

Analysis of Mask-Based Nanowire Decoders

Eric Rachlin, *Student Member, IEEE*, and John E. Savage, *Life Fellow, IEEE*

Abstract—Stochastically assembled nanoscale architectures have the potential to achieve device densities 100 times greater than today's CMOS. A key challenge facing nanotechnologies is controlling parallel sets of nanowires (NWs), such as those in crossbars, using a moderate number of mesoscale wires. Three similar methods have been proposed to control NWs using a set of perpendicular mesoscale wires. The first is based on NW differentiation during manufacture, the second makes random connections between NWs and mesoscale wires, and the third, a mask-based approach, interposes high-K dielectric regions between NWs and mesoscale wires. Each of these addressing schemes involves a stochastic step in their implementation. In this paper, we analyze the mask-based approach and show that, when compared to the other two schemes, a large number of mesoscale control wires are necessary for its realization.

Index Terms—Emerging technologies, memory structures, stochastic processes.

1 INTRODUCTION

THE crossbar, a simple but well-known connection network, consists of two orthogonal sets of parallel wires (see Fig. 1). Switches are positioned at the crosspoints defined by the intersections of pairs of wires. Crossbars can be used as switching networks, memories, and programmed logic arrays.

Chemists have developed methods to assemble nanowires (NWs) into crossbars [1], [2], [3], [4]. They have realized switches by placing a thin layer of bistable molecules such as rotaxanes or [2]-catenanes between two orthogonal sets of NWs [5], [6], [7]. When a large positive or negative electric field is applied between two orthogonal NWs, the molecules at their crosspoint become either conducting or nonconducting. A smaller electric field can then measure the conductivity of a crosspoint without changing it.

A number of methods have been devised to produce NWs using vapor-liquid-solid (VLS) processes [8], [9], nanoimprinting [10], superlattice NW pattern transfer (SNAP) [2], and nanolithography [11], [3]. NWs produced through VLS can be differentiated. They can be grown with different electrical or chemical properties before being stochastically assembled into crossbars. NWs produced by the other three methods are undifferentiated.

Both differentiated and undifferentiated NWs must interface with the larger mesoscale technology. An important challenge is to control individual NWs with mesoscale wires (MWs) without losing the high crosspoint density NWs allow. This challenge can be met by 1) positioning MWs at right angles to the NWs and 2) using MWs to apply electric fields to lightly doped regions of

NWs. The application of an electric field to an exposed lightly doped region drives the conductance of that NW to be low (see Fig. 2). In other words, NWs combined with MWs form field-effect transistors (FETs). Turning on a subset of the MWs turns off some subset of the NWs.

An **NW decoder** is a circuit that addresses (leaves conducting) one NW (or a desired subset of NWs) by associating it with some subset of MWs. As explained in Section 2, the following three methods to control NWs with MWs have been proposed and analyzed: 1) grow differentiated NWs containing lightly doped regions and then place a random subset of the NWs on a chip using fluidic self-assembly [13], [14], 2) make random contacts between MWs and undifferentiated NWs [15], [16], and 3) randomly place lithographically defined high-K dielectric regions between MWs and undifferentiated lightly doped NWs [17]. The last one is analyzed here. Three other methods for controlling NWs have been proposed (see Section 2.4), but the effect of assembly variation on the amount of control circuitry they require has not been analyzed.

The number of MWs needed to control N NWs with high probability has been determined analytically for the first [18], [19], [20] and second methods [15], [20], [16]. Here, we analyze the third method, which is implemented using a randomized mask-based decoder. We show that, for very reasonable assumptions, it requires a large number of MWs.

Mask-based decoders (see Section 2) are designed to work with undifferentiated NWs produced by nanoimprinting [10] or the SNAP process [2], [11]. These decoders use lithographically defined mesoscale rectangular regions of high-K dielectric (we call these **LRs**) to allow each MW to make some subsets of the NWs nonconducting (see Fig. 3). Lithography, however, puts a lower limit on the size of such regions. The smallest regions must be randomly shifted to ensure that all pairs of adjacent NWs are controlled by different subsets of MWs. This makes it possible to address individual NWs but requires a large number of MWs, as we show.

1.1 Crossbar-Based Architectures

It is anticipated that crossbar-based architectures will be assembled modularly from many small crossbars [21], [22]

• The authors are with the Department of Computer Science, Brown University, 115 Waterman St., ROOM: CIT 503, Providence, RI 02912-1910. E-mail: {eerc, jes}@cs.brown.edu.

Manuscript received 14 Apr. 2006; revised 1 June 2007; accepted 25 June 2007; published online 13 Aug. 2007.

Recommended for acceptance by J. Lach.

For information on obtaining reprints of this article, please send e-mail to: tc@computer.org, and reference IEEECS Log Number TC-0146-0406.

Digital Object Identifier no. 10.1109/TC.2007.70795.

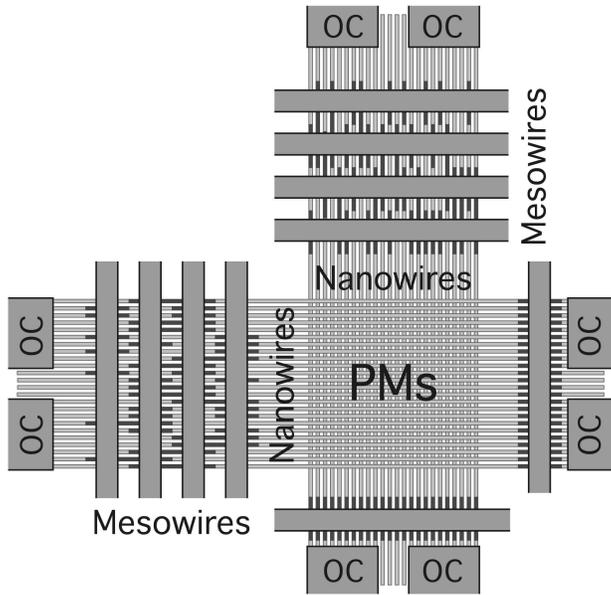


Fig. 1. A crossbar formed from two orthogonal sets of NWs with programmable molecules (PMs) at crosspoints defined by intersecting orthogonal NWs. NWs are segmented into contact groups connected to pairs of OCs. To activate an NW in one dimension, a contact group is activated and MWs are used to deactivate all but one NW in that group. Data is stored at a crosspoint by applying a large electric field across it. Data is sensed with a smaller field after disconnecting OCs using the two additional MWs.

(say, $1,000 \times 1,000$ NWs). Crossbars can be used not only as memories but also as programmable architectures [23]. In either case, manufacturing uncertainties will require that each crossbar be tested for storage capacity and addressability; some will be used and others will be ignored.

In the case of memories, one can imagine a grid of crossbars and associated control logic that only accesses individual crossbars whose storage capacity exceeds a threshold. The design of NW decoders will have a substantial impact on both the area occupied by the individual crossbars (if many MWs are needed to control NWs, the area of individual crossbars will be large) and the

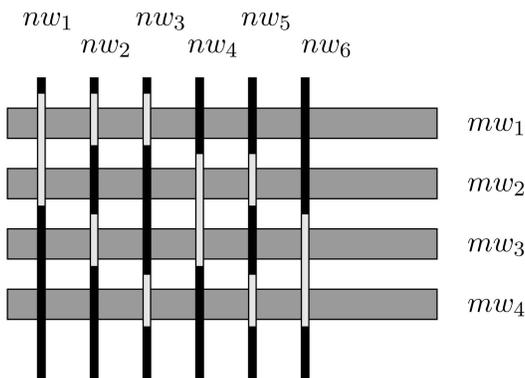


Fig. 2. A method for addressing six differentiated modulation-doped NWs $\{nw_1, \dots, nw_6\}$ with four large MWs $\{mw_1, \dots, mw_4\}$. The lightly doped regions of each NW are highlighted. An NW is nonconducting if it has a doped region adjacent to an MW carrying a high electric field. If exactly two of the four MWs carry a high field, then exactly one of the six NWs is conducting. This idea was developed in [12].

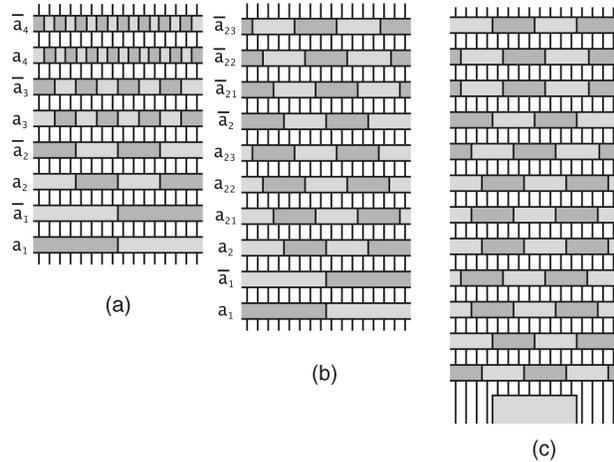


Fig. 3. Three mask-based decoders in which horizontal MWs lie across vertical NWs. The dark gray rectangles indicate the locations of lithographically defined dielectric regions (LRs) under the MWs. The regions under an MW determine which NWs become nonconducting when that MW carries an electric field. (a) A logarithmic mask-based decoder that uses $2 \log N$ MWs to select one of N NWs. (b) A hybrid mask-based decoder in which the rows of the smallest manufacturable LRs are shifted and repeated to provide control over individual NWs. (c) A randomized mask-based decoder in which small groups of NWs are connected to a OCs; one OC is shown.

likelihood that their storage capacity is above the desired threshold. A designer of nanoscale architectures will need to take into account the characteristics of the decoders of individual nanoscale crossbars. This motivates our investigation into the size requirements of mask-based decoders.

1.2 Overview of the Paper

In Section 2, we describe three related methods of addressing NWs with MWs, 1) “encoded-NW decoders,” 2) “randomized-contact decoders,” and 3) “mask-based decoders,” and briefly review three other methods whose control overhead has not been thoroughly analyzed under assumptions on variation in assembly. In Section 3, we model the control that MWs exert over NWs. We then focus specifically on masked-based decoders, giving a condition that LRs must satisfy in order for each NW to be addressed individually.

