

Reconfigurable Interconnect for next generation systems

Ingrid Verbauwhede, M.-C. Frank Chang

UCLA, EE Dept.

7440 B Boelter Hall

Box 951594, Los Angeles CA 90095

ABSTRACT

This paper describes our vision on the architectures required to build next generation systems. Next generation systems will not be PC centric anymore, but they will be built based on distributed, networked, power constraint embedded systems on a chip (SOC) or systems on a multi chip module (SOM). These architectures will consist of large set of heterogeneous building blocks, many of them reconfigurable at different levels of abstraction. The paper will describe new forms of reconfigurable interconnect and it will include a description of the critical position of reconfigurable interconnect in these architectures. The challenge is not to provide general reconfigurability but to tune it to optimize the energy efficiency.

Categories and Subject Descriptors

C.3 [Computer Systems Organization]: Special purpose and application based systems – *real-time and embedded systems*.

General Terms

Design

Keywords

Architectures, interconnect, reconfiguration, power efficiency, design methods

1. INTRODUCTION

Next generation applications will be deployed based on wireless, networked, power constraint, embedded systems on a chip or systems on a multi-chip module (MCM). Computations will be spatially distributed close to the source of data instead of moving large amounts of unprocessed data to a central CPU unit. The applications will require processing in real time of a diverse set of data, including multimedia events. The systems and architectures need to adapt to changing applications. Yet at the same time they need to operate on an extremely low power budget. Thus there is a fundamental energy flexibility trade-off. This will be illustrated in section 2.

General programmability, such as provided by general purpose micro processors, or general reconfiguration as provided by FPGA's are too power-hungry. Thus, to address this energy-flexi-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SLIP '02, April 6-7, 2002, San Diego, California, USA.

Copyright 2002 ACM 1-58113-481-9/02/0004...\$5.00.

bility trade-off, reconfiguration has to be introduced selectively. Therefore, these systems will consist of a heterogeneous set of domain specific processors and reconfigurable components. The domain specific processors are introduced in section 3.

Tuning the components of an architecture to the application domain, is an exploration of a design space, which we call the reconfiguration hierarchy. It is an extra design axis next to power, area or throughput optimizations. This design space is introduced in section 4.

Interconnect is a fundamental component to provide flexibility, adaptability, reconfiguration, etc. A new vision on interconnect and reconfigurable interconnect is introduced in section 5.

Possible applications for reconfigurable interconnect are introduced in section 6. Conclusions are given in section 7.

2. ENERGY-FLEXIBILITY TRADE-OFF

The traditional approach, i.e. the development of applications in general purpose languages (such as C, C++ or Java), running on general purpose platforms, such as general purpose micro processors or FPGA's, might provide the required reconfiguration or reprogrammability, but it fails in power efficiency and it might even fail to deliver the required throughput.

The root of this problem is the energy inefficiency of general purpose solutions, as illustrated in Table 1. The table illustrates an AES encryption algorithm (a function likely to be part of a modern embedded information system) that is implemented on three different platforms: a specialized processor with CMOS standard cell technology, a FPGA, and a Pentium III microprocessor.

Table 1: Figure of Merit for three AES implementations

AES 128 bit key 128 bit data	Throughput	Power	Figure of Merit (Gb/s/W)
AES Proc 0.18 μ m CMOS [5]	1.28 Gb/s	56 mW	228 (100%)
AES on FPGA ^a	640 Mb/s	1.63 W	0.39 (1.7%)
AES on Pen- tium III ^{b c}	607 Mb/s	41.4W	0.015 (0.06%)

a.Xilinx Webpack + Xilinx Virtex Power Estimator on AES Design

b.Helger Lipmaa, <http://www.cs.tut.fi/~helger/aes/rijndael.html>: PIII assembly optimized.

c.Intel Pentium III datasheet, 1.13GHz (Vcc=1.8V, Icc =23A)

While the three implementations have approximately the same processing throughput, their power consumption varies three orders of magnitude. Therefore the energy efficiency (throughput normalized to power consumption) varies with three orders of magnitude also. We call the AES processor a *domain specific processor* because it is highly tuned towards the implementation of one particular type of function, in this case AES encryption. The AES processor is designed in such a way that it still can cover a family of functions which have different operating parameters and different coarse grain architecture organization.

