

# A Fine-Grain, Uniform, Energy-Efficient Delay Element for 2-Phase Bundled-Data Circuits

AJAY SINGHVI, Stanford University

MATHEUS T. MOREIRA, Pontifícia Universidade Católica do Rio Grande do Sul

RAMY N. TADROS, University of Southern California

NEY L. V. CALAZANS, Pontifícia Universidade Católica do Rio Grande do Sul

PETER A. BEEREL, University of Southern California

Contemporary digitally controlled delay elements (DEs) trade off power overheads and delay quantization error (DQE). This article proposes a new programmable DE that provides a balanced design that yields low power with moderate DQE even under process, voltage, and temperature variations. The element employs and leverages the advantages offered by a 28nm fully depleted silicon on insulator technology, using back body biasing to add an extra dimension to its programmability. To do so, a novel generic delay shift block is proposed, which enables incorporating both fine and coarse delays in a single DE that can be easily integrated into digital systems, which is an advantage over hybrid DEs that rely on analog design.

CCS Concepts: • **Hardware** → **Circuits power issues**; **Asynchronous circuits**;

Additional Key Words and Phrases: Delay elements, fine-grain delay, delay quantization error, low power, FD-SOI

## ACM Reference Format:

Ajay Singhvi, Matheus T. Moreira, Ramy N. Tadros, Ney L. V. Calazans, and Peter A. Beerel. 2016. A fine-grain, uniform, energy-efficient delay element for 2-phase bundled-data circuits. *J. Emerg. Technol. Comput. Syst.* 13, 2, Article 15 (November 2016), 23 pages.

DOI: <http://dx.doi.org/10.1145/2948067>

## 1. INTRODUCTION

Delay elements (DEs) are used in a variety of applications in VLSI systems and are typically employed to provide precise timing control and/or satisfy timing constraints, which can be strict in nanoelectronic design. In synchronous systems, DEs support clock distribution and synchronization across different blocks dealing with clock skew and jitter problems [Chakraborty et al. 2008; Jung et al. 2001]. Other uses include phase-locked loops, digitally controlled oscillators [Moon et al. 2008], time-to-digital converters [Li and Chou 2007], and polyphase clock generators [Lin and Chen 2001]. DEs are also widely used in bundled-data asynchronous systems to control the timing of request and acknowledge signals between different blocks [Heck et al. 2015]. For

---

M. T. Moreira and N. L. V. Calazans were supported by CNPq (grants 401839/2013-3 and 312556/2014-4) and CAPES (grant 2129/14-0). R. N. Tadros and P. A. Beerel were supported by a gift from Qualcomm.

Authors' addresses: A. Singhvi, Paul G Allen Building, 420 Via Palou Mall, Stanford University, Stanford, CA 94305; email: [asinghvi@stanford.edu](mailto:asinghvi@stanford.edu); M. T. Moreira, Faculty of Informatics - PUCRS, Av. Ipiranga, 6681, Building 32 Suite 726, Porto Alegre, RS, Brazil, 90619-900; email: [matheus.moreira@pucrs.br](mailto:matheus.moreira@pucrs.br); N. L. V. Calazans, Faculty of Informatics - PUCRS, Av. Ipiranga, 6681, Building 32 Suite 718, Porto Alegre, RS, Brazil, 90619-900; email: [ney.calazans@pucrs.br](mailto:ney.calazans@pucrs.br); R. N. Tadros, Electrical Engineering Building (EEB), 3740 McClintock Ave, Los Angeles, CA 90089, USA; email: [rtadros@usc.edu](mailto:rtadros@usc.edu); P. A. Beerel, EEB 350, MC 2562, Ming Hsieh EE Dept., USC, 3740 McClintock Ave, Los Angeles, CA 90089; email: [pabeerel@usc.edu](mailto:pabeerel@usc.edu).

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).

© 2016 ACM 1550-4832/2016/11-ART15 \$15.00

DOI: <http://dx.doi.org/10.1145/2948067>

some of these applications, like control circuits of 2-phase bundled-data asynchronous designs, DEs require balanced rise and fall delays [Heck et al. 2015; Beerel et al. 2010]. Moreover, a typical concern in the design of DEs in modern technologies is the effect of process, voltage, and temperature (PVT) variations. To account for those, DEs must be conservatively designed to have extra timing margins, which can compromise performance. The alternative is to use programmable DEs.

Programmable DEs alleviate the detrimental effects of PVT variations in deep sub-micron technologies by providing a range of attainable delays to which the DE can be tuned post-silicon. The delay granularity provided by programmable DEs is an important concern. For instance, systems that require precise timing control, such as phase shift compensators [Dogsa et al. 2014], timing generators [Ryu et al. 2013], and timing verniers [Arkin 2004] used for delay fault testing in automatic testing equipment, employ fine-grain DEs to ensure correct operation. In essence, the precision to which these DEs can be tuned affects the amount of timing margin that they can effectively avoid. DEs can be controlled by either analog voltages (or currents) or digitally. Traditionally, analog-controlled DEs provide fine delay tuning, whereas digitally controlled DEs provide coarse-grain delays, with their combination forming hybrid DEs. However, since this work deals primarily with energy-efficient digital VLSI circuits, the use of hybrid DEs is not considered, to avoid the high power consumption of the required analog circuitry, the switching noise at high frequencies, and the challenges in the distribution of global analog signals in predominantly digital systems [Heck et al. 2015]. The target application of this DE is low-power 2-phase bundled-data asynchronous circuits in which energy efficiency is the primary concern. This is mainly because these circuits target low-power applications like mobile computers.

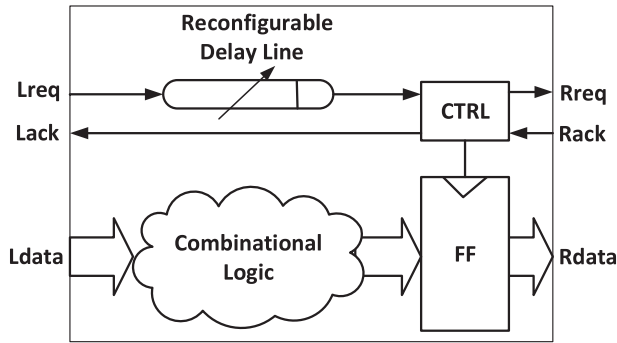
This article's contributions can be summarized as follows:<sup>1</sup>

- It proposes a new DE architecture providing low power and low delay quantization error (DQE) with balanced rise and fall delays.
- It proposes the delay shift inverter (DSI) employing fully depleted silicon on insulator (FD-SOI) back body biasing [Pelloux-Prayer et al. 2013] to equip the DE with an additional dimension of programmability, achieving fine-grain delay.
- It details the available options for designers to generate the required biasing voltages and proposes the use of the contention mitigated level shifter (CMLS) proposed in Tran et al. [2005].
- It discusses the various trade-offs and optimizations in designing DEs, including those that target performance, leakage, dynamic energy, area, and both local and global variability.

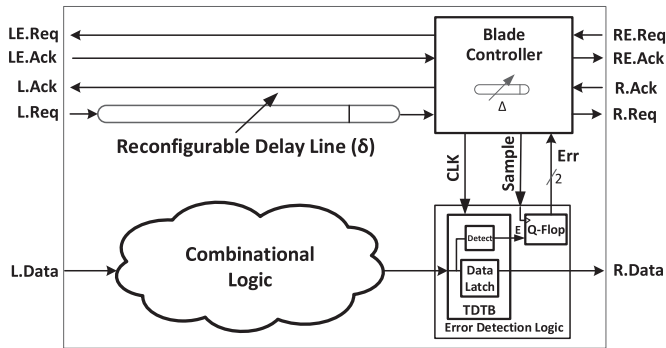
The remainder of the article is organized as follows. Section 2 discusses the target application and introduces important metrics for the required DEs. It also reviews the state of the art in digitally controlled DEs and provides an overview of back-body biasing in FD-SOI technologies. Next, Section 3 explains the design of the proposed DE to provide a quasilinear and monotonic delay characteristic, reducing the DQE to 12.57% from 269.92% presented by a state-of-the-art DE. It also proposes the architecture of the DSI to provide fine-grain delays in a single DE structure that can be easily incorporated into digital systems without any of the problems posed by hybrid DEs. Moreover, it discusses the methodology adopted for optimizing power consumption of the proposed design, resulting in significantly lower energy consumption than existing DEs [Maymandi-Nejad and Sachdev 2003, 2005]. Section 4 presents and discusses our simulation results, which include Monte Carlo analysis to assess and compare the

---

<sup>1</sup>This article is an extension of the work originally published in Singhvi et al. [2015], where the DE proposition explored here first appeared.



(a) Bundled-data circuit stage with a reconfigurable DE



(b) Resilient bundled-data circuit stage

Fig. 1. Two examples of asynchronous templates that use programmable DEs.

impact of PVT variations on the compared DEs. Finally, Section 5 summarizes and concludes the article.

