21.2 An Energy-Efficient IEEE 1363-based Reconfigurable Public-Key Cryptography Processor

James Goodman[1], Anantha P. Chandrakasan[2]

[1]Chrysalis-ITS, Ottawa, Canada
[2]Massachusetts Institute of Technology, Cambridge, MA

One of the main consumer concerns, and outstanding issues, confronting the widespread acceptance of wireless networks is security, as demonstrated by the widespread fraud that plagues the cellular phone industry today. The portability of these systems requires an energy-efficient implementation to maximize battery lifetime. However, the lack of a coherent global security standard requires algorithm agility to allow the user to dynamically adjust the type and strength of cryptographic algorithms being used to ensure compatibility.

Conventional software-based solutions are computationally inefficient for the algorithms being used in public-key cryptography, leading to flexible solutions that are energy inefficient due to the high control overhead per instruction. For example, a single public-key operation in software consumes the same amount of energy as encrypting 10s of Mb of data using symmetric algorithms. Significant gains can be made by migrating to programmable logic (e.g., FPGAs) which is more efficient than software for bit-level processing, leading to improved energy efficiency. Unfortunately, FPGAs suffer from a great deal of overhead due primarily to their programmable interconnect which accounts for approximately 65% of the total energy dissipation of the device [1].

For a specific domain of operation (e.g., public key cryptography), it is possible to provide the required level of algorithm agility by using a limited amount of reconfigurability which avoids the overhead associated with generic programmable logic. The resulting Domain Specific Reconfigurable Cryptographic Processor (DSRCP) represents a prototype implementation for the domain of public key cryptography. The processor ISA is based upon the recently adopted IEEE 1363-2000 Public Key Cryptography Standard. The DSRCP is the first reported hardware-based solution that performs conventional arithmetic, modular integer arithmetic, binary Galois Field arithmetic (i.e., GF(2^n)), and elliptic curve arithmetic over GF(2^n). In addition, it does so in an energy efficient manner.

The processor is dynamically scalable using the SET_LENGTH(length$_{9:0}$) instruction to set the current width of the processor and configure the I/O interface to schedule the correct size of data transfers into and out of the processor. The DSRCP can compute a complete set of modular arithmetic functions given any valid odd modulus ranging from 8b to 1024b. The processor is also capable of performing exponentiation, multiplication, inversion, and addition over GF(2^n), 8 n 1024, for any valid n-th degree field polynomial. Similarly, the processor is capable of performing elliptic curve point addition, doubling and multiplication operations over any valid elliptic curve of characteristic 2 of the form $y^2 + xy = x^3 + ax^2 + b$. This flexibility enables the DSRCP to be as algorithm-agile as software-based solutions in the desired design space of public-key cryptography, while providing much of the efficiency of a dedicated hardware solution.

The processor architecture is shown in Figure 21.2.1. The dominant feature, both in terms of area and complexity, is the reconfigurable datapath. The datapath (Figure 21.2.2) consists of a register file, conventional integer add/subtract unit, fast magnitude comparator, and reconfigurable logic block with embedded local memory. The eight-word register file is a one-write, two-read ported memory implemented using edge-triggered flip-flops. The adder utilizes a fast hybrid carry-bypass/skip adder [2] that maps to an efficient bit-sliced implementation. Magnitude comparisons use a fast tree-based magnitude comparator. The comparator utilizes a pre-conditioning circuit (Figure 21.2.3) to generate pairwise greater-than and equal-to flags that are then encoded to determine a full 1024b comparison in a single cycle. Clock gating is used extensively within the processor to minimize unnecessary switched capacitance by dynamically shutting down parts of the processor that are not currently in use. Gating is performed along 32-bit boundaries using the scheme depicted in Figure 21.2.4, allowing the processor to automatically disable its datapath to accommodate the current operand sizes being processed.

The reconfigurable logic block performs modular/GF(2^n) multiplication and GF(2^n) inversion by using the bitslice circuit of Figure 21.2.5 and exploiting the fact that a full-adder circuit sum output is a three input addition over GF(2). Reconfigurability is achieved using transmission-gate multiplexors which allow single-cycle reconfigurability. Modular multiplication is implemented using a radix-2 Montgomery Multiplication algorithm. GF(2^n) multiplication is implemented using a n-cycle MSB-first approach that performs a conventional double-and-add partial product accumulation with reduction via the programmable field polynomial. GF(2^n) inversion is performed using an optimized parallelized implementation of the binary extended euclidean algorithm that is capable of performing a GF(2^n) inversion and multiplication concurrently in at most 4n cycles, and on average 3.3n cycles. The processor utilizes constant execution time exponentiation and point multiplication algorithms to provide resistance to both timing attacks and simple power analysis [3]. A summary of performance of the various processor functions is shown in Figure 21.2.6 and Figure 21.2.7. Local memory is used to store both operating parameters (e.g., moduli and field polynomials) and temporary values generated during execution, allowing locality to be exploited in order to minimize power consumption. A regular, bit-sliced implementation is used to achieve a area-efficient implementation that minimizes global interconnect.

