FPGA-based Design of an Evolutionary Controller for Collision-free Robot Navigation

M. A. H. B. Azhar, K. R. Dimond
Department of Electronics, University of Kent, Canterbury, Kent,CT2 7NT, UK
{maha2, k.r.dimond}@ukc.ac.uk

The employment of field programmable gate arrays (FPGAs) to a robot controller is very attractive, since it allows for fast IC prototyping and low cost modifications. The speedup is achieved because of pipelining and dedicated functions in hardware that are customized to the problem. The self learning ability and the adaptive nature of an Artificial Neural Network (ANN) makes it a good candidate for the control structure of a robot's navigation. An evolutionary approach in designing robots can evolve the architecture of ANNs and yields automatic creation of the controller while the robot moves in task environments. The poster briefly describes the important hardware issues involved with the FPGA based design of an evolutionary robot controller for the collision free navigation of mobile robots

FPGA Implementation of a Fast Hadamard Transformer for WCDMA

Sanat Kamal Bahl, Jim Plusquellic
Dept. of Computer Science and Electrical Engineering, University of Maryland, Baltimore County,
1000 Hilltop Circle, Baltimore, MD 21227
{skamal2, plusquel}@cs.umbc.edu
http://www.gl.umbc.edu/~skamal2

In code division multiple access (CDMA) systems the base station identifies each user in a cell by unique orthogonal (Walsh) codes. The Walsh codes are generated at the transmitter using a Walsh-Hadamard function. A Fast Hadamard Transformer (FHT) is used at the receiver to decode the transmitted codes. The purpose of this study is to design a FHT which utilizes less hardware resources as compared to the existing designs and also suggest means for reducing the input length of the Walsh sequence. Our study results indicate that the FHT design using 16-chip sequence achieves 90% reduction in hardware resources (equivalent gate count) as compared to the design which uses 256-chip sequence. Also, the maximum frequency of operation of the 16-chip FHT (35.679 MHz) is more than double as compared to the 256-chip FHT (16.025 MHz).

Making Area-Performance Tradeoffs at the High Level Using the AccelFPGA Compiler for FPGAs

P. Banerjee, V. Saxena, J. Uribe, M. Haldar, A. Nayak, V. Kim, D. Bagchi, S. Pal,
N. Tripathi, R. Anderson
AccelChip, Inc., 999 Plaza Drive, Suite 340, Schaumburg, IL-60103
www.accelchip.com

Applications such as digital cell phones, 3G wireless receivers, and voice over IP, require DSP functions that are typically mapped onto general purpose DSP processors. With the introduction of advanced FPGA architectures which provide built-in DSP support such as the Xilinx Virtex-II, and the Altera Stratix, a new hardware alternative is available for DSP designers. DSP design has traditionally been divided into algorithm development and hardware/software implementation. The majority of DSP algorithm developers use the MATLAB language for prototyping their DSP algorithm. Hardware design teams take the specifications in MATLAB code and manually create an RTL model in VHDL or Verilog. This paper describes how area-performance tradeoffs can be performed quickly at the high-level using a behavioral synthesis tool called AccelFPGA which reads in high-level descriptions of DSP applications written in MATLAB, and automatically generates synthesizable RTL models in VHDL or Verilog. Experimental results are reported with the AccelFPGA compiler on a set of 8 MATLAB benchmarks that are mapped onto the Xilinx Virtex II and Altera Stratix FPGAs.

Design Framework for the Implementation of the 2-D Orthogonal Discrete Wavelet Transform on FPGA

A. Benkrid, D. Crookes, K. Benkrid
School of Computer Science,The Queen's University of Belfast, Belfast BT7 1NN, UK
a.benkrid@qub.ac.uk

This paper gives a design framework for the implementation of the 2-D Orthogonal Discrete Wavelet Transform (DWT) on FPGA. The architecture is based on the Pyramid Algorithm Analysis. Our architecture spatially maps the multistage filter banks of the DWT onto the Xilinx Virtex-E FPGA family. In this paper we propose a novel FIR structure to handle the computation along the borders using symmetric extension. The paper includes a new detailed mathematical approach to determine the architecture's dynamic range as well as predicting and reducing the error dynamic range due to wordlength rounding. For an NxN image size input, our architecture has a period of N2 clock cycles, and requires only the minimum storage size. The architecture is highly scalable for different filter lengths and number of octaves. The implementation results for a specific 2-D Daubechies-4 Wavelet Transform are included.

