

An Interconnection Network For A Novel Reconfigurable Circuit Board

Richard Norman, Etienne Lepercq¹, Yves Blaquière², Olivier Valorge¹, Yan Basile-Bellavance¹,
Richard Prytula³, Yvon Savaria¹

GR2M (1), École Polytechnique de Montréal,
Département d'Informatique (2), Université du Québec à Montréal
Gestion TechnoCap Inc.(3), Montréal

Abstract—This paper presents a programmable interconnection network for a novel multi-reticle integrated circuit providing a reconfigurable circuit board for rapid system prototyping. This multi-dimensional mesh grid network, called WaferNet™, can actively interconnect any pair of pins of integrated circuits deposited on the configurable system board. Two crossbar architectures are implemented and compared, one based on crosspoints and one based on standard cell multiplexers. Implementation results show the feasibility of this proposed cell-based array network that could interconnect a very large number of nodes, spread over an area that could fill a whole wafer, using a typical 6-metal 0.18 μm CMOS technology.

I. INTRODUCTION

To be competitive in today's market, electronic systems must incorporate sets of components of growing complexity while meeting increasingly tight constraints on size, power efficiency and time-to-market. Some of the main difficulties do not come from within the components themselves, where simulation and prototyping platforms are very mature, but in their interconnections. This requires new techniques for rapid development and prototyping of compact systems that densely interconnect a wide variety of component types at ever-increasing bandwidths.

Field Programmable Interconnect Chips (FPICs) proposed by Mohsen [1] are capable of reconfigurably interconnecting electronic components, but a printed circuit board must be built to match the contact patterns of those components. Furthermore, the limited capacity of FPICs involves multiple stages of these chips for high-end circuit boards. A high degree of connectivity can be achieved, but full connectivity would require several times more area in FPICs than in the chips of the target system, greatly limiting the density of the system chips and increasing the cost of each prototype.

The International Technology Roadmap for Electronic Interconnections [2] foresees large increases in PCB interconnect density in the coming years, where by 2009, PCB line/space width will be 100/85 μm , while I/O pads diameters will be of 200 μm on boards with 8 layers. For a high-end FPGA using a 1 mm pitch with 0.4 mm space between pins [3], the wire density is evaluated at 45 cm/cm^2 [4]. New

approaches are needed to ease the development of systems with these upcoming technologies.

A novel high-density programmable system board, WaferBoard™ is proposed to reduce the cost and time of developing complex electronic systems. It is based on a large Integrated Circuit called WaferIC™, which can be larger than a reticle and can ultimately fill a whole wafer [5, 6]. The WaferIC™ can have several millions of tiny (e.g. 50 $\mu\text{m} \times 50 \mu\text{m}$) conducting pads (NanoPads) on its surface, with a compliant contact layer (e.g. Z-axis film, stamped metal spring contacts, anisotropic conductive film (ACF), wire embedded in elastomer, or sea of leads) [7] on top. This surface can thus have over three orders of magnitude more contacts than the most complex integrated circuit (IC) packages. Integrated circuit packages with a variety of contact spacings and patterns can be placed anywhere on WaferIC™ surface, and their balls will each achieve electrical contact with one or more NanoPads. A small array of NanoPads is grouped into each Unit-Cell, these cells are tiled within a reticle, and the WaferIC™ is built from repetition of such sea-of-cells reticles, with inter-reticle stitching for connections between reticles [5, 6]. This regular cell-based architecture allows creating a repeatable pattern that dramatically simplifies the design. It complies with several design rules and techniques that were proposed for building wafer-scale integrated circuits, such as defect tolerance and structural regularity [e.g. 8, 9].

This paper presents the defect-tolerant interconnection network, called WaferNet™ that makes this system prototyping technology feasible. This network is used to programmably interconnect any NanoPad to any other according to a netlist provided by the user. The following section describes the network architecture. It shows that its logic area is dominated by crossbars, which also defines the maximum number of configurable interconnects supported by the interconnection network. The third section describes and compares two different crossbar architectures, based on their logic area and the number of supported links. Performance results of WaferNet™ laid out using a 6-metal 0.18 μm CMOS technology are presented in the last section.

II. PROPOSED DEFECT TOLERANT INTERCONNECT NETWORK

The target system prototyping technology should support any standard type of ICs, especially current high-end FPGAs, processors and memories, with high pin-count packages of several balls per square millimeter [10]. The configurable prototyping board is covered with an array of NanoPads, providing contacts with IC balls laid over it without needing precise alignment.

