac{\text{REF}}{V_{\text{dd}}} \right)^2. \quad (6)$$

1) *Level Converter with Low- V_t Device (LCLVD)*: Fig. 7 shows the schematic diagram of the LCLVD scheme. In this scheme, the receiver is the same as the conventional level converter, except that it uses low- V_t devices for N1, N2, and the internal inverter. Because inb is slower than $in2$, the two branches are designed asymmetrically to balance the switching delays in different directions, say, N2 is sized larger than N1 and P1 is larger than P2.

In our simulation, REF is set at 0.7 V, and V_{tn} and $|V_{tp}|$ of the low- V_t devices are set at 0.3 V. Simulation at the process corners proves that this circuit can operate reliably against supply noise and process variations. The receiver behaves like a differential sense amplifier by regenerating a complementary input signal internally. The increase of leakage currents of those low- V_t devices is negligible compared to the dominant wire switching power since they are sized much smaller than the driver.

2) *Capacitive-Coupled Level Converter (CCLC)*: Without using low- V_t devices, the high end of a signal can barely turn on an NMOS and turn off a PMOS. In the CCLC scheme, shown in Fig. 8(a), a coupling capacitor is used to boost the low-swing signal so that the NMOS transistor of the receiver can be turned on. Shown in the waveforms in Fig. 8(b), the input to N3 (node A) has a swing from Ground to REF, while the input to P3 (node B) has a swing from REF2 to V_{dd} , where REF2 is set to be less than $(\text{REF} + V_{tn})$. Its operation is explained as follows: When A switches from high to low, pass transistor N3 is turned on, hence pulling node C to Ground. Out is pulled up to V_{dd} with transistor N2 turned off and P2 on. With pass transistor P4 conducting, B is set to REF2. Since the gate-source voltage across P3 is less than its threshold voltage, P3 is not conducting, and therefore no static current path exists. When A goes from low to high, the coupling capacitor C_c couples the voltage step onto B. Meanwhile, pass transistor N3 is turned off, and C rises up by

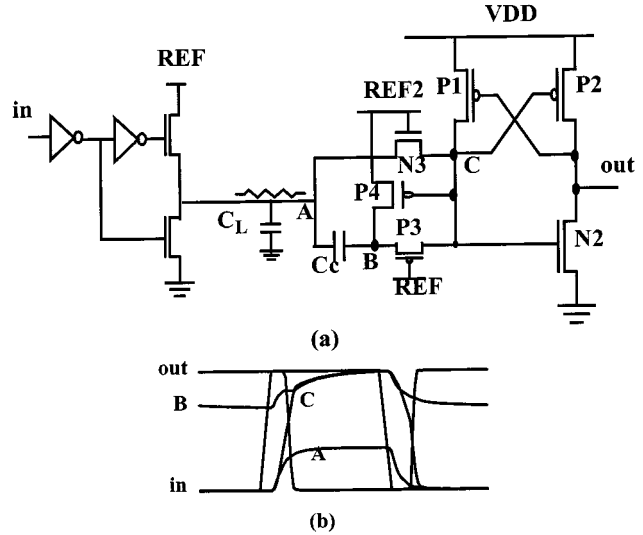


Fig. 8. (a) Capacitive-coupled level converter. (b) Simulated waveforms.

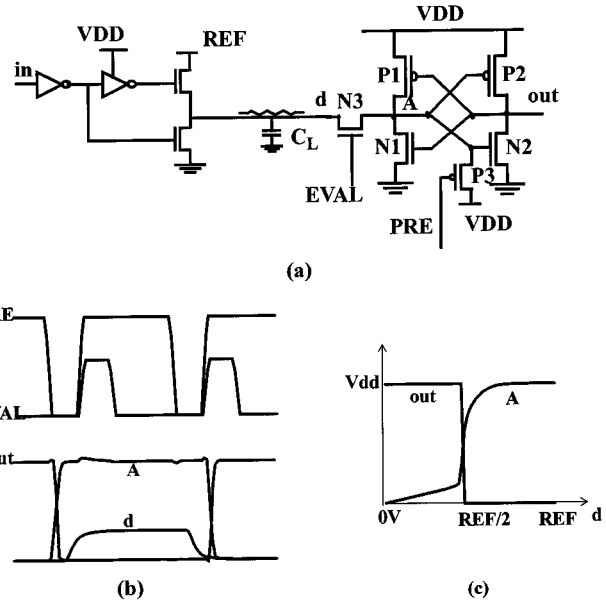


Fig. 9. (a) Level-converting register. (b) Simulated waveforms. (c) Voltage transform curve.

charge sharing with B through P3, as shown in Fig. 8(b). With Out being pulled low by N2, P1 pulls C and B further up to V_{dd} .