Methods for manufacturing undifferentiated NWs and mask-based decoders are described in Section 4. The limitations on photolithography that lead to uncertainties in the placement of LRs are examined and modeled probabilistically. This model is used in Section 5 to begin an analysis of the “randomized n -cycle mask-based decoder.” In this decoder, groups of adjacent NWs are connected to ohmic contacts (OCs). A mask-based decoder is then used to control individual NWs within each group. Our goal is to determine how many MWs are required to individually address each NW.

In Section 6, we present and analyze three models for the random placement of LRs that capture the variation in LR placement, as well as the choices that a designer has in transferring LRs to a chip using one or several masks. These models are abstracted into three generalizations of the standard coupon collector problem: 1) the “coupon collector problem with failures”, 2) the “targeted coupon collector

problem,” and 3) the “multistage targeted coupon collector problem.” In the standard coupon collector problem, one of C coupons is selected with probability $1/C$ at each trial. The problem is to determine how many independent trials are required to collect all coupons with high probability. Here, we consider coupon collection with nonuniform probability distributions.

Drawing on the results in Section 6, the performance of the randomized n -cycle mask-based decoder is summarized in Section 7. Section 8 discusses several practical considerations for designing mask-based decoders. Conclusions are drawn in Section 9.

2 ADDRESSING NWS WITH MWS

This section describes three related methods for addressing NWS with MWS, as well as three additional methods. Each of the first three methods assumes that NWS are divided into groups of approximately 10 adjacent NWS, called **contact groups**. Each contact group is connected to a separate OC. When an NW is addressed, a contact group is selected using standard CMOS circuitry; then, all but one NW within the group is made nonconducting by turning on some subset of the MWS. These three methods of addressing NWS with MWS introduce uncertainty with regard to which MWS address which NWS. Thus, each method requires programmable circuitry to map external binary addresses to subsets of MWS. The other three methods described in Section 2.4 will also require programmable circuitry due to assembly variation. However, the amount of such circuitry has not yet been quantified.

2.1 Differentiated NW Decoders

Lieber et al. have shown that differentiated NWS can be assembled into crossbars using fluidic methods [13], [14]. When NWS are manufactured using a VLS process, they can be grown with a pattern of lightly and heavily doped sections along their length [24], [25], [26], a process known as “modulation doping.” Many copies of differently patterned NWS are collected in a large ensemble and deposited on a chip. As a result, the NWS in each contact group have doping patterns selected at random from the larger ensemble. NWS with distinct doping patterns can be individually addressed with MWS, as described in Fig. 2. This **encoded-NW decoder** is analyzed in [18], [27], [19].

DeHon et al. show that all N NWS can be individually addressed more than 99 percent of the time using M MWS when $M \geq \lceil 2.2 \log_2 N \rceil + 11$ [18]. Three other addressing strategies are explored by Gojman et al. [27], [19]. These include 1) individually addressing half of the NWS in each contact group and 2) addressing each NW doping pattern in at least p contact groups. Their analysis indicates that these two strategies require less area than the strategy that requires all NWS within each contact group to be addressable [18]. Unfortunately, NWS encoded through modulation doping may suffer from axial misalignment when deposited on a chip. A new technique for encoding NWS through the use of shells of different types has also been proposed [28]. This approach has been shown to be competitive with modulation doping; it does not suffer from misalignment but may require slightly larger NWS.

2.2 Randomized-Contact NW Decoders

The **randomized-contact decoder** was proposed by Williams and Kuekes [15]. It controls N NWS with M MWS by making random contacts between them with a probability of $1/2$. They state that $M \geq 5 \log_2 N$ MWS suffice to provide unique addresses to all N NWS [15]. This method has been analyzed empirically and approximately by Hogg et al. [16], who show that the probability that all NWS in a contact group are controllable rises rapidly to near 1 as M increases from slightly less than $4 \log_2 N$ to slightly more than $6 \log_2 N$. They also explore the number of MWS needed when contacts are imperfect. Rachlin and Savage have done a mathematical analysis of this model and derived bounds on the number of MWS needed to ensure that all NWS are controllable with a probability of at least $1 - \epsilon$ [20].

2.3 Mask-Based NW Decoders

The third decoder [29], [17], called a **mask-based decoder**, places lithographically defined high-K dielectric regions in between MWS and lightly doped NWS. If an LR lies between a set of NWS and an MW, those NWS are made nonconducting by that MW when it carries a sufficiently strong electric field. Manufacturing constraints limit the precision with which LRs can be placed. These manufacturing constraints are described below.

An idealized **logarithmic mask-based decoder** is shown in Fig. 3a [30]. This decoder uses k pairs of MWS to control $N = 2^k$ NWS. For $1 \leq j \leq \log_2 N$, the two MWS in the j th pair each lie over a row of 2^{j-1} evenly spaced LRs. These two rows of LRs cover complementary halves of the NWS (see Fig. 3a). When a field is applied to one of the two MWS in a pair, exactly half of the NWS are turned off. Each pair of MWS turns off half of the NWS left on by the previous pair. This allows a logarithmic mask-based decoder to select exactly one NW to remain conducting when a field is applied to one MW in each of the k pairs of MWS. A logarithmic mask-based decoder thus assigns a unique address to each of the N undifferentiated NWS using $2 \log_2 N$ MWS.

Unfortunately, the logarithmic mask-based decoder is not feasible. It requires that LRs have lengths that are equal to the pitch of NWS and that the position of their boundaries be tightly controlled, which are characteristics that cannot be met with lithography. As an alternative, Beckman et al. have proposed the **hybrid decoder** [17] to cope with this uncertainty (see Fig. 3b). This decoder has linear and logarithmic portions.

The logarithmic portion is a logarithmic mask-based decoder that resolves the set of active NWS down to a small contact group of w NWS. In the linear portion of the hybrid decoder, the goal is to have one LR left boundary and one LR right boundary fall in the space between each pair of NWS. If each LR has a length that is exactly equal to w NW pitches, $2w$ rows of LRs would suffice to allow fields to be applied to MWS so that one NW in a set of w NWS is active and the rest are inactive, a condition derived in Section 3.1.

We refer to the $2w$ rows of evenly spaced LRs, where each row is offset by one NW pitch from the previous rows, as **one cycle**. Because this type of precision is not yet possible at the nanometer scale, multiple cycles are needed to ensure that, with high probability, both left and right LR

boundaries fall between pairs of NWs (see Fig. 3c). We refer to these cycles as a **randomized n -cycle linear decoder**.

Since the placement of LRs is difficult to control at the nanometer scale, it is more likely that the logarithmic portion of the mask-based decoder would be replaced by a conventional lithographic-scale decoder in which the contact groups of w NWs are connected to pairs of OCs and one contact group is activated at a time by this decoder. A randomized n -cycle linear decoder is then created for each contact group, all of which share the same set of MWs. We call this a **randomized n -cycle mask-based decoder** and analyze its performance in Section 7.

2.4 Other Proposed Decoder Technologies

Other methods of controlling NWs with MWs have been proposed. Each has uncertainties in their construction that have not been fully analyzed.

Ma et al. have proposed that NWs in an NW crossbar be controlled by MWs in an MW crossbar [31]. Pins with nanoscale diameter tips are formed at crosspoints in the MW crossbar and make contact with NWs in the NW crossbar. One crossbar is rotated by a small angle relative to the other such that pins make contact with both sets of parallel NWs. Pins must pass through gaps in one NW layer to make contact with NWs on the second layer. This method is sensitive to small changes in the angle of rotation between the two crossbars.

Di Spigna et al. [32] propose creating NW decoders in a parallel set of insulated NWs by exposing an NW-width region that cuts diagonally across multiple NWs. The angle of the cut must be chosen so that each NW has an exposed portion under a different MW. This method is sensitive to small changes in the angle of the cut. An angular accuracy of less than a few degrees is necessary when MWs are 10 times the width of NWs, which is expected to be typical. Also, a translational misalignment of the cut relative to MWs may result in the ambiguous control of NWs by MWs, as is the case with modulation-doped NWs [28].

Gopalakrishnan et al. assemble a small set of parallel NWs using nanolithography, attach OCs to both ends of NWs, position gate electrodes near one end of the NWs, and control NWs by applying potentials via electrodes that deplete the NWs nearest to the electrodes with the objective of leaving one undepleted NW [33]. As yet, no modeling of variability in the assembly of NWs or of the effect on the voltages needed to control individual NWs is known to the authors. It is unclear if this method can control more than three or four NWs with a diameter of 5 to 10 nm.

3 CRITERIA FOR NW ADDRESSABILITY

In the decoders described in Sections 2.1 through 2.3, turning on an MW increases the resistance of some random subset of NWs. When multiple MWs are turned on, the increases in resistance introduced by each MW accumulate. When an NW is addressed, its resistance must be much less than the resistance of all $w - 1$ other NWs in the same contact group when combined in parallel.

When all MWs are turned off, let r_{low} denote the maximum resistance of any one NW. An MW **controls** a section of an NW if, when it is turned on, it increases the

NW's total resistance by an amount much larger than wr_{low} . This ensures that, when an MW is turned on, the combined resistance of the NWs it controls is greater than the resistance of the NW being addressed. The section of the NW under that MW is said to be **controllable**. Conversely, the section of an NW under an MW is **noncontrollable** if the MW increases an NW's resistance by an amount much less than r_{low} . Finally, the section is **ambiguous** if it is neither controllable nor noncontrollable. A more mathematically rigorous discussion of these definitions is given in [35].

An MW **controls, does not affect, or is ambiguous with respect to** an NW if the NW has a section that is controllable, noncontrollable, or ambiguous underneath that MW. An NW n_i is **individually addressable** if there exists some subset S of MWs such that every MW in S does not affect n_i and every other NW in the same contact group as n_i is controlled by at least one NW in S .

