

Reconfigurable Computing: Architectures, Design Methods, and Applications

T.J. Todman*, G.A. Constantinides**, S.J.E. Wilton***, O. Mencer*, W. Luk* and P.Y.K. Cheung**

*Department of Computing, Imperial College London, UK

**Department of Electrical and Electronic Engineering, Imperial College London, UK

***Department of Electrical and Computer Engineering, University of British Columbia, Canada

Abstract—Reconfigurable computing is becoming increasingly attractive for many applications. This survey covers three aspects of reconfigurable computing: architectures, design methods, and applications. Our paper includes recent advances in reconfigurable architectures, such as the Altera Stratix II and Xilinx Virtex 4 FPGA devices. We identify major trends in general-purpose and special-purpose design methods. We describe a variety of driving applications for reconfigurable technology, including data encryption, video processing, network security, and image generation. It is shown that reconfigurable computing designs are capable of achieving up to 500 times speedup and 70% energy savings over microprocessor implementations for specific applications.

I. INTRODUCTION

Reconfigurable computing is rapidly establishing itself as a major discipline that covers various subjects of learning, including both computing science and electronic engineering. Reconfigurable computing involves the use of reconfigurable devices, such as Field Programmable Gate Arrays (FPGAs), for computing purposes. Reconfigurable computing is also known as configurable computing or custom computing, since many of the design techniques can be seen as customising a computational fabric for specific applications [99].

Reconfigurable computing systems often have impressive performance. Consider, as an example, the point multiplication operation in Elliptic Curve cryptography. For a key size of 270 bits, it has been reported [166] that a point multiplication can be computed in 0.36 ms with a reconfigurable computing design implemented in an XC2V6000 FPGA at 66 MHz. In contrast, an optimised software implementation requires 196.71 ms on a dual-Xeon computer at 2.6 GHz; so the reconfigurable computing design is more than 540 times faster, while its clock speed is almost 40 times slower than the Xeon processors. Further material on reconfigurable computing for encryption applications can be found in Section V-A.

Is this speed advantage of reconfigurable computing over traditional microprocessors a one-off or a sustainable trend? Recent research suggests that it is a trend rather than a one-off for a wide variety of applications: from image processing [64] to floating-point operations [173].

Sheer speed, while important, is not the only strength of reconfigurable computing. Another compelling advantage is reduced energy and power consumption. In a reconfigurable system, the circuitry is optimized for the application, such that the power consumption will tend to be much lower than that

for a general-purpose processor. A recent study [157] reports that moving critical software loops to reconfigurable hardware results in average energy savings of 35% to 70% with an average speedup of 3 to 7 times, depending on the particular device used.

Other advantages of reconfigurable computing include a reduction in size and component count (and hence cost), improved time-to-market, and improved flexibility and upgradability. These advantages are especially important for embedded applications. Indeed, there is evidence [174] that embedded systems developers show a growing interest in reconfigurable computing systems, especially with the introduction of soft cores which can contain one or more instruction processors [7], [144], [54], [97], [145], [190].

In this paper, we present a survey of modern reconfigurable system architectures, design methods, and applications. Although we also provide background information on notable aspects of older technologies, our focus is on the most recent architectures, design methods, and applications, as well as the trends that will drive each of these areas in the near future. In other words, we intend to complement other survey papers [21], [35], [100], [141], [167] by:

- 1) providing an up-to-date survey of material that appears after the publication of the papers mentioned above;
- 2) identifying explicitly the main trends in architectures, design methods and applications for reconfigurable computing;
- 3) examining reconfigurable computing from a perspective different from existing surveys, for instance classifying design methods as special-purpose and general-purpose;
- 4) offering various direct comparisons of technology options according to a selected set of metrics from different perspectives.

The rest of the paper is organised as follows. Section 2 contains background material that motivates the reconfigurable computing approach. Section 3 describes the structure of reconfigurable fabrics, showing how various researchers and vendors have developed fabrics that can efficiently accelerate time-critical portions of applications. Section 4 then covers recent advances in the development of design methods that map applications to these fabrics, and distinguishes between those which employ special-purpose and general-purpose optimization methods. Section 5 presents specific applications that

are suitable for reconfigurable computing, and shows how the applications can be designed with the target platforms in mind. Finally, Section 6 concludes the paper and summarises the main trends in architectures, design methods and applications of reconfigurable computing.

II. BACKGROUND

Many of today's compute-intensive applications require more processing power than ever before. Applications such as streaming video, image recognition and processing, and highly interactive services are placing new demands on the computation units that implement these applications. At the same time, the power consumption targets, the acceptable packaging and manufacturing costs, and the time-to-market requirements of these computation units are all decreasing rapidly, especially in the embedded hand-held devices market. Meeting these performance requirements under the power, cost, and time-to-market constraints is becoming increasingly challenging.

In the following, we describe three ways of supporting such processing requirements: high-performance microprocessors, application-specific integrated circuits, and reconfigurable computing systems.

High-performance microprocessors provide an off-the-shelf means of addressing processing requirements described earlier. Unfortunately for many applications, a single processor, even an expensive state-of-the-art processor, is not fast enough. In addition, the power consumption and cost of state-of-the-art processors place them out-of-reach for many embedded applications. Even if microprocessors continue to follow Moore's Law so that their density doubles every 18 months, they may still be unable to keep up with the requirements of some of the most aggressive embedded applications.

Application-specific integrated circuits (ASICs) provide another means of addressing these processing requirements. Unlike a software implementation, an ASIC implementation provides a natural mechanism for implementing the large amount of parallelism found in many of these applications. In addition, an ASIC circuit does not need to suffer from the serial (and often slow and power-hungry) instruction fetch, decode, and execute cycle that is at the heart of all microprocessors. Finally, an ASIC can contain just the right mix of functional units for a particular application; in contrast, an off-the-shelf microprocessor contains a fixed set of functional units which must be selected to satisfy a wide variety of applications.

Despite the advantages of an ASIC, they are often infeasible or uneconomical for many embedded systems. This is primarily due to two factors: the cost of producing an ASIC often due to the masks cost, and the time to develop a custom integrated circuit, can both be unacceptable. Only the very highest-volume applications would the improved performance and lower per-unit price warrant the high non-recurring engineering (NRE) cost of designing an ASIC.

A third means of providing this processing power is a reconfigurable computing system. A reconfigurable comput-

ing system typically contains one or more processors and a reconfigurable fabric upon which custom functional units can be built. The processor(s) executes sequential and non-critical code, while code that can be efficiently mapped to hardware can be "executed" by processing units that have been mapped to the reconfigurable fabric. Like a custom integrated circuit, the functions that have been mapped to the reconfigurable fabric can take advantage of the parallelism achievable in a hardware implementation. Also like an ASIC, the embedded system designer can produce the right mix of functional and storage units in the reconfigurable fabric, providing a computing structure that matches the application.

Unlike an ASIC, however, a new fabric need not be designed for each application. A given fabric can implement a wide variety of functional units. This means that a reconfigurable computing system can be built out of off-the-shelf components, significantly reducing the long design-time inherent in an ASIC implementation. Also unlike an ASIC, the functional units implemented in the reconfigurable fabric can change over time. This means that as the environment or usage of the embedded system changes, the mix of functional units can adapt to better match the new environment. The reconfigurable fabric in a handheld device, for instance, might implement large matrix multiply operations when the device is used in one mode, and large signal processing functions when the device is used in another mode.

Typically, not all of the embedded system functionality needs to be implemented by the reconfigurable fabric. Only those parts of the computation that are time-critical and contain a high degree of parallelism need to be mapped to the reconfigurable fabric, while the remainder of the computation can be implemented by a standard instruction processor. The interface between the processor and the fabric, as well as the interface between the memory and the fabric, are therefore of the utmost importance.

Despite the compelling promise of reconfigurable computing, it has limitations of which designers should be aware. For instance, the flexible routing on the bit level tends to produce large silicon area and performance overhead when compared with ASIC technology. Hence for large volume production of designs in applications without the need for field upgrade, ASIC technology or gate array technology can still deliver higher performance design at lower unit cost than reconfigurable computing technology. However, since FPGA technology tracks advances in memory technology and has demonstrated impressive advances in the last few years, many are confident that the current rapid progress in FPGA speed, capacity and capability will continue, together with the reduction in price.

It should be noted that the development of reconfigurable systems is still a maturing field. There are a number of challenges in developing a reconfigurable system. We describe three of such challenges below.

First, the structure of the reconfigurable fabric and the interfaces between the fabric, processor(s), and memory must be very efficient. Some reconfigurable computing systems

use a standard Field-Programmable Gate Array [3], [6], [89], [112], [131], [189] as a reconfigurable fabric, while others adopt custom-designed fabrics [34], [51], [52], [63], [71], [103], [106], [110], [138], [147], [150], [164].

Another challenge is the development of computer-aided design and compilation tools that map an application to a reconfigurable computing system. This involves determining which parts of the application should be mapped to the fabric and which should be mapped to the processor, determining when and how often the reconfigurable fabric should be reconfigured, which changes the functional units implemented in the fabric, as well as the specification of algorithms for efficient mappings to the reconfigurable system.

Finally, the applications themselves need to be designed with the target reconfigurable system in mind. No matter how good the architecture, and how good the tools, if the application is not designed in such a way that the characteristics of the reconfigurable fabric can be exploited, the reconfigurable system may become no faster – or even sometimes slower – than a standard microprocessor implementation.

In this paper, we provide a survey of reconfigurable computing, focusing our discussion on each of the three issues described above. In the next section, we provide a survey of various architectures that are found useful for reconfigurable computing; material on design methods and applications will follow.


III. ARCHITECTURES

We shall first describe system-level architectures for reconfigurable computing. We then present various flavours of reconfigurable fabric. Finally we identify and summarise the main trends.


A. System-level architectures

A reconfigurable system typically consists of one or more processors, one or more reconfigurable fabrics, and one or more memories. Reconfigurable systems are often classified according to the degree of coupling between the reconfigurable fabric and the CPU. Compton et al [35] present the four classifications shown in Figure 1(a-d). In Figure 1(a), the reconfigurable fabric is in the form of one or more stand-alone devices. The existing input and output mechanisms of the processor are used to communicate with the reconfigurable fabric. In this configuration, the data transfer between the fabric and the processor is relatively slow, so this architecture only makes sense for applications in which a significant amount of processing can be done by the fabric without processor intervention. Emulation systems often take on this sort of architecture [24], [109].


Figure 1(b) and Figure 1(c) shows two intermediate structures. In both cases, the cost of communication is lower than that of the architecture in Figure 1(a). Architectures of these types are described in [8], [63], [71], [90], [138], [150], [175], [185]. Next, Figure 1(d) shows an architecture in which the processor and the fabric are very tightly coupled; in this case, the reconfigurable fabric is part of the processor itself; perhaps




(a) External stand-alone processing unit.




(b) Attached processing unit.



(c) Co-processor.



(d) Reconfigurable functional unit.



(e) Processor embedded in a reconfigurable fabric.

Fig. 1. Five classes of reconfigurable systems. The first four are adapted from [35].

TABLE I
SUMMARY OF SYSTEM ARCHITECTURES.

Class	CPU to memory bandwidth	Shared memory size	Fine Grained or Coarse Grained	Example Application
(a) External stand-alone processing unit RC2000 [30]	528MB/s	152MB	Fine Grained	Video processing
(b) / (c) Attached processing unit / Co-processor Garp [71] Morphosys [150]	2128MB/s 800MB/s	16KB cache + 4GB 2048 bytes	Fine Grained Coarse Grained	Encryption, sorting Video compression
(d) Reconfigurable functional unit Chess [103]	6400MB/s	12288 bytes	Coarse Grained	Video processing
(e) Processor embedded in a reconfigurable fabric Xilinx Virtex II Pro [189]	1600MB/s	1172KB	Fine Grained	Video compression

forming a reconfigurable sub-unit that allows for the creation of custom instructions. Examples of this sort of architecture have been described in [103], [110], [134], [164].

Figure 1(e) shows a fifth organization. In this case, the processor is embedded in the programmable fabric. The processor can either be a “hard” core [5], [188], or can be a “soft” core which is implemented using the resources of the programmable fabric itself [7], [144], [54], [97], [145], [190].

A summary of the above organisations can be found in Table I. Note the bandwidth is the theoretical maximum available to the CPU: for example, in Chess [103], we assume that each block RAM is being accessed at its maximum rate.


B. Reconfigurable fabric

The heart of any reconfigurable system is the reconfigurable fabric. The reconfigurable fabric consists of a set of reconfigurable functional units, a reconfigurable interconnect, and a flexible interface to connect the fabric to the rest of the system. In this section, we review each of these components, and show how they have been used in both commercial and academic reconfigurable systems.


A common theme runs through this entire section: in each component of the fabric, there is a tradeoff between flexibility and efficiency. A highly flexible fabric is typically much larger and much slower than a less flexible fabric. On the other hand, a more flexible fabric is better able to adapt to the application requirements.

In the following discussions, we will see how this tradeoff has influenced the design of every part of every reconfigurable system. A summary of the main features of various architectures can be found in Table II.

1) *Reconfigurable functional units*: Reconfigurable functional units can be classified as either coarse-grained or fine-grained. A fine-grained functional unit can typically implement a single function on a single (or small number) of bits. The most common kind of fine-grained functional units are the small lookup tables that are used to implement the bulk of the logic in a commercial field-programmable gate array. A coarse-grained functional unit, on the other hand, is typically much larger, and may consist of arithmetic and logic units



(a) Three-input lookup table.



(b) Cluster of lookup tables.

Fig. 2. Fine-grained reconfigurable functional units.

(ALUs) and possibly even a significant amount of storage. In this section, we describe the two types of functional units in more detail.

Many reconfigurable systems use commercial FPGAs as a reconfigurable fabric. These commercial FPGAs contain many three to six input lookup tables, each of which can be thought of as a very fine-grained functional unit. Figure 2(a)

TABLE II
COMPARISON OF RECONFIGURABLE FABRICS AND DEVICES.

Fabric or Device	Fine Grained or Coarse Grained	Base Logic Component	Routing Architecture	Embedded Memory	Special Features
Actel ProASIC+ [3]	Fine	3-input block	Horizontal and vertical tracks	256x9 bit blocks	Flash-based
Altera Excalibur [5]	Fine	4-input Lookup Tables	Horizontal and vertical tracks	2Kbit memory blocks	ARMv4T Embedded Processor
Altera Stratix II [6]	Fine/Coarse	8-input Adaptive Logic Module	Horizontal and vertical tracks	512 bits, 4Kbits, and 512Kbit blocks	DSP blocks
Garp [71]	Fine	logic or arithmetic functions on four 2-bit input words	2-bit Buses in horizontal and vertical columns	External to fabric	
Xilinx Virtex II Pro [188]	Fine	4-input Lookup Tables	Horizontal and vertical tracks	18Kbit blocks	Embedded Multipliers, PowerPC 405 Processor
Xilinx Virtex II [189]	Fine	4-input Lookup Tables	Horizontal and vertical tracks	18Kbit blocks	Embedded Multipliers
DReAM [12]	Coarse	8-bit ALUs	16-bit local and global buses	Two 16x8 Dual Port memory	Targets mobile applications
Elixent D-fabrix [52]	Coarse	4-bit ALUs	4-bit buses	256x8 memory blocks	
HP Chess [103]	Coarse	4-bit ALUs	4-bit buses	256x8 bit memories	
IMEC ADRES [106]	Coarse	32-bit ALUs	32-bit buses	Small register files in each logic component	
Matrix [110]	Coarse	8-bit ALUs	Hierarchical 8-bit buses	256x8 bit memories	
MorpoSys [150]	Coarse	ALU and Multiplier, and Shift Units	Buses	External to fabric	
Piperench [63]	Coarse	8-bit ALUs	8-bit Buses	External to fabric	Functional units arranged in 'stripes'
RaPiD [51]	Coarse	ALUs	buses	Embedded memory blocks	
Silicon Hive Avispa [147]	Coarse	ALUs, Shifters, Accumulators, and Multipliers	Buses	Five embedded memories	

illustrates a lookup table; by shifting in the correct pattern of bits, this functional unit can implement any single function of up to three inputs – the extension to lookup tables with larger numbers of inputs is clear. Typically, lookup tables are combined into clusters, as shown in Figure 2(b). Figure 3 shows clusters in two popular FPGA families. Figure 3(a) shows a cluster in the Altera Stratix device; Altera calls these clusters “Logic Array Blocks [6]. Figure 3(b) shows a cluster in the Xilinx architecture [189]; Xilinx calls these clusters “Configurable Logic Blocks” (CLBs). In the Altera diagram, each block labeled “LE” is a lookup table, while in the Xilinx diagram, each “slice” contains two lookup tables. Other commercial FPGAs are described in [3], [89], [112], [131].