In the simulations, REF and REF2 are set as 0.8 and 1.2 V, respectively. The coupling capacitor C_c has to be big enough (0.2 pF in our simulations) to provide enough coupling effect in the presence of charge sharing between C_c and parasitic capacitances. Nonetheless, the operation of this circuit is not too sensitive to variations in C_c . Overall, this receiver can bootstrap a very low swing signal to a full one without special low- V_t devices and timing signals, but on the other hand, it suffers from a relatively small noise margin due to its susceptibility to the device variations.

3) *Level-Converting Register (LCR)*: In the next two schemes, extra timing signals are provided to help the receivers to detect the low-swing signal more effectively. Fig. 9(a) shows

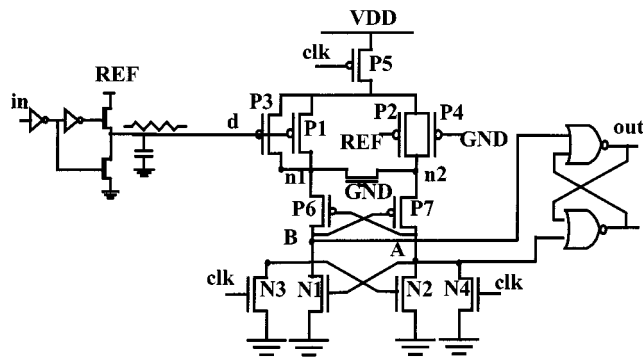


Fig. 10. Pseudodifferential interconnect.

the circuit diagram of the LCR scheme. The receiver consists of a cross-coupled inverter pair, with one precharge transistor P3 and one pass transistor N3, whose gates are controlled by two timing signals: PRE and EVAL, respectively. Typical waveforms are shown in Fig. 9(b). Initially, a negative pulse PRE is applied to P3 to precharge node *A* to V_{dd} and discharge node out to Ground. After the signal at node *d* stabilizes, a positive pulse EVAL is applied to N3. The high value of the voltage swing of EVAL is set to be less than $REF + V_{tn}(N3)$. If *d* is high, N3 stays off, and the state of the inverter pair remains the same. In the case of *d* being low, N3 starts conducting, and pulls *A* low, hence flipping the state of the inverter pair. After EVAL switches back to low, N3 is cut off, and the inverter pair keeps the data as a static register. The receiver is level sensitive. Consequently, the inverter pair will switch its state by a high to low glitch on the interconnect when EVAL is active. This cannot be remedied by returning the input to high. Therefore, the EVAL pulse has to be made as narrow as possible to avoid such an error. Fig. 9(c) illustrates the dc voltage transform curves of the receiver, when the gate voltages of the feedback transistors P1 and N1 are set to Ground.

A major advantage of this simple receiver is that it combines the functions of a level converter and a register. It has little area overhead, although the extra timing signals increase its complexity. The matching of the current drive capabilities of the P1-N3 pair is critical to the receiver's noise margin, which is susceptible to supply noise and V_t variations. Nevertheless, the receiver is fast and reliable as long as EVAL is applied after the input of the receiver reaches stable point. This circuit can be used for both synchronous and asynchronous signaling, assuming that the timing signals PRE and EVAL are generated correctly.

4) *Pseudodifferential Interconnect (PDIFF)*: Finally, we present a PDIFF scheme. Fig. 10 shows the circuit diagram of the PDIFF scheme. The receiver is a clocked sense amplifier followed by a static flip-flop. It has double pairs of input transistors, with the gates of P1 and P3 being connected to *d*, while the gates of P4 and P2 being biased at Ground and REF, respectively. Initially, *A* and *B* are discharged to Ground, and *n1* and *n2* are equalized. After *d* reaches the desired level, the receiver is enabled by a negative pulse of *clk*. If *d* is low, the current drive of P3 is same as that of P4, while the current drive of P1 is larger than that of P2. As a result, *B* is pulled high and *A* is kept low by the cross-coupled inverter pair

(N1-N2-P6-P7). An opposite transition is triggered when *d* is high. The following static flip-flop will retain the data value even after the sense amplifier is initialized again.

