

NanoFabrics: Extending Moore's Law Beyond the CMOS Era

Seth Copen Goldstein, School of Computer Science, Carnegie Mellon University

The remarkable success of modern computing is based in large part on advances in CMOS-based integrated circuits. Although we have come to expect, and plan for, the exponential increase in processing power in our everyday lives, today Moore's Law faces imminent challenges both from the physics of deep-submicron CMOS devices and economics of manufacturing. In particular, the large costs of fabrication plants and increased cost of chip masks will prohibit all but the highest volume chips from being manufactured economically.

The increasing importance of highly specialized embedded systems and the changing workloads on the desktop further threaten the continued success of general-purpose processors. Modern workloads focus more on streaming data with highly predictable, parallel, fine-grained operations. These workloads obviate many of the advances in computer architecture leading to inefficient usages of silicon and power. Reconfigurable computing offers one solution to this change in workloads. A reconfigurable device is one in which the hardware can be customized repeatedly at runtime to suit a particular application. A program for a reconfigurable computer is essentially a description of a circuit. The price of creating a customized circuit for each application is that the underlying hardware must support post-fabrication configuration. This introduces an overhead in space (SRAM cells at every programmable device and the wires to address these configuration bits) and time (extra delay to go through reconfigurable switches). This overhead reduces the application domains in which reconfigurable computing can offer an advantage.

A promising solution to both the manufacturing and architectural problems presented above lies in exploiting a new device technology: chemically assembled electronic nanotechnology (CAEN), a form of electronic nanotechnology (EN), which uses self-alignment to construct electronic circuits out of nanometer-scale devices. The fundamental strategy of CAEN is to substitute compile time (which is inexpensive) for manufacturing precision (which is expensive). CAEN-based devices are inherently small, low power, and inexpensive to manufacture. Furthermore, a CAEN-based reconfigurable switch has recently been demonstrated. The key benefit of this switch is that it holds its own configuration state, eliminating the need for both configuration wires and SRAM cells found in CMOS-based configurable devices. Thus CAEN based reconfigurable fabrics will be considerably more dense and efficient than CMOS ones.

There are, however, many seemingly inherent drawbacks to CAEN-based devices. The two most important are: (1) It will not be possible to fabricate anything but the most symmetric and homogeneous circuits and (2) The manufacturing process is inherently prone to defects. A further hurdle is that currently there appears to be no way to fabricate and interconnect devices with gain, i.e., there are no CAEN transistors.

To address these drawbacks and provide a reconfigurable fabric with more than 10^{12} transistor equivalents per centimeter we propose a nanoFabric, an architecture driven by the limited set of CAEN fabrication primitives. It uses post-fabrication reconfiguration to implement circuits. The nanoFabric consists of a two-dimensional mesh of nanoBlocks. Each nanoBlock is itself a two-dimensional mesh of wires with configurable diodes at the intersections. Each intersection can be configured to be either a diode or an open circuit. The wires are either connected to other nanoBlocks through configurable switches or through resistors to power or ground. This allows a nanoBlock to implement a combinational circuit using diode-resistor logic. The wires that are interconnected to other nanoBlocks go through a novel latch composed of negative differential resistors. This latch provides both clocked logic and signal restoration. Preliminary Spice simulations show that even with today's early devices a nanoFabric will operate in the hundreds of megahertz at less than tens of nanoWatts per binary operation.

The highly configurable nature of the nanoBlock allows us to overcome the inherent limitations of CAEN devices. Defect tolerance is provided in two stages. First, a post-fabrication step detects the defects by downloading self-test circuits onto the device that create a map of the defects. Next, the defect map is used to create a chip specific configuration that avoids the defects. The lack of fabrication complexity is overcome by post-fabrication configuration of any circuit on the homogeneous substrate.

To support compilation that will scale to devices with trillions of configurable components we propose a four-level hierarchy of architectural abstractions. The highest level of abstraction is a split-phase abstract machine. Programs are decomposed at split-phase operations, which increases latency tolerance (thereby easing the place-and-route problem) and reduces the compilation problem to many independent smaller problems. The next three levels in the hierarchy introduce global place-and-route, local place-and-route and circuits, and then finally, defect tolerance.

Reconfigurable computing with a CAEN-based nanoFabric has the potential to not only continue, but increase, the pace of improvement known as Moore's Law. Manufacturing costs will drop due to the use of guided chemical self-assembly and defects can be avoided through reconfiguration. Mask costs will be eliminated as reconfiguration will allow any circuit to be implemented post-fabrication. Additionally, the elimination of the overhead to support configuration will widen the applications that can benefit from the reconfigurable computing.