

Fault-Tolerant Nanoscale Processors on Semiconductor Nanowire Grids

Csaba Andras Moritz, Teng Wang, Pritish Narayanan, Michael Leuchtenburg, Yao Guo, Catherine Dezan, and Mahmoud Bennis

Abstract—Nanoscale processor designs pose new challenges not encountered in the world of conventional CMOS designs and manufacturing. Nanoscale devices based on crossed semiconductor nanowires (NWs) have promising characteristics in addition to providing great density advantage over conventional CMOS devices. This density advantage could, however, be easily lost when assembled into nanoscale systems and especially after techniques dealing with high defect rates and manufacturing related layout/doping constraints are incorporated. Most conventional defect/fault-tolerance techniques are not suitable in nanoscale designs because they are designed for very small defect rates and assume arbitrary layouts for required circuits. Reconfigurable approaches face fundamental challenges including a complex interface between the micro and nano components required for programming. In this paper, we present our work on adding fault-tolerance to all components of a processor implemented on a 2-D semiconductor NW fabric called nanoscale application specific integrated circuits (NASICs). We combine and explore structural redundancy, built-in nanoscale error correcting circuitry, and system-level redundancy techniques and adapt the techniques to the NASIC fabric. Faulty signals caused by defects and other error sources are masked on-the-fly at various levels of granularity. Faults can be masked at up to 15% rates, while maintaining a $7\times$ density advantage compared to an equivalent CMOS processor at projected 18-nm technology. Detailed analysis of yield, density, and area tradeoffs is provided for different error sources and fault distributions.

Index Terms—Defect tolerance, fault tolerance, nanoscale fabrics, nanoscale processors, nanoscale application specific integrated circuit (NASIC), semiconductor nanowires (NWs).

I. INTRODUCTION

THE RECENT progress on manufacturing and assembling of semiconductor nanowires (NWs) is driving researchers to explore possible circuits and architectures. Examples of proposed architectures include [7]–[10].

Manuscript received January 14, 2007; revised May 25, 2007. This work was supported in part by the Center for Hierarchical Manufacturing (CHM), and by NSF under Award CCR:0105516, Award NER:0508382, and Award CCR:0541066. This paper was recommended by Guest Editor C. Lau

C. A. Moritz is with the Department of Electrical and Computer Engineering, University of Massachusetts, Amherst, MA 01003 USA and also with BlueRISC Inc, Amherst, MA 01002 USA (e-mail: andras@ecs.umass.edu).

T. Wang, P. Narayanan, M. Leuchtenburg, and M. Bennis are with the Department of Electrical and Computer Engineering, University of Massachusetts, Amherst, MA 01003 USA.

Y. Guo was with the Department of Electrical and Computer Engineering, University of Massachusetts, Amherst, MA 01003 USA. He is now with Operating Systems Lab, School of Electronics Engineering and Computer Science, Peking University, Beijing 100871, China.

C. Dezan was with the Department of Electrical and Computer Engineering, University of Massachusetts, Amherst, MA 01003 USA. She is now with the Université de Bretagne Occidentale, Brest 93837, France.

Digital Object Identifier 10.1109/TCSI.2007.907839

A fabric architecture proposed based on NWs and targeting datapaths is the nanoscale application specific integrated circuit (NASIC) [13]. NASIC is a tile-based fabric built on 2-D NW grids. Based on NASICs, nanoscale processors are being explored. For example, the Wire Streaming Processor (WISP-0) [14] is a processor design that exercises several NASIC design principles and optimizations. In this paper we use NASIC as the underlying fabric and evaluate the impact of built-in fault-tolerance techniques on WISP-0's yield and area. Additionally, WISP-0's density is compared with an equivalent CMOS version developed with state-of-the-art conventional CAD tools and scaled to projected technologies at the end of the International Technology Roadmap for Semiconductors (ITRS)-defined semiconductor roadmap.

Two main directions have been proposed to handle defects/faults at nanoscale: reconfiguration and built-in fault tolerance. Most conventional built-in defect/fault-tolerance techniques, however, are not suitable in nanoscale designs because they were designed for very small defect rates and assume arbitrary layouts for required circuits. Moreover, the circuits used for fault correction are often assumed to be defect free, which cannot be guaranteed in nanoscale fabrics.

Secondly, if reconfigurable devices are available, defective devices might be replaceable after manufacturing. Reconfiguration based approaches, however, include significant technical challenges: i) highly complex interfaces are required between micro and nano circuits for accessing defect maps and reprogramming around defects—this is considered by many researchers a serious manufacturing challenge due to the alignment requirement of a large number of NWs with programming microwires (MWs); ii) special reconfigurable nanodevices are needed requiring unique materials with programmable and reversible characteristics; iii) an accurate defect map has to be extracted through a limited number of pins from a fabric with perhaps orders of magnitude more devices than in conventional designs.

In addition to the potentially intractable complexity, it might not always be possible to correctly extract such a map from a fabric with very high defect rates. Reconfiguration has been proposed at higher levels (e.g., node level in [15]) where it may not require a fully accurate defect map, assuming that self-checking at node-level is supported. However, the complexity of a node might make the nanoscale implementation almost always defective.

Furthermore, reconfiguration-based approaches would primarily address permanent defects; it might be difficult, if not unfeasible, to work around faults caused by device parameter variations visible only for certain input combinations, or internal/external noise related faults that are transient.

Alternatively, as shown in this paper, we can introduce fault tolerance at various granularities, such as fabric, circuit, and architecture levels, to make nanoscale designs functional even in

the presence of errors, while carefully managing area tradeoffs. Such built-in fault tolerance could possibly address more than just permanent defects. Faults caused by speed irregularities due to device parameter variations, noise, and other transient errors could be potentially masked. Compared with reconfiguration based approaches, this strategy also simplifies the micro-nano interfacing: no access to every crosspoint in the nanoarray is necessary. Furthermore, a defect map is not needed and the devices used do not have to be reconfigurable.

In this paper, we introduce several fault-tolerance techniques into all parts of WISP-0 while simultaneously managing their area efficiency. The fault tolerance approach used is based on both structural/fabric redundancy, built-in error-correcting (EC) circuitry at nanoscale, and CMOS-based voting at the architectural level. Error correction in general has been proposed by other researchers for nanoscale designs [17], [18], however, error correction was used either in memory or at the interface between micro and nanoscale circuits. When used in arithmetic circuits, e.g., with *residue codes* [27], [47], components of the correcting circuitry are often assumed to be defect free, and/or, such as in the case of arithmetic with *stochastic computing* and *serialized data* [36] operand lengths are increased prohibitively.

We are the first to apply an EC technique directly on a logic and fabric with significant layout constraints and the first group to evaluate a nanoscale processor design with a combination of EC, structural, and system-level techniques. The combined techniques make redundant circuits more tuned for specific designs and better tradeoff between area overhead and fault tolerance can be achieved. For example, our simulation results show that a hybrid fault tolerance approach is up to $11\times$ better than 2-way structural redundancy alone in terms of its achieved yield on WISP-0. It gives a 12% improvement at 2% defect rate, a 103% improvement at 5% defect rate, and $11\times$ at 10% defect rate. The improvement in the density-yield product compared to 2-way redundancy alone is 52% at 5% defect rate and $4.2\times$ at 10% defect rate.

We found that the yield of WISP-0 is as high as 20% at 10% defective devices while the density of this design is still $7\times$ denser¹ than of the 18-nm equivalent CMOS processor. Much additional experimental data for various fault rates and error sources is provided. The paper is organized as follows.

In Section II, we provide a brief overview of NASICs and WISP-0. The fault model is described in Section III. Section IV describes the built-in fault tolerance techniques. The yield and density simulation results for WISP-0 with uniformly distributed and clustered faults are provided in Section V. A detailed comparison with a CMOS WISP-0 designed with conventional CAD tools is shown in Section V-B. Section VI shows a sensitivity analysis including the impact of a larger NW pitch on the density of WISP-0. Section VII estimates delay and power consumption. Section VIII discusses related work. Section IX concludes the paper.

II. NASICs AND WISP-0 PROCESSOR

A. Overview of NASICs

NASIC designs use field-effect transistor (FETs) on 2-D semiconductor NWs to implement logic functions; various

¹ $3\times$ when structural redundancy is combined with CMOS TMR. The NW pitch assumed is 10 nm.

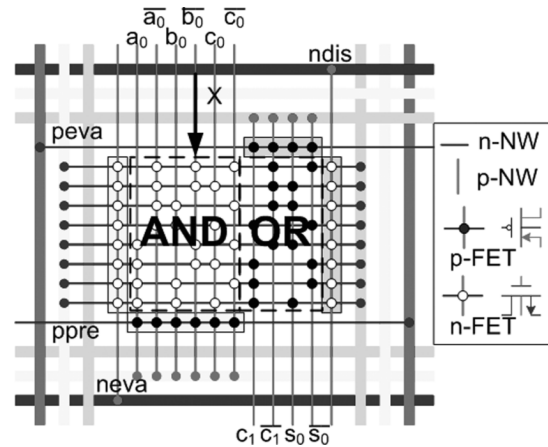


Fig. 1. 1-bit NASIC full adder in dynamic style.

optimizations are applied to work around layout and manufacturing constraints as well as defects [10], [13]. While still based on cascaded 2-level logic style, e.g., AND–OR, NASIC designs are optimized according to specific applications to achieve higher density and defect/fault-masking. The selection of this logic family is due to its simplicity and applicability on a 2-D style fabric where arbitrary placement and routing is not possible. Furthermore, due to manufacturing constraints (such as layout and uniform doping in each NW dimension) it may be impossible to use, for example, complementary devices close to each other, such as in CMOS or orient devices in arbitrary ways. By using dynamic circuits and pipelining on the wires, NASICs eliminate the need for explicit flip-flops in many areas of the design and therefore can improve the density considerably [12].

Fig. 1 demonstrates the design of a simple 1-bit NASIC full adder in dynamic style [14]. The signals *ndis*, *neva*, *ppre*, and *peva*, correspond to discharge, evaluation, precharge and evaluation phases on the different NWs. Each nanotile is surrounded by microwires (MWs) (thicker wires in the figure), which carry V_{dd} , Gnd (or V_{ss}) and control signals for the dynamic style evaluation of outputs. In multi-tile designs, local communication between adjacent nanotiles is provided by NWs. For more details, please refer to [10]–[14].

