

CTL: A Platform-Independent Crypto Tools Library Based on Dataflow Programming Paradigm^{*}

Junaid Jameel Ahmad¹, Shujun Li^{1,2}, Ahmad-Reza Sadeghi^{3,4}, and Thomas Schneider³

¹ University of Konstanz, Germany ^{**}

² University of Surrey, UK

³ TU Darmstadt, Germany

⁴ Fraunhofer SIT, Germany

Abstract. The diversity of computing platforms is increasing rapidly. In order to allow security applications to run on such diverse platforms, implementing and optimizing the same cryptographic primitives for multiple target platforms and heterogeneous systems can result in high costs. In this paper, we report our efforts in developing and benchmarking a platform-independent Crypto Tools Library (CTL). CTL is based on a dataflow programming framework called Reconfigurable Video Coding (RVC), which was recently standardized by ISO/IEC for building complicated reconfigurable video codecs. CTL benefits from various properties of the RVC framework including tools to 1) simulate the platform-independent designs, 2) automatically generate implementations in different target programming languages (e.g., C/C++, Java, LLVM, and Verilog/VHDL) for deployment on different platforms as software and/or hardware modules, and 3) design space exploitation such as automatic parallelization for multi- and many-core systems. We benchmarked the performance of the SHA-256 implementation in CTL on single-core target platforms and demonstrated that implementations automatically generated from platform-independent RVC applications can achieve a runtime performance comparable to reference implementations manually written in C and Java. For a quad-core target platform, we benchmarked a 4-adic hash tree application based on SHA-256 that achieves a performance gain of up to 300% for hashing messages of size 8 MB.

Keywords: Crypto Tools Library (CTL), Reconfigurable Video Coding (RVC), dataflow programming, reconfigurability, platform independence, multi-core.

^{*} Extended edition of this paper is available at <http://eprint.iacr.org/2011/679>.

^{**} Junaid Jameel Ahmad and Shujun Li were partially supported by the Zukunftscolleg of the University of Konstanz, Germany, which is part of the “Excellence Initiative” program of the DFG (German Research Foundation). The first author would like to thank International Financial Cryptography Association (IFCA) and Google Research for awarding him the “Google Student Award”, which covered his registration fees to FC 2012 and associated workshops.

1 Introduction

Nowadays we are living in a fully digitized and networked world. The ubiquitous transmission of data over the open network has made security one of the most important concerns in almost all modern digital systems, being privacy another. Both security and privacy concerns call for support from applied cryptography. However, the great diversity of today’s computing hardware and software platforms is creating a big challenge for applied cryptography since we need building blocks that should ideally be reused at various platforms without reprogramming. For instance, a large-scale video surveillance system (like those we have already been seeing in many big cities) involves many different kinds of hardware and software platforms: scalar sensors, video sensors, audio sensors, mobile sensors (e.g. mobile phones), sensor motor controller, storage hub, data sink, cloud storage servers, etc. [11]. Supporting so many different devices in a single system or cross the boundary of multiple systems is a very challenging task. Many cryptographic libraries have been built over the years to partly meet this challenge, but most of them are written in a particular programming language (e.g. C, C++, Java and VHDL) thus their applications are limited in nature. While it is always possible to port a library written in one language to the other, the process requires significant human involvement on reprogramming and/or re-optimization, which may not be less easier than designing a new library from scratch.

In this paper, we propose to meet the above-mentioned technical challenges by building a platform-independent⁵ library based on a recently-established ISO / IEC standard called RVC (Reconfigurable Video Coding) [33, 34]. Unlike its name suggests, the RVC standard offers a general development framework for all data-driven systems including cryptosystems, which is not surprising because video codecs are among the most complicated data-driven systems we can have. The RVC framework follows the dataflow paradigm, and enjoys the following nice features at the level of programming language: *modularity*, *reusability*, *re-configuration*, *code analyzability* and *parallelism exploitability*. Modularity and reusability help to simplify the design of complicated programs by having functionally separated and reusable computational blocks; reconfigurability makes reconfiguration of complicated programs easier by offering an interface to configure and replace computational blocks; code analyzability allows automatic analysis of both the source code and the functional behavior of each computational block so that code conversion and program optimization can be done in a more systematic manner. The automated code analysis enables to conduct a fully-/semi-automated design-space exploitation to find critical paths and/or parallel data-flows, which suggests different optimization refactorings (merging or splitting) of different computational blocks [43], and/or to achieve concurrency

⁵ In the context of MPEG RVC framework, the word “platform” has a broader meaning. Basically, it denotes any computing environment that can execute/interpret code or compile code to produce executable programs, which includes both hardware and software platforms and also hybrid hardware-software systems.

by mapping different computational blocks to different computing resources [20]. In contrast to the traditional sequential programming paradigm, the dataflow programming paradigm is ideally suited for such optimizations thanks to its data-driven nature as described next.

The dataflow programming paradigm, invented in the 1960s [61], allows programs to be defined as a directed graph in which the nodes correspond to computational units and edges represent the direction of the data flowing among nodes [25, 40]. The modularity, reusability and reconfigurability are achieved by making each computational unit’s functional behavior independent of other computational units. In other words, the only interface between two computational units is the data exchanged. The separation of functionality and interface allows different computational units to run in parallel, thus easing parallelism exploitation. The dataflow programming paradigm is suited ideally for applications with a data-driven nature like signal processing systems, multimedia applications, and as we show in this paper also for cryptosystems.