WaferNet™ takes advantage of the large interconnect density of multiple metal layers in state-of-the-art CMOS technologies, producing a defect-tolerant interconnection network capable of implementing any interconnection scheme specified by a user's netlist. These links must be compatible with IC-to-IC interconnections such as point-to-point, point-to-multipoint and busses. Furthermore, routing capability of the proposed network must be sufficient to route any reasonable sets of those signals without any routing conflict, and must support the routing density of future high-end printed circuit boards.

WaferNet™ is a scalable multi-dimensional mesh network. A two dimensional mesh is efficient for establishing connections between near neighbors but inefficient for longer interconnections, with signals having to pass through hundreds of elements to cross a whole wafer. The grid mesh interconnection structure therefore includes longer interconnects, where each node is linked to K other nodes in each physical direction (N-S-E-W). The K links in a given direction have lengths growing according to a geometric series. As shown with the example in Figure 1, when $K=3$, each node is connected to the 2nd and the 4th neighbor

nodes (only outgoing links on South and East directions are shown for the 0;0 cell).

This network architecture greatly improves the efficiency over a simple mesh when interconnections between distant neighbors are required. It also provides defect tolerance on interconnections as well. Defect tolerance improves as K becomes larger as each crossbar supports more links than are normally needed. Defective crossbars and links are treated as 'already used' resources by the routing software. Simply changing the order of links will route around defects on a longer paths, while one-link paths have shortest-length links added. Defects at the network endpoints are handled by having an array of NanoPads dense enough that each component contact intersects a plurality of NanoPads, and by having cells able to control NanoPads of defective neighbors. Further analysis of the robustness to defects of this network is beyond the scope of the present paper.

Due to Wafer-Scale integration constraints, the whole system, network included, is implemented with a regular architecture based on a *Unit-Cell* elementary tile. For simplicity, the number of cells in a reticle is a power-of-two in each dimension. Each Unit-Cell has one $M \times N$ crossbar to route its $4 \times K$ (4 directions) incoming signals to its $4 \times K$ outgoing signals. Each network node corresponds to a Unit-Cell that can manage up to B IC balls. Each input IC ball supported by a cell implies one more crossbar input. Each bidirectional IC ball implies two more crossbar output where one signal is required to control the signal direction. The size of the crossbar therefore depends on B and the number of IC balls supported by a cell and its neighbor cell (for defect tolerance), where $N \geq 4K+B$ and $M \geq 4K+2B$. Increasing M or N makes the network more robust to faults or defects. Increasing K , i.e. adding longer links, improves routability at the expense of complexity, as each extra dimension of longer links requires as much line length as all links of shorter lengths added together, thus doubling the total bus width required for links (assuming a fixed pitch).

This network is easily mapped to a regular sea-of-cells array. This tile-based cell array implies $2^{K-1} + \dots + 2^1 + 2^0 = 2^K - K - 1$ wires pass through each cell in one direction, i.e. $4(2^K - K - 1)$ wires for both directions and dimensions (vertical and horizontal). Note that beyond some threshold length, determined by performance objectives and technology used, each long line could require repeaters to reduce signal propagation time. The number of metal layers, the wire pitch as well as the physical dimension of the cells therefore limits the value of K . A detailed analysis of these parameters is beyond the scope of this paper, as the analysis depends on other factors such as power consumption, logic complexity, and power distribution.

Our on-going research has shown that the crossbar tends to be Unit Cell's most expensive block in terms of logic area. The following section analyzes the area and efficiency of different crossbar implementation techniques, which is a focal point of this paper.