B. Single-Type Versus Complementary-Type NASICs

In order to produce complementary FETs, two different types of doped NWs must be used. Complementary FETs have been demonstrated in zinc oxide [35], silicon [33], and germanium [34], but in all cases differences in transport properties were found between the two types, sometimes much greater than those seen in today's traditional CMOS FETs. By suitably modifying the NASIC dynamic control scheme and circuit style, we can implement arbitrary logic functions with one type of FETs in NASICs. A design using only n-type FETs will implement a NAND–NAND cascaded scheme whereas a design using only p-type FETs will implement a NOR–NOR scheme. Fundamentally, these are equivalent with the original AND–OR. These schemes may thus be used with manufacturing processes where complementary devices are difficult to achieve. The 1-bit adder example with nFETs is shown in Fig. 2. A detailed analysis of the control scheme for this circuit is beyond the scope of this paper; we refer the interested reader to [16] for more details.

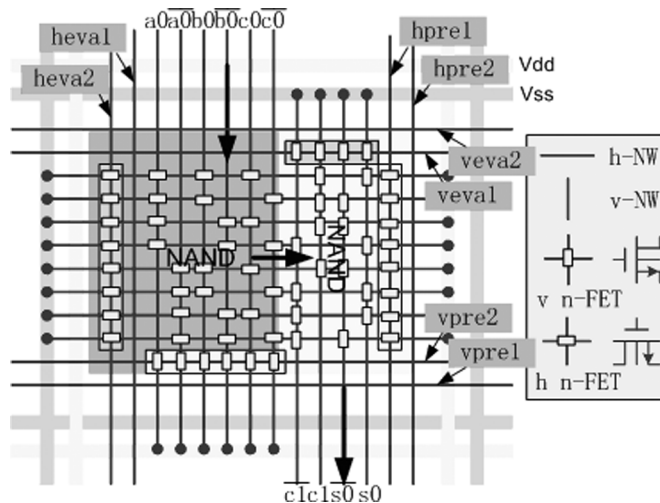


Fig. 2. nFET only version of a 1-bit adder using the NAND–NAND cascaded scheme. The FET channel is oriented along the length of the rectangle in both horizontal and vertical NWs in the figure; arrows show propagation of data through the tile.

C. Overview of the WISP-0 Processor

WISP-0 is a stream processor that implements a 5-stage pipelined streaming architecture. Each stage is implemented in its own nanotile. NWs are used to provide communication between adjacent nanotiles. Each nanotile is surrounded by MWs which carry ground, power supply voltage, and some control signals. Additionally, in order to preserve the density advantages of nanodevices, data is streamed through with minimal control/feedback paths. With the help of dynamic Nanolatches [12], intermediate values during processing are stored on the wire without requiring explicit latching. Support is assumed in the compiler to avoid hazards. WISP-0 uses a 3-bit opcode and 2-bit operands. It supports many different arithmetic operations including multiplication.

Fig. 3 shows the layout. A nanotile is shown as a box surrounded by dashed lines. More details about the various circuits used can be found in [12]–[14]. In this paper, we use WISP-0 to evaluate the efficiency of our fault-tolerance techniques which are added to all circuits.

D. Manufacturing of NASICs

NASIC manufacturing can be done with a combination of self-assembly and more conventional top-down manufacturing steps. It is useful to review this before a fault model can be discussed. NASICs do not require reconfigurable devices.²

The interfacing between the micro and the nano components is therefore limited to input–output (I/O) signals as no programming related interfacing and decoders are needed. Nevertheless, there are a number of other key manufacturing challenges that still remain. To manufacture NASIC fabrics, we envision the following main process steps.

Prepare and align NWs as follows.

- Grow NWs to a certain diameter under the control of seed catalysts [1] or by other methods. During the growth NWs are lightly doped for semiconductivity [2]. For single-type FET NASICs, only one type of doping

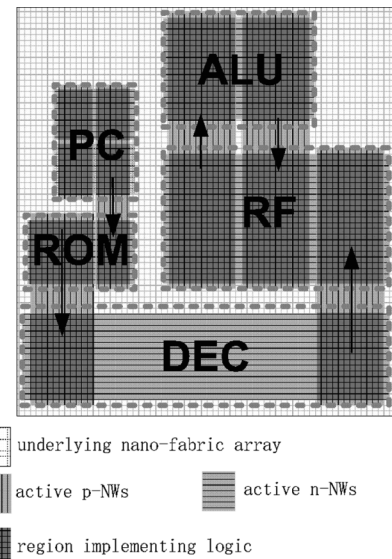


Fig. 3. Floorplan of the WISP-0 processor.

is used for both horizontal and vertical NWs. For NASICs with both types of FETs, each NW set (horizontal versus vertical) will need to be differently doped.

- NWs can be aligned into parallel horizontal and vertical sets with Langmuir–Blodgett techniques [3]. Depending on the NW pitch assumed, other approaches relying on soft lithographic techniques [37] or based on using grooves to align NWs on a substrate might be possible.

Create FETs, metallic interconnect between FET channels, gate regions, and form 2-D NW grid.

- Regions on both the horizontal and vertical NWs—where there should be no FET channels—are first metalized over with the help of a lithographic mask. The resolution required is 2NW pitches (e.g., 20 nm \times 20 nm at a 10-nm NW pitch). While this resolution can be fairly demanding depending on the size of the NW pitch, the shape and size of these regions do not have to be precise. A crosspoint area has a rectangular shape proportional with the NW width—as opposed to the typically larger NW pitch. A metalized crosspoint region can, therefore, be of any shape up to a 2NW \times 2NW square area—beyond that size another crosspoint could be covered causing a defect. This process step is, therefore, likely less challenging than a lithographic process in conventional CMOS with a similar feature size requiring exact shapes, sizes, and straight edges. Lithographic techniques with a resolution required for this step have been reported in [4], [5]. Nevertheless, we expect this process step to be a key factor in determining the actual NW pitch that can be manufactured. The misalignment of this lithographic mask could generate stuck-short defects, e.g., when some FET channels, that should normally be part of the design, are metalized over. As will be shown in the following sections, these defects can be masked fairly well with a combination of built-in fault-tolerance techniques. In the evaluation section, we also explore the impact of larger NW pitches on the density of the WISP-0 design. A larger NW pitch could facilitate

²Some of our earlier papers on NASICs assumed reconfigurable FETs. However, if built-in fault tolerance is added that is not necessary.

manufacturing designs, even before all process steps are worked out.

- Metallization of the NW gate regions can be done for each set of NWs in conjunction with the previous metallization step. The required resolution for gate regions is fairly low as each logic plane will have either its entire horizontal or all its vertical NWs acting as gates. After being metalized, the gate regions will need to be covered with an oxide shell. Once this step is completed, a 2-D NW grid can be assembled by moving one NW set on top of the other.
- A fine-grained metallization step is essentially responsible for creating the FET channels, creating the metallic interconnects between the FETs, and extending the metallic segments created in the earlier metallization step. Before this step, the assembled 2-D NW grid contains some metallic regions corresponding to: i) crosspoints where no FET channels are needed and ii) gate regions; other segments of the NWs remain doped as required for the FET channels. FET channels can be distinguished at the crosspoints by using one layer of NWs as a fine-grained mask over the other layer during a final metallization step. This step needs to be completed for both dimensions of a nanogrid—flipping of the structure might be required. After this, channels are formed at grid crosspoints (see, for example, the process in [6] with NiSi), in both dimensions, because the top layer protects the bottom NW from being metalized over; at the same time, the FET channels become automatically connected with small metallic NW segments. Crosspoint regions that have already been metalized in the previous step would remain metallic and would not be affected by this step.

MWs and contacts.

- Can be added with lithographic process steps.

As discussed in this section, while key individual steps have been demonstrated in laboratory settings (e.g., FETs at NW crosspoints, NW growth and specialization, NW alignment, and fine-grained metallization with the help of NWs to create FET channels), combining the necessary manufacturing steps remains a challenging and unproven process. By working on nanoscale fabrics and architectures, the research community can, however, expose these requirements and tradeoffs between manufacturability and system-level capabilities, fueling more focused research on manufacturing techniques required for assembling nanoscale systems. More on the manufacturing related differences between various proposed nanoscale fabrics is discussed in Section VIII-B.

III. SOURCES OF ERROR AND FAULT MODEL IN NASICS

A. Types and Sources of Error

Sources of error include permanent defects, process and environmental variation related errors, transient errors, as well as internal and external noise related ones.

Permanent defects are mainly caused by the manufacturing process. The small NW dimensions combined with the self-assembly process, driven by the promise of cheaper manufacturing, is expected to contribute to high defect rates in nanoscale designs. Examples of permanent defects in NASIC fabrics would include malfunctioning FET devices, broken

NWs, bridging faults between NWs, and contact problems between controlling MWs and NWs. For example, in a process that requires the metallization of segments connecting NASIC FETs, the channels of transistors could be metalized over and therefore stuck-on. The NWs used as gate control have a core-shell structure [22] and, therefore, if a shell is thicker than expected, the FETs controlled by these gates may have no bias applied. Prevalent defect types are also dependent on the types of transistors used. The FET channels will be conducting for depletion mode FETs [19] but will be cut off for enhancement-mode FETs [20]. This means that when the FET has no bias applied it would be either always conducting (easier to tolerate) or would be cut off (much harder to tolerate) depending on its type.

Process variation related errors are caused by speed deviations due to device parameter variations. These errors occur typically for certain input combinations as a result of larger than expected circuit delays for those input combinations. While the actual parameter variation in NASIC depends on the manufacturing process ultimately used (so this data is currently not available), research from deep sub-micron CMOS technology underlines the seriousness of this problem. We project that delay variations in NASICs would be caused by doping variations on the NWs used for channels and by channel length variations caused by the metallization process that separates FETs from each other (by creating small metallic interconnects between them) and they could be fairly significant.

Internal noise related faults caused by higher frequency and crosstalk between NWs are to be expected in fabrics like NASICs where NWs are placed close to each other. The NASIC control and the dynamic logic used could also affect noise margins. External noise factors such as radiation could be also present: with small dimensions, there might be an increasing likelihood that an α -particle, neutron or proton hitting the chip would cause transient faults. Other noise sources such as electromagnetic interference and electrostatic discharge could cause permanent faults [37].

Overall, we expect that these faults and process variation related ones will be less of a problem in NASICs compared to manufacturing defects, but factors to account for nevertheless. Our objective in the NASIC project is to address all these different sources of errors in a uniform manner with built-in fault tolerance techniques at fabric, circuit, and architecture levels. This paper is a snapshot of our efforts to date.