The processor is fabricated in a 0.25µm CMOS technology with 5 levels of metallization. Figure 21.2.9 depicts a micrograph of the processor whose core contains 880k devices and measures 2.9x2.9mm^2. The datapath consists of 1024 processing bit-slices, each of which measures 30x150µm^2. At 50MHz the processor operates at 2V supply and consumes at most 75 mW (the power consumption of the processor is a function of both the instruction being executed and the operand sizes processed). In ultra low power mode (3MHz at 0.7V) the processor consumes at most 525µW.

Figure 21.2.8 demonstrates the energy efficiency of the DSRCP relative to optimized software-based implementations on the StrongARM SA-1100 and programmable logic-based implementations ([4], [5]) on Xilinx XC4000 parts. The DSRCP is approximately 2 to 3 orders of magnitude more energy efficient than both software and programmable logic based solutions, while providing the same degree of flexibility and algorithm-agility.

Acknowledgements:
The authors thank the National Semiconductor Corp. for funding and fabrication facilities for this work. This research is sponsored by the Defense Advanced Research Project Agency (DARPA) Power Aware Computing/Communication Program and the Air Force Research Laboratory, Air Force Materiel Command, USAF, under agreement number F30602-00-2-0551.

References:
[1] E. Kusse and J. Rabaey, "Low-energy embedded FPGA structures," Proceedings of ISPLED'98, pp. 155-160.
[2] A. Satoh et. al., "A high-speed small RSA encryption LSI with low power dissipation," Proceedings of the 1997 Information Security Workshop.
[3] P. Kocher, "Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems," Proceedings of CRYPTO'96, pp. 104-113.
[4] T. Blum and C. Paar, "Modular exponentiation on reconfigurable hardware," 14th IEEE Symposium on Computer Arithmetic, 1999.
[5] M. Rosner, Elliptic Curve Cryptosystems in Reconfigurable Hardware, Master's Thesis, Worcester Polytechnic Institute, 1998.

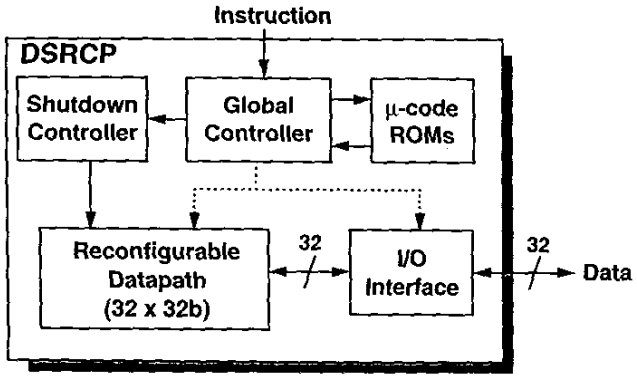

Figure 21.2.1: Top-level system architecture.

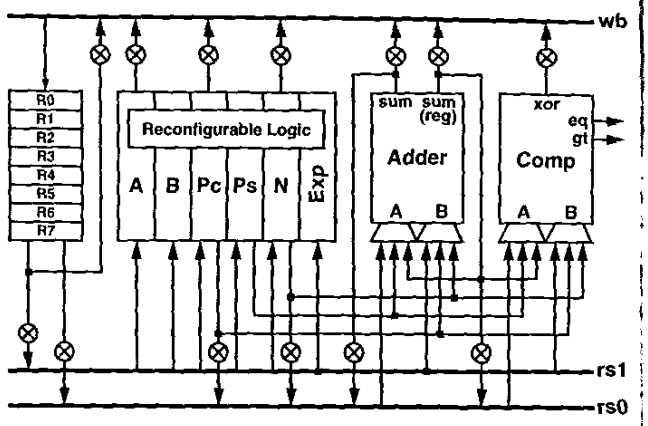

Figure 21.2.2: Reconfigurable data path architecture block diagram.