3.1 Conditions for NW Control in the Linear Decoder

In a mask-based decoder, the locations of the LRs determine which MWs control which NWs. Consider adjacent NWs n_a and n_b , where n_a is to the left of n_b . If an LR under an MW has a left boundary between n_a and n_b , the section of n_a under the MW is uncontrollable and the section of n_b under the MW is controllable. As the LR's left boundary moves rightward, there is a point at which the section of n_b goes from being controllable to ambiguous. Similarly, as the boundary moves leftward, there is a point at which the section of n_a goes from being noncontrollable to ambiguous. The region between these two limits is called the **interNW region**. The following condition ensures that all pairs of NWs in a group of consecutive NWs are individually addressable.

Lemma 3.1. *Assume that the length of and separation between LRs both span at least w NWs. All NWs in a group of w consecutive NWs are addressable if and only if the left boundary and right boundary of two different LRs fall in the interNW region associated with each of the $w - 1$ pairs of consecutive NWs.*

Proof. An NW n_i is individually addressable if and only if there exists a subset of MWs, denoted S_i , such that no MW in S_i affects n_i and all $w - 1$ other NWs are controlled by at least one MW in S_i .

For the "if" case, assume that all consecutive pairs of NWs have left and right LR boundaries in the interNW regions between them and consider an arbitrary NW n_a . There exists an MW m_1 that lies on top of an LR whose left boundary is in the interNW region to the right of n_a . Since the LR must have a length spanning at least w NWs, MW m_1 controls all NWs in question to the right of n_a . Similarly, there exists an MW m_2 that lies on top of an LR whose right boundary is in the interNW region to the left of n_a . This MW controls all of the NWs in question to the left of n_a . The set $S_a = \{m_1, m_2\}$ individually addresses n_a .

For the "only if" case, assume that all NWs are independently addressable. Consider any two adjacent NWs n_a and n_b , where n_a is to the left of n_b and I_{ab} is the interNW region between them. If n_a is individually

addressable, then there must be an MW in S_a that controls n_b but not n_a . This implies that the LR under this MW has its left boundary in I_{ab} . Similarly, since n_b is individually addressable, there exists an MW that controls n_a but not n_b and, thus, some LR has its right boundary in I_{ab} as well. \square

This lemma proves that w consecutive NWs are controllable when right and left LR boundaries lie in each of the $w - 1$ interNW regions. As explained in Section 4, LR boundaries are placed stochastically. Consequently, many rows of LRs are necessary to ensure that these conditions hold with high probability.

This closely resembles the classic **coupon collector problem** in which a random “coupon” (here, an interNW region) is collected at each of T trials (here, LRs). One then asks how large T must be for each of the C coupons to be collected with high probability. It is well known that T must be proportional to $C \ln C$. In Section 6, we introduce variants of the coupon collector problem that are relevant to the randomized mask-based decoder.

4 STOCHASTIC ASSEMBLY OF MASK-BASED DECODERS

The randomized mask-based decoder can be used to control any type of long straight uniformly spaced lightly doped semiconducting NWs. This decoder was first proposed for use with NWs produced by the superlattice NW pattern transfer method (SNAP) [2]. It can also be used with NWs grown by nanoimprinting [10], [7].

SNAP uses molecular beam epitaxy (MBE) to make a GaAs/AlGaAs superlattice from which the AlGaAs layer is etched back, creating a sawtoothed block face. Metal is deposited through evaporation on edges and pressed onto an adhesive layer on silicon. After the superlattice is removed, metallic NWs remain attached to the silicon. These metallic NWs are used as a nanometer-scale mask for a thin silicon layer residing on top of silicon oxide to produce silicon NWs [11]. SNAP has been shown to produce very long (2-3 mm) thin (8-10 nm) largely defect-free NWs having a uniform pitch (16-60 nm) that can be deposited on a chip with each application of SNAP.

In more recent experiments [17], SNAP has been used to create an array of 150 silicon NWs with a width of 13 nm and a pitch of 34 nm. To produce lightly doped NWs, the silicon is doped before metallic wire deposition and silicon NW etching. After exposing silicon NWs, a light etching is done to remove the top few nanometers so that the dopant concentration is reduced to a controllable level. Control over groups of consecutive of NWs was demonstrated using lithographically produced high-K dielectric regions.

4.1 LR Manufacture

To deposit LRs on a chip using lithography, one or more masks are constructed containing multiple rectangular openings. When openings are first made in masks, a one-time process, the separation between the rectangles, as well as their size can vary somewhat from their intended values. Additionally, when masks are used, it is difficult to control

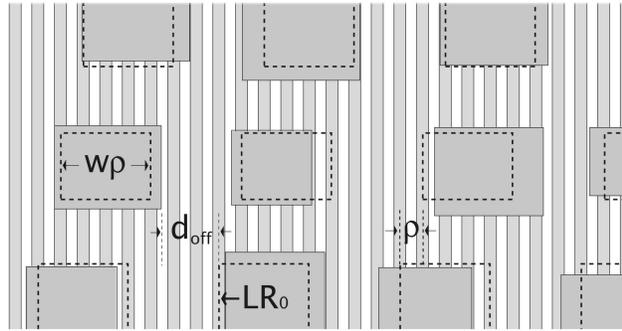


Fig. 4. Each LR has a nominal location on a mask indicated by dashed lines. Its actual location depends on random perturbances in end points denoted by random variables $\{d_k\}$. The location of a mask containing LRs is specified by the relative offset d_{off} to a canonical LR denoted LR_0 . Also, note that, in this figure, $\theta = \rho/2$, $w = 4$, and the full cycle of LRs would consist of $2w = 8$ rows.

the precise alignment of the openings with the NWs. The offset of a mask from its intended location may be large.

After light is passed through the openings in a mask onto a photoresist, an etching process removes either the lithographically defined regions (positive photoresist) or their complement (negative photoresist). The duration of the etching process, which cannot be precisely controlled, causes variation in the length and width of the LRs.

Let ρ denote the pitch of the NWs. We refer to the intended location of an LR’s right or left boundary, relative to the NWs, as its nominal location. Variation in mask manufacture, mask placement, and mask application all cause an LR’s end point to vary from its nominal location. In the absence of variation, $2w$ left and $2w$ right LR boundaries suffice to create a perfect one-cycle linear decoder (see Lemma 3.1). Variation, however, introduces the need for multiple cycles, which we assume are placed using one or more masks.

E-beam lithography is currently too expensive for mass production, but it sets a limit on the best possible conditions. Using it, 1) masks can be offset by 50 to 100 nm from their intended locations, 2) the length and relative placement of rectangular mask openings can vary by 5 to 10 nm from their intended locations on a mask, and 3) the etching of photoresist can increase the length of LRs by up to 5 nm on a chip [34]. If photolithography is used, the longer wavelength of the radiation results in larger variations in these parameters. Uncertainty in mask placement and variation in mask manufacture are independent of the type of lithography employed.

4.2 Modeling Variation in Mask Placement

Let d_{off} be the offset of a mask from its ideal location, which we assume places the nominal locations of LR boundaries at the midpoint between NWs. d_{off} is defined in terms of the location of a particular but arbitrary LR boundary that we call the **canonical LR boundary** LR_0 (see Fig. 4). If d_{off} can be large relative to an NW pitch ρ , as we assume is the case, then the assumed uniformity in the placement of NWs allows us to replace d_{off} by the **phase difference** θ , which is restricted to the interval $-\rho/2 \leq \theta \leq \rho/2$. Note that $\theta = 0$ corresponds to the boundary LR_0 being at the middle of the space between

two NWs. It is not important which two NWs it lies between.

Because we assume that the variation of d_{off} is large relative to ρ , we model θ as a uniform random variable (r.v.) over the interval $[-\rho/2, \rho/2]$. If the variation in d_{off} is small, as would be the case when the spacing between NWs is large, a nonuniform distribution in d_{off} would be appropriate, a case that we ignore.

4.3 Modeling Variation in LR Boundary Placement

When θ is fixed, uncertainties in the LR boundary locations result from uncertainties in 1) the inscribing of rectangles on masks, 2) the exposure of photoresist by electromagnetic radiation through mask rectangles, and 3) the photoresist etching time. We collect all these variations in an r.v. d associated with each LR boundary. The actual location of an LR boundary is determined by θ , the offset of the nominal location of the boundary relative to the adjacent NWs, and d , the change in the position of the boundary relative to its nominal location.

We assume that d has a symmetric probability distribution $f(d)$ that decreases monotonically with d from $d = 0$. This reflects the fact that small variations in d are expected, and variations are equally likely to be positive or negative. We also assume that the r.v.s $\{d_k\}$ associated with left and right LR boundaries are statistically independent and identically distributed.

Lemma 3.1 states that, for all NWs to be controllable, an LR right boundary and an LR left boundary must fall in each interNW region between each pair of consecutive NWs. If an LR boundary does not fall into an interNW region, the LR boundary is said to **fail**. If a boundary does not fail, it may fall in the interNW region closest to its nominal location or some other interNW region. We refer to the interNW region closest to the nominal location as the **targeted interNW region**.

For each LR boundary, we let $p_i(\theta)$ be the **probability, given a mask phase difference of θ , that an LR boundary moves i regions to the right (left) from its targeted interNW region, when i is positive (negative)**. Because the r.v.s $\{d_k\}$ are statistically independent when θ is fixed, the conditional joint probability that LR boundaries on a given mask fall into particular interNW regions is the product of $p_i(\theta)$.

The facts cited in Section 4.1 suggest that an LR boundary will vary by at most a few NW pitches when the mask offset d_{off} is fixed. That is, $q_i(\theta)$ will be nonzero only for small absolute values of i . We assume that $p_i(\theta) = 0$ for $i \geq w$. Since the right (and left) boundaries of LRs under the same MW are separated by $2w\rho$, only one such boundary has a nonzero probability of landing in any particular interNW region.