A Logic Based Approach to Hardware Abstraction

K. Benkrid, S. Belkacemi, D. Crookes
School of Computer Science,The Queen's University of Belfast, Belfast BT7 1NN, UK
k.benkrid@qub.ac.uk, http://www.cs.qub.ac.uk/~K.Benkrid/

This paper presents a novel approach to hardware abstraction based on the logic programming language Prolog. This is an attempt to satisfy the dual requirement of abstract hardware design and hardware efficiency. Central to this approach is a hardware description environment called HIDE, which provides more abstract hardware descriptions and compositions than are possible in traditional hardware description languages such as VHDL or Verilog. HIDE enables highly scaleable and parameterised composition of blocks using a small set of abstract constructors such as the horizontal and vertical constructors for 2D circuit abstractions, and the novel above constructor for 3D circuit abstractions. It also generates pre-placed configurations in EDIF (and VHDL) format for Xilinx FPGAs. The paper presents the syntax and semantics of HIDE and illustrates our logic-based approach in the construction of a high performance Matrix Multiplier core for Xilinx Virtex FPGAs.

A Single-FPGA Implementation of Image Connected Component Labelling

K. Benkrid, S. Sukhsawas, D. Crookes, S. Belkacemi
The Queens Univ. of Belfast, School of Computer Science, Belfast BT7 1NN Northern Ireland, UK
k.benkrid@qub.ac.uk, http://www.cs.qub.ac.uk/~K.Benkrid/

This paper describes an architecture based on a serial iterative algorithm for Image Connected Component Labelling with a hardware complexity O(N) for an NxN image. The algorithm iteratively scans the input image, performing a recursive non-zero maximum neighbourhood operation. A complete forward pass is followed by an inverse pass in which the image is scanned in reverse order. The process is repeated until no change in the image occurs. The algorithm has been coded in Handel C language and targeted to a Celoxica RC1000-PP PCI board, which is based on a Virtex XCV2000E-6 FPGA. The whole design was fully implemented and tested on real hardware in less than 24 man-hours of work. The implementation speed (pixel throughput) is virtually independent of the image size and is equal to ~34MHz. For 1024x1024
input images, the whole circuit consumes 566 Slices and 5 BlockRAMs and can run at 34 MHz, leading to a 32 pass/sec performance.

An Estimation and Exploration Methodology from System-Level Specifications: Application to FPGAs

Sebastien Bilavarn[1], Guy Gogniat[2], Jean Luc Philippe[2]
[1]EPFL-ITS-LTS2, Swiss Federal Institute of Technology
sebastien.bilavarn@epfl.ch
[2]LESTER, South Britany University, Lorient, France
{guy.gogniat, jean-luc.philippe}@univ-ubs.fr

Rapid evaluation and design space exploration from early specifications are important issues in the design cycle. We propose an original area vs. delay estimation methodology that targets reconfigurable architectures. Two main steps compose the estimation flow: i) structural estimations where architectural solutions are defined at the RT level, this step is technological independent and performs an automatic design space exploration and ii) physical estimations which perform technology mapping to the target reconfigurable architecture. Experiments conducted on Xilinx (XC4000, Virtex) and Altera (Flex10K, Apex) components for a 2D DWT and a speech coder lead to an average error of about 10% for temporal values and 18% for area estimations. The originality of this work is mainly a complete and realistic cost characterization that takes care of the processing, memory and control units, and supplies architectural information for the design of each solution.

A Granularity-based Classification Model for Systems-on-a-Chip

Stephan Bingemer, Peter Zipf, Manfred Glesner
Institute of Microelectronic Systems, Department of Electrical Engineering and Information Technology, Darmstadt University of Technology, Germany
bingemer@bwl.tu-darmstadt.de, {zipf, glesner}@mes.tu-darmstadt.de

Field-programmable logic has become an increasingly important technology for the design of digital circuits. One interesting point in the field of reconfigurable logic is its classification within the implementation space of other technologies. Such a classification gains importance if FPGA technology becomes an integral part of Systems-on-a-Chip (SoC). The poster discusses an approach to classify technologies based on their granularity. Therefore, a new distinction into homogeneous and heterogeneous granularity is proposed. This leads to a theoretic classification method for the combination of two or more technologies. Using this method the result of a combination of two or more technologies can be determined qualitatively. The poster further shows opportunities to improve standard CAD-software based on the model proposed.