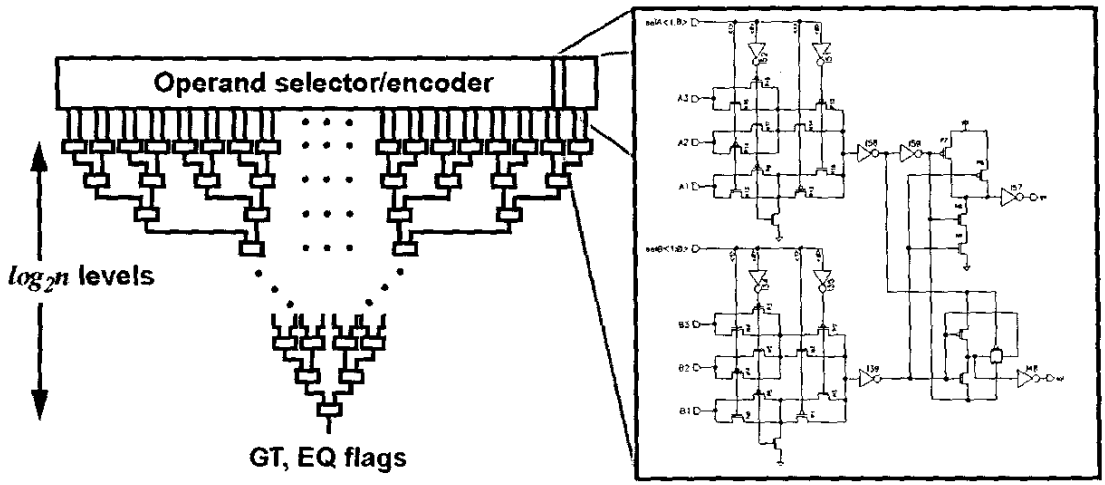

Figure 21.2.3: Tree-based magnitude comparator topology used in the DSRCP.

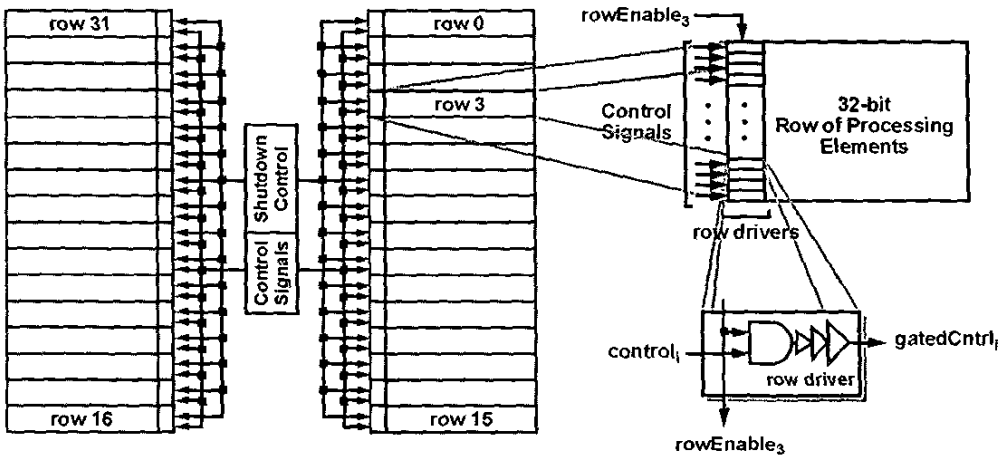

Figure 21.2.4: Shutdown circuitry used in the DSRCP.

Continued on Page 461

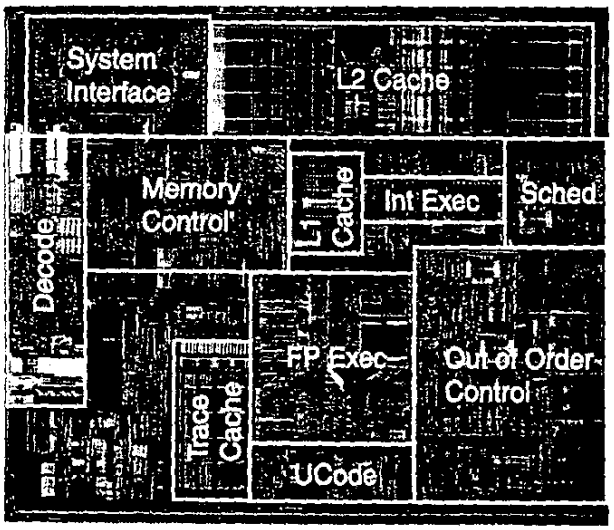

Figure 20.6.6: Micrograph.