B. Fault Model Assumed

In NASICs we consider a fairly generic model with both uniform and clustered defects and three main types of permanent defects: NWs may be broken, the transistors at the crosspoints may be stuck-on (no active transistor at crosspoint) or stuck-off (channel is switched off). A stuck-off transistor can also be treated as a broken NW. The initial thinking is that the more common defect type is due to stuck-on FETs as a consequence of the metallization process used. NASIC fabrics require a mask at a 2NW pitch for one of their metallization steps (to avoid channels at crosspoints where no FETs are placed). Stuck-off FETs are also less likely especially in depletion mode fabrics. Recent thinking from [24] suggests that we will be able to control the reliability of NWs fairly well so broken NWs will be likely less frequent than stuck-on FETs.

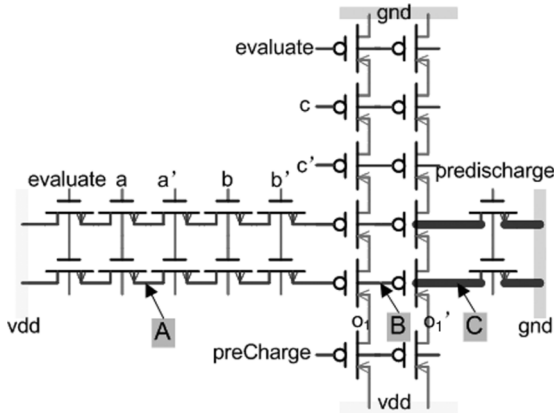


Fig. 4. Simple NASIC circuit with built-in redundancy.

In this paper, we consider defect rates of up to 15%. As suggested by other researchers, the defect levels in nanofabrics are in a few percent range [2]. During our initial work we found that defect rates greater than 15% would likely eliminate the density benefits of nanoscale fabrics compared to projected CMOS, in the context of microprocessor designs. Fabrics with higher defect rates might still be applicable as replacement technology for field-programmable gate arrays (FPGAs) and structured ASICs: e.g., if look-up tables for programming of interconnect in FPGAs could be replaced with programmable devices; the lost density due to high defect rates will likely be offset.

In addition to permanent defects, other error sources such as due to process variation and transient faults are also discussed. Both uniformly distributed and clustered faults are modeled.

IV. BUILT-IN FAULT-TOLERANCE IN NASICS

A. Circuit-Level and Structural Redundancy

Fig. 4 shows a simple example of a NASIC circuit implementing an AND–OR logic function with built-in redundancy: redundant copies of NWs are added and redundant signals are created and logically merged in the logic planes with the regular signals. To make the masking mechanism work, we also modify the dynamic circuit style reported in our prior work [12]. We use different clocking schemes for horizontal and vertical NWs: this, we have found empirically to yield better results. As shown in Fig. 4, horizontal NWs are predischarged to “0” and then evaluated. Vertical NWs are instead precharged to “1” and then evaluated. The circuit implements the logic function $o_1 = ab + c$; a' is the redundant copy of a and so on. Signal a and a' are called a NW pair.

A NASIC design is effectively a connected chain of AND–OR (or equivalent) logic planes. Our objective is to mask defects/faults either in the logic stage where they occur or following ones. For example, a break on a horizontal NW in the AND plane (see, for example, position “A” in the figure) causes the signal on the NW to be “0.” This is because the NW is disconnected from V_{dd} . The faulty “0” signal can, however, be masked by the following logic OR plane if the corresponding duplicated/redundant NW is not defective.

A NW break at position “B” can be masked by the AND plane in the next stage. Similar masking can be achieved for breaks on vertical NWs. Stuck-off FETs can be modeled as broken NWs; the defect tolerance would work as described above. For

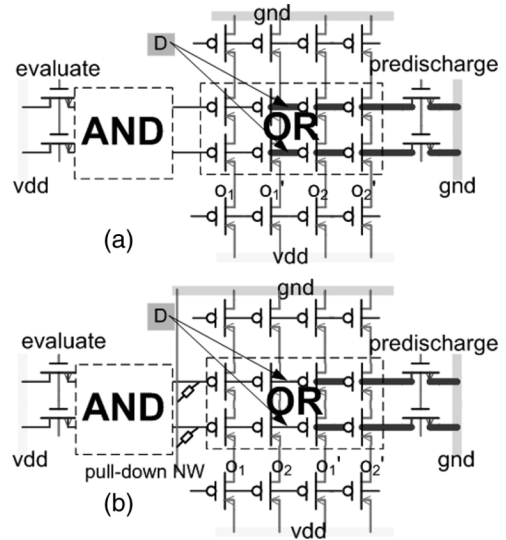


Fig. 5. Interleaving NWs and adding weak pull-up/down NWs to reduce hard-to-mask regions. The bottom circuit has interleaved vertical NWs and weak pull-down NW between the AND and OR planes.

stuck-on FETs, the situation is relatively simpler as each FET has its redundant copy: if one of the two transistors is stuck-on, the circuit still works.

B. Improving Fault-Tolerance by Interleaving NWs

While the previous technique can mask many types of defects, faults at certain positions are difficult to mask. For example, if there is a break at position “C” in Fig. 4, the bottom horizontal NW is disconnected from ground preventing predischarge. The signal on this NW may potentially retain a logic “1” from a previous evaluation. Because of OR logic on the vertical NWs, the two vertical NWs would then be set to logic “1.” Since both outputs on the vertical NWs are faulty, the error cannot be masked in the next stage. In Fig. 4, the thicker segments along the horizontal NWs show the locations at which faults are difficult to mask. We call these segments *hard-to-mask segments*.

For nanotiles with multiple outputs, a particular arrangement of output NWs and their redundant copies could significantly reduce the size of hard-to mask segments. This is shown in Fig. 5: Fig. 5(a) presents a design in which each output NW and its redundant copy are adjacent to each other. In this arrangement, all segments to the right of the leftmost output NW pair [o_1 and o'_1 in Fig. 5(a)] are hard to mask. Alternatively, the interleaved version in Fig. 5(b), shows an arrangement in which the output NWs and their redundant copies are separated into two groups (o_1 and o_2 form one group; o'_1 and o'_2 form another group). In this case, the size of the hard-to-mask segments is reduced. In general, the size of hard-to-mask segments can be reduced in larger scale designs to half, i.e., to half of the region covered by the vertical NWs plus the segment related to the control FET. This latter region is fixed and for most designs adds a negligible area. Interleaving is also helpful in masking clustered defects because duplicated NWs are set apart from one another.

C. Adding Weak Pull-Up/Down NWs

Even after built-in redundancy and careful interleaving, there are still some hard-to-mask segments remaining: for example, the thick lines in Fig. 5(b). A possible solution to mitigate this

problem is to insert weak pull-down vertical wires between the AND and OR planes. The idea is to pull down (or up depending on logic plane) floating inputs, due to broken NWs, that would cause logic faults: e.g., a floating “1” input to an OR plane that would make the OR logic always compute “1.” Modifying floating signals to a preferred logic level would allow masking in following logic planes.

A weak pull-down NW does not affect operation if there are no defects, but introduces a performance tradeoff when there are defects, by slowing the circuit down somewhat. It also contributes to leakage power. At each crosspoint between a vertical pull-down wire and horizontal NWs there is a resistance created. This resistance has to be made larger than the switch-on resistance (estimated to be smaller than $10\text{ M}\Omega$ according to [2], [3]) of a depletion-mode FET and smaller than the switch-off resistance (over $10\text{ G}\Omega$). We are currently building a detailed SPICE simulator that would enable us to explore the performance tradeoffs due to these added wires in more detail. To ease manufacturing one could also use MWs instead of the NWs to implement weak pull-up/down wires.

D. Adding CMOS TMR

Voting-based techniques such as triple modular redundancy (TMR) [30] have been used extensively before. To be efficient, voting requires that the probability of a defect in the voting circuit is much smaller than in the design it is applied to. This is clearly the case in conventional technology. TMR is not applicable as is in NASIC designs because at 5%–15% fabric defect rates the TMR circuits themselves would be likely defective.

Nevertheless, in pipelined processor designs one could add TMR, e.g., with majority voting, at certain points in a design in CMOS, without affecting throughput significantly. If each nanotile has two extra identical replicas, we could vote either at each stage or on the final outputs. Voting helps where the other nanoscale techniques leave faulty outputs.

E. Nanoscale EC Circuits

1) *Hamming Distance*: The Hamming distance between two input codes is defined as the number of bits that is different. For example, the Hamming distance between “000” and “001” is 1. For the simple 1-bit adder design in Fig. 1, the minimum Hamming distance between the input codes is 1. Therefore, in that example, we cannot tolerate any defect on vertical NWs.

By adding redundant bits to the input signals, we are able to increase the minimum Hamming distance of input codes. In the 2-way redundancy example shown in Fig. 4, the input codes are simply duplicated and the Hamming distance is increased to 2. With a minimum Hamming distance of 2, the design with 2-way redundancy can tolerate 1-bit error on the input signals. In the following subsection, we will show the required circuit-level modification to achieve error-correction with built-in EC circuits and redundant code signals, for a more efficient defect masking.

2) *EC Code Background*: Achieving a certain Hamming distance between codes with minimum redundant bits is a well-known problem in the communication area. These codes called as EC codes are widely used to correct signal errors in noisy

channels. Various kinds of EC codes have been proposed and used; the Hamming code is one of the most popular codes due to its simplicity [23].

Considering a set of 3-bit codes {“000,” “001,” “010,” “011,” “100,” “101,” “110,” “111”}, the minimum Hamming distance between these codes is 1. By adding three redundant bits to the codes, we can achieve a Hamming distance of 3. The redundant bits (shown in parentheses below) are not unique according to the coding theory. An example of a new code set is {“(000)000,” “(011)001,” “(101)010,” “(110)011,” “(110)100,” “(101)101,” “(011)110,” “(000)111”}. Obviously, this code set is more efficient than the one created by a simple signal duplication used in 2-way redundancy—which achieves a Hamming distance of 2 similarly with 3 added redundant bits.

In general, the number of required redundant bits is determined by the desired Hamming distance and the code width. For a given Hamming distance, the EC code rate, defined as the ratio between the original signal width and the width of all signals including redundant ones, approaches 1 as the original signal width gets large [23]—which means the relative overhead goes down. For example, 11-bit wide signals would only need 4 redundant bits to achieve a Hamming distance of 3.

Note that in traditional coding theory, codes for a 1-bit error correcting require a Hamming distance of 3; codes for 2-bit error correcting require a Hamming distance of 5. In general, codes for n -bit error correcting require a Hamming distance of $2n + 1$ [23]. In NASICs, however, with Hamming distance of n we can tolerate $n - 1$ defects on vertical NWs. This is because in the case of permanent defects any input combination can only be impacted in the same bit positions.