3. DOMAIN SPECIFIC PROCESSOR

A domain specific processor holds the middle ground between a fully fixed function design (ASIC) and a general purpose, fully programmable architecture. The latter can be either a FPGA or a processor. Both a general purpose FPGA and micro processor differ only in their programming style (spatial as opposed to sequential). As illustrated in Figure 1, a programmable system consists of a platform and an application. Thus we are considering designs in which programmability is reduced to a certain extent, i.e. a certain application domain. Reducing programmability implies a more efficient use of implementation technology, which directly impacts energy efficiency. As an example of this, the programming efficiency of a contemporary high end FPGA is between 5 and 10 bits per system gate (sources: Xilinx Virtex II, Altera Apex II). These programming bits are used to configure circuit topology using multiplexers, pass gates and so on. Per useful functional gate on an FPGA, there are 5 to 10 gates required for fixing that functionality.

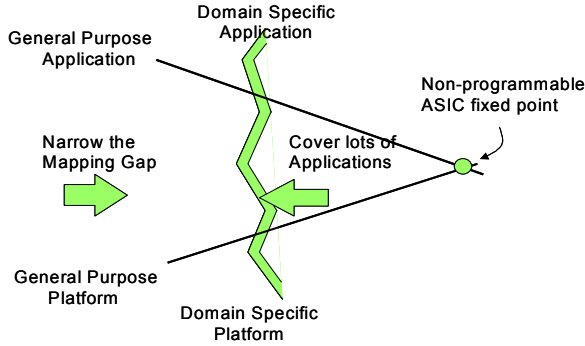


Figure 1: Domain specific processor

4. RECONFIGURATION HIERARCHY

The fundamental reason for the energy inefficiency of general purpose solutions is that the architecture is not matched towards the application or application domain. Tuning the reconfiguration towards the application domain is the task of deciding which component of an architecture at what abstraction level at what binding rate needs to be fixed or left reconfigurable. This corresponds to choosing a design point into a design space, which we call the reconfiguration hierarchy [8].

This is illustrated in table 2: a processor has the following fundamental components: execution units (including control), communication and storage. Each of these components can be tuned separately at different abstraction levels. Several examples are given in table 2. The third axis is the binding rate: it describes the when reconfiguration data is sent to the processing part, ranging from implementation time to design time binding [4].

The reconfiguration for execution units and instruction sets has received attention by the research community. One example is described in [7]: it uses parametrizable blocks which are at a higher abstraction level than CLB units of FPGA's. Another example is the introduction of a specialized instruction sets to improve the performance of multimedia applications on general purpose CPU's, such as the MMX instruction set extension. Similarly, DSP processors have datapaths and even more important memory and interconnect architectures that are optimized to the data processing nature of signal and communication algorithms [9].

Yet, the exploration of the energy flexibility curve for interconnect has not been explored.

5. RECONFIGURABLE INTERCONNECT

Interconnect is used to transport data over a transmission medium between the different components of a system, called senders and receivers. Current interconnect schemes are all based on space or time division. Similarly, reconfiguration is confined to a space and a time axis[4]. New interconnect schemes are proposed that introduce frequency and code division or a combination of all above.

5.1 SDMA - Space division multiple access

If every sender, receiver pair has its own physical transmission medium, e.g. its own metal wire, there will be no access conflicts. But the amount of wires will grow exponential with the required number of sender, receiver pairs. Hence, to keep the space, i.e. number of wires, under control, space and time division multiple

Table 2: Reconfiguration design space [8]

	Communication	Storage	Processing
Implementation	Switches Muxes	RAM Organization	CLB Parametrizable IP-block
Micro-Architecture	Crossbar Busses	Register File Size Cache Architecture	Execution Unit Type Interpreter Levels
Instruction Set Architecture	Size of address/data bus	Register Set Memory Architecture	Custom Instructions Interrupt Architecture
Process Architecture/ Systems Architecture	RF Reconfigurable Interconnection	Buffer Size	Number and type of asynchronous processes and tasks

Figure 2: Reconfiguration design space. A 3D plot showing the relationship between Computation Abstraction Level (Y-axis), Binding rate (X-axis), and Reconfiguration feature (Z-axis). The plot shows a green arrow pointing from the General Purpose Application towards the Domain Specific Application, and another green arrow pointing from the Domain Specific Application towards the Non-programmable ASIC fixed point.

access schemes are introduced. This has been the reason to add ever more layers of metal.