Reconfigurable fabrics containing lookup tables are very flexible, and can be used to implement any digital circuit. However, compared to the coarse-grained structures in the next subsection, these fine-grained structures have significantly more area, delay, and power overhead. Recognizing that these fabrics are often used for arithmetic purposes, FPGA companies have added additional features such as carry-chains and cascade-chains to reduce the overhead when implementing common arithmetic and logic functions. Figure 4 shows how the carry and cascade chains, as well as the ability to break a 4-input lookup table into four two-input lookup tables can be exploited to efficiently implement carry-select adders [6]. The

multiplexers and the exclusive-or gate in Figure 4 are included as part of each logic array block, and need not be implemented using other lookup tables.

The example in Figure 4 shows how the efficiency of commercial FPGAs can be improved by adding architectural support for common functions. We can go much further than this though, and embed significantly larger, but far less flexible, reconfigurable functional units. There are two kinds of devices that contain coarse-grained functional units; modern FPGAs, which are primarily composed of fine-grained functional units, are increasingly being enhanced by the inclusion of larger blocks. As an example, the Xilinx Virtex device contains embedded 18-bit by 18-bit multiplier units [189]. When implementing algorithms requiring a large amount of multiplication, these embedded blocks can significantly improve the density, speed and power of the device. On the other hand, for algorithms which do not perform multiplication, these blocks are rarely useful. The Altera Stratix devices contain a larger, but more flexible embedded block, called a DSP block, shown in Figure 5 [6]. Each of these blocks can perform accumulate functions as well as multiply operations. The comparison between the two devices clearly illustrates the flexibility and overhead tradeoff; the Altera DSP block may be more flexible than the Xilinx multiplier, however, it consumes more chip area and runs somewhat slower.

The commercial FPGAs described above contain both fine-

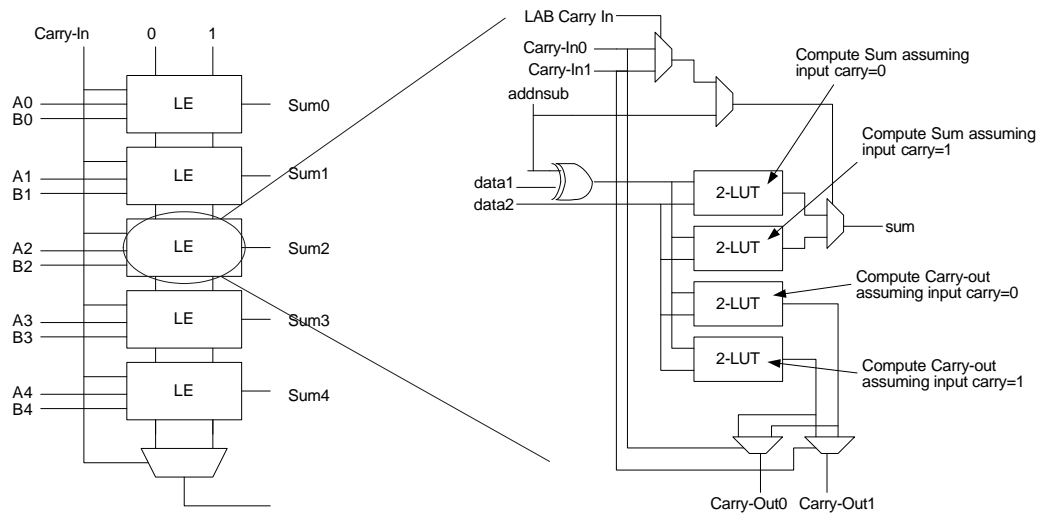
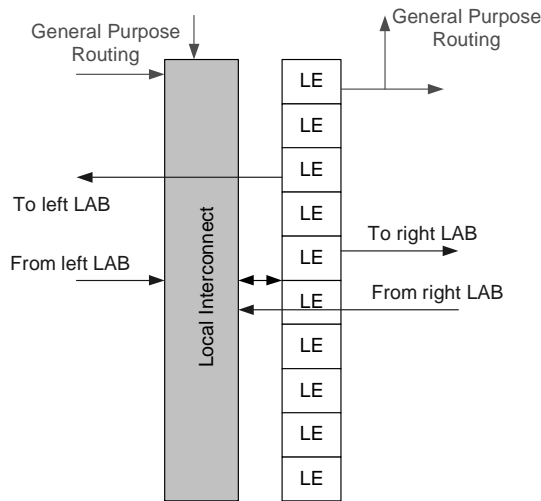
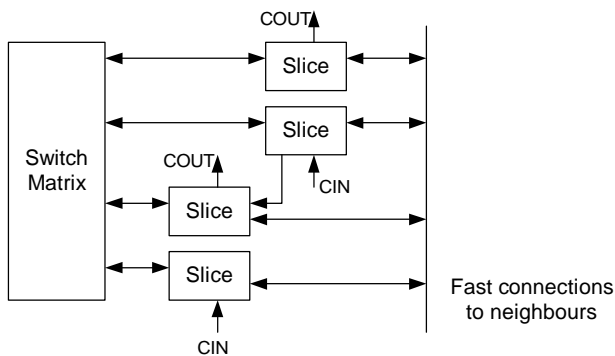


Fig. 4. Implementing a carry-select adder in an Altera Stratix device [6]. 'LUT' denotes 'Lookup Table'.



(a) Altera Logic Array Block [6]



(b) Xilinx Configurable Logic Block [189]

Fig. 3. Commercial Logic Block Architectures.

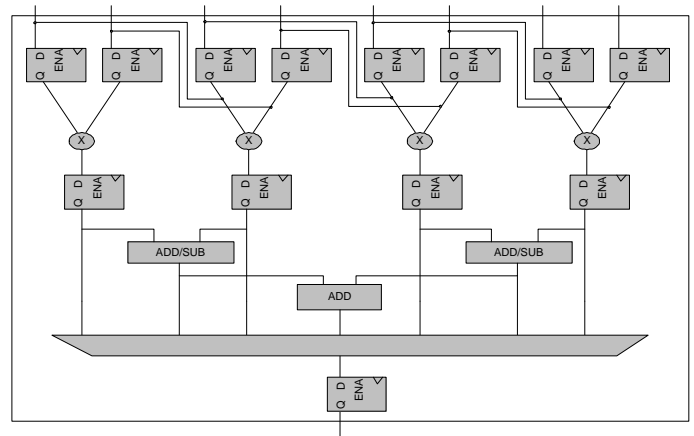


Fig. 5. Altera DSP Block [6].

grained and coarse-grained blocks. There are also devices which contain only coarse-grained blocks [34], [51], [63], [103], [106], [150]. An example of a coarse-grained architecture is the ADRES architecture which is shown in Figure 6 [106]. Each reconfigurable functional unit in this device contains a 32-bit ALU which can be configured to implement one of several functions including addition, multiplication, and logic functions, with two small register files. Clearly, such a functional unit is far less flexible than the fine-grained functional units described earlier; however if the application requires functions which match the capabilities of the ALU, these functions can be very efficiently implemented in this architecture.

2) *Reconfigurable interconnects*: Regardless of whether a device contains fine-grained functional units, coarse-grained functional units, or a mixture of the two, the functional units needed to be connected in a flexible way. Again, there is a tradeoff between the flexibility of the interconnect (and hence the reconfigurable fabric) and the speed, area, and power-efficiency of the architecture.

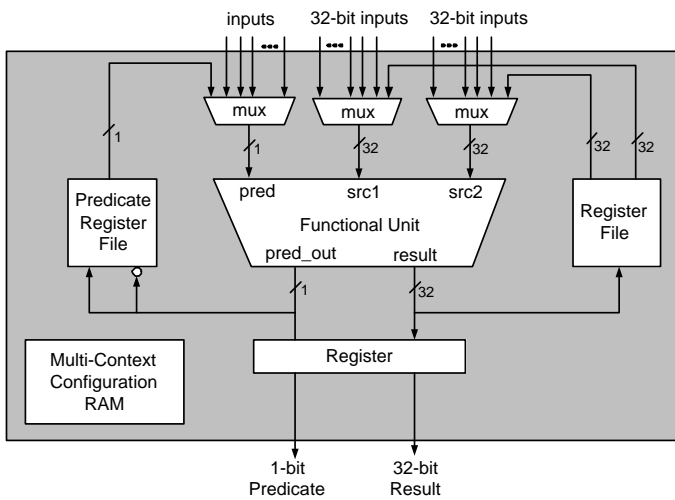


Fig. 6. ADRES Reconfigurable Functional Unit [106].





Fig. 7. Routing architectures.


As before, reconfigurable interconnect architectures can be classified as fine-grained or coarse-grained. The distinction is based on the granularity with which wires are switched. This is illustrated in Figure 7, which shows a flexible interconnect between two buses. In the fine-grained architecture in Figure 7(a), each wire can be switched independently, while in Figure 7(b), the entire bus is switched as a unit. The fine-grained routing architecture in Figure 7(a) is more flexible, since not every bit needs to be routed in the same way; however, the coarse-grained architecture in Figure 7(b) contains far fewer programming bits, and hence suffers much less overhead.

Fine-grained routing architectures are usually found in commercial FPGAs. In these devices, the functional units are typically arranged in a grid pattern, and they are connected using horizontal and vertical channels. Significant research has been performed in the optimization of the topology of this interconnect [16], [95].

Coarse-grained routing architectures are commonly used in devices containing coarse-grained functional units. Figure 8 shows two examples of coarse-grained routing architectures. The routing architecture in Figure 8(a) is used in the Totem reconfigurable system [34]; the interconnect is designed to



(a) Totem coarse-grained routing architecture [34]



(b) Silicon Hive coarse-grained routing architecture [147]

Fig. 8. Example coarse-grained routing architectures.

be very flexible, and to provide arbitrary connection patterns between functional units. On the other hand, the routing architecture in Figure 8(b), which is used in the Silicon Hive reconfigurable system, is less flexible, and hence faster and smaller [147]. In the Silicon Hive architecture, only connections between units that are likely to communicate are provided.

3) *Emerging directions:* Several emerging directions will be covered in the following. These directions include low-power techniques, asynchronous architectures, and molecular microelectronics.

- Low-power techniques. Early work explores the use of low-swing circuit techniques to reduce the power consumption in a hierarchical interconnect for a low-energy FPGA [59]. Recent work involves: (a) activity reduction in power-aware design tools, with energy saving of 23% [88]; (b) leakage current reduction methods such as gate biasing and multiple supply-voltage integration, with up to two times leakage power reduction [132]; and (c) dual supply-voltage methods with the lower voltage assigned to non-critical paths, resulting in an average power reduction of 60% [58].
- Asynchronous architectures. There is an emerging interest in asynchronous FPGA architectures. An asynchronous version of Pipherch [63] is estimated to improve performance by 80%, at the expense of a significant increase in configurable storage and wire count [82]. Other efforts in this direction include fine-grained asyn-

chronous pipelines [165], quasi delay-insensitive architectures [186], and globally-asynchronous locally-synchronous techniques [137].

- Molecular microelectronics. In the long term, molecular techniques offer a promising opportunity for increasing the capacity and performance of reconfigurable computing architectures [23]. Current work is focused on developing Programmable Logic Arrays based on molecular-scale nano-wires [47], [183].

C. Architectures: main trends

The following summarises the main trends in architectures for reconfigurable computing.

1) *Coarse-grained fabrics*: As reconfigurable fabrics are migrated to more advanced technologies, the cost (in terms of both speed and power) of the interconnect part of a reconfigurable fabric is growing. Designers are responding to this by increasing the granularity of their logic units, thereby reducing the amount of interconnect needed. In the Stratix II device, Altera moved away from simple 4-input lookup tables, and used a more complex logic block which can implement functions of up to 7 inputs. We should expect to see a slow migration to more complex logic blocks, even in stand-alone FPGAs.

2) *Heterogeneous functions*: As devices are migrated to more advanced technologies, the number of transistors that can be devoted to the reconfigurable logic fabric increases. This provides new opportunities to embed complex non-programmable (or semi-programmable) functions, creating heterogeneous architectures with both general-purpose logic resources and fixed-function embedded blocks. Modern Xilinx parts have embedded 18 by 18 bit multipliers, while modern Altera parts have embedded DSP units which can perform a variety of multiply/accumulate functions. Again, we should expect to see a migration to more heterogeneous architectures in the near future.

3) *Soft cores*: The use of “soft” cores, particularly for instruction processors, is increasing. A “soft” core is one in which the vendor provides a synthesizable version of the function, and the user implements the function using the reconfigurable fabric. Although this is less area- and speed-efficient than a hard embedded core, the flexibility and the ease of integrating these soft cores makes them attractive. The extra overhead becomes less of a hindrance as the number of transistors devoted to the reconfigurable fabric increases. Altera and Xilinx both provide numerous soft cores, including soft instruction processors such as NIOS [7] and Microblaze [190]. Soft instruction processors have also been developed by a number of researchers, ranging from customisable JVM and MIPS processors [144] to ones specialised for machine learning [54] and data encryption [97].

IV. DESIGN METHODS

Hardware compilers for high-level descriptions are increasingly recognised to be the key to reducing the productivity gap for advanced circuit development in general, and for

reconfigurable designs in particular. This section looks at high-level design methods from two perspectives: special-purpose design and general-purpose design. Low-level design methods and tools, covering topics such as technology mapping, floor-planning, and place and route, are beyond the scope of this paper – interested readers are referred to [35].

A. General-purpose design

This section describes design methods and tools based on a general-purpose programming language such as C, possibly adapted to facilitate hardware development.

A number of compilers from C to hardware have been developed. Some of the significant ones are reviewed here. These range from compilers which only target hardware, to those which target complete hardware/software systems; some would also perform partitioning between hardware and software themselves.

We can classify different design methods into two approaches: the annotation and constraint-driven approach, and the source-directed compilation approach. The first approach preserves the source programs in C or C++ as much as possible and makes use of annotation and constraint files to drive the compilation process. The second approach involves modifying the source language to enable the designer to specify, for instance, the amount of parallelism or the size of data variables.

1) *Annotation and constraint-driven approach*: The systems mentioned below employ annotations in the source-code and constraint files to control the optimisation process. Their strength is that usually only minor changes are needed to produce a compilable program from a software description – no extensive re-structuring is required. Three representative methods are Streams-C [61], SPARK [66] and SPC [177].

SPC [177] combines vectorisation, loop transformations and retiming with an automatic memory allocation technique to improve design performance. The design flow associated with this method is shown in Figure 9. The purpose is to accelerate C loop nests with data dependency restrictions, compiling them into pipelines. Based on the SUIF framework [184], this approach extends work on loop transformations from software compilation field to hardware, and can take advantage of runtime reconfiguration and memory access optimisation. Similar methods have been advocated by other researchers [72], [142].