This paper focuses on Hamming codes; we are currently also exploring a variety of other techniques such as based on BCH codes [44], [45].

3) *EC in NASICs*: To apply the EC technique in NASICs, redundant bits are added to original input signals for the desired Hamming distance. Next, EC related FETs are added so as to keep the output signals the same as outputs in original designs.

The following simple circuit in Fig. 6 (OR plane is omitted for clarity) shows how to add error-correction to a NASIC circuit. The AND logic outputs $\sim c$ on the top horizontal NW and c on the bottom NW. It is easy to see that one single defect makes the output faulty: e.g., the defect shown on the right vertical NW forces the output on the top horizontal NW to logic “1” [Fig. 6(a)] for all input values. The output is set during evaluation (*neva* is turned on).

To apply EC, as shown in Fig. 6(b), we add 2 redundant bits (a and b and their complementary forms) to the original input signals c and $\sim c$. The values of a and b are related to the value of c . In this example, we choose “110” and “001” as possible input combinations with a Hamming distance of 3. We then add redundant vertical NWs for the redundant inputs. At each new crosspoint [shown in the shadowed area in Fig. 6(b)], we place a FET only if it does not impact the correct outputs. For example, the output signal on the top horizontal NW should be “1” when c is “0.” Based on the input combinations we choose, a and b are “1,” so we place 2 FETs at the corresponding crosspoints [shown as nFETs on the top horizontal NW in the shadowed area in Fig. 6(b)]. We can similarly set the crosspoints for the second horizontal NW in the shadowed area. As mentioned, the added overhead is of course smaller for larger designs.




Fig. 6. Simple NASIC circuit. (a) Original design without defect-tolerance. (b) Design with the built-in EC technique.




Fig. 7. 1-bit NASIC full adder with EC. The circuits in the shadowed area are redundant circuits added for the purpose of error correction.

Let us analyze why this design can tolerate 2-bit errors on vertical NWs. For example, assuming the input combination is “001,” the output signal on the top horizontal NW should be equal to “0” ($\sim c$). If we, however, add 2 breaks on the vertical NWs b and $\sim c$ [indicated by “X”s in Fig. 6(b)], the signals on NWs b and $\sim c$ will be set to faulty “1” because they are disconnected from Gnd. As a result, the FETs shown in the circles in Fig. 6(b) will be switched on permanently. Without the added circuits, the output signal on the top horizontal NW would be forced to faulty “1.” However, the redundant signal a (“0” in the example) forces the output signal on the top horizontal NW to a correct “0.” Similar analysis can be made for other input combinations. Clearly, we can guarantee the correct output signals on horizontal NWs even when any two vertical NWs have defects. The key insight here is that *the added FETs in the EC circuit take over the role of any of the original FETs in case they would become faulty or have incorrect input(s) and would because of that not be able to affect the output.*

With a Hamming distance of 3, the circuit in Fig. 6(b) can tolerate any 2 defects on vertical NWs.

4) *1-bit NASIC Adder With EC:* We apply EC on the 1-bit NASIC adder using the method described above. The new adder is shown in Fig. 7.

Three redundant bits (r_1, r_2 and r_3) are added for a Hamming distance of 3. EC related FETs for these 3 redundant bits

are shown in the left-side shadowed area. Circuits in the left-side shadowed area are actually helping in providing the correct output on each horizontal NW (input to the OR plane); the right-hand side shadowed area is used to generate redundant signals for the EC circuits in the next stage. This example also shows how EC can be applied in cascaded circuits.

5) *EC Combined With 2-Way Redundancy:* There is one issue with the EC technique: complementary signals are required for proper functionality. However the product-term signals on horizontal NWs are not complementary. Thus, it may not be feasible to apply the EC technique for defects on horizontal NWs. Creating a complementary version for each product-term is not feasible on a 2-D fabric with this type of 2-level logic—we are currently investigating other logic style based on mixed AND/NAND–OR/NOR logic in the same tile where this might be possible. For the time being, we therefore apply 2-way redundancy techniques on horizontal NWs. As will be shown in the next section, the yield of WISP-0 can still be improved considerably with this hybrid approach.

V. EVALUATION

Using the design approaches described in Section IV, we can incorporate the techniques into all circuits of WISP-0 [14]. We used our NASIC CAD tools to modify WISP-0. To verify the efficiency of our fault-tolerance approaches, we developed a sim-

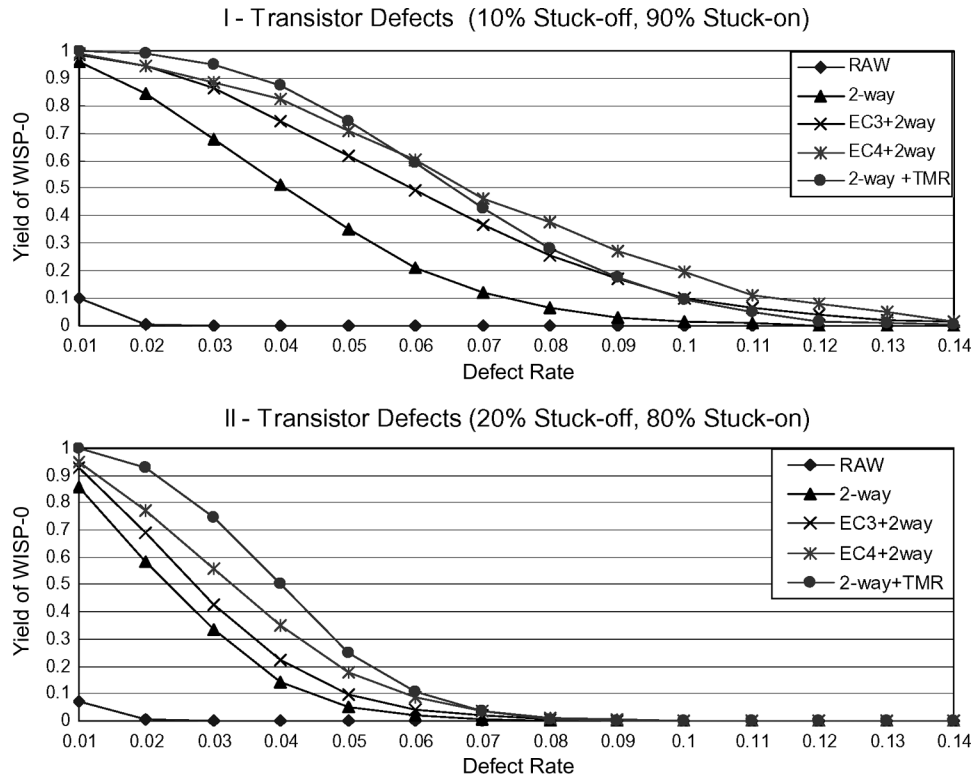


Fig. 8. Yield achieved for WISP-0 with different techniques when only considering defective transistors.

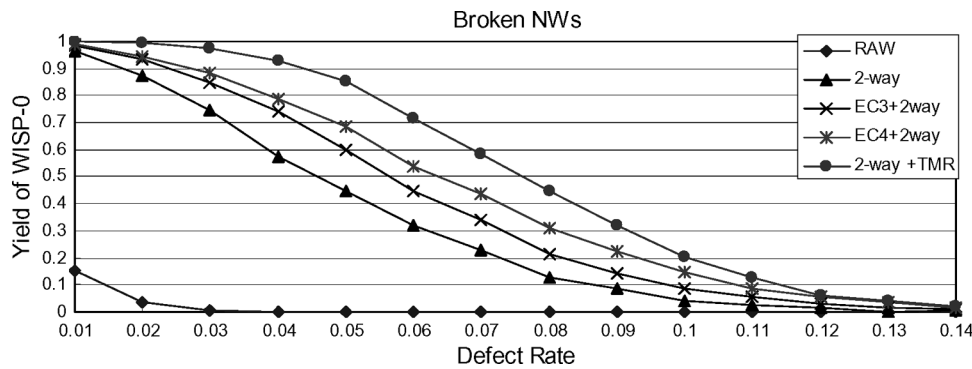


Fig. 9. Yield achieved for WISP-0 with different techniques when only considering broken NWs.

ulator to estimate the yield of WISP-0 for different defect rates and also considered other error sources.

A. Yield Evaluation of WISP-0

The simulation results for permanent defects are provided in Fig. 8 (assumes defective FETs) and Fig. 9 (assumes broken NWs). First we present results assuming defects are uniformly distributed. Clustered defects are addressed in separate subsequent subsections.

The notation used is: RAW stands for WISP-0 without redundancy (or baseline); 2-way stands for WISP-0 with 2-way redundancy; 2-way + TMR stands for 2-way redundancy plus micro-scale TMR on the WISP-0 result; EC3 + 2way denotes a design with EC using a Hamming distance of 3 on vertical NWs and 2-way redundancy on horizontal NWs; and EC4 + 2way denotes EC with a Hamming distance of 4 on vertical NWs and

2-way redundancy on horizontal NWs. While other combinations are possible, we found these to be most insightful and representative. The 2-way redundancy techniques also incorporate the techniques discussed in Sections IV-B and IV-C

From the results, we can see that EC-based techniques achieve the best overall yield. Compared with a 2-way redundancy approach, the improvement of the hybrid approach (EC3 + 2-way) on the yield of WISP-0 is 12% when the defect rate of transistors is at 2%, 76% at 5% defect rate, and 5x at 10%. Note that the improvement is greater for higher defect rates.

As expected, EC with a Hamming distance of 4 (EC4) on vertical NWs, achieves a better yield compared to EC3. The improvement compared to 2-way is 12% when the defect rate of transistors is at 2%, 103% at 5% defect rate, and 11x at 10%. However, the rate of improvement is not as significant as for the EC3 version—especially when the defect rate of transistors is less than 10%.

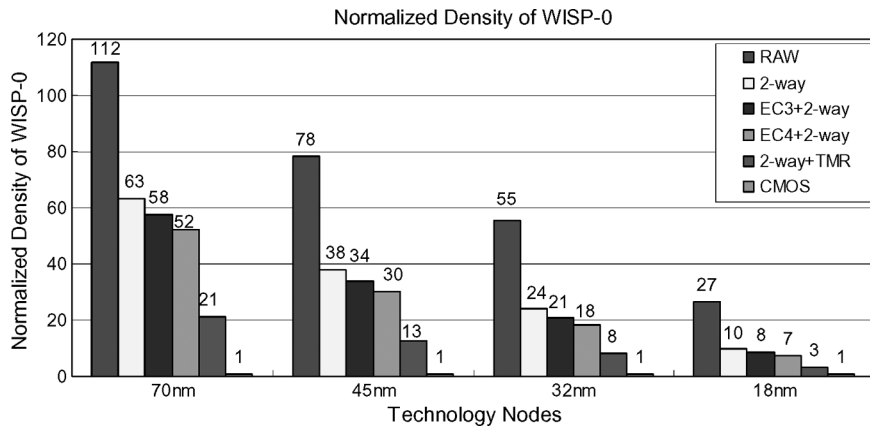


Fig. 10. WISP-0 density with different defect tolerance techniques.

