

# A Comparison of Circuits for On-Chip Programmable Crossbar Switches

Enrique Coen-Alfaro and Gregory W. Donohoe  
Institute of Advanced Microelectronics  
at the  
University of New Mexico

ecoen@mrc.unm.edu

**ABSTRACT.** Many modern integrated circuits employ programmable on-chip interconnects to enable the chip to be configured in the field. The interconnect structure can have a significant effect on area, speed, and power consumption. This is especially true as process geometries shrink, and interconnects, which don't scale as aggressively as transistors, become proportionately more significant. The performance advantages of multilevel, hierarchical interconnect structures built around multi-input, multi-output switching nodes (crossbars) have been demonstrated in several studies. This paper presents a study of two approaches to implementing on-chip crossbar switches: a fully-connected square switching matrix, and a multilevel switching network based on two-port pass gate elements. We consider transistor count and area, and comment on switching speed, and power consumption, with particular emphasis on sub-threshold leakage current.

## I. Introduction

The recent growth of programmable logic and reconfigurable computing has created interest in on-chip programmable interconnects. As CMOS processes shrink and the number of gates per chip grows, interconnects play an increasingly important role in area, performance, and power consumption. Interconnects take up to 90% of the area of a field-programmable gate array [DeHon96]. In order to improve performance and reduce power consumption, programmable devices are built on a hierarchical architecture, with multiple partitions sharing central crossbars [Lai97, Zhang99, Zhang99a].

A basic crossbar, or crosspoint, switch has a set of inputs, a corresponding set of outputs, and a set of addresses mapping inputs to outputs. Crossbars can be classified according to following:

1. Blocking or non-blocking
2. Bi-directional or unidirectional
3. One-to-one or multicasting
4. Static or dynamic configuration
5. Fully connected or partially connected

Recent changes in CMOS processes have changed the design parameters for crossbar networks. As process geometries shrink, supply voltages are also reduced, from 5 V a few years ago to 3.3V, then 2.5V, and less. Recent work in ultra-low-power CMOS has produced devices operating with a supply voltage of 0.5V [BURR91]. As the supply voltage is reduced, the MOS transistor thresholds are also reduced. This causes an increase in sub-threshold leakage current, which affects power consumption and switching performance. The effects of leakage are most pronounced in circuits with large numbers of parallel transistors, such as memories and crossbars [Hass01].

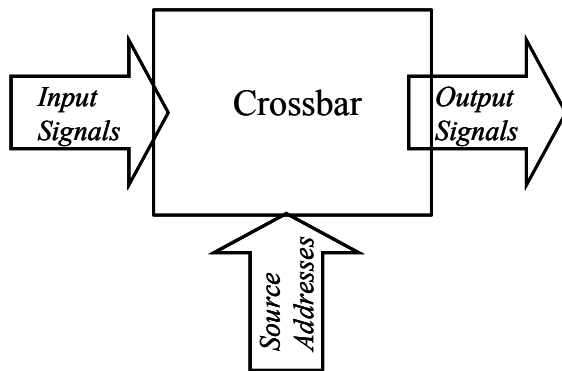


Figure 1: Basic crossbar switch.

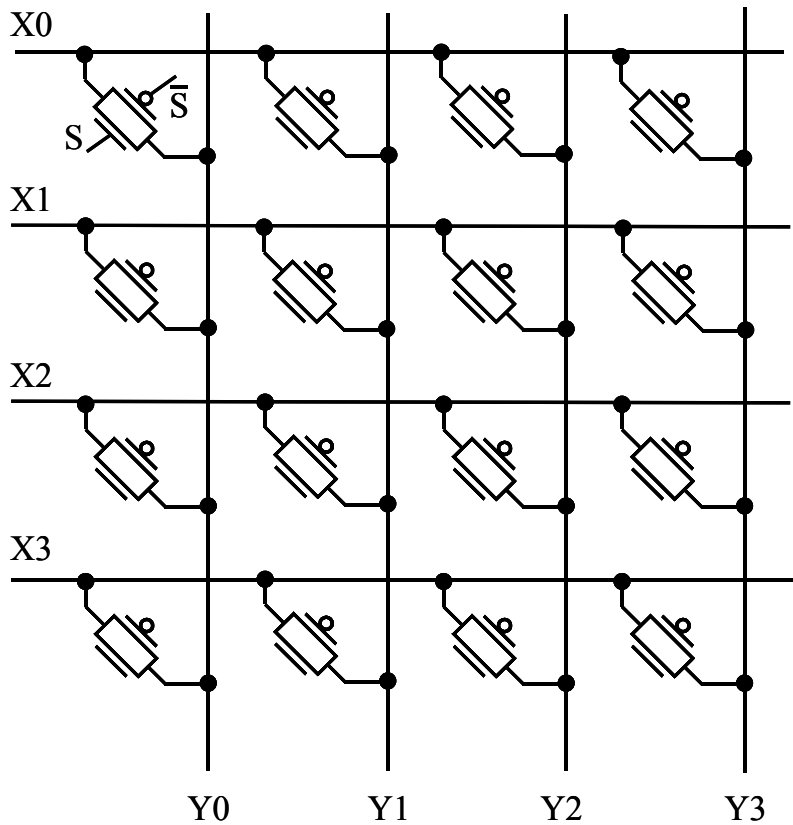
This paper explores several alternative crossbar designs, and compares their performance with regard to transistor count and area, with qualitative comments on throughput, power consumption and routability.