Streams-C [61] takes a design description in the form of an ANSI C program as input, and generates synthesizable VHDL. Streams-C exploits coarse-grained parallelism in stream-based computations, while low-level optimisations such as pipelining are performed automatically by the compiler.

SPARK [66] is a high-level synthesis framework targeting multimedia and image processing applications. It compiles behavioral ANSI-C code with the following steps: (a) a list scheduling pass based on speculative code motions and loop transformations, (b) a resource binding pass with minimisation of interconnect, (c) generation of a finite state machine controller for the scheduled datapath, (d) a code generation pass producing synthesizable register-transfer level VHDL.

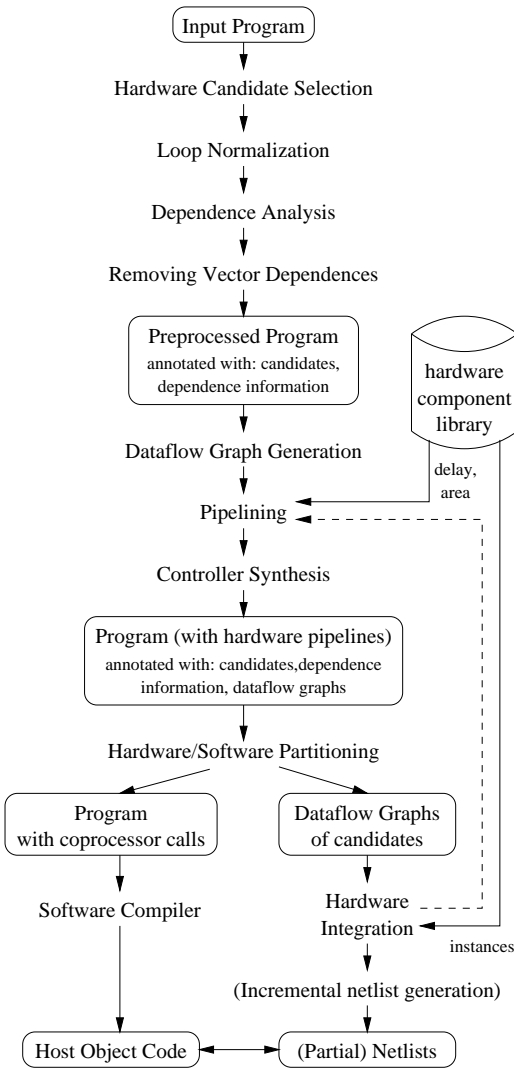


Fig. 9. Design flow for pipeline vectorisation [177].

The VHDL can then be synthesized using logic synthesis tools into a reconfigurable design.

2) *Source-directed compilation approach*: A different approach involves adapting the source language to enable explicit description of parallelism, communication and other customisable hardware resources such as variable size. Examples of design methods following this approach include ASC [108], Handel-C [29], Haydn-C [46], Bach-C [191], Hardware Promela [178].

ASC [108] adopts C++ custom types and operators to provide a programming interface on the algorithm level, the architecture-level, the arithmetic-level and the gate-level. As a unique feature, all levels of abstraction are accessible from C++. This enables the user to program on the desired level for each part of the application. Semi-automated design space exploration further increases design productivity, while supporting the optimisation process on all available levels of abstraction. Object-oriented design enables efficient code-

reuse, and includes an integrated arithmetic unit generation library, PAM-Blox II [107]. A floating-point library [98] with corresponding types and operators provide the ASC user with over 200 different floating point units, each with custom bitwidths for mantissa and exponent.

Handel-C [29] is based on ANSI-C extended to support flexible width variables, signals, parallel blocks, bit-manipulation operations, and channel communication. A distinctive feature is that timing of the compiled circuit is fixed at one cycle per C statement. This makes it easy for programmers to know in which cycle a statement will be executed at the expense of greatly reducing the scope for optimisation by restructuring. It gives application developers the ability to schedule hardware resources manually, and Handel-C tools generate the resulting designs automatically.

Handel-C compiles to a *one-hot* state machine, where each assignment of the program maps to exactly one control flip-flop in the state machine. These control flip-flops capture the flow of control in the program: if the control flip-flop corresponding to a particular statement is in a true state, then control has passed to that statement, and the circuitry compiled for that statement is activated. These control flip-flops essentially provide a token-passing compilation scheme, where each statement is only active when it has the token, which happens when the corresponding control flip-flop is activated. When the statement has finished execution, it passes the token to the next statement after it in the program. The compiler generates circuitry to control the flow of the token, implementing the control structures for Handel-C. Handel-C's token-passing scheme is based on work by Page and Luk in compiling Occam onto FPGAs [124].

Haydn-C [46] is a language extending Handel-C for component-based design. Like Handel-C, it supports description of parallel blocks, bit-manipulation operations, and channel communication. However, there is no distinction between using components written in Haydn-C and external components defined in a component library. Haydn-C also supports the entity construct to specify component templates. These templates can be optionally parametrised to generate components with different behaviours and/or structures based on a single description. The principal innovation of Haydn-C is a framework of optional annotations to enable users to describe design constraints, and to direct source-level transformations such as scheduling and resource allocation. There are automated transformations so that a single high-level design can be used to produce many implementations with different design trade-offs. The effectiveness of this approach has been evaluated using various case studies, including FIR filters, fractal generators, and morphological operators. For instance, the fastest morphological erosion design is 129 times faster and 3.4 times larger than the smallest design.

Bach-C [191] is similar to Handel-C but has an untimed semantics, only synchronising between parallel threads on synchronous communications between them, possibly giving greater scope for optimisation. It also allows asynchronous communications but otherwise resembles Handel-C, using the

same basic one-hot compilation scheme.

There has been work on compiling the Promela language to hardware. Promela, a non-deterministic language based on Dijkstra's guarded command language and CSP (Communicating Sequential Processes [75], also the basis of Occam and hence Handel-C). Promela (Process Metalanguage) is used to specify systems for the Spin formal software verification toolkit [78], [153]. Hardware Promela is a subset of Promela which drops the less-easily synthesisable parts of Promela, but retains the guarded commands and CSP-style communication [178]. The Hardware Promela language and compiler are reminiscent of Handel-C – the syntax is somewhat different but the semantics are similar, though the compilation scheme supports an exception mechanism in addition to the basic scheme used by Handel-C like compilers. The Promela compiler has a separate compilation and linking facility (as do more recent versions of Handel-C), allowing software-style linking of separately compiled (and clocked) hardware parts. A graphical development system [179] is also provided which allows software-style debugging. Multi-threaded C++ code can be generated from Promela specifications to run on the host for interfacing to the generated circuitry.

Table III summarises the various compilers discussed in this section, showing their approach, source and target languages, target architecture and some example applications. Note that the compilers discussed are not necessarily restricted to the architectures reported; some can usually be ported to a different architecture by using a different library of hardware primitives.

B. Special-purpose design

Within the wide variety of problems to which reconfigurable computing can be applied, there are many specific problem domains which deserve special consideration. The motivation is to exploit domain-specific properties: (a) to describe the computation, such as using MATLAB for digital signal processing, and (b) to optimise the implementation, such as using word-length optimisation techniques described later.

We shall begin with an overview of digital signal processing and relevant tools which target reconfigurable implementations. We then describe the word-length optimisation problem, the solution to which promises rich rewards; an example of such a solution will be covered. Finally we summarise other domain-specific design methods which have been proposed for video and image processing and networking.

1) *Digital signal processing*: One of the most successful applications for reconfigurable computing is Real-time Digital Signal Processing (DSP). This is illustrated by the inclusion of hardware support for DSP in the latest FPGA devices, such as the embedded DSP blocks in Altera Stratix II chips [6].

DSP problems tend to share the following properties: design latency is usually less of an issue than design throughput, algorithms tend to be numerically intensive but have very simple control structures, controlled numerical error is acceptable, and standard metrics, such as signal-to-noise ratio, exist for measuring numerical precision quality.

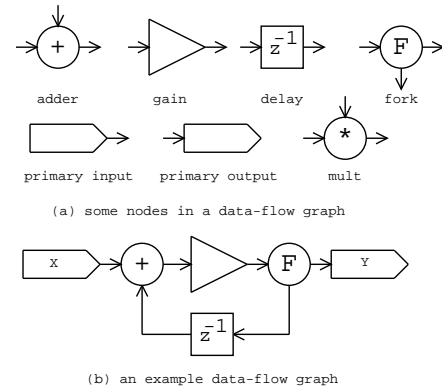


Fig. 10. The graphical representation of a data-flow graph.

DSP algorithm design is often initially performed directly in a graphical programming environment such as Mathworks' MATLAB Simulink [149]. Simulink is widely used within the DSP community, and has been recently incorporated into the Xilinx System Generator [79] and Altera DSP builder [4] design flows. Design approaches such as this are based on the idea of data-flow graphs (DFGs) [92]. A DFG $G(V, S)$ can be thought of as the formal representation of an algorithm. V is a set of graph nodes, each representing an atomic computation or input/output port, and S , a subset of $V \times V$, is a set of directed edges representing the data flow. In signal processing applications, elements of S are referred to as *signals*. Some standard node types and an example DFG are illustrated in Figure 10.

Tools working with this form of description vary in the level of user intervention required to specify the numerical properties of the implementation. For example, in the Xilinx System Generator flow [79], it is necessary to specify the number of bits used to represent each signal, the scaling of each signal (*namely* the binary point location), and whether to use saturating or wrap-around arithmetic [38].

Ideally, these implementation details could be automated. Beyond a standard DFG-based algorithm description, only one piece of information should be required: a lower-bound on the *output* signal to quantization noise acceptable to the user. Such a design tool would thus represent a truly 'behavioural' synthesis route, exposing to the DSP engineer only those aspects of design naturally expressed in the DSP application domain.

2) *The word-length optimization problem*: Unlike microprocessor-based implementations where the word-length is defined *a-priori* by the hard-wired architecture of the processor, reconfigurable computing based on FPGAs allows the size of each variable to be customised to produce the best trade-offs in numerical accuracy, design size, speed, and power consumption. The use of such custom data representation for optimising designs is one of the main strengths of reconfigurable computing.

Given this flexibility, it is desirable to automate the process of finding a good custom data representation. The most impor-

TABLE III
SUMMARY OF GENERAL-PURPOSE HARDWARE COMPILERS.

System	Approach	Source Language	Target Language	Target Architecture	Example Applications
Streams-C [61]	Annotation / Constraint-driven	C + library	RTL VHDL	Xilinx FPGA	Image contrast enhancement, Pulsar detection [57]
SPARK [66]	Annotation / Constraint-driven	C	RTL VHDL	LSI, Altera FPGAs	MPEG-1 predictor, Image tiling
SPC [177]	Annotation / Constraint-driven	C	EDIF	Xilinx FPGAs	String pattern matching, Image skeletonisation
ASC [108]	Source-directed Compilation	C++ using class library	EDIF	Xilinx FPGAs	Wavelet compression, Encryption
Handel-C [29]	Source-directed Compilation	Extended C	Structural VHDL, Verilog, EDIF	Actel, Altera, Xilinx FPGAs	Image processing, Polygon rendering [158]
Haydn-C [46]	Source-directed Compilation	Extended C	extended C (Handel-C)	Xilinx FPGAs	FIR filter, Image erosion
Bach-C [191]	Source-directed Compilation	Extended C	Behavioural and RTL VHDL	LSI FPGAs	Viterbi decoders, Image processing
Hardware Promela [178]	Source-directed Compilation	Extended Promela subset	Structural VHDL, EDIF	Altera FPGAs	Alternating-bit protocol

tant implementation decision to automate is the selection of an appropriate word-length and scaling for each signal [40] in a DSP system. Unlike microprocessor-based implementations, where the word-length is defined *a-priori* by the hard-wired architecture of the processor, reconfigurable computing allows the word-length of each signal to be customised to produce the best trade-offs in numerical accuracy, design size, speed, and power consumption. The use of custom data representation is one of the greatest strengths

It has been argued that, often, the most efficient hardware implementation of an algorithm is one in which a wide variety of finite precision representations of different sizes are used for different internal variables [36]. The accuracy observable at the outputs of a DSP system is a function of the word-lengths used to represent all intermediate variables in the algorithm. However, accuracy is less sensitive to some variables than to others, as is implementation area. It is demonstrated in [40] that by considering error and area information in a structured way using analytical and semi-analytical noise models, it is possible to achieve highly efficient DSP implementations.

In [43] it has been demonstrated that the problem of word-length optimization is NP-hard, even for systems with special mathematical properties that simplify the problem from a practical perspective [39]. There are, however, several published approaches to word-length optimization. These can be classified as heuristics offering an area / signal quality tradeoff [36], [87], [176], approaches that make some simplifying assumptions on error properties [26], [87], or optimal approaches that can be applied to algorithms with particular mathematical properties [37].

Some published approaches to the word-length optimization problem use an analytic approach to scaling and/or error estimation [119], [155], [176], some use simulation [26], [87], and some use a hybrid of the two [33]. The advantage of analytic techniques is that they do not require representative simulation stimulus, and can be faster, however they tend to be more pessimistic. There is little analytical work on supporting

data-flow graphs containing cycles, although in [155] finite loop bounds are supported, while [39] supports cyclic data-flow when the nodes are of a restricted set of types, extended to semi-analytic technique with fewer restrictions in [42].

Some published approaches use worst-case instantaneous error as a measure of signal quality [26], [119], [176] whereas some use signal-to-noise ratio [36], [87].

The remainder of this section reviews in some detail particular research approaches in the field.

The Bitwise Project [155] proposes propagation of integer variable ranges backwards and forwards through data-flow graphs. The focus is on removing unwanted most-significant bits (MSBs). Results from integration in a synthesis flow indicate that area savings of between 15% and 86% combined with speed increases of up to 65% can be achieved compared to using 32-bit integers for all variables.

The MATCH Project [119] also uses range propagation through data-flow graphs, except variables with a fractional component are allowed. All signals in the model of [119] must have equal fractional precision; the authors propose an analytic worst-case error model in order to estimate the required number of fractional bits. Area reductions of 80% combined with speed increases of 20% are reported when compared to a uniform 32-bit representation.

Wadekar and Parker [176] have also proposed a methodology for word-length optimization. Like [119], this technique also allows controlled worst-case error at system outputs, however each intermediate variable is allowed to take a word-length appropriate to the sensitivity of the output errors to quantization errors on that particular variable. Results indicate area reductions of between 15% and 40% over the optimum uniform word-length implementation.

Kum and Sung [87] and Cantin *et al.* [26] have proposed several word-length optimization techniques to trade-off system area against system error. These techniques are heuristics based on bit-true simulation of the design under various internal word-lengths.

In Bitsize [1], [2], Abdul Gaffar et al. propose a hybrid method based on the mathematical technique known as automatic differentiation to perform bitwidth optimisation. In this technique, the gradients of outputs with respect to the internal variables are calculated and then used to determine the sensitivities of the outputs to the precision of the internal variables. The results show that it is possible to achieve an area reduction of 20% for floating-point designs, and 30% for fixed-point designs, when given an output error specification of 0.75% against a reference design.

A useful survey of algorithmic procedures for word-length determination has been provided by Cantin *et al.* [27]. In this work, existing heuristics are classified under various categories. However the ‘exhaustive’ and ‘branch-and-bound’ procedures described in [27] do not necessarily capture the optimum solution to the word-length determination problem, due to non-convexity in the constraint space: it is actually possible to have a *lower* error at a system output by reducing the word-length at an internal node [41]. Such an effect is modeled in the MILP approach proposed in [37].

A comparative summary of existing optimization systems is provided in Table IV. Each system is classified according to the several defining features described below.