4.4 Additional Sources of LR Boundary Variation

LRs can also be placed using a stamping process [34]. The LRs in a stamp could then be inscribed using E-beam lithography and the stamp could be used multiple times. Two issues arise in the use of a stamp: 1) Uncertainties in the length and separation of LRs grow with the number of stampings and 2) large uncertainties arise in the angular orientation of a stamp relative to NWs. It is estimated that

the latter could be as large as 20 to 30 degrees. E-beam lithography may also introduce a small amount of angular uncertainty.

We do not explicitly model either the degradation of the stamps or the angular uncertainty introduced by both stamping and E-beam lithography in this paper. We believe, however, that these sources of variation can still be analyzed using our methods. Both have the effect of increasing the length of LR and decreasing the amount of space between NWs. As a result, the width of an interNW region shrinks because sections of NWs that would otherwise be uncontrollable become ambiguous. This in turn reduces each $p_i(\theta)$.

5 ANALYZING THE RANDOMIZED MASK-BASED DECODER

The randomized n -cycle mask-based decoder uses a standard CMOS decoder to activate a contact group of w NWs. The high- K dielectric regions are then used to turn off all but one NW in a group. As described in Section 2.3, the regions are arranged in n cycles, where a cycle requires $2w$ MWs. The randomized n -cycle decoder is designed to activate one of w NWs with high probability. As shown in Lemma 3.1, this requires both a left LR boundary and a right LR boundary to fall into each of the $w - 1$ interNW regions.

During manufacture, the n cycles of the decoder are placed using some number of masks. Associated with each mask is a phase difference θ . The θ s are uniformly distributed independent r.v.s. Given θ , we know the nominal positions of all LR boundaries produced by that mask. We assume that each LR boundary varies independently about its nominal position according to some unimodal symmetric distribution centered at 0.

We consider two models for assignment of cycles to masks. In the first, which is the **coarse-grained model**, we assume that the LRs under each MW are on separate masks. Thus, this model has $2nw$ different masks and one phase difference r.v. per mask, $\{\theta_t | 1 \leq t \leq 2nw\}$. In the second, which is the **fine-grained model**, we assume that each mask places one or more cycles.

A randomized n -cycle mask-based decoder has N/w groups of w NWs. The decoder controls all N NWs if each NW in each set of w NWs is individually addressable. If there are m masks, let $\underline{\theta} = (\theta_1, \dots, \theta_m)$ denote the set of m phase differences of these masks. For $1 \leq l \leq N/w$, let $F_l(\underline{\theta})$ denote the **failure to control all w NWs associated with the l th set of NWs given a value for $\underline{\theta}$** . Let $F(\underline{\theta})$ be the event that some NW in some set of N/w NWs is not controllable. It follows that $F(\underline{\theta})$ is the union of the events $F_l(\underline{\theta})$, $1 \leq l \leq N/w$. That is,

$$F(\underline{\theta}) = F_1(\underline{\theta}) \cup \dots \cup F_{N/w}(\underline{\theta}).$$

The unconditional probability of failure to control all N NWs $P_r(F)$ is the average of $P_r(F(\underline{\theta}))$ over all m phase differences:

$$P_r(F) = \left(\frac{1}{\rho}\right)^m \int_{-\rho/2}^{\rho/2} \dots \int_{-\rho/2}^{\rho/2} P_r(F(\underline{\theta})) d\theta_1 \dots d\theta_m.$$

Below, we use the principle of inclusion and exclusion to bound $P_r(F)$.

Theorem 5.1. *The probability $P_r(F)$ has the following bounds:*

$$Q(1 - Q/2) < P_r(F) \leq Q,$$

where $Q = \rho^{-m} (N/w) \int_{-\rho/2}^{\rho/2} \cdots \int_{-\rho/2}^{\rho/2} P_r(F_1(\underline{\theta})) d\theta_1 \cdots d\theta_m$.

Proof. When the principle of inclusion and exclusion is used, the conditional probability $P_r(F(\underline{\theta}))$ has the following bounds:

$$Q(\underline{\theta}) - \sum_{l < m} P_r(F_l(\underline{\theta}) \cap F_m(\underline{\theta})) \leq P_r(F(\underline{\theta})) \leq Q(\underline{\theta}).$$

Here, $Q(\underline{\theta}) = \sum_{l=1}^{N/w} P_r(F_l(\underline{\theta}))$. Because the conditioned events $F_l(\underline{\theta})$ are assumed to be statistically independent, $P_r(F_l(\underline{\theta}) \cap F_m(\underline{\theta})) = P_r(F_l(\underline{\theta}))P_r(F_m(\underline{\theta}))$.

Let Q be the average of $Q(\underline{\theta})$, that is,

$$Q = \rho^{-m} \int_{-\rho/2}^{\rho/2} \cdots \int_{-\rho/2}^{\rho/2} Q(\underline{\theta}) d\theta_1 \cdots d\theta_m.$$

Because the events $F_l(\underline{\theta})$ are identically distributed, $Q = (N/w) \overline{P_r(F_1(\underline{\theta}))}$, where $\overline{P_r(F_1(\underline{\theta}))}$ is defined below:

$$\overline{P_r(F_1(\underline{\theta}))} = \rho^{-m} \int_{-\rho/2}^{\rho/2} \cdots \int_{-\rho/2}^{\rho/2} P_r(F_1(\underline{\theta})) d\theta_1 \cdots d\theta_m.$$

The sum in the above lower bound has $(N/w)(N/w - 1)/2$ terms. Each term $P_r(F_l(\underline{\theta}))P_r(F_m(\underline{\theta}))$ is a product of statistically independent identically distributed r.v.s. Thus, its average over $\underline{\theta}$ is

$$(N/w)(N/w - 1) \left(\overline{P_r(F_1(\underline{\theta}))} \right)^2 / 2.$$

Because $Q = (N/w) \overline{P_r(F_1(\underline{\theta}))}$, this average becomes $((N/w - 1)/(N/w))Q^2/2$, which is less than $Q^2/2$, giving the desired result. \square

Since the goal is to make Q very small, Q and $P_r(F)$ are very close. In the remainder of this paper, we approximate the probability of failure to control all N NWs by Q .

Recalling that $F_l(\underline{\theta})$ is the event that between every pair of w NWs, we collect at least one left LR boundary and one right LR boundary given the phase differences $\underline{\theta}$. Let L (R) be the event that some left (right) LR boundary fails to be collected. Then, $P_r(L \cup R)$ is the probability that one or the other type of boundary fails to be collected. It follows that

$$\max(P_r(L), P_r(R)) \leq P_r(L \cup R) \leq P_r(L) + P_r(R).$$

Lemma 5.1. *The probability of failure to collect both left LR and right LR boundaries between every pair of N NWs is within a factor of two of the probability of failure to collect just left (or right) LR boundaries between every pair of N NWs.*

In light of the above fact, we consider only the collection of left LR boundaries.

In the next section, we model the collection of LR left boundaries as variants of the coupon collector problem. When there is one mask for each LR under each MW, this problem is modeled by the Coupon Collector Problem with Failures (Section 6.1). When all LRs are produced by one

mask, this is modeled by the targeted coupon collector problem (Section 6.2). In the final case, when multiple cycles are produced by multiple masks, the problem is a multi-stage version of the latter problem (Section 6.3).

6 COUPON COLLECTION

In this section, we analyze three increasingly general variants of the standard coupon collector problem: 1) the Coupon Collector Problem with Failures (CCF), 2) the targeted coupon collector problem, and 3) the multistage targeted coupon collector problem. These generalizations are motivated by the cyclic placement of LRs in mask-based decoders. They are used in Section 7 to analyze the randomized n -cycle mask-based decoder.

6.1 The Coupon Collector Problem with Failures

In the classic coupon collector problem, one of C coupons is randomly collected during each of the T trials. Trials are independent and each coupon is selected with probability $1/C$. We introduce the CCF in which, on each trial, either a coupon fails to be collected with probability p_f (this models an LR boundary that falls outside of an interNW region) or a coupon is collected with probability $(1 - p_f)/C$. T is chosen so that all coupons are collected with high probability.

Theorem 6.1. *Let Γ_{CCF} be the probability of failing to collect all C coupons in T trials when each trial has probability of failure $p_f = 1 - p_s$ and the probability of selecting the i th coupon is $p_i = p_s/C$ for $1 \leq i \leq C$. Then, Γ_{CCF} and T satisfy the following bounds:*

$$z(1 - z/2) \leq \Gamma_{CCF} \leq z,$$

where $z = C(1 - p_s/C)^T$. Let $\phi_{CCF} = -C \ln(1 - p_s/C)$. When z is small, minimizing z minimizes the bound on Γ_{CCF} . Then,

$$\frac{C}{\phi_{CCF}} \ln \left(\frac{C}{\Gamma_{CCF}(1 + \Gamma_{CCF})} \right) \leq T \leq \frac{C}{\phi_{CCF}} \ln \left(\frac{C}{\Gamma_{CCF}} \right)$$

when $\Gamma_{CCF} \leq \sqrt{2} - 1$. ϕ_{CCF} satisfies $p_s \leq \phi_{CCF} \leq p_s(1 + p_s/C)$ if $C \geq 2$.

Proof. Theorem 6.1 is a special case of Theorem 6.2 below.

When $p_r = p_s/C$ for all r , z and ϕ are the same as defined above. \square

6.2 The Targeted Coupon Collector Problem

We further generalize the coupon collector problem by allowing each trial to “target” a certain coupon. We call this the **targeted coupon collector problem**. As before, trials fail with probability p_f , but, when a failure does not occur, each coupon is collected with a probability that is a function of the distance of the coupon from the targeted location. Let p_0, p_1, \dots, p_{C-1} be these probabilities. Clearly, $p_f + \sum_{r=0}^{C-1} p_r = 1$. The targeted coupon collector problem reduces to the CCF when $p_r = p_s/C$ for all r .

Associated with each trial is a coupon t_j , $1 \leq j \leq T$, that is targeted. The probability that the j th trial collects the i th coupon is $p_{r(i,j)}$, where $r(i,j) = (i - t_j) \bmod C$. This has the effect of targeting the coupons in a cyclic fashion.