Our Contributions: In this paper, we present the Crypto Tools Library (CTL) as the first (to the best of our knowledge) open and platform-independent cryptographic library based on a dataflow programming framework (in our case the RVC framework). In particular, the CTL achieves the following goals:

- **Fast development/prototyping:** By adapting the dataflow programming paradigm the CTL components are inherently *modular*, *reusable*, and easily *reconfigurable*. These properties do not only help to quickly develop/prototype security algorithms but also make their maintenance easier.
- **Multiple target languages:** The CTL cryptosystems are programmed only once, but can be used to automatically generate source code for multiple programming languages (C, C++, Java, LLVM, Verilog, VHDL, XLIM, and PROMELA at the time for this writing⁶).
- **Automatic code analyzability and optimization:** An automated design-space exploitation process can be performed at the algorithmic level, which can help to optimize the algorithmic structure by refactoring (merging or splitting) selected computational blocks, and by exploiting multi-/many-core computing resources to run different computational blocks in parallel.
- **Hardware/Software co-design:** Heterogenous systems involving software, hardware, and various I/O devices/channels can be developed in the RVC framework [62].
- **Adequate run-time performance:** Although CTL cryptosystems are highly abstract programs, the run-time performance of automatically synthesized implementations is still adequate compared to non-RVC reference implementations.

In this paper, along with the development of the CTL itself, we report some performance benchmarking results of CTL that confirm that the highly abstract

⁶ More code generation backends are to be developed in the future, especially OpenCL for GPUs.

nature of the RVC code does not compromise the run-time performance. In addition, we also briefly discuss how different key attributes of the RVC framework can be used to develop different cryptographic algorithms and security applications.

Outline: The rest of the paper is organized as follows. In Sec. 2 we will give a brief overview of related work, focusing on a comparison between RVC and other existing dataflow solutions. Sec. 3 gives an overview of the building blocks of the RVC framework and Sec. 4 describes the design principles of CTL and the cryptosystems that are already implemented. In Sec. 5, we benchmark the performance of SHA-256 implemented in CTL on a single-core machine and a quad-core one. In Sec. 6, we conclude the paper by giving directions for future work.

2 Related Work

Many cryptographic libraries have been developed over the years (e.g., [16,24,30,41,46,56,57,63,64]), but very few can support multiple programming languages. Some libraries do support more than one programming language, but often in the form of separate sets of source code and separate programming interfaces/APIs [63], or available as commercial software only [8,41]. There is also a large body of optimized implementations of cryptosystems in the literature [17,18,21,44,45,55,67], which normally depend even more on the platforms (e.g., the processor architecture and/or special instruction sets [28,45,66,67]).

Despite being a rather new standard, the RVC framework has been successfully used to develop different kinds of data-driven systems especially multimedia (video, audio, image and graphics) codecs [12–14,19,35] and multimedia security applications [10]. In [10], we highlighted some challenges being faced by developers while building multimedia security applications in imperative languages and discussed how those challenges can be addressed by developing multimedia security applications in the RVC framework. In addition, we presented three multimedia security applications (joint H.264/MPEG-4 video encoding and decoding, joint JPEG image encoding and decoding and compressed domain JPEG image watermark embedding and detecting) developed using the CTL cryptosystems and the RVC implementations of H.264/MPEG-4 and JPEG codecs. Considering the focus of that paper, we only used and briefly summarized CTL. In this paper, we give a detailed discussion on CTL, its design principles, features and benefits, and performance benchmarking results.

The wide usage of RVC for developing multimedia applications is not the only reason why we chose it for developing CTL. A summary of advantages of RVC over other solutions is given in Table 1 (this is an extension of the table in [10]). We emphasize that this comparison focuses on the features relevant to achieve the goals of CTL, so it should not be considered as an exhaustive overview of all pros and cons of the solutions compared.

Table 1: Comparison of RVC framework with other candidate solutions. Candidates with similar characteristics are grouped together. These categories include 1) high-level specification languages for hardware programming languages, 2) frameworks for hardware/software co-design, 3) commercial products, and 4) other cryptographic libraries. The columns in the table represent the following features: A) high-level (abstract) modeling and simulation; B) platform independence; C) code analyzability (i.e., semi-automated design-space exploitation); D) hardware code generation; E) software code generation; F) hardware/software co-design; G) supported target languages; H) open-source or free implementations; I) international standard.

Cat.	Candidate	A	B	C	D	E	F	G	H	I
	RVC	✓	✓	✓	✓	✓	✓	C, C++, Java, LLVM, Verilog, VHDL, XLIM, PROMELA	✓	✓
1	Handel-C [39]	✗	✗	✗	✓	✗	✗	VHDL	✗	✗
	ImpulseC [15]	✗	✗	✗	✓	✗	✓	VHDL	✗	✗
	Spark [29]	✗	✗	✗	✓	✗	✓	VHDL	✗	✗
2	BlueSpec [49]	✓	✗	✓	✓	✓	✗	C, Verilog	✗	✗
	Daedalus [65]	✓	✓	✓	✓	✓	✓	C, C++, VHDL	✓	✗
	Koski [38]	✓	✓	✓	✓	✓	✓	C, XML, VHDL	✗	✗
	PeaCE [31]	✓	✓	✓	✓	✓	✓	C, C++, VHDL	✓	✗
3	CoWare [58]	✓	✓	✗	✓	✓	✓	C, VHDL	✗	✗
	Esterel [1]	✗	✓	✗	✓	✓	✗	C, VHDL	✓	✗
	LabVIEW [3]	✓	✓	✓	✗	✗	✗	-	✗	✗
	Simulink [4]	✓	✓	✓	✓	✓	✗	C, C++, Verilog, VHDL	✗	✗
	Synopsys System Studio [7]	✓	✓	✓	✓	✓	✓	C++, SystemC, SystemVerilog	✗	✗
4	CAO [9,47]	✓	✓	✗	✗	✓	✗	C, x86-64 assembly, ARM	✗	✗
	Cryptol [8,41]	✓	✓	✓	✓	✓	✗	C, C++, Haskell, VHDL, Verilog	✗	✗

