Multiwire Phase Encoding: A Signaling Strategy for High-Bandwidth, Low-Power Data Movement

Prashansa Mukim and Forrest Brewer, *Senior Member, IEEE*

***Abstract*— This article presents multiwire phase encoding (MWPE), a transition signaling technique aimed at chip-to-chip communication on silicon interposer technology, where multiple, relatively low-bandwidth transmission lines can be easily routed between high-performance dies. The encoding exploits timing correlation between transitions on multiple band-limited wires to achieve high ensemble bandwidth, potentially exceeding that of parallel conventionally encoded NRZ links. The unambiguous encoding enables instantaneous bit synchronization, resulting in low-power, low-latency, PLL-/DLL-free on-chip data movement. Theoretical and practical bandwidths achievable by phase-encoded links are evaluated as a function of channel properties. Link timing, driver, and receiver circuits are implemented to evaluate the link performance and power costs associated with moving MWPE data. The Hspice simulation-based estimates indicate that a 2-mm-long MWPE link can achieve 126-Gb/s bandwidth on a lossy, dispersive transmission line medium, with an energy cost of 0.24 pJ/bit in 22-nm FDX technology.**

***Index Terms*— High bandwidth, low-power on-chip link, transiting signaling.**

I. Introduction

IT HAS been well established that the energy cost of moving data across a large scale System-on-Chip (SoC) or Network-on-Chip (NoC) is orders of magnitude higher than the cost of computation [1]–[3]. Most widely used high-performance I/O interfaces are based on either a source synchronous architecture [4], where the clock is transmitted on a separate channel along with parallel data channels or an embedded clock architecture [5], where the clock is recovered from the transmitted data edges. These techniques generally transmit non-return-to-zero (NRZ) encoded data and utilize delay-locked loops (DLLs) or phase-locked loops (PLLs) at the receiver for clock and data recovery. While techniques to reduce the dynamic power consumption of NRZ links have been proposed [6] and [7], these links also incur static power consumption due to an always-ON PLL or DLL and require hundreds to tens of thousands of clock cycles to achieve a lock. This increases the link overhead, especially if the communication is burst-mode or is switched from different sources.

Transition signaling protocols, in contrast, entirely encode a clock in signal transitions and, thus, do not require a sampling clock to recover the transmitted data. This eliminates most of the static power consumption associated with a DLL- or PLL-based link receiver and also makes the link "fire-and-forget" as the data can be recovered instantaneously. However, currently known transition signaling techniques, such as 1-of-n level encoded transition signaling (LETS) [8], m-of-n RTZ encoding [9], relative order-based encoding [10], and relative time-based encoding [11], achieve this at the cost of substantially lower effective data bandwidth compared to conventional NRZ level encoded links. Furthermore, their use has been limited to short-distance and low-speed (<2 Gb/s) bus-based communication.

One differentiating work is a two-phase level encoded dual-rail (LEDR) link proposed by Dobkin *et al.* [12], which has been shown to operate at bandwidths up to 67 Gb/s in 65-nm technology. However, this technique achieves a high data bandwidth by transmitting pulses as narrow as 15 ps and may not be viable for communication over band-limited, lossy, and dispersive on-chip wires with a reach of several millimeters. In this article, we propose a general transition signaling protocol called multiwire phase encoding (MWPE), which is designed to exploit timing correlations between multiple signaling wires to transmit data at effective rates: 1) higher than the bandwidth limit of a single wire; 2) higher than the rate of other known transition signaling techniques; and 3) comparable to or higher than parallel conventional NRZ encoded links. Both passive [13] and active [14] interposer links have been successfully implemented and shown promise, as has the use of transition signaling for communication between multiple locally clocked but globally asynchronous domains on an NoC [15]. On-chip or on-interposer links would be the ideal target applications for the proposed MWPE encoding.

The organization of this article is given as follows. The MWPE protocol, theoretical bandwidth limits, and bandwidth comparisons are presented in Section II. The link architecture and performance limits of a six-wire MWPE link on practical on-chip channels are presented in Section III. Circuits for the six-wire MWPE link are implemented in 22-nm FDX technology, and the link is characterized in terms of

Manuscript received December 17, 2020; revised April 25, 2021; accepted May 16, 2021. Date of publication June 18, 2021; date of current version August 2, 2021. *(Corresponding author: Prashansa Mukim.)*

Prashansa Mukim was with the Department of Electrical and Computer Engineering, University of California at Santa Barbara, Santa Barbara, CA 93106 USA. She is now with the Department of Chemistry and Biochemistry, University of Maryland, College Park, MD 20742 USA, and also with the National Institute of Standards and Technology (NIST), Gaithersburg, MD 20899 USA (e-mail: prashansa.mukim@nist.gov).

Forrest Brewer is with the Department of Electrical and Computer Engineering, University of California at Santa Barbara, Santa Barbara, CA 93106 USA (e-mail: forrest@ece.ucsb.edu).

Color versions of one or more figures in this article are available at https://doi.org/10.1109/TVLSI.2021.3087585.

Digital Object Identifier 10.1109/TVLSI.2021.3087585

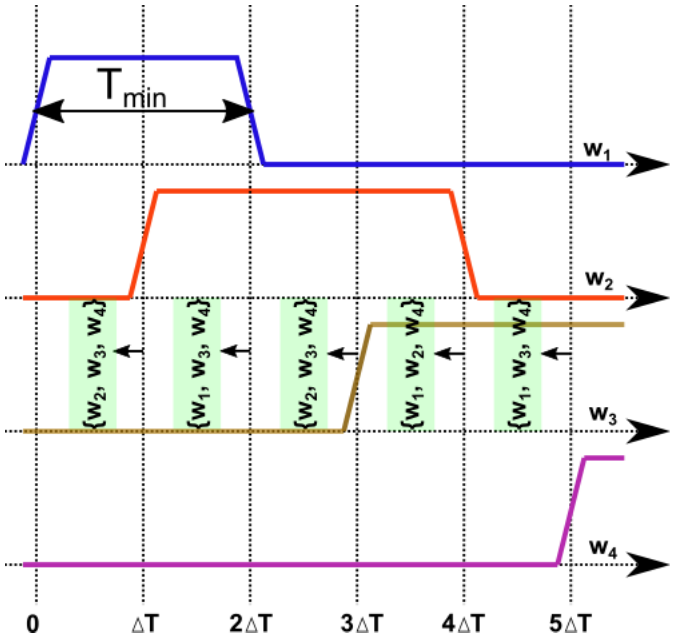

Fig. 1. Transitions on an $n = 4, k = 2$ S-MWPE link.

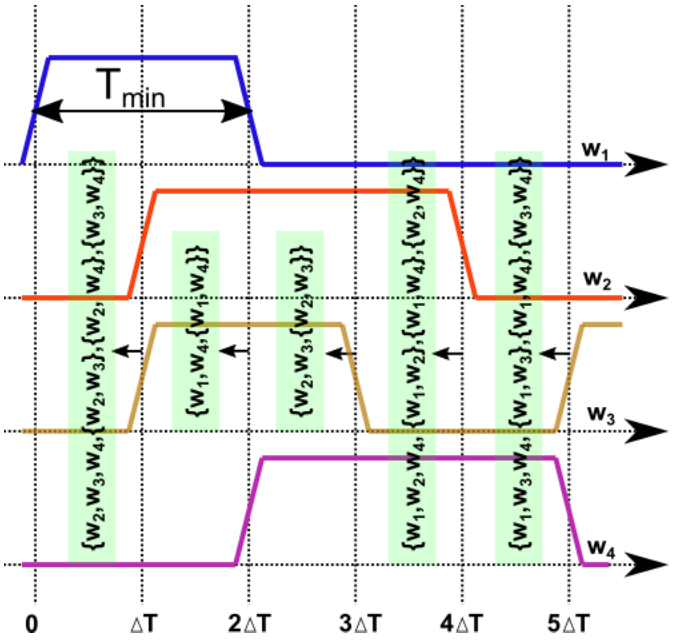

Fig. 2. Transitions on an $n = 4, k = 2$ M-MWPE link.

power, performance, and reliability against PVT variations in section IV.