5.2 TDMA - Time division multiple access

Given a set of metal layers and a set of senders and receivers, the current approach to solve the interconnect demand is to introduce time division. This has been done in multiple approaches and multiple levels of hierarchy.

Examples of bus architectures are the Amba bus [1] or the MicroNetwork [10]. Since resources are limited there is always a latency, throughput, bandwidth flexibility trade-off. To address these issues, newer approaches look at modeling the interconnect as a switched network [3].

5.3 CDMA - Code division multiple access

In [2], a novel RF/wireless interconnect scheme is proposed. Unlike the “passive” metal interconnect, the “active” RF/wireless interconnect is based on low loss and dispersion free microwave signal transmission and multiple access algorithms, well known in the radio communication infrastructure.

The miniature LAN is illustrated in Figure 2. Two VLSI circuits are housed in a MCM package. The I/O pads are the users, the capacitive couplers are the near-field antennas and the micro-strip line (MTL) or coplanar waveguide (CPW) are used as a shared transmission medium. Code division multiple access (CDMA) or frequency division multiple access (FDMA) communication algorithms can be used to address the cross-channel interference associated with a shared medium.

An example of a CDMA based interconnect at baseband level is shown in Figure 3. The data of two users are multiplied with two

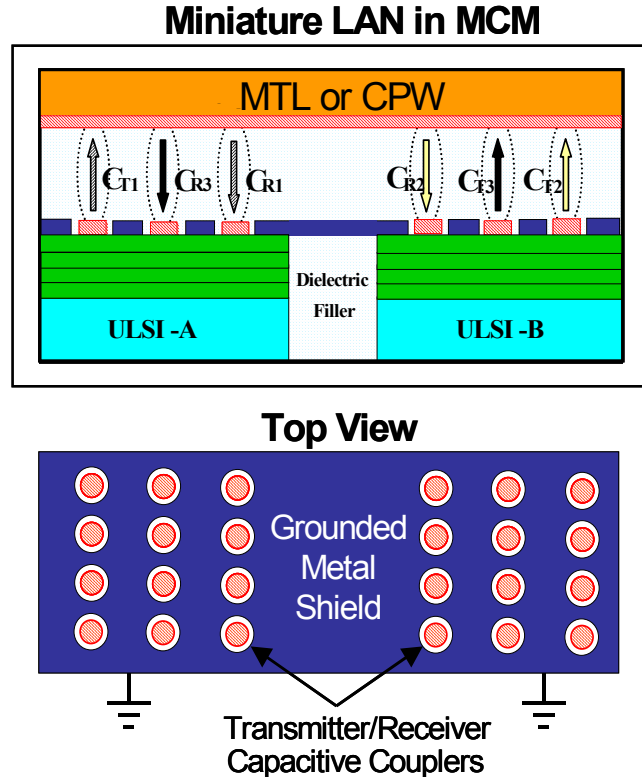


Figure 2: RF Interconnect

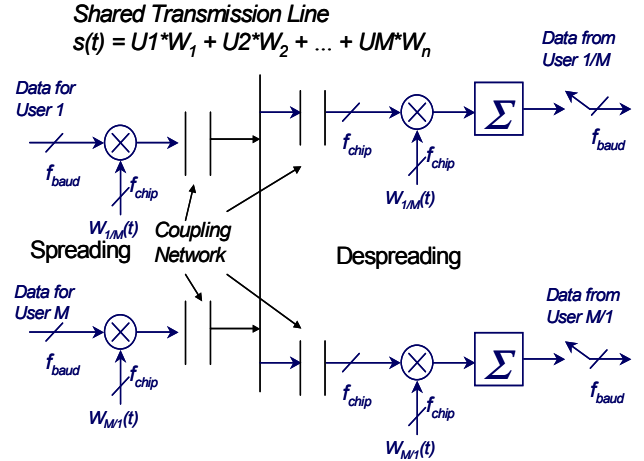


Figure 3: CDMA based interconnect

orthogonal spreading codes, e.g. Walsh codes. Each receiver has its own Walsh code and thus can retrieve its data from the superposed signal. By reprogramming the spreading codes, the interconnect can easily be reconfigured.

