

Extracting and Verifying Cryptographic Models from C Protocol Code by Symbolic Execution

Mihhail Aizatulin
The Open University

Andrew D. Gordon
Microsoft Research

Jan Jürjens
TU Dortmund & Fraunhofer ISST

ABSTRACT

Consider the problem of verifying security properties of a cryptographic protocol coded in C. We propose an automatic solution that needs neither a pre-existing protocol description nor manual annotation of source code. First, symbolically execute the C program to obtain symbolic descriptions for the network messages sent by the protocol. Second, apply algebraic rewriting to obtain a process calculus description. Third, run an existing protocol analyser (ProVerif) to prove security properties or find attacks. We formalise our algorithm and appeal to existing results for ProVerif to establish computational soundness under suitable circumstances. We analyse only a single execution path, so our results are limited to protocols with no significant branching. The results in this paper provide the first computationally sound verification of weak secrecy and authentication for (single execution paths of) C code.

Categories and Subject Descriptors

C.2.2 [Computer-Communication Networks]: Network Protocols—*Protocol verification*; D.2.4 [Software Engineering]: Software/Program Verification

General Terms

Security, Verification

1. INTRODUCTION

Recent years have seen great progress in formal verification of cryptographic protocols, as illustrated by powerful tools like ProVerif [8], CryptoVerif [7] or AVISPA [4]. There remains, however, a large gap between what we verify (formal descriptions of protocols, say, in the pi calculus) and what we rely on (protocol implementations, often in low-level languages like C). The need to start the verification from C code has been recognised before and implemented in tools like CSur [20] and ASPIER [13], but the methods

```
void * key; size_t keylen;
readenv("k", &key, &keylen);
size_t len;
read(&len, sizeof(len));
if(len > 1000) exit();
void * buf = malloc(len + 2 * MACLEN);
read(buf, len);
mac(buf, len, key, keylen, buf + len);
read(buf + len + MACLEN, MACLEN);
if(memcmp(buf + len,
          buf + len + MACLEN,
          MACLEN) == 0)
    event("accept", buf, len);
```

$\text{in}(x_1); \text{in}(x_2); \text{if } x_2 = \text{mac}(k, x_1) \text{ then event } \text{accept}(x_1)$

Figure 1: An example C fragment together with the extracted model.

proposed there are still rather limited. Consider, for example, the small piece of C code in fig. 1 that checks whether a message received from the network matches a message authentication code. Intuitively, if the key is honestly chosen and kept secret from the attacker then with overwhelming probability the event will be triggered only if another honest participant (with access to the key) generated the message. Unfortunately, previous approaches cannot prove this property: the analysis of CSur is too coarse to deal with authentication properties like this and ASPIER cannot directly deal with code manipulating memory through pointers. Furthermore the previous works do not offer a definition of security directly for C code, i.e. they do not formally state what it means for a C program to satisfy a security property, which makes it difficult to evaluate their overall soundness. The goal of our work is to improve upon this situation by giving a formal definition of security straight for C code and proposing a method that can verify secrecy and authentication for typical memory-manipulating implementations like the one in fig. 1 in a fully automatic and scalable manner, without relying on a pre-existing protocol specification.

Our method proceeds by extracting a high-level model from the C code that can then be verified using existing tools (we use ProVerif in our work). Currently we restrict our analysis to code in which all network outputs happen on a single execution path, but otherwise we do not require use of any specific programming style, with the aim of applying our methods to legacy implementations. In particular, we do not assume memory safety, but instead explicitly verify it during model extraction. The method still assumes that the cryptographic primitives such as encryption or hashing

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

CCS'11, October 17–21, 2011, Chicago, Illinois, USA.

Copyright 2011 ACM 978-1-4503-0948-6/11/10 ...\$10.00.

are implemented correctly—verification of these is difficult even when done manually [3].

The two main contributions of our work are:

- formal definition of security properties for source code;
- an algorithm that computes a high-level model of the protocol implemented by a C program.

We implement and evaluate the algorithm as well as give a proof of its soundness with respect to our security definition. Our definition of security for source code is given by linking the semantics of a programming language, expressed as a transition system, to a computational security definition in the spirit of [10, 19, 30]. We allow an arbitrary number of sessions. We restrict our definition to trace properties (such as weak secrecy or authentication), but do not consider observational equivalence (for strong secrecy, say).

Due to the complexity of the C language we give the formal semantics for a simple assembler-like language into which C code can be easily compiled, as in other symbolic execution approaches such as [14]. The soundness of this step can be obtained by using well-known methods, as outlined in the technical report [2].

Our model-extraction algorithm produces a model in an intermediate language without memory access or destructive updates, while still preserving our security definition. The algorithm is based on symbolic execution [21] of the C program, using symbolic expressions to over-approximate the sets of values that may be stored in memory during concrete execution. The main difference from existing symbolic execution algorithms (such as [12] or [18]) is that our variables represent bitstrings of potentially unknown length, whereas in previous algorithms a single variable corresponds to a single byte.

We show how the extracted models can be further simplified into the form understood by ProVerif. We apply the computational soundness result from [5] to obtain conditions where the symbolic security definition checked by ProVerif corresponds to our computational security definition. Combined with the security-preserving property of the model extraction algorithm this provides a computationally sound verification of weak secrecy and authentication for C.

Outline of our Method. The verification proceeds in several steps, as outlined in fig. 2. The method takes as input:

- the C implementations of the protocol participants, containing calls to a special function `event` as in fig. 1,
- an environment process (in the modelling language) which spawns the participants, distributes keys, etc.,
- symbolic models of cryptographic functions used by the implementation,
- a property that event traces in the execution are supposed to satisfy with overwhelming probability.

We start by compiling the program down to a simple stack-based instruction language (CVM) using CIL [25] to parse and simplify the C input. The syntax and semantics of CVM are presented in section 2. Due to space constraints the translation from C to CVM is in the technical report [2].

In the next step we symbolically execute CVM programs to eliminate memory accesses and destructive updates, thus obtaining an equivalent program in an intermediate model language (IML)—a version of the applied pi calculus extended with bitstring manipulation primitives. For each al-

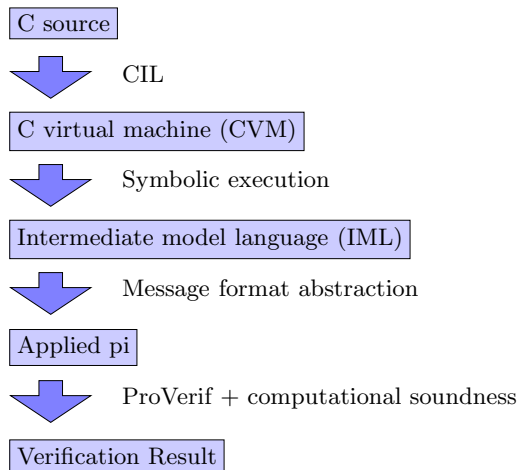


Figure 2: An outline of the method

located memory area the symbolic execution stores an expression describing how the contents of the memory area have been computed. For instance a certain memory area might be associated with an expression $hmac(01\|x, k)$, where x is known to originate from the network, k is known to be an environment variable, and $\|$ denotes concatenation. The symbolic execution does not enter the functions that implement the cryptographic primitives, it uses the provided symbolic models instead. These models thus form the trusted base of the verification. An example of the symbolic execution output is shown at the bottom of fig. 1. We define the syntax and semantics of IML in section 3 and describe the symbolic execution in section 5.