## 2. BACKGROUND

### 2.1. Bundled-Data Design

Asynchronous bundled-data circuits are traditionally designed with a single DE per pipeline stage that is matched to the critical path through the associated combinational logic, as illustrated in Figure 1(a). This critical path is the path with the longest delay over all input conditions and is typically determined via static timing analysis. The DE can be static, providing a single delay, or programmable, providing a range of delays. In the former case, any expected PVT variations must be accounted for in the design of the DE, because if PVT variations result in the DE being too short, the chip becomes inoperable. For many technologies, these variations can be large and eliminate much of the advantages of asynchronous circuits. In the latter case, post-silicon tuning can mitigate the variations, making these preferable [Diamant et al. 2015].

During chip characterization, worst-case vectors can be repeatedly applied with different DE codeword settings, finding the codeword that corresponds to the smallest delay for which the DE is sufficiently long for the test to pass with sufficient yield. During a manufacturing test, this codeword can be used to determine chip disposal. However, chips may also be shipped with a slightly longer programmed delay, adding a test margin [Xiong et al. 2009] to account for differences between a manufacturing test and normal operating conditions.

The traditional bundled-delay template determines five important desirable characteristics of a well-designed programmable DE:

- Energy efficiency.* Asynchronous bundled-data design has several advantages over traditional synchronous design, including improving modularity, avoiding generation and distribution of a global clock, and lower electromagnetic interference [Beerel et al. 2010]. However, perhaps one of the most important aspects is its potential for achieving lower power than synchronous alternatives. To achieve this advantage, the extra circuitry necessary to support the asynchronous handshaking, including the DEs, must be efficiently designed. Because these designs may have low duty cycles and be idle for long periods of time, minimizing their dynamic power, measured via energy per transition (EPT) for a given delay, as well as controlling leakage power, are both important issues.
- Appropriate programmable delay range.* The tunability of a DE’s delay should be sufficiently large to mitigate the expected PVT variations. However, at the same time, supporting too much tunability makes the DE overly complex, wasting area and power.
- Fine-grain programmable delay.* The discrete nature of the codewords implies that not all delays are achievable. The difference between the desired delay and the closest discretely achievable delay becomes margin, some of which may be unwanted, yielding unnecessary performance degradations. Thus, it is desirable to provide fine-grain programmable control of the delay, to minimize margins.
- Low DQE.* The DQE is the maximum ratio of the achieved increase in delay between adjacent codewords over the average increase in delay. More formally,

$$DDe = \frac{DR}{N - 1}, \quad (1)$$

$$DQE = \frac{\max_i(|DD_i - DDe|)}{DDe} * 100\%, \quad (2)$$

where  $DR$ , the delay range, is the delay difference between the minimum and maximum delay settings, and  $N$  is an integer representing the number of codewords employed by a particular DE. In addition,  $DD_i$  is the delay difference between the  $i^{th}$  and  $(i + 1)^{th}$  adjacent codewords as observed in simulations, and  $DDe$  is the ideal expected delay difference, computed by (1).

A small DQE is desirable to enable the DE to be used efficiently across all codewords and possible delay values. Moreover, a small DQE is needed to ensure a consistent test margin across all chips. In other words, when the codeword is adjusted after a successful manufacturing test, a similar amount of increased delay should be added to the delay regardless of the initial codeword. Notice that the DQE notion encompasses the features of monotonicity, a uniform delay distribution across codewords, and the ability to predict the amount of delay provided by a particular codeword or by a change in codewords.

- Equal rise and fall times.* Depending on the target handshaking protocol, a bundled-data circuit has different requirements on the rise and fall times of the DEs [Beerel et al. 2010]. For the 4-phase handshaking protocol, the DEs should have a programmable rise delay to match the worst-case delay through the combinational logic but a fast fall time to minimize delay overhead associated with the control circuit’s return-to-zero phases. This is because in 4-phase control, the DE performs both transitions (rise and fall) for each token of data sent through the combinational logic, but only the rising transition signals valid data availability. For 2-phase handshaking protocols, on the other hand, both rising and falling transitions of the DEs are associated with transmitting data tokens through the combinational logic. Thus, both the rise and fall delays must match the worst-case combinational logic delay [Heck

et al. 2015]. Compared to their 4-phase counterparts, 2-phase designs can be faster and more energy efficient, because they completely avoid return-to-zero phases.

Interestingly, a recently proposed resilient bundled-data design called *Blade* [Hand et al. 2015] uses two DEs per pipeline stage, as illustrated in Figure 1(b). One DE, referred to as  $\delta$ , targets the average-case delay of the pipeline stage's combinational logic, and a second, referred to as  $\Delta$ , extends the delay of the next stage upon detecting a delay in this stage larger than the average-case delay. In particular, this is achieved using special error-detecting logic and asynchronous Blade controllers. When the error-detecting logic identifies that a setup timing violation occurred, it tells the Blade controller to delay the handshake with the next pipeline stage by an extra  $\Delta$  delay to maintain correct downstream operation. As in the regular bundled-data template, 2-phase handshaking is proposed to minimize control overhead. Consequently, maintaining equal rise and fall delays for the DE remains an important feature. However,  $\delta$  may need to be less fine grain than in traditional bundled data because differences from the desired delay in  $\delta$  only cause the design to deviate from the optimal error rate. Because Blade has a flat optimal curve (its optimal performance is not a strong function of  $\delta$ ), a 5% deviation in  $\delta$  can lead to less than a 2% drop in average performance [Hand et al. 2015]. Moreover, differences from the desired delay in  $\Delta$  impact performance only when errors occur. Since optimal error rates are often less than 25%, the  $\Delta$  DEs are only activated less than a quarter of the time and thus also need not be fine grain. A typical design will contain both error-detecting and non-error-detecting pipeline stages, and thus a range of granularities is desirable. However, a relatively low DQE for the  $\Delta$  DE is still useful, because even in resilient templates, adding a consistent amount of test margin is desired. Finally, as in the bundled-data template, the delay range of these DEs should be sufficiently large to mitigate PVT variations, but in fact a larger delay range could be desirable to support design for test and debug. In particular, a useful test mode of operation is to increase  $\delta$  to the point where no timing errors will occur.

## 2.2. State-of-the-Art Programmable DEs

Different digitally controlled DE architectures exist in the literature, exploring trade-offs in terms of delay range, power consumption, and area. Some existing DEs are the thyristor-based [Kim et al. 1996], transmission gate-based [Mahapatra et al. 2000], current starved [Maymandi-Nejad and Sachdev 2003], and cascaded inverter-based [Mahapatra et al. 2002] designs. Among these, thyristor-based designs provide delays in ranges from a few microseconds to a few milliseconds. However, their discussion is beyond the scope of this article, which focuses on DEs that provide shorter delay ranges (in the order of a few picoseconds to nanoseconds). Moreover, it is difficult to accurately control both rise and fall transitions in thyristor-based designs. For example, the transmission gate-based DE suffers from poor signal integrity, and modifications that alleviate the problem [Mahapatra et al. 2000] add significant costs in terms of area and power.

Therefore, the focus here is on cascaded and current-starved inverter (CSI)-based designs. The simplest and perhaps most common cascaded inverter-based design is the multiplexer-based DE (MUX-DE), depicted in Figure 2. Its popularity arises from the fact that it has a relatively simple design that can be implemented using standard cells. The codeword provided to the MUXes fixes the number of inverters in the signal path and hence its delay.

Maymandi-Nejad and Sachdev proposed a CSI-based design using a programmable current mirror to control the current and thus the delay through an inverter. Their CMCS-DE has low area [Maymandi-Nejad and Sachdev 2005], but the current mirror suffers from very large static power consumption and cannot be employed in low-power applications [Heck et al. 2015]. Another CSI-based DE is the directly controlled current starved DE (DCCS-DE), analyzed in Maymandi-Nejad and Sachdev [2003]

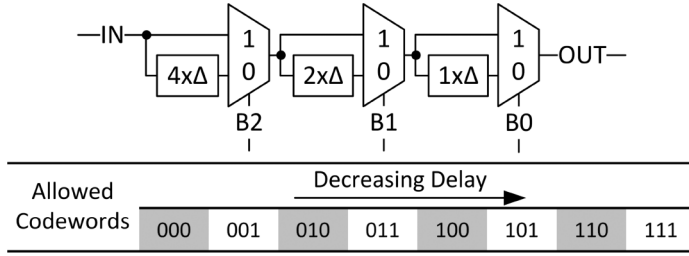


Fig. 2. Multiplier-based DE (MUX-DE).