II. Multiwire Phase Encoding: Protocol and Theoretical Bandwidth Limits

Assume $T_{\min}$ to be the minimum pulsewidth that can be reliably transmitted on a given channel, chosen based on length, loss, and dispersion limits. This dictates that consecutive transitions on any particular wire should be apart in time by at least $T_{\min}$. However, for a multiwire link, transitions on *different* wires have no such constraint. MWPE on an n-wire link is defined by the following three rules.

1) Consecutive transitions on any wire must be apart in time by at least $T_{\min}$.
2) Consecutive transitions on different wires can happen at a minimum interval of ΔT, which is a fraction of $T_{\min}$, i.e., $\Delta T = T_{\min}/k$. We define k to be the number of phases, and ΔT is the resolution of consecutive transitions on different wires.
3) Data are encoded in the identity of switching wire(s).

MWPE can be of the single-transition type [or single-transition MWPE (S-MWPE)] if only a single wire is allowed to transition at any given time or of the multitransition type [or multitransition MWPE (M-MWPE)] if multiple wires are allowed to transition simultaneously. Fig. 1 shows an example of an S-MWPE link with $n = 4$ wires and $k = 2$ phases such that a single wire can switch at intervals $\Delta T = T_{\min}/2$, as long as consecutive transitions on the same wire are apart by at least $T_{\min}$. Fig. 2 shows an example of an M-MWPE link with $n = 4$ wires and $k = 2$ phases such that one or two wires can switch at intervals $\Delta T = T_{\min}/2$, and switching times between consecutive transitions on the same wire have a lower limit of $T_{\min}$. The data transmission bandwidth for each type of encoding can be estimated as follows.

S-MWPE: Considering the example of Fig. 1, after wire-1 (w_1) switches at time 0, for the next word transmission at time ΔT, any wire except w_1 can switch without violating $T_{\min}$. As $k = 2$, this constraint is released at time $2\Delta T$. Thus, on each phase, there are three available wires that can switch, enabling transmission of $\log_2 3$ bits of data at a rate $= \Delta T = T_{\min}/k$.

By the same argument, an n-wire, k-phase S-MWPE link can transmit $\log_2(n - k + 1)$ bits of data at a rate $= \Delta T = T_{\min}/k$, resulting in an effective data transmission bandwidth of

$$\mathrm{BW}_{\text{S-MWPE}} = \frac{k \times \log_2(n - k + 1)}{T_{\min}}. \tag{1}$$

$k \geq 2$ ensures that ΔT is, indeed, less than $T_{\min}$, and $k \leq n-1$ ensures that at least two wires are available to be switched on every phase and at least $\log_2 2 = 1$ bit of data can be transmitted by means of transitions on each of the k-phases.

For bandwidth estimation of M-MWPE, consider the example of Fig. 2. At every ΔT step, either one wire can switch (leaving three available wires to switch in the next time step) or two wires can switch (leaving two available wires to switch in the next time step, allowing at least 1 bit of data transmission). Let the state where one wire switched in the previous time step be called S_1 and the state where two wires switched in the previous time step be called S_2. When the state is S_1, there are six combinations of transitions possible on the three available wires: three single-wire transitions + $\binom{3}{2}$ double-wire transitions, and thus, $\log_2 6 = 2.58$ bits can be transmitted. On the other hand, when the state is S_2, there are three combinations of transitions possible on the two available wires: two single-wire transitions + 1 double-wire transition, and $\log_2 3 = 1.58$ bits can be transmitted. To estimate the *average* number of transmitted bits, we need to know the probability of being in each state. If each combination is

equally likely, when in S_1, three of the six combinations result in the next state being S_1, and when in S_2, two of the three combinations result in the next state being S_1. Therefore, the probability of being in each state is $S_1 =$

$$p_1 = p_1 \times (3/6) + (1 - p_1) \times (2/3). \tag{2}$$

On solving, $p_1 = 4/7$, and the probability of being in each state is $S_2 = p_2 = 1 - p_1 = 3/7$. The average data transmission bandwidth for $n = 4, k = 2$ M-MWPE is then

$$\frac{p_1 * \log_2 6 + p_2 * \log_2 3}{\Delta T} = \frac{4.3}{T_{\min}}. \tag{3}$$

The average bandwidth for other combinations of n and k can be obtained similarly and will be of the form shown in the following equation:

$$\mathrm{BW}_{\text{M-MWPE}} = \frac{k \times \sum_{i=1}^{S} (p_i \times \log_2 c_i)}{T_{\min}} \tag{4}$$

where S in the total number of states for that encoding, p_i is the probability of being in state i, and c_i is the number of switching combinations available when in state i.

Algorithm 1 describes the steps for estimating the bandwidth of any n, k M-MWPE link. First, all possible states are calculated, where the length of each state is $k - 1$ (as every transitioning wire is busy for $k - 1$ time steps), and the elements of each state are ordered to denote how many wires transitioned in the previous $k - 1$ time steps. A maximum of $n - 2$ wires can be busy at any given time (to allow transmission of at least one bit in the next cycle). All permutations of lists of length $k - 1$ whose elements lie between 1 and $n - 2$, with the sum of all elements less than or equal to $n - 2$ is the set of valid states for the given n and k. Next, the number of transitions possible for every state is counted. The count depends on: 1) the number of available wires to switch given the history of the current state and 2) the maximum number of wires that can switch in the current time step, given the history of the current state, such that at least two wires are available to switch in the next time step. The maximum number of wires allowed to switch is the minimum of these two values, and count_{ij} stores the number of combinations of transitions for state i when j wires switch. Then, for each count_{ij} value, the next state is formed by the history of the current state except the first element, as it will no longer be busy, and the number of wires that switch in the current state (j). Now, the probability p_i of being in each state S_i can be calculated by scanning through all the count_{mj} values for each state S_m. If the next state for a count_{mj} value is the same as the state for which the probability is being calculated (S_i), the ratio of this count_{mj} value to the sum of all count_{mj} values for the current state S_m is multiplied with the probability p_m of being in the state being iterated (S_m) and added to the total probability p_i of being in state S_i. Linear equations for the probability of each state create S (= number of states) linear equations in S variables. Solving these equations determines the state probabilities. The bandwidth for the encoding is calculated by calculating the average number of bits transmitted in every time step ($= T_{\min}/k$).