Our definition of security for source code is given in section 4. The definition is generic in that it does not assume a particular programming language. We simply require that the semantics of a language is given as a set of transitions of a certain form, and define a computational execution of the resulting transition system in the presence of an attacker and the corresponding notion of security. This allows one to apply the same security definition to protocols expressed both in the low-level implementation language and in the high-level model-description language, and to formulate a correspondence between the two.

Given that the transition systems generated by different languages are required to be of the same form, we can mix them in the same execution. This allows us to use CVM to specify a single executing participant, but at the same time use IML to describe an environment process that spawns multiple participants and allows them to interact. In particular, CVM need not be concerned with concurrency, thus making symbolic execution easier. Given an environment process P_E with n holes, we write $P_E[P_1, \dots, P_n]$ for a process where the i th hole is filled with P_i , which can be either a CVM or an IML process. The soundness result for symbolic execution (theorem 1) states that if P_1, \dots, P_n are CVM processes and $\tilde{P}_1, \dots, \tilde{P}_n$ are IML models resulting from their symbolic execution then for any environment process P_E the security of $P_E[\tilde{P}_1, \dots, \tilde{P}_n]$ with respect to a trace property ρ relates to the security of $P_E[P_1, \dots, P_n]$ with respect to ρ .

To verify the security of an IML process, we replace its bitstring-manipulating expressions by applications of constructor and destructor functions, thus obtaining a process in the applied pi-calculus (the version proposed in [9] and

augmented with events). We can then apply a computational soundness result, such as the one from [5], to specify conditions under which such a substitution is computationally sound: if the resulting pi calculus process is secure in a symbolic model (as can be checked by ProVerif) then it is asymptotically secure in our computational model (the probability of a successful attack is negligible [5]). The correctness of translation from IML to pi is captured by theorem 2 and the computational soundness for resulting pi processes is captured by theorem 3. The verification of IML (and these two theorems in particular) is described in section 6.

Theoretical and Practical Evaluation. Theorems 1 to 3 establish the correctness of our approach. In a nutshell, their significance is as follows: given implementations P_1, \dots, P_n of protocol participants in CVM, which are automatically obtained from the corresponding C code, and an IML process P_E that describes an execution environment, if P_1, \dots, P_n are successfully symbolically executed with resulting models $\tilde{P}_1, \dots, \tilde{P}_n$, the IML process $P_E[\tilde{P}_1, \dots, \tilde{P}_n]$ is successfully translated to a pi process P_π , and ProVerif successfully verifies P_π against a trace property ρ then P_1, \dots, P_n form a secure protocol implementation with respect to the environment P_E and property ρ . We did not aim for completeness; our implementation work has been driven by examples and in our theory we have aimed to show soundness of our implementation.

We are aiming to apply our method to large legacy code bases like OpenSSL. As a step towards this goal we evaluated it on a range of protocol implementations, including recent code for smart electricity meters [27]. We were able to find bugs in preexisting implementations or to verify them without having to modify the code. Section 7 provides details.

The current restriction of analysis to a single execution path may seem prohibitive at first sight. In fact, a great majority of protocols (such as those in the extensive SPORE repository [26]) follow a fixed narration of messages between participants, where any deviation from the expected message leads to termination. For such protocols, our method allows us to capture and analyse the fixed narration directly from the C code. In the future we plan to extend the analysis to more sophisticated control flow by generalising our symbolic execution to an abstract interpretation [15].

Related Work. We mention particularly relevant works, a broader survey can be found in [2]. One of the first attempts at cryptographic verification of C code is contained in [20], where a C program is used to generate a set of Horn clauses that are then solved using a theorem prover. The method is implemented in the tool CSur. We improve upon CSur in two ways in particular.

First, we have an explicit attacker model with a standard computational attacker. The attacker in CSur is essentially symbolic—it is allowed to apply cryptographic operations, but cannot perform any arithmetic computations.

Second, we handle authentication properties in addition to secrecy properties. Adding authentication to CSur would be non-trivial, due to a rather coarse over-approximation of C code. For instance, the order of instructions in CSur is ignored, and writing a single byte into an array with unknown length is treated the same as overwriting the whole array. Authentication, however, crucially depends on the order of events in the execution trace as well as making sure that the

authenticity of a whole message is preserved and not only of a single byte of it.

ASPIER [13] uses model checking to verify implementations of cryptographic protocols. The model checking operates on a protocol description language, which is rather more abstract than C; for instance, it does not contain pointers and cannot express variable message lengths. The translation from C to the protocol language is not described in the paper. Our method applies directly to C code with pointers, so that we expect it to provide much greater automation.

Corin and Manzano [14] report an extension of the KLEE test-generation tool [12] that allows KLEE to be applied to cryptographic protocol implementations (but not to extract models, as in our work). They do not extend the class of properties that KLEE is able to test for; in particular, testing for trace properties is not yet supported. Similarly to our work, KLEE is based on symbolic execution; the main difference is that [14] treats every byte in a memory buffer separately and thus only supports buffers of fixed length.

An advantage of the aforementioned tools (CSur, ASPIER, and KLEE) is that they do not share our limitation to single execution paths.

An online technical report includes more details and proofs of all results stated in this paper [2]. The implementation can be found at <http://github.com/tari3x/csec-tools>.

2. C VIRTUAL MACHINE (CVM)

This section describes our low-level source language CVM (C Virtual Machine). The language is simple enough to formalise, while at the same time the operations of CVM are closely aligned with the operations performed by C programs, so that it is easy to translate from C to CVM.

The model of execution of CVM is a stack-based machine with random memory access. All operations with values are performed on the stack, and values can be loaded from memory and stored back to memory. The language contains primitive operations that are necessary for implementing security protocols: reading values from the network or the execution environment, choosing random values, writing values to the network and signalling events. The only kind of conditional that CVM supports is a testing operation that checks a boolean condition and aborts execution immediately if it is not satisfied.

The fact that CVM permits no looping or recursion in the program allows us to inline all function calls, so that we do not need to add a call operation to the language itself. For simplicity of presentation we omit some aspects of the C language that are not essential for describing the approach, such as global variable initialisation and structures. We also restrict program variables to all be of the same size: for the rest of the paper we choose a fixed but arbitrary $N \in \mathbb{N}$ and assume $\text{sizeof}(v) = N$ for all program variables v . Our implementation does not have these restrictions and deals with the full C language.

Let $BS = \{0, 1\}^*$ be the set of finite bitstrings with the empty bitstring denoted by ε . For a bitstring b let $|b|$ be the length of b in bits. Let Var be a countably infinite set of variables. We write $f: X \rightarrow Y$ to denote a partial function and let $\text{dom}(f) \subseteq X$ be the set of x for which $f(x)$ is defined. We write $f(x) = \perp$ when f is not defined on x and use the notation $f\{x \mapsto a\}$ to update functions. A *partial algorithm* is an algorithm that implements a partial function.

Let \mathbf{Ops} be a finite set of operation symbols such that

$b \in BS, v \in Var, op \in \mathbf{Ops}$	
$src ::= \mathbf{read} \mid \mathbf{rnd}$	input source
$dest ::= \mathbf{write} \mid \mathbf{event}$	output destination
$instr ::=$	instruction
$\mathbf{Const} \ b$	constant value
$\mathbf{Ref} \ v$	pointer to variable
\mathbf{Malloc}	pointer to fresh memory
\mathbf{Load}	load from memory
$\mathbf{In} \ v \ src$	input
$\mathbf{Env} \ v$	environment variable
$\mathbf{Apply} \ op$	operation
$\mathbf{Out} \ dest$	output
\mathbf{Test}	test a condition
\mathbf{Store}	write to memory
$P \in CVM ::= \{instr\}^*$	program

Figure 3: The syntax of CVM.

each $op \in \mathbf{Ops}$ has an associated arity $\text{ar}(op)$ and an efficient partial algorithm $A_{op}: BS^{\text{ar}(op)} \rightarrow BS$. The set \mathbf{Ops} is meant to contain both the primitive operations of the language (such as the arithmetic or comparison operators of C) and the cryptographic primitives that are used by the implementation. The security definitions of this paper (given later) assume an arbitrary security parameter. Since real-life cryptographic protocols are typically designed and implemented for a fixed value of the security parameter, for the rest of the paper we let $k_0 \in \mathbb{N}$ be the security parameter with respect to which the operations in \mathbf{Ops} are chosen.