3 Reconfigurable Video Coding (RVC)

The RVC framework was standardized by the ISO/IEC (via its working group JTC1 / SG29 / WG11, better known as MPEG – Motion Picture Experts Group [48]) to meet the technical challenges of developing more and more complicated video codecs [33,34]. One main concern of the MPEG is how to make video codecs more reconfigurable, meaning that codecs with different configurations (e.g., different video coding standards, different profiles and/or levels, different system requirements) can be built on the basis of a single set of platform-independent building blocks. To achieve this goal, the RVC standard defines a framework that covers different steps of the whole life cycle of video codec development.

The RVC community has developed supporting tools [2, 5, 6] to make the RVC framework not only a standard, but also a real development environment.

While the RVC framework is developed in the context of video coding, it is actually a general-purpose framework that can model any data-driven applications such as cryptosystems. It allows developers to work with a single platform-independent design at a higher level of abstraction while still being able to generate multiple editions of the same design that target different platforms like embedded systems, general-purpose PCs, and FPGAs. In principle, the RVC framework also supports hardware-software co-design by converting parts of a design into software and other parts into hardware. Additionally, the RVC framework is based on two languages that allow automatic code analysis to facilitate large-scale design-space exploitation like enhancing parallelism of implementations running on multi-core and many-core systems [14, 20, 43].

The RVC standard is composed of two parts: MPEG-B Part 4 [34] and MPEG-C Part 4 [33]. MPEG-B Part 4 specifies the dataflow framework for designing and/or reconfiguring video codecs, and MPEG-C Part 4 defines a video tool library that contains a number of Functional Units (FUs) as platform-independent building blocks of MPEG standard compliant video codecs [33]. To support the RVC dataflow framework, MPEG-B Part 4 specifies three different languages: a dataflow programming language called RVC-CAL for describing platform-independent FUs, an XML dialect called FNL (FU Network Language) for describing connections between FUs, and another XML dialect called RVC-BSDL for describing the syntax format of video bitstreams. RVC-BSDL is not involved in this work, so we will not discuss it further.

The real core of the RVC framework is RVC-CAL, a general-purpose dataflow programming language for specifying platform-independent FUs. RVC-CAL is a subset of an existing dataflow programming language CAL (Caltrop Actor Language) [26]. In RVC-CAL, FUs are implemented as actors containing a number of fireable actions and internal states. In the RVC-CAL's term, the data exchanged among actions are called tokens. Each actor can contain both input and output port(s) that receive input token(s) and produce output token(s), respectively. Each action may fire depending on four different conditions: 1) input token availability; 2) guard conditions; 3) finite-state machine based action scheduling; 4) action priorities. In RVC-CAL, actors are the basic functional entities that can run in parallel, but actions in an actor are atomic, meaning that only one action can fire at one time. This structure gives a balance between modularity and parallelism, and makes automatic analysis of actor merging/splitting possible.

Figure 1 illustrates how an application can be modeled and how target implementations can be generated with the RVC framework. At the design stage, different FUs (if not implemented in any standard library) are first written in RVC-CAL to describe their I/O behavior, and then an FU network is built to represent the functionality of a whole application. The FU network can be built by simply connecting all FUs involved graphically via a supporting tool called Graphiti Editor [2], which translates the graphical FU network description into

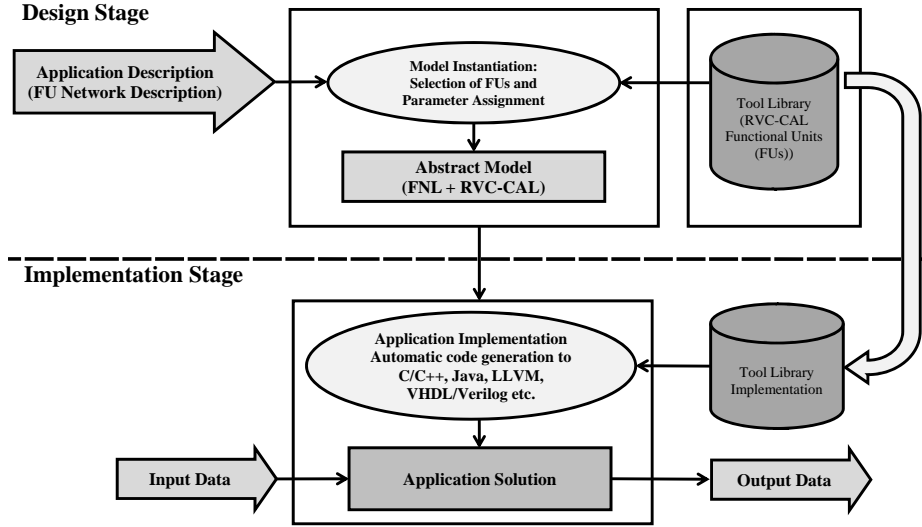


Fig. 1: Process of application implementation generation in the RVC framework.