Algorithm 1 Calculating Bandwidth for M-MWPE

function BANDWIDTH(n, k, T_{min})
  $states \leftarrow$ STATES(n, k)
  for all $S_i \in states$ **do**
    $count_i \leftarrow$ COUNT(n, k, S_i)
    $c_i = \sum_{allj} count_{ij}$
    $nextState_i \leftarrow$ MAP_NEXT_STATE($S_i, count_i$)
  end for
  $P \leftarrow$ PROB($states, count, nextState$)
  $bandwidth_{M-MWPE} = \dfrac{\sum(p_i \times log_2 c_i)}{T_{min}/k}$
  return $bandwidth_{M-MWPE}$
end function

---

function STATES(n, k)
  $states \leftarrow \{$all $S_i \in \{1, n-2\}^{k-1}$ such that
$\sum_{j=1}^{k-1} S_{ij} \leq n-2\}$
  return $states$
end function

---

function COUNT(n, k, S_i)
  $avail\_switch \leftarrow n - \sum_{j=1}^{k-1} S_{ij}$
  $next\_switch \leftarrow n - 2 - \sum_{j=2}^{k-1} S_{ij}$
  $max\_switch \leftarrow min(avail\_switch, next\_switch)$
  for $j := 1 : max\_switch$ **do**
    $count_{ij} \leftarrow \binom{avail\_switch}{j}$
  end for
  return $count_i$
end function

---

function MAP_NEXT_STATE($S_i, count_i$)
  for all $count_{ij} \in count_i$ **do**
    $nextState_{ij} \leftarrow \{\{$ all $S_{ij} \in S_i$ such that $j \geq 2\}, j\}$
  end for
  return $nextState_i$
end function

---

function PROB($states, count, nextState$)
  for all $S_i \in states$ **do**
    for all $count_m \in count$ **do**
      for all $count_{mj} \in count_m$ **do**
        if $nextState_{mj} == S_i$ **then**
          $p_i \leftarrow p_i + p_m \times \dfrac{count_{mj}}{\sum_{allj} count_{mj}}$
        end if
      end for
    end for
  end for
  $\sum p_i = 1$
  Solve $S \times S$ matrix ($S = |states|$) to get all p_i
  $P \leftarrow \{p_i\}_{i=1}^{S}$
  return P
end function

The receiver decodes data by determining which of the n-wires switched, as well as maintaining the history of the previous $k - 1$ transitions. It should be noted that ΔT

TABLE I
N-WIRE BANDWIDTH COMPARISONS

Encoding	Effective Bandwidth
Conventional level-based NRZ	$\frac{n}{T_{min}}$
1-of-n NRZ LETS [8]	$\frac{log_2(n)}{T_{min}}$
m-of-n RTZ [9]	$\frac{\log_2 \binom{n}{m}}{2 \times T_{min}}$
Order Encoding [10]	$\frac{log_2(n!)}{T_{min} + (n-1)\Delta T}$
Proposed S-MWPE	$\frac{k \times log_2(n-k+1)}{T_{min}}$
Proposed M-MWPE	$\frac{k \times \sum_{i=1}^{S}(p_i \times log_2 c_i)}{T_{min}}$

cannot be arbitrarily small, as the received signal will show timing deviations caused by channel dispersion, intersymbol interference, crosstalk, device noise, and PVT variations. Ensuring that ΔT is greater than the total expected timing noise of the received signal will result in successful data reception.

Table I compares the effective data bandwidths for conventional level-based NRZ encoding, other transition signaling protocols, and phase encoding. The 1-of-n LETS [8] is similar to S-MWPE in which it is one-hot, but the transitions on any two wires are constrained to be apart by T_{min} and not a fraction of T_{min}. The m-of-n RTZ codes [9] involve transitions on $\binom{n}{m}$ wires every bit period and are denser than the 1-of-n LETS codes, but the bit-period is twice of T_{min} as it is a return-to-zero protocol and transmits data on only one edge. Order codes [10] further improve the code density by encoding data in the relative order of transitions on n-wires, resulting in $n!$ available codes transmitted over a bit period $= T_{\text{min}} + (n-1)\Delta T$. Although counterintuitive, the efficacy of order codes actually reduces as n increases because the time to transmit a symbol is also $\propto n$.

It is clear from Table I that the proposed S-MWPE protocol outperforms the other transition signaling protocols in terms of the effective bandwidth for sufficiently large values of k. Fig. 3 shows a comparison of the per-wire bandwidth of conventional NRZ encoding and MWPE as a function of the number of wires, n. For $n \geq 6$, S-MWPE can outperform the bandwidth of parallel conventional NRZ encoding without requiring clock-based decoding. Fig. 4 shows the bandwidth comparison of different M-MWPE encodings. M-MWPE results in a higher effective bandwidth than certain S-MWPE cases; for example, when $n = 4$ and $k = 2$, S-SWPE achieves a bandwidth that is 80% of an equivalent NRZ encoded link, whereas M-MWPE achieves a bandwidth that is 110% of NRZ. However, from a switching point of view, M-MWPE is also prone to larger crosstalk induced time deviations, as it

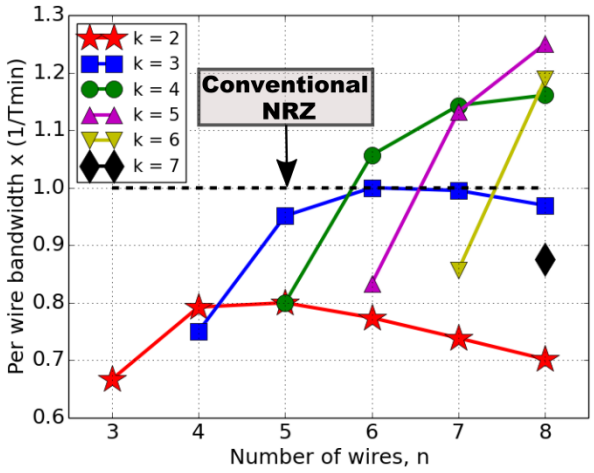

Fig. 3. Bandwidth comparisons of S-MWPE with conventional NRZ links.

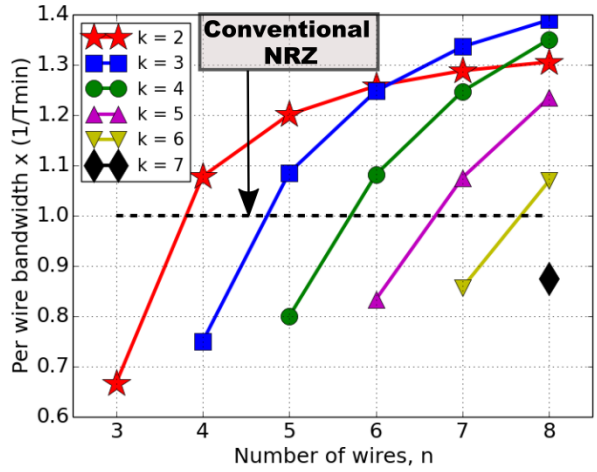

Fig. 4. Bandwidth comparisons of M-MWPE with conventional NRZ links.

allows simultaneous switching of wires. Although choosing a larger n may lead to larger per wire bandwidths, wire length matching requirements are imposed across a larger set of wires in order to preserve phase relations of the transmitted signals. For a fixed value of T_{min}, choosing a larger k may also lead to larger per wire bandwidths. However, this leads to smaller values of ΔT, tightening the timing error margins between transitions. In Section III, the performance of S-MWPE and M-MWPE protocols on practical on-chip channels will be evaluated to determine the actually achievable bandwidth given channel distortion. The analysis will be focused on $n = 6$, $k = 2$ MWPE links as they offer reasonably high per-wire and total bandwidths with the largest possible ΔT and low encoding/decoding complexity.