The syntax of CVM is shown in fig. 3. To define the semantics of CVM we choose two functions that relate bitstrings to integer values, $\text{val}: BS \rightarrow \mathbb{N}$ and $\text{bs}: \mathbb{N} \rightarrow BS$ and require that for $n < 2^N$ the value $\text{bs}(n)$ is a bitstring of length N such that $\text{val}(\text{bs}(n)) = n$. We allow bs to have arbitrary behaviour for larger numbers. The functions val and bs encapsulate architecture-specific details of integer representation such as the endianness. Even though these functions capture an unsigned interpretation of bitstrings, we only use them when accessing memory cells and otherwise place no restriction on how the bitstrings are interpreted by the program operations. For instance, the set \mathbf{Ops} can contain both a signed and an unsigned arithmetic and comparison operators. Bitstring representations of integer constants shall be written as $i1, i20$, etc, for instance, $i10 = \text{bs}(10)$.

We let $\text{Addr} = \{1, \dots, 2^N - 1\}$ be the set of valid memory addresses. The reason we exclude 0 is to allow the length of the memory to be represented in N bits. The semantic configurations of CVM are of the form $(\mathcal{A}^c, \mathcal{M}^c, \mathcal{S}^c, P)$, where

- $\mathcal{M}^c: \text{Addr} \rightarrow \{0, 1\}$ is a partial function that represents concrete memory and is undefined for uninitialised cells,
- $\mathcal{A}^c \subseteq \text{Addr}$ is the set of allocated memory addresses,
- \mathcal{S}^c is a list of bitstrings representing the execution stack,
- $P \in CVM$ is the executing program.

Semantic transitions are of the form $(\eta, s) \xrightarrow{l} (\eta', s')$, where s and s' are semantic configurations, η and η' are environments (mappings from variables to bitstrings) and l is a protocol action such as reading or writing values from the

attacker or a random number generator, or raising events. We give an informal overview of the semantics of CVM; the details are in the technical report [2]. Before the program is executed, each referenced variable v is allocated an address $\text{addr}(v)$ in \mathcal{M}^c such that all allocations are non-overlapping. If the program contains too many variables to fit in memory, the execution does not proceed. Next, the instructions in the program are executed one by one as described below. For $a, b \in \mathbb{N}$ we define $\{a\}_b = \{a, \dots, a + b - 1\}$.

- **Const** b places b on the stack.
- **Ref** v places $\text{bs}(\text{addr}(v))$ on the stack.
- **Malloc** takes a value s from the stack, reads a value p from the attacker, and if the range $\{\text{val}(p)\}_{\text{val}(s)}$ does not contain allocated cells, it becomes allocated and the value p is placed on the stack. Thus the attacker gets to choose the beginning of the allocated memory area.
- **Load** takes values l and p from the stack. In case $\{\text{val}(p)\}_{\text{val}(l)}$ is a completely initialised range in memory, the contents of that range are placed on the stack. In case some of the bits are not initialised, the value for those bits is read from the attacker.
- **In** v **read** or **In** v **rnd** takes a value l from the stack. **In** v **read** reads a value of length $\text{val}(l)$ from the attacker and **In** v **rnd** requests a random value of length $\text{val}(l)$. The resulting value b is then placed on the stack. The environment η is extended by the binding $v \mapsto b$.
- **Env** v places $\eta(v)$ and $\text{bs}(|\eta(v)|)$ on the stack.
- **Apply** op with $\text{ar}(op) = n$ applies A_{op} to n values on the stack, replacing them by the result.
- **Out write** sends the top of the stack to the attacker and **Out event** raises an event with the top of the stack as payload. Events with multiple arguments can be represented using a suitable bitstring pairing operation. Both commands remove the top of the stack.
- **Test** takes the top of the stack and checks whether it is $i1$. If yes, the execution proceeds, otherwise it stops.
- **Store** takes values p and b from the stack and writes b into memory at position starting with $\text{val}(p)$.

3. INTERMEDIATE MODEL LANGUAGE

This section presents the intermediate model language (IML) that we use both to express the models extracted from CVM programs and to describe the environment in which the protocol participants execute. IML borrows most of its structure from the pi calculus [1, 9]. In addition it has access both to the set \mathbf{Ops} of operations used by CVM programs and to primitive operations on bitstrings: concatenation, substring extraction, and computing lengths of bitstrings. Unlike CVM, IML does not access memory or perform destructive updates.

The syntax of IML is presented in fig. 4. In contrast to the standard pi calculus we do not include channel names, but implicitly use a single public channel instead. This corresponds to our security model in which all communication happens through the attacker. The nonce sampling operation $(\nu x[e])$ takes an expression as a parameter that specifies the length of the nonce to be sampled—this is necessary in the computational setting in order to obtain a probability distribution. We introduce a special abbreviation for programs that choose randomness of length equal to the secu-

$b \in BS, x \in Var, op \in \mathbf{Ops}$	
$e \in IExp ::=$	expression
b	concrete bitstring
x	variable
$op(e_1, \dots, e_n)$	computation
$e_1 \parallel e_2$	concatenation
$e\{e_o, e_l\}$	substring extraction
$len(e)$	length
$P, Q \in IML ::=$	process
0	nil
$!P$	replication
$P Q$	parallel composition
$(\nu x[e]); P$	randomness
$\mathbf{in}(x); P$	input
$\mathbf{out}(e); P$	output
$\mathbf{event}(e); P$	event
$\mathbf{if } e \mathbf{ then } P \mathbf{ [else } Q\mathbf{]}$	conditional
$\mathbf{let } x = e \mathbf{ in } P \mathbf{ [else } Q\mathbf{]}$	evaluation

Figure 4: The syntax of IML.

$$\begin{aligned}
\llbracket b \rrbracket &= b, \text{ for } b \in BS, \\
\llbracket x \rrbracket &= \perp, \text{ for } x \in Var, \\
\llbracket op(e_1, \dots, e_n) \rrbracket &= A_{op}(\llbracket e_1 \rrbracket, \dots, \llbracket e_n \rrbracket), \\
\llbracket e_1 \parallel e_2 \rrbracket &= \llbracket e_1 \rrbracket \parallel \llbracket e_2 \rrbracket, \\
\llbracket e\{e_o, e_l\} \rrbracket &= \text{sub}(\llbracket e \rrbracket, \text{val}(\llbracket e_o \rrbracket), \text{val}(\llbracket e_l \rrbracket)), \\
\llbracket len(e) \rrbracket &= \text{bs}(\llbracket e \rrbracket).
\end{aligned}$$

Figure 5: The evaluation of IML expressions, whereby \perp propagates.

rity parameter k_0 introduced in section 2: $\mathbf{let } (\tilde{\nu}x); P$ stand for $(\nu \tilde{x}[k_0]); \mathbf{let } x = \mathbf{nonce}(\tilde{x}) \mathbf{ in } P$, where $\mathbf{nonce} \in \mathbf{Ops}$. Using \mathbf{nonce} allows us to have tagged nonces, which will be necessary to link to the pi calculus semantics from [5].