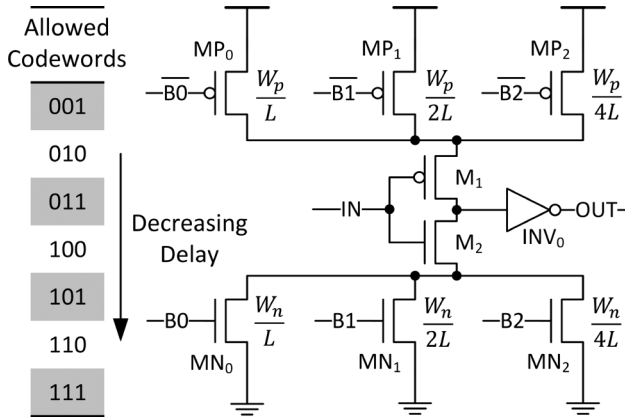


Fig. 3. Directly controlled current-starved DE (DCCS-DE) [Maymandi-Nejad and Sachdev 2003].

and shown in Figure 3. This DE has current source transistors with different lengths that determine the current through an inverter. These transistors reduce the current through the inverter (or starve it), thereby increasing the delay of a signal propagating through it. Compared to the MUX-DE, it has much lower EPT, making it attractive for low-power applications [Heck et al. 2015]. However, as analyzed in Maymandi-Nejad and Sachdev [2003], it displays a few drawbacks, including unequal rise and fall delays, as well as a nonuniform and nonmonotonic delay difference between successive codewords, leading to a relatively large DQE. Heck et al. [2015] modified the DCCS-DE design to yield balanced rise and fall delays. In particular, their design comprises two replicated DCCS-DEs in series, with signal conditioning inverters added to their inputs, to provide an acceptable slew rate, and inverters at their outputs, to provide the same load to each of the replicated DCCS-DEs. Unfortunately, this modified DCCS-DE still exhibits somewhat poor DQE, a problem that will be addressed by the DE that we propose in Section 3.1.

### 2.3. FD-SOI Technology

New VLSI technologies enable different optimization strategies and novel circuit techniques that can be used in the design of low-power DEs. In particular, some VLSI foundries have moved from dopant-controlled transistors to fully depleted (FD) devices [Hook 2012] typically in silicon on insulator (SOI), as illustrated in Figure 4. Because the Si film is very thin, the depletion region continues to its end (i.e., is FD). Compared to bulk devices, the resulting FD-SOI devices provide better electrostatic control and rely on the geometry of the structures instead of the doping to limit the drain-induced barrier lowering (DIBL) and subthreshold slopes. Moreover, the technology




Fig. 4. Reproduction of the cross section of a UTBB FD-SOI transistor.

reduces the random doping fluctuations and hence threshold voltage variations, making FD solutions well suited to low-voltage applications [Hook 2012].

In contrast, other VLSI foundries have adopted three-dimensional transistors known as FinFETs, either in SOI or bulk, to reduce leakage and mitigate PVT variations [Lin et al. 2011]. Although FinFETs have a lower contact resistance (or access resistance) than FD-SOI FETs, they also have an additional source of variation resulting from the device’s quantized width determined by the number of transistor fins drawn [Lin et al. 2011]. Moreover, the body effect in the FinFET is absent for both bulk or SOI types [Hook 2012]. Meanwhile, in planar thin-body FD devices, the threshold voltage strongly depends on the body potential, providing circuit designers with an additional dimension of control.

In particular, this work uses STMicroelectronics 28nm UTBB FD-SOI technology, where *UTBB* stands for ultrathin body and BOX, where *BOX* stands for buried oxide. Transistors are normally controlled by the high- $\kappa$  metal gate, which is called the *front gate*. In addition, due to the very small width of the UTBB, applying a potential from the back body (or the back gate) has a large influence on the transistor’s threshold voltage [Pelloux-Prayer et al. 2013]. This is what is called *back body biasing*, or just *body biasing*. There are two ways to employ body biasing: (i) forward body biasing (FBB), which decreases the threshold voltage and is used to boost performance, and (ii) reverse body biasing (RBB), which increases the threshold voltage and is typically used to lower dynamic and leakage power. For regular-Vt FETs in this technology (1V supply), the RBB range for pMOS and nMOS is 1 to 3V and 0 to  $-3$ V, respectively, whereas the FBB range is 1 to  $-0.3$ V and 0 to 0.3V, respectively [Flatresse et al. 2013].

More specifically, this work uses reverse back body biasing not to lower power but instead to provide fine-grain control to our proposed new DEs.

### 3. ENERGY-EFFICIENT, FINE-GRAIN, UNIFORM DES DESIGN

#### 3.1. Proposed Low DQE DE Design

As mentioned previously, both the DCCS design of Maymandi-Nejad and Sachdev [2003] and the modified one in Heck et al. [2015] have variable and nonmonotonic delay behavior, which results in a large DQE. To minimize DQE while taking into consideration low power and high density results in state-of-the-art DEs, a new version of the DCCS-DE is proposed herein, based on the design presented in Heck et al. [2015]. The new architecture, illustrated in Figure 5, uses a one-hot coding scheme instead of the binary codes used in Maymandi-Nejad and Sachdev [2003] and Heck et al. [2015].

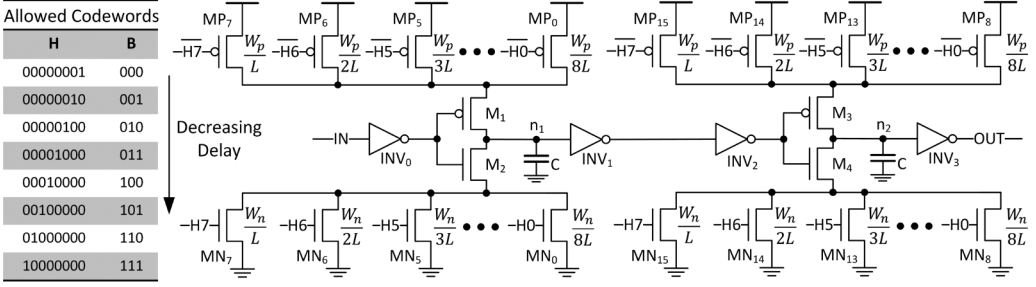


Fig. 5. Proposed one-hot DCCS-DE.

This imposes the constraint that for a particular codeword, only one of the current source transistors ( $MN_x$  and  $MP_x$ ) is ON. It is worth mentioning that due to the large capacitance of the internal nodes compared to  $n_1$  and  $n_2$ , the voltages of the virtual rails experience only a negligible disturbance when  $Mx$  switch, guaranteeing that  $MN_x$  and  $MP_x$  remain in the linear operating region. In moving from a binary to one-hot code, the lengths of the current source transistors were altered to increase linearly ( $1L, 2L, 3L, 4L, \dots, nL$ ) instead of exponentially ( $1L, 2L, 4L, 8L, \dots, 2^{(n-1)}L$ ), where  $n$  is the number of current source transistors, chosen on the basis of the amount of delay needed. The preceding changes ensure a constant delay difference between any two adjacent codewords, thereby minimizing DQE. This can be demonstrated mathematically as follows:

$$t_{pd} \propto C_L \frac{V_{ds}}{I_{ds}} \text{ and } I_{ds} \propto \frac{1}{R_{ds}}, \quad (3)$$

with

$$R_{ds} \propto L \Rightarrow t_{pd} \propto L. \quad (4)$$

Thus, as  $L$  increases linearly for different codewords, the delay also increases linearly. This is different from the binary scheme used in previous works, where multiple parallel current source transistors could simultaneously be active. This implied summing currents together, which produced a nonlinear relation between the total current and the codewords, and hence resulted in a nonlinear delay behavior in previously proposed DEs. Using this simple first-order analysis of (3) and (4), the expected calculated delay and DQE are plotted in Figures 6 and 7 for the binary DCCS-DE and the proposed one-hot DCCS-DE, respectively. The expected values are shown along with the simulated values in 28nm FD-SOI. These results demonstrate our claim about the large DQE associated with the binary DCCS-DE, and that using the proposed architecture in Figure 5 solves it successfully.

Regarding the MUX-DE, the design proposed in Heck et al. [2015] uses a sum-of-products MUX implementation. This has an intrinsically linear delay behavior, because changes in the number of cascaded inverters from one codeword to the next are constant, ensuring a low DQE. Note that the MUX-DE still utilizes a binary codeword, as opposed to the one-hot scheme employed in the proposed DCCS-DE design.

### 3.2. Proposed Fine-Grain Tuning DE Design

Body biasing is conventionally used either to reduce power consumption or to provide a performance boost. This article explores a new use of body biasing, which is to provide fine-grain delay control. Accordingly, we focus on RBB, as increasing the threshold voltage enables not only increasing the delay of transistors but also reducing their leakage power, a side benefit to our techniques. But RBB is not applied to all transistors



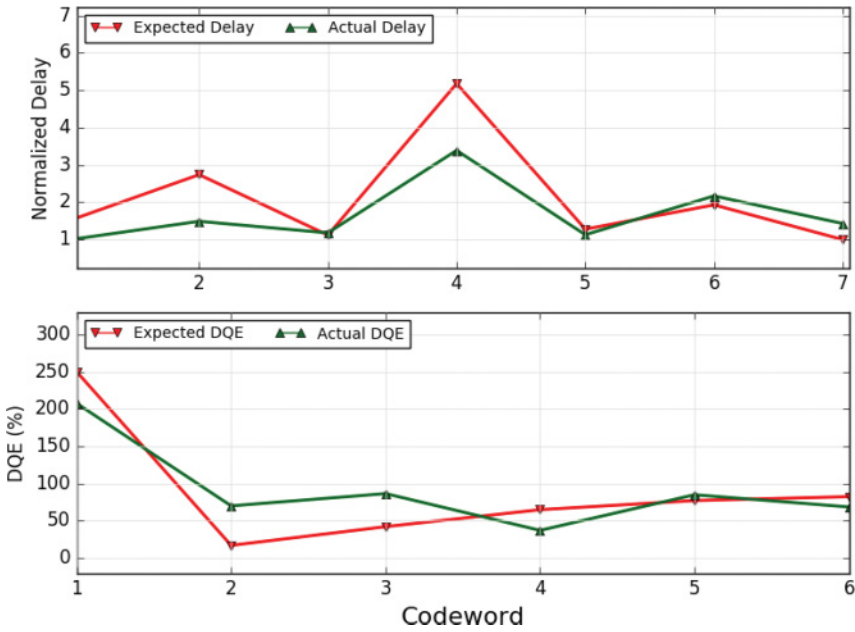


Fig. 6. Expected calculated delay and DQE versus the actual simulated values for the original DCCS-DE in Figure 3.