III. LINK ARCHITECTURE AND PERFORMANCE ON PRACTICAL CHANNELS

Six-wire, two-phase S-MWPE and M-MWPE links were implemented in 22-nm FDX technology, and the link circuits and coupled transmission lines using W-models [16] were simulated in Hspice. The link architecture is shown in Fig. 5.

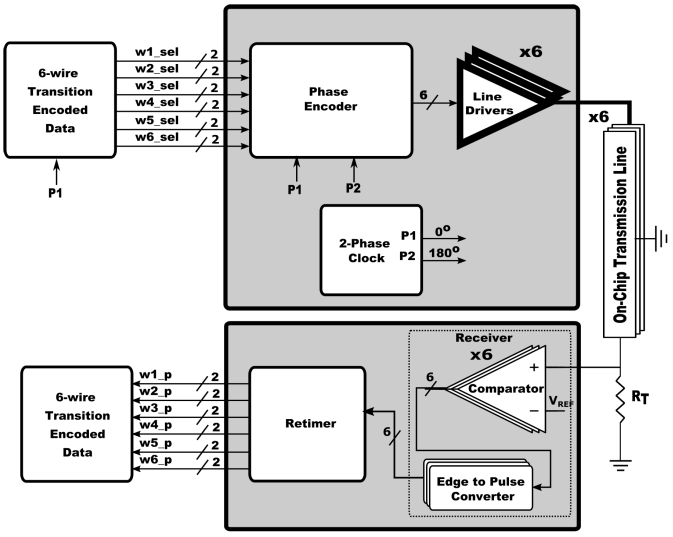

Fig. 5. Six-wire MWPE link architecture.

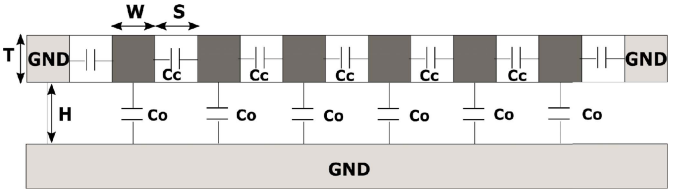

Fig. 6. Six-wire coupled microstrip line topology: $W = S = 3\ \mu$m, $H = 1\ \mu$m, $T = 90$ nm; $C_o = 230$ pF/m, $L_o = 165$ nH/m, $C_c = 56$ pF/m, and $R_o = 1.93$ KΩ/m.

At the transmit end, a two-phase clock is utilized to read from a shift register block and random six-wire transition encoded data, which follows the protocol rules described in Section II. The data for each of the six wires are one-hot encoded and transmitted on two wires. For example, a high value on $w1$_sel[0] would indicate that wire-1 is selected to switch on clock phase P1, whereas a high value on $w1$_sel[1] would indicate that wire-1 is selected to switch on clock phase P2. The phase encoder block uses the values on the wire select lines to switch one or more of the six wires on clock phase P1 or P2. The line drivers drive six coupled single-ended transmission lines. Fig. 8 shows the transient waveforms of various signals at the transmitter end. The lines are terminated at the receiver, where a comparator first amplifies the received data and converts rising and falling edges to pulses, which are finally retimed to occur at $k \times \Delta T$ intervals by the retimer block. The output of the retimer consists of pulses on two channels for each wire, whose presence indicates that the corresponding wire is transitioned in the respective clock phase. Fig. 9 shows the transient waveforms of various signals at the receiver end.

In order to evaluate the maximum allowable bandwidths for $n = 6, k = 2$ S-MWPE and M-MWPE links on practical on-chip channels, the link was driven by random data (following the rules specified by the encoding protocol) with the minimum pulsewidth on any wire ($T_{\min}$) ranging between 40 and 80 ps. Timing deviations in the received data due to channel dispersion, intersymbol interference, and crosstalk

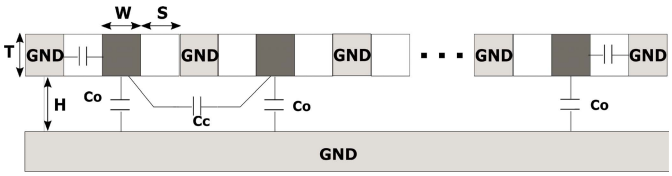

Fig. 7. Six-wire shielded coplanar waveguide line topology: $W = S = 3\ \mu$m, $H = 1\ \mu$m, $T = 90$ nm; $C_o = 230$ pF/m, $L_o = 160$ nH/m, $C_c = 8.5$ pF/m, and $R_o = 1.93$ KΩ/m.

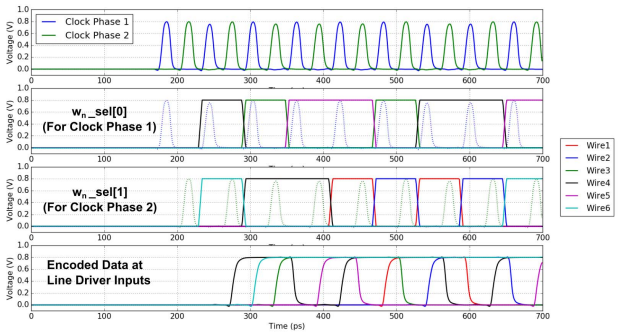

Fig. 8. Transient waveforms for signals at the transmitter.

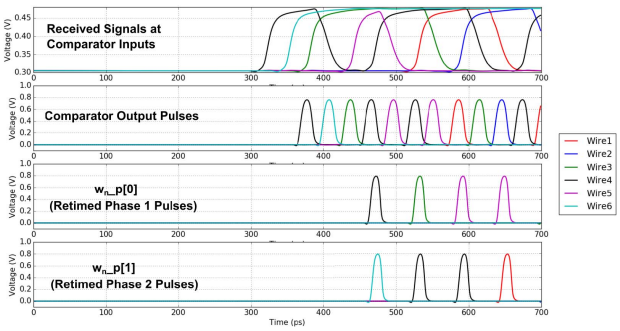

Fig. 9. Transient waveforms for signals at the receiver.

effects were then extracted to analyze the reliability of the link at different lengths. The first channel that was simulated is shown in Fig. 6. It consists of six unshielded local chip wire microstrip lines with the parameters shown in the figure. The outer two wires have a larger C_o/C_c value than the inner two wires and the values shown in the figure are the ones that lead to a minimum C_o/C_c. A channel comprising of six shielded coplanar waveguides, as shown in Fig. 7, was also simulated, which shows significantly smaller crosstalk effects but consumes nearly double the wiring resources. The insertion loss for these channels at 20 GHz for lengths of 1–5 mm lies in the range of 1.6–6 dB. While eye diagrams are a useful metric to evaluate the reliability of NRZ encoded links involving a sampling receiver clock, their direct use is not appropriate for evaluating MWPE links. MWPE links do not make use of any sampling clock at the receiver, and their reliability only depends on the timing deviations of *local successive* transitions, as opposed to global timing deviations across all transitions. For example, as shown in Fig. 10, a 4-mm-long shielded S-MWPE link shows closed eyes on all wires. However, when the eye diagrams are realigned such