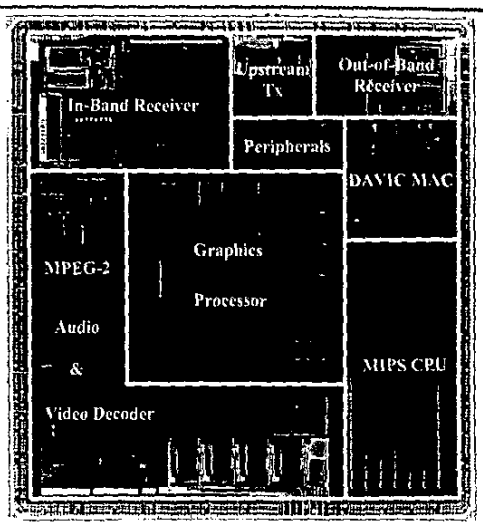

Figure 21.1.7: Die micrograph.

Process Technology	0.22μm CMOS, 5-layer metal, single-poly
Maximum Clock Rate	81MHz
Package	456-PBGA
Power Dissipation	3.0W at 3.3V/2.5V
Transistor Count	14.6M
Die Size	9.4x10 mm

Figure 21.1.8: Chip physical specifications summary.

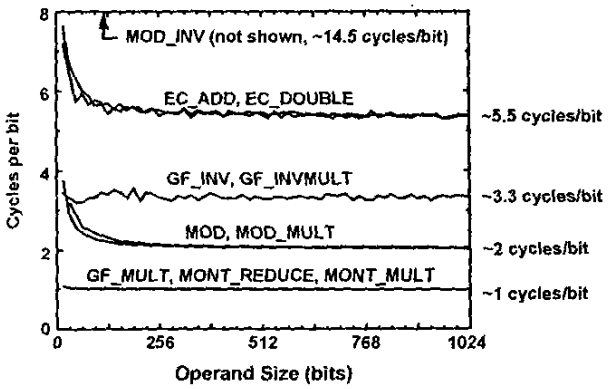

Figure 21.2.6: Performance of several OSRCP arithmetic instruction.

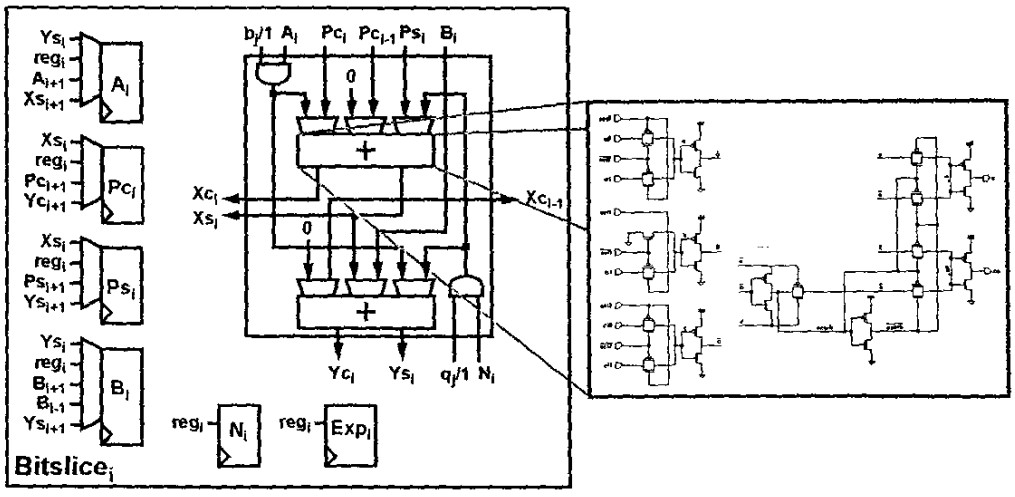

Figure 21.2.5: Reconfigurable logic bitslice.

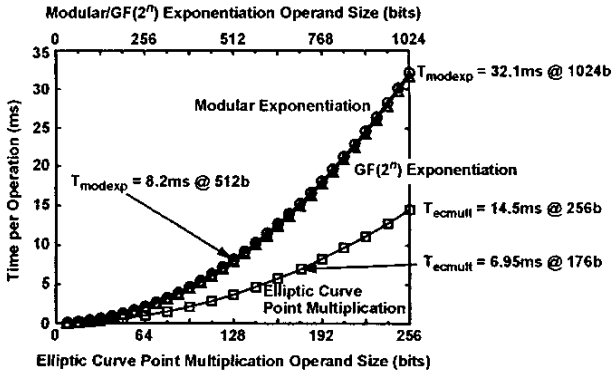

Figure 21.2.7: Performance of various crytographic primitives using the DSRCP at 50MHz.