Fig. 7. Expected calculated delay and DQE versus the actual simulated values for the proposed one-hot DCCS-DE in Figure 5.




Fig. 8. DSIs with RBB applied to both pMOS and nMOS (a), only pMOS (b), and only nMOS (c).

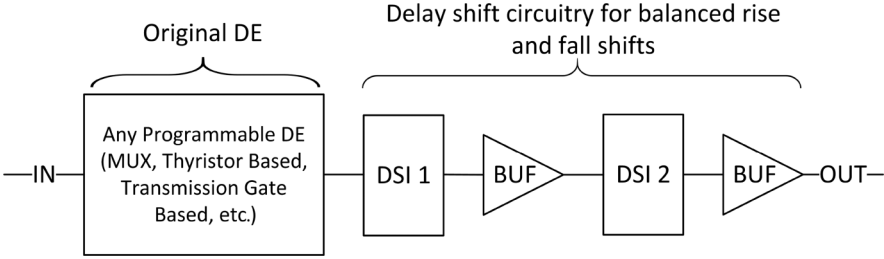


Fig. 9. Proposed architecture of fine-grain DEs.

of the DE, as each of them would get affected differently depending on its size. Moreover, this adds to the complexity of the design, and hence the delay characteristic can change significantly. In addition, this would increase the load that the bias voltage generating circuitry has to drive, resulting in more power consumption. Therefore, instead of employing RBB to each separate transistor, we propose the use of DSIs, as shown in Figure 8. The DSI is a conventional CMOS inverter with a programmable back body voltage that adjusts the threshold voltage of the inverter transistors, altering the current flowing through the inverter and changing its delay. Under normal operating conditions, the back gate of the inverter pMOS transistor is connected to the core supply, whereas the back gate of the nMOS is connected to ground. As illustrated in Figure 8, depending on availability, additional body biasing voltages can be applied to (a) both pMOS and nMOS, (b) only pMOS, or (c) only nMOS transistors. The delay shift provided by a DSI depends on two factors: (i) the change in the back body voltage and (ii) transistor sizes. The number of delay shifts can be increased by additional body biasing voltages or by using differently sized DSIs. Section 4 explores this further.

DSIs can be easily incorporated into any existing DE architecture, as Figure 9 illustrates. The intrinsic rise and fall delay characteristic of the original DE can also be maintained by cascading two DSIs in series, as shown in Figure 9, with buffers used to provide identical loads to both DSIs. The novelty of using the DSI is thus threefold: (i) it does not alter the original delay characteristics, (ii) it leads to less overhead in terms of area as compared to replicating the DE architectures to increase the delay range, and (iii) it can be applied to any DE architecture. Hence, it serves as a good candidate to cope with the problems of using hybrid DE architectures to achieve precise and fine-grain delays. Moreover, applying body biasing to specific inverters is well suited to the proposed DCCS-DE, because the biasing can be directly applied to the existing signal conditioning inverters (INV0 and INV2) of the DCCS-DE design (see Figure 5) instead of using additional area- and power-expensive DSI blocks.

Despite its advantages, a complication for the DSI is the generation and control of voltages from a domain other than the core supply and ground. As mentioned in Section 2.3, the RBB range is 1 to 3 V and 0 to -3 V in pMOS and nMOS, respectively. Even though DC-DC converter design is a rich field due to the importance of power

management, most architectures are optimized for load current, power efficiency, and bucking the voltage, not boosting it, all of which are not important when providing a voltage to the back gate of MOS transistors for RBB. Therefore, it is impractical to use bulky inductor-based converters, even the ones optimized for low power as in Choi et al. [2007] or McShane and Shenai [2001]. Switched capacitor-based DC-DC converters may seem more convenient to use for low-power applications [Kwong et al. 2009], which for boosting are usually called *charge pumps* (CPs) [Palumbo and Pappalardo 2010]. Level shifters seem also handy to translate the control to the target voltage domain [Tran et al. 2005]. Other architectures are also available, like the voltage dithering approach in Calhoun and Chandrakasan [2005] or the voltage down converter (VDC) with no passive components from Lai and Lee [2007] and Lee et al. [2007], but using them has their own disadvantages. DACs may be one of the obvious solutions, but they are too expensive and complex for small low-power applications.

Thus, the two most practical solutions are voltage CPs and level shifters. CPs are circuits that use a complex arrangement of switching capacitors to pump charges up and generate higher voltage levels using only the input supply. The output voltage value can be controlled by the number of pumping stages and the voltage gain obtained per stage, which mainly depends on the ratio between the capacitors and the parasitic capacitances [Dickson 1976]. Many architectures have been proposed to reduce the power consumption and enhance the functionality of CPs, including the charge transfer switches of Wu and Chang [1998] and the voltage doubler of Phang and Johns [2001]. In particular, Ker et al. [2006] achieved good performance while solving all of the common CP problems: (i) the threshold drop affecting the stage gain, (ii) the dynamic switching without any additional or external control, and (iii) the gate oxide reliability that was solved by keeping the voltage drop across any transistor within the supply range during operation. However, from our application's perspective, even the circuit in Ker et al. [2006] still suffers from the fundamental problems of CPs: (i) a significant amount of area due to the need for capacitors, (ii) significant power consumption due to the need for always switching clocks, (iii) considerable delay due to the time required to pump charges through several stages of capacitors, and (iv) output voltage that suffers from ripples [Palumbo and Pappalardo 2010]. The main advantage of CPs is that they do not need any reference voltages, and even the negative values for RBBing the nMOS transistors can be obtained if the circuit is connected in a discharging configuration.

On the other hand, the level shifter is a simple circuit that shifts an input signal from its voltage domain to the provided reference domain, and it is commonly used to interface off-chip and on-chip voltage domains. They are basically small, simple, and fast, and they do not use any passive components. There are not a variety of level shifters that can be employed to the application under investigation, but two main issues should be pointed out: (i) level shifters architecturally need a reference voltage, and (ii) gate oxide breakdown should be considered carefully. Pan et al. [2003] proposed a stacking architecture to solve the gate oxide breakdown issue at the expense of larger delay response. To solve issue (i), our target DSI employs RBB only in the pMOS transistors (Figure 8(b)) driven by the already available 1.8V I/O voltage for the 28nm FD-SOI technology. This solution was also used in Yamaoka et al. [2006] for an SRAM design in FD-SOI technology. To summarize, we propose using level shifters to actively switch the pMOS back body from the normal supply  $V_{DD} = 1V$  to the I/O voltage  $V_{high} = 1.8V$ .

Figure 10 shows the used level shifter, which is the CMLS proposed in Tran et al. [2005]. A similar circuit was used in Hamon and Beigné [2013] for body biasing the LVT (flip-well) devices in FD-SOI technology. However, since the devices used in the design of this work's DEs are RVT (normal well), the low voltage connected to the CMOS inverter is changed to  $V_{DD}$  instead of ground as in Hamon and Beigné [2013].






Fig. 11. Leakage versus delay trade-off in gate length biasing.

must be done conservatively, ensuring that the CSIs do not dominate the current source transistors by constraining the maximum current that can flow. This leaves two options, both implying the increase of the output capacitance. An added advantage of using parameters (iii) and (iv) is that these help to mitigate the charge sharing problem present in the DCCS-DE of Maymandi-Nejad and Sachdev [2003].

From experiments, we conclude that increasing the input capacitance of the signal conditioning inverters yields the largest energy overhead. This is because increasing the input capacitance for these results in a slow slew rate, which in turn generates more short circuit current through the inverters, and hence leads to larger EPT. Thus, adding an external shunt capacitance at the output node is the preferred approach. In fact, an optimal combination of (i) and (iii) helps to achieve the best energy-delay trade-off.

As for the MUX-DE, its delay range depends on two factors: (i) the number of cascaded inverters and (ii) the  $W$  and  $L$  of these. The optimization of EPT for the MUX-DE is better done using approach (ii)—that is, increasing the length of the nMOS and pMOS transistors to meet the desired delay range rather than adding more cascaded inverters. Approach (i) is used only after reaching the maximum allowable  $L$  for a transistor, because a larger  $L$  would result in less current flowing through the inverters, thus increasing delay. Cascading smaller inverters results in additive current flowing through the DE, consuming more EPT.