Design of a Fingerprint System Using a Hardware/Software Environment

Lee Vanderlei Bonato[1], Rolf Fredi Molz[1], João Carlos Furtado[1], Marcos Flores Ferrão[1], Fernando G. Moraes[2]
[1]UNISC – Departamento de Informática, Av. Independência, 2293 Bairro Universitário. CEP: 96815-900, Santa Cruz do Sul – RS – Brazil
ferrao@dquimfis.unisc.br, rolf@dinf.unisc.br , jcarlosf@dinf.unisc.br, vbonato74@hotmail.com
[2]PUCRS – Faculdade de Informática, Av. Ipiranga, 6681 - Prédio 30., CEP: 90619-900 – Porto Alegre - Brazil
moraes@inf.pucrs.br

Processing system of fingerprint are CPU time intensive, being normally implemented in software. This paper present a new algorithm for fingerprint features localization, that can be easily implemented in hardware (system-on-a-chip, FPGA). This algorithm is composed by 3 stages, first stage read a fingerprint image (255x255pixels, ash tones) and apply a Gaussian Filter, after this, apply a absolute difference mask (ADM) for detector the edges in the image filtered and the last stage look for fingerprint features into the image. The information showed by 3th stage are the coordinate X and Y for each feature detected, asked minutiae. For localization the minutiae, the system pursue the edge detected by ADM, this edge represent ridge edge, and analyzing the information from each pixel pursued is possible to locate the minutiae. The average time for localization all minutiae into the fingerprint image, implemented in hardware (FLEX10KE Family, Altera), was 306 milliseconds. Beyond hardware implementation be fast, is possible create embedded systems.

Customized Regular Channel Design in FPGAs

Elaheh Bozorgzadeh, Majid Sarrafzadeh
Computer Science Department, UCLA, Los Angeles, USA
elib@cs.ucla.edu , http://www.cs.ucla.edu/~elib/

In this paper, we study the problem of customized regular segmentation design in FPGA routing channels. We propose a deterministic algorithm for segmentation design problem in which each interval is assigned to only one segment (1-Segmentation). We solve the problem of maximum number of incremental track assignment of intervals by mincost network flow technique for 1-Segmentation design. The general K-Segmentation design problem can also be solved by some modifications in our algorithm. We have experimented our algorithm on a set of MCNC benchmarks and compared the routability of the segmented channels with the routability of different segmentations in FPGA routing architectures, one used by industrial Xilinx TM 4000 FPGA series. The experimental results show that the routability of the segmentation in general-purpose architecture can be as low as 18.4% while the routability of our proposed customized segmentation can be as high as 91%. This result shows the gap between the general FPGA routing architecture and customized architecture for a given application can be very significant and our method is capable of generating such optimized segmentation with high routability for a given application.

A High-speed Successive Erasure BCH Decoder Architecture

Thomas Buerner
Friedrich-Alexander University of Erlangen-Nuremberg, Institute for Computer Aided Circuit Design,
Paul-Gordan-Str. 5, 91052 Erlangen, Germany
buerner@lrs.eei.uni-erlangen.de

A new high speed architecture for a BCH successive erasure decoder is presented. The Berlekamp-Massey based decoder by Sarwate and Shanbhag is extended to handle successive erasures. The critical path in the calculation submodules is increased from $T_{add}+T_{mult}$ to $T_{add}+T_{mult}+T_{mux}$. The proposed architecture is implemented exemplary for a BCH(63,45,7) code with up to two erasures on a XILINX Spartan2E300-7. Thus a clock frequency of 95 MHz is reached using 47% of the available slices instead of 105 MHz with 30% slices for the unextended architecture. The additional performance gain accessible by successive erasure decoding is in this case about 0.5 dB at a bit error rate of 10^{-7} and below over an additive white Gaussian noise channel and binary phase shift keying.