For a bitstring b let $b[i]$ be the i th bit of b counting from 0. The concatenation of two bitstrings b_1 and b_2 is written as $b_1 \parallel b_2$.

Just as for CVM, the semantics of IML is parameterised by functions bs and val . The semantics of expressions is given by the partial function $\llbracket \cdot \rrbracket : IExp \rightarrow BS$ described in fig. 5. The partial function $\text{sub} : BS \times \mathbb{N} \times \mathbb{N} \rightarrow BS$ extracts a substring of a given bitstring such that $\text{sub}(b, o, l)$ is the substring of b starting at offset o of length l :

$$\text{sub}(b, o, l) = \begin{cases} b[o] \dots b[o+l-1] & \text{if } o+l \leq |b|, \\ \perp & \text{otherwise.} \end{cases}$$

For a *valuation* $\eta : Var \rightarrow BS$ we denote with $\llbracket e \rrbracket_\eta$ the result of substituting all variables v in e by $\eta(v)$ (if defined) and then applying $\llbracket \cdot \rrbracket$.

The technical report [2] describes the detailed semantics of IML.

4. SECURITY OF PROTOCOLS

To define security for protocols implemented by CVM and IML programs we need to specify what a protocol is and give a mapping from programs to protocols. The notion of a protocol is formally captured by a *protocol transition system (PTS)*, which describes how processes evolve and in-

teract with the attacker. A PTS is a set of transitions of the form $(\eta, s) \xrightarrow{l} \{(\eta_1, s_1), \dots, (\eta_n, s_n)\}$, where η and η_i are environments (modelled as valuations), s and s_i are semantic configurations of the underlying programming language, and l is an action label. Actions can include reading values from the attacker, generating random values, sending values to the attacker, or raising events. We call a pair (η, s) an *executing process*. Multiple processes on the right hand side capture replication. Full details are in [2].

The semantics of CVM and IML are given in terms of the PTS that are implemented by programs. For a CVM program P we denote with $\llbracket P \rrbracket_C$ the PTS that is implemented by P . Similarly, for an IML process P the corresponding PTS is denoted by $\llbracket P \rrbracket_I$.

Given a PTS T and a probabilistic machine E (an attacker) we can execute T in the presence of E . The state of the executing protocol is essentially a multiset of executing processes. The attacker repeatedly chooses a process from the multiset which is then allowed to perform an action according to T . The result of the execution is a sequence of raised events. For a resource bound $t \in \mathbb{N}$ we denote with $\text{Events}(T, E, t)$ the sequence of events raised during the first t steps of the execution. We shall be interested in the probability that this sequence of events belongs to a certain “safe” set. This is formally captured by the following definition:

Definition 1 (Protocol security) We define a *trace property* as a prefix-closed set of event sequences with efficiently decidable membership. For a PTS T , an attacker E , a trace property ρ and a resource bound $t \in \mathbb{N}$ let $\text{insec}(T, E, \rho, t)$ be the probability $\Pr[\text{Events}(T, E, t) \notin \rho]$. \square

Since the semantics of CVM and IML are in the same formalism, we may combine the sets of semantic rules and obtain semantics $\llbracket \cdot \rrbracket_{CI}$ for mixed programs, where a CVM program can be a subprocess of a larger IML process. We add an additional syntactic form $\llbracket \cdot \rrbracket_i$ (a hole) with $i \in \mathbb{N}$ and no reductions to IML. For an IML process P_E with n holes and CVM or IML processes P_1, \dots, P_n we write $P_E[P_1, \dots, P_n]$ to denote process P_E where each hole $\llbracket \cdot \rrbracket_i$ is replaced by P_i .

Being able to embed a CVM program within an IML process is useful for modelling. As an example, let P_1 be the CVM program resulting from the translation of the C code in fig. 1 and let P_2 be a description of another participant of the protocol, in either CVM or IML. Then we might be interested in the security of the following process:

$$P_E[P_1, P_2] = !((\tilde{\nu}k); (!P_1)!(P_2)),$$

where k is the free variable of P_1 (used as the key). A trace property ρ of interest might be, for instance, “Each event of the form $\mathbf{accept}(x)$ is preceded by an event of the form $\mathbf{request}(x)$ ”, where $\mathbf{request}$ is an event possibly raised in P_2 . The next section shows how we can relate the security of $P_E[P_1, P_2]$ to the security of $P_E[\tilde{P}_1, P_2]$, where IML process \tilde{P}_1 is a model of the CVM process P_1 , extracted by symbolic execution.

5. CVM TO IML: SYMBOLIC EXECUTION

We describe how to automatically extract an IML model from a CVM program while preserving security properties. The key idea is to execute a CVM program in a symbolic semantics, where, instead of concrete bitstrings, memory locations contain IML expressions representing the set of all possible concrete values at a given execution point.

$v \in Var, i \in \mathbb{N}$	
$pb \in PBase ::=$	pointer base
stack v	stack pointer to variable v
heap i	heap pointer with id i
$e \in SExp ::=$	symbolic expression
ptr(pb, e)	pointer
...	same as <i>IExp</i> in fig. 4

Figure 6: Symbolic expressions.

$\text{apply}(+_b, \text{ptr}(pb, e_o), e) = \text{ptr}(pb, e_o +_b e),$
 for $e \in IExp,$
 $\text{apply}(-_b, \text{ptr}(pb, e_o), \text{ptr}(pb, e'_o)) = e_o -_b e'_o,$
 $\text{apply}(op, e_1, \dots, e_n) = op(e_1, \dots, e_n),$
 for $e_1, \dots, e_n \in IExp,$
 $\text{apply}(\dots) = \perp,$ otherwise.

Figure 7: Operations on symbolic expressions.

To track the values used as pointers during CVM execution, we extend IML expressions with an additional construct, resulting in the class of *symbolic expressions* shown in fig. 6. An expression of the form $\text{ptr}(pb, e_o)$ represents a pointer into the memory location identified by the *pointer base* pb with an offset e_o relative to the beginning of the location. We require that $e_o \in IExp$, so that pointer offsets do not contain pointers themselves. Pointer bases are of two kinds: a base of the form *stack* v represents a pointer to the program variable v and a base of the form *heap* i represents the result of a **Malloc**.

Symbolic execution makes certain assumptions about the arithmetic operations that are available in **Ops**. We assume that programs use operators for bitwise addition and subtraction (with overflow) that we shall write as $+_b$ and $-_b$. We also make use of addition and subtraction without overflow—the addition operator (written as $+_{\mathbb{N}}$) is expected to make the result wide enough to prevent truncation or overflow (say, $1 +_{\mathbb{N}} 1 = 10$) and the negation operator (written as $-_{\mathbb{N}}$) returns \perp instead of a negative result. We assume that **Ops** contains comparison operators $=, \leq,$ and $<$ such that $A_{=} (a, b)$ returns $i1$ if $\text{val}(a) = \text{val}(b)$ and $i0$ otherwise, similarly for the other operators. This way \leq and $<$ capture unsigned comparisons on bitstring values. We assume **Ops** contains logical connectives \neg and \vee that interpret $i0$ as false value and $i1$ as true value. These operators may or may not be the ones used by the program itself.

To evaluate symbolic expressions concretely, we need concrete values for pointer bases as well as concrete values for variables. Given an *extended valuation* $\eta: Var \cup PBase \rightarrow BS$, we extend the function $\llbracket \cdot \rrbracket_{\eta}$ from fig. 5 by the rule:

$$\llbracket \text{ptr}(pb, e_o) \rrbracket_{\eta} = \eta(pb) +_b \llbracket e_o \rrbracket_{\eta}.$$

When applying arithmetic operations to pointers, we need to make sure that the operation is applied to the pointer offset and the base is kept intact. Also we can only apply $-_b$ to pointers with (syntactically) equal bases. This behaviour is encoded by the function apply , shown in fig. 7.