### 3.4. Leakage Reduction

Gate length biasing [Lazzari et al. 2009; Gupta et al. 2006] is a promising technique for achieving substantial leakage reduction and also requires no additional process steps. It involves increasing the length of the transistors to reduce leakage at the cost of a slight delay increase. Gupta et al. [2006] suggested a 10% upper bound on the increase in length to achieve the best trade-off for a bulk 130nm process. This limitation relies on consideration of timing constraints in critical paths. In the current work, experiments were run on an inverter in a bulk 65nm process and in the 28nm FD-SOI process to decide on a bound. The trade-off can be seen in Figure 11, with greater reduction of leakage in the 28nm FD-SOI process as compared to the 65nm bulk CMOS technology at the expense of an increase in delay.

An important observation is that the leakage versus delay trade-off in DEs is not subject to the same constraints as critical paths, as the delay range can always be tuned using other parameters like an external shunt capacitance or changing the number of stages used to build the DE. Thus, this work advocates overlooking the 10% limitation

Table I. Trade-Offs Between DEs for a 400ps Range

| DE                              | Original Binary DCCS-DE | Proposed One-Hot DCCS-DE | MUX-DE |
|---------------------------------|-------------------------|--------------------------|--------|
| DQE (%)                         | 269.92                  | 18.94                    | 3.19   |
| Avg. EPT (fJ)                   | 0.73                    | 1.01                     | 2.71   |
| Avg. Idle Power (nW)            | 0.30                    | 0.12                     | 0.28   |
| Active Area ( $\mu\text{m}^2$ ) | 1.18                    | 1.96                     | 0.35   |

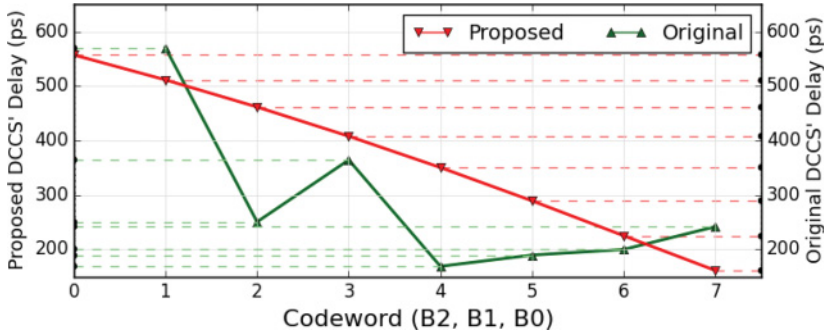


Fig. 12. Comparison of delay characteristics for the proposed DCCS-DE and the original DCCS-DE.

suggested in Lazzari et al. [2009], since it can be observed from Figure 11 that after roughly 40% increase in  $L$ , one obtains high leakage reduction, after which leakage reduction stagnates. While designing any of the DEs mentioned here, the minimum  $L$  chosen is 40% greater than the technology's smallest  $L$ . Experiments on the DCCS-DE and MUX-DE showed the same trend with substantial leakage reduction.

#### 4. EXPERIMENTS AND DISCUSSION

A 28nm FD-SOI CMOS technology with 1V core supply and 1.8V I/O supply was used for the DEs. All simulations employed the Cadence Spectre simulator with the same environment across all designs for fair comparisons. Unless otherwise stated, simulations employ an operating temperature of 27°C and use transistors in the typical process corner. The DEs employed the techniques proposed in Section 3.2, as well as the power reduction optimizations of Section 3.3. Each DE was designed to have eight different delay settings and provide an identical delay range at nominal conditions.

##### 4.1. Performance, Energy, and Area Trade-Offs

Table I summarizes the trade-offs between the DEs regarding performance, energy, and area. The techniques from Section 3.1 significantly improve the delay characteristic of the proposed one-hot DCCS-DE over the DCCS-DEs proposed in Maymandi-Nejad and Sachdev [2003] and Heck et al. [2015]. Improvement is quantified using the definition of DQE in Equation (2). As Figure 12 shows, the original DCCS-DE has a nonmonotonic delay, which is problematic, as certain codewords might provide delays that are too close or too far from that of their neighbor codeword. This characteristic translates into a large DQE of 269.92%, making the original DCCS-DE unreliable for building a programmable DE. Note that the DQE for the original DCCS-DE was calculated after reordering the codewords to provide a monotonically increasing delay characteristic; still, it presented high DQE. On the other hand, the proposed DCCS-DE has an almost linear delay characteristic, with nearly uniform delay difference between codewords, and does not require codeword reordering. This uniform delay difference enabled a much smaller DQE of 18.94%. Moreover, this DQE improvement comes without




Fig. 13. Comparison of proposed DCCS-DE and MUX-DE: EPT (a), *Energy/Delay* (b), and leakage (c).

significant power overhead. The active area values in Table I are the sum of the  $W \cdot L$  of all transistors in the design. Furthermore, to make the comparison between the proposed and original DEs more conservative, any area and power overheads of the circuitry needed to reorder the original DE codewords are not considered when presenting results for the original binary DCCS-DE.

As Table I shows, the MUX-DE displays a better DQE of 3.19% as opposed to 18.94% of the proposed DCCS. This is due to the fact that the technique presented in Section 3.1 does not take nonidealities into account. For the 400ps delay range programmed using eight codewords, this translates to an absolute error of 1.82ps, whereas the proposed DCCS-DE has a max deviation of 10.08ps from the ideal characteristic of having a uniform delay difference of  $400/7$  ps between adjacent codewords. However, with the aforementioned technique as a basis, the DQE achieved by the DCCS-DE can be improved by iteratively adjusting the  $L$ 's of those current source transistors that contribute to the larger DQE. Moreover, as discussed later in this section, the proposed fine-graining technique further improves the DQE and can also alleviate any issues arising due to minor deviations from the ideal characteristic. On the other hand, the proposed DCCS-DE still consumes 2.68 times less energy than the MUX-DE.

The metric used for comparing energy efficiency is the average EPT for all codewords measured for a particular delay range. As Figure 13(a) shows, the MUX-DE consumes nearly five times more energy than the proposed DCCS-DE for small delay ranges due to more current being drawn by the cascaded inverters in the MUX-DE than the CSIs of the DCCS-DE. The disparity decreases as  $L$  of the cascaded inverters increases to improve the delay range of the MUX-DE, with the energy advantage of the proposed DCCS-DE reducing by a factor of two for ranges larger than 2ns.

To better understand the relationship between delay range and EPT, Figure 13(b) shows the *Energy/Delay* relation between DEs. The results are consistent with the preceding discussion, because for delay ranges bigger than 2ns, the energy spent per unit of delay becomes nearly equal for both. Next, the DEs idle power is compared. Leakage reduction is achieved for both DEs using the gate length biasing strategy from Section 3.4. As can be seen from Figure 13(c), the DCCS-DE has a very low leakage

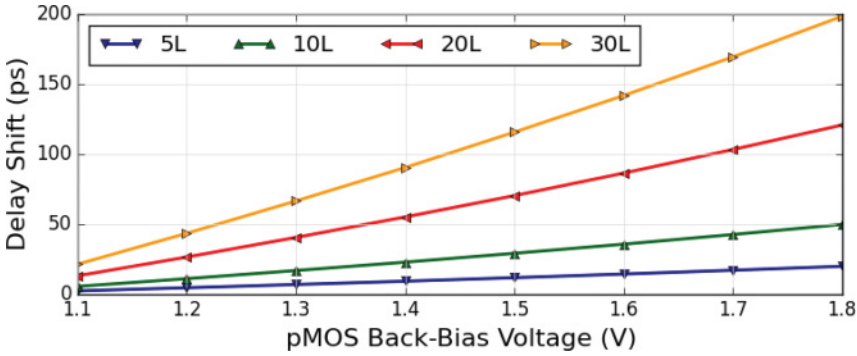


Fig. 14. Effect of L and body biasing voltages on delay shift.

power consumption of 0.12nW, which remains constant across delay ranges due to the fact that the extended delay ranges were met using external shunt capacitors rather than more transistors. On the other hand, it was noticed that the MUX-DE has substantially higher leakage power consumption when compared to the DCCS-DE. This is attributable to the large transistor count of the MUX-DE compared to the DCCS-DE.

The next set of experiments targets enabling a fine-grain delay range of 400ps for the DCCS-DE and MUX-DE. In other words, the idea is to reduce the delay difference between two adjacent delay settings. As mentioned in Section 3.2, the amount of delay shift achieved by the DSI shown in Figure 8 is controlled by both the size and the magnitude of the additional body biasing voltage. Experiments were run to determine the optimal sizing and voltage. As Figure 14 shows, the amount of delay shift increases as the body biasing voltage or the length of the transistor increases. Depending on the delay range and on the application, the appropriate number and magnitude of body biasing voltages and transistor sizes can be chosen.