Implementation of Digital Fixed-Point Approximations to Continuous-Time IIR Filters

J. E. Carletta, R. J. Veillette, F. W. Krach, Z. Fang
The Univ. of Akron, Department of Electrical and Computer Engineering, Akron, OH 44325-3904
http://www.ecgf.uakron.edu/~elec/

An analytical framework for the implementation of digital infinite impulse response filters in fixed-point hardware on FPGAs is presented. It presumes that a continuous-time filter with the desired response is given. Within the framework, the constant coefficient bit widths are determined by accounting for the sensitivity of the filter's pole and zero locations with respect to the coefficient perturbations. The internal signal bit widths are determined by calculating theoretical bounds on the ranges of the signals, and on the errors introduced by truncation in the fixed-point hardware. The bounds form the basis for a methodology for the fixed-point digital implementation of a given continuous-time filter. The methodology avoids overflow, and guarantees a prescribed degree of accuracy in the filter output. The methodology is applied to a second-order filter used as a compensator in a magnetic bearing control system.

Track Placement: Orchestrating Routing Structures to Maximize Routability

Katherine Compton[1], Scott Hauck[2]
[1]Northwestern University, Evanston, IL USA
kati@ece.northwestern.edu
[2]University of Washington, Seattle, WA USA
hauck@ee.washington.edu

The design of a routing channel for an FPGA is a complex process, requiring the careful balance of flexibility with silicon efficiency. With the growing move towards embedding FPGAs into SoC designs, and the opportunity to automatically generate FPGA architectures, this problem becomes even more critical. The design of a routing channel requires determining the number of routing tracks, the length of the wires in those tracks, and the positioning of the breaks on the tracks. This paper focuses on the last of these problems, the placement of breaks in tracks to maximize overall flexibility. We have developed both an optimal algorithm and a number of heuristics to solve the track placement problem. The optimal algorithm finds a best solution provided the problem meets a number of restrictions. Most of the heuristics are without restrictions, and the most promising of these find solutions on average within 1.13% of optimal. Full text available as UWEE Technical Report #UWEETR-2002-0013 at http://www.ee.washington.edu/

Recursive Circuit Clustering for Minimum Delay and Area

Mehrdad Eslami Dehkordi, Stephen D. Brown
Dept. of Electrical & Computer Engineering, Univ. of Toronto, Toronto, Ontario, Canada M5S 3G4
{eslami, brown}@eecg.toronto.edu

We present an effective recursive algorithm for circuit clustering for delay and area minimization, which is applicable to FPGAs. At the highest level of clustering, the circuit is clustered using a modified single-level clustering algorithm. A cluster to netlist transformation technique is proposed, which converts each cluster into a new subcircuit. The algorithm then continues recursively by clustering the generated subcircuits into further levels of clusters. To reduce the amount of node duplication and the number of clusters at each level of clustering, we propose a node removal algorithm based on the node slack along with a simple cluster-packing algorithm. Experimental results on the two-level clustering problem using Quartus Design System from Altera show that our algorithm achieves, on average, 7.3% more delay reduction when compared to the latest published work on the problem. Also total FPGA compile time reported by Quartus is reduced by 36%.

Using FPGAs for Data and Reorganization Engines: Preliminary Results for Spatial Pointer-based Data Structures

Pedro C. Diniz, Joonseok Park
University of Southern California / Information Sciences Institute
pedro@isi.edu
http://www.isi.edu/~pedro/projects/SLATE/slate.html

FPGAs have appealing features such as customizable internal and external bandwidth and the ability to exploit vast amounts of fine-grain instruction-level parallelism. In this paper we explore the applicability of these features in using FPGAs as data search and reorganization engines for performing search and reorganization computations over spatial pointer-based data structures for which traditional computing platforms perform poorly. The preliminary experiments, for a set of simple spatial queries over spatial sparse-mesh and quad-tree data structures, reveal that 3 year-old FPGA devices can deliver performance that is on par and in some instances even superior to that of today's workstations. This experience suggests that the integration in memory of FPGA-like fabrics for implementing smart memory engines should be performance-wise very advantageous.

On Hiding Latency in Reconfigurable Systems: The Case of Merge-Sort for an FPGA-Based System

Hossam ElGindy, George Ferizis
http://www.cse.unsw.edu.au/~hossam/officialpage/index.html

Recursive solutions are effective software techniques that are difficult to map into hardware due to their dependency on input size and data values. As a result, most high-level design tools do not allow for recursive calls. In this paper we present a technique for mapping the merge-sort algorithm, as a case study, into a reconfigurable system. Our mapping employs an on-line prediction method to reconfigure the necessary hardware only when the need arises, and to hide the reconfiguration delay. As a result, our implementation uses the smallest possible size hardware to sort an input data stream without prior knowledge of its length and eliminates the reconfiguration delay penalty. We outline a reconfigurable system with self-organizing multiple-buses as the communication subsystem. The processing elements and memory modules are connected to the multiple-buses as a linear array. We also demonstrate the effectiveness of adding simple LUTs to the multiple-buses in improving the throughput by allowing for pipelining at the word level.