Consider C bins placed in a circle. At each of the T trials, a ball is thrown from directly overhead. A trial collects the

i th coupon if it lands in the i th bin. Each throw is aimed at a particular bin t_j . The likelihood that a ball hits its target is always p_0 . The probability that a ball deviates one bin to the right is p_1 . The probability that a ball deviates one bin to the left is p_{C-1} . The probability that a ball fails to land in any bin at all is p_f . Clearly, in this model, these probabilities are independent of t_j .

As before, we wish to know how large T must be so that all coupons are collected with high probability. We are free to assign any value to each t_j , but we require these values to be chosen in advance. Each t_j cannot be based on the outcomes of previous trials. In our model, we assume that each value of t_j is chosen an equal number of times and that T is a multiple of C . This is equivalent to cycling through all C coupons multiple times. Thus, we let $t_j = j \bmod C$ and call this the **cyclic coupon collector problem (CCC)**.

Theorem 6.2. *Let Γ_{CCC} be the probability of failing to collect all C coupons in T trials, T being a multiple of C in the CCC when each trial has probability of failure $p_f = 1 - p_s$ and the probability of collecting the i th coupon on the j th trial is $p_{r(i,j)}$, where $r(i,j) = (i-j) \bmod C$. Then, Γ_{CCC} and T satisfy the following bounds:*

$$z(1 - z/2) \leq \Gamma_{CCC} \leq z,$$

where $z = C \prod_{r=0}^{C-1} (1 - p_r)^{T/C} = C e^{-\phi_{CCC} T/C}$ and

$$\phi_{CCC} = - \sum_{r=0}^{C-1} \ln(1 - p_r).$$

When z is small, minimizing z minimizes the bound on Γ_{CCC} . Then,

$$\frac{C}{\phi_{CCC}} \ln\left(\frac{C}{\Gamma_{CCC}(1 + \Gamma_{CCC})}\right) \leq T \leq \frac{C}{\phi_{CCC}} \ln\left(\frac{C}{\Gamma_{CCC}}\right)$$

when $\Gamma_{CCC} \leq \sqrt{2} - 1$. $p_s \leq \phi_{CCC} \leq p_s + \sum_{r=0}^{C-1} p_r^2$ when $p_r \leq 0.5$, where $p_s = \sum_{r=0}^{C-1} p_r$. The bounds on T are minimized by maximizing ϕ_{CCC} .

Proof. We use the principle of inclusion/exclusion. Let E_i be the event that the i th coupon is not collected after T trials and let $\Gamma_{CCC} = P(E_0 \cup \dots \cup E_{C-1})$.

We assume that coupons are targeted in a cyclic fashion. Let E'_i be the event that the i th coupon is not collected after C trials. The probability that the i th coupon is not collected on the j th trial is $(1 - p_{r(i,j)})$, where $r(i,j) = (i-j) \bmod C$. In C consecutive trials, $r(i,j)$ will take on every value from 0 to $C-1$. Since trials are independent,

$$P(E'_i) = \prod_{r=0}^{C-1} (1 - p_r).$$

Now, let E_i be the event that the i th coupon is not collected in any of the T trials, T being a multiple of C . Since $P(E_i) = P(E'_i)^{T/C}$,

$$P(E_i) = \prod_{r=0}^{C-1} (1 - p_r)^{T/C},$$

which is independent of i .

Now, bound $P(E_h \cap E_i)$. Observe that the h th and i th coupons are not collected on the j th trial with probability $(1 - p_{r(h,j)} - p_{r(i,j)})$. Since $(1 - a - b) \leq (1 - a)(1 - b)$,

$$(1 - p_{r(h,j)} - p_{r(i,j)}) \leq (1 - p_{r(h,j)})(1 - p_{r(i,j)}).$$

As before, over C consecutive trials, $r(h,j)$ and $r(i,j)$ range over all values from 0 to $C-1$. Reordering terms allows us to write

$$\begin{aligned} P(E_h \cap E_i) &= P(E'_h \cap E'_i)^{T/C} \\ &\leq \left[\prod_{r=0}^{C-1} (1 - p_r) \prod_{r=0}^{C-1} (1 - p_r) \right]^{T/C} \leq P(E_i)^2. \end{aligned}$$

Applying the principle of inclusion and exclusion, we have

$$\sum_{i=0}^{C-1} P(E_i) - \sum_{h<i} P(E_i)^2 \leq \Gamma_{CCC} \leq \sum_{i=0}^{C-1} P(E_i).$$

Since $\sum_{h<i} P(E_i)^2 \leq (\sum_{i=0}^{C-1} P(E_i))^2 / 2$, this yields the following bounds:

$$z(1 - z/2) \leq \Gamma_{CCC} \leq z,$$

where $z = \sum_{i=0}^{C-1} P(E_i)$. The inequality $z(1 - z/2) \leq \delta$ implies that $z \leq 1 - \sqrt{1 - 2\delta}$. In turn, this implies that $z \leq \delta(1 + \delta)$ when $\delta \leq \sqrt{2} - 1$. Thus, if $\Gamma_{CCC} \leq \sqrt{2} - 1$,

$$\Gamma_{CCC} \leq z \leq \Gamma_{CCC}(1 + \Gamma_{CCC}).$$

Substituting in $z = C \prod_{r=0}^{C-1} (1 - p_r)^{T/C} = C e^{-\phi_{CCC} T/C}$, where $\phi_{CCC} = - \sum_{r=0}^{C-1} \ln(1 - p_r)$, gives

$$\frac{C}{\phi_{CCC}} \ln\left(\frac{C}{\Gamma_{CCC}(1 + \Gamma_{CCC})}\right) \leq T \leq \frac{C}{\phi_{CCC}} \ln\left(\frac{C}{\Gamma_{CCC}}\right).$$

Finally, since $-x(1+x) \leq \ln(1-x) \leq -x$ when $x \leq 0.5$, $\sum_{r=0}^{C-1} p_r \leq \phi_{CCC} \leq \sum_{r=0}^{C-1} (p_r + p_r^2)$. Thus, $p_s \leq \phi_{CCC} \leq p_s + \sum_{r=0}^{C-1} p_r^2$ when $p_r \leq 0.5$, where $p_s = \sum_{r=0}^{C-1} p_r$. \square

It is of interest to know how sensitive the bounds on T are to the probability distribution $\{p_0, p_1, \dots, p_{C-1}\}$. When all probabilities are the same, that is, $p_i = p_s/C$, the cyclic coupon collector problem is equivalent to the standard CCF. In this case, $\phi_{CCC} = -C \ln(1 - p_s/C)$ and the bounds are the same.

Now, consider a distribution that is far from uniform, one that is concentrated on just $C = 3$ points. If $p_0 = p_1 = p_2 = 1/4$ and $p_s = 3/4$, then $\phi_{CCC} = 3 \ln(4/3) = 0.86$. On the other hand, $\phi_{CCF} = -C \ln(1 - p_s/C) \approx p_s$ when $C \geq 10$. In this case, $\phi_{CCF} \approx \phi_{CCC}$ and the two bounds differ by a constant factor close to 1. Even if our ability to target specific coupons is good, the bounds on T continue to grow as $C \ln(C/\delta)$, where δ is the probability of failing to collect all coupons. Collecting all coupons remains difficult.

6.3 The Multistage Targeted Coupon Collector Problem

The targeted coupon collector problem is now generalized to m "stages," where each stage captures the variation introduced by using a new mask. It is an extension of the

CCC. In this problem, for some integer T_μ divisible by C , a **stage** is a set of T_μ trials, where the j th coupon t_j is targeted T_μ/C times. Associated with each stage is a uniformly distributed r.v. $\theta \in [-\rho/2, \rho/2]$ such that the probability of collecting a coupon targeted at a location i places away is $p_i(\theta)$, $0 \leq i \leq C-1$, a continuous function of θ . Also, $p_s(\theta) = 1 - p_f(\theta) = p_0(\theta) + \dots + p_{C-1}(\theta)$, where $p_f(\theta)$ is the failure to collect any coupon on one trial. In addition, the stage r.v.s $\underline{\theta} = (\theta_1, \theta_2, \dots, \theta_m)$ are statistically independent. We call this the **multistage targeted coupon collector problem**.

Because this problem models an n -cycle randomized mask-based decoder, we are free to consider putting either one or multiple cycles on one stage. Thus, we would like to know how the failure probability $\Gamma_{MM} = P(E_0 \cup E_1 \cup \dots \cup E_{C-1})$ depends on the number of cycles per stage. We show that it is smallest when each stage contains one cycle.

Theorem 6.3. *Let Γ_{MS} be the probability of failure to collect all coupons in the multistage targeted coupon collection problem with m stages in T trials when there are T_μ cycles in the μ th stage, T_μ being a multiple of C , $1 \leq \mu \leq m$, and $T = T_1 + \dots + T_m$, where the stage r.v.s $\underline{\theta}$ are statistically independent. Then, Γ_{MS} and T satisfy the following bounds:*

$$z(1 - z/2) \leq \Gamma_{MS} \leq z,$$

$$\text{where } z = Ce^{-\phi_{MS}T/C}$$

$$\phi_{MS} = -\ln \prod_{\mu=1}^m \left(\frac{1}{\rho} \int_{-\rho/2}^{\rho/2} \left(\prod_{r=0}^{C-1} (1 - p_r(\theta_\mu)) \right)^{T_\mu/C} d\theta_\mu \right)$$

and

$$\frac{C}{\phi_{MS}} \ln \left(\frac{C}{\Gamma_{MS}(1 + \Gamma_{MS})} \right) \leq T \leq \frac{C}{\phi_{MS}} \ln \left(\frac{C}{\Gamma_{MS}} \right)$$

when $\Gamma_{CCC} \leq \sqrt{2} - 1$.

When z is small, minimizing z (maximizing ϕ_{MS}), minimizes the bound on Γ_{MS} . The quantity z is minimized by placing each cycle in a separate stage, in which case, z satisfies the following bound:

$$z \geq C \left(\frac{1}{\rho} \int_{-\rho/2}^{\rho/2} \prod_{r=0}^{C-1} (1 - p_r(\theta_\mu)) d\theta_\mu \right)^{T/C}.$$

Proof. We use the principle of inclusion/exclusion in which E_i is the event that the i th coupon is not collected after T trials and we let $\Gamma_{MS} = P(E_0 \cup \dots \cup E_{C-1})$.