PDIFF scheme only uses single wire per bit while still retaining most advantages of differential amplifier such as low-input offset and good sensitivity. This is because its major reliability degradation comes from the local device mismatch between the double input transistor pairs, which usually can be controlled very well. The variation between distant REF's of the driver and the receiver also contributes some reliability degradation. The operation of the receiver is not sensitive to the VDD supply noise, as opposed to other schemes.

C. Simulation Results and Comparison

The six proposed circuits are optimized individually against the testing benchmark to get a fair comparison. The performances of them along with the full swing case are tabulated in Table III for the parameter settings of $V_{dd} = 2$ V, $C_L = 1$ pF (with the exception of SSDLC where V_{dd} is set to 2.8 V). Their total delay numbers are in a similar range. The low-swing receivers have longer delays than the simple inverter, and introduce more short-circuit power. These are dominated by the big savings from reducing the swing on the wire though. As shown in the results, the ASDLC scheme can reduce the energy consumption with 55% (same ratio for SSDLC scheme if scaled down to the same supply voltage), while with very little complexity overhead. LCLVD can achieve energy savings by a factor of almost five, with the help of low- V_t devices. CCLC can reduce the energy by a factor of more than four, with two extra reference supplies and a large coupling capacitor. LCR has a very simple receiver and can achieve the same energy savings as the LCLVD scheme, but requires two reference supplies and additional timing signals. PDIFF operates with the lowest signal swing at 0.5 V, which results in an energy reduction by a factor of six.

Noise analysis is performed for each of the schemes. Because every scheme uses static single-ended signaling, the total proportional noise coefficient K_N can be derived as 0.13 from Table I. The receiver input offset is assessed for each scheme by conducting dc voltage transform curve (VTC) simulations on different process corners. The receiver input sensitivity is also derived from VTC curves. Signal-unrelated power supply noise is assumed to be 5% of the supply magnitude. The power supply attenuation coefficients are derived from VTC curves at different supply voltages. The transmitter offset results from either the V_t variation at the driver side (for SSDLC and ASDLC) or the reference supply noise (assumed to be 5% of the reference magnitude) for the rest schemes. Table IV summarizes the noise sources for every scheme and shows the signal-to-noise-ratio numbers. From the results, it can be seen that all the schemes with the exception of CCLC have an SNR larger than one. PDIFF presents a SNR even higher than that of the full-swing case and has a noise margin of 92%. LCR and LCLVD have noise margins around 20%, while both SSDLC and ASDLC have 8%. The important observation is that, for low-swing signaling, independent noise sources play a dominant role. Therefore, to enhance the signal integrity, well-thoughtout power distribution schemes, device matching,

TABLE III
PERFORMANCE COMPARISON OF PROPOSED SCHEMES ($V_{dd} = 2$ V, $C_L = 1$ pF)

Schemes	Energy (PJ)			Delay (ns)			E•D (PJ•ns)	Swing (V)	Complexity
	Driver/Wire	Receiver	Total	Driver/Wire	Receiver	Total			
CMOS	11.45	0.15	11.6	1.64	0.47	2.11	24.5	2.0	least area overhead
SSDLC ($V_{dd}=2.8$ V)	7.93	0.82	8.75	1.62	0.87	2.49	21.8	0.8	little area overhead
ASDLC	4.80	0.42	5.22	1.35	1.05	2.40	12.5	0.8	1 additional REF
LCLVD	2.18	0.25	2.40	1.08	1.42	2.50	6.00	0.7	Low-Vt devices, 1 REF
CCLC	2.25	0.42	2.67	1.13	1.47	2.60	6.94	0.8	coupling capacitor, 2 REFs
LCR	2.19	0.25	2.44	1.79	0.80	2.59	6.32	0.8	timing, 2 REFs
PDIFP	1.32	0.60	1.92	1.65	0.75	2.40	4.6	0.5	timing, 1 REF