TABLE I
TECHNOLOGY PARAMETERS

NW pitch	10nm
NW width	3–4nm
Technology Node (ITRS 2005)	MW pitch
70-nm	170nm
45-nm	108nm
32-nm	76nm
18-nm	42nm

One possible explanation is that the likelihood of 3-bit errors on vertical NWs is relatively small compared to 1 or 2-bit errors for these rates, so the approach starts to have diminishing returns despite the greater Hamming distance.

We simulated two different distributions of defective transistors; we assumed that the stuck-on FETs are more prevalent and simulated a relatively smaller fraction of stuck-off defects (10% and 20% respectively) for the reasons we discussed in Section III. In Fig. 8 (bottom graph), we can see that our techniques are more efficient for stuck-on defects than for stuck-off defects. EC based approaches perform well for defects based on broken NWs but not as good as the 2-way + TMR combination. Similar to the case with 20% stuck-off FETs, broken NWs are difficult to mask. However, as discussed in Section III, we project stuck-off FET defects and broken NWs to be less prevalent than stuck-on FETs.

Some defect-masking techniques provide good yield improvement but require relatively large area overhead. For example, as shown in Fig. 8 and Fig. 9 micro-scale TMR (implemented in CMOS at the output of WISP-0) combined with 2-way redundancy achieves a somewhat higher yield than EC3 + 2-way in some scenarios. This comes, however, at a cost of a $2.67\times$ larger area than with EC3 + 2-way (density results will be detailed in the following section). Therefore, it is important to understand the area overhead (or impact on density) of the different fault-tolerance techniques in conjunction with their fault masking ability.

B. Comparison With Equivalent CMOS Processor

The normalized density of WISP-0 for the various scenarios is shown in Fig. 10. Technology parameters used in the calculations are listed in Table I. To get a better sense of what the densities actually mean we show the density of an equivalent WISP-0 processor. We designed this processor in Verilog, synthesized it

to 180-nm CMOS. We derive the area with the help of the Synopsys Design Compiler tool. Next, we scaled it to various projected technology nodes based on the predicted parameters by ITRS, assuming area scales down quadratically. For the purpose of this paper, we assume that the CMOS version of WISP-0 is defect-free and no fault-tolerance technique is applied.

We can see from the results that the area overhead of adding 2-way redundancy for the nanoscale designs is roughly $3\times$ when MWs in NASICs are assumed to have the same dimensions as MWs would have in 18-nm CMOS technology. TMR-related overhead added to the nanoscale design brings an extra $3\times$ overhead because TMR requires 3 copies of nanoscale blocks. A WISP-0 design based on EC3 + 2-way requires around 20% more area than one based on 2-way redundancy for both horizontal and vertical NWs, but achieves a much better yield.

Overall, the density of a NASIC based WISP-0 remains at least $3\times$ (without EC but with TMR) or $7\times$ (with EC) greater than the density of the corresponding CMOS processor at 18 nm.

C. WISP-0 Density-Yield Product Evaluation

To evaluate the tradeoff between yield improvement and area, we also consider the yield and density together in a combined metric. The yield–density product is a comprehensive indicator for the efficiency of different defect-tolerance techniques; it represents the ratio between the benefit (yield of designs) and its cost (area overhead).

The yield–density product results for various defect rates are presented in Figs. 11 and 12, respectively. We can see that the EC-based approaches, EC3 + 2way and EC4 + 2way, are significantly more efficient than the other approaches, except for relatively small defect rates. Compared to 2-way redundancy, an approach based on EC3 + 2way improves the yield–density product by 52% when the defect rate of FETs is 5% and by $4.2\times$ for a 10% rate. Clearly, different levels of defect rates may require different defect-tolerance techniques: for defect rates lower than 3%, 2-way redundancy appears to be sufficient.

When defect rates increase beyond 3%, EC with a Hamming distance of 3 is desirable. If the defect rate is larger than 5%, EC with a Hamming distance of 4 is the best choice. Future NASIC CAD tools can take advantage of this and insert appropriate levels of defect tolerance depending on expected defect rates.




Fig. 11. WISP-0 yield-density products considering defective FETs.




Fig. 12. WISP-0 yield-density products considering broken NWs.




Fig. 13. A simple defect model for clustered defects; shows how defect probabilities are decreasing for FETs and NWs further away from a cluster center.

D. NASICs With Clustered Defects

In our previous results we assumed that all defects are uniformly distributed. However, defects can also be clustered as a group of adjacent FETs or a group of adjacent NWs could be damaged during the manufacturing process. In a 2-way redundancy scheme, if clustered defects make two redundant signals faulty, these faults cannot be masked. However, if the same two redundant signals are placed far-enough apart, clustered defects will unlikely make them faulty simultaneously.

To evaluate the impact of clustered defects, we first introduce a model for clustered defects. First, we set a probability for defect clusters or *cluster rate*. FETs belonging to clusters would have greater probabilities to be defective than in defect models based on uniformly distributed defects. Intuitively, the probability of a FET being defective decreases with increasing distances from the center of the cluster it belongs to.

Fig. 13 shows how the probability of defects is modeled in a cluster. Parameters of this model include a , representing the probability of defects in nodes adjacent to cluster centers, and

n representing the maximum distance between the outmost defective transistors or NWs and the center; n also determines the size of clusters.

E. WISP-0 Yield With Clustered Defects

Fig. 14 shows the yield of WISP-0 assuming clustered transistor defects; Fig. 15 shows the yield with clustered broken NWs. The results indicate that our defect-tolerance techniques also work for clustered defects/faults: the yield remains at around 20% even when the cluster rate of transistors is 5% for the parameters simulated. Note that each defect cluster may have multiple defects.

The yield-density product of WISP-0 for clustered defects is shown in Figs. 16 and 17. While the microscale TMR combined with 2-way redundancy (2-way + TMR) gives a somewhat higher yield than EC3+2-way (see Figs. 14 and 15), it achieves a lower yield-density product due to its significantly higher area overhead.

F. Impact of Transient Errors

We extended the yield simulator to provide an initial analysis on the benefits of the built-in fault tolerance techniques for transient errors. This is shown in Fig. 18. The results indicate that we could tolerate transient faults fairly well although the masking is less effective than for permanent defects. On the other hand, we expect these errors to be much less frequent than those caused by permanent defects. One insight is that the system-level TMR appears to have the best overall benefit for these types of errors. The reason is that as these errors are random and transient, if an error does not occur at the same time and same position across at

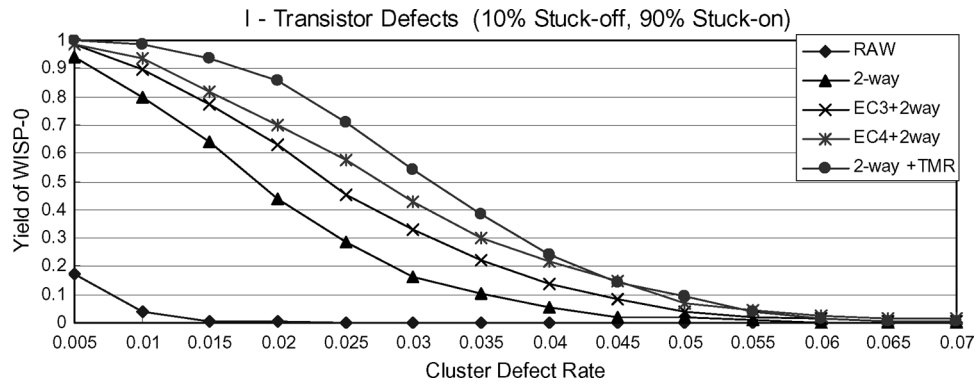


Fig. 14. WISP-0's yield for various cluster rates assuming defective transistors; clustered defects with parameters $a = 0.2$ and $n = 2$.




Fig. 15. WISP-0's yield for various cluster rates when considering broken NWs; clustered defects with parameters $a = 0.2$ and $n = 2$.

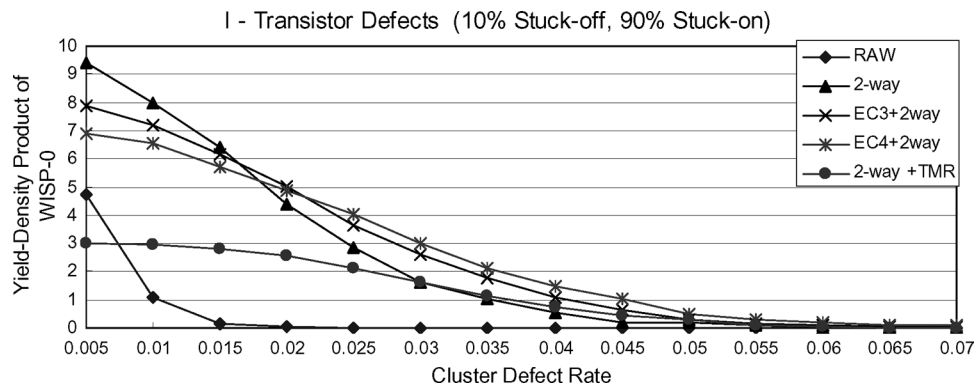


Fig. 16. Yield-density product achieved for WISP-0 considering defective transistors; clustered defects with parameters $a = 0.2$ and $n = 2$.

least 2 copies, the system-level TMR voting could mask it—assuming that other errors are corrected.

G. Impact of Device Parameter Variation

The actual parameter variation for devices used in NASICs is not known as yet. We can predict, however, based on deep sub-micron CMOS processes, that process variation could cause significant variations in the parameters of semiconductor NW devices. Device parameter variation can impact a circuit's speed/delay, by making certain execution paths longer than expected. Delay variation related faults are in many ways similar to those caused by permanent defects except that they would be limited to certain input combinations (using the circuit paths with longer than acceptable delays). One can argue that the techniques presented in this paper would therefore be able to address such faults. In fact, we estimate that we would be able to mask a higher rate of faults caused by device parameter variations than

due to permanent defects, as only a subset of inputs would cause errors as opposed to all inputs. As part of our future work, we plan to model delay variation in NASIC circuits for an exact analysis of the built-in fault tolerance techniques for these types of faults.