We derive bounds on the failure event conditioned on the r.v.s $\underline{\theta}$, namely, $\Gamma_{MS}(\underline{\theta}) = P(E_0 \cup E_1 \cup \dots \cup E_{C-1} | \underline{\theta})$, and then average the bounds over all values of $\underline{\theta}$.

Let E_i^μ be the event that the i th coupon fails to be collected during T_μ trials in the μ th stage. It follows that $E_i = E_i^1 \cap \dots \cap E_i^m$, where $\{E_i^1, E_i^2, \dots, E_i^m\}$ are statistically independent given the parameters $\underline{\theta}$. It follows that the conditional probabilities factor is as stated below:

$$P(E_i | \underline{\theta}) = P_r(E_i^1 | \theta_1) \dots P_r(E_i^m | \theta_m).$$

To employ the principle of inclusion/exclusion, we derive a bound on the conditional probability

$P(E_h \cap E_i | \underline{\theta})$. Using the definition of these two events and the reasoning employed in the proof of Theorem 6.2, we have the following bound:

$$P(E_h \cap E_i | \theta_1, \theta_2, \dots, \theta_m) \leq \prod_{\mu=1}^m P^2(E_i^\mu | \theta_\mu).$$

Here, $P(E_i^\mu | \theta_\mu)$ is independent of i , although it is dependent on θ_μ .

Averaging the bounds over $\underline{\theta}$ and applying the reasoning of the proof of Theorem 6.2, we have that $z(1 - z/2) \leq \Gamma_{MS} \leq z$, where

$$\begin{aligned} z &= \left(\frac{1}{\rho} \right)^m \int_{-\rho/2}^{\rho/2} \dots \int_{-\rho/2}^{\rho/2} \sum_{i=1}^C P(E_i | \underline{\theta}) d\underline{\theta} \\ &= \sum_{i=1}^C \prod_{\mu=1}^m \left(\frac{1}{\rho} \int_{-\rho/2}^{\rho/2} P(E_i^\mu | \theta_\mu) d\theta_\mu \right) \\ &= C \prod_{\mu=1}^m \left(\frac{1}{\rho} \int_{-\rho/2}^{\rho/2} \left(\prod_{r=0}^{C-1} (1 - p_r(\theta_\mu)) \right)^{T_\mu/C} d\theta_\mu \right) \\ &= Ce^{-\phi_{MS}T/C}. \end{aligned}$$

The latter result follows because $P(E_i^\mu | \theta_\mu)$ is independent of i .

A lower bound to z follows from a lower bound to $\frac{1}{\rho} \int_{-\rho/2}^{\rho/2} \left(\prod_{r=0}^{C-1} (1 - p_r(\theta_\mu)) \right)^{T_\mu/C} d\theta_\mu$. Holder's inequality is stated below, where $1/p + 1/q = 1$ and p and $q \geq 1$:

$$\int_X |f(y)g(y)| dy \leq \left(\int_X |f(y)|^p dy \right)^{1/p} \left(\int_X |g(y)|^q dy \right)^{1/q}.$$

Let $X = [-\rho/2, \rho/2]$, $f(y) = \left(\prod_{r=0}^{C-1} (1 - p_r(\theta_\mu)) \right)$ and $g(y) = 1/\rho$. Then, the inequality becomes the following:

$$\begin{aligned} \int_{-\rho/2}^{\rho/2} \frac{1}{\rho} f(y) dy &\leq \left(\int_{-\rho/2}^{\rho/2} f(y)^p dy \right)^{1/p} \left(\int_{-\rho/2}^{\rho/2} \rho^{-q} dy \right)^{1/q} \\ &= \left(\int_{-\rho/2}^{\rho/2} \frac{1}{\rho} f(y)^p dy \right)^{1/p}. \end{aligned}$$

Here, we have used the fact that $(1/q) - 1 = -1/p$. Consequently, when $p = T_\mu/C$,

$$\begin{aligned} \frac{1}{\rho} \int_{-\rho/2}^{\rho/2} \left(\prod_{r=0}^{C-1} (1 - p_r(\theta_\mu)) \right)^{T_\mu/C} d\theta_\mu &\geq \\ \left(\frac{1}{\rho} \int_{-\rho/2}^{\rho/2} \prod_{r=0}^{C-1} (1 - p_r(\theta_\mu)) d\theta_\mu \right)^{T_\mu/C} &. \end{aligned}$$

This implies the following lower bound to z :

$$\begin{aligned} z &\geq C \prod_{\mu=1}^m \left(\frac{1}{\rho} \int_{-\rho/2}^{\rho/2} \prod_{r=0}^{C-1} (1 - p_r(\theta_\mu)) d\theta_\mu \right)^{T_\mu/C} \\ &= C \left(\frac{1}{\rho} \int_{-\rho/2}^{\rho/2} \prod_{r=0}^{C-1} (1 - p_r(\theta_\mu)) d\theta_\mu \right)^{T/C}. \end{aligned}$$

However, this is the bound that applies when each cycle is placed on a separate stage. \square

7 PERFORMANCE OF THE RANDOMIZED n -CYCLE MASK-BASED DECODER

In this section, we bound the number of MWs required to control all NWs in a randomized n -cycle mask-based decoder. We consider two models for the random placement of LRs: 1) the **course-grained model**, in which each LR is placed independently using a separate mask, and 2) the **fine-grained model**, in which LRs are placed using masks that contain one or more cycles. The course-grained model provides a conservative upper bound on the number of MWs required to control all NWs with high probability. The fine-grained model provides a lower bound on the number of MWs required using more optimistic assumptions.

7.1 The Coarse-Grained Model

In the coarse-grained model, each LR is placed on a separate mask. Since we assume that mask displacement can be at least 50-100 nm, this is comparable to the number of NWs (which might be as small as 10 but could be larger) that are expected to fall under the smallest LR. Thus, one can view the LR boundaries as equally likely to fall between any pair of w NWs. Only one of the two boundaries of a given LR falls within a set of w NWs. Thus, we can treat each boundary displacement as a uniformly distributed r.v. because all of its variation is in the displacement of the mask.

The randomized n -cycle mask-based decoder activates one set of w NWs in a group and contains one linear decoder with n cycles for each group. Each linear decoder addresses one NW by deactivating all but one of these w NWs. Theorem 5.1 provides tight bounds on the probability $P(F)$ that not all NWs can be addressed. This bound is the sum of the probabilities of failure to have an NW be addressable in one or more of the N/w sets of w NWs. For each NW to be addressable, a left NW boundary and a right NW boundary must fall in the interNW region between every adjacent pair of NWs. We consider only the collection of left LR boundaries and incur a penalty of at most a factor of two, as explained at the end of Section 5. The probability that there is an LR left NW boundary between each pair of NWs is modeled by the CCF. The probability of failure to collect all coupons is bounded in Theorem 6.1. We use the upper bound on T to obtain the following bound on the total probability of failure $P(F_{cg})$ for the coarse-grained case.

Theorem 7.1. *The probability of failure to address all NWs in the coarse-grained model $P(F_{cg})$ satisfies $P(F_{cg}) \leq \epsilon$ when T , the number of MWs in the linear portion of the randomized mask-based decoder, is chosen as follows:*

$$T = \frac{w-1}{p_s} \ln \left[\left(\frac{2N}{\epsilon} \right) \left(\frac{w-1}{w} \right) \right].$$

The smallest value of T that satisfies $P(F_{cg}) \leq \epsilon$ is close to this value when ϵ is small.

Proof. As shown in Section 5, $P(F_{cg})$ is at most twice the sum of the probabilities of failing to collect all LR left boundaries in N/w sets of w NWs. That is, $P(F_{cg}) \leq 2(N/w)\Gamma_{CCF}$, where Γ_{CCF} is the probability of failure to collect $C = w - 1$ coupons when the i th coupon

is collected with probability $p_i = p_s/C$ and $p_s = 1 - p_f$, where p_s and p_f are the probabilities of success and failure in collecting coupons. If T is chosen so that $\Gamma_{CCF} = (\epsilon w)/(2N)$, then $P(F_{cg}) \leq \epsilon$. We use the bounds of Theorem 7.1 to bound T when $\Gamma_{CCF} = (\epsilon w)/(2N)$. In particular, if $T = \frac{C}{p_s} \ln \left(\frac{C}{\Gamma_{CCF}} \right)$, $P(F_{cg}) \leq \epsilon$. By examining the steps in the approximations, it is clear that this bound is tight when ϵ is small. \square

Performance of the model. The number T of MWs in the linear portion of the decoder to ensure that the probability of failing to address all N NWs in the coarse-grained model is very close to $((w-1)/p_s) \ln(2N/\epsilon)$ when $w \geq 10$, which is logarithmic in $2N$ with an additive term proportional to $-\ln \epsilon$. The denominator p_s is the probability that an LR boundary succeeds in falling into an interNW region. Because an interNW region is slightly more than the space between two NWs, $p_s \geq 0.5$. Hence, $T \geq 2(w-1) \ln(2N/\epsilon)$.

Consider a concrete example in which there are $w = 10$ NWs per group, $N = 1,000$, and $\epsilon = 0.01$, that is, success is achieved in controlling all N NWs with a probability of 0.99 or higher. In this case, $T \geq 220$. This is a very large number of MWs.

This value for T should be compared to $T_{all.diff}$, the number of trials for the ‘‘all different’’ encoded-NW decoder described in [18], where it is shown that $T_{all.diff} \geq [2.2 \log_2 N] + 11$ suffices to control N NWs with a failure probability of $\epsilon = 0.01$. When $w = 10$ and $N = 1,000$, $T_{all.diff} = 33$ MWs can control 1,000 NWs with a probability of 0.99. The method in [18] requires a very large number C of differently encoded NW types. In particular, C may be more than 10,000. This number can be greatly reduced with a small effect on the number of addressable NWs using decoding strategies analyzed in [19].