a textual description written in FU Network Language (FNL). The FUs and the FU network are instantiated to form an abstract model. This abstract model can be simulated to test its functionality without going to any specific platform. Two available supporting tools allowing the simulation are OpenDF [5] and ORCC [6]. At the implementation stage, the source code written in other target programming languages can be generated from the abstract application description *automatically*. OpenDF includes a Verilog HDL code generation backend, and ORCC contains a number of code generation backends for C, C++, Java, LLVM, VHDL, XLIM, and PROMELA. ORCC is currently more widely used in the RVC community and it is also the choice of our work reported in this paper.

4 Crypto Tools Library (CTL)

Crypto Tools Library (CTL) is a collection of RVC-CAL actors and XDF networks for cryptographic primitives such as block ciphers, stream ciphers, cryptographic hash functions and PRNGs (see Sec. 4.2 for a list of currently implemented algorithms). Being an open project, the source code and documentation of CTL is available at <http://www.hooklee.com/default.asp?t=CTL>.

As mentioned in Sec. 1, most existing cryptographic libraries are developed based on a single programming language (mostly C/C++ or Java) that can hardly be converted to other languages. In contrast, CTL is a platform-independent solution whose source code is written in RVC-CAL and FNL that can be automatically translated into multiple programming languages (C, C++, Java, LLVM, Verilog, VHDL, XLIM, PROMELA). More programming languages can be supported by developing new code generation tools for RVC applications.

4.1 Design Principles

The CTL is developed by strictly following the specifications/standards defining the implemented cryptosystems. For block ciphers, both enciphers and decipherers are implemented so that a complete security solution can be built. When it is possible, the CTL FUs are designed to exploit inherent parallelism in the implemented cryptosystems. For instance, for block ciphers based on multiple rounds, the round number is also transmitted among different FUs so that encryption/decryption of different blocks can be parallelized.

The CTL is designed so that different cryptosystems can share common FUs. We believe that this can help enhance code reusability and ease reconfigurability of the CTL cryptosystems. In addition, CTL includes *complete* solutions (e.g., both encipher and decipher) of the implemented cryptosystems, normally a set of CAL and XDF files.

4.2 Cryptosystems Covered

CTL contains some standard and frequently used cryptosystems. In the following, we list the cryptosystems currently implemented in CTL. The correctness of all cryptosystems has been validated using the test vectors given in the respective standards.

- Block Ciphers:
 - AES-128/192/256 [51],
 - DES [50] and Triple DES [50, 52],
 - Blowfish [59],
 - Modes of operations: CBC, CFB, OFB, CTR.
- Stream Ciphers: ARC4 [60] and Rabbit [23].
- Cryptographic hash functions: SHA-1, SHA-2 (SHA-224, SHA-256) [53].
- PSNRs: 32-bit and 64-bit LCG [60] and LFSR-based PRNG [60].

CTL also includes some common utility FUs (e.g., multiplexing/demultiplexing of dataflows, conversion of bytes to bits and vice versa etc.) that are shared among different cryptosystems and can also find applications in non-cryptography systems. Due to the space limitation, we refer the reader to the extended edition of this paper for a list of the utility FUs and more discussions of the cryptosystems implemented in CTL.

5 Performance Benchmarking of CTL

Previous work has demonstrated that the RVC framework can outperform other sequential programming languages in terms of implementing highly complex and highly parallelizable systems like video codecs [19]. However, there are still doubts about if the high-level abstraction of RVC-CAL and the automated code generation process may compromise the overall performance to some extent at the platform level. In this section, we clarify those doubts by showing that the

Table 2: Configuration of the test machine.

Machine	Hardware and Operating System Details
Desktop PC:	<ul style="list-style-type: none"> – Model: HP Centurion – CPU: Intel(R) Core(TM)2 Quad CPU Q9550 2.83GHz – Memory: 8GB RAM – OS1: Windows Vista Business with Service Pack 2 (64-bit Edition) – OS2: Ubuntu Linux (Kernel version: 2.6.27.11)

automatically generated implementations from a typical RVC-based application can usually achieve a performance comparable to manually-written implementations in the target programming language. This was verified on AES and SHA-256 applications in CTL. In this section, we take SHA-256 as an example to show how we did the benchmarking on a single-core machine and a quad-core one. The main purpose of getting the quad-core machine involved is to show how easy one can divide an FU network and map different parts to different cores to make a better use of the computing resources. In the given example, the partitioning and mapping were both done manually, but they can be automated for large applications thanks to the code analyzability of RVC-CAL.

Run-Time Performance Metric We ran our experiments on both Microsoft Windows and Linux OSs (see Table 2 for details). Both operating systems support high-resolution timers to measure time in nanoseconds. More specifically, we used the `QueryPerformanceCounter()` and `QueryPerformanceFrequency()` functions (available from Windows API) on Windows, and the `clock_gettime()` and `clock_getres()` functions with `CLOCK_MONOTONIC` clock (available from the Higher Resolution Timer [22] package) on Linux. In addition, to circumvent the caching problem, we conducted 100 independent runs (with random input data) of each configuration and used the average value as the final performance metric.