The study focuses on statically configured, bi-directional crossbars with one-to-one input/output mapping, as would be used in reconfigurable logic devices.

## II. Crossbar Implementations

The simplest crossbar structure consists of a crosspoint matrix, a two-dimensional grid of crossed wires. Inputs run horizontally, for example, and outputs run vertically. Each wire crossing is linked by a transmission. To complete the connection between an input and output pair, the transmission gate is turned on. Figure 2 illustrates the concept for a  $4 \times 4$  crosspoint matrix (the control signals  $S$  is shown for only one transmission gate).


The crosspoint matrix is simple and regular, and the signal only passes through one switch. The matrix poses several concerns for low-voltage CMOS implementations, however. The switch requires  $O(N^2)$  switching elements, where  $N$  is the number of inputs. Each input and output line is connected to a large number of parallel transistors, which in turn are connected to the pull-up and pull-down transistors of the drivers. This presents a large load capacitance and multiple low-resistance leakage paths to ground and  $V_{DD}$ . This consumes static power and reduces signal throughput.



**Figure 2: 4x4 crosspoint matrix.**

The  $N \times N$  crosspoint matrix can be implemented with static logic gates instead of transmission gates, which helps alleviate the parallel leakage issue, but requires a very large number of gates.

An alternative design approach is based on multistage interconnection networks, which date back to the early days of electronic telephone switches. Multistage interconnects based on electromechanical relays were proposed and studied by Clos [Clos52] and Benes [Benes65].



**Figure 3: Binary switching element.**

Multistage networks can be synthesized from the binary switching element, shown in Figure 3. The element has two one-bit data inputs,  $X0$  and  $X1$ , and two outputs,  $Y0$  and  $Y1$ , as well as a selection switch,  $S$ . When  $S=0$ ,  $Y0$  is connected to  $X0$  and  $Y1$  is connected to  $X1$ , as shown in Fig. 3 (a). When  $S=1$ , the connects are swapped, as seen in Fig. 3 (b).

Figure 4 shows a  $4 \times 4$  switch implemented in two stages from the binary switches. This is an example of a shuffle-exchange network. In general, we can map  $N$  inputs to  $N$  outputs in a non-blocking fashion with  $O(N \log N)$  switches, as opposed to




Figure 4:  $4 \times 4$  fully-connected, non-blocking multistage network.

$O(N^2)$  switching points for the fully-connected crossbar. The number of intermediate stages grows with the number of inputs and outputs. Many different, but functionally equivalent shuffle-exchange switches have been proposed. All can implement all of the  $N!$  permutations of the input-output mappings. However, some formulations may have advantages in terms of local on-chip interconnect routing. Figure 8 depicts an  $8 \times 8$  switch based on the Benes shuffle-exchange design, requiring five successive stages [Duato97].

### III. Circuit Design and Layout

Both crosspoint matrix and multistage implementations were designed for two CMOS processes: a commercial TSMC 0.25-micron digital process operating with a supply voltage of 2.5V and an experimental ultra-low-power process based on an AMI 0.35-micron technology, operating at 0.5V. The designs were laid out, and SPICE circuits extracted.

The binary switch can be implemented with static logic or with transmission gates. A simple implementation requires four transmission gates, or 10 transistors (8 for the transmission gates, and two to invert the select signal.) Using restoring logic, 26 transistors are required per binary switch. (This does not include the source address decoders.)