- Is the word-length and scaling selection performed through analytic or simulation-based means?
- Can the system support algorithms exhibiting cyclic data flow? (such as infinite impulse response filters).
- What mechanisms are supported for Most Significant Bit (MSB) optimizations? (Such as ignoring MSBs that are known to contain no useful information, a technique determined by the scaling approach used).
- What mechanisms are supported for Least Significant Bit (LSB) optimizations? These involve the monitoring of word-length growth. In addition, for those systems which support error-tradeoffs, further optimizations include the quantization (truncation or rounding) of unwanted LSBs.
- Does the system allow the user to trade-off numerical accuracy for a more efficient implementation?

3) *An example optimization flow:* One possible design flow for word-length optimization, used in the Right-Size system [42] is illustrated in Figure 11 for Xilinx FPGAs. The inputs to this system are a specification of the system behaviour (e.g. using Simulink), a specification of the acceptable signal-to-noise ratio at each output, and a set of representative input signals. From these inputs, the tool automatically generates a synthesizable structural description of the architecture and a bit-true behavioural VHDL testbench, together with a set of expected outputs for the provided set of representative inputs. Also generated is a makefile which can be used to automate the post-Right-Size synthesis process.

This design flow is illustrated for a simple DFG in Figure 12. The original DFG is illustrated in Figure 12(a).

The first step of the optimization is to determine the sensitivity of each output to quantization errors on each signal in the design; this will be illustrated for the particular signals marked (*) and (**) in this figure. Firstly nonlinear components are

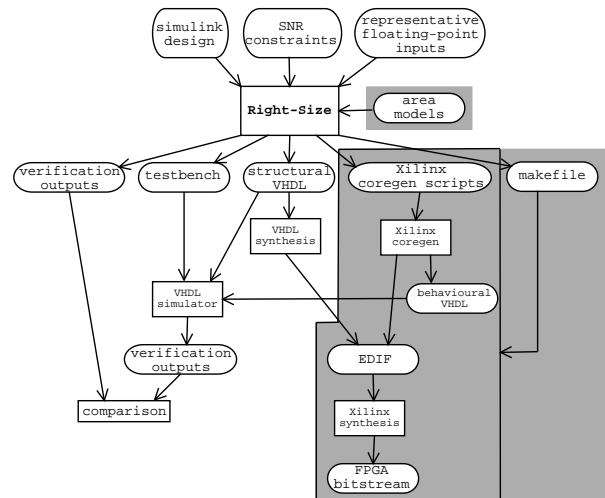


Fig. 11. Design flow for the Right-Size tool [42]. The shaded portions are FPGA vendor-specific.

identified by the design tool, and extra system outputs are inserted to monitor the derivative of the output of each these nonlinear components with respect of each of the inputs. After inserting derivative monitors for nonlinear components, the transformed DFG is shown in Figure 12(b).

The DFG may then be linearized, by automatically replacing each nonlinear component by its first-order Taylor series expansion. The linearized DFG is shown in Figure 12(c).

Once linearized, additional inputs can be inserted to model the quantization noise due to roundoff error; the two variants for the signals (*) and (**) are illustrated in Figures 12(d) and (e) respectively. To complete the first-order Taylor model, it is necessary to propagate zeros from the original primary system inputs. The corresponding simplified DFGs after zero-propagation are shown in Figures 12(f) and (g) respectively. These two systems may then be simulated, and the mean-square output can be taken as a measure of the sensitivity of the system output to quantization on the signal under consideration.

Once this process has been repeated for all signals, sensitivity measures can be combined with area models within a discrete optimization framework to decide an appropriate word-length for each signal [40].

Application of the proposed procedure to various adaptive filters implemented in a Xilinx Virtex FPGA has resulted in area reduction of up to 80%. Other encouraging results include power reduction of up to 98%, and speed-up of up to 36% over common alternative design methods without word-length optimisation.

4) *Other design methods:* Besides signal processing, video and image processing is another area that can benefit from special-purpose design methods. Three examples will be given to provide a flavour of this approach. First, the CHAMPION system [120] maps designs captured in the Cantata graphical programming environment to multiple reconfigurable computing platforms. Second, the IGOL framework [168] provides

TABLE IV
A COMPARISON OF WORDLENGTH AND SCALING OPTIMIZATION SYSTEMS AND METHODS.

System	Analytic / Simulation	Cyclic Data Flow?	MSB-optimization	LSB-optimizations	Error Trade off	Comments
Benedetti [13]	analytic	none	through interval arithmetic	through 'multi-interval' approach	no error	can be pessimistic
Stephenson [155], [156]	analytic	for finite loop bounds	through forward and backward range propagation	none	no error	less pessimistic than [13] due to backwards propagation
Nayak [119]	analytic	not supported for error analysis	through forward and backward range propagation	through fixing number of fractional bits for all variables	user-specified or inferred absolute bounds on error	fractional parts have equal wordlength
Wadekar [176]	analytic	none	through forward range propagation	through genetic algorithm search for suitable wordlengths	user-specified absolute bounds	uses Taylor series at limiting values to determine error propagation
Keding [84], [181]	hybrid	with user intervention	through user-annotations and forward range propagation	through user-annotations and forward wordlength propagation	not automated	possible truncation error
Cmar [33]	hybrid for scaling simulation for error	with user intervention only	through combined simulation and forward range propagation	wordlength bounded through hybrid fixed or floating simulation	not automated	less pessimistic than [13] due to input error propagation
Kum [86], [87], [161], [162]	simulation (hybrid for multiply-accumulate signals in [86], [87])	yes	through measurement of variable mean and standard deviation	through heuristics based on simulation results	user-specified bounds and metric	long simulation time possible
Constantinides [40]	analytic	yes	through tight analytic bounds on signal range and automatic design of saturation arithmetic	through heuristics based on an analytic noise model	user-specified bounds on noise power and spectrum	only applicable to linear time-invariant systems
Constantinides [42]	hybrid	yes	through simulation	through heuristics based on a hybrid noise model	user-specified bounds on noise power and spectrum	only applicable to differentiable non-linear systems
Abdul Gaffar [1], [2]	hybrid	with user intervention	through simulation based range propagation	through automatic differentiation based dynamic analysis	user-specified bounds and metric	covers both fixed-point and floating-point

a layered architecture for facilitating hardware plug-ins to be incorporated in various applications in the Microsoft Windows operating system, such as Premiere, Winamp, VirtualDub and DirectShow. Third, the SA-C compiler [19] maps a high-level single-assignment language specialised for image processing description into hardware, using various optimisation methods including loop unrolling, array value propagation, loop-carried array elimination, and multi-dimensional stripmining.

Recent work indicates that another application area that can benefit from special-purpose techniques is networking. Two examples will be given. First, a framework has been developed to enable description of designs in the network policy language Ponder [44], into reconfigurable hardware implementations [93]. Second, it is shown [85] how descriptions in the Click networking language can produce efficient reconfigurable designs.

C. Other design methods

In the following, we describe various design methods in brief.

1) *Run-time customisation*: Many aspects of run-time reconfiguration have been explored [35], including the use of directives in high-level descriptions [94]. Effective run-time

customisation hinges on appropriate design-time preparation for such customisation. To illustrate this point, consider a run-time customisable system that supports partial reconfiguration: one part of the system continues to be operational, while another part is being reconfigured. As FPGAs are getting larger, partial reconfiguration is becoming increasingly important as a means of reducing reconfiguration time. To support partial reconfiguration, appropriate circuits must be built at fabrication time as part of the FPGA fabric. Then at compile time, an initial configuration bitstream and incremental bitstreams have to be produced, together with run-time customisation facilities which can be executed, for instance, on a micro-processor serving as part of the run-time system [146]. Run-time customisation facilities can include support for condition monitoring, design optimisation and reconfiguration control.

Opportunities for run-time design optimisation include: (a) run-time constant propagation [48], which produces a smaller circuit with higher performance by treating run-time data as constant, and optimising them principally by boolean algebra; (b) library-based compilation – the DISC compiler [32] makes use of a library of precompiled logic modules which can be loaded into reconfigurable resources by the procedure call mechanism; (c) exploiting information

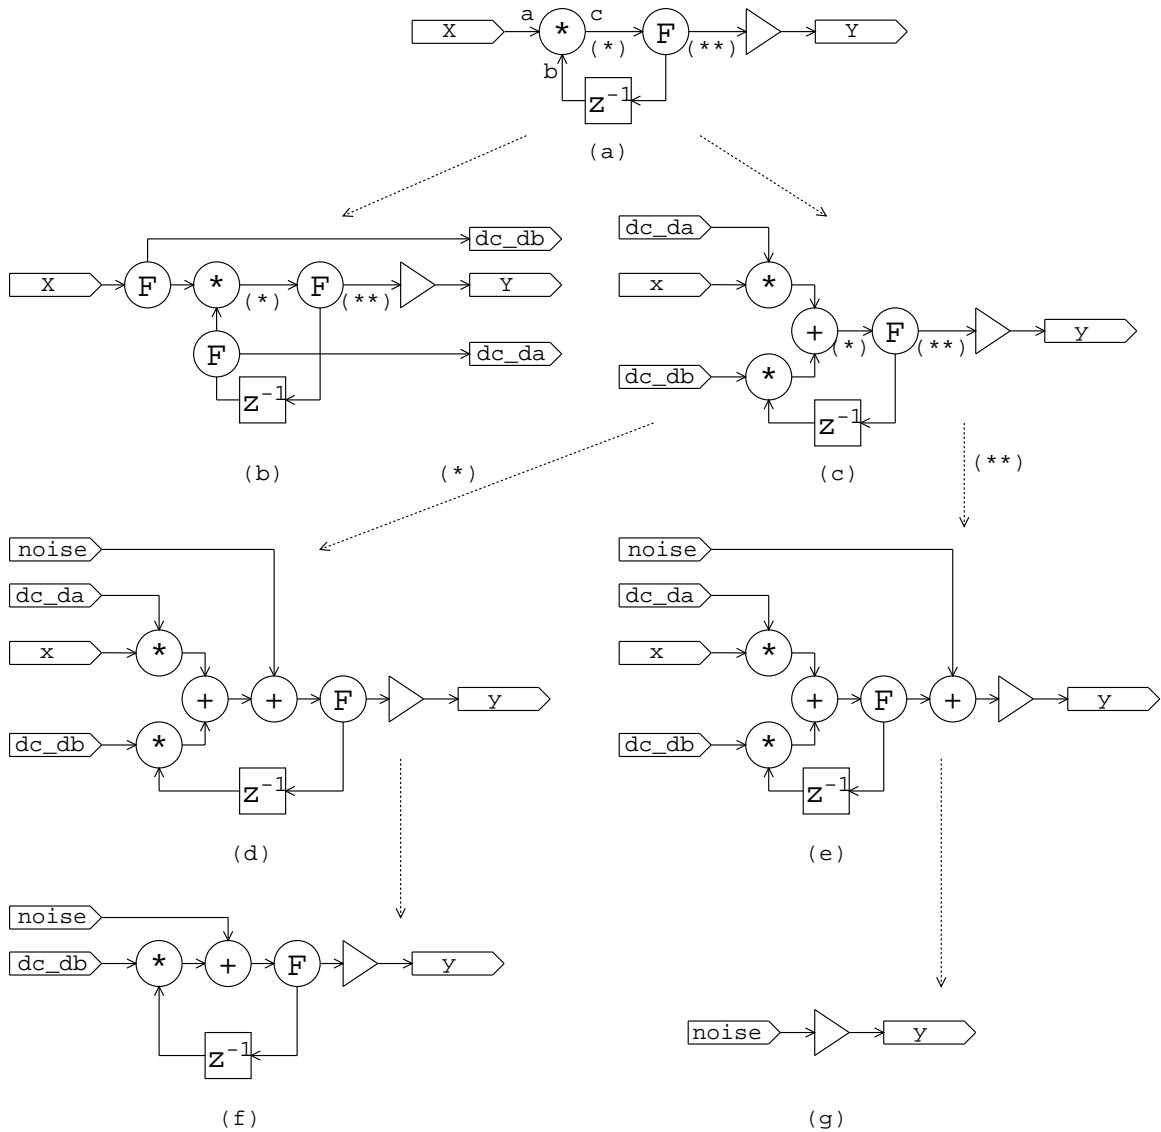


Fig. 12. An example of the perturbation analysis process [42]. The graphical representations are explained in Figure 10.

about program branch probabilities [160]; the idea is to promote utilisation by dedicating more resources to branches which execute more frequently. A hardware compiler has been developed to produce a collection of designs, each optimised for a particular branch probability; the best can be selected at run time by incorporating observed branch probability information from a queueing network performance model.

2) *Soft instruction processors*: FPGA technology can now support one or more instruction processors implemented using reconfigurable resources on a single chip; proprietary instruction processors, like MicroBlaze and Nios, are now available from FPGA vendors. Often such soft instruction processors support customisation of resources and custom instructions. Custom instructions have two main benefits. First, they reduce the time for instruction fetch and decode, provided that each custom instruction replaces one or more regular instructions.

Second, additional resources can be assigned to a custom instruction to improve performance. Bit-width optimisation, described in Section IV-B, can also be applied to customise instruction processors at compile time. A challenge of customising instruction processors is that the tools for producing and analysing instructions would also need to be customised. For instance, the *flexible instruction processor* framework [144] has been developed to automate the steps in customising an instruction processor and the corresponding tools. Other researchers have proposed similar approaches [83].

It is interesting to note that one can develop an instruction processor to support a declarative language. For instance, a scalable architecture [54], consisting of multiple processors based on the Warren Abstract Machine, has been developed to support the execution of the Progol system [114] which is based on the declarative language Prolog. Its effectiveness has

been demonstrated using the mutagenesis data set containing 12000 facts about chemical compounds.

3) *Multi-FPGA compilation*: Peterson et al. have developed a C compiler which compiles to multi-FPGA systems [128]. The available FPGAs and other units are specified in a library file, allowing portability. The compiler can generate designs using speculative and lazy execution to improve performance and ultimately they aim to partition a single program between host and reconfigurable resource (hardware/software codesign). Duncan et al. have developed a system with similar capabilities [50]. This is also retargetable, using hierarchical architecture descriptions. It synthesises a VLIW architecture that can be partitioned across multiple FPGAs. Both methods can split designs across several FPGAs, and are retargetable via hardware description libraries. Other C-like languages that have been developed include MoPL-3, a large C extension supporting data procedural compilation for the Xputer architecture which comprises an array of reconfigurable ALUs [9], and spC, a systolic parallel C variant for the Enable++ board [77].

4) *Hardware/software codesign*: Several research groups have studied the problem of compiling C code to both hardware and software. The work of the Garp compiler group [25] intends to accelerate plain C, with no annotations to help the compiler, making it more widely applicable. The work targets one architecture only: the Garp chip, which integrates a RISC core and reconfigurable logic. This compiler also uses the SUIF framework. The compiler uses a technique first developed for VLIW architectures called hyperblock scheduling, which optimises for instruction-level parallelism across several common paths, at the expense of rarer paths. Infeasible or rare paths are implemented on the processor with the more common, easily parallelisable paths synthesised into logic for the reconfigurable resource. Another platform-specific compiler is the NAPA C compiler for the NAPA architecture [60] – another integration of a RISC processor with an FPGA-like logic array. SUIF was again used to develop the compiler. This compiler can also work on plain C code but the programmer can add C pragmas to indicate large-scale parallelism and the bit-widths of variables to the code. The compiler can synthesise pipelines from loops, doing the necessary analysis to detect that this is allowable.