VI. SENSITIVITY ANALYSIS

A. Impact of NW Pitch on Density

In the previous analyses, we assumed that the pitch between NWs is 10 nm. While this has been demonstrated in the laboratory, it will take time until we can reliably manufacture larger designs at this scale (the same way as it took the semiconductor industry decades to refine lithography to today's resolution). A larger NW pitch may come with lower defect rates and it will also be significantly easier to manufacture. For example, a 20-nm pitch design would require the NASIC metallization

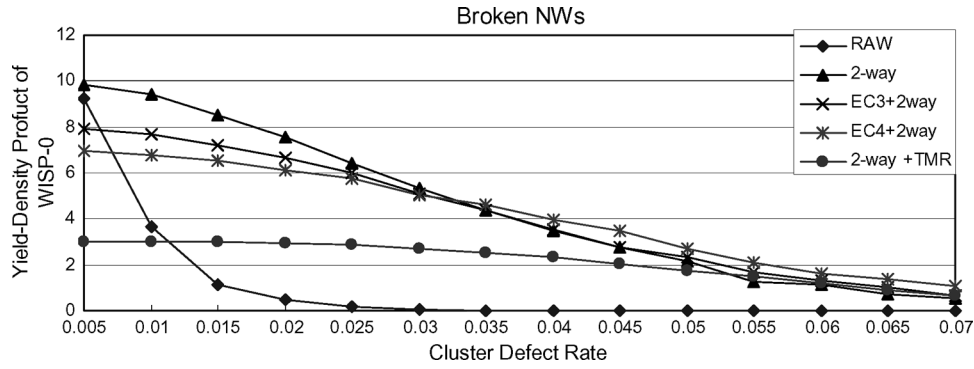


Fig. 17. Yield-density product achieved for WISP-0 when only considering broken NWs; clustered defects with parameters $a = 0.2$ and $n = 2$.

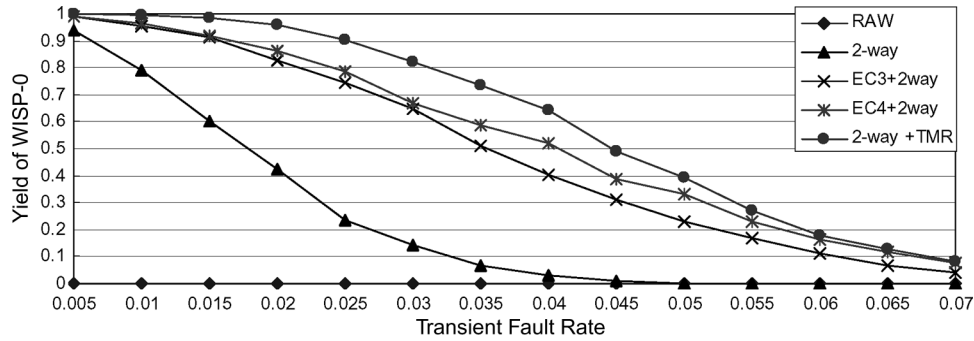


Fig. 18. Yield achieved assuming transient faults.

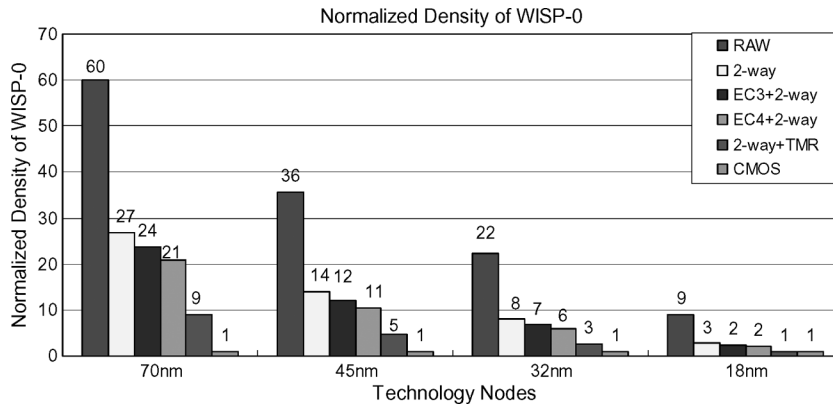


Fig. 19. Density comparison between NASIC WISP-0 assuming a 20-nm NW pitch and an equivalent CMOS WISP-0.

masks at 40-nm resolution: a much more doable undertaking than 20 nm. On the other hand, as expected, a larger NW pitch will result in lower overall density so it is important to understand its impact at the system level.

The impact of a 20-nm NW pitch on density is presented in Fig. 19. Note that the density of WISP-0 with any of the EC-based approaches is still $2\times$ better than 18-nm CMOS technology. This is a result of a high density interconnect structure combined with high-density logic in a NW-based fabric. A plausible option might be to start manufacturing at a relatively lower density and gradually scale with improvements in nano manufacturing.

VII. DELAY AND POWER ESTIMATES

Delay and power estimation was done for the WISP-0 processor built on Silicon NWs.

A NW-MW contact resistance of $10\text{ k}\Omega$ and resistivity values of $10^{-7}\ \Omega\cdot\text{m}$ and $10^{-5}\ \Omega\cdot\text{m}$ for NiSi and Si respectively were

used in these calculations [21]. R_{ON} for a transistor of length 5 nm and width 4 nm was calculated to be around 4 k Ω . An R_{OFF} resistance of 10 G Ω was used [7]. A NW pitch of 10 nm, an oxide layer thickness of 1 nm, and a dielectric constant of 2.2 were assumed. Table II summarizes all the parameter values used in these calculations.

A. Delay Calculations

A lumped RC model was used for the worst-case delay analysis. Expressions from [7] were used for capacitance estimation. These calculations take into account NW-NW junction capacitances and relatively realistic coupling scenarios. The coupling capacitance per unit length was found to be 39.04 pF/m. The junction capacitance was found to be 0.652 aF.

Table III indicates the capacitive loading on each tile of WISP-0 for different clock phases. During each phase, there is one control NW and one or more datapath NWs switching. In the table “Control NW (H)” refers to a Horizontal

TABLE II
PARAMETER VALUES

NW-pitch	10nm
NW-shell thickness (t_{sh})	1nm
NW-width (w)	4nm
Dielectric Constant of SiO ₂ (ϵ_r)	2.2
Resistivity of Si (ρ_{Si})	$10^{-5} \Omega m$
Resistivity of NiSi (ρ_{NiSi})	$10^{-7} \Omega m$
NW-MW contact Resistance (R_c)	10 k Ω
Transistor ON Resistance (R_{ON})	4 k Ω
Transistor OFF Resistance (R_{OFF})	10 G Ω
Supply Voltage	3V-4.5V

TABLE III
CAPACITIVE LOADING (IN aF)

	Control NW(H)		Datapath NW(H)		Control NW(V)		Datapath NW(V)	
	pre/eva	pre	eva	pre/eva	pre	eva	pre	eva
PC	14.99	9.78	25.27	11.08	4.56	32.43		
ROM	8.48	11.08	33.47	9.78	20.12	82.68		
DEC	11.74	20.21	83.33	11.74	55.42	143.1		
RF	27.38	26.73	98.21	9.13	42.38	167.6		
ALU	29.34	18.26	37.78	16.95	30.64	138.7		

precharge/evaluate signal. Since the precharge and evaluate control NWs in one plane are geometrically identical, the capacitive loading on these NWs is the same. “Datapath NW (V)” refers to datapath NWs in the vertical plane. The capacitive loading during precharge and evaluate is dissimilar for datapaths owing to different lengths and coupling effects.

The lumped capacitance is in the range of atto-Farads, and as expected, larger components such as the Register File (RF) are more heavily loaded. Table IV shows the maximum delay for the tiles of WISP-0 for a MW-NW contact resistance of 10 k Ω . “H-pre” and “V-pre” stand for horizontal and vertical precharge phases respectively, “H-eva” and “V-eva” are horizontal and vertical evaluate phases. All delays are in picoseconds.

In WISP-0, datapath lengths and the number of transistors on each datapath are different. Consequently the delay varies over a wide range of values. However, the performance of a pipeline is determined by the slowest segment; in this case it is the vertical plane of the RF (delay = 4.778 ps). The operating frequency assuming a 33% duty cycle (reflecting a clock needed for a precharge-evaluate-hold control) is easily shown to be 69 GHz. It is expected that the frequency will be lower in practical designs with longer datapaths and larger bitwidths.

The contact resistance of 10 k Ω is a large contributor to the overall delay for all nanotiles. It is expected that with improvements in manufacturing, this value may be significantly reduced. Table V tabulates the delay for all nanotiles without any contact resistance.

When compared with the values in Table IV, it is clear that even on the larger nanotiles, a large portion of the delay is due to the contact resistance. For example, for the slowest segment (“V-eva” of RF), the contact resistance contributes 25% of the delay. On smaller nanotiles this effect is far more prominent (75% for “H-pre” of the Program Counter tile or PC). The operating frequency for the nanotile without contact resistance is estimated to be 93 GHz.

TABLE IV
DELAY (ps)—ASSUMES CONTACT RESISTANCE

	H-pre	H-eva	V-pre	V-eva
PC	0.227	0.463	0.141	0.536
ROM	0.215	0.796	0.302	3.785
DEC	0.375	1.485	0.934	2.742
RF	0.596	2.135	0.615	4.778
ALU	0.481	1.415	0.667	3.667

TABLE V
DELAY (ps)—NO CONTACT RESISTANCE

	H-pre	H-eva	V-pre	V-eva
PC	0.56	0.186	0.33	0.236
ROM	0.80	0.508	0.96	3.147
DEC	0.155	0.830	0.471	1.674
RF	0.222	1.268	0.260	3.558
ALU	0.153	0.952	0.339	2.593

TABLE VI
DYNAMIC POWER CONSUMPTION (μW)

	3V	3.5V	4V	4.5V
PC	213	290	380	481
ROM	377	509	665	841
DEC	977	1330	1738	2199
RF	2780	3784	4942	6254
ALU	447	609	795	1007

B. Power Estimation

The average dynamic power and the leakage power were estimated for the tiles of WISP-0. Dynamic power calculations were done for a 69 GHz operating frequency for a range of typical operating voltages between 3 V–4.5 V—the voltage is estimated based on the original NW FET papers. The expression used is

$$P_{dyn} = \sum_{pre,eva} (C_{L1} + N * C_{L2}) * V_{dd}^2 * f$$

where f is the operating frequency, C_{L1} is the capacitance on the control NW and C_{L2} is the capacitance on a datapath NW. N is the number of datapath NWs switching simultaneously. In cases where N is variable (e.g., application specific), an average value is chosen assuming a 50% switching probability.