TABLE IV
NOISE ANALYSIS OF PROPOSED SCHEMES ($V_{dd} = 2$ V, $C_L = 1$ pF)

Schemes	V_S (V)	$K_N V_S$	R_{x_O} (V)	R_{x_S} (V)	PS (V)	Atn	T_{x_O} (V)	V_N (V)	SNR
CMOS	2.0	0.26	0.15	0.15	0.1	1	0	0.66	1.52
SSDLC	0.8	0.104	0.1	0.05	0.1	0.7	0.05	0.37	1.08
ASDLC	0.8	0.104	0.1	0.05	0.1	0.7	0.05	0.37	1.08
LCLVD	0.7	0.091	0.1	0.03	0.1	0.5	0.03	0.30	1.17
CCLC	0.8	0.104	0.15	0.1	0.1	1	0.04	0.494	0.81
LCR	0.8	0.104	0.1	0.05	0.1	0.5	0.04	0.324	1.23
PDIFP	0.5	0.065	0.02	0.01	0.1	0.1	0.025	0.13	1.92

and carefully selected receivers should be employed. Cross-talk noise should also be handled with care, with good isolation between low-swing and full-swing signals.

To further compare the proposed schemes, two sets of simulations were performed. In the first set of simulations, V_{dd} is set at 2 V for all the schemes except for *SSDLC* ($V_{dd} = 2.8$ V), and the capacitive load on the interconnect is swept from 0 to 5 pF with the transistor sizes kept constant. The simulation results of four representing schemes (CMOS, ASDLC, LCLVD, and PDIFP) are shown in Fig. 11. All the proposed schemes have similar speed performances and their delays increase linearly with C_L . From the *energy* versus C_L plots, it can be observed that the energy values increase linearly against C_L , but with different slopes for different schemes. Low-swing schemes show increasing energy savings with increasing capacitive load, since the receiver energy overhead remains constant while the savings from the driver and wire become more and more dominant (e.g., PDIFP shows a factor of nine energy savings at $C_L = 5$ pF). Fig. 12 shows the second set of simulations, where C_L is set to 1 pF, while the supply voltage is swept from 1.5 to 3.3 V.

Rank ordering among the circuits is similar to Table III, while low-swing circuits can achieve higher energy efficiencies with increasing supply voltage. For instance, PDIFP has shown almost flat energy and energy-delay-product curves for the entire range, and it achieves energy savings of a factor of ten at 3.3 V.

V. CONCLUSION

Existing low-swing interconnect interface-circuit schemes show a wide variety of problems in both efficiency, performance, and reliability. We have introduced a number of novel or improved circuits to address some of these problems, or to get even higher energy savings. The schemes using threshold voltage drops can reduce the energy consumption by 55% with little overhead. Several schemes with very low driver supplies can reduce the energy consumption by a factor of four–six. The pseudodifferential scheme combines the best performance and greatest energy savings, with the best reliability. In summary, reducing the swing on interconnect is an effective and powerful tool for the minimization of energy dissipation, but requires a judicious optimization with respect to robustness, design complexity, and energy reduction.

APPENDIX

In Section II, we introduced briefly the worst case noise analysis method [12] to measure the reliability of each circuit. Here, we would like to elaborate the physical explanations of the noise sources for interested readers.

A. Cross Talk

Cross talk is noise induced by one signal that interferes with another signal. On-chip cross talk primarily comes from capacitive coupling of nearby signals (Fig. 13). The cross-talk





 Fig. 11. Delay, energy, and energy-delay product versus capacitive load of interconnect at $V_{dd} = 2$ V.

 Fig. 12. Delay, energy, and energy-delay product versus supply voltage at $C_L = 1$ pF.


Fig. 13. Cross-talk noise. (a) Coupling to a floating interconnect and (b) coupling to a driven interconnect.

coupling coefficient K_C is derived from the ratio between coupling capacitance and wire load capacitance ($K_C = 0.4$ for the targeted test bed). For the case of coupling to a floating interconnect, a ΔV of the aggressor (line A) will cause a ΔV on the victim (line B), and $\Delta V_B = K_C \Delta V_A$. If line B is driven with an output impedance of R [Fig. 13 (b)], then ΔV_B becomes a transient, which will decay with a time constant, $\tau = R(C_C + C_B)$. Therefore, the cross-talk noise attenuation for the static driver scenario can be achieved by increasing the timing budget for the signal so that the charge loss due to the cross-talk noise can be recovered by the driver. In Table I, we set $A_{tn_C} = 1$ for dynamic drivers, and $A_{tn_C} = 0.2$ for static ones.