As well as tracking the expressions stored in memory, we also track logical facts discovered during symbolic execution. To record these facts, we use symbolic expressions themselves, interpreted as logical formulas with $=, \leq,$ and

$<$ as relations and \neg and \vee as connectives. We allow quantifiers in formulas, with straightforward interpretation. Given a set Σ of formulas and a formula ϕ we write $\Sigma \vdash \phi$ iff for each Σ -consistent valuation η (that is, a valuation such that $\llbracket \psi \rrbracket_{\eta} = i1$ for all $\psi \in \Sigma$) we also have $\llbracket \phi \rrbracket_{\eta} = i1$.

To check the entailment relation, our implementation relies on the SMT solver *Yices* [17], by replacing unsupported operations, such as string concatenation or substring extraction, with uninterpreted functions. This works well for our purpose—the conditions that we need to check during the symbolic execution are purely arithmetic and are supported by Yices’ theory.

The function getLen returns for each symbolic expression an expression representing its length:

$\text{getLen}(\text{ptr}(\dots)) = \text{bs}(N),$
 $\text{getLen}(\text{len}(\dots)) = \text{bs}(N),$
 $\text{getLen}(b) = \text{bs}(|b|),$ for $b \in BS,$
 $\text{getLen}(x) = \text{len}(x),$ for $x \in Var,$
 $\text{getLen}(op(e_1, \dots, e_n)) = \text{len}(op(e_1, \dots, e_n)),$
 $\text{getLen}(e_1 || e_2) = \text{getLen}(e_1) +_{\mathbb{N}} \text{getLen}(e_2),$
 $\text{getLen}(e\{e_o, e_l\}) = e_l.$

We assume that the knowledge about the return lengths of operation applications is encoded in a fact set Σ_{op} . As an example, Σ_{op} might contain the facts:

$$\forall x, y, a: \text{len}(x) = a \wedge \text{len}(y) = a \Rightarrow \text{len}(x +_b y) = a,$$

$$\forall x: \text{len}(\text{sha1}(x)) = i20.$$

We assume that Σ_{op} is consistent: $\emptyset \vdash \phi$ for all $\phi \in \Sigma_{op}$.

The transformations prescribed by the symbolic semantic rules would quickly lead to very large expressions. Thus the symbolic execution is parametrised by a simplification function simplify that is allowed to make use of the collected fact set Σ . We demand that the simplification function is sound in the following sense: for each fact set Σ , expression e and a Σ -consistent valuation η we have

$$\llbracket e \rrbracket_{\eta} \neq \perp \implies \llbracket \text{simplify}_{\Sigma}(e) \rrbracket_{\eta} = \llbracket e \rrbracket_{\eta}.$$

The simplifications employed in our algorithm are described in the technical report [2].

The algorithm for symbolic execution is determined by the set of semantic rules presented in fig. 8. The initial semantic configuration has the form (Init, P) with the executing program $P \in CVM$. The other semantic configurations have the form $(\Sigma, \mathcal{A}^s, \mathcal{M}^s, \mathcal{S}^s, P)$, where

- $\Sigma \subseteq SExp$ is a set of formulas (the path condition),
- $\mathcal{A}^s: PBase \rightarrow SExp$ is the symbolic allocation table that for each memory location stores its allocated size,
- $\mathcal{M}^s: PBase \rightarrow SExp$ is the symbolic memory. We require that $\text{dom}(\mathcal{M}^s) = \text{dom}(\mathcal{A}^s)$,
- \mathcal{S}^s is a list of symbolic expressions representing the execution stack,
- $P \in CVM$ is the executing program.

The crucial rules are **(S-Load)** and **(S-Store)** that reflect the effect of storing and loading memory values on the symbolic level. The rule **(S-Load)** is quite simple—it tries to deduce from Σ that the extraction is performed from a defined memory range given by \mathcal{M}^s (which is, in general, smaller than the allocated memory range given by \mathcal{A}^s), after which it represents the result of the extraction using an IML range

$(\text{Init}, P) \rightarrow (\Sigma_{op}, \{\text{stack } v \mapsto \text{bs}(N) \mid v \in \text{var}(P)\}, \{\text{stack } v \mapsto \varepsilon \mid v \in \text{var}(P)\}, [], P)$	(S-Init)
$(\Sigma, \mathcal{A}^s, \mathcal{M}^s, \mathcal{S}^s, \text{Const } b; P) \rightarrow (\Sigma, \mathcal{A}^s, \mathcal{M}^s, b :: \mathcal{S}^s, P)$	(S-Const)
$(\Sigma, \mathcal{A}^s, \mathcal{M}^s, \mathcal{S}^s, \text{Ref } v; P) \rightarrow (\Sigma, \mathcal{A}^s, \mathcal{M}^s, \text{ptr}(\text{stack } v, i0) :: \mathcal{S}^s, P)$	(S-Ref)
$\frac{e_l \in \text{IExp} \quad i \in \mathbb{N} \text{ minimal s.t. } pb = \text{heap } i \notin \text{dom}(\mathcal{M}^s)}{(\Sigma, \mathcal{A}^s, \mathcal{M}^s, e_l :: \mathcal{S}^s, \text{Malloc}; P) \rightarrow (\Sigma, \mathcal{A}^s \{pb \mapsto e_l\}, \mathcal{M}^s \{pb \mapsto \varepsilon\}, \text{ptr}(pb, i0) :: \mathcal{S}^s, P)}$	(S-Malloc)
$\frac{pb \in \text{dom}(\mathcal{M}^s) \quad e = \text{simplify}_{\Sigma}(\mathcal{M}^s(pb)\{e_o, e_l\}) \quad \Sigma \vdash (e_o +_{\mathbb{N}} e_l \leq \text{getLen}(\mathcal{M}^s(pb)))}{(\Sigma, \mathcal{A}^s, \mathcal{M}^s, e_l :: \text{ptr}(pb, e_o) :: \mathcal{S}^s, \text{Load}; P) \rightarrow (\Sigma, \mathcal{A}^s, \mathcal{M}^s, e :: \mathcal{S}^s, P)}$	(S-Load)
$\frac{e_l \in \text{IExp} \quad l = (\text{if } \text{src} = \text{read} \text{ then } \mathbf{in}(v); \text{ else } (\nu v[e_l]));}{(\Sigma, \mathcal{A}^s, \mathcal{M}^s, e_l :: \mathcal{S}^s, \text{In } v \text{ src}; P) \xrightarrow{l} (\Sigma \cup \{\text{len}(v) = e_l\}, \mathcal{A}^s, \mathcal{M}^s, v :: \mathcal{S}^s, P)}$	(S-In)
$(\Sigma, \mathcal{A}^s, \mathcal{M}^s, \mathcal{S}^s, \text{Env } v; P) \rightarrow (\Sigma, \mathcal{A}^s, \mathcal{M}^s, \text{len}(v) :: v :: \mathcal{S}^s, P)$	(S-Env)
$\frac{e = \text{apply}(op, e_1, \dots, e_n) \neq \perp}{(\Sigma, \mathcal{A}^s, \mathcal{M}^s, e_1 :: \dots :: e_n :: \mathcal{S}^s, \text{Apply } op; P) \rightarrow (\Sigma, \mathcal{A}^s, \mathcal{M}^s, \text{len}(e) :: e :: \mathcal{S}^s, P)}$	(S-Apply)
$\frac{e \in \text{IExp} \quad l = (\text{if } \text{dest} = \text{write} \text{ then } \mathbf{out}(e); \text{ else } \mathbf{event}(e));}{(\Sigma, \mathcal{A}^s, \mathcal{M}^s, e :: \mathcal{S}^s, \text{Out } \text{dest}; P) \xrightarrow{l} (\Sigma, \mathcal{A}^s, \mathcal{M}^s, \mathcal{S}^s, P)}$	(S-Out)
$\frac{e \in \text{IExp}}{(\Sigma, \mathcal{A}^s, \mathcal{M}^s, e :: \mathcal{S}^s, \text{Test}; P) \xrightarrow{\text{if } e \text{ then}} (\Sigma \cup \{e\}, \mathcal{A}^s, \mathcal{M}^s, \mathcal{S}^s, P)}$	(S-Test)
$\frac{\begin{array}{l} e_h = \mathcal{M}^s(pb) \neq \perp \quad e_s = \mathcal{A}^s(pb) \neq \perp \quad e_{lh} = \text{getLen}(e_h) \quad e_l = \text{getLen}(e) \\ \text{either } \Sigma \vdash (e_o +_{\mathbb{N}} e_l < e_{lh}) \text{ and } e'_h = \text{simplify}_{\Sigma}(e_h\{i0, e_o\}) \parallel e \parallel e_h\{e_o +_{\mathbb{N}} e_l, e_{lh} -_{\mathbb{N}} (e_o +_{\mathbb{N}} e_l)\} \\ \text{or } \Sigma \vdash (e_o +_{\mathbb{N}} e_l \geq e_{lh}) \wedge (e_o \leq e_{lh}) \wedge (e_o +_{\mathbb{N}} e_l \leq e_s) \text{ and } e'_h = \text{simplify}_{\Sigma}(e_h\{i0, e_o\}) \parallel e \end{array}}{(\Sigma, \mathcal{A}^s, \mathcal{M}^s, \text{ptr}(pb, e_o) :: e :: \mathcal{S}^s, \text{Store}; P) \rightarrow (\Sigma, \mathcal{A}^s, \mathcal{M}^s \{pb \mapsto e'_h\}, \mathcal{S}^s, P)}$	(S-Store)