Table VI shows the dynamic power consumption (in μW) for the components of WISP-0 at the 69 GHz frequency. It is seen that the Register File consumes maximum average dynamic power. This is due to a relatively large capacitive load owing to the relatively large size of the tile. The power consumption trends on the whole are orders of magnitude lower than those seen in conventional CMOS technologies.

Leakage power consumption of NASIC tiles was estimated for a supply of 4.5 V. An R_{OFF} resistance of 10 G Ω [7] was used for the calculations. Table VII enumerates the calculated values

TABLE VII
LEAKAGE POWER AT 4.5 V (nW)

PC	10.8
ROM	10.1
DEC	24.3
RF	38.6
ALU	14.0

for WISP-0. The high R_{OFF} implies that the leakage power in these circuits is negligibly small (in the order of nanowatts).

VIII. RELATED WORK

A. Nanoscale Devices for Computing

Some of the most promising underlying nanodevices today targeting digital applications, potentially applicable in 2-D computing fabrics, are based on semiconductor NWs (such as in NASICs) and carbon nanotubes (CNTs). The diameters of NWs and CNTs are in the order of a few nanometers, and their density can be as high as 100 billion switches/cm² [39]. The electrical characteristics of NWs can be more reliably controlled than those of nanotubes [2]; many researchers believe therefore that NW-based devices are easier to assemble into grids and computing systems in general. Current control in NWs or CNTs is realized by using gates formed in various ways, or by forming diode junctions. FET behavior has been achieved using metallic gates [40], [41] and crossing NWs or CNTs [2], [41]. By varying the amount of oxide grown at their intersection, crossing CNTs or NWs can be made such that one NW forms a diode with the other, or one acts as a FET gate to the other, or they do not couple at all [2]. Rapid progress is being made in the development of feasible logic devices. Diode resistor logic was demonstrated. At the same time restoring logic was introduced with NW FET-resistor logic [2]. Avouris from IBM made important progress toward low power logic by developing complementary devices on the same nanotube and demonstrated a CMOS-like nano-inverter [43].

B. Nanoscale Computing Fabrics

Table VIII shows the comparison of four recent fabric styles. These include NASICs, NanoPLA [7], CMOL [9], and a fabric proposed by HP/UCLA [31], [32]. Hewlett-Packard Research has patented a molecular crossbar latch (Kuekes, patent #6,586,965). NASICs use field-effect transistors (FETs) at nanocrossbar junctions to implement logic, rather than diodes or molecular switches such as proposed by NanoPLA and CMOL. With exception of CMOL—that implements part of the logic functions with CMOS cells connected with vertical pins to a nanogrid implementing wired-OR logic—all other fabrics assume the availability of FETs for either logic or signal restoration. NanoPLA uses the FETs in the decoder logic: this is required for addressing grid crosspoints and for reprogramming the fabric around faults. NASIC is also different from the other fabric schemes in the areas of fault tolerance and applications targeted. While most fabrics rely on reconfigurable devices, defect map extraction, and reconfiguration around defects, NASICs use built-in fault-tolerance techniques at various levels to mask faults. Only the NASIC approach might provide a solution to address faults that are caused by non-permanent

defects such as device parameter variation related ones and transient faults.

Most other fabrics are targeted and evaluated for logic applications targeting FPGAs and comparison is often done with CMOS FPGA logic. In contrast, the NASIC project and fabric focuses on processor designs and datapath. All proposals face various manufacturing difficulties at this time. The CMOL fabric has lower requirements on alignment but uses a somewhat challenging 2-level interconnect solution—with different height vertical pins that need to connect the CMOS cells to the nano grid. The NanoPLA approach requires complex defect map extraction and addressing decoder where all crosspoints need to be reached. All fabrics with exception of NASICs assume the availability of reconfigurable devices. All designs use a variant of 2-level logic as underlying logic family.

C. Built-in Nanoscale Fault Tolerance

While there has been little work done on fault-tolerance techniques for nanoscale fabrics, there has been a considerable amount of work done in the field of coding for fault masking in logic in the past. Much of it is based on restoring logic following logic in which faults may occur [25], [26], [46]. These approaches are problematic when working with crossed NW fabrics because the fault rates are expected to be so high that the restoring logic would itself have faults in it. Systems using residue codes either can only be used to detect errors [27], or require complicated iterative processing to correct a limited number of errors [28]. The most representative recent related work (likely developed in parallel with this work) at nanoscale is [18]; it focuses on built-in defect-tolerance at the nanomicro interface. A comprehensive overview of fault tolerance techniques focusing primarily on deep sub-micron CMOS is presented in [36]. In terms of the logic structure proposed, the interwoven logic in [29] is the closest to the one used in our work and the theory regarding critical and non-critical errors in regular logic structures appears applicable.

IX. CONCLUSIONS AND FUTURE WORK

In this paper we demonstrated a variety of built-in fault tolerance techniques on a NASIC-based processor. Our simulation results show that we can tolerate faults from a variety of sources and still achieve considerably higher density than in an equivalent CMOS design at the end of the projected ITRS roadmap. NASIC-based processors show great promise due to the combination of fault-masking, high density, and scalability. The density of NASIC-based designs scales with improvements in nanomanufacturing. Our current focus is on exploring additional techniques for fault tolerance and addressing manufacturability issues. We are working on a second nano processor with a larger bitwidth than WISP-0, incorporating additional NASIC-related architectural innovations and circuit optimizations.

ACKNOWLEDGMENT

The authors would like to acknowledge the fruitful collaboration on NASIC CAD tools with Drs Pottier and Lagadec from the Universite de Bretagne Occidentale, France. Furthermore, the authors have also received valuable input from Drs Krishna, Koren, Jackson, Anderson, Cieselski, and Tuominen from the

TABLE VIII
COMPARISON OF NASIC WITH OTHER NANOSCALE FABRICS

Design	Nano Devices	Targeted Applications	Defect Tolerance	CMOS Roles	Manufacturing Difficulties
NASIC	Single or complementary types of FETs	ASIC-style logic, processors	Built-in defect tolerance at various levels of granularity	Providing Vdd/Gnd and dynamic logic control signals	Alignment during metallization of crosspoints with no FETs for logic customization
NanoPLA	Diodes + FETs as restoration	FPGA logic	Reconfiguration	Vdd/Gnd, extraction of defect maps, reconfiguration	Decoder imprint implementation or stochastic decoder; addressing all crosspoints
CMOL	Molecular switches	FPGA logic, Memory	Reconfiguration	Logic functions, signal restoration and reconfiguration	Nano-micro interface; pins with different heights required; some alignment between nano grid and CMOS cells
HP/UCLA	Diodes + two types of reconfigurable FETs	Logic	Reconfiguration	Providing Vdd/Gnd	Reconfigurable FETs

University of Massachusetts in at Amherst, Dr Kiehl, University of Minnesota, Dr Likharev, Stony Brook University, and Dr Mircea Stan, University of Virginia. The authors would also like to acknowledge the support of Dr Avouris, IBM, who encouraged our early efforts in exploring nanoscale processors.