5) *Annotation-free compilation*: Some researchers aim to compile a sequential program, without any annotations, into efficient hardware design. This requires analysis of the source program to extract parallelism for an efficient result, which is necessary if compilation from languages such as C is to compete with traditional methods for designing hardware. One example is the work of Babb et al. [11], which is targeting custom, fixed-logic implementation while also applicable to reconfigurable hardware. The compiler uses the SUIF infrastructure to do several analyses to find what computations affect exactly what data, as far as possible. A tiled architecture is synthesised, where all computation is kept as local as possible to one tile. More recently, Ziegler et al. [192] have used loop transformations in mapping loop nests onto a pipeline spanning several FPGAs. A further effort is given by the Garp

project [25].

D. Emerging directions

1) *Verification*: As designs are becoming more complex, techniques for verifying their correctness are becoming increasingly important. Four approaches are described below. First, the InterSim framework [135] provides a means of combining software simulation and hardware prototyping. Second, the Lava system [17] can convert designs into a form suitable for input to a model checker; a number of FPGA design libraries have been verified in this way [151]. Third, the Ruby language [65] supports correctness-preserving transformations, and a wide variety of hardware designs have been produced. Fourth, the Pebble [101] hardware design language has been specified in a formal way [105], so that provably-correct design tools can be developed.

2) *Customisable hardware compilation*: Recent work [171] explains how customisable frameworks for hardware compilation can enable rapid design exploration, and reusable and extensible hardware optimisation. It is shown how such a framework can be based on a parallel imperative language, which supports multiple levels of design abstraction, transformational development, optimisation by compiler passes, and metalanguage facilities. The approach has been used in producing designs for applications such as signal and image processing, with different trade-offs in performance and resource usage.

E. Design methods: main trends

We summarise the main trends in design methods for reconfigurable computing below.

1) *Special-purpose design*: As explained earlier, special-purpose design methods and tools enable both high-level design as well as domain-specific optimisation. Existing methods, such as those compiling MATLAB Simulink descriptions into reconfigurable computing implementations [1], [4], [42], [79], [123], allow application developers without electronic design experience to produce efficient hardware implementations quickly and effectively. This is an area that would assume further importance in future.

2) *Low-power design*: Several hardware compilers have demonstrated their capabilities in minimising power consumption of their generated designs. Examples include special-purpose design methods such as Right-Size [42] and Py-Gen [123], and general-purpose methods that target loops for configurable hardware implementation [157]. These design methods, when combined with low-power architectures [58] and power-aware low-level tools [88], can provide significant reduction in power consumption.

3) *High-level transformations*: Many hardware design methods [19], [66], [177] involve high-level transformations: loop unrolling, loop restructuring and static single assignment are three examples. The development of powerful transformations for design optimisation would continue for both special-purpose and general-purpose designs.

V. APPLICATIONS

This section provides a general overview of applications to which reconfigurable computing has been applied. Characteristics of such applications include:

- 1) significant amount of parallelism, which can be exploited by providing multiple custom data processing units on a reconfigurable computing platform;
- 2) non-standard operations not well supported by micro-processor instruction sets, which can be supported by custom data processors;
- 3) the need for upgrading to meet new standards, or for run-time customisation to adapt to operating conditions; such need can be met by exploiting reconfigurability.

Table V summaries such characteristics for various applications such as data encryption, video processing, network security, and image generation.

A. Encryption

Reconfigurable computing has shown to be well-suited to data encryption applications, for both private-key and public-key systems. There are many opportunities for parallelism in encryption applications; for instance messages are typically encrypted as a block, so successive blocks can be encrypted in parallel. Also many cryptographic operations require finite-field arithmetic, which is not supported by standard microprocessor instruction sets. Finally, it is possible to obtain a smaller and faster key-specific design by treating the key as constant and applying boolean optimisation; reconfiguration (see also Section IV-C.1) is then used to produce a new design if another key is required [96].

In the following, we would first describe reconfigurable computing approaches to private-key systems: DES and AES. We then review reconfigurable computing approaches to public-key systems: RSA and ECC.

1) *DES*: DES [117], which stands for Data Encryption Standard, is probably the most widely used private-key encryption method today. The basic computation involves a 56-bit private key and encrypts blocks of 64-bits independently. The algorithm contains 16 iterations, or rounds, of identical operations, performed with a set of 48-bit subkeys. An overview of the DES algorithm is shown in Figure 13. The encryption begins with an initial permutation, encrypts in 16 rounds with an inverse of the initial permutation. Within each round, the left and right 32-bits are swapped with specific functions.

A closer look at a specific round of operation is shown in Figure 14. The swapping involves permutations and S-box substitution which is implemented with 32 CLBs for 6-input lookup tables (LUTs) targeting Xilinx Virtex V150-6 FPGA [127]. Specific S-Boxes can be specified as a table with a 6-bit input and a 4-bit output, and can be implemented in a single CLB using the SRL16 shift register mode. This structure enables high degree of pipelining for the datapath.

Pipelining is commonly used in reconfigurable device to enhance performance. Both DES and AES have S-boxes for data substitution. For instance, a 3-stage pipeline design is

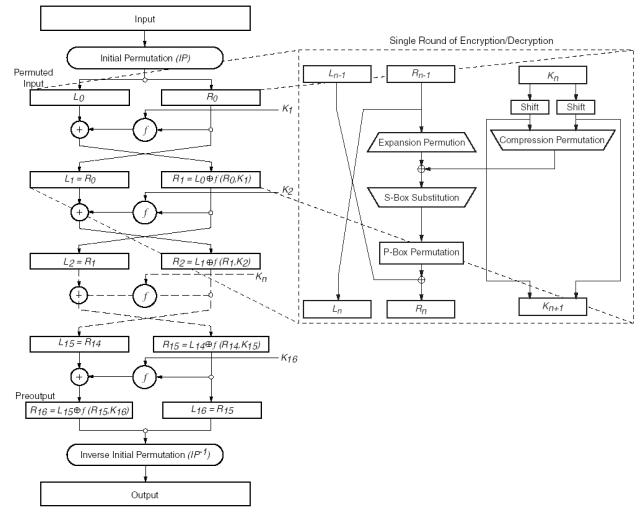


Fig. 13. DES Algorithm overview [126].

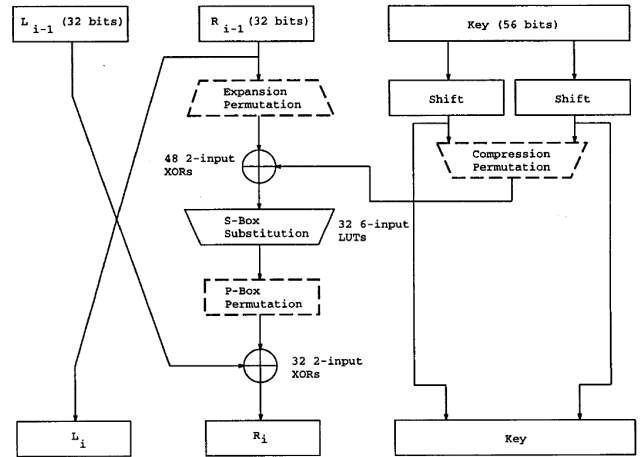


Fig. 14. A single DES round [127].

TABLE VI
FEATURES OF DIFFERENT DES HARDWARE DESIGNS.

Developer	Year	Platform	Clock (MHz)	Throughput (Mbits/s)	Design Capture
Leonard [96]	1997	Xilinx 4013-4	7	26	VHDL
Wilcox [180]	1999	4 Altera 10K100-3	20	1280	AHDL
Wilcox [180]	1999	Sandia ASIC	145	9280	VHDL
Patterson [127]	2000	Xilinx V150-6	168	10752	JBits
Pasham [126]	2001	Xilinx XC2V1000	237	15100	Verilog

depicted in Figure 15 using 5 flip-flops. We can make use of the parallelism by using 8 parallel S-boxes in the substitution. Run-time reconfigurable designs for both DES and AES can also be developed [48].

Table VI shows the features of some recent DES designs. It can be seen that two reconfigurable designs are faster than the ASIC design.

2) *AES*: AES [118], short for Advanced Encryption Standard which is also called Rijndael, has four different steps: byte substitution, shift row, mix column, and key addition (Fig-

TABLE V
SUMMARY OF SYSTEM ARCHITECTURES.

Application	Cryptography	Video	Networking
Parallelism	between message blocks, unroll loop	pixel-level (spatial) and frame-level (temporal)	between fields and between packets
Non-standard operations	data-substitution, finite-field	multi-pixel operations in parallel	customised content addressible memory
Exploitation of reconfigurability	key-specific designs	coefficient-specific designs	adapt to noise or security breach




Fig. 15. Single round data path [126].

ure 16). This block cipher standard includes three block sizes: 128 (AES-128), 192 (AES-192) and 256 (AES-256) bits. The whole block encryption is divided into several different rounds. The AES-128 standard consists of 10 rounds. The entire design can be fully pipelined (Figure 16) and mapped into a single FPGA achieving over 21Gbit/s throughput. In recent years, many high performance AES designs have been published, and most of them have explored specific architectures in current FPGAs, such as the use of block RAM and block multiplier. These designs [139], [154] can usually achieve over 10Gbit/s.

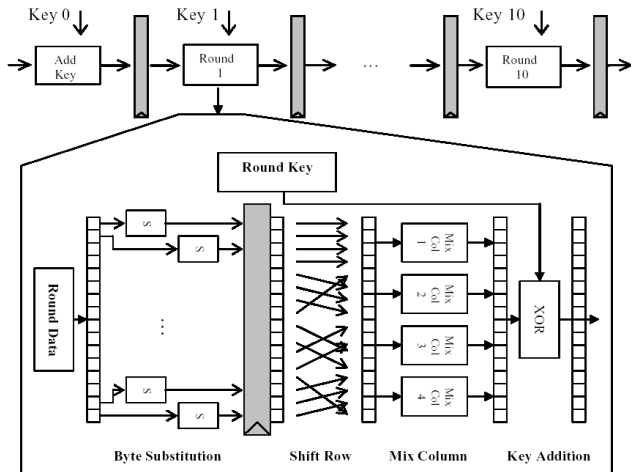


Fig. 16. Fully pipelined AES design [76].

3) *RSA*: RSA is probably the most widely used public-key encryption method today. The key operation in RSA [136] is the modular exponentiation, which can be divided into a series of modular multiplication. A common approach is to use the Montgomery method [111] to carry out the multiplication. Recent work tends to follow an instruction-processor design approach [104] and mixed public-key cryptosystem such as ECC and RSA. The Montgomery multiplier is also commonly adopted in ECC cryptosystems, using prime field representation [122] on a XCV1000E FPGA.

Since the Montgomery modular operation is not supported by current instruction sets, much research has focused on the design of the multiplier. A recent study proposes a semi-systolic structure for the modular exponentiation unit [163], which can be developed for reconfigurable device which makes use of the embedded 18 by 18-bit multipliers. This unit offers an attractive result: $0.66ms$ for a 1,024-bit RSA decryption with the Chinese Remainder Theorem. Moreover, users can easily replace the FPGAs with an existing faster chip without re-engineering the design.

4) *ECC*: Elliptic Curve Cryptography [69] is a public-key cryptosystem that has been shown to offer significantly higher security than the RSA method for the same key size. As an example, we describe a framework [166] for producing ECC hardware designs over finite field $GF(2^m)$, using the optimal normal basis for the representation of numbers. The field multiplier design is based on a parallel architecture containing m -bit serial multipliers. A design generator has been developed which can automatically produce a customised ECC hardware design that meets user-defined requirements. This method enables designers to rapidly explore and implement a design with the best trade-offs in speed, size and level of security. The resulting hardware implementations are among the fastest reported, and can often run several orders of magnitude faster than software implementations. For instance, given a key size of 270 bits, a point multiplication can be computed in 0.36 ms with a reconfigurable computing design implemented in an XC2V6000 FPGA at 66 MHz [166]. In contrast, an optimised software implementation requires 196.71 ms on a dual-Xeon computer at 2.6 GHz. Hence the reconfigurable computing design is more than 540 times faster, although its clock speed is almost 40 times slower than the Xeon processors.

Table VII shows the features of some recent ECC designs. As with DES, a reconfigurable computing design is faster than the ASIC design.

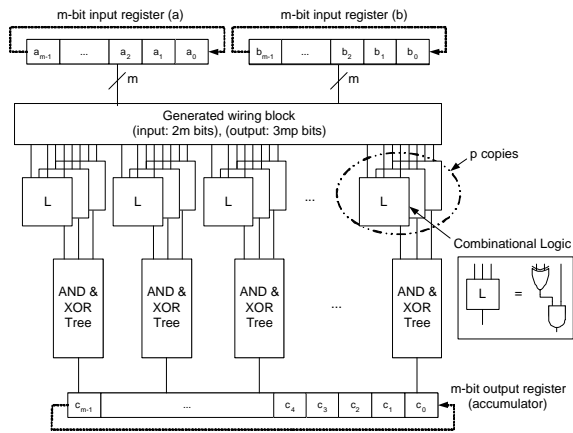


Fig. 17. Parametric finite field multiplier design [166].

TABLE VII

SUMMARY OF DIFFERENT ECC HARDWARE DESIGNS. PB STANDS FOR POLYNOMIAL BASIS AND ONB STANDS FOR OPTIMAL NORMAL BASIS.

Developer	Year	Platform	Basis	m	Timing
Orlando [121]	2000	XCV400E	PB	167	0.21 ms
Leong [97]	2002	XCV1000	ONB	113	0.75 ms
Gura [67]	2002	XCV2000E	PB	163	0.14 ms
Satoh [140]	2003	0.13- μ m ASIC	PB	160	0.19 ms
Telle [166]	2004	XC2V6000	ONB	270	0.36 ms

B. Video image processing

Many industrial video image processing systems use a mixture of high performance computers and application-specific integrated circuits. Implementations based on digital signal processors such as the TriMedia [133] are often energy inefficient and may result in unacceptable dataflow bottlenecks, limiting their performance. On the other hand, ASIC-based solutions are not only expensive to develop, but are also inflexible. Reconfigurable computers combine the flexibility of a software approach found in a workstation, with the performance of a hardware solution found in an ASIC. Video Image processing algorithms are particularly suitable for implementation on such machines because of their inherent spatial and temporal parallelism [10]. Data-flow often exhibits stream-like behaviour [22] with mostly local communications between neighbouring processing elements. Therefore mapping such algorithms onto reconfigurable hardware is often relatively straight forward.

An interesting early example of using reconfigurable computer for real-time video processing is found in [130]. They report a very coarse grain application-specific Field-Programmable Operator Array (FPOA) where each FPOA includes two configurable data-path cells (CDPs) that implements 8-bit and 16-bit operators. A 3D array of FPOAs are connected together to form a Data-Flow Functional Computer which could implement many real-time video applications such as coloured object tracking, connected component labeling and non-linear filtering. Bergmann [14] and Woods [187]

independently report the use of SRAM-based FPGAs as a co-processor resource to a conventional CPU, and demonstrate significant speedup over a workstation when implementing 2-D Discrete Cosine Transforms of 8 by 8 pixels blocks within a 512 by 512 pixel image. Reconfigurable computer can also be found onboard remote sensing satellites to perform real-time cloud detection [182], reducing the delay between image capture, analysis and action, and also reducing onboard storage and downlink capability requirements.

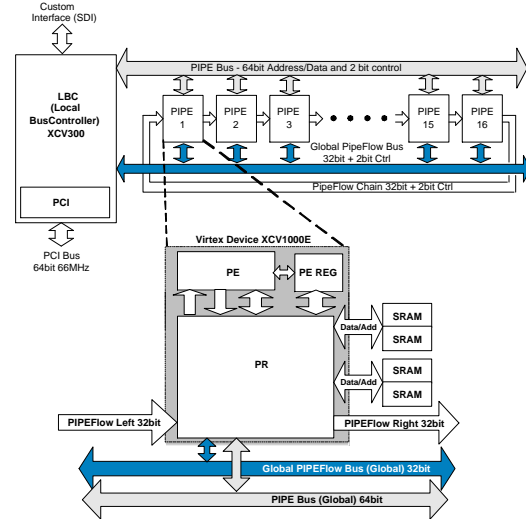


Fig. 18. The UltraSONIC reconfigurable computing system [74]. PE denotes PIPE Engine and PR denotes PIPE Router.