As these calculations illustrate, the randomized mask-based decoder for the coarse-grained model requires many more MWs to decode N NWs than other decoders when N is 1,000 or more. We now explore the case when multiple cycles are placed on one mask.

7.2 The Fine-Grained Model

In the fine-grained model, several masks may be used. The mask phase differences are independent uniformly distributed r.v.s. The displacement of LR boundaries is small and modeled by the multistage targeted coupon collection problem. As with the coarse-grained model, the probability of failure to address all N NWs $P(F_{fg})$ is closely approximated by $2(N/w)\Gamma_{MS}$, where Γ_{MS} is the probability of failure to collect $C = w - 1$ coupons when the i th coupon is collected with probability $p_i(\theta)$ on a mask with phase difference θ .

As shown in Theorem 6.3, the probability of failure to collect all coupons in the multistage coupon collection problem is smallest when each cycle of the MWs occurs in a different stage. We summarize the result below.

Theorem 7.2. *The probability of failure to address all NWs in the fine-grained model satisfies $P(F_{fg}) \leq \epsilon$ when T , the number of MWs in the linear portion of the randomized mask-based decoder, is chosen as follows:*

$$T = \frac{w-1}{\phi_{MS}} \ln \left[\left(\frac{2N}{\epsilon} \right) \left(\frac{w-1}{w} \right) \right],$$

where ϕ_{MS} is defined below:

$$\phi_{MS} = -\ln \left(\frac{1}{\rho} \int_{-\rho/2}^{\rho/2} \prod_{\tau=0}^{w-2} (1 - p_r(\theta_\mu)) d\theta_\mu \right).$$

The smallest value of T that satisfies $P(F_{cg}) \leq \epsilon$ is close to this value when ϵ is small.

Proof. As with the previous proof, we observe that $P(F_{cg})$ is at most twice the sum of the probabilities of failing to collect all LR left boundaries in N/w sets of w NWs. That is, $P(F_{cg}) \leq 2(N/w)\Gamma_{MS}$, where Γ_{MS} is the probability of failure to collect $C = w - 1$ coupons in the multistage coupon collection problem when the i th coupon on a mask with phase difference θ is collected with probability $p_i(\theta)$. The bounds of Theorem 6.3 on Γ_{MS} are $z(1 - z/2) \leq \Gamma_{MS} \leq z$, where $z = (w - 1)e^{-\Gamma_{MS}T/(w-1)}$. When z is small, Γ_{MS} is approximated by z , which provides the desired result. \square

The bound on T for the fine-grained case is identical to that given for the coarse-grained model except that the denominator term p_s is replaced by ϕ_{MS} . Observe that ϕ_{MS} is increased and T is decreased if the product term in the definition of ϕ_{MS} is reduced.

Lemma 7.1. *The factor ϕ_{MS} satisfies the following bound, where $p_s(\theta)$ is the probability that an LR left boundary falls into an interNW region:*

$$\phi_{MS} \leq -\ln \left(1 - \frac{1}{\rho} \int_{-\rho/2}^{\rho/2} p_s(\theta) d\theta \right).$$

Proof. The proof follows from the fact that $(1 - a)(1 - b) \geq (1 - a - b)$. \square

Performance of the model. Given that θ is uniformly distributed, $\frac{1}{\rho} \int_{-\rho/2}^{\rho/2} p_s(\theta) d\theta$ is close to $1/2$ if the width and space of NWs are equal to $1/2$ of the NW pitch. Thus, ϕ_{MS} is close to $\ln 2 = 0.7$. Since the p_s from the bound for the coarse-grained case is about 0.5, we conclude that the number of MWs required is approximately $(0.5/0.7) * 220 \approx 157$. The randomized mask-based decoder is inefficient even when the location of LR boundaries can be tightly controlled.

8 PRACTICAL CONSIDERATIONS

The mask-based decoder requires a large number of MWs to control all NWs with high probability. In this section, we describe several practical considerations that may make mask-based decoding more attractive.

8.1 Address Translation Circuitry

To use an NW crossbar as a memory, each external binary address must be mapped to a different pair of orthogonal NWs. All three types of decoders described in Section 2 introduce uncertainty with regard to which MWs address which NWs. As a result, programmable address translation

circuitry (ATC) is required to map binary addresses to subsets of MWs.

When performing this mapping, we assume that each external binary address is divided into high and low-order bits. Each of these binary sequences is used to separately address an NW along each dimension of the crossbar. The ATC for each dimension maps the supplied binary sequence B to a contact group σ and a subset of MWs \mathcal{M} . When the MWs in \mathcal{M} are turned on, an NW in σ is addressed. The NW addressed by each B must be unique.

If each contact group has exactly w addressable NWs, the mapping from B to σ is fixed; it does not vary from decoder to decoder. Furthermore, if w is a power of 2, we can simply take σ to be the high-order bits of B . For \mathcal{M} , however, we cannot use the low-order bits of B . The subsets of MWs used to address individual NWs vary from contact group to contact group and from decoder to decoder.

For each B , the ATC must store a value for \mathcal{M} . The number of bits required for each \mathcal{M} is at most M since any subset of M MWs can be specified using M bits. M bits are necessary if most of the 2^M subsets appear with approximately equal frequency. This holds for both differentiated NW decoders and randomized contact decoders, which use $\Omega(M) = \Omega(\log N)$ bits per address.

In mask-based decoders, however, each NW can be addressed using just two MWs, one MW to turn off all NWs to its left and another to turn off all NWs to its right. Since each \mathcal{M} is a subset of two MWs, it can be stored using $2 \log(M) = \Omega(\log N)$ bits. Even though mask-based decoders require a large number of MWs, they do not require significantly larger ATC than other decoders.

8.2 Alternative Addressing Strategies

We have computed bounds on the number of MWs required so that every NW in every contact group is addressable with high probability $1 - \epsilon$. As explained in Section 7, this is equivalent to requiring any given contact group to have all NWs addressable with a probability of approximately $\epsilon/(N/w)$. Here, N/w , the number of contact groups, is on the order of 100.

As explained in [19] and [20], the number of MWs can be reduced if we relax the requirement that all NWs in all contact groups be addressable and modify the ATC accordingly. One approach is to require only most contact groups to have every NW addressable. If only a small number of contact groups fail to have every NW addressable, we can store each group that has failed in the ATC and have it skip these groups when mapping binary addresses to contact groups.

This alternative addressing strategy is illustrated in the following example: Suppose every contact group has every NW addressable with a probability of 0.955. By computing the tail of a binomial distribution with $p = 0.955$ and $N = 100$, one can show that, with a probability of 0.99, no more than 10 of the 100 contact groups fail. This only decreases the number of addressable NWs by a factor of 10 (from Nw to $0.9Nw$), but, since each contact group need only have every NW addressable with probability 0.045, that is, $\epsilon = 0.045$, from the theorems in Section 7, the number of MWs is reduced by a factor of two (compare $\ln(2 * 9 * 100/0.01)$ to $\ln(2 * 9/0.045)$). This still implies that more than 70 MWs are required, which

is significantly more than the number required by other decoding technologies when using the same addressing strategy. Under the same conditions, less than 30 MWs are necessary when either an encoded-NW or a randomized-contact decoder is used [19], [35].

9 CONCLUSION

We have analyzed the randomized n -cycle mask-based decoder, a new method for addressing NWs by interposing lithographically defined high-K dielectric regions between NWs and MWs [17]. The process of placing LRs is stochastic due to two factors: 1) the absolute location of masks relative to NWs is difficult to control and 2) small random variations will occur in the placement of LRs relative to one another. We have created models for the stochastic assembly of this decoder to account for these variations.

We have established conditions that LR boundaries must satisfy to ensure that all NWs in a set of w NWs can be individually addressed, namely, both an LR right boundary and an LR left boundary must fall between each pair of NWs.

We have modeled the satisfaction of this condition as the collection of coupons in variants of the classical coupon collector problem. We have introduced three models, the *coupon collector problem with failures*, the *targeted coupon collector problem*, and the *multistage targeted coupon collector problem*. The first problem is the classical problem except that coupons may fail to be collected. The second is like the first except that, over a series of trials, each coupon is targeted the same number of times, although nearby coupons may be collected instead. The third is the same as the second except that the trials are grouped into a series of stages wherein the probabilities associated with collecting coupons in a stage are parameterized with a different r.v. for each stage. The coupon collector problems that we present are of interest in their own right and may be useful in studying problems unrelated to mask-based decoding.

When our bounds are converted into numerical values representing typical cases, we find that the randomized mask-based decoder requires almost an order of magnitude more MWs to address all NWs than the encoded-NW decoder. In both [19] and [20], it was demonstrated that relaxing the requirement that all NWs in all contact groups be addressable results in a substantial reduction in the number of MWs. Although this is also true for mask-based decoders, they still require significantly more MWs than either an encoded-NW or a randomized-contact decoder.

A key lesson to be learned from these results is that it is difficult to individually address NWs if it is very likely that two adjacent NWs are either both controlled or both not controlled by any given MW. A strong correlation of this kind is a key characteristic of the randomized mask-based decoder. When NWs are differentiated before their random selection for deposition on a chip, this correlation disappears. A strong lack of correlation is also exhibited by the randomized-contact decoder [15], [16], which requires fewer NWs.

ACKNOWLEDGMENTS

The authors are pleased to acknowledge fruitful conversations with Benjamin Gojman. This research was funded in part by US National Science Foundation Grant

CCF-0403674. J.E. Savage was also supported in part by the École Polytechnique. An early version of this paper appeared in the *Proceedings of the International Symposium on VLSI 2005*.