5.4 FDMA - Frequency division multiple access

To increase the capacity of the system the baseband signals can be modulated with a radio frequency. Combining this with the introduction of frequency bands, a frequency division multiple access (FDMA) scheme is realized. This is illustrated in Figure 4. The data stream of the k -th user is multiplied by its sinusoidal carrier, $A_k \cos(2\pi f_k t)$. The resulting signal is filtered through a band pass filter BPF_{TK}, which is then coupled on the shared medium. The receiver has to demodulate the signal to recover the data. Similar to radio receivers, pre-amplifiers, low pass or band pass filters, threshold comparator modules, etc. can be added.

Similar to cellular systems, FDMA can be combined with TDMA (such as the GSM system) or with CDMA and even more possibilities open up.

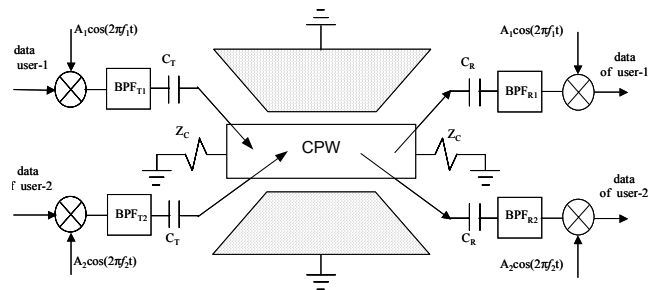


Figure 4: Top level schematic of a FDMA interconnect

6. APPLICATIONS FOR RECONFIGURABLE INTERCONNECT

The importance of reconfigurable interconnect has been recognized in many systems. Yet, providing general multiplexer based reconfigurable interconnect architectures can be very expensive in terms of area and power. It is described in [6], that for a Xilinx XC4003A FPGA, 65% of the power is attributed to interconnect,

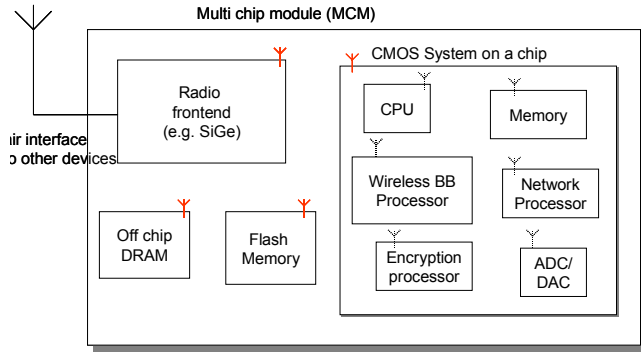


Figure 5: Intra and Inter chip RF interconnect

21% to clock power, 9% to I/O power and only 5% to the actual calculations (CLB) power. Although this is an older FPGA device, new devices such as the Virtex II [11], still focus on providing high speed switching matrices. They also include a large set of components to address on-chip and off-chip interconnect, such as multi-gigabit serial links.

In our opinion, RF/Wireless interconnect can be used to connect components in one IC device, one MCM module, or it can be used at the board level. This is illustrated in Figures 5, 6 and 7.

In Figure 5, an architecture, typical for the next generation systems on chip, is shown. It might e.g. be used as a node in a distributed wired or wireless network. At the IC level, it will consist of a heterogeneous set of reconfigurable and reprogrammable components. These components are connected together by a first level of RF interconnect. This IC is a component in a MCM module, that might again include a set of heterogeneous devices. It will have its own level of RF reconfigurable interconnect.

The RF wireless interconnect is also useful at the board level. One example is given in Figures 6 and 7. Figure 7 shows a conventional SDRAM interface between a CPU and a DRAM memory. Figure 6 shows how multiple PCB lines can be replaced by one multi-signal line.