An example of a powerful reconfigurable computer system that has successfully moved from the research laboratory to industry is based around the SONIC architecture [73]. The UltraSONIC [74] is a reconfigurable computing system designed to cope with the computational power and the high data throughput demanded by real-time video applications at broadcast quality (see Figure 18). The system consists of Plug-In Processing Elements (PIPEs) interconnected by local and global buses. The architecture fully exploits the spatial and the temporal parallelism found in video processing algorithms. It also facilitates design reuse and supports the software plug-in methodology. The basic PIPE consists of three parts (Figure 18): PIPE engine, PIPE router, and PIPE memory. The PIPE engine handles computation while the PIPE router is responsible for image data movement and formatting. The PIPE memory provides local buffering of video data therefore reducing bus traffic. UltraSONIC uses a Virtex XCV1000E device implementing both the PIPE engine and the PIPE router. The PIPE memory is implemented with synchronous SRAM, which allows fast and easy memory access. The communication between the host system (which provides the software resource) and the UltraSONIC board is via a 64-bit PCI bus running at 66MHz. The SONIC architecture is currently being adapted as a system-on-a-chip solution for embedded video applications [143].

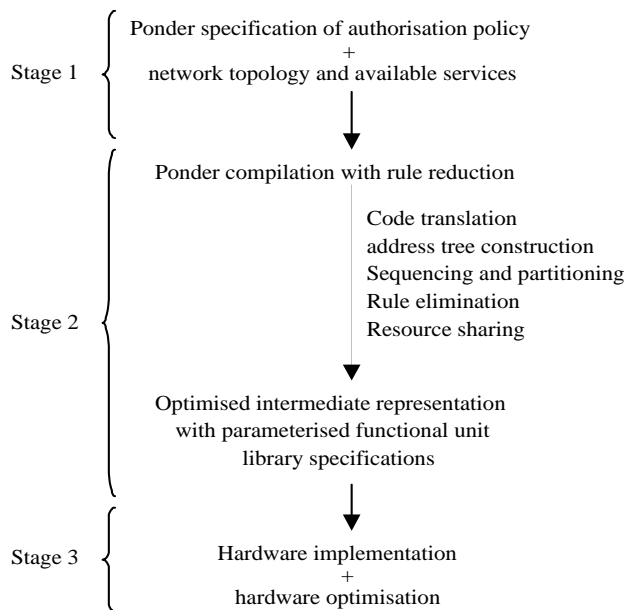


Fig. 19. Development stages for firewall processors [93].

C. Network security

In the following we describe two applications of reconfigurable computing to network security: firewall processors, and network intrusion detection.

1) *Firewall processors*: This study concerns application customisation for producing network firewall designs [93], which involves three stages (Figure 19). In the first stage, an authorisation policy is specified in the Ponder language [44] together with the information about the organisation’s network topology and services. In the second stage, this specification is translated into a platform-independent intermediate representation. A series of customisations, including construction of IP address trees, sequencing, rule elimination, and resource sharing, are then performed to optimise the representation. In the third stage, this optimised representation is used to target a specific hardware platform. Hardware packet filter designs are captured using the Handel-C language (see Section IV-A). Optimisations can be applied to improve performance or resource usage. One advantage of this approach is that multiple levels of customisation, based on various platform-independent and platform-specific criteria, can improve performance, portability and design reuse, since platform-independent customisations are applicable to a large variety of designs.

Other reconfigurable computing designs for firewalls are mainly based on content addressable memories [49], [80].

2) *Network intrusion detection*: Network Intrusion Detection Systems, or NIDS for short, is an increasingly common approach to securing a network from intruders who constantly adjust their approach to keep ahead of advances in security. Once common packet-header-based security systems have now been found to be lacking where contemporary network security is concerned. Conventional NIDS search through a

packet payload, seeking out known attack signatures. The time taken to search through each packet for a matching string is computationally expensive, and can adversely affect the network speed. Most software-based NIDS suffer from this drawback, being able to process traffic at only hundreds of Mbps, and this has led to popular interest in research into hardware-based NIDS as a solution to software drawbacks, and providing systems capable of handling traffic at Gbps rates.

Hardware-based NIDS are able to achieve relatively high speed by making use of the true parallelism offered by hardware. Their review focuses on the varying techniques used to match strings in hardware-based NIDS. Using hardware, however, limits in terms of the number of strings which can be stored in the system.

Table VIII summarises various NIDS designs. It is clear from this table that the most efficient implementations in terms of area resource usage make use of regular expressions, while the most efficient implementations in terms of speed make use of CAM based techniques. However, to determine how effective the system is as a whole, we use a metric of the throughput per LUTs/byte ratio (efficiency), and in this case the larger the ratio, the more effective the system. By using this ratio, we determine that the most balanced system is the one developed by Clark and Schimmel [31], as that design is relatively fast compared to the size of the design.

D. Image generation

Realistic image generation is known to require significant computational power. In the following we review several efforts in this direction.

Dao et. al. have applied reconfigurable computing to ray casting: ray tracing without reflections and refractions [45]. This is a system for the visualisation of volume data, such as that produced by measurements in medicine or fluid dynamics. The system uses parallel projection of discrete rays from a screen to the data to be visualised, using an FPGA to accelerate this process.

An alternative method is proposed by Todman and Luk [169]. Their method is based on breadth-first strategies, which have previously been used by Hanrahan [70] for exploiting the coherence between similar rays, and by Muller et al [113] and Nakamura and Ohno [115] to reduce disk thrashing when accessing very large scenes and for ray-tracing on vector and parallel processors [125], [129]. This method has two main features. First, it implements ray tracing with reflection and refraction. The motivation for this breadth-first transformation comes from reusing the results of previous rays, which is something Dao et. al. do not consider. Second, Dao et. al. consider volume visualisation, while Todman and Luk consider discrete objects.

Recently Fender and Rose [53] also describe ray tracing using a reconfigurable computing system that can outperform a 2.4GHz Pentium processor by more than 30 times, and has the potential of accelerating designs up to two orders of magnitude. Their approach is based on barycentric coordinates with a hierarchical data structure for fixed-point numbers.

TABLE VIII
A COMPARISON OF METHODS FOR INTRUSION DETECTION.

Developer	Technique	Device	Throughput (Gbps)	# Bytes	# LUTs	LUTs/Byte (LpB)	Efficiency Gbps/(LpB)
Clark et al. [31]	Regular Expression	Virtex2 6000	1.5	17,537	18,924	1.08	1.39
Franklin et al. [55]	Regular Expression	VirtexE 2000	0.4	8,003	20,618	2.58	0.16
Gokhale et al. [62]	CAM	Virtex 1000	2	640	9,722	15.2	0.13
Sourdis et al. [152]	CAM	Virtex2 6000	8.064	2,457	47,686	19.4	0.42

Styles and Luk [159] have developed a ray-triangle intersector on an FPGA to accelerate the radiosity algorithm, which is related to ray-tracing. The key component in this design is the hardware intersection tester, which consists of a long hardware pipeline containing a single branch, 28 multiply and 12 add/subtract operations.

Another description of FPGAs for non-standard operations in image generation is the work of Cantle *et al.* [28], who implement particle systems on the Nallatech DIME modular PCI platform [116]. They implement particle systems – computer graphics techniques producing images by simulation of interactions between thousands of particles and their environment – on a Xilinx Virtex FPGA on a DIME module. Their parameterised particle systems can be used to animate smoke, fire, exhaust plumes and related effects. Utilising the massive parallelism available in the FPGA, they animate thousands of particles at 100 Hz frame rate, for a 1024 by 1024 image with four times oversampling. The parameterisation means the system can be adapted to non-standard uses, for instance infra-red simulations for defence applications.

E. Applications: main trends

As their capacity, capability and performance continue to rise while their price continues to fall, reconfigurable computing devices such as FPGAs will be used for an increasing range of applications. Such devices will strengthen their position in existing applications, such as signal and media processing, where they already have a healthy share of the market. At the same time, they would expand into new areas where their improved capability has potential to bring significant benefits. The three main trends that we can report are discussed below.

1) *Media processing*: Reconfigurable computing has already shown much benefit for media processing. It appears to continue to be applied successfully to existing and new areas of media processing, including video, image and audio data transformation, compression and transmission. Recent studies indicate that, for instance, an FPGA-augmented TriMedia device can perform MPEG2-compliant pel reconstruction with an average speedup of 1.4 times over the standard TriMedia device [148]. Another example involves performing LDPC encoding: an implementation of 16 instances of a hardware LDPC encoder on the same Xilinx Virtex-II XC2V4000-6 FPGA at 82MHz is capable of 410 million codeword bits per second, 80 times faster than an Intel Pentium-IV 2.4GHz PC [91].

2) *Numerical computations*: A few years ago a view emerges that reconfigurable computing may not provide a

competitive solution for numerical computations which involve a significant amount of floating-point operations. Recent progress [98], [173] suggest that reconfigurable computing has great potential in this area, particularly if operations can be avoided or adapted appropriately. For instance, if it is known that overflow and underflow would not occur, then there is no need to include the resources to deal with them, which can result in a smaller and faster design. Reconfigurable computing techniques are beginning to apply to areas that involve floating-point operations, and early results from medical imaging [81] appear promising.

3) *Embedded Applications*: As the capacity of reconfigurable devices increases, more functions can be supported by a single device which becomes a system-on-chip. An example that we explained before is the SONIC architecture (Section V-B). A SONIC PIPE takes two Altera FLEX 10K parts to implement in 1999 [73]. Now five years later, a single Xilinx Virtex-II Pro device can accommodate multiple PIPES [143]. This additional capacity, when combined with low-power architecture and design methods, can lead to new applications of reconfigurable computing in areas such as mobile systems or pervasive computing, where small size, high performance and low power and energy consumption are key. Indeed, a recent study [157] reports that moving critical software loops to reconfigurable hardware results in average energy savings of 35% to 70% with an average speedup of 3 to 7 times, depending on the particular device used.

VI. SUMMARY

This paper surveys three aspects of reconfigurable computing: architectures, design methods, and applications. The main trends in architectures are coarse-grained fabrics, heterogeneous functions, and soft cores. The main trends in design methods are special-purpose design methods, low-power techniques, and high-level transformations. The main trends in applications are media processing, numerical computations, and embedded applications. We wonder what a survey paper on reconfigurable computing, written in 2015, will cover?

ACKNOWLEDGMENTS

Our thanks to Ray Cheung and Sherif Yusuf for their support in preparing this paper. The support of Celoxica, Xilinx and UK EPSRC (Grant number GR/R 31409, GR/R 55931, GR/N 66599) is gratefully acknowledged.

REFERENCES

- [1] A. Abdul Gaffar, O. Mencer, W. Luk, P.Y.K. Cheung and N. Shirazi, "Floating-point bitwidth analysis via automatic differentiation", *Proc. Int. Conf. on Field-Programmable Technology*, IEEE, 2002.