REFERENCES

- [1] Y. Chen, G.-Y. Jung, D.A.A. Ohlberg, X. Li, D.R. Stewart, J.O. Jeppeson, K.A. Nielson, J.F. Stoddart, and R.S. Williams, "Nano scale Molecular-Switch Crossbar Circuits," *Nanotechnology*, vol. 14, pp. 462-468, 2003.
- [2] N.A. Melosh, A. Boukai, F. Diana, B. Gerardot, A. Badolato, P.M. Petroff, and J.R. Heath, "Ultra-high-Density Nanowire Lattices and Circuits," *Science*, vol. 300, pp. 112-115, Apr. 2003.
- [3] D. Whang, S. Jin, and C.M. Lieber, "Nanolithography Using Hierarchically Assembled Nanowire Masks," *Nano Letters*, vol. 3, no. 7, pp. 951-954, 2003.
- [4] Z. Zhong, D. Wang, Y. Cui, M.W. Bockrath, and C.M. Lieber, "Nanowire Crossbar Arrays as Address Decoders for Integrated Nanosystems," *Science*, vol. 302, pp. 1377-1379, 2003.
- [5] C.P. Collier, E.W. Wong, M. Belohradský, F.M. Raymo, J.F. Stoddart, P.J. Kuekes, R.S. Williams, and J.R. Heath, "Electronically Configurable Molecular-Based Logic Gates," *Science*, vol. 285, pp. 391-394, 1999.
- [6] C.P. Collier, G. Mattersteig, E.W. Wong, Y. Luo, K. Beverly, J. Sampaio, F. Raymo, J. Fraser Stoddart, and J.R. Heath, "A [2]Catenate-Based Solid State Electronically Reconfigurable Switch," *Science*, vol. 290, pp. 1172-1175, 2000.
- [7] G.Y. Jung, S. Ganapathiappan, A.A. Ohlberg, L. Olynick, Y. Chen, W.M. Tong, and R.S. Williams, "Fabrication of a 34×34 Crossbar Structure at 50 nm Half-Pitch by UV-Based Nanoimprint Lithography," *Nano Letters*, vol. 4, no. 7, pp. 1225-1229, 2004.
- [8] A.M. Morales and C.M. Lieber, "A Laser Ablation Method for Synthesis of Crystalline Semiconductor Nanowires," *Science*, vol. 279, pp. 208-211, 1998.
- [9] Y. Cui, L. Lauhon, M. Gudiksen, J. Wang, and C.M. Lieber, "Diameter-Controlled Synthesis of Single Crystal Silicon Nanowires," *Applied Physics Letters*, vol. 78, no. 15, pp. 2214-2216, 2001.
- [10] S.Y. Chou, P.R. Krauss, and P.J. Renstrom, "Imprint Lithography with 25-Nanometer Resolution," *Science*, vol. 272, pp. 85-87, 1996.
- [11] E. Johnston-Halperin, R. Beckman, Y. Luo, N. Melosh, J. Green, and J.R. Heath, "Fabrication of Conducting Silicon Nanowire Arrays," *J. Applied Physics Letters*, vol. 96, no. 10, pp. 5921-5923, 2004.
- [12] A. DeHon, "Array-Based Architecture for FET-Based, Nanoscale Electronics," *IEEE Trans. Nanotechnology*, vol. 2, no. 1, pp. 23-32, Mar. 2003.
- [13] Y. Huang, X. Duan, Q. Wei, and C.M. Lieber, "Directed Assembly of One-Dimensional Nanostructures into Functional Networks," *Science*, vol. 291, pp. 630-633, 2001.
- [14] D. Whang, S. Jin, Y. Wu, and C.M. Lieber, "Large-Scale Hierarchical Organization of Nanowire Arrays for Integrated Nanosystems," *Nano Letters*, vol. 3, no. 9, pp. 1255-1259, 2003.
- [15] R.S. Williams and P.J. Kuekes, "Demultiplexer for a Molecular Wire Crossbar Network," US patent 6,256,767, July 2001.
- [16] T. Hogg, Y. Chen, and P.J. Kuekes, "Assembling Nanoscale Circuits with Randomized Connections," *IEEE Trans. Nanotechnology*, vol. 5, no. 2, pp. 110-122, 2006.
- [17] R. Beckman, E. Johnston-Halperin, Y. Luo, J.E. Green, and J.R. Heath, "Bridging Dimensions: Demultiplexing Ultra-high-Density Nanowire Circuits," *Science*, vol. 310, pp. 465-468, 2005.
- [18] A. DeHon, P. Lincoln, and J.E. Savage, "Stochastic Assembly of Sublithographic Nanoscale Interfaces," *IEEE Trans. Nanotechnology*, vol. 2, no. 3, pp. 165-174, 2003.
- [19] B. Gojman, E. Rachlin, and J.E. Savage, "Evaluation of Design Strategies for Stochastically Assembled Nanoarray Memories," *J. Emerging Technologies in Computing Systems*, vol. 1, no. 2, pp. 73-108, 2005.
- [20] E. Rachlin and J.E. Savage, "Nanowire Addressing in the Face of Uncertainty," *Proc. IEEE Int'l Symp. VLSI*, pp. 225-230, Mar. 2006.
- [21] S.C. Goldstein and M. Budiu, "NanoFabrics: Spatial Computing Using Molecular Electronics," *Proc. 28th Ann. Int'l Symp. Computer Architecture*, pp. 178-189, <http://www.cs.cmu.edu/seth/papers/isca01.pdf>, June 2001.

- [22] M.R. Stan, P.D. Franzon, S.C. Goldstein, J.C. Lach, and M.M. Ziegler, "Molecular Electronics: From Devices and Interconnect to Circuits and Architecture," *Proc. IEEE*, vol. 91, no. 11, Nov. 2003.
- [23] A. DeHon, "Nanowire-Based Programmable Architectures," *J. Emerging Technologies in Computing Systems*, vol. 1, no. 2, pp. 109-162, 2005.
- [24] M.S. Gudixsen, L.J. Lauhon, J. Wang, D.C. Smith, and C.M. Lieber, "Growth of Nanowire Superlattice Structures for Nanoscale Photonics and Electronics," *Nature*, vol. 415, pp. 617-620, Feb. 2002.
- [25] Y. Wu, R. Fan, and P. Yang, "Block-by-Block Growth of Single-Crystal Si/SiGe Superlattice Nanowires," *Nano Letters*, vol. 2, no. 2, pp. 83-86, 2002.
- [26] M.T. Björk, B.J. Ohlsson, T. Sass, A.I. Persson, C. Thelander, M.H. Magnusson, K. Deppert, L.R. Wallenberg, and L. Samuelson, "One-Dimensional Steeplechase for Electrons Realized," *Nano Letters*, vol. 2, no. 2, pp. 87-89, 2002.
- [27] B. Gojman, E. Rachlin, and J.E. Savage, "Decoding of Stochastically Assembled Nanoarrays," *Proc. 2004 Int'l Symp. VLSI*, Feb. 2004.
- [28] J.E. Savage, E. Rachlin, A. DeHon, C.M. Lieber, and Y. Wu, "Radial Addressing of Nanowires," *J. Emerging Technologies in Computing Systems*, vol. 2, no. 2, pp. 129-154, 2006.
- [29] J.R. Heath, Y. Luo, and R. Beckman, "System and Method Based on Field-Effect Transistors for Addressing Nanometer-Scale Devices," US patent application 20050006671, Jan. 2005.
- [30] J.R. Heath and M.A. Ratner, "Molecular Electronics," *Physics Today*, vol. 56, no. 5, pp. 43-49, 2003.
- [31] X. Ma, D.B. Strukov, J.H. Lee, and K.K. Likharev, "Afterlife for Silicon: CMOL Circuit Architectures," *Proc. Fifth IEEE Conf. Nanotechnology*, 2005.
- [32] N.H. Di Spigna, D.P. Nackashi, C.J. Amsinck, S.R. Sonkusale, and P.D. Franzon, "Deterministic Nanowire Fanout and Interconnect without Any Critical Translational Alignment," *IEEE Trans. Nanotechnology*, vol. 5, no. 4, pp. 356-361, 2006.
- [33] K. Gopalakrishnan, R.S. Shenoy, C. Rettner, R. King, Y. Zhang, B. Kurdi, L.D. Bozano, J.J. Welsler, M.B. Rothwell, M. Jurich, M.I. Sanchez, M. Hernandez, P.M. Rice, W.P. Risk, and H.K. Wickramasinghe, "The Micro to Nano Addressing Block (MNAB)," *IEEE Electron Devices Meeting, IEDM Technical Digest*, pp. 471-474, Dec. 2005.
- [34] Dr. Rob Beckman, Caltech Dept. of Chemistry, personal comm., 2005.
- [35] E. Rachlin and J.E. Savage, "Nanowire Addressing with Randomized-Contact Decoders," *Proc. Int'l Conf. Computer-Aided Design*, Nov. 2006.



Eric Rachlin received the ScB degree with a concentration in applied mathematics and computer science from Brown University, Providence, Rhode Island, in 2003. He is currently working toward the PhD degree in computer science at Brown University. His research interests include computational nanotechnology, automatic music transcription, computer vision, and inference problems involving noisy data. He is a student member of the IEEE.



John E. Savage received the SB, SM, and PhD degrees from the Massachusetts Institute of Technology in 1961, 1962, and 1965, respectively. He joined Bell Laboratories, Holmdel, New Jersey, in January 1965, leaving to join the faculty of Brown University, Providence, Rhode Island, in 1967. He is a founder of the Department of Computer Science at Brown University and was its second chair. He has done research in information theory, coding, communication theory, circuit complexity, space-time trade-offs, I/O complexity, VLSI algorithms and analysis, silicon compilation, parallel algorithms, scientific computing, nanotechnology, and human cognition. His current interests are in computational nanotechnology. He is the author of *The Complexity of Computing* (John Wiley & Sons, 1976) and *Models of Computation* (Addison-Wesley, 1998), and a coauthor, with S. Magidson and A. Stein, of *The Mystical Machine* (Addison-Wesley, 1986). He is a coeditor, with Tom Knight, of *Advanced Research in VLSI and Parallel Systems* (MIT Press, 1992). He is a member of the editorial board of the *Journal of Computer and Systems Sciences*. He is a Guggenheim Fellow, a fellow of the ACM and AAAS, and a life fellow of the IEEE.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.