B. Supply Noise

The IR drop of the power and ground distribution networks and the ringing of LC components of these networks cause the power rails of both drivers and receivers to vary in both time and space. The noise induced by the currents from all of the drivers is proportional to signal swing. Using the estimation techniques introduced in [12], the signal-induced supply noise is estimated to be 5% of the signal swing for the case of single-ended signaling across the chip (10 mm apart). Differential signaling will induce double size the noise onto the power rails, but since it has

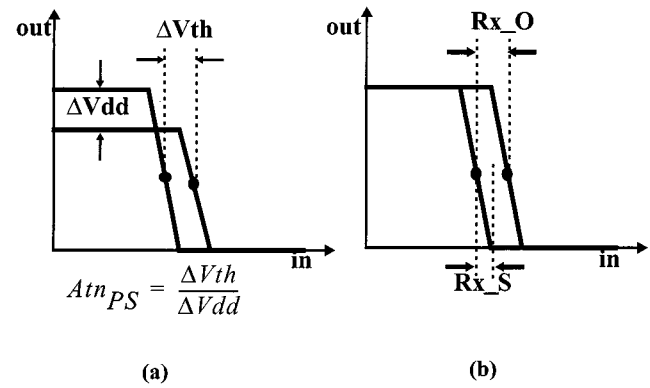


Fig. 14. Voltage transform curves. (a) Receiver input threshold varies with supply noise. (b) Receiver input offset due to process variation; receiver sensitivity.

good common-mode rejection of power supply noise (the attenuation factor is estimated as 10%), the effective signal-induced supply noise will be 1% of the signal swing.

For a well-designed power distribution network the signal-unrelated power supply noise is assumed to be 5% of the magnitude of power supply. The power supply attenuation coefficient is defined as the change of the receiver switching threshold voltage induced by an unit change of the supply voltage [see Fig. 14 (a)].

C. Receiver Input Offset, Receiver Sensitivity, and Transmitter Offset

Process variations (e.g., transistor threshold voltage variation, device size mismatch, etc.) will induce receiver input offset noise [Rx_O in Fig. 14 (b)]. For each of the receivers, every process corner case is simulated to get the worst difference of the input threshold (e.g., an inverter has 150-mV input offset). A differential source-coupled pair has a relatively small input

offset (20 mV in our circuits) because it only depends on the local mismatch of transistor V_t and sizes.

Fig. 14 (b) also shows the definition of receiver sensitivity as a half of the transient region of the VTC (e.g., an inverter has a sensitivity of 150 mV while a differential pair has only 10 mV). The transmitter offset results from the parameter mismatch between the transmitter and receiver, such as threshold voltage mismatch and reference voltage variation (estimated as 5% in our test circuits).

ACKNOWLEDGMENT

The authors acknowledge the efforts of the UC Berkeley ee241 class of Spring 1997, which contributed greatly to the analysis of some of the low-swing circuit schemes.