- [2] A. Abdul Gaffar, O. Mencer, W. Luk, and P.Y.K. Cheung, "Unifying bit-width optimisation for fixed-point and floating-point designs", *Proc. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 2004.
- [3] Actel Corp., *ProASIC Plus Family Flash FPGAs*, v3.5, April 2004.
- [4] Altera Corp., *DSP Builder User Guide*, Version 2.1.3 rev.1, July 2003.
- [5] Altera Corp., *Excalibur Device Overview*, May 2002.
- [6] Altera Corp., *Stratix II Device Handbook*, February 2004.
- [7] Altera Corp., *Nios II Processor Reference Handbook*, May 2004.
- [8] Annapolis Microsystems, Inc., *Wildfire Reference Manual*, 1998.
- [9] A. Ast, J. Becker, R. Hartenstein, R. Kress, H. Reinig and K. Schmidt, "Data-procedural languages for FPL-based machines", *Field-Programmable Logic and Applications*, LNCS 849, Springer, 1994.
- [10] P.M. Athanas and A.L. Abbott, "Real-time image processing on a custom computing platform", *IEEE Computer*, Vol. 28, pp. 16–24, 1995.
- [11] J. Babb, M. Reinard, C. Andras Moritz, W. Lee, M. Frank, R. Barwa and S. Amarasinghe, "Parallelizing applications into silicon", *Proc. Symp. on FPGAs for Custom Computing Machines*, IEEE Computer Society Press, 1999.
- [12] J. Becker and M. Glesner, "A parallel dynamically reconfigurable architecture designed for flexible application-tailored hardware/software systems in future mobile communication", *The Journal of Supercomputing*, Vol. 19, No. 1, 2001, pp. 105–127.
- [13] A. Benedetti and B. Perona, "Bit-width optimization for configurable DSP's by multi-interval analysis", *Proc. 34th Asilomar Conference on Signals, Systems and Computers*, 2000.
- [14] N.W. Bergmann and Y.Y. Chung, "Video compression with custom computers", *IEEE Transactions on Consumer Electronics*, Vol. 43, Aug. 1997, pp. 925–933.
- [15] G. Berry and M. Kishinevsky, "Hardware Esterel Language Extension Proposal", http://www.esterel-technologies.com/v2/solutions/iso_album/hardest.pdf.
- [16] V. Betz, J. Rose and A. Marquardt, *Architecture and CAD for Deep-Submicron FPGAs*, Kluwer Academic Publishers, February 1999.
- [17] P. Bjesse, K. Claessen, M. Sheeran and S. Singh, "Lava: hardware design in Haskell", *Proc. ACM Int. Conf. on Functional Programming*, ACM Press, 1998.
- [18] T. Blum and C. Paar, "High-radix Montgomery modular exponentiation on reconfigurable hardware", *IEEE Transactions on Computers*, Vol. 50, No. 7, pp. 759–764, 2001.
- [19] W. Bohm, J. Hammes, B. Draper, M. Chawathe, C. Ross, R. Rinker and W. Najjar, "Mapping a single assignment programming language to reconfigurable systems", *The Journal of Supercomputing*, Vol. 21, 2002, pp. 117–130.
- [20] K. Bondalapati and V.K. Prasanna, "Dynamic precision management for loop computations on reconfigurable architectures", *Proc. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 1999.
- [21] K. Bondalapati and V.K. Prasanna, "Reconfigurable computing systems", *Proc. IEEE*, Vol. 90, No. 7, pp. 1201–1217, July 2002.
- [22] V.M. Bove, J. Watlington and J. Cheops, "A reconfigurable data-flow system for video processing", *IEEE Transactions on Circuits and Systems for Video Technology*, Vol. 5, 1995, pp. 140–149.
- [23] M. Butts, A. DeHon and S. Goldstein, "Molecular electronics: devices, systems and tools for gigagate, gigabit chips", *Proc. Int. Conference on Computer-Aided Design*, IEEE, 2002.
- [24] Cadence Design Systems Inc, *Palladium Datasheet*, 2004.
- [25] T. Callahan and J. Wawrzynak, "Instruction-Level parallelism for reconfigurable computing", *Field-Programmable Logic and Applications*, LNCS 1482, Springer, 1998.
- [26] M.-A. Cantin, Y. Savaria and P. Lavoie, "An automatic word length determination method", *Proc. IEEE International Symposium on Circuits and Systems*, p.V-53-V-56, 2001.
- [27] M.-A. Cantin, Y. Savaria and P. Lavoie, "A comparison of automatic word length optimization procedures", *Proc. IEEE Int. Symp. on Circuits and Systems*, 2002.
- [28] A.J. Cantle, M. Devlin, E. Lord and R. Chamberlain, "High frame rate, low latency hardware-in-the-loop image generation? An illustration of the Particle Method and DIME", *Proc. Aerosense Conference*, SPIE, 2004.
- [29] Celoxica, *Handel-C Language Reference Manual for DK2.0*, Document RM-1003-4.0, 2003.
- [30] Celoxica, *RC2000 Development and evaluation board data sheet*.
- [31] C.R. Clark and D.E. Schimmel, "A pattern-matching co-processor for network intrusion detection systems", *Proc. Int. Conf. on Field Programmable Technology*, IEEE, 2003.
- [32] D. Clark and B. Hutchings, "The DISC programming environment", *Proc. Symp. on FPGAs for Custom Computing Machines*, IEEE Computer Society Press, 1996.
- [33] R. Cmar, L. Rijnders, P. Schaumont, S. Vernalde and I. Bolsens, "A methodology and design environment for DSP ASIC fixed point refinement", *Proc. Design Automation and Test in Europe*, 1999.
- [34] K. Compton and S. Hauck, "Totem: Custom reconfigurable array generation", *Proc. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 2001.
- [35] K. Compton and S. Hauck, "Reconfigurable computing: a survey of systems and software", *ACM Computing Surveys*, Vol. 34, No. 2, pp. 171–210, June 2002.
- [36] G.A. Constantinides, P.Y.K. Cheung and W. Luk, "The multiple wordlength paradigm", *Proc. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 2001.
- [37] G.A. Constantinides, P.Y.K. Cheung and W. Luk, "Optimum wordlength allocation", *Proc. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 2002.
- [38] G.A. Constantinides, P.Y.K. Cheung and W. Luk, "Optimum and heuristic synthesis of multiple wordlength architectures", *IEEE Trans. on Computer Aided Design of Integrated Circuits and Systems*, Vol. 22, No. 10, pp. 1432–1442, October 2003.
- [39] G.A. Constantinides, P.Y.K. Cheung and W. Luk, "Synthesis of saturation arithmetic architectures", *ACM Trans. on Design Automation of Electronic Systems*, Vol. 8, No. 3, July 2003.
- [40] G.A. Constantinides, P.Y.K. Cheung and W. Luk, *Synthesis and Optimization of DSP Algorithms*, Kluwer Academic, Dordrecht, 2004.
- [41] G.A. Constantinides, *High Level Synthesis and Word Length Optimization of Digital Signal Processing Systems*, PhD thesis, Imperial college London, 2001.
- [42] G.A. Constantinides, "Perturbation analysis for word-length optimization", *Proc. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 2003.
- [43] G.A. Constantinides and G.J. Woeginger, "The complexity of multiple wordlength assignment", *Applied Mathematics Letters*, Vol. 15, No. 2, pp. 137–140, February 2002.
- [44] N. Damianou, N. Dulay, E. Lupu and M. Sloman, "The Ponder policy specification language", *Proc. Workshop on Policies for Distributed Systems and Networks*, LNCS 1995, Springer, 2001.
- [45] M. Dao, T. Cook, D. Silver and P. D'Urbano, "Acceleration of template-based ray casting for volume visualization using FPGAs", in *Proc IEEE Symposium on Field-Programmable Custom Computing Machines*, 1995.
- [46] J.G. De Figueiredo Coutinho and W. Luk, "Source-directed transformations for hardware compilation", *Proc. Int. Conf. on Field-Programmable Technology*, IEEE, 2003.
- [47] A. DeHon and M.J. Wilson, "Nanowire-based sublithographic programmable logic arrays", *Proc. Int. Symp. on FPGAs*, ACM Press, 2004.
- [48] A. Derbyshire and W. Luk, "Compiling run-time parametrisable designs", *Proc. Int. Conf. on Field-Programmable Technology*, IEEE, 2002.
- [49] J. Ditmar, K. Torkelsson and A. Jantsch, "A dynamically reconfigurable FPGA-based content addressable memory for internet protocol characterization", *Field Programmable Logic and Applications*, LNCS 1896, Springer, 2000.
- [50] A. Duncan, D. Hendry and P. Gray, "An overview of the COBRA-ABS high-level synthesis system for multi-FPGA systems", *Proc. IEEE Symposium on FPGAs for Custom Computing Machines*, IEEE Computer Society Press, 1998.
- [51] C. Ebeling, D. Conquist and P. Franklin, "RaPiD – reconfigurable pipelined datapath", *Field-Programmable Logic and Applications*, LNCS 1142, Springer, 1996.
- [52] Elitex Corporation, *DFA 1000 Accelerator Datasheet*, 2003.
- [53] J. Fender and J. Rose, "A high-speed ray tracing engine built on a field-programmable system", *Proc. Int. conf. on Field-Programmable Technology*, IEEE, 2003.

- [54] A. Fidjeland, W. Luk and S. Muggleton, "Scalable acceleration of inductive logic programs", *Proc. Int. Conf. on Field-Programmable Technology*, IEEE, 2002.
- [55] R. Franklin, D. Carver and B.L. Hutchings, "Assisting network intrusion detection with reconfigurable hardware", *Proc. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 2002.
- [56] C. Fraser and D. Hanson, *A Retargetable C Compiler: Design and Implementation*, Benjamin/Cummings Pub. Co, 1995.
- [57] J. Frigo, D. Palmer, M. Gokhale and M. Popkin-Paine, "Gamma-Ray pulsar detection using reconfigurable computing hardware", *Proc. Symp. on Field Programmable Custom Computing Machines*, IEEE Computer Society Press, 2003.
- [58] A. Gayasen, K. Lee, N. Vijaykrishnan, M. Kandemir, M. J. Irwin and T. Tuan, "A dual- V_{DD} low power FPGA architecture", *Field Programmable Logic and Applications*, LNCS Series, Springer, 2004.
- [59] V. George, H. Zhang, and J. Rabaey, "The design of a low energy FPGA", *Proc. Int. Symp. on Low Power Electronics and Design*, 1999.
- [60] M. Gokhale and J. Stone, "NAPA-C: compiling for a hybrid RISC/FPGA architecture", *Proc. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 1998.
- [61] M. Gokhale, J.M. Stone, J. Arnold and M. Kalinowski, "Stream-oriented FPGA computing in the Streams-C high level language", *Proc. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 2000.
- [62] M. Gokhale, D. Dubois, A. Dubois, M. Boorman, S. Poole and V. Hogsett, *Granidt: towards gigabit rate network intrusion detection technology*, *Field Programmable Logic and Applications*, LNCS 2438, Springer, 2002.
- [63] S.C. Goldstein, H. Schmit, M. Budiu, S. Cadambi, M. Moe and R. Taylor, "PipeRench: a reconfigurable architecture and compiler", *IEEE Computer*, Vol. 33, No. 4, 2000, pp. 70-77.
- [64] Z. Guo, W. Najjar, F. Vahid and K. Vissers, "A quantitative analysis of the speedup factors of FPGAs over processors", *Proc. Int. Symp. on FPGAs*, ACM Press, 2004.
- [65] S. Guo and W. Luk, "An integrated system for developing regular array design", *Journal of Systems Architecture*, Vol. 47, 2001.
- [66] S. Gupta, N.D. Dutt, R.K. Gupta and A. Nicolau, "SPARK: a high-level synthesis framework for applying parallelizing compiler transformations", *Proc. International Conference on VLSI Design*, January 2003.
- [67] N. Gura, S.C. Shantz, H. Eberle, S. Gupta, V. Gupta, D. Finchelstein, E. Goupy and D. Stebila, "An end-to-end systems approach to Elliptic Curve Cryptography", *Proc. Cryptographic Hardware and Embedded Systems*, LNCS 2523, 2002.
- [68] P. Haglund, O. Mencer, W. Luk and B. Tai, "PyHDL: hardware scripting with Python", *Proc. Int. Conf. on Engineering of Reconfigurable Systems and Algorithms*, 2003.
- [69] D. Hankerson, A. Menezes and S. Vanstone, *Guide to Elliptic Curve Cryptography*, Springer, 2004.
- [70] P. Hanrahan, "Using caching and breadth-first search to speed up ray-tracing", *Proc. Graphics Interface '86*, May 1986.
- [71] J.R. Hauser and J. Wawrzyniek, "Garp: a MIPS processor with a reconfigurable processor", *IEEE Symposium on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 1997.
- [72] T. Harriss, R. Walke, B. Kienhuis and E. Deprettere, "Compilation from Matlab to process networks realized in FPGA", *Design Automation of Embedded Systems*, Vol. 7, No. 4, 2002.
- [73] S.D. Haynes, J. Stone, P.Y.K. Cheung and W. Luk, "Video image processing with the SONIC architecture", *IEEE Computer*, Vol. 33, 2000, pp. 50-57.
- [74] S.D. Haynes, H.G. Epsom, R.J. Cooper and P.L. McAlpine, "Ultra-SONIC: a reconfigurable architecture for video image processing", *Field Programmable Logic and Applications*, LNCS 2438, Springer, 2002.
- [75] C.A.R. Hoare, *Communicating Sequential Processes*, Prentice Hall, 1985.
- [76] A. Hodjat and I. Verbauwhede, "A 21.54 Gbit/s fully pipelined AES processor on FPGA", *Proc. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 2004.
- [77] H. Högl, A. Kugel, J. Ludvig, R. Männer, K. Noffz and R. Zoz, "Enable++: a second-generation FPGA processor", *IEEE Symposium on FPGAs for Custom Computing Machines*, IEEE Computer Society Press, 1995.
- [78] G. Holzmann, "The model checker SPIN", in *IEEE Transactions on Software Engineering*, Vol. 23, No. 5, May 1997, pp. 279-295.
- [79] J. Hwang, B. Milne, N. Shirazi and J.D. Stroomer, "System level tools for DSP in FPGAs", *Field-Programmable Logic and Applications*, LNCS 2147, Springer, 2001.
- [80] P.B. James-Roxby and D.J. Downs, "An efficient contentaddressable memory implementation using Dnamic routing", *Proc. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 2001.
- [81] J. Jiang, W. Luk and D. Rueckert, "FPGA-based computation of free-form deformations in medical image registration", *Proc. Int. Conf. on Field-Programmable Technology*, IEEE, 2003.
- [82] H. Kagotani and H. Schmit, "Asynchronous PipeRench: architecture and performance evaluations", *Proc. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 2003.
- [83] V. Kathail et. al., "PICO: automatically designing custom computers", *Computer*, Vol. 35, No. 9, Sept. 2002.
- [84] H. Keding, M. willems, M. Coors and H. Meyr, "FRIDGE: A fixed-point design and simulation environment", *Proc. Design Automation and Test in Europe*, 1998.
- [85] C. Kulkarni, G. Brebner and G. Schelle, "Mapping a domain specific language to a platform FPGA", *Proc. Design Automation Conference*, 2004.
- [86] K. Kum and W. Sung, "Word-length optimization for high-level synthesis of digital signal processing systems", *Proc. IEEE Int. Workshop on Signal Processing Systems*, 1998.
- [87] K.-I. Kum and W. Sung, "Combined word-length optimization and high-level synthesis of digital signal processing systems", *IEEE Trans. Computer Aided Design*, Vol. 20, No. 8, pp. 921-930, August 2001.
- [88] J. Lamoureux and S.J.E. Wilton, "On the interaction between power-aware FPGA CAD algorithms", *Int. Conf. on Computer-Aided Design*, IEEE, 2003.
- [89] Lattice Semiconductor Corp, *ispXPGA Family*, January 2004.
- [90] R. Laufer, R. Taylor and H. Schmit, "PCI-PipeRench and the Swor-dAPI: a system for stream-based reconfigurable computing", *Proc. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 1999.
- [91] D.U. Lee, W. Luk, C. Wang, C. Jones, M. Smith and J. Villasenor, "A flexible hardware encoder for Low-Density Parity-Check Codes", *Proc. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 2004.
- [92] E.A. Lee and D.G. Messerschmitt, "Static scheduling of synchronous data flow programs for digital signal processing", *IEEE Trans. Computers*, January 1987.
- [93] T.K. Lee, S. Yusuf, W. Luk, M. Sloman, E. Lupu and N. Dulay, "Compiling policy descriptions into reconfigurable firewall processors", *Proc. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 2003.
- [94] T.K. Lee, A. Derbyshire, W. Luk and P.Y.K. Cheung, "High-level language extensions for run-time reconfigurable systems", *Proc. Int. Conf. on Field-Programmable Technology*, IEEE, 2003.
- [95] G. Lemieux and D. Lewis, *Design of Interconnect Networks for Programmable Logic*, Kluwer Academic Publishers, Feb. 2004.
- [96] J. Leonard and W.H. Mangione-Smith, "A case study of partially evaluated hardware circuits: key-specific DES", *Field Programmable Logic and Applications*, LNCS 1304, Springer 1997.
- [97] P.H.W. Leong and K.H. Leung, "A microcoded Elliptic Curve Processor using FPGA technology", *IEEE Transactions on Very Large Scale Integration Systems*, Vol. 10, No. 5, 2002, pp. 550-559.
- [98] J. Liang, R. Tessier and O. Mencer, "Floating point unit generation and evaluation for FPGAs", *Proc. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 2003.
- [99] W. Luk, "Customising processors: design-time and run-time opportunities", *Computer Systems: Architectures, Modeling, and Simulation*, LNCS 3133, Springer, 2004.
- [100] W. Luk, P.Y.K. Cheung and N. Shirazi, "Configurable computing", *Electrical Engineer's Handbook*, W.K. Chen (ed.), Academic Press, 2004.
- [101] W. Luk and S.W. McKeever, "Pebble: a language for parametrised and reconfigurable hardware design", *Field-Programmable Logic and Applications*, LNCS 1482, Springer, 1998.
- [102] S. Mahlke et. al., "Bitwidth cognizant architecture synthesis of custom hardware accelerators", *IEEE Trans. on Computer-Aided Design*, Vol. 20, No. 11, November 2001.