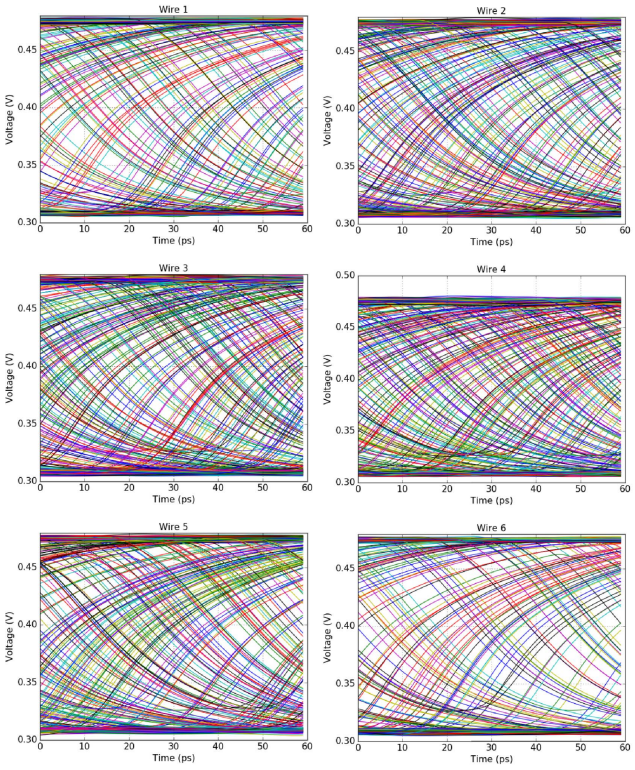

Fig. 10. Eye diagrams for $n = 6,\ k = 2,\ T_{\min} = 60$-ps S-MWPE shielded links of 4-mm length.

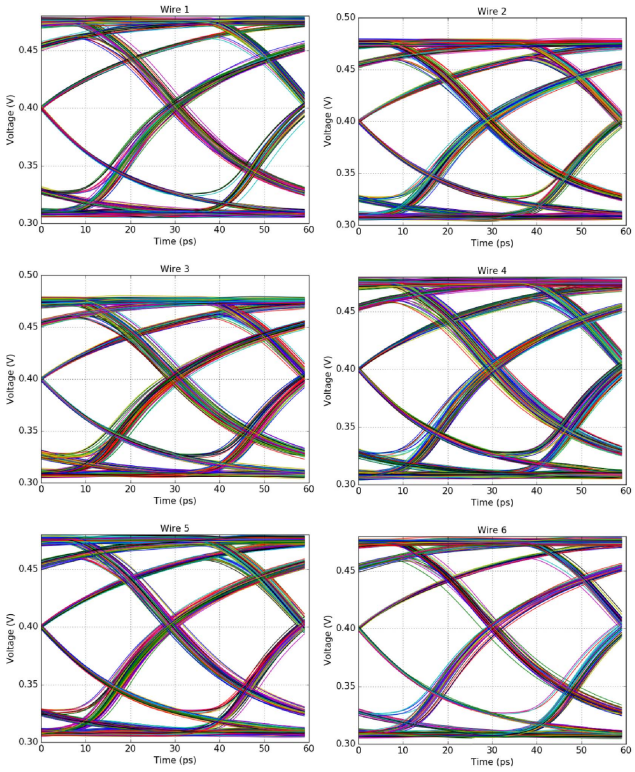

Fig. 11. Aligned eye diagrams for $n = 6,\ k = 2,\ T_{\min} = 60$-ps shielded S-MWPE links of 4-mm length.

that for each wire, all zero crossings of the waveforms are considered to be the origin, and voltages on all wires are plotted for a bit period starting from this origin (as shown in Fig. 11), the resulting eye diagrams show open eyes. This shows that, even though the actual eye diagrams show closed eyes, local edge-to-edge timing of deviations is much smaller, and decoding based on the timing of consecutive edges is possible. Thus, *jitter*, as defined in Fig. 12 to be the deviation in the timing of consecutive transitions from the ideal timing of consecutive transitions (ΔT), is used to assess the reliability of MWPE links.

Fig. 13 shows the jitter for $n = 6, k = 2, T_{\min} = 60$-ps S-MWPE links implemented using the unshielded and shielded transmission line topologies of Figs. 6 and 7, with line lengths of 4 mm. The following constraint is used to determine if the link operation would fail with the given jitter: the retiming circuit at the receiver is designed to reliably trigger at $minimum$ consecutive transition times of $(2/3)\Delta T$, it is known to not retrigger if the spacing between consecutive transition is between $0 - (1/3)\Delta T$, and the window between $(1/3)\Delta T - (2/3)\Delta T$ is the "unreliability" window where the retimer may trigger but produce pulses too narrow to drive the subsequent logic, and hence, the window should be avoided. Thus, for the S-MWPE link with a single transitioning wire on every phase, the link would not fail as long as the time between consecutive transitions is $\geq (2/3)\Delta T$, or if the jitter is $< -\Delta T/3 = -10$ ps. As shown in Fig. 13, the S-MWPE link has enough margin when shielded signal wires are used

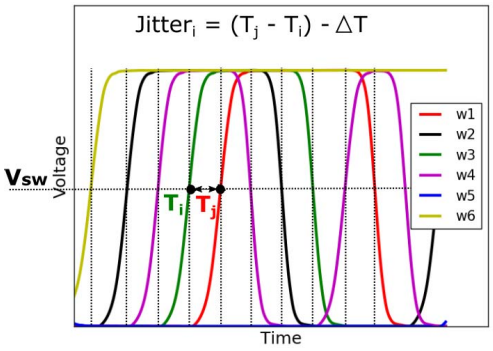

Fig. 12. Jitter for an MWPE link computed from relative timing of consecutive transitions.

and is very close to failure with unshielded wires for wire lengths of 4 mm.

Fig. 14 shows the jitter for $n = 6, k = 2, T_{\min} = 60$-ps M-MWPE links implemented using the same unshielded and shielded transmission line topologies. The reliability of M-MWPE links depends on one additional constraint: since multiple wires can transition in the same clock phase, the jitter for the same phase transitions should be $<(1/3)\Delta T$ to avoid retriggering the retiming circuit, as the retiming circuit is designed to trigger *once* every clock phase (working described in Section IV). The M-MWPE link has enough margin when shielded signal wires are used, and it fails with unshielded wires of 4-mm length, as shown in Fig. 14. Fig. 15 shows the maximum wire lengths over which data can be reli-

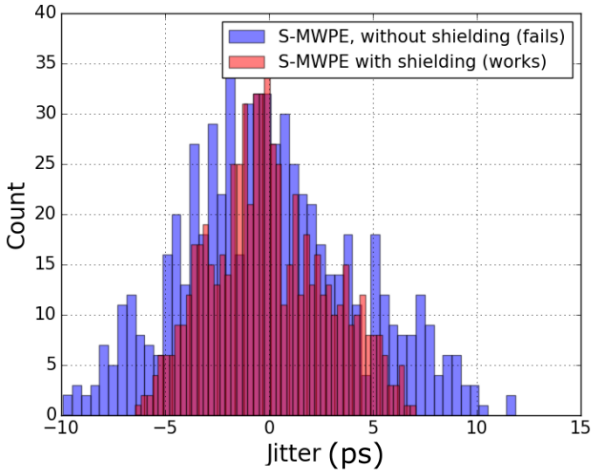

Fig. 13. Jitter for $n = 6, k = 2, t_{\min} = 60$ ps S-MWPE links of 4-mm length.