7. CONCLUSIONS

In this contribution we have shown that flexibility is not for free. For embedded systems, energy flexibility is a trade off. We approach this trade-off by introducing domain specific processors and by exploring the associated reconfiguration hierarchy. Recon-

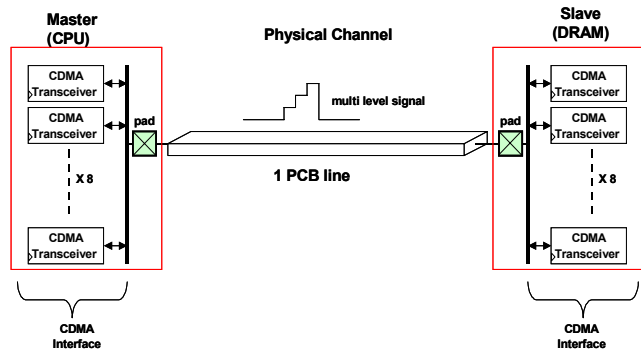


Figure 6: CDMA based interconnect for off chip CPU - DRAM interface

figurative interconnect is a key component for domain specific processors. The usage of new CDMA and FDMA based RF interconnect has been shown for practical inter and intra chip reconfigurable systems.

8. ACKNOWLEDGMENTS

The authors would like to acknowledge the contributions from and stimulating discussions with Patrick Schaumont and Jongsun Kim.

9. REFERENCES

- [1] ARM, Amba Specification, available from www.arm.com
- [2] M.-C. F. Chang, V. Roychowdhury, L. Zhang, H. Shin, Y. Qian, "RF/Wireless Interconnect for Inter and Intra-chip communications," Proceedings of IEEE, Vol. 89, No.4, April 2001, pg. 456-466.
- [3] W. Dally, B. Towles, "Route packets, not wires: On-Chip Interconnection networks," Proc. DAC 2001, pg. 684-689.
- [4] A. Dehon, J. Wawrzynek, "Reconfigurable computing: what, why, and implications for design automation," Proc. DAC 1999, pg. 610-615.
- [5] H. Kuo, I. Verbauwhede, P. Schaumont, "A 2.29 Gb/s, 56 mW non-pipelined Rijndael AES Encryption IC in a 1.8V, 0.18 μ m CMOS technology," Proceedings IEEE Custom Integrated Circuits Conference, CICC 2002, May 2002.
- [6] E. Kusse, J. Rabaey, "Low-energy embedded FPGA structures," Proc. 1998 International Symposium on Low Power Electronics and Design, ISLPED 1998, pg. 155 -160.
- [7] S. Ogrenci, E. Bozorgzadeh, R. Kastner, M. Sarrafzadeh, "SPS: A Strategically programmable system," Proc. Reconfigurable Architectures Workshop, RAW 2001, April 2001.
- [8] P. Schaumont, I. Verbauwhede, K. Keutzer, M. Sarrafzadeh, "A Quick Safari through the Reconfiguration Jungle," Proceedings Design Automation Conference, DAC-2001, Las Vegas, June 2001, pg. 172-177.
- [9] I. Verbauwhede, C. Nicol, "Low Power DSP's for Wireless Communications," Proc. ISLPED 2000, pg. 303-310.
- [10] D. Wingard, "MicroNetwork-Based Integration for SOC," Proc. DAC 2001, pg. 673-677.
- [11] Xilinx Virtex II, www.xilinx.com

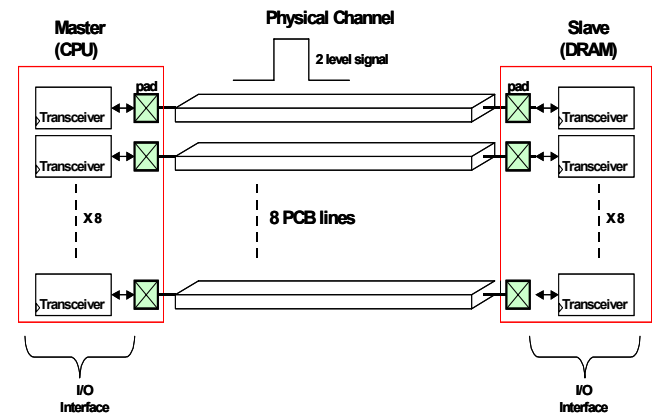


Figure 7: Conventional SDRAM interconnect for off chip CPU - DRAM interface.