Testing for Bit Error Rate in FPGA Communication Interfaces

Yongquan Fan, Zeljko Zilic
Department of Electrical and Computer Engineering, McGill University, Montreal, Canada
yfan@macs.ece.mcgill.ca, zeljko@macs.ece.mcgill.ca

FPGAs have witnessed an increased use of dedicated communication interfaces. With their increased use, it is becoming critical to test and properly characterize all such interfaces. Bit error rate (BER) characteristic is one of the basic measures of the performance of any digital communication system. We propose a scheme for BER testing in FPGAs, which exhibits a few orders of magnitude speedup compared to traditional software simulation methods. In this scheme, we include a novel implementation of an additive white Gaussian noise (AWGN) generator with high speed and high accuracy for channel emulator. Compared to traditional BER test products, our scheme can test BER under different noise conditions. The whole system is implemented as an IP core, suitable for a single FPGA device.

On Computation and Resource Management in an FPGA-based Computation Environment

Soheil Ghiasi, Karlene Nguyen, Elaheh Bozorgzadeh, Majid Sarrafzadeh
Computer Science Dept, UCLA
{soheil, karlene, elib, majid}@cs.ucla.edu
http://er.cs.ucla.edu

The idea of managing the comprising computations of an application executed in an FPGA-based system is presented. An efficient algorithm for exploiting the timing slack of building blocks of the application is proposed. The slack of these blocks can be utilized by replacing them with slower but "cheaper" modules and by assigning the computations to the proper resources. Thus, our approach manages the comprising computations and system resources at the same time. This is performed without compromising the timing constraints of the application and can lead to significant improvements in power dissipation, computation accuracy or other design metrics based on the application domain. Our algorithm is well-suited for arbitrary tree computations. Moreover, it delivers solutions that are desirably close to the optimal solution. Experimental results for a number of object tracking applications executed on resources embedded in cameras, show a significant amount of slack utilization

A SC-based Novel Configurable Analog Cell

Binlin Guo, Jiarong Tong
ASIC & System State-Key-Lab, Microelectronics Dept., Fudan University, Shanghai, China 200433
blguo2000@yahoo.com.cn

This paper presents a high performance Configurable Analog Cell (CAC) which is made up of a Basic Configurable Analog Cell (BCAC) and a digital converter block. The CAC can be used either for Field Programmable Analog Array (FPAA) or for Field Programmable Digital-Analog Mixed Array (FPMA). The BCAC include three innovative Programmable Switch Blocks (PSBs), three Programmable Capacitor Arrays (PCAs), and an amplifier. PSB and PCA can be programmed to generate many equivalent components. In addition, digital converter block is used for enhancing the function of CAC. The results of the configuration demonstrate that the CAC architecture can not only implement linear functions including gain amplifier, integrator, and filter, etc., but also realize complex nonlinear functions such as Voltage Control Oscillator (VCO), Analog-Digital Convertor (ADC), and so on. Compared with other architectures, the CAC architecture can offer substantially more analog functions.

FPGAs in Critical Hardware / Software Systems

Adrian J. Hilton, Gemma Townson, Jon G. Hall
{adrian.Hilton, gemma.townson}@praxis-cs.co.uk, j.g.hall@open.ac.uk
http://www.suslik.org/FPL/

FPGAs are being used in increasingly complex roles in critical systems, interacting with conventional critical software. Established safety standards require rigorous justification of safety and correctness of the conventional software in such systems. Newer standards now make similar requirements for safety-related electronic hardware, such as FPGAs, in these systems. In this paper we examine the current state-of-the-art in programming FPGAs, and their use in conventional (low-criticality) hardware/software systems. We discuss the impact that the safety standards requirements have on the co-development of hardware/software combinations in critical systems and suggest adaptations of existing best practice in software development that could discharge them. We pay particular attention to the development and analysis of high-level language programs for FPGAs designed to interact with conventional software

Power-aware Architectures and Circuits for FPGA-based Signal Processing

Frank Honoré, Ben Calhoun, Anantha Chandrakasan
MIT Department of Electrical and Computer Engineering
honore@mtl.mit.edu
http://www-mtl.mit.edu/research/icsystems/