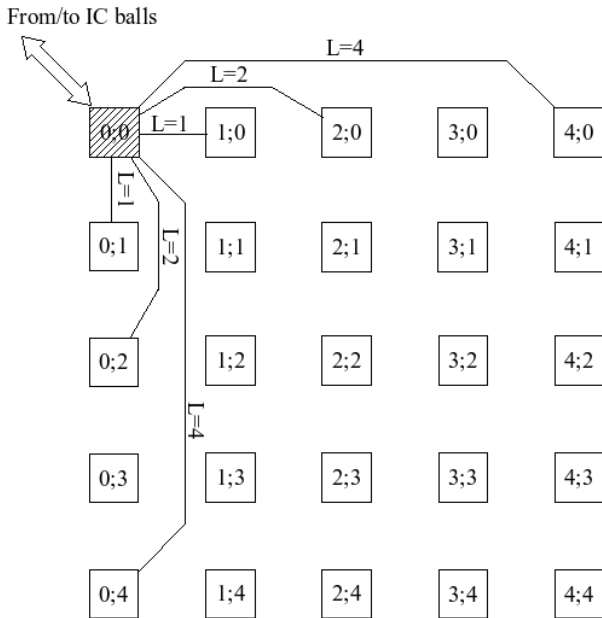


Figure 1. WaferNet™ with $K=3$, with interconnects shown for Unit-Cell 0,0 on East and South directions only (L represents the link length in terms of Unit-Cells).

III. CROSSBAR IMPLEMENTATIONS

A crossbar is an entity able to route M inputs to N outputs [11]. The system cell size is dominated by the crossbar required by the network architecture. Several existing techniques [12] considered in other applications are compared to identify one leading to acceptable areas for the reconfigurable interconnect network explored in this paper. To implement the equivalent of a given printed circuit board, the network is configured according to a user netlist. The netlist is compared to detected chip positions, and translated to a configuration bit-stream by external software, and the configuration is downloaded through scan chain registers, similar to that found in found in FPGAs [13].

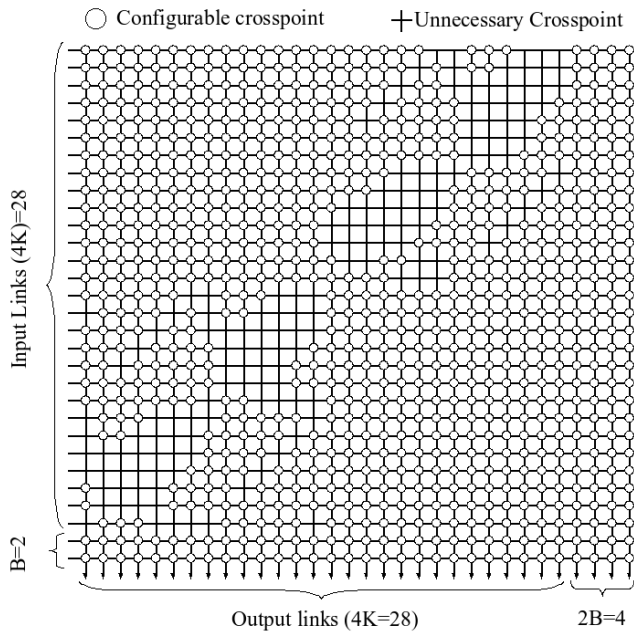


Figure 2. An $M \times N$ crosspoint-based crossbar ($K=7$).

A. Crosspoint-Based Crossbar

The first approach considered to implement crossbars is to use crosspoints. A crosspoint-based crossbar (Fig. 2) implies a memory element per crosspoint that lets a row signal be propagated or not to some column. Note that some crosspoints are labeled as unnecessary. Indeed in our application, as bidirectional communications are implemented in the network, crosspoints returning a signal where it comes from are not useful. However, in some regular layout techniques, not implementing such crosspoints could lead to more complex layout and no area savings. This depends on specific design choices, and exploring the impact of these choices is one object of this research. In the rest of this paper, we will call *partial* the implementations where the unnecessary connections are pruned from the design.

Pass transistor switches (Fig. 3 (b)) are commonly used in FPGAs to implement such crosspoints. This advantageously supports bi-directional links. Implementing crosspoints with pass transistors demands careful design, and possibly using a multi-threshold technology to alleviate threshold voltage losses. Using transmission gate switches alleviates these losses at the cost of more silicon area and parasitic capacitances. The large resistivity of pass transistors or transmission gates may

null the advantages of this switch-based crosspoint technique if repeaters are needed to regenerate the signals.

The power-up and configuration sequence of this crosspoint-based crossbar must be carefully managed to avoid any crosspoint short on column lines. These shorts could generate current surges that could significantly reduce WaferBoard™ life time or consume excessive power.