For the reasons elaborated on in Section 3.2, only an additional body biasing voltage of 1.8V is generated, using the CMLS shown in Figure 10 to add an extra dimension of programmability. Moreover, for this application, only the length of the DSI transistors was increased, as it would be the more energy-efficient solution. For a 400ps delay range, across eight codewords, the normal delay difference between each codeword would be  $400/7 \text{ ps} \approx 57\text{ps}$ . Thus, the length of the transistor chosen is one that corresponds to a delay shift of  $400/14 \text{ ps} \approx 29\text{ps}$  for each codeword, in this case  $6.3 \times$  the minimum L. The final fine-grain delay characteristic for the DCCS-DE can be seen in Figure 15. Similar results were also observed for the MUX-DE. As the figure shows, the addition of a single body biasing voltage level doubles the resolution of the discrete delays offered by the DE. In the preceding experiment, a delay step of  $\approx 29\text{ps}$  is achieved as one moves from one setting to the next. Additional body biasing voltages can be used to further reduce the delay step size and make the achieved delay granularity finer.

To study the effect of using the DSI on the DQE, experiments were conducted on the DCCS-DE and MUX-DE. To ensure a fair comparison, the fine-grain structure was implemented in two flavors: one with the DSI and another without it. The two finer-grain DCCS-DE designs were implemented by (a) using 16 current source transistors sized as (1L, 1.5L, 2L, ..., 7.5L, 8L) instead of the original 8 sized as (1L, 2L, ..., 7L, 8L), and (b) body biasing the signal conditioning inverters INV0 and INV2 of Figure 5 while still using only 8 current source transistors. The MUX-DE design was reimplemented to have 16 codewords instead of the original 8 (Figure 2) to get a fair comparison and






Fig. 15. Fine-grain delay characteristic.

Table II. Trade-Offs Between Fine-Grain DEs for a 400ps Range

| DE                              | One-Hot DCCS-DE With 16 Current Sources | One-Hot DCCS-DE With 8 Current Sources + DSI | MUX-DE With 15 Buffers |
|---------------------------------|---|--|------------------------|
| DQE (%)                         | 26.81                                   | 12.57  | 4.55                   |
| Avg. EPT (fJ)                   | 1.03                                    | 1.57   | 5.08                   |
| Avg. Idle Power (nW)            | 0.22                                    | 0.16   | 0.40                   |
| Active Area ( $\mu\text{m}^2$ ) | 3.72                                    | 3.42   | 0.28                   |

Table III. PVT Analysis of DE (%) for the Compared DEs

|               | Process Variations |        |        | VDD Variations |        |        | Temperature Variations |        |        |
|---------------|--------------------|--------|--------|----------------|--------|--------|------------------------|--------|--------|
|               | FF                 | TT     | SS     | 1.2V           | 1V     | 0.8V   | 125°C                  | 25°C   | -50°C  |
| MUX-DE        | 3.24               | 3.22   | 3.30   | 4.90           | 3.22   | 1.23   | 3.30                   | 3.22   | 3.37   |
| Orig. DCCS-DE | 256.00             | 259.47 | 262.45 | 237.89         | 259.47 | 283.12 | 252.57                 | 259.47 | 263.95 |
| Prop. DCCS-DE | 20.85              | 18.94  | 17.24  | 26.54          | 18.94  | 12.13  | 23.17                  | 18.94  | 15.56  |

Note: "Orig." stands for the original binary, and "Prop." stands for the proposed one-hot.

to reduce the delay difference between adjacent codewords. The comparison of these designs appears in Table II.

As Table II shows, using the DSI with the DCCS-DE enables a DQE of 12.57%, which is less than half of what was achieved using additional current source transistors. Compared to the DCCS-DE with 16 current source transistors, the one with the DSI has lower area and also consumes 3.23 times less energy than the fine-grain MUX-DE, further improving the energy efficiency of the DCCS-DE over the MUX-DE that was observed in Table I. Thus, adding the DSI to the DCCS not only presents a better DQE but also does not result in excessive overhead.

#### 4.2. PVT Variations

To further analyze the trade-offs between the DEs discussed in this article, we conducted PVT analyses. We used three foundry-provided global process corners (slow (SS), typical (TT), and fast (FF)), three voltages, and three temperatures, as well as Monte Carlo analysis [Cadence 2016] using foundry-provided local variations models that capture threshold and transistor-size mismatches due to process variations. The PVT analysis in Tables III and IV compares the binary and the one-hot DCCS-DE with a MUX-DE with the same resolution. The designs are the same evaluated in Section 4.1, and the experiments extend the results from Table I.

Table IV. PVT Analysis of Worst-Case Rise/Fall Mismatch (%) for the Compared DEs

|               | Process Variations |       |       | VDD Variations |       |       | Temperature Variations |       |       |
|---------------|--------------------|-------|-------|----------------|-------|-------|------------------------|-------|-------|
|               | FF                 | TT    | SS    | 1.2V           | 1V    | 0.8V  | 125°C                  | 25°C  | -50°C |
| MUX-DE        | 14.76              | 15.65 | 16.46 | 12.55          | 15.64 | 20.97 | 10.83                  | 15.64 | 20.52 |
| Orig. DCCS-DE | 5.71               | 6.2   | 6.7   | 2.72           | 6.19  | 15.07 | 4.11                   | 6.19  | 8.7   |
| Prop. DCCS-DE | 7.07               | 7.42  | 7.82  | 5              | 7.43  | 12.1  | 5.43                   | 7.43  | 9.96  |

Note: “Orig.” stands for the original binary, and “Prop.” stands for the proposed one-hot.

First, we explored the impact of PVT variations on DQE (see Table III). The results show that global process variations have little impact on the DQE of the MUX-DE and binary DCCS-DE. However, faster transistors do increase the DQE of the proposed one-hot DCCS-DE by 21% comparing SS to FF corners. Similarly, temperature variations have little impact on the DQE of the MUX-DE and binary DCCS-DE, but higher temperatures increase the DQE of the proposed one-hot DCCS-DE by 49% comparing  $-50^{\circ}\text{C}$  to  $125^{\circ}\text{C}$ .<sup>2</sup> These results suggest that the assumed linear relationship between gate length and propagation delay in the one-hot DCCS-DE actually changes significantly as we vary process and temperature. Higher values of VDD increase the DQE of the MUX-DE and one-hot DCCS-DE but have the opposite effect on the binary DCCS-DE. In this way, results indicate that for the one-hot DCCS-DE, the desired linear relationship between gate length and delay is more accurate when the design is slower (low VDD, low temperature, and SS corner).

Results for the binary DCCS-DE design, on the other hand, can be explained by a distinct phenomenon. In particular, the desired linear relationship between delay and codewords for the binary DCCS-DE may be improved when the design runs faster than typical conditions (high VDD, high temperature, and FF corner), because under these conditions, the nonlinear charge-sharing effects in the current sources of this design are minimized. As for the MUX-DE, we notice that the DQE remains under 5% irrespective of any PVT variations. This translates to at most a 3ps deviation from linearity across codewords for the MUX-DE designed for a resolution of 400ps. As a reference, 3ps is less than half the propagation delay of a minimum-size inverter in 28nm FD-SOI technology. For this reason, the trends seen in DQE for the MUX-DE may be due to harder-to-analyze second-order effects such as DIBL, mobility degradation, velocity saturation, or other short/narrow channel effects, or reverse short/narrow channel effects [Tsividis 1999; Pelgrom et al. 1989; Rozeau et al. 2012; Yang et al. 2010].

Next, we explored the impact of PVT variations on worst-case rise/fall mismatch (see Table IV). Global process variations have little impact on the worst-case rise/fall mismatch of the compared DEs, but cooler temperatures and lower VDDs consistently increase the rise/fall mismatch of all three designs. This is because global process variations affect nMOS and pMOS in the same way, resulting in the rise/fall mismatch to remain constant across different process corners. On the other hand, temperature affects both nMOS and pMOS transistors differently, which further exacerbates the original mismatch in rise/fall caused by imbalances that arise as the signal travels through internal nodes of the DE. For example, the threshold voltage varies differently in nMOS and pMOS transistors, as shown in Shin et al. [2014], with the difference being more pronounced at lower temperatures and then becoming constant at higher temperatures, which is consistent with the trend seen in Table IV. The same explanation also holds true for changes in VDD, as a lower VDD worsens the difference in slew

<sup>2</sup>Note that in the FD-SOI technology used in this work, delay decreases as temperature increases, unlike what happens in conventional longer channels technologies [Shin et al. 2014; Sasaki et al. 2015]. This occurs because the higher temperatures cause a decrease in threshold voltage that more than compensates for the reduction in transistor transconductance.