Figure 5: 8x8 fully connected, non-blocking, multistage Benes network.

Using transmission gates, the binary switch becomes a two-port network. The large number of leakage paths to  $V_{DD}$  and ground which occurred with the crosspoint matrix no longer exists.

## IV. Implementation Results

Table 1: Transistor Count

Transistor Count		
	Crosspoint Matrix	Multistage Network
4 x 4	64	60
8 x 8	256	200
16 x 16	1024	320

Table 1 lists the transistor count for the two different implementations, and Table 2 shows the area for the two different processes. As we can see, the multilevel network leads to compact designs. As expected, the crosspoint matrix implementation grows approximately quadratically with the size, while the multistage network grows much more slowly. The ULP process has extra contacts for back bias, and this affects the area: in the ULP processes, the multistage network is always smaller.

**Table 2: Area**

Area ( $\mu^2$ )				
	TSMC25 0.25 $\mu$		ULP 0.35 $\mu$	
	Crosspoint Matrix		Multistage Network	
	Crosspoint Matrix	Multistage Network	Crosspoint Matrix	Multistage Network
4 x 4	356	440	1006	896
8 x 8	1491	2002	3981	3460
16 x 16	4811	3602	16111	6316

Designing with series pass gates presents potential problems with charge-sharing. However, the number of series pass gates the signal must traverse is the same for all switching paths, and is known at design time, so we can insert buffers as necessary to ensure acceptable operation. In fact, the number of series gates is small enough so that such buffers should not be necessary except in very large crossbars.

## V. Conclusion

Non-blocking, multistage crosspoint switches offer some important benefits for programmable on-chip interconnect blocks. Brute-force, fully-connected crosspoint matrix implementations are reasonable for small switches, but since the number of transistors required increases quadratically with the number of inputs, the area, power consumption and throughput of large crosspoint switches can be significant. With low-voltage, high-leakage CMOS processes, the large number of parallel, low-resistance leakage paths of simple crosspoint matrix presents electrical challenges. However, multistage networks offer a regular decomposition of the design, permitting large crosspoint switches to be constructed from simple two-port switching elements, using  $O(N \log N)$  transistors for an  $N \times N$  switch. The parasitic leakage paths are isolate to simple blocks, and so become less significant.

## References

- [Benes65] Benes, V.E., *Mathematical Theory of Connecting Networks and Telephone Traffic*, Academic Press, 1965.
- [BURR91] Burr, James. B., and Allen M. Peerson, "Ultra Low Power CMOS Technology", NASA Symposium on VLSI Design, pp. 4.2.1-4.2.13, 1991.

[Cloc52] Clos, Charles, "A Study of Non-Blocking Switching Networks", Bell System Tech Journal, Vol. 32, pp. 406-424, March 1953.

[DeHon96] André DeHon, "Reconfigurable Architectures for General-Purpose Computing", *Massachusetts Institute of Technology Artificial Intelligence Laboratory Report No. 1586*, sponsored under contract F30602-94-C-0252, October 1996.

[Duato97] Duato, José, S. Yalamanchili, and L. Ni, *Interconnection Networks: an Engineering Approach*, IEEE Computer Society Press, 1997.

[Hass01] Hass, K. Joseph, J. Venbrux and P. Bhatia, "Logic Design Considerations for 0.5-Volt CMOS", 2001 Conference on Advanced Research in VLSI, March 14-16, 2001.

[Lau97] Lai, Yen-Tai and Ping-Tsung Wang, "Hierarchical Interconnection Structures for Field Programmable Gate Arrays", *IEEE Trans. VLSI Systems*, Vol. 5, No. 2, June 1977.

[Zhang99] Ahang, Hui, Marlene Wan, Varghese George, Jan Rabaey, "Interconnect Architecture Exploration for Low-Energy Reconfigurable Single-Chip DSPs", *Proc. IEEE Computer Society Workshop on VLSI '99*, 8-9 April 1999, pp. 2-8.

[Zhang99a] Ahang, Yan and Mary Jane Irwin, "Power and Performance Comparison of Crossbars and Busses as On-Chip Interconnect Structures", *Proc. 33<sup>rd</sup> Asilomar Conf. on Signals, Systems and Computers*, Oct. 24-27, 1999, pp. 378-383.