B. Multiplexer-Based Crossbar

A second approach to implement the crossbars is to use logic multiplexers, commonly available in ASIC libraries. A full $M \times N$ crossbar requires M N -input multiplexers (one per column). As $\log_2(M)$ memory elements are required for each column multiplexer, a total of $N \log_2(M)$ configuration memory elements are needed. This approach leads to a significant reduction in the number of memory elements, especially for large K . Using multiplexers also prevents the short-circuit issue encountered with crosspoint-based crossbars.

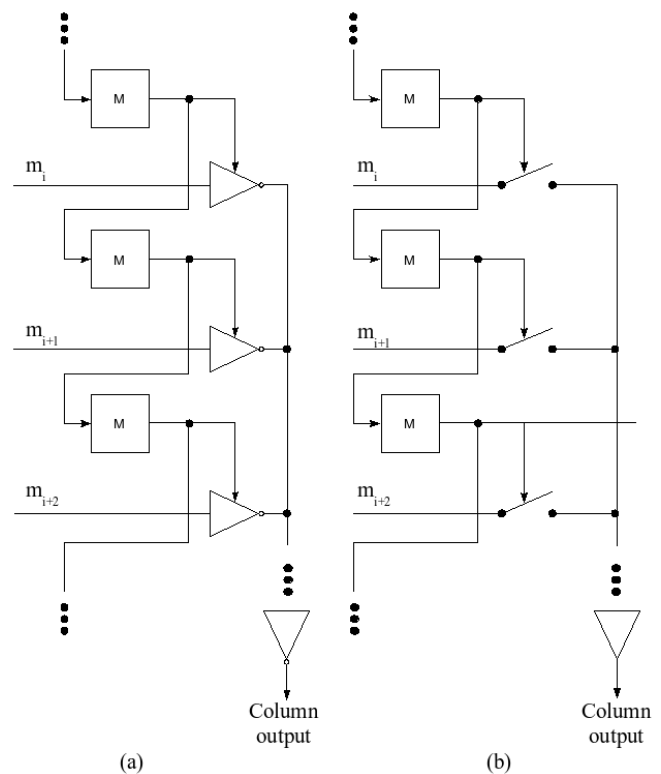


Figure 3. Examples of crosspoint-based implementations: (a) Tri-state based crosspoint, (b) switch-based (pass transistor or transmission gate) crosspoint.

IV. NETWORK IMPLEMENTATIONS AREA RESULTS

The complete network was implemented and laid out using a popular 6-metal layer 0.18 μm CMOS technology, with a commercial static CMOS standard cell library. The pass-transistor or transmission gate crosspoint implementations are not compatible with the standard cell design approach and are therefore not included in the remainder of this analysis as they imply custom layout. For reference purposes, with that library, the area of a 32-to-1 synthesized multiplexer (one crossbar

column) is roughly half the size of 30 minimum size tri-state inverters, and the multiplexer takes six times less configuration memory per column (5 bits rather than 30 bits).

Results are presented for a specific network comprising 30×32 crossbars for which $K=7$. Each Unit-Cell in this network is therefore directly connected to its 1st, 2nd and up to the 64th nearest neighbor cells for a total of 28 input and 28 output links connected directly to its crossbar. More than 480 long wires pass through it, i.e. $2^K - K - 1 = 120$ wires in each direction, horizontally and vertically.

Table 1 summarizes the logic area extracted from RTL synthesis for four implementations. The reported full crossbar connects all input-output pairs, whereas the reported partial crossbar implements the sparse crossbar shown in figure 2. The tri-state based crossbar was implemented with minimum size tri-state inverters followed by one inverter per crossbar column. Results obtained under previous assumptions show that the partial multiplexer-based crossbar occupies $43538 \mu\text{m}^2$, i.e. slightly more than 15% of the total area of a targeted $540 \mu\text{m} \times 540 \mu\text{m}$ cell, while the full tri-state based crossbar, including the configuration memory, is three times larger.

Memory modules for the tri-state based crossbar occupy nearly 50% of the total area, which further explains the gains with the second implementation, the mux-based crossbar. As expected, the partial mux-based area is slightly smaller than the full one, where the number of multiplexer inputs varies from 30 to 22 inputs (according to the crossbar column).