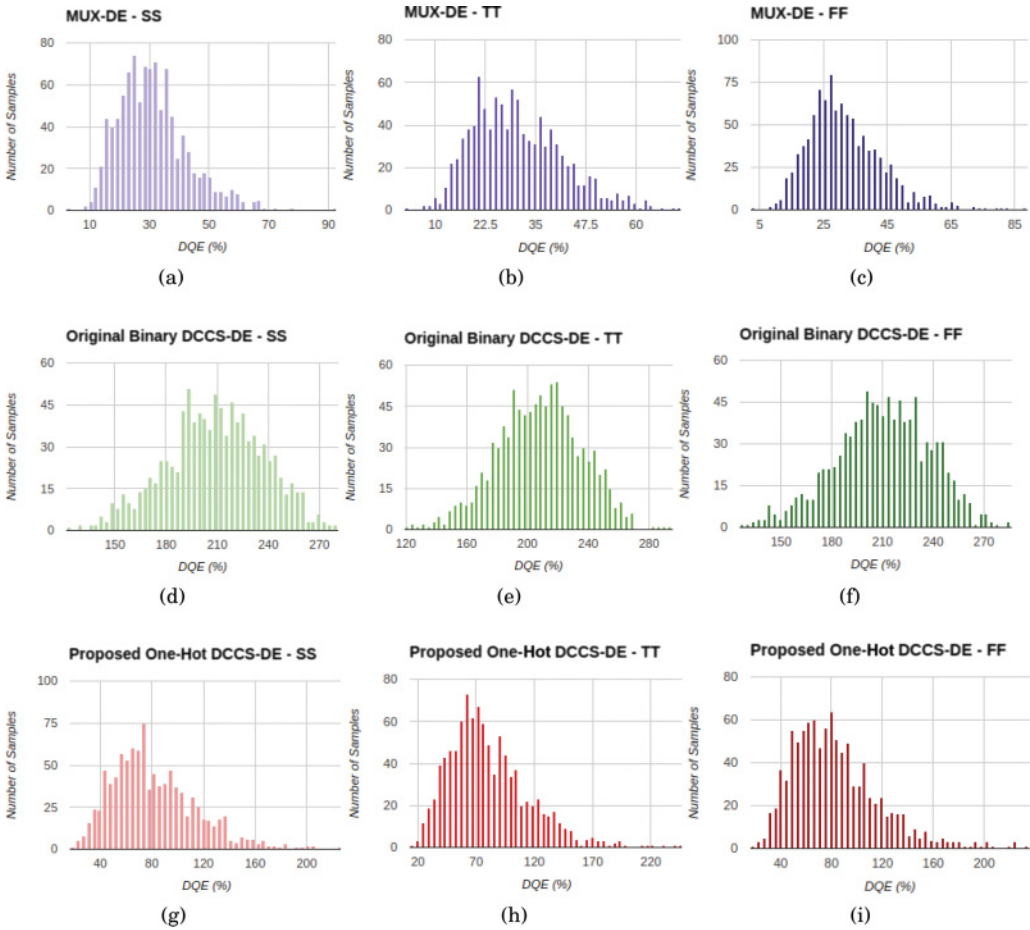


Fig. 16. Monte Carlo analysis of compared DEs.

Table V. Monte Carlo Process Variations Analysis on DQE

|                    |                          | DQE (%) | FF     | TT     | SS     |
|--------------------|--------------------------|---------|--------|--------|--------|
| Mean               | MUX-DE                   |         | 28.68  | 30.00  | 30.84  |
|                    | Original binary DCCS-DE  |         | 203.74 | 206.62 | 209.34 |
|                    | Proposed one-hot DCCS-DE |         | 80.77  | 78.97  | 84.14  |
| Standard Deviation | MUX-DE                   |         | 10.65  | 11.23  | 11.94  |
|                    | Original binary DCCS-DE  |         | 28.26  | 27.89  | 29.19  |
|                    | Proposed one-hot DCCS-DE |         | 32.09  | 34.97  | 34.70  |

rates of the signal (due to the increased propagation delay) across the various internal nodes of the DE, which in turn cascades into a worse rise/fall mismatch.

Monte Carlo analysis was performed on these three designs to show the effect of local process variations at the three different global corners, as illustrated in Figure 16. According to the charts, in all cases, local variations increase the average DQE significantly, with the binary DCCS-DE being impacted the most. As summarized in Table V, the binary DCCS-DE has a mean DQE of greater than 200% in all corners with a significant standard deviation—almost 30%. This means that even decreasing

Table VI. Monte Carlo Process Variations Analysis of the DSI

| DSI Delay (ps)     | FF    | TT    | SS    |
|--------------------|-------|-------|-------|
| Mean               | 22.05 | 25.16 | 29.12 |
| Standard deviation | 1.78  | 2.17  | 2.48  |

the codeword by two does not guarantee a delay increment. The one-hot DCCS-DE behaves significantly better but still can have greater than 100% DQE for some samples, compromising delay monotonicity. Only the MUX-DE with a mean DQE of 30% and a standard deviation of less than 12% guarantees monotonicity of delay in most chips. This observed difference in DQE may be due to the difference in impact of mismatch on differently sized transistors. In particular, larger transistors are effected by mismatch less than smaller transistors [Pelgrom et al. 1989]. The MUX-DE uses identically sized larger-than-minimum transistors, yielding smaller DQE, and the DCCS-DE use transistors of dramatically different sizes that are impacted by process variations differently, yielding higher DQE.

Finally, we analyzed the shift in delay of the DSI under process variations, which is illustrated in Table VI. Results show the expected trend in delay and a very low standard deviation. This means that changes in codewords that cause the DSI to become back body biased are quite reliable even under process variations. This supports the claim in Section 3.2 that the DSI can be applied independently of the DE architecture for constructing fine-grain resolution DEs.

## 5. CONCLUSIONS

This work presents and analyzes design modifications to the DCCS-DE. The proposed design has a lower DQE and is more robust to PVT variations than the previously discussed DCCS-DEs in Maymandi-Nejad and Sachdev [2003] and Heck et al. [2015]. Additionally, the proposed DCCS-DE is significantly more energy efficient than the current mirror-based design proposed in Maymandi-Nejad and Sachdev [2003, 2005]. Compared to the MUX-DE, it consumes less energy for delay ranges smaller than 2ns. However, this DE is less robust to local process variations, which makes adding a reliable test margin to these designs more challenging.

The article also proposes a generic DSI architecture that utilizes the body biasing feature in 28nm UTBB FD-SOI technology to obtain fine-grain delays in a single DE structure. Note that this feature allows the proposed architecture to be easily integrated into digital systems. Such advances enable leveraging the advantages of UTBB FD-SOI technologies for circuit design and allow better design space exploration for applications that need low-power DEs.

Our simulation results indicate that the MUX-DE is the only design that is robust enough against process variations to ensure delay monotonicity. This means that the other two DEs require larger delay test margins. For example, testing these with one delay setting and then shipping them with test margin to cover discrepancies caused by variations may require more margin than what can be ensured by simply decreasing the codeword by one. Furthermore, the results obtained for the impact of process variations in the characteristics of the DSI suggest that it provides a reliable way of adding test margin to DEs with moderate or high DQE. Namely, if we add the DSI to such a DE, program and test the DE with the DSI normally biased, and then reverse body bias the DSI to add test margin, the resulting margin will be reliable even under process variations.