Figure 8: The symbolic execution of CVM.

expression. The rule (S-Store) distinguishes between two cases depending on how the expression e to be stored is aligned with the expression e_h that is already present in memory. If e needs to be stored completely within the bounds of e_h then we replace the contents of the memory location by $e_h\{\dots\} \parallel e \parallel e_h\{\dots\}$ where the first and the second range expression represent the pieces of e_h that are not covered by e . In case e needs to be stored past the end of e_h , the new expression is of the form $e_h\{\dots\} \parallel e$. The rule still requires that the beginning of e is positioned before the end of e_h , and hence it is impossible to write in the middle of an uninitialised memory location. This is for simplicity of presentation—the rule used in our implementation does not have this limitation (it creates an explicit “undefined” expression in these cases).

Since all semantic rules are deterministic there is only one symbolic execution trace. The sequence of transition labels produces an IML process that simulates the behaviour of the original CVM program. Formally, for a CVM program P , let L be the symbolic execution trace starting from the state (Init, P) . If L ends in a state with an empty program, let $\lambda_1, \dots, \lambda_n$ be the sequence of labels of L and set $\llbracket P \rrbracket_S = \lambda_1 \dots \lambda_n 0 \in \text{IML}$, otherwise set $\llbracket P \rrbracket_S = \perp$.

We shall say that a polynomial is *fixed* iff it is independent of the arbitrary values assumed in this paper, such as N or the properties of the set **Ops**. Our main result relates the security of P to the security of $\llbracket P \rrbracket_S$.

Theorem 1 (Symbolic Execution is Sound) *There exists a fixed polynomial p such that if P_1, \dots, P_n are CVM*

processes and $\tilde{P}_i := \llbracket P_i \rrbracket_S \neq \perp$ for all i then for any IML process P_E , attacker E , trace property ρ , and resource bound $t \in \mathbb{N}$:

$$\begin{aligned} & \text{insec}(\llbracket P_E[P_1, \dots, P_n] \rrbracket_{CI}, E, \rho, t) \\ & \leq \text{insec}(\llbracket P_E[\tilde{P}_1, \dots, \tilde{P}_n] \rrbracket_I, E, \rho, p(t)). \quad \square \end{aligned}$$

The condition that p is fixed is important—otherwise p could be large enough to give the attacker the time to enumerate all the $2^{2^N - 1}$ memory configurations. For practical use the actual shape of p can be recovered from the proof of the theorem given in the technical report [2]. In particular, the proof shows that the execution time of the IML process is polynomial with respect to the execution time of the original C program.

Fig. 9 illustrates our method by showing how the symbolic execution proceeds for our example in fig. 1. For each line of the C program we show updates to the symbolic memory, the set of new facts, and the generated IML code if any. In our example **MAC_LEN** is assumed to be 20 and N is equal to **sizeof(size_t)**. The variables l , x_1 , and x_2 are arbitrary fresh variables chosen during the translation from C to CVM. Below we mention details for some particularly interesting steps (numbers correspond to line numbers in fig. 9).

1. The call to **readenv** redirects to a proxy function that generates CVM instructions for retrieving the environment variable k and storing it in memory.

Line no.	C line	symbolic memory updates	new facts	generated IML line
1.	<code>readenv("k", &key, &keylen);</code>	stack $key \Rightarrow \text{ptr}(\text{heap } 1, i0)$ heap $1 \Rightarrow k$ stack $keylen \Rightarrow \text{len}(k)$		
2.	<code>read(&len, sizeof(len));</code>	stack $len \Rightarrow l$	$\text{len}(l) = iN$	in (l)
3.	<code>if(len > 1000) exit();</code>		$\neg(l > i1000)$	
4.	<code>void * buf = malloc(len + 2 * MAC_LEN);</code>	stack $buf \Rightarrow \text{ptr}(\text{heap } 2, i0)$ heap $2 \Rightarrow \varepsilon$		
5.	<code>read(buf, len);</code>	heap $2 \Rightarrow x_1$	$\text{len}(x_1) = l$	in (x_1)
6.	<code>mac(buf, len, key, keylen, buf + len);</code>	heap $2 \Rightarrow x_1 \parallel \text{mac}(k, x_1)$		
7.	<code>read(buf + len + MAC_LEN, MAC_LEN);</code>	heap $2 \Rightarrow x_1 \parallel \text{mac}(k, x_1) \parallel x_2$	$\text{len}(x_2) = i20$	in (x_2)
8.	<code>if(memcmp(...) == 0)</code>			if $\text{mac}(k, x_1) = x_2$ then
9.	<code>event("accept", buf, len);</code>			event $\text{accept}(x_1)$

Figure 9: Symbolic execution of the example in fig. 1.

4. A new empty memory location is created and the pointer to it is stored in `buf`. We make an entry in the allocation table \mathcal{A}^s with the length of the new memory location $(l +_b i2 * i20)$.
5. We check that the stored value fits within the allocated memory area, that is, $l \leq l +_b i2 * i20$. This is in general not true due to possibility of integer overflow, but in this case succeeds due to the condition $\neg(l > i1000)$ recorded before (assuming that the maximum integer value $2^N - 1$ is much larger than 1000). Similar checks are performed for all subsequent writes to memory.
7. The memory update is performed through an intermediate pointer value of the form $\text{ptr}(\text{heap } 2, l +_b i20)$. The set of collected facts is enough to deduce that this pointer points exactly at the end of $x_1 \parallel \text{mac}(k, x_1)$.
8. The proxy function for `memcmp` extracts values $e_1 = e\{l, i20\}$ and $e_2 = e\{l +_b i20, i20\}$, where e is the contents of memory at heap 2, and puts $\text{cmp}(e_1, e_2)$ on the stack. With the facts collected so far e_1 simplifies to $\text{mac}(k, x_1)$ and e_2 simplifies to x_2 . With some special comprehension for the meaning of `cmp` we generate IML **if** $e_1 = e_2$ **then**.

