An Automatically Verified Prototype of the Tokeneer ID Station Specification

Maximiliano Cristiá and Gianfranco Rossi

Abstract—The Tokeneer project was an initiative set forth by the National Security Agency (NSA, USA) to be used as a demonstration that developing highly secure systems can be made by applying rigorous methods in a cost effective manner. Altran Praxis (UK) was selected by NSA to carry out the development of the Tokeneer ID Station. The company wrote a Z specification later implemented in the SPARK Ada programming language, which was verified using the SPARK Examiner toolset. In this paper, we show that the Z specification can be easily and naturally encoded in the $\{log\}$ set constraint language, thus generating a functional prototype. Furthermore, we show that $\{log\}$'s automated proving capabilities can discharge all the proof obligations concerning state invariants as well as important security properties. As a consequence, the prototype can be regarded as correct with respect to the verified properties. This provides empirical evidence that Z users can use {log} to generate correct prototypes from their Z specifications. In turn, these prototypes enable or simplify some verification activities discussed in the paper.

Index Terms—Tokeneer ID Station specification, Z notation, $\{log\}$, constraint programming, prototyping.

arXiv:2009.00999v1 [cs.SE] 2 Sep 2020

1 Introduction

Formal methods (FM) are still questioned with respect to their actual value to deliver software at a reasonable cost. Formal methods researchers and practitioners have proved many times that formal methods can deliver software of unmatched quality, e.g. [1], [2]. The FM community has also shown that high quality is not an impediment to keep costs low when total cost of ownership and critical systems are considered [3], [4]. However, most of the software industry is still reluctant to apply FM, and frequently is unaware of their possible value. Even critical software providers do not always apply FM and are not obligated to do so as standards do not always mandate a formal approach, e.g. IEC 61508 Safety Integrity Level 4.

In this context the National Security Agency (NSA) of the U.S.A. conducted a project aiming at providing evidence to software vendors that formal techniques can deliver high quality software in time and within budget. In particular NSA chose the Tokeneer ID Station (TIS) which makes part of the Tokeneer system. Tokeneer provides protection to secure information held on a network of workstations situated in a physically secure enclave. TIS, in turn, is a standalone trusted entity responsible for performing biometric verification of the user. NSA asked Altran Praxis (then Praxis Critical Systems) to provide an implementation of TIS conforming to Common Criteria's EAL5-level [5] and to disclose to the public domain all the deliverables produced during system development. As said, the ultimate goal of NSA was to show that this kind of development efforts are feasible (i.e., they achieve reliable software) and cost-effective (i.e., they are not more expensive than traditional development processes) [6], [7].

Altran Praxis applied its own Correctness by Construction development process to the TIS software. The second phase of this development process consists in writing a Z [8], [9] formal specification of the user requirements. The Z specification is central to the development process as it is used as the correctness criteria for many verification activities.

The work described in this paper starts from this Z specification. More precisely, we first *encode* the Z specification in the $\{log\}$ set constraint programming language [10], [11], [12]. This provides a functional prototype of the Z specification. We say 'encode' and not 'implement' due to the close resemblance between the $\{log\}$ language and Z; however, the encoding provides an implementation in the form of a prototype. In a second step, we use $\{log\}$'s constraint solving capabilities to automatically prove that the prototype verifies *all the state invariants* defined in the Z specification as well as *all but one of the security properties* stated by Altran Praxis team—this amounts to discharge 523 proof obligations. This constitutes a step forward with respect to the original project as these invariants and properties were not machine-checked by Altran Praxis at the specification level. Besides, this implies that the $\{log\}$ program is an automatically verified functional prototype of the Z specification. This provides empirical evidence that Z users can use {log} to generate correct prototypes from their Z specifications. The paper discusses some verification activities that can be carried out or simplified once the $\{log\}$ prototype is available.

The present work aims at providing more empirical evidence that FM-based tools such as $\{log\}$ can effectively be used in industrial projects willing to provide high quality software without incurring in increased costs or delayed schedules.

The paper assumes the reader has some exposure to

- *M. Cristiá is with Universidad Nacional de Rosario and CIFASIS, Argentina.*
  E-mail: cristia@cifasis-conicet.gov.ar
- *G. Rossi is with Universitá di Parma, Italy.*
  E-mail: gianfranco.rossi@unipr.it

to be completed

Z specification—otherwise the reader can consult any Z textbook, e.g. [9].

The paper is structured as follows. The Z specification of the Tokeneer project is briefly introduced in Section 2, along with the security properties the system should enforce. $\{log\}$ is presented in Section 3 by means of examples. Section 4 discusses the encoding of the Z specification in $\{log\}$. Section 5 shows how the $\{log\}$ prototype is verified by running automated proofs of state invariants and security properties; Section 5.3 discusses some other verification activities that can be done with the prototype. A quantitative report of some key aspects of this endeavor is presented in Section 6. In Section 7 we put our results in the context of other works that used the Tokeneer project as a case study. In Section 8 we present our conclusions.

2 The Z Specification of the TIS

The Tokeneer project carried out be Altran Praxis has been thoroughly documented in part due to the requirements of the Common Criteria [6], [13]. In this section we will focus on the Z specification generated during the project. The goal is for the reader to have an idea of the complexity and peculiarities of the Z specification. We will not explain it in full in part because, precisely, encoding it in $\{log\}$ can be done in a completely formal manner; understanding what the specification is about is unnecessary.

The Z specification is a 117 pages long document containing formal Z code plus informal explanatory statements about it [14]. There is also an 11 pages long document stating security properties the Z specification must verify [15]. Whenever we refer to "the (Z) specification" we mean these two documents, unless stated differently; and by "the team" we mean the Altran Praxis development team that worked out the specification.

The Z specification is written in a