## REFERENCES

- B. Arkin. 2004. Realizing a production ATE custom processor and timing IC containing 400 independent low-power and high-linearity timing verniers. In *Proceedings of the IEEE International Solid-State Circuits Conference (ISSCC'04)*. 348–349.
- P. A. Beerel, R. O. Ozdag, and M. Ferretti. 2010. *A Designer's Guide to Asynchronous VLSI*. Cambridge University Press.
- Cadence. 2016. Virtuoso Analog Design Environment Family. Available at <https://www.cadence.com>.
- B. H. Calhoun and A. P. Chandrakasan. 2005. Ultra-dynamic voltage scaling using sub-threshold operation and local voltage dithering in 90nm CMOS. In *Proceedings of the IEEE International Solid-State Circuits Conference (ISSCC'05)*. Vol. 1. 300–599.
- A. Chakraborty, K. Duraisami, A. Sathanur, P. Sithambaram, L. Benini, A. Macii, E. Macii, and M. Poncino. 2008. Dynamic thermal clock skew compensation using tunable delay buffers. *IEEE Transactions on VLSI Systems* 16, 6, 639–649.
- Y. Choi, N. Chang, and T. Kim. 2007. DC-DC converter-aware power management for low-power embedded systems. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 26, 8, 1367–1381.
- R. Diamant, R. Ginosar, and C. Sotiriou. 2015. Asynchronous sub-threshold ultra-low power processor. In *Proceedings of the International Workshop on Power and Timing Modeling, Optimization, and Simulation (PATMOS'15)*. 89–96.
- J. F. Dickson. 1976. On-chip high-voltage generation in MNOS integrated circuits using an improved voltage multiplier technique. *IEEE Journal of Solid-State Circuits* 11, 3, 374–378.
- T. Dogsa, M. Solar, and B. Jarc. 2014. Precision delay circuit for analog quadrature signals in Sin/Cos encoders. *IEEE Transactions on Instrumentation and Measurement* 63, 12, 2795–2803.
- P. Flatresse, B. Giraud, J. Noel, B. Pelloux-Prayer, F. Giner, D. Arora, F. Arnaud, et al. 2013. Ultra-wide body-bias range LDPC decoder in 28nm UTBB FDSOI technology. In *Proceedings of the IEEE International Solid-State Circuits Conference (ISSCC'13)*. 424–425.
- P. Gupta, A. B. Kahng, P. Sharma, and D. Sylvester. 2006. Gate-length biasing for runtime-leakage control. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 25, 8, 1475–1485.
- J. Hamon and E. Beigné. 2013. Automatic leakage control for wide range performance QDI asynchronous circuits in FD-SOI technology. In *Proceedings of the IEEE International Symposium on Asynchronous Circuits and Systems (ASYNC'13)*. 142–149.
- D. Hand, H. Huang, B. Cheng, Y. Zhang, M. T. Moreira, M. Breuer, N. L. V. Calazans, and P. A. Beerel. 2015. Performance optimization and analysis of blade designs under delay variability. In *Proceedings of the IEEE International Symposium on Asynchronous Circuits and Systems (ASYNC'15)*. 61–68.
- G. Heck, L. S. Heck, A. Singhvi, M. T. Moreira, P. A. Beerel, and N. L. V. Calazans. 2015. Analysis and optimization of programmable delay elements for 2-phase bundled-data circuits. In *Proceedings of the International Conference on VLSI Design (VLSID'15)*. 321–326.
- T. B. Hook. 2012. Fully depleted devices for designers: FDSOI and FinFETs. In *Proceedings of the IEEE Custom Integrated Circuits Conference (CICC'12)*. 7.
- Y. J. Jung, S. W. Lee, D. Shim, W. Kim, C. Kim, and S. I. Cho. 2001. A dual-loop delay-locked loop using multiple voltage-controlled delay lines. *IEEE Journal of Solid State Circuits* 36, 5, 784–791.
- M. D. Ker, S. L. Chen, and C. S. Tsai. 2006. Design of charge pump circuit with consideration of gate-oxide reliability in low-voltage CMOS processes. *IEEE Journal of Solid-State Circuits* 41, 5, 1100–1107.
- G. Kim, M. K. Kim, B. S. Chang, and W. Kim. 1996. A low-voltage, low-power CMOS delay element. *IEEE Journal of Solid-State Circuits* 31, 7, 966–971.
- J. Kwong, Y. K. Ramadass, N. Verma, and A. P. Chandrakasan. 2009. A 65 nm sub- $V_t$  microcontroller with integrated SRAM and switched capacitor DC-DC converter. *IEEE Journal of Solid-State Circuits* 44, 1, 115–126.
- F. S. Lai and C. F. Lee. 2007. On-chip voltage down converter to improve SRAM read/write margin and static power for sub-nano CMOS technology. *IEEE Journal of Solid-State Circuits* 42, 9, 2061–2070.
- C. Lazzari, A. Ziesemer, and R. A. L. Reis. 2009. An automated design methodology for layout generation targeting power leakage minimization. In *Proceedings of the IEEE International Conference on Electronics, Circuits, and Systems (ICECS'09)*. 81–84.
- C. F. Lee, W. Lin, F. S. Lai, and S. C. Lin. 2007. On-chip VDC circuit for SRAM power management. In *Proceedings of the International Symposium on VLSI Design, Automation, and Test (VLSI-DAT'07)*. 4.

- G. H. Li and H. P. Chou. 2007. A high resolution time-to-digital converter using two-level vernier delay line technique. In *Proceedings of the IEEE Nuclear Science Symposium Conference Record (NSS'07)*. 276–280.
- C. H. Lin, W. Haensch, P. Oldiges, H. Wang, R. Williams, J. Chang, M. Guillorn, et al. 2011. Modeling of width-quantization-induced variations in logic FinFETs for 22nm and beyond. In *Proceedings of the Symposium on VLSI Technology (VLSIT'11)*. 16–17.
- H. Lin and N. H. C. Chen. 2001. New four-phase generation circuits for low-voltage charge pumps. In *Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS'01)*. 504–507.
- N. R. Mahapatra, S. V. Garimella, and A. Tareen. 2000. An empirical and analytical comparison of delay elements and a new delay element design. In *Proceedings of the IEEE Computer Society Workshop on VLSI (VLSI'00)*. 81–86.
- N. R. Mahapatra, S. V. Garimella, and A. Tareen. 2002. Comparison and analysis of delay elements. In *Proceedings of the Midwest Symposium on Circuits and Systems (MWSCAS'02)*. 473–476.
- M. Maymandi-Nejad and M. Sachdev. 2003. A digitally programmable delay element: Design and analysis. *IEEE Transactions on VLSI Systems* 11, 5, 871–878.
- M. Maymandi-Nejad and M. Sachdev. 2005. A monotonic digitally controlled delay element. *IEEE Journal of Solid-State Circuits* 40, 11, 2212–2219.
- E. McShane and K. Shenai. 2001. A CMOS monolithic 5-MHz, 5-V, 250-mA, 56% efficiency DC/DC switch-mode boost converter with dynamic PWM for embedded power management. In *Proceedings of the IEEE Industry Application Society 36th Annual Meeting (IAS'01)*, Vol. 1. 653–657.
- B. M. Moon, Y. J. Park, and D. K. Jeong. 2008. Monotonic wide-range digitally controlled oscillator compensated for supply voltage variation. *IEEE Transactions on Circuits and Systems II: Express Briefs* 55, 10, 1036–1040.
- G. Palumbo and D. Pappalardo. 2010. Charge pump circuits: An overview on design strategies and topologies. *IEEE Circuits and Systems Magazine* 10, 1, 31–45.
- D. Pan, H. W. Li, and B. M. Wilamowski. 2003. A low voltage to high voltage level shifter circuit for MEMS application. In *Proceedings of the 15th Biennial University/Government/Industry Microelectronics Symposium (UGIM'03)*. 128–131.
- M. J. M. Pelgrom, A. C. J. Duinmaijer, and A. P. G. Welbers. 1989. Matching properties of MOS transistors. *IEEE Journal of Solid-State Circuits* 24, 5, 1433–1439.
- B. Pelloux-Prayer, A. Valentian, B. Giraud, Y. Thonnart, J. P. Noel, P. Flatresse, and E. Beigné. 2013. Fine grain multi-VT co-integration methodology in UTBB FD-SOI technology. In *Proceedings of the International Conference on Very Large Scale Integration (VLSI-SoC'13)*. 168–173.
- K. Phang and D. A. Johns. 2001. A 1V 1mW CMOS front-end with on-chip dynamic gate biasing for a 75Mb/s optical receiver. In *Proceedings of the IEEE International Solid-State Circuits Conference (ISSCC'01)*. 218–219.
- O. Rozeau, M. A. Jaud, T. Poiroux, and M. Benosman. 2012. *UTSOI Model 1.1.3*. Laboratoire d'électronique et de technologie de l'information (Leti).
- K. Ryu, D. H. Jung, and S. O. Jung. 2013. All-digital process-variation-calibrated timing generator for ATE with 1.95-ps resolution and a maximum 1.2-GHz test rate. In *Proceedings of the European Solid-State Circuits Conference (ESSCIRC'13)*. 41–44.
- K. R. A. Sasaki, J. A. Martino, M. Aoulaiche, E. Simoen, and C. Claeys. 2015. Enhanced dynamic threshold UTBB SOI at high temperature. In *Proceedings of the Joint International EUROSOCI Workshop and the International Conference on Ultimate Integration on Silicon (EUROSOCI-ULIS'15)*. 261–264.
- M. Shin, M. Shi, M. Mouis, A. Cros, E. Josse, G. T. Kim, and G. Ghibaudo. 2014. Low temperature characterization of 14nm FDSOI CMOS devices. In *Proceedings of the International Workshop on Low Temperature Electronics (WOLTE'14)*. 29–32.
- A. Singhvi, M. T. Moreira, R. Tadros, N. L. V. Calazans, and P. A. Beerel. 2015. A fine-grained, uniform, energy-efficient delay element for FD-SOI technologies. In *IEEE Computer Society Annual Symposium on VLSI (ISVLSI'15)*.
- C. Q. Tran, H. Kawaguchi, and T. Sakurai. 2005. Low-power high-speed level shifter design for block-level dynamic voltage scaling environment. In *Proceedings of the International Conference on Integrated Circuit Design and Technology (ICICDT'05)*. 229–232.
- Y. Tsidividis. 1999. *Operation and Modeling of the MOS Transistor*. Oxford University Press.
- J. T. Wu and K. L. Chang. 1998. MOS charge pumps for low-voltage operation. *IEEE Journal of Solid-State Circuits* 33, 4, 592–597.

- J. Xiong, V. Zolotov, C. Visweswariah, and P. A. Habitz. 2009. Optimal test margin computation for at-speed structural test. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 28, 9, 1414–1423.
- M. Yamaoka, R. Tsuchiya, and T. Kawahara. 2006. SRAM circuit with expanded operating margin and reduced stand-by leakage current using thin-BOX FD-SOI transistors. *IEEE Journal of Solid-State Circuits* 41, 11, 2366–2372.
- W. Yang, C. H. Lin, T. H. Morshed, D. Lu, A. Niknejad, and C. Hu. 2010. *BSIMSOIv4.4 MOSFET MODEL Users' Manual*. Regents of the University of California. Available at [http://www-device.eecs.berkeley.edu/bsim/Files/BSIMSOI/bsimsoi4p4/BSIMSOIv4.4\\_UsersManual.pdf](http://www-device.eecs.berkeley.edu/bsim/Files/BSIMSOI/bsimsoi4p4/BSIMSOIv4.4_UsersManual.pdf).

Received September 2015; revised March 2016; accepted May 2016