6. VERIFICATION OF IML

The symbolic model extracted in fig. 9 does not contain any bitstring operations, so it can readily be given to ProVerif for verification. In general this is not the case and some further simplifications are required. In a nutshell, the simplifications are based on the observation that the bitstring expressions (concatenation and substring extraction) are meant to represent pairing and projection operations, so we can replace them by new symbolic operations that behave as pairing constructs in ProVerif. We then check that the expressions indeed satisfy the algebraic properties expected of such operations.

We outline the main results regarding the translation to ProVerif. The technical report [2] contains the details. The pi calculus used by ProVerif can be described as a subset of IML from which the bitstring operations have been removed. Unlike CVM and IML, the semantics of pi is given with respect to an arbitrary security parameter: we write $\llbracket P \rrbracket_\pi^k$ for the semantics of a pi process P with respect to the parameter $k \in \mathbb{N}$. In contrast, we consider IML as executing with respect to a fixed security parameter $k_0 \in \mathbb{N}$. For an IML process P we specify conditions under which it is translatable to a pi process \tilde{P} .

Theorem 2 (Soundness of the translation)

There exists a fixed polynomial p such that for any $P \in \text{IML}$

translatable to a pi process \tilde{P} , any attacker E , trace property ρ , and resource bound $t \in \mathbb{N}$:

$$\text{insec}(\llbracket P \rrbracket_I, E, \rho, t) \leq \text{insec}(\llbracket \tilde{P} \rrbracket_\pi^{k_0}, E, \rho, p(t)). \quad \square$$

Backes et al. [5] provide an example of a set of operations \mathbf{Ops}^S and a set of *soundness conditions* restricting their implementations that are sufficient for establishing computational soundness. The set \mathbf{Ops}^S contains a public key encryption operation that is required to be IND-CCA2 secure. The soundness result is established for the class of the so-called *key-safe* processes that always use fresh randomness for encryption and key generation, only use honestly generated decryption keys and never send decryption keys.

Theorem 3 (Computational soundness) *Let P be a pi process using only operations in \mathbf{Ops}^S such that the soundness conditions are satisfied. If P is key-safe and symbolically secure with respect to a trace property ρ (as checked by ProVerif) then for every polynomial p and attacker E the following function is negligible in k : $\text{insec}(\llbracket P \rrbracket_\pi^k, E, \rho, p(k))$.* \square

Overall, theorems 1 to 3 can be interpreted as follows: let P_1, \dots, P_n be implementations of protocol participants in CVM and let P_E be an IML process that describes an execution environment. Assume that P_1, \dots, P_n are successfully symbolically executed with resulting models $\tilde{P}_1, \dots, \tilde{P}_n$, the IML process $P_E[\tilde{P}_1, \dots, \tilde{P}_n]$ is successfully translated to a pi process P_π , and ProVerif successfully verifies P_π against a trace property ρ . Then we know by theorem 3 that P_π is a pi protocol model that is (asymptotically) secure with respect to ρ . By theorems 1 and 2 we know that P_1, \dots, P_n form a secure implementation of the protocol described by P_π for the security parameter k_0 .

7. IMPLEMENTATION & EXPERIMENTS

We have implemented our approach and successfully tested it on several examples. Our implementation performs the conversion from C to CVM at runtime—the C program is instrumented using CIL so that it outputs its own CVM representation when run. This allows us to identify and compile the main path of the protocol easily. Apart from information about the path taken we do not use any runtime information and we plan to make the analysis fully static in future. The idea of instrumenting a program to emit a low-level set of instructions for symbolic execution at runtime as well as some initial implementation code were borrowed from the CREST symbolic execution tool [11].

Currently we omit certain memory safety checks and assume that there are no integer overflows. This allows us

	C LOC	IML LOC	outcome	result type	time
simple mac	~ 250	12	verified	symbolic	4s
RPC	~ 600	35	verified	symbolic	5s
NSL	~ 450	40	verified	computat.	5s
CSur	~ 600	20	flaw: fig. 11	—	5s
minexplib	~ 1000	51	flaw: fig. 12	—	15s

Figure 10: Summary of analysed implementations.

```
read(conn_fd, temp, 128);
// BN_hex2bn expects zero-terminated string
temp[128] = 0;
BN_hex2bn(&cipher_2, temp);
// decrypt and parse cipher_2
// to obtain message fields
```

Figure 11: A flaw in the CSur example: input may be too short.

to use the more efficient theory of mathematical integers in Yices, but we are planning to move to exact bitvector treatment in future.

The implementation comprises about 4600 lines of OCaml code. The symbolic proxies for over 80 of the cryptographic functions in the OpenSSL library comprise further 2000 lines of C code.

Fig. 10 shows a list of protocol implementations on which we tested our method. Some of the verified programs did not satisfy the conditions of computational soundness (mostly because they use cryptographic primitives other than public key encryption and signatures supported by the result that we rely on [5]), so we list the verification type as “symbolic”, referring to Dolev and Yao’s standard “symbolic” (or formal) model of cryptography [16]. In all three cases marked “verified” in fig. 10, we successfully verified the IML process using the ProVerif verifier for this formal model. In the NSL case, we can appeal to the computational soundness result to obtain a result in the computational model. In all cases we provide a guarantee that the C program correctly implements a much simpler IML process (theorem 1).

The “simple mac” is an implementation of a protocol similar to the example in fig. 1. RPC is an implementation of the remote procedure call protocol in [6] that authenticates a server response to a client using a message authentication code. It was written by a colleague without being intended for verification using our method, but we were still able to verify it without any further modifications to the code.

The NSL example is an implementation of the Needham-Schroeder-Lowe protocol written by us to obtain a fully computationally sound verification result. Masking the second participant’s identity check triggers Lowe’s attack [24] as expected. The technical report [2] shows the source code and the extracted models.

The CSur example is the code analysed in a predecessor paper on C verification [20]. It is an implementation of a protocol similar to Needham-Schroeder-Lowe. During our verification attempt we discovered a flaw, shown in fig. 11: the received message in buffer `temp` is being converted to a `BIGNUM` structure `cipher_2` without checking that enough bytes were received. Later a `BIGNUM` structure derived from `cipher_2` is converted to a bitstring without checking that the length of the bitstring is sufficient to fill the message buffer. In both cases the code does not make sure that the information in memory actually comes from the network, which makes it impossible to prove authentication proper-

```
unsigned char session_key[256 / 8];
...
// Use the 4 first bytes as a pad
// to encrypt the reading
encrypted_reading =
  ((unsigned int) *session_key) ^ *reading;
```

Figure 12: A flaw in the minexplib code: only one byte of the pad is used.

ties. The CSur example has been verified in [20], but only for secrecy, and secrecy is not affected by the flaw we discovered. The code reinterprets network messages as C structures (an unsafe practise due to architecture dependence), which is not yet supported by our analysis and so we were not able to verify a fixed version of it.