- [103] A. Marshall, T. Stansfield, I Kostarnov, J. Vuillemin and B. Hutchings, "A reconfigurable arithmetic array for multimedia applications", *ACM/SIGDA International Symposium on FPGAs*, Feb 1999, pp. 135-143.
- [104] A. Mazzeo, L. Romano and G.P. Saggese, "FPGA-based implementation of a serial RSA processor", *Proc. Design, Automation and Test in Europe Conf.*, IEEE, 2003.
- [105] S.W. McKeever, W. Luk and A. Derbyshire, "Compiling hardware descriptions with relative placement information for parametrised libraries", *Formal Methods in Computer-Aided Design*, LNCS 2517, Springer, 2002.
- [106] B. Mei, S. Vernalde, D. Verkest, H. De Man and R. Lauwereins, "ADRES: An architecture with tightly coupled VLIW processor and coarse-grained reconfigurable matrix", *Field-Programmable Logic and Applications*, LNCS 2778, Springer, 2003.
- [107] O. Mencer, "PAM-Blox II: design and evaluation of C++ module generation for computing with FPGAs", *Proc. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 2002.
- [108] O. Mencer, D.J. Pearce, L.W. Howes and W. Luk, "Design space exploration with A Stream Compiler", *Proc. Int. Conf. on Field Programmable Technology*, IEEE, 2003.
- [109] Mentor Graphics, *Vstation Pro: High Performance System Verification*, 2003.
- [110] E. Mirsky and A. DeHon, "MATRIX: a reconfigurable computing architecture with configurable instruction distribution and deployable resources", *Proc. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 1996.
- [111] P. Montgomery, "Modular multiplication without trial division", *Math. of Computation*, Vol. 44, 1985, pp. 519-521.
- [112] K. Morris, "Virtex 4: Xilinx details its next generation", *FPGA and Programmable Logic Journal*, June 2004.
- [113] H. Muller and J. Winckler, "Distributed image synthesis with breadth-first ray tracing and the Ray-Z-buffer", *Data Structures and Efficient Algorithms - Final Report on the DFG Special Initiative*, LNCS 594, Springer, 1992.
- [114] S.H. Muggleton, "Inverse entailment and Progol", *New Generation Computing*, Vol. 13, 1995.
- [115] K. Nakamaru and Y. Ohno, "Breadth-first ray tracing using uniform spatial subdivision", *IEEE Transactions on Visualization and Computer Graphics*, Vol. 3, No. 4, IEEE, 1997.
- [116] Nallatech, *Ballynuey 2: Full-length PCI Card DIME Motherboard with four DIME slots*, Nallatech, NT190-0045, 2002.
- [117] National Bureau of Standards, *Announcing the Data Encryption Standard*, Technical Report FIPS Publication 46, US Commerce Department, National Bureau of Standards, January 1977.
- [118] National Institute of Standards and Technology, *Advanced Encryption Standard*, <http://csrc.nist.gov/publication/drafts/dfips-AES.pdf>
- [119] A. Nayak, M. Haldar, A. Choudhary and P. Banerjee, "Precision and error analysis of MATLAB applications during automated hardware synthesis for FPGAs", *Proc. Design Automation and Test in Europe*, 2001.
- [120] S. Ong, N. Kerkiz, B. Srijanto, C. Tan, M. Langston, D. Newport and D. Bouldin, "Automatic mapping of multiple applications to multiple adaptive computing systems", *Proc. Int. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 2001.
- [121] G. Orlando, and C. Paar, "A high performance reconfigurable Elliptic Curve for $GF(2^m)$ ", *Proc. Workshop on Cryptographic Hardware and Embedded Systems*, LNCS 1965, 2000.
- [122] G. Orlando, and C. Paar, "A scalable $GF(p)$ Elliptic Curve processor architecture for programmable hardware", *Proc. Workshop on Cryptographic Hardware and Embedded Systems*, LNCS 2162, 2001.
- [123] J. Ou and V. Prasanna, "PyGen: a MATLAB/Simulink based tool for synthesizing parameterized and energy efficient designs using FPGAs", *Proc. Int. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 2004.
- [124] I. Page and W. Luk, "Compiling occam into FPGAs", *FPGAs*, Abingdon EE&CS Books, 1991.
- [125] S. Parker, W. Martin, P. Sloan, P. Shirley, B. Smits and C. Hansen, "Interactive ray tracing", *Proc. 1999 Symposium on Interactive 3D Graphics*, ACM Press, April 1999.
- [126] V. Pasham, and S. Trimberger, "High-speed DES and Triple DES encryptor/decryptor", Xilinx Application note: XAPP270, 2001. <http://www.xilinx.com/bvdocs/appnotes/xapp270.pdf>
- [127] C. Patterson, "High performance DES encryption in Virtex FPGAs using JBits", *Proc. Int. Conf. on Field-Programmable Custom Computing Machines*, IEEE, pp. 113-121, 2000.
- [128] J. Peterson, B. O'Connor and P. Athanas, "Scheduling and partitioning ANSI-C programs onto multi-FPGA CCM architectures", *Int. Symp. on FPGAs for Custom Computing Machines*, IEEE Computer Society Press, 1996.
- [129] D. Plunkett and M. Bailey, "The vectorization of a ray-racing algorithm for increased speed", *IEEE Computer Graphics and Applications*, Vol. 5, No. 8, 1985.
- [130] G.M. Quenot, I.C. Kraljic, J. Serot and B. Zavidovique, "A reconfigurable compute engine for real-time vision automata prototyping", *Proc. IEEE Workshop on FPGAs for Custom Computing Machines*, IEEE Computer Society Press, 1994.
- [131] Quicklogic Corp., *Eclipse-II Family Datasheet*, January 2004.
- [132] A. Rahman and V. Polavarapuv, "Evaluation of low-leakage design techniques for Field Programmable Gate Arrays", *Proc. Int. Symp. on Field-Programmable Gate Arrays*, ACM Press, 2004.
- [133] S. Rathman and G. Slavenburg, "Processing the new world of interactive media", *IEEE Signal Processing Magazine*, Vol. 15, Issue 2, March 1998, pp. 108-117.
- [134] R. Razdan and M.D. Smith, "A high performance microarchitecture with hardware programmable functional units", *International Symposium on Microarchitecture*, 1994, pp. 172-180.
- [135] T. Rissa, W. Luk and P.Y.K. Cheung, "Automated combination of simulation and hardware prototyping", *Proc. Int. Conf. on Engineering of Reconfigurable systems and algorithms*, CSREA Press, 2004..
- [136] R.L. Rivest, A. Shamir and L. Adleman, "A method for obtaining digital signatures and public-key cryptosystems", *Communications of ACM*, Vol. 21, No. 2, 1978, pp. 120-126.
- [137] A. Royal and P.Y.K. Cheung, "Globally asynchronous locally synchronous FPGA architectures", *Field Programmable Logic and Applications*, LNCS 2778, Springer, 2003.
- [138] C.R. Rupp, M. Landguth, T. Garverick, E. Gomersall, H. Holt, J. Arnold and M. Gokhale, "The NAPA adaptive processing architecture", *IEEE Symposium on Field-Programmable Custom Computing Machines*, May 1998, pp. 28-37.
- [139] G.P. Saggese, A. Mazzeo, N. Mazzocca and A.G.M. Strollo, "An FPGA-based performance analysis of the unrolling, tiling, and pipelining of the AES Algorithm", *Field-Programmable Logic and Applications*, LNCS 2778, 2003.
- [140] A. Satoh and K. Takano, "A scalable dual-field Elliptic Curve Cryptographic Processor", *IEEE Transactions on Computers*, Vol. 52, No. 4, 2003, pp. 449-460.
- [141] P. Schaumont, I. Verbauwhede, K. Keutzer and M. Sarrafzadeh, "A quick safari through the reconfiguration jungle", *Proc. Design Automation Conf.*, ACM Press, 2001.
- [142] R. Schreiber et. al., "PICO-NPA: High-level synthesis of nonprogrammable hardware accelerators", *Journal of VLSI Signal Processing Systems*, Vol. 31, No. 2, June 2002.
- [143] P. Sedcole, P.Y.K. Cheung, G.A. Constantinides and W. Luk, "A reconfigurable platform for real-time embedded video image processing", *Field Programmable Logic and Applications*, LNCS 2778, Springer, 2003.
- [144] S.P. Seng, W. Luk and P.Y.K. Cheung, "Flexible instruction processors", *Proc. Int. Conf. on Compilers, Arch. and Syn. for Embedded Systems*, ACM Press, 2000.
- [145] S.P. Seng, W. Luk and P.Y.K. Cheung, "Run-time adaptive flexible instruction processors", *Field-Programmable Logic and Applications*, LNCS 2438, Springer, 2002.
- [146] N. Shirazi, W. Luk and P.Y.K. Cheung, "Framework and tools for runtime reconfigurable designs", *IEE Proc.-Comput. Digit. Tech.*, May, 2000.
- [147] Silicon Hive, "Avispa Block Accelerator", Product Brief, 2003.
- [148] M. Sima, S.D. Cotozana, S. Vassiliadis, J.T.J. van Eijndhoven and K.A. Vissers, "Pel reconstruction on FPGA-augmented TriMedia", *IEEE Trans. on VLSI*, Vol. 12, No. 6, June 2004, pp. 622-635.
- [149] Simulink, <http://www.mathworks.com>
- [150] H. Singh, M-H Lee, G. Lu, F. Kurdahi, N. Bagherzadeh and E. Chaves, "MorphoSys: An integrated reconfigurable system for data-parallel and compute intensive applications", *IEEE Trans. on Computers*, Vol. 49, No. 5, May 2000, pp. 465-481.

- [151] S. Singh and C.J. Lillieroth, "Formal verification of reconfigurable cores", *Proc. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 1999.
- [152] I. Sourdis and D. Pnevmatikatos, "Fast, large-scale string matching for a 10Gbps FPGA-based network intrusion system", *Field Programmable Logic and Applications*, LNCS 2778, Springer, 2003.
- [153] Spin, <http://spinroot.com/spin/whatispin.html>.
- [154] F.X. Standaert, G. Rouvroy, J.J. Quisquater and J.D. Legat, "Efficient implementation of Rijndael encryption in reconfigurable hardware: improvements and design tradeoffs", *Proc. Cryptographic Hardware and Embedded Systems*, LNCS 2779, 2003.
- [155] M. Stephenson, J. Babb and S. Amarasinghe, "Bitwidth analysis with application to silicon compilation", *Proc. SIGPLAN Programming Language Design and Implementation*, June 2000.
- [156] M.W. Stephenson, *Bitwise: Optimizing bitwidths using data-range propagation*, Master's Thesis, Massachusetts Institute of Technology, Dept. Electrical engineering and Computer Science, May 2000.
- [157] G. Stitt, F. Vahid and S. Nematbakhsh, "Energy savings and speedups from partitioning critical software loops to hardware in embedded systems", *ACM Trans. on Embedded Computing Systems*, Vol. 3, No. 1, Feb. 2004, pp. 218–232.
- [158] H. Styles and W. Luk, "Customising graphics applications: techniques and programming interface", *Proc. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 2000.
- [159] H. Styles and W. Luk, "Accelerating radiosity calculations using reconfigurable platforms", *Proc. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 2002.
- [160] H. Styles and W. Luk, "Branch optimisation techniques for hardware compilation", *Field-Programmable Logic and Applications*, LNCS 2778, Springer, 2003.
- [161] W. Sung and K. Kum, "Word-length determination and scaling software for a signal flow block diagram", *Proc. IEEE Int. Conf. on Acoustics Speech and Signal Processing*, 1994.
- [162] W. Sung and K. Kum, "Simulation-based word-length optimization method for fixed-point digital signal processing systems", *IEEE Trans. Signal Processing*, Vol. 43, No. 12, December 1995, pp. 3087–3090.
- [163] S.H. Tang, K.S. Tsui and P.H.W. Leong, "Modular exponentiation using parallel multipliers", *Proc. Int. Conf. on Field Programmable Technology*, IEEE, 2003.
- [164] M. Taylor et al, "The RAW microprocessor: a computational fabric for software circuits and general purpose programs", *IEEE Micro*, vol 22, No. 2, March/April 2002, pp. 25–35.
- [165] J. Teife and R. Manohar, "Programmable asynchronous pipeline arrays", *Field Programmable Logic and Applications*, LNCS 2778, Springer, 2003.
- [166] N. Telle, C.C. Cheung and W. Luk, "Customising hardware designs for Elliptic Curve Cryptography", *Computer Systems: Architectures, Modeling, and Simulation*, LNCS 3133, Springer, 2004.
- [167] R. Tessier and W. Burleson, "Reconfigurable computing and digital signal processing: a survey", *Journal of VLSI Signal Processing*, Vol. 28, May/June 2001, pp. 7–27.
- [168] D. Thomas and W. Luk, "A framework for development and distribution of hardware acceleration", *Reconfigurable Technology: FPGAs and Reconfigurable Processors for Computing and Communications*, *Proc. SPIE*, Vol. 4867, 2002.
- [169] T. Todman and W. Luk, "Reconfigurable designs for ray tracing", *Proc. Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 2001.
- [170] T. Todman and W. Luk, "Combining imperative and declarative hardware descriptions", *Proc. 36th Hawaii Int. Conf. on System Sciences*, IEEE, 2003.
- [171] T. Todman, J.G.F. Coutinho and W. Luk, "Customisable hardware compilation", *Proc. Int. Conf. on Engineering of Reconfigurable Systems and Algorithms*, CSREA Press, 2004.
- [172] M.M. Uddin, Y. Cao and H. Yasuura, "An accelerated datapath width optimization scheme for area reduction of embedded systems", *Proc. Int. Symp. on Systems Synthesis*, ACM Press, 2002.
- [173] K. Underwood, "FPGAs vs. CPUs: trends in peak floating-point performance", *Proc. int. Symp. on FPGAs*, ACM Press, 2004.
- [174] L. Vereen, "Soft FPGA cores attract embedded developers", *Embedded Systems Programming*, 23 April 2004, <http://www.embedded.com/showArticle.jhtml?articleID=19200183>.
- [175] J. Vuillemin, P. Bertin, D. Roncin, M. Shand, H. Touati and P. Boucard, "Programmable Active Memories: Reconfigurable Systems come of age", *IEEE Transactions on VLSI Systems*, Vol. 4, No. 1, March 1996, pp. 56–69.
- [176] S.A. Wadekar and A.C. Parker, "Accuracy sensitive word-length selection for algorithm optimization", *Proc. International Conference on Computer Design*, 1998.
- [177] M. Weinhardt and W. Luk, "Pipeline vectorization", *IEEE Trans. on Computer-Aided Design*, Vol. 20, No. 2, 2001.
- [178] A. Wenban, J.W. O'Leary and G.M. Brown, "Codesign of communication protocols", *IEEE Computer*, Vol. 26, No. 12, pp. 46–52, December 1993.
- [179] A. Wenban and G. Brown, "A software development system for FPGA-based data acquisition systems", *IEEE Symposium on FPGAs for Custom Computing Machines*, IEEE Computer Society Press, 1996.
- [180] D.C. Wilcox, L.G. Pierson, P.J. Robertson, E.L. Witzke and K. Gass, "A DES ASIC suitable for network encryption at 10 Gbps and beyond", *Proc. Cryptographic Hardware and Embedded Systems*, LNCS 1717, 1999.
- [181] M. Willems, V. Bürgens, H. Keding, T. Grötter and M. Meyer, "System-level fixed-point design based on an interpolative approach", *Proc 34th Design Automation Conference*, June 1997.
- [182] J.A. Williams, A.S. Dawood and S.J. Visser, "FPGA-based cloud detection for real-time onboard remote sensing", *Proc. Int. Conf. on Field-Programmable Technology*, IEEE, 2002.
- [183] R.S. Williams and P.J. Kuekes, "Molecular nanoelectronics", *Proc. Int. Symp. on Circuits and Systems*, IEEE, 2000.
- [184] R.P. Wilson, R.S. French, C.S. Wilson, S.P. Amarasinghe, J.M. Anderson, S.W.K. Tjiang, S.-W. Liao, C.-W. Tseng, M.W. Hall, M.S. Lam and J.L. Hennessy, "SUIF: an infrastructure for research on parallelizing and optimizing compilers", *ACM SIGPLAN Notices*, Vol. 29, No. 12, December 1994.
- [185] R.D. Wittig and P. Chow, "OneChip: an FPGA processor with reconfigurable logic", *IEEE Symposium on FPGAs for Custom Computing Machines*, 1996.
- [186] C.G. Wong, A.J. Martin and P. Thomas, "An architecture for asynchronous FPGAs", *Proc. Int. Conf. on Field-Programmable Technology*, IEEE, 2003.
- [187] R. Woods, D. Trainor and J.P. Heron, "Applying an XC6200 to real-time image processing", *IEEE Design and Test of Computers*, Vol. 15, Jan.-March 1998, pp. 30–38.
- [188] Xilinx, Inc., *PowerPC 405 Processor Block Reference Guide*, October 2003.
- [189] Xilinx, Inc., *Virtex II Datasheet*, June 2004.
- [190] Xilinx, Inc., *Microblaze Processor Reference Guide*, June 2004.
- [191] A. Yamada, K. Nishida, R. Sakurai, A. Kay, T. Nomura and T. Kambe, "Hardware synthesis with the Bach system", *Proc. ISCAS*, IEEE, 1999.
- [192] H. Ziegler, B. So, M. Hall and P. Diniz, "Coarse-grain pipelining on multiple-FPGA architectures", in *IEEE Symposium on Field-Programmable Custom Computing Machines*, IEEE, 2002, pp77–88.