REFERENCES

- [1] Y. Cui, L. J. Lauhon, M. S. Gudiksen, J. Wang, and C. M. Lieber, "Diameter-controlled synthesis of single crystal silicon nanowires," *Appl. Phys. Lett.*, vol. 15, pp. 2214–2216, 2001.
- [2] Y. Huang, X. F. Duan, Y. Cui, L. J. Lauhon, K. H. Kim, and C. M. Lieber, "Logic gates and computation from assembled nanowire building blocks," *Sci*, vol. 294, no. 5545, pp. 1313–1317, 2001.
- [3] Y. Huang, X. Duan, Q. Wei, and C. M. Lieber, "Directed assembly of one-dimensional nanostructures into functional networks," *Science*, vol. 291, p. 5504, 2001.
- [4] A. J. Bourdillon, G. P. Williams, Y. Vladimirovsky, and C. B. Boothroyd, "22-nm lithography using near field x-rays," in *Proc. SPIE Emerging Lithographic Technol. VII*, Jun. 2003, vol. 5037, pp. 622–633.
- [5] A. J. Bourdillon, G. P. Williams, Y. Vladimirovsky, and C. B. Boothroyd, "Near field X-ray lithography to 15 nm," in *Proc. SPIE Emerging Lithographic Technol. VIII*, 2004, vol. 5374, pp. 546–557.
- [6] Y. Wu, J. Xiang, C. Yang, W. Lu, and C. M. Lieber, "Single-crystal metallic nanowires and metal/semiconductor nanowire heterostructures," *Nature*, vol. 430, pp. 61–65, 2004.
- [7] A. DeHon, "Nanowire-based programmable architectures," *ACM J. Emerging Technol. Comput. Syst.*, vol. 1, no. 2, 2005.
- [8] S. C. Goldstein and M. Budiu, "Nanofabrics: Spatial computing using molecular electronics," in *Proc. 28th Annu. Int. Symp. Comput. Architecture (ISCA'01)*, Jun. 2001, vol. , p. 178.
- [9] K. K. Likharev and D. B. Strukov, "CMOL: Devices, Circuits, and Architectures," in *Introducing Molecular Electronics*. Stony Brook, NY: Stony Brook Univ., 2004.
- [10] C. A. Moritz and T. Wang, "Towards defect-tolerant nanoscale architectures," in *Proc. IEEE Nano2006 Conf.*, Cincinnati, OH, 2006.
- [11] C. A. Moritz, "Exploring nasics and a comparison with CMOL: An architect's perspective," in *Proc. 3rd Adv. Res. Dev. Agency (ARDA) Workshop*, Tampa, FL, 2006.
- [12] C. A. Moritz and T. Wang, "Latching on the wire and pipelining in nanoscale designs," in *Proc. 3rd Non-Silicon Comput. Workshop, NSC-3, Organized Conjunction 31st Int. Symp. Comput. Archit. (ISCA 2004)*, Munich, Germany, 2004.
- [13] T. Wang, Z. Qi, and C. A. Moritz, "Opportunities and challenges in application-tuned circuits and architectures based on nanodevices," in *Proc. 1st ACM Int. Conf. Comput. Frontiers (CF'04)*, Ischia, Italy, 2004.
- [14] T. Wang, M. Bannaser, Y. Guo, and C. A. Moritz, "Wire-streaming processors on 2-D nanowire fabrics," in *Proc. Nanotech 2005, Nano Sci. Technol. Inst.*, Anaheim, Ca, 2005.
- [15] J. P. Patwardhan, V. Johri, C. Dwyer, and A. R. Lebeck, "A defect tolerant self-organizing nanoscale SIMD architecture," in *Proc. ASPLOS'06*, San Jose, CA, 2006.
- [16] P. Narayanan, M. Leuchtenburg, T. Wang, and C. A. Moritz, CMOS enabled single-type FET NASIC Univ. Massachusetts, Amherst, UMASS Tech. Rep..
- [17] D. B. Strukov and K. K. Likharev, "Defect-tolerant architecture for nanoelectronic crossbar memories," *J. Nanosci. Nanotechnology*, vol. 7, pp. 840–853, 2006.
- [18] P. J. Kuekes, W. Robinett, G. Seroussi, and R. S. Williams, "Defect-tolerant interconnect to nanoelectronic circuits: Internally redundant demultiplexers based on error-correcting codes," *Nanotechnology*, vol. 16, pp. 869–882, 2005.
- [19] Y. W. Heo, L. C. Tien, Y. Kwon, D. P. Norton, and S. J. Pearton, "Depletion-mode ZnO nanowire field-effect transistor," *Appl. Phys. Lett.*, vol. 85, pp. 2274–2276, 2004.
- [20] S. Koo, M. D. Edelstein, Q. Li, C. A. Richterand, and E. M. Vogel, "Silicon nanowires as enhancement-mode schottky barrier field-effect transistors," *Nanotechnology*, vol. 16, pp. 1482–1485, 2005.
- [21] Y. Wu, J. Xiang, C. Yang, W. Lu, and C. M. Lieber, "Single-crystal metallic nanowires and metal/semiconductor nanowire heterostructures," *Nature*, vol. 430, pp. 61–65, 2004.
- [22] Y. Huang, X. Duan, Y. Cui, L. J. Lauhon, K.-Y. Kim, and C. M. Lieber, "Logic gates and computation from assembled nanowire building blocks," *Science*, vol. 1313, no. 294, 2001.
- [23] A. A. Bruen and M. A. Forcinito, *Cryptography, Information Theory, and Error-Correction*. Singapore: Wiley-Interscience, 2005.
- [24] Y. Li, F. Qian, J. Xiang, and C. M. Lieber, "Nanowire electronic and optoelectronic devices," *Mater. Today*, vol. 9, pp. 18–27, 2006.
- [25] J. Von Neumann, "Probabilistic logics and the synthesis of reliable organisms from unreliable components," in *Automata Studies (Annals of Math. Studies No. 34)*, C. E. Shannon and J. McCarthy, Eds. Princeton, NJ: Princeton Univ. Press, 1956, pp. 43–98.
- [26] D. B. Armstrong, "A general method of applying error correction to synchronous digital systems," *Bell Syst. Tech. J.*, vol. 40, pp. 477–593, 1961.
- [27] I. L. Sayers and D. J. Kinniment, "Low-cost residue codes and their applications to self-checking VLSI systems," *IEE Proc.*, vol. 132, no. 4, pt. E, Jul. 1985.
- [28] H. Krishna and J. D. Sun, "On theory and fast algorithms for error correction in residue number system product codes," *IEEE Trans. Comput.*, vol. 42, no. 7, pp. 840–853, Jul. 1993.
- [29] W. H. Pierce, *Interconnection Structure For Redundant Logic, Failure-Tolerant Computer Design*. New York: Academic, 1965.
- [30] R. E. Lyons and W. Vanderkulk, "The use of triple modular redundancy to improve computer reliability," *IBM J. Res. Dev.*, vol. 6, no. 2, pp. 200–209, 1962.
- [31] J. R. Heath, P. J. Kuekes, G. S. Snider, and R. S. Williams, "A defect-tolerant computer architecture: Opportunities for nanotechnology," *Science*, vol. 280, no. 5370, pp. 1716–1721, Jun 12, 1998.


- [32] Y. Luo, C. P. Collier, J. O. Jeppesen, K. A. Nielsen, E. DeIono, G. Ho, J. Perkins, H.-R. Tseng, T. Yamamoto, J. F. Stoddart, and J. R. Heath, "Two-dimensional molecular electronics circuits," *Chem. Phys. Chem.*, vol. 3, pp. 519–525, 2002.
- [33] Y. Cui, X. Duan, J. Hu1, and C. M. Lieber, "Doping and electrical transport in silicon nanowires," *J. Phys. Chem.*, vol. 104, pp. 5213–5216, 2000.
- [34] A. B. Greytak, L. J. Lauhon, M. S. Gudixsen, and C. M. Lieber, "Growth and transport properties of complementary germanium nanowire field-effect transistors," *Appl. Phys. Lett.*, vol. 84, no. 21, pp. 4176–4178, May 2004.
- [35] H. T. Ng, J. Han, T. Yamada, P. Nguyen, Y. P. Chen, and M. Meyyappan, "Single crystal nanowire vertical surround-gate field-effect transistor," *Nano Lett.*, vol. 4, no. 7, pp. 1247–1252, 2004.
- [36] T. Lehtonen, J. Plosila, and J. Isoaho, "On fault tolerance techniques towards nanoscale circuits and systems," Turku Centre for Computer Science, Turku, Finland, TUCS Tech. Rep. 708, Aug. 2005.
- [37] B. D. Gates, Q. Xu, J. C. Love, D. B. Wolfe, and G. M. Whitesides, "Unconventional nanofabrication," *Annu. Review. Mater. Res.*, pp. 339–372, 2004.
- [38] C. Constantinescu, "Trends and challenges in VLSI circuit reliability," *IEEE Micro*, vol. 23, no. 4, pp. 930–931, Apr. 2004.
- [39] T. Rueckes, K. Kim, E. Joselevich, G. Y. Tseng, C.-L. Cheung, and C. M. Lieber, "Carbon nanotube-based nonvolatile random access memory for molecular computing," *Science*, vol. 289, pp. 94–97, 2000.
- [40] R. Martel, V. Derycke, J. Appenzeller, S. Wind, and P. Avouris, "Carbon nanotube field-effect transistors and logic circuits," in *DAC 2002*, New Orleans, 2002.
- [41] A. Bachtold, P. Hadley, T. Nakanishi, and C. Dekker, "Logic circuits with carbon nanotube transistors," *Science*, vol. 294, no. 5545, pp. 1317–1320, Nov. 9, 2001.
- [42] A. B. Greytak, L. J. Lauhon, M. S. Gudixsen, and C. M. Lieber, "Growth and transport properties of complementary germanium nanowire field effect transistors," *Appl. Phys. Lett.*, vol. 84, p. 4176, 2004.
- [43] P. Avouris, R. Martel, V. Derycke, and J. Appenzeller, "Carbon nanotube transistors and logic circuits," *Phys. B Condensed Matter*, vol. 323, no. 1–4, pp. 6–14, Oct. 2002.
- [44] R. C. Rose and D. K. Ray-Chaudhuri, "On a class of error-correcting binary group codes," *Inform. Contr.*, vol. 3, pp. 68–79, Mar. 1960.
- [45] A. Hocquengham, "Codes correcteurs d'erreurs," *Chiffre*, vol. 2, pp. 147–156, Sep. 1959.
- [46] T. R. N. Rao, *Error Coding For Arithmetic Processor*. New York: Academic, 1974.
- [47] B. W. Johnson, *Design and Analysis of Fault-Tolerant Digital Systems*. Reading, MA: Addison-Wesley, 1989.



Csaba Andras Moritz received the Ph.D. degree in computer systems from the Royal Institute of Technology, Stockholm, Sweden, in 1985.

He is a tenured Associate Professor in the Department of Electrical and Computer Engineering, University of Massachusetts, Amherst. From 1997 to 2000, he was a Research Scientist at MIT, Laboratory for Computer Science, Cambridge. He has consulted for several technology companies in Scandinavia and held industrial positions ranging from CEO, to CTO, and to founder. His most recent startup company,

BlueRISC Inc, develops security microprocessors and hardware-assisted security solutions. His research interests include computer architecture, compilers, low power design, security, and nanoscale systems.



Teng Wang received the B.S. degree in electronic engineering and information science from the University of Science and Technology of China (USTC), Hefei, China, and the M.S. degree from Chinese Academy of Sciences in 1999 and 2002, respectively. He is currently working toward the Ph.D. degree in electrical and computer engineering at University of Massachusetts, Amherst.

His research interests include nanoscale systems, computer architecture, and VLSI design.



Prithiv Narayanan received the B.E. (Hons) degree in electrical and electronics engineering and M.Sc. (Hons) degree in chemistry from the Birla Institute of Technology and Science, Pilani, India in 2005. He is currently working toward the Master's degree in electrical and computer engineering at the University of Massachusetts, Amherst.

He was previously employed as a Research and Development Engineer at IBM, where he worked on process variation and statistical timing analysis. His interests include nano-architectures and neural net-

works.



Michael Leuchtenburg is currently working toward the M.S. degree in computer architecture at the University of Massachusetts.

His research interests are nanoscale computer architecture, quantum computing, and parallel algorithms.



Yao Guo received the Ph.D. degree in computer engineering from University of Massachusetts at Amherst in 2007.

He is an Assistant Professor in the Operating Systems Lab, School of Electronics Engineering and Computer Science, Peking University, Beijing, China. His research interests include operating systems, low-power design, compilers, embedded systems and software engineering.



Catherine Dezan received the Master's degree in electrical engineering from the Ecole Nationale Supérieure des Télécommunications de Bretagne, Brest, France, in 1989, and the Ph.D. degree in computer science from the Université de Rennes I, Rennes, France, in 1993.

Since 1994, she has been an Associate Professor of Computer Science at the Université de Bretagne Occidentale, Brest, France. Her main interests are on design methodologies and applications for hardware implementations with a special focus on parallel, re-

configurable and recently nanoscale architectures.



Mahmoud Bannaser received the B.S. degree in computer engineering from Kuwait University, Kuwait, in 1999 and the M.S. degree in computer engineering from Brown University, New Providence, RI, in 2002. He is currently working toward the Ph.D. degree in computer engineering at University of Massachusetts, Amherst.

His research interests include computer architecture, and low-power circuit design.

Mr. Bannaser received the certificate of academic excellence for the years 1995–1998. In 2000, he received a scholarship from Kuwait University to pursue M.S. and Ph.D. degrees in computer engineering.

received a scholarship from Kuwait University to pursue M.S. and Ph.D. degrees in computer engineering.