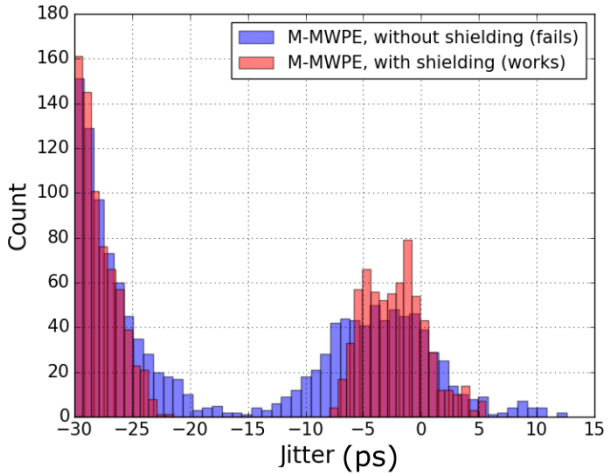

Fig. 14. Jitter for $n = 6, k = 2, T_{\min} = 60$ ps M-MWPE links of 4-mm length.

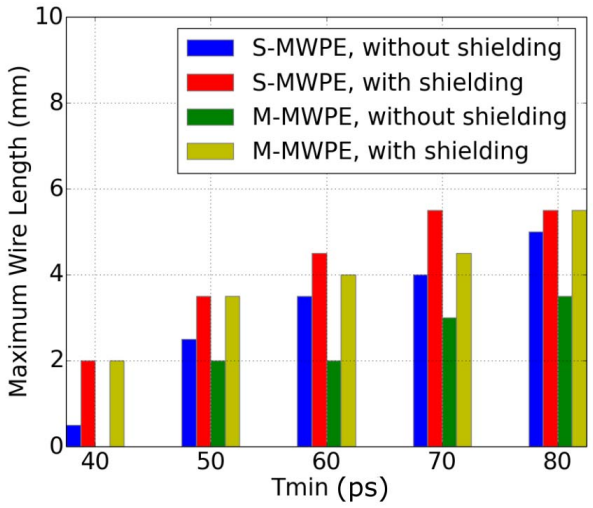

Fig. 15. Maximum link length for which $n = 6, k = 2$ MWPEs are successful for the given channel for different $T_{\min}$'s.

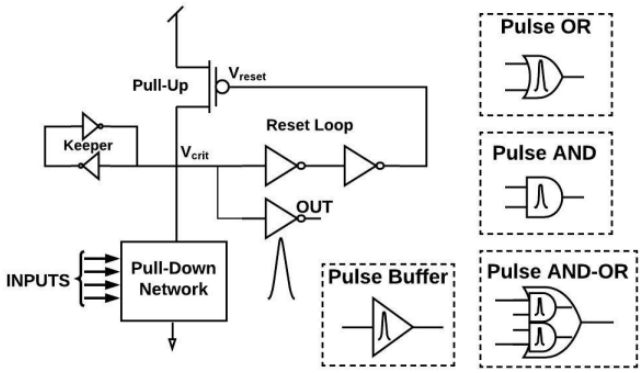

Fig. 16. Pulse gates used for implementing the link.

ably transmitted using shielded and unshielded S-MWPE and M-MWPE links with varying minimum pulse widths $T_{\min}$. When $T_{\min}$ is 40 ps, the unshielded M-MWPE link fails even for line lengths of 0.5 mm. Unshielded S-MWPE links are reliable for longer distances than M-MWPE links as they exhibit lower worst case crosstalk-induced timing deviations. Intersymbol interference, loss, and dispersion effects limit the minimum $T_{\min}$ that can be used. Note that, as these arguments are worst case, restricted M-MWPE codes avoiding the worst case crosstalk transition pairs might well achieve higher performance; however, the selection of such codes is beyond the scope of this article. Section IV presents link circuits that are implemented in 22-nm FDX technology to evaluate the maximum performance achievable in this technology, the power cost of moving data using MWPE, and factors that govern the tolerance of the link under PVT variations.

IV. Link Implementation

The two-phase clock, phase encoder, and retimer circuits were implemented using pulse logic [17], a self-resetting logic family of gates, typically used for the design of high-performance asynchronous circuits. Pulse logic is only rising-edge sensitive and uses atomic pulses to encode both data and timing, allowing the efficient realization of high-performance transition signaling protocols, such as MWPE. Fig. 16 shows a pulse gate, which accepts a pulse at the pull-down network input to produce an output pulse, whose width is primarily set by the delay of the reset loop. Different logic functions, such as OR, AND, and AND-OR, can be implemented by modifying the pull-down network. Pulse AND in particular allows gated pulse triggering, as one input is a pulse and the other input is a level. To convert a pulse to a level, latches shown in Fig. 17 are used. The rules for constructing pulse gate circuits are described in detail in [18].

In 22-nm FDX technology, the FO4 delay of an inverter implemented using low-V_t transistors is $\approx$5 ps, and for a pulse gate, the propagation delay is $\approx$7 ps, and the reset time is $\approx$20 ps. Thus, the retimer can retrigger at minimum intervals of 20 ps, and we use a $\Delta T = 30$ ps in order to meet the timing constraints described in Section III. The clock used in this implementation is a 12-gate, one-pulse collective pulse oscillator (CPO) [19], designed to have a period $= T_{\min} =$ 60 ps, and the two phases used in the link are tapped from the

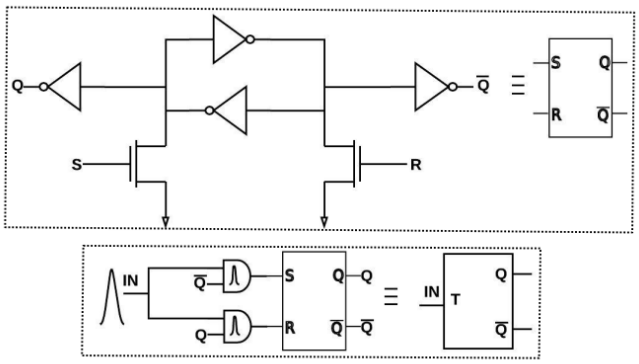

Fig. 17. Latches used for implementing the link.

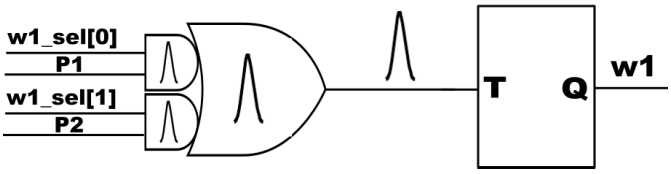

Fig. 18. Phase encoder circuit.

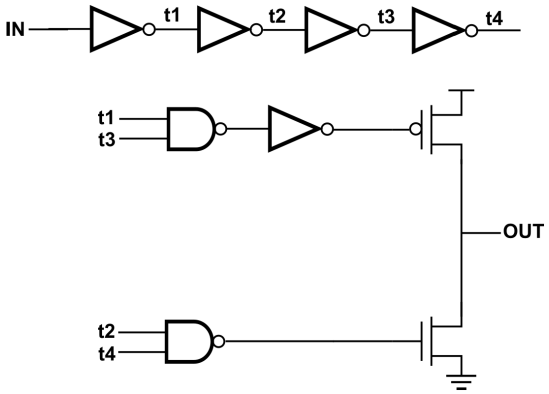

Fig. 19. Driver circuit.