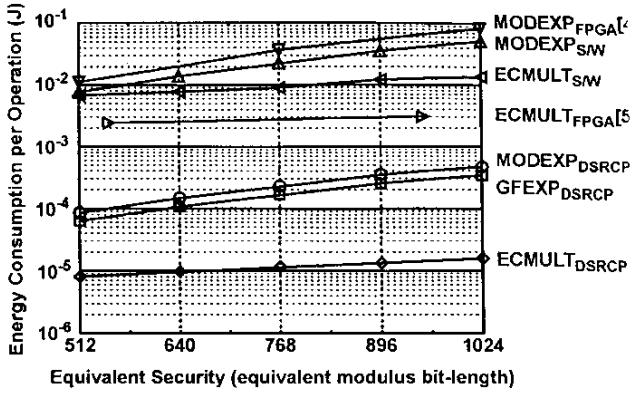

Figure 21.2.8: Comparison of energy consumption per operation for software, FPGA and DSRCP.

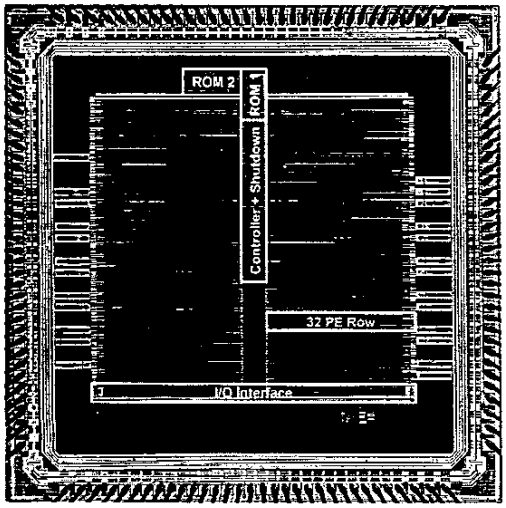

Figure 21.2.9: DSRCP die micrograph.

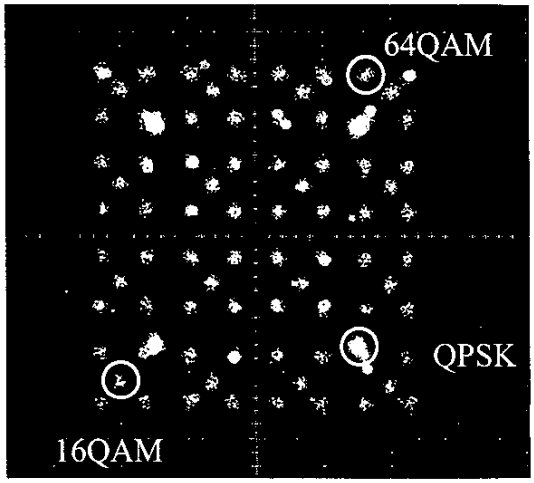

Figure 21.4.6: Constellation diagram for stream transmitted with three moduation schemes; QPSK, 16QAM and 64QAM.

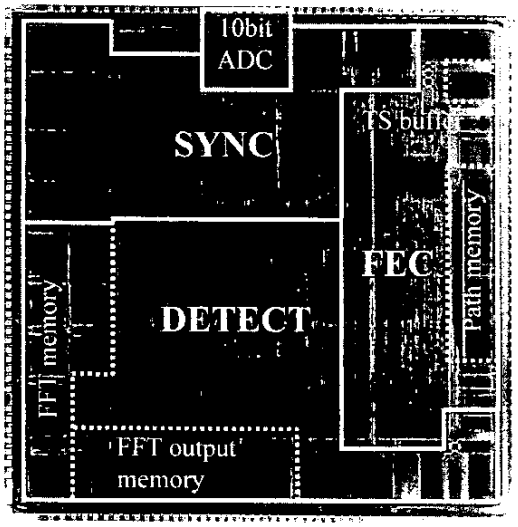

Figure 21.4.7: Chip micrograph.

Technology	CMOS 0.18$_{\mu}$m 1.8V core, 3.3V I/O		
Nominal Clock Frequency (= sampling frequency)	20 MHz		
Package	160-PQFP		
Equ. Gate Count (core)	431,000 = 100.0 %		
- active in reception	416,000 = 96.5 %		
- active in transmission	79,000 = 18.0 %		
- of which equalizer	270,000 = 62.6 %		
- of which FFT	42,000 = 9.7 %		
- of which RAMs	78,000 = 18.1 %		
Die Size	20.8mm^2		
Measured power consumption in IEEE mode @ 20MHz	3.3V I/O	1.8V core	total
- transmission mode	156mW	43mW	199mW
- reception mode	66mW	146mW	212mW
- programming mode	35mW	81mW	116mW

Figure 21.5.5: IC key figures.

0-7803-6608-5/$10.00 ©2001 IEEE