This work showcases a power-aware system design methodology for DSP applications on reconfigurable hardware platforms. In particular, an enhanced FPGA architecture is proposed and analyzed for a deep submicron process technology. These enhancements reduce Configurable Logic Block (CLB) usage for distributed arithmetic implementations of signal processing applications by 50% or more thereby reducing the load on interconnect resources. Multi-Threshold CMOS (MTCMOS) circuit design techniques are aggressively applied to reduce subthreshold leakage using an auto power-down feature for unused logic. Results show a 14x reduction in leakage current for unused CLBs or CLBs in deep sleep mode. CLBs in active mode see up to 2.8x steady-state power reduction. A testchip demonstrating these techniques in 0.13 micron technology has been sent out for fabrication.

An Automated and Power-Aware Framework for Utilization of IP Cores in Hardware Generated from C Descriptions Targeting FPGAs

Alex Jones, Prith Banerjee
Center for Parallel & Distributed Computing, Dept. of Electrical & Computer Engineering, Technological Institute, Northwestern University, 2145 Sheridan Road, Evanston, IL 60208-3118
{akjones, banerjee}@ece.northwestern.edu

Use of hand optimized Intellectual Property (IP) logic cores is prolific in hardware design. While IP cores remain a standard way to utilize the improvement in FPGA technology and contend with time to market pressure through reuse, popularity of tools generating hardware descriptions from high-level languages is also increasing in popularity. PACT HDL combines these two methods within a power-aware framework. The PACT HDL compiler generates power optimized VHDL/Verilog from a C language description. This work presents an automated framework to incorporate arbitrary IP logic cores into the C to HDL flow. Thus, PACT HDL can leverage IP cores corresponding to C intrinsic operators. The logic cores to be used are specified by the user through compiler directives or compiler command line options. The framework is power-aware through use of an automated clock-gating technique to “turn off” the IP cores when not in use. The validity of the approach is demonstrated using several image and signal processing benchmarks with a variety of multiplier and divider implementations on a Xilinx Virtex FPGA.

Reconfigurable Randomized K-way Graph Partitioning

Fatih Kocan
Southern Methodist University
http://engr.smu.edu/~kocan

In this paper, a randomized k-way graph partitioning algorithm is mapped onto reconfigurable hardware. The randomized algorithm relies on repetitive running of the same algorithm with different random number sequences to achieve the (near-)optimal solution. The run-time and hardware requirements of this reconfigurable solution per a random number sequence are O(|V|-K) cycles and O(|V|log|V|+|E|) gates and flip-flops, respectively. Performance is improved further at the expense of more hardware by running multiple copies of the partitioning algorithm with different random number sequences concurrently, and/or splitting a random sequence into subsequences and running them in parallel. Furthermore, in the context of this mapping, dynamic randomly configurable pattern-generation-based random number generation methods are introduced.

Synthetic Circuit Generation Using Clustering and Iteration

Paul D. Kundarewich, Jonathan Rose
Dept. of Electrical & Computer Engineering, Univ. of Toronto, Toronto, Ontario, M5S 3G4, Canada
{kundarew, jayar}@eecg.utoronto.ca
http://www.eecg.toronto.edu/~jayar/software/Cgen/Cgen.htm

The development of next-generation CAD tools and FPGA architectures requires benchmark circuits to experiment with new algorithms and architectures. There has always been a shortage of good public benchmarks for these purposes, and even companies that have access to proprietary customer designs could benefit from designs that meet size and other particular specifications. In this paper, we present a new method of generating realistic synthetic benchmark circuits to help alleviate this shortage. The method significantly improves the quality of previous work by imposing the natural hierarchy of circuits through clustering and by using a simpler method of characterizing the nature of sequential circuits. Also, in contrast to current constructive generation methods, we employ new iterative techniques in the generation that provide better control over the generated circuit s characteristics. As in previous work, we assess the realism of the generated circuits by comparing properties of real circuits and generated "clones" of the real circuit after placement and routing. On average, the real and clone circuits' total detailed wirelength differed by only 14%, a major improvement over previous results. In addition, the minimum track count was within 14% and the critical path delay was within 10%.