The concrete specifications of our test machine can be found in Table 2. Due to the multi-tasking nature of Windows and Linux operating systems, the benchmarking result can be influenced by other tasks running in parallel. In order to minimize this effect, we conducted all our experiments under the safe mode of both OSs. We used Microsoft Visual Studio 2008 and GCC 4.3.2 as C compilers for the Windows and the Linux operating systems, respectively. Both compilers were configured to maximize the speed of generated executables. For Java programs, we used Eclipse SDK 3.6.1 and Java(TM) SE Runtime Environment (build 1.6.0.12-b04).

Benchmarking of SHA-256 on Single-Core Platform In this subsection, we present the results of benchmarking a single SHA-256 FU against some non-RVC reference implementations in C (OpenSSL [64], OGay [27], and sphlib [56]) and Java (Java Cryptography Architecture (JCA) [54]). Figure 2 shows the results of our benchmarking under Windows operating system while our test machine was configured to run only one CPU core. One can see that the run-time

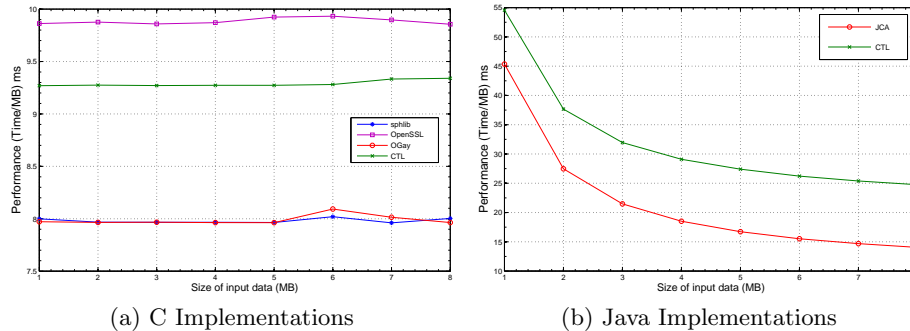


Fig. 2: Benchmarking of CTL’s SHA-256 implementation.

performance of CTL implementation is better than OpenSSL but inferior to carefully optimized implementations (OGay and sphlib). In addition, the CTL’s Java implementation of SHA-256 does not outperform the JCA implementation. This can probably be explained by the fact that the current edition of the ORCC Java backend does not generate very efficient code. These results indicate that the CTL’s SHA-256 implementation can achieve a performance similar to reference implementations. We also did similar benchmarking experiments on the AES block cipher in CTL (included in the extended edition of the paper) and came to a similar conclusion.

Benchmarking of SHA-256 on Multi-Core Platform On a platform with multiple CPU cores, one can map different parts of an FU network to different CPU cores so that the overall run-time performance of the application can be improved. The C backend of the RVC supporting tool ORCC [6] supports multi-core mapping, so one can easily allocate different FUs or FU sub-networks to different CPU cores. To see how much benefit we can get from a multi-core platform, we devised a very simple RVC application called *HashTree* that implements the following functionality using five hash H operations: given an input signal $x = x_1 \parallel x_2 \parallel x_3 \parallel x_4$ consisting of four blocks x_i , hash each block $h_i = H(x_i)$ and then output $H(h_1 \parallel h_2 \parallel h_3 \parallel h_4)$. In our implementation of *HashTree*, we instantiated H with SHA-256. By comparing this application with the simple single-core SHA-256 application computing H on the same input (i.e., $H(x_1 \parallel x_2 \parallel x_3 \parallel x_4)$), we can roughly estimate the performance gain.

In the benchmarking process, we considered three different configurations:

- **Single SHA-256:** This configuration represents a single SHA-256 FU running on a single-core, which processes an input x and produces the hash. We used this configuration as the reference point to evaluate the performance gain of the following two configurations, which implement *HashTree* using five SHA-256 instances.
- **5-thread with manual mapping:** In this configuration, each SHA-256 instance is programmatically mapped to run as a separate thread on a specific

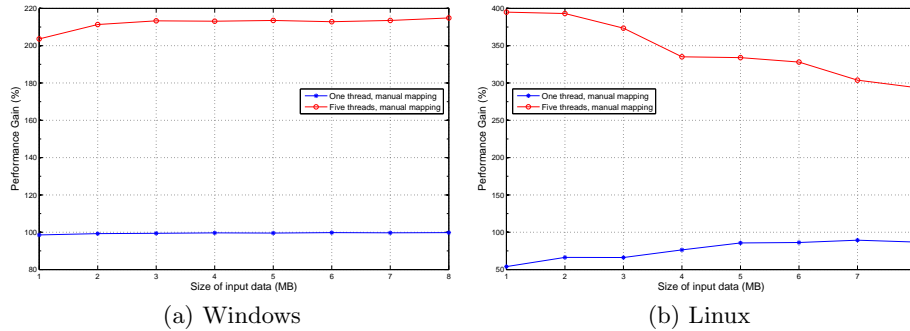


Fig. 3: The performance gain we got from the benchmarked configurations.

CPU core of our quad core machine. At the start of the hashing process, we manually mapped the four threads (processing $h_i = H(x_i)$) to four CPU cores. The fifth thread performing the final hashing operation is created and mapped after the preceding four threads are finished with their execution.

- **1-thread with manual mapping:** Similar to above configuration, this configuration also implements HashTree. However, all five SHA-256 instances are bounded to run in a single thread on a specific CPU core of our quad core machine.