The minexplib example is an implementation of a privacy-friendly protocol for smart electricity meters [27] developed at Microsoft Research. The model that we obtained uncovered a flaw shown in fig. 12: incorrect use of pointer dereferencing results in three bytes of each four-byte reading being sent unencrypted. We found two further flaws: one could lead to contents of uninitialised memory being sent on the network, the other resulted in 0 being sent (and accepted) in place of the actual number of readings. All flaws have been acknowledged and fixed. An F# implementation of the protocol has been previously verified [28], which highlights the fact that C implementations can be tricky and can easily introduce new bugs, even for correctly specified and proven protocols. The protocol uses low-level cryptographic operations such as XOR and modular exponentiation. Under certain restrictions these primitives can be modelled in ProVerif [22, 23], but in general it is impossible to obtain a computational soundness result for XOR [29]. Thus we could not use ProVerif to verify the metering protocol, but we are investigating the use of CryptoVerif for this purpose.

8. CONCLUSION

We presented methods and tools for the automated verification of cryptographic security properties of protocol implementations in C. More specifically, we provided a computationally sound verification of weak secrecy and authentication for (single execution paths of) C code. Despite the limitation of analysing single execution paths, the method often suffices to prove security of authentication protocols, many of which are non-branching. We plan to extend the analysis to more sophisticated control flow.

In future, we aim to provide better feedback in case verification fails. In our case this is rather easy to do as symbolic execution proceeds line by line. If a condition check fails for a certain symbolic expression, it is straightforward to print out a computation tree for the expression together with source code locations in which every node of the tree was computed.

Acknowledgements. Discussions with Bruno Blanchet, François Dupressoir, Bashar Nuseibeh, and Dominique Unruh were useful. We also thank George Danezis, François Dupressoir, and Jean Goubault-Larrecq for giving us access to the code of minexplib, RPC, and CSur, respectively. Bruno Blanchet, Ricardo Corin and François Dupressoir commented on a draft.

9. REFERENCES

- [1] M. Abadi and C. Fournet. Mobile values, new names, and secure communication. In *ACM POPL*, pages 104–115, 2001.
- [2] M. Aizatulin, A. D. Gordon, and J. Jürjens. Extracting and verifying cryptographic models from C protocol code by symbolic execution (long version). Available at <http://arxiv.org/abs/1107.1017>, 2011.
- [3] J. B. Almeida, M. Barbosa, J. S. Pinto, and B. Vieira. Deductive verification of cryptographic software. In *NASA Formal Methods Symposium 2009*, 2009.
- [4] A. Armando, D. A. Basin, Y. Boichut, Y. Chevalier, L. Compagna, J. Cuéllar, P. H. Drielsma, P.-C. Héam, O. Kouchnarenko, J. Mantovani, S. Mödersheim, D. von Oheimb, M. Rusinowitch, J. Santiago, M. Turuani, L. Viganò, and L. Vigneron. The AVISPA tool for the automated validation of internet security protocols and applications. In *CAV*, volume 3576 of *Lecture Notes in Computer Science*, pages 281–285. Springer, 2005.
- [5] M. Backes, D. Hofheinz, and D. Unruh. CoSP: A general framework for computational soundness proofs. In *ACM CCS 2009*, pages 66–78, November 2009. Preprint on IACR ePrint 2009/080.
- [6] J. Bengtson, K. Bhargavan, C. Fournet, A. D. Gordon, and S. Maffei. Refinement types for secure implementations. In *CSF '08: Proceedings of the 2008 21st IEEE Computer Security Foundations Symposium*, pages 17–32. IEEE Computer Society, 2008.
- [7] B. Blanchet. A computationally sound mechanized prover for security protocols. In *IEEE Symposium on Security and Privacy*, pages 140–154. IEEE Computer Society, 2006.
- [8] B. Blanchet. Automatic verification of correspondences for security protocols. *Journal of Computer Security*, 17(4):363–434, 2009.
- [9] B. Blanchet, M. Abadi, and C. Fournet. Automated verification of selected equivalences for security protocols. *Journal of Logic and Algebraic Programming*, 75(1):3–51, Feb.–Mar. 2008.
- [10] M. Blum and S. Micali. How to generate cryptographically strong sequences of pseudo-random bits. *SIAM J. Comput.*, 13(4):850–864, 1984.
- [11] J. Burnim and K. Sen. Heuristics for scalable dynamic test generation. In *ASE '08: Proceedings of the 2008 23rd IEEE/ACM International Conference on Automated Software Engineering*, pages 443–446. IEEE Computer Society, 2008.
- [12] C. Cadar, D. Dunbar, and D. Engler. KLEE: Unassisted and automatic generation of high-coverage tests for complex systems programs. In *USENIX Symposium on Operating Systems Design and Implementation (OSDI 2008)*, San Diego, CA, Dec. 2008.
- [13] S. Chaki and A. Datta. SPIER: An automated framework for verifying security protocol implementations. In *Computer Security Foundations Workshop*, pages 172–185, 2009.
- [14] R. Corin and F. A. Manzano. Efficient symbolic execution for analysing cryptographic protocol implementations. In *International Symposium on Engineering Secure Software and Systems (ESSOS'11)*, LNCS. Springer, 2011.
- [15] P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Conference Record of the Fourth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 238–252, Los Angeles, California, 1977. ACM Press, New York, NY.
- [16] D. Dolev and A. Yao. On the Security of Public-Key Protocols. *IEEE Transactions on Information Theory*, 29(2):198–208, 1983.
- [17] B. Dutertre and L. D. Moura. The Yices SMT Solver. Technical report, 2006.
- [18] P. Godefroid, M. Y. Levin, and D. A. Molnar. Automated whitebox fuzz testing. In *Proceedings of the Network and Distributed System Security Symposium, NDSS 2008*. The Internet Society, 2008.
- [19] S. Goldwasser and S. Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28:270–299, 1984.
- [20] J. Goubault-Larrecq and F. Parrennes. Cryptographic protocol analysis on real C code. In *Proceedings of the 6th International Conference on Verification, Model Checking and Abstract Interpretation (VMCAI'05)*, volume 3385 of *Lecture Notes in Computer Science*, pages 363–379. Springer, 2005.
- [21] J. C. King. Symbolic execution and program testing. *Commun. ACM*, 19(7):385–394, 1976.
- [22] R. Küsters and T. Truderung. Using ProVerif to Analyze Protocols with Diffie-Hellman Exponentiation. In *Proceedings of the 22nd IEEE Computer Security Foundations Symposium (CSF 2009)*, pages 157–171. IEEE Computer Society, 2009.
- [23] R. Küsters and T. Truderung. Reducing protocol analysis with XOR to the XOR-free case in the horn theory based approach. *Journal of Automated Reasoning*, 46(3):325–352, 2011.
- [24] G. Lowe. An attack on the Needham-Schroeder public-key authentication protocol. *Inf. Process. Lett.*, 56:131–133, November 1995.
- [25] G. C. Necula, S. McPeak, S. P. Rahul, and W. Weimer. CIL: Intermediate Language and Tools for Analysis and Transformation of C Programs. In *Proceedings of the 11th International Conference on Compiler Construction, CC '02*, pages 213–228, London, UK, 2002. Springer-Verlag.
- [26] Project EVA. Security protocols open repository, 2007. <http://www.lsv.ens-cachan.fr/spore/>.
- [27] A. Rial and G. Danezis. Privacy-friendly smart metering. Technical Report MSR-TR-2010-150, 2010.
- [28] N. Swamy, J. Chen, C. Fournet, K. Bhargavan, and J. Yang. Security programming with refinement types and mobile proofs. Technical Report MSR-TR-2010-149, 2010.
- [29] D. Unruh. The impossibility of computationally sound XOR, July 2010. Preprint on IACR ePrint 2010/389.
- [30] A. C. Yao. Theory and application of trapdoor functions. In *SFCS '82: Proceedings of the 23rd Annual Symposium on Foundations of Computer Science*, pages 80–91. IEEE Computer Society, 1982.