TABLE I. SYNTHESIS RESULTS OF CROSSBAR IMPLEMENTATIONS

Crossbar Type	Total Area (μm^2)	Memory Area (μm^2)	Memory proportion (%)
Full Tri-state based	125400	54240	43.3
Partial Tri-state based	91521	41764	45.6
Full mux based	45050	9040	20.2
Partial mux based	43538	9040	20.8

V. CONCLUSION

This paper has presented a programmable interconnection network for an innovative reconfigurable rapid-prototyping system technology, implemented in a large integrated circuit that can scale up to a whole wafer. This network can interconnect specified sets of component contacts according to a netlist provided by a user, without conflicting with paths between other such sets of connections. This cell-based mesh network takes advantage of the multi-metal layers of current technologies. It adds long interconnects to other nodes in each physical direction with length growing exponentially for efficient propagation to distant neighbors. This regular sea-of-cells architecture is defect tolerant, a required property for wafer-scale circuits. This network has been successfully implemented as a register transfer level model. It has been thoroughly verified and laid out as a tile-based regular cell array scalable to arbitrarily large arrays. Each cell has

unidirectional links of length 1, 2, 4, 8, 16, 32 and 64 in each direction and both dimensions. Our results show that each network node tends to be dominated by a large crossbar connecting all incoming links to all outgoing links. Two different crossbar implementations have been studied and compared using a standard $0.18 \mu\text{m}$ CMOS process and an automatic workflow leveraging a commercial static CMOS standard cell library. Our results show that the multiplexer-based crossbar is the most efficient in terms of area cost, mainly due to the fact that the number of configuration memory element per output grows logarithmically with the number of crossbar inputs (or $O(N \log_2 M)$ in total) compared to a linear growth ($O(N \times M)$ in total) with crosspoint-based crossbars.

Further research on this innovative reconfigurable circuit board for fast system prototyping and its wafer-scale interconnection network is required. An analysis of the defect tolerance capabilities of this network is ongoing. We also investigate the optimal scheme for signal propagation under power dissipation and area constraints.

ACKNOWLEDGMENTS

The authors would like to thank NSERC and Gestion TechnoCap Inc. for financial support, and CMC Microsystems for providing the design tools and access to technologies.

REFERENCES

- [1] Mohsen, Field programmable printed circuit board, Patents, US 5,661,409, US 5,377,124, 1995.
- [2] IPC, *Volume 1*, International Technology Roadmap for Electronic Interconnections 2006 - 2007, 2006, p 31.
- [3] Xilinx, *Virtex-5 Packaging and Pinout Specification*, 14 august 2007.
- [4] George Messner, Wim Smit, *Equations for selection of Cost-Efficient Interconnection Design*, AMP-AKZO, 1987.
- [5] Norman, R. U.S. Patent Application Number 11/611,263.
- [6] R. Norman, O. Valorge, Y. Blaquiere, E. Lepercq, Y.B. Basile-Bellavance, Y. El-Alaoui, R. Prytula, Y. Savaria, *An Active Reconfigurable Circuit Board*, NEWCAS, 2008.
- [7] Bakir, M.S. *et al*, "Sea of Leads Compliant I/O Interconnect Process Integration for the Ultimate Enabling of Chips With Low- k Interlayer Dielectrics," *Advanced Packaging, IEEE Transactions on*, vol.28, no.3, pp. 488-494, Aug. 2005.
- [8] Brewer, J.E., "Promises and Pitfalls of WSI", in *Wafer Scale Integration*, Kluwer, pp.1-29, 1989.
- [9] Koren, I., Koren, Z., "Defect Tolerance in VLSI Circuits: Techniques and Yield Analysis", *IEEE*, 86(9):1819-1838, 1998.
- [10] ITRS, *Assembly and Packaging*, The International Technology Roadmap For Semiconductors, 2003.
- [11] Hongbing Fan; Yu-Liang Wu, Crossbar based design schemes for switch boxes and programmable interconnection networks, *Design Automation Conference, 2005. Proceedings of the ASP-DAC 2005. Asia and South Pacific*, vol.2, no., pp. 910-915 Vol. 2, Jan. 2005.
- [12] Ting Wu; Chi-Ying Tsui; Hamdi, M., "A 2 Gb/s 256×256 CMOS crossbar switch fabric core design using pipelined MUX," *IEEE International Symposium on Circuits and Systems*, vol.2, pp. II-568-II-571, 2002.
- [13] S. D. Brown, R. J. Francis, J. Rose, and Z. G. Vranesic, *Field-Programmable Gate Arrays*. Norwell, MA: Kluwer, 1992.