It should be noted that thread creation and mapping also consume some CPU time, which is the cost one has to pay to achieve concurrency. Therefore, in order to make the study judicial, we also count the times spent on thread creation and thread mapping. The benchmarking results are shown in Fig. 3. One can see that the performance gain is between 200% to 300% when five threads are used.

6 Future Work

In order to allow researchers from different fields to extend CTL and use it for more applications, we have published CTL as an open-source project at <http://www.hooklee.com/default.asp?t=CTL>. In our future work, we plan to continue our research on the following possible directions.

Cryptographic Primitives. The CTL can be enriched by including more cryptographic primitives (especially public-key cryptography), which will allow creation of more multimedia security applications and security protocols. Another direction is to develop optimized versions of CTL cryptosystems. For instance, bit slicing can be used to optimize parallelism in many block ciphers [28, 45].

Security Protocols. Another direction is to use the RVC framework for the design and development of security protocols and systems with heterogenous components and interfaces. While RVC itself is platform independent, “wrappers” [62] can be developed to bridge the platform-independent FUs with physical I/O

devices/channels (e.g., a device attached to USB port, a host connected via LAN/WLAN, a website URL, etc.). Although there are many candidate protocols that can be considered, as a first step we plan to implement the hPIN/hTAN e-banking security protocol [42], which is a typical (but small-scale) heterogeneous system involving a hardware token, a web browser plugin on the user's computer, and a web service running on the remote e-banking server. We have already implemented an hPIN/hTAN prototype system without using RVC, so the new RVC-based implementation can be benchmarked against the existing system.

Cryptographic Protocols. Many cryptographic protocols require a high amount of computations. One example are garbled circuit protocols [68] that allow secure evaluation of an arbitrary function on sensitive data. These protocols can be used as basis for various privacy-preserving applications. At a high level, the protocol works by one party first generating an encrypted form of the function to be evaluated (called garbled circuit) which is then sent to the other party who finally decrypts the function using the encrypted input data of both parties and finally obtains the correct result. Recent implementation results show that such garbled circuit based protocols can be implemented in a highly efficient way in software [32]. However, until now, there exist no software implementations that exploit multi-core architectures. It was shown that such protocols can be optimized when using both software and hardware together: For generation of the garbled circuit, a trusted hardware token can generate the garbled circuit locally and hence remove the need to transfer it over the Internet [36]. Here, the encrypted versions of the gates, which require four invocations of a cryptographic hash function can be computed in parallel similar to the 4-adic hash tree we have shown in Sec. 5. Furthermore, the evaluation of garbled circuits can be improved when using hardware accelerations as shown in [37]. We believe that the RVC framework can serve as an ideal basis for hardware-software co-designed systems with parallelized and/or hardware-assisted garbled circuit-based protocols.

References

1. Esterel Synchronous Language. <http://www-sop.inria.fr/esterel.org/files/>
2. Graphiti. <http://graphiti-editor.sf.net>
3. LabVIEW. <http://www.ni.com/labview/whatis/>
4. Mathworks Simulink: Simulation and Model-Based Design. <http://www.mathworks.com/products/simulink/>
5. Open Data Flow (OpenDF). <http://sourceforge.net/projects/opendf>
6. Open RVC-CAL Compiler (ORCC). <http://sourceforge.net/projects/orcc>
7. Synopsys Studio. <http://www.synopsys.com/SYSTEMS/BLOCKDESIGN/DIGITALSIGNALPROCESSING/Pages/SystemStudio.aspx>
8. Cryptol: The Language of Cryptography. Case Study, http://corp.galois.com/downloads/cryptography/Cryptol_Casestudy.pdf (2008)
9. CAO and qhasm compiler tools. EU Project CACE deliverable D1.3, Revision 1.1, http://www.cace-project.eu/downloads/deliverables-y3/32_CACE_D1.3_CAO_and_qhasm_compiler_tools_Jan11.pdf (2011)