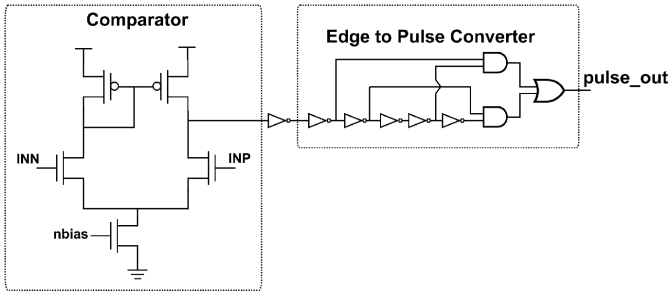

Fig. 20. Receiver circuit.

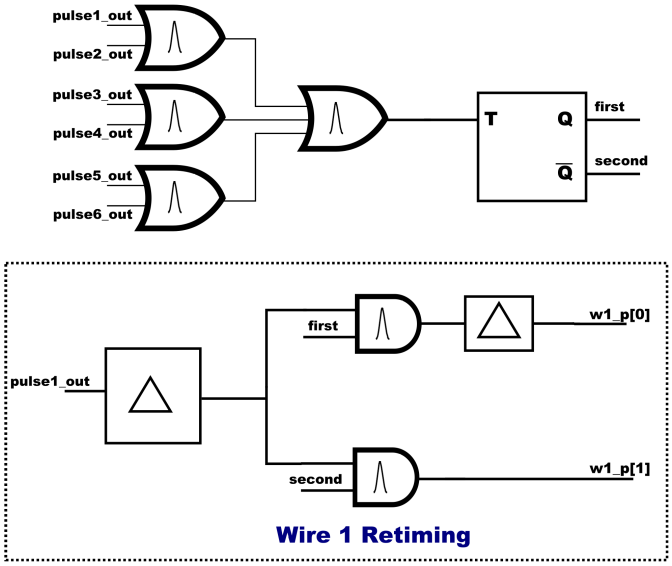

Fig. 21. Retimer circuit.

first and seventh gates of the ring. CPOs have been shown to exhibit better noise rejection properties than similarly powered inverter-based ring oscillators [19], and their use for MWPE links helps to lower the timing noise at the transmit end.

The phase encoder circuit shown in Fig. 18 is a pulsed multiplexer utilizing pulse AND-OR gates. It multiplexes either clock phase P1 or P2 to the input of a toggle (T)-latch, based on the values (levels) of the $w1\_sel[0:1]$ signals. The T-latch, thus, produces a transition on wire-1 in one of the clock phases. The encoder circuit of Fig. 18 is replicated six times for the six wires and drives the line driver shown in Fig. 19. The driver is a single-ended make-before-break driver [20] driving a terminated transmission line of $Zo \approx 27\ \Omega$, with the signal swing on each line being ≈ 170 mV.

At the receiver, the signal is amplified, and rising and falling transitions are converted to pulses (see Fig. 20), which are retimed by the circuit shown in Fig. 21 to occur at intervals $= 2 \times \Delta T = T_{\min}$. The retimer combines transition events (represented by pulses) on the six wires using a tree of pulse-OR gates, whose output drives a T-latch (or a 1-bit counter) to indicate the clock phase in which the wire is transitioned. In parallel, delayed versions of transitions on all wires (encoded as pulses) are split into two channels based on the phase in which they occurred using pulse-AND gates, and the output of the channel containing phase-1 pulses is delayed by one clock phase to be aligned with the pulses on the second channel. In general, for a k-phase MWPE link, a k-bit counter can be used to split events on each wire to k-channels identifying their phase of occurrence. Thereafter, the data can be decoded by evaluating groups of k-consecutive transitions. The following two-sided timing constraints should be met for the retimer circuit to work as intended:

Reset-Time Constraint: Minimum separation between consecutive events on different clock phases should be longer than the reset time of the OR-gates used in the combining tree to ensure that all distinct events are captured.

Setup-Time Constraint: The T-latch should be set before the next event arrival, and therefore, the delay of the T-latch should be less than the minimum expected separation between events on distinct clock phases.

Hold-Time Constraint: The delay of the T-latch should not be smaller than the width of pulses; else, both pulse-AND gates in the retiming logic would trigger. Due to these constraints, the maximum allowable jitter due to channel distortions in Section III was limited to $\Delta T/3$.

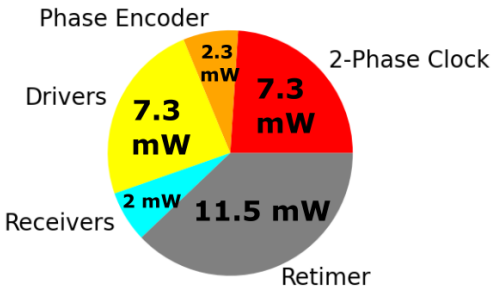

Fig. 22. Power breakdown.

TABLE II
COMPARISONS WITH STATE OF THE ART

Reference	[21]	[22]	[23]	This work
BW/pin (Gbps)	30	21	10-20	12-21
Energy (pJ/b)	1.11	0.94	0.68-1.36	0.24-0.39
Technology	28nm	28nm	65nm	22nm
Signaling	PAM-3	CNRZ-5	NRZ	MWPE
Measured	Yes	Yes	Yes	No

The same circuits can be used for both S-MWPE and M-MWPE links. The S-MWPE link achieves a total bandwidth = 77.4 Gb/s ($\log_2 5 = 2.32$ bits transmitted every 30 ps), and the M-MWPE link achieves 125.8 Gb/s (3.77 average bits transmitted every 30 ps) with total power consumption of 30.4 mW according to the power breakdown shown in Fig. 22. Thus, an $n = 6$, $k = 2$, $T_{\min} = 60$ ps S-MWPE link allows data movement across wires up to 3.5–4.5-mm long (see Fig. 15) at an energy cost of 0.39 pJ/b, whereas an $n = 6$, $k = 2$, $T_{\min} = 60$ ps M-MWPE link allows data movement across wires up to 2–4-mm long at an energy cost of 0.24 pJ/b. Table II compares the performance of MWPE links with related state-of-the-art works. By using transmit preemphasis and/or receive equalization, the transmission length could be further extended. Since the receiver for MWPE links does not make use of a PLL/DLL, these links eliminate the static power dissipation associated with an always-ON PLL/DLL. In a system where these links are used with an activity factor α, i.e., the links are active for a fraction of time $= \alpha$, the total power consumption would be $= \alpha \times P_{\text{dynamic}} + P_{\text{static}}$, where P_{dynamic} and P_{static} are the link dynamic and static power consumptions, respectively. P_{static} in MWPE links corresponds to the bias current of the receiver comparators ($\approx$ 2mW), whereas, for PLL-/DLL-based links, this term would be much higher corresponding to the PLL/DLL power. The other system-level benefit of MWPE links is due to the instantaneous link recovery. If the lock time of a PLL-/DLL-based link receiver is L bit periods and payloads of length P bit periods are transmitted, the link efficiency would be $P/(P + L)$, whereas, for MWPE links, the efficiency would be close to 100% for all values of P as the data can be recovered instantaneously.

Finally, these encodings are also robust against PVT variations, as long as the retiming circuit can be guaranteed to work even at the slowest corner as the transmitter works at the fastest corner, and vice versa. Fig. 23 shows how the delay of a pulse gate varies with local and global process variations, $\pm 10\% V_{\text{dd}}$ variations, and temperature variations between 0 °C and 100 °C. With the tunable body bias offered by 22-nm FDX technology, the n and p body biases of the transmitter circuits, particularly the clock, and the retiming circuit at the receiver end can be tuned to lower the TX/RX speed disparity and further improve the link robustness.