An FPGA Architecture with Built-in Error Correction Capability

P. K. Lala, B. Kiran Kumar
Dept. of Computer Science & Computer Engineering, Univ. of Arkansas, Fayetteville, AR 72701
lala@engr.uark.edu

The use of very deep submicron technology makes VLSI-based digital systems more susceptible to transient or soft errors, and thus compromises their reliability. This paper proposes an FPGA architecture inspired by the human immune system that allows tolerance of transient errors. The architecture is composed of a two-dimensional array of identical functional cells with different genetic codes. These codes are chosen based on the required functions to be performed by the functional cells. An error in a FPGA based digital system designed using the proposed architecture is treated as an antigen by the system. Using its distributed defense mechanism the system heals itself from the effect of the error. A major advantage of this architecture is that, the outputs of functional cells are connected to the inputs of other physically adjacent functional cells without having to go through complicated routing. Thus, lengthy communication paths between cells are avoided.

Lattice Adaptive Filter Implementation for FPGA

Zdeněk Pohl[1], Rudolf Matoušek[1], Jiří Kadlec[1], Milan Tichý[1], Miroslav Líčko[2]
[1]Department of Signal Processing, Institute of Information Theory and Automation, CAS
http://www.utia.cas.cz/ZS
[2]Center for Applied Cybernetics
Faculty of Electrical Engineering, CTU
http://www.c-a-k.cz

Our poster introduces an innovative RLS Lattice filter implementation for FPGAs. The signal processing applications typically require wide numeric range, and that poses a problem when using an FPGA implementation. Our approach is based on arithmetic using logarithmic numeric representation (LNS). The test application - an adaptive noise canceller - has been optimized for the Xilinx Virtex devices. It consumes roughly 70% of all logic resources of the XCV800 device and all block memory cells. The filter orders up to 252 at 7 kHz sampling frequency have been achieved using the 19-bit LNS arithmetic. The maximum performance of the application is about 70 MFLOPs including i/o conversions. It has been shown that the custom filter design can be 10 times faster than commonly used solutions. This work was supported by the Ministry of Education of the Czech Republic under Project No. LN00B06 .

Wireless Sensor Networks: a Power-Scalable Motion Estimation IP for Hybrid Video Coding

Federico Quaglio, Maurzio Martina, Fabrizio Vacca, Guido Masera, Andrea Molino, Gianluca Piccinini, Maurizio Zamboni
Dipartimento di Elettronica, Politecnico di Torino, Corso Duca degli Abruzzi 24, I-10129 TORINO, ITALY
quaglio@mbox.vlsilab.polito.it
http://www.vlsilab.polito.it

Wireless Sensor Networks are an emerging phenomenon in the research community. The design and development of network architectures and nodes implementation are fostering many research activities. Due to their wide application fields and pervasive employment possibilities, the investigation of novel classes of wireless sensor nodes is of great concern. In this paper we presented a novel Power-Scalable Motion Estimation IP suitable for video-surveillance over Wireless Sensor Networks. The proposed architecture can achieve low dynamic power, good video quality and reasonable frame-rates. Moreover, it can operate on different video format and can be reconfigured both off-line and on-line. Further researches have to be accomplished in order to reduce the internal critical path, achieving better maximum frequencies performances. Moreover, new FPGA architectures must be exploited searching low-static-power devices suitable for an actual implementation of our IP on a WSN.

Design Strategies and Modified Descriptions to Optimize Cipher FPGA Implementations: Fast and Compact Results for DES and Triple-DES

Gaël Rouvroy, Francois-Xavier Standaert, Jean-Jacques Quisquater, Jean-Didier Legat
Université catholique de Louvain
Louvain-La-Nouve, Belgium
{rouvroy, standaert, quisquater, legat}@dice.ucl.ac.be

We propose a new mathematical DES description that allows optimized implementations. It also provides the best DES and triple-DES FPGA implementations known in term of ratio throughput/area, where area means the number of FPGA slices used. First, we get a less resource consuming unrolled DES implementation that works at data rates of 21.3 Gbps (333 MHz), using VIRTEX II technology. In this design, the plaintext, the key and the mode (encryption/decrytion) can be changed on a cycle-by-cycle basis with no dead cycles. In addition, we also propose sequential DES and triple-DES designs that are currently the most efficient ones in term of resources used as well as in term of throughput. Based on our DES and triple-DES results, we also set up conclusions for optimized FPGA design choices and possible improvement of cipher implementations with a modified structure description.