10. Ahmad, J.J., Li, S., Amer, I., Mattavelli, M.: Building multimedia security applications in the MPEG Reconfigurable Video Coding (RVC) framework. In: Proc. 2011 ACM SIGMM Multimedia and Security Workshop (MM&Sec 2011) (2011)
11. Akyildiz, I.F., Melodia, T., Chowdhury, K.R.: Wireless multimedia sensor networks: Applications and testbeds. *Proc. IEEE* 96(10), 1588–1605 (2008)
12. Ali, H.I.A.A., Patoary, M.N.I.: Design and Implementation of an Audio Codec (AMR-WB) using Dataflow Programming Language CAL in the OpenDF Environment. TR: IDE1009, Halmstad University, Sweden (2010)
13. Aman-Allah, H., Maarouf, K., Hanna, E., Amer, I., Mattavelli, M.: CAL dataflow components for an MPEG RVC AVC baseline encoder. *J. Signal Processing Systems* 63(2), 227–239 (2011)
14. Amer, I., Lucarz, C., Roquier, G., Mattavelli, M., Raulet, M., Nezan, J., Déforges, O.: Reconfigurable Video Coding on multicore: An overview of its main objectives. *IEEE Signal Processing Magazine* 26(6), 113–123 (2009)
15. Antola, A., Fracassi, M., Gotti, P., Sandionigi, C., Santambrogio, M.: A novel hardware/software codesign methodology based on dynamic reconfiguration with Impulse C and CoDeveloper. In: Proc. 2007 3rd Southern Conference on Programmable Logic (SPL 2007). pp. 221–224 (2007)
16. Barbosa, M., Noad, R., Page, D., Smart, N.P.: First steps toward a cryptography-aware language and compiler. Cryptology ePrint Archive: Report 2005/160, <http://eprint.iacr.org/2005/160.pdf> (2005)
17. Bernstein, D.J., Schwabe, P.: New AES software speed records. In: Progress in Cryptology – INDOCRYPT 2008. LNCS, vol. 5365, pp. 322–336 (2008)
18. Bertoni, G., Breveglieri, L., Fragneto, P., Macchetti, M., Marchesin, S.: Efficient software implementation of AES on 32-bit platforms. In: Cryptographic Hardware and Embedded Systems – CHES 2002. LNCS, vol. 2523, pp. 159–171 (2002)
19. Bhattacharyya, S., Eker, J., Janneck, J.W., Lucarz, C., Mattavelli, M., Raulet, M.: Overview of the MPEG Reconfigurable Video Coding framework. *J. Signal Processing Systems* 63(2), 251–263 (2011)
20. Boutellier, J., Gomez, V.M., Silvén, O., Lucarz, C., Mattavelli, M.: Multiprocessor scheduling of dataflow models within the Reconfigurable Video Coding framework. In: Proc. 2009 Conference on Design and Architectures for Signal and Image Processing (DASIP 2009) (2009)
21. Canright, D., Osvik, D.A.: A more compact AES. In: Selected Areas in Cryptography (SAC 2009). LNCS, vol. 5867, pp. 157–169 (2009)
22. Corbet, J.: The high-resolution timer (API). <http://lwn.net/Articles/167897> (2006)
23. Cryptico A/S: Rabbit stream cipher, performance evaluation. White Paper, Version 1.4, available online at <http://www.cryptico.com/DWSDownload.asp?File=Files%2FFiler%2FWP%5FRabbit%5FPerformance%2Epdf> (2005)
24. Dai, W.: Crypto++ library. <http://www.cryptopp.com>
25. Dennis, J.: First version of a data flow procedure language. In: Programming Symposium, Proceedings Colloque sur la Programmation Paris, April 9-11, 1974, LNCS, vol. 19, pp. 362–376 (1974)
26. Eker, J., Janneck, J.W.: CAL language report: Specification of the CAL actor language. Technical Memo UCB/ERL M03/48, Electronics Research Laboratory, UC Berkeley (2003)
27. Gay, O.: SHA-2: Fast Software Implementation. <http://www.ouah.org/ogay/sha2>
28. Grabher, P., Großschädl, J., Page, D.: Light-weight instruction set extensions for bit-sliced cryptography. In: Cryptographic Hardware and Embedded Systems – CHES 2008. LNCS, vol. 5154, pp. 331–345 (2008)

29. Gupta, S., Dutt, N., Gupta, R., Nicolau, A.: SPARK: A high-level synthesis framework for applying parallelizing compiler transformations. In: Proc. 2003 16th International Conference on VLSI Design (VLSI Design 2003) (2003)
30. Gutmann, P.: Cryptlib. <http://www.cs.auckland.ac.nz/~pgut001/cryptlib>
31. Ha, S., Kim, S., Lee, C., Yi, Y., Kwon, S., Joo, Y.P.: PeaCE: A hardware-software codesign environment for multimedia embedded systems. *ACM Trans. on Design Automation of Electronic Systems* 12(3), Article 24 (2007)
32. Huang, Y., Evans, D., Katz, J., Malka, L.: Faster secure two-party computation using garbled circuits. In: Proc. 20th USENIX Security Symposium (2011)
33. ISO/IEC: Information technology – MPEG video technologies – Part 4: Video tool library. ISO/IEC 23002-4 (2009)
34. ISO/IEC: Information technology - MPEG systems technologies - Part 4: Codec configuration representation. ISO/IEC 23001-4 (2009)
35. Janneck, J., Miller, I., Parlour, D., Roquier, G., Wipliez, M., Raulet, M.: Synthesizing hardware from dataflow programs: An MPEG-4 Simple Profile decoder case study. *J. Signal Processing Systems* 63(2), 241–249 (2011)
36. Järvinen, K., Kolesnikov, V., Sadeghi, A.R., Schneider, T.: Embedded SFE: Offloading server and network using hardware tokens. In: Financial Cryptography and Data Security (FC 2010). LNCS, vol. 6052, pp. 207–221 (2010)
37. Järvinen, K., Kolesnikov, V., Sadeghi, A.R., Schneider, T.: Garbled circuits for leakage-resilience: Hardware implementation and evaluation of one-time programs. In: Cryptographic Hardware and Embedded Systems – CHES 2010. LNCS, vol. 6225, pp. 383–397 (2010)
38. Kangas, T., Kukkala, P., Orsila, H., Salminen, E., Hännikäinen, M., Hämäläinen, T.D., Riihimäki, J., Kuusilinna, K.: UML-based multiprocessor SoC design framework. *ACM Trans. on Embedded Computer Systems* 5, 281–320 (2006)
39. Khan, E., El-Kharashi, M.W., Gebali, F., Abd-El-Barr, M.: Applying the Handel-C design flow in designing an HMAC-hash unit on FPGAs. *Computers and Digital Techniques* 153(5), 323–334 (2006)
40. Lee, E.A., Messerschmitt, D.G.: Synchronous data flow. *Proc. IEEE* 75(9), 1235–1245 (1987)
41. Lewis, J.R., Martin, B.: Cryptol: High assurance, retargetable crypto development and validation. In: Proc. 2003 IEEE Military Communication Conference (MILCOM 2003). pp. 820–825 (2003)
42. Li, S., Sadeghi, A.R., Heisrat, S., Schmitz, R., Ahmad, J.J.: hPIN/hTAN: A lightweight and low-cost e-banking solution against untrusted computers. In: Financial Cryptography and Data Security (FC 2011), LNCS, vol. 7035, pp. 235–249 (2011)
43. Lucarz, C., Mattavelli, M., Dubois, J.: A co-design platform for algorithm/architecture design exploration. In: Proc. 2008 IEEE International Conference on Multimedia and Expo (ICME 2008). pp. 1069–1072 (2008)
44. Manley, R., Gregg, D.: A program generator for intel AES-NI instructions. In: Progress in Cryptology – INDOCRYPT 2010. LNCS, vol. 6498, pp. 311–327 (2010)
45. Matsui, M., Nakajima, J.: On the power of bitslice implementation on Intel Core2 processor. In: Cryptographic Hardware and Embedded Systems – CHES 2007. LNCS, vol. 4727, pp. 121–134 (2007)
46. Moran, T.: The Qilin Crypto SDK: An open-source Java SDK for rapid prototyping of cryptographic protocols. <http://qilin.seas.harvard.edu/>
47. Moss, A., Page, D.: Bridging the gap between symbolic and efficient AES implementations. In: Proc. 2010 ACM SIGPLAN Workshop on Partial Evaluation and Program Manipulation (PEPM 2010). pp. 101–110 (2010)