REFERENCES

- [1] D. Liu *et al.*, "Power consumption estimation in CMOS VLSI chips," *IEEE J. Solid-State Circuits*, vol. 29, pp. 663–670, June 1994.
- [2] E. Kusse, "Analysis and circuit design for low power programmable logic modules," M.S. thesis, Univ. Calif., Berkeley, 1997.
- [3] B. Gunning *et al.*, "A CMOS low-voltage-swing transmission-line transceiver," *ISSCC Dig. Tech. Papers*, pp. 58–59, Feb. 1992.
- [4] E. Musoll *et al.*, "Working-zone encoding for reducing the energy in microprocessor address buses," *IEEE Trans. VLSI Syst.*, vol. 6, pp. 568–572, Dec. 1998.
- [5] M. R. Stan and W. P. Burleson, "Bus-invert coding for low-power I/O," *IEEE Trans. VLSI Syst.*, vol. 3, pp. 49–58, Mar. 1995.
- [6] H. Zhang and J. Rabaey, "Low-swing interconnect interface circuits," in *Proc. 1998 Int. Symp. Low Power Electronic Devices*, Monterey, CA, Aug. 1998, pp. 161–166.
- [7] Y. Nakagome *et al.*, "Sub-1-V swing internal bus architecture for future low-power ULSI's," *IEEE J. Solid-State Circuits*, vol. 28, pp. 414–419, Apr. 1993.
- [8] M. Hiraki *et al.*, "Data-dependent logic swing internal bus architecture for ultralow-power LSI's," *IEEE J. Solid-State Circuits*, vol. 30, pp. 397–402, Apr. 1995.
- [9] H. Yamauchi *et al.*, "An asymptotically zero power charge-recycling bus architecture for battery-operated ultrahigh data rate ULSI's," *IEEE J. Solid-State Circuits*, vol. 30, pp. 423–431, Apr. 1995.
- [10] R. Colshan and B. Jaroun, "A novel reduced swing CMOS BUS interface circuit for high speed low power VLSI systems," *Proc. IEEE Int. Symp. Circuits and Systems*, vol. 4, pp. 351–354, May 1994.
- [11] J. Rabaey, *Digital Integrated Circuits*. Englewood Cliffs, NJ: Prentice-Hall, 1996.
- [12] W. Dally and J. Poulton, *Digital Systems Engineering*. Cambridge, U.K.: Cambridge Univ. Press, 1998.
- [13] A. J. Stratakos, "High-efficiency low-voltage dc-dc conversion for portable applications," Ph.D. dissertation, Univ. Calif., Berkeley, 1998.
- [14] T. Burd, "Energy efficient processor system design," Ph.D. dissertation, Univ. Calif., Berkeley, 1998.



Hui Zhang (S'95) received the B.S. degree in physics from the University of Science and Technology of China in 1993. He is currently working toward the Ph.D. degree in electrical engineering at the University of California at Berkeley.

His research interests include low-power circuits, low-power interconnect architectures, and reconfigurable DSP architectures for wireless applications. He has also been working on superconducting electronics and low-temperature CMOS circuits.



Varghese George (S'94) received the M.Tech. degree in electronics in 1993 from the Cochin University of Science and Technology, Kerala, India. He is currently working toward the Ph.D. degree in electrical engineering at the University of California at Berkeley.

From 1993 until 1994, he was a Research Engineer at the Raman Research Institute, India. He is now working on low-energy embedded reconfigurable architectures. His research interests include low-power techniques at the architecture and circuit levels.

Jan M. Rabaey (S'80–M'83–SM'92–F'95) received the E.E. and Ph.D. degrees in applied sciences from the Katholieke Universiteit Leuven, Belgium, in 1978 and 1983, respectively.

From 1983 to 1985, he was with the University of California at Berkeley as a Visiting Research Engineer. From 1985 to 1987, he was a Research Manager at IMEC, Belgium, where he pioneered the development of the CATHEDRAL II synthesis system for digital signal processing. In 1987, he joined the faculty of the Electrical Engineering and Computer Science Department, University of California at Berkeley, where he is now a Professor and the Vice Chair as well as the Scientific Codirector of the newly formed Berkeley Wireless Research Center (BWRC). He has authored or coauthored a wide range of papers in the area of signal processing and design automation. His current research interests include the exploration and synthesis of architectures and algorithms for digital signal processing systems and their interaction. He is also active in various aspects of portable distributed multimedia systems, including low-power design, networking, and design automation. He has served as Associate Editor of the *IEEE JOURNAL OF SOLID-STATE CIRCUITS*.

Dr. Rabaey received numerous scientific awards, including the 1985 IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN Best Paper Award (Circuits and Systems Society), the 1989 Presidential Young Investigator Award, and the 1994 Signal Processing Society Senior Award. He has served as Associate Editor of the *TODAES ACM Journal*. He is/has been on the program committee of the ISSCC, EDAC, ICCD, ICCAD, ASP-DAC, High Level Synthesis, and VLSI Signal Processing conferences. He is also the Vice Chair of the 2000 Design Automation Conference to be held in Los Angeles, CA, in June 2000.