A Physical Retiming Algorithm for Field Programmable Gate Arrays

Peter Suaris, Dongsheng Wang, Pei-Ning Guo, Nan-Chi Chou
Mentor Graphics Corporation, 8005 Boeckman Road, Wilsonville, OR 97007
{peter_suaris, dongsheng_wang, pn_guo, nanchi_chou}@mentorg.com

In this paper, we present a physical retiming algorithm for sequential circuits implemented in field programmable gate arrays (FPGAs). This algorithm can speed up the sequential circuits by reducing delay of all critical paths with negative slacks. By taking advantage of the physical information provided by placed circuits, this algorithm integrates two operations: retiming and register duplication. Retiming moves registers across combinational components. Register duplication moves registers across interconnect. Circuit functionality remains unchanged while performing these operations. A concept of timing budget is proposed to guide register moves. An accurate delay estimator for FPGA designs is developed to make register moves much more acceptable compared to conventional retiming technique. Experiments show that the physical retiming algorithm proposed in this paper can effectively reduce the critical path delays, increasing maximum frequency by 13.54%, on average, and improving by 59.54%, on average, the minimum slack for sequential circuits implemented in FPGAs.

Application-Dependent Testing of FPGAs for Bridging Faults

Mehdi Baradaran Tahoori

Center for Reliable Computing, Stanford University, Stanford, CA 94305
mtahoori@crc.stanford.edu

A new technique is presented for testing for bridging faults in the interconnects of an arbitrary design implemented in an FPGA. The configuration of the routing resources used in the original design remains unchanged in the test configurations. Only the logic blocks used in the design are reprogrammed in order to implement single-term functions, logic functions with only one minterm or one maxterm. As shown by formal proofs, all activated faults are detected when single-term functions and appropriate test vectors are used. As presented in the paper, only two test configurations are necessary to detect all bridging faults, achieving 100% fault coverage. An algorithm is presented for test configuration and test vector generation for the entire FPGA. Also, test vector and configuration generation problem is systematically converted to a satisfiability problem, and state of the art SAT-solvers are exploited for automatic test vector and configuration generation.

A High Resolution Diagnosis Technique for Open and Short Defects in FPGA Interconnects

Mehdi Baradaran Tahoori

Center for Reliable Computing, Stanford University, Stanford, CA 94305
mtahoori@crc.stanford.edu

A two-step diagnosis flow, coarse-grain and fine-grain, is presented in order to identify a faulty element in the FPGA interconnects. The fault models used for interconnect are open, resistive-open, and bridging fault. The coarse-grain phase identifies the faulty net, the routing between two consecutive sequential elements in the FPGA. This phase is performed by just post-processing tester results for the test configurations used for interconnect testing. During the fine-grain step, the faulty net is rerouted without using some of the resources used in the original routing. The faulty element, programmable switch or wire segment, is uniquely identified based on the tester output for the rerouted configurations. Effective search methods are exploited in order to minimize the number of test configurations and the total diagnosis time. This method is implemented on real FPGA chips and verified using hardware fault emulation.

A Four-bit Full Adder Implemented on Fast SiGe FPGAs with Novel Power Control Scheme

K. Zhou, M. Chu, C. You, J.-R. Guo, Channakeshav, J. Mayega, B.S. Goda, R.P. Kraft, J.F. McDonald

Center for Integrated Electronics, Rensselaer Polytechnic Institute
{zhouk,chum2}@rpi.edu, mcdonald@unix.cie.rpi.edu

The low operating speed of current CMOS Field Programmable Gate Arrays (FPGAs), i.e., 10-220 MHz, has prevented their use in high-speed digital applications. With the advent of IBM Silicon Germanium (SiGe) 7HP technology, designers have been able to design FPGAs operating in the gigahertz range. This paper is going to elaborate on the implementation of a 4-bit ripple-carry full adder (FA) on the new SiGe FPGA with new architectures and a novel power management strategy. The 1-bit FA can be realized in three Configurable Logic Blocks (CLBs). Apart from these, the FA can operate in multiple modes: FAST, NON-CRITICAL, SLOW and OFF. The propagation delays of the 1-bit FA and 4-bit ripple-carry FA are 240 ps and 675 ps respectively in the FAST mode. All the simulation and layouts were done using Cadence 4.4.6 and IBM SiGe 7HP design kit version 1.1.1.0.