48. Moving Picture Experts Group (MPEG): Who we are. http://mpeg.chiariglione.org/who_we_are.htm
49. Nikhil, R.: Tutorial – BlueSpec SystemVerilog: Efficient, correct RTL from high-level specifications. In: Proc. 2nd ACM/IEEE International Conference on Formal Methods and Models for Co-Design (MEMOCODE 2004). pp. 69–70 (2004)
50. NIST: Data Encryption Standard (DES). FIPS PUB 46-3 (1999)
51. NIST: Specification for the Advanced Encryption Standard (AES). FIPS PUB 197 (2001)
52. NIST: Recommendation for the Triple Data Encryption Algorithm (TDEA) block cipher. Special Publication 800-67, Version 1.1 (2008)
53. NIST: Secure Hash Standard (SHS). FIPS PUB 180-3 (2008)
54. Oracle®: Java™ Cryptography Architecture (JCA) Reference Guide. <http://download.oracle.com/javase/6/docs/technotes/guides/security/crypto/CryptoSpec.html>
55. Osvik, D.A., Bos, J.W., Stefan, D., Canright, D.: Fast software AES encryption. In: Fast Software Encryption (FSE 2010). LNCS, vol. 6147, pp. 75–93 (2010)
56. Pornin, T.: sphlib 3.0. <http://www.saphir2.com/sphlib>
57. PureNoise Ltd Vaduz: PureNoise CryptoLib. <http://cryptolib.com/crypto>
58. Rompaey, K.V., Verkest, D., Bolsens, I., Man, H.D.: CoWare – a design environment for heterogeneous hardware/software systems. Design Automation for Embedded Systems 1(4), 357–386 (1996)
59. Schneier, B.: Description of a New Variable-Length Key, 64-bit Block Cipher (Blowfish). In: Fast Software Encryption (FSE'94). LNCS, vol. 809, pp. 191–204 (1994)
60. Schneier, B.: Applied Cryptography: Protocols, algorithms, and source code in C. John Wiley & Sons, Inc., New York, second edn. (1996)
61. Sutherland, W.R.: The On-Line Graphical Specification of Computer Procedures. Ph.D. thesis, MIT (1966)
62. Thavot, R., Mosqueron, R., Dubois, J., Mattavelli, M.: Hardware synthesis of complex standard interfaces using CAL dataflow descriptions. In: Proc. 2009 Conference on Design and Architectures for Signal and Image Processing (DASIP 2009) (2009)
63. The Legion of the Bouncy Castle: Bouncy Castle Crypto APIs. <http://www.bouncycastle.org>
64. The OpenSSL Project: OpenSSL cryptographic library. <http://www.openssl.org/docs/crypto/crypto.html>
65. Thompson, M., Nikolov, H., Stefanov, T., Pimentel, A.D., Erbas, C., Polstra, S., Deprettere, E.F.: A framework for rapid system-level exploration, synthesis, and programming of multimedia MP-SoCs. In: Proc. 5th IEEE/ACM International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS 2007). pp. 9–14 (2007)
66. Tillich, S., Großschädl, J.: Instruction set extensions for efficient AES implementation on 32-bit processors. In: Cryptographic Hardware and Embedded Systems – CHES 2006. LNCS, vol. 4249, pp. 270–284 (2006)
67. Tillich, S., Herbst, C.: Boosting AES performance on a tiny processor core. In: Topics in Cryptology – CT-RSA 2008. LNCS, vol. 4964, pp. 170–186 (2008)
68. Yao, A.C.: How to generate and exchange secrets. In: Proc. 27th Annual Symposium on Foundations of Computer Science (FOCS'86). pp. 162–167 (1986)