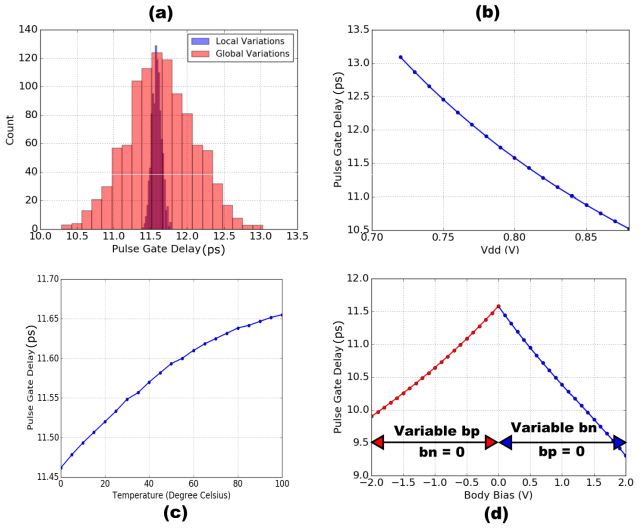

Fig. 23. Pulse gate delay variance due to (a) process variations, (b) voltage variations, (c) temperature variations, and (d) body bias tuning.

V. CONCLUSION

A general class of MWPE protocols, which encode data in the timing correlation of transitions on multiple signaling wires as opposed to voltage levels, was presented in this work. They offer the advantage of very high-bandwidth data transmission on lossy and dispersive on-chip channels at rates higher than the effective bandwidth of the channel and comparable to or higher than conventional level encoded NRZ links. Furthermore, these links do not use PLL-/DLL-based recovery mechanisms, leading to instantaneous data recovery and low idle-state power dissipation. In this article, theoretical bandwidth bounds for S-MWPE and M-MWPE were derived, and the performance of these encodings on practical on-chip/on-interposer channels was evaluated. Simulation of link circuits for moving MWPE data across wires of 2–5-mm long in 22-nm FDX technology showed that the energy cost of these links is as low as 0.2–0.4 pJ/b.

REFERENCES

[1] A. Boroumand *et al.*, "Google workloads for consumer devices: Mitigating data movement bottlenecks," in *Proc. Archit. Support Program. Lang. Operating Syst. (ASPLOS)*, 2018, pp. 316–331.
[2] S. Borkar and A. A. Chien, "The future of microprocessors," *Commun. ACM*, vol. 54, no. 5, pp. 67–77, 2011.
[3] M. Horowitz, "1.1 computing's energy problem (and what we can do about it)," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2014, pp. 10–14.
[4] B. Casper *et al.*, "A 20Gb/s forwarded clock transceiver in 90nm CMOS B," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2006, pp. 263–272.

[5] J. Lee and B. Razavi, "A 40-Gb/s clock and data recovery circuit in 0.18-μm CMOS technology," *IEEE J. Solid-State Circuits*, vol. 38, no. 12, pp. 2181–2190, Dec. 2013.

[6] M. R. Stan and W. P. Burleson, "Bus-invert coding for low-power I/O," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 3, no. 1, pp. 49–58, Mar. 1995.

[7] P. Behnam and M. N. Bojnordi, "STFL: Energy-efficient data movement with slow transition fast level signaling," in *Proc. 56th Annu. Design Autom. Conf.*, Jun. 2019, pp. 1–6.

[8] P. B. McGee, M. Y. Agyekum, M. A. Mohamed, and S. M. Nowick, "A level-encoded transition signaling protocol for high-throughput asynchronous global communication," in *Proc. 14th IEEE Int. Symp. Asynchronous Circuits Syst.*, Apr. 2008, pp. 116–127.

[9] W. J. Bainbridge, W. B. Toms, D. A. Edwards, and S. B. Furber, "Delay-insensitive, point-to-point interconnect using m-of-n codes," in *Proc. 9th Int. Symp. Asynchronous Circuits Syst.*, 2003, pp. 132–140.

[10] C. D'Alessandro, D. Shang, A. Bystrov, A. Yakovlev, and O. Maevsky, "Multiple-rail phase-encoding for NoC," in *Proc. 12th IEEE Int. Symp. Asynchronous Circuits Syst. (ASYNC)*, Mar. 2006, pp. 107–116.

[11] C. D'Alessandro, N. Minas, K. Heron, D. Kinniment, and A. Yakovlev, "NoC communication strategies using time-to-digital conversion," in *Proc. 1st Int. Symp. Netw.-on-Chip (NOCS)*, May 2007, pp. 65–74.

[12] R. Dobkin, R. Ginosar, and A. Kolodny, "Fast asynchronous shift register for bit-serial communication," in *Proc. 12th IEEE Int. Symp. Asynchronous Circuits Syst. (ASYNC)*, Mar. 2006, p. 10.

[13] M. W. Chaudhary and A. Heinig, "High speed interfaces for chip to chip communication on interposer based integration," in *Proc. 21st Eur. Microelectron. Packag. Conf. (EMPC) Exhib.*, Sep. 2017, pp. 1–4.

[14] M. S. McCorquodale, J. D. O'Day, S. M. Pernia, G. A. Carichner, S. Kubba, and R. B. Brown, "A monolithic and self-referenced RF LC clock generator compliant with USB 2.0," *IEEE J. Solid-State Circuits*, vol. 42, no. 2, pp. 385–399, Feb. 2007.

[15] D. Bertozzi, G. Miorandi, M. Tala, and S. M. Nowick, "Cost-effective and flexible asynchronous interconnect technology for GALS networks-on-chip," in *Proc. New Gener. CAS (NGCAS)*, Sep. 2017, pp. 77–80.

[16] *HSPICE User Guide: Signal Integrity Modeling and Analysis, Version Q-2020.03*, Synopsys, Research Triangle, NC, USA, May 2013.

[17] M. Miller, G. Hoover, and F. Brewer, "Pulse-mode link for robust, high speed communications," in *Proc. IEEE Int. Symp. Circuits Syst.*, vol. 5, May 2008, pp. 3073–3077.

[18] M. Miller, C. Segal, D. Mc Carthy, A. Dalakoti, P. Mukim, and F. Brewer, "Impolite high speed interfaces with asynchronous pulse logic," in *Proc. Great Lakes Symp. VLSI*, May 2018, pp. 99–104.

[19] P. Mukim *et al.*, "Design and analysis of collective pulse oscillators," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 28, no. 5, pp. 1242–1255, May 2020.

[20] W. J. Dally and J. W. Poulton, *Digital Systems Engineering*. Cambridge, U.K.: Cambridge Univ. Press, Jun. 1998.

[21] H. Park *et al.*, "30-Gb/s 1.11-pJ/bit single-ended PAM-3 transceiver for high-speed memory links," *IEEE J. Solid-State Circuits*, vol. 56, no. 2, pp. 581–590, Feb. 2021.

[22] A. Shokrollahi *et al.*, "A pin-efficient 20.83Gb/s/wire 0.94pJ/bit forwarded clock CNRZ-5-coded SerDes up to 12 mm for MCM packages in 28nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Jan. 2016, pp. 182–183.

[23] H. G. Rhew, M. P. Flynn, and J. Park, "A 22Gb/s, 10mm on-chip serial link over lossy transmission line with resistive termination," in *Proc. Eur. Solid-State Circuits Conf.*, Sep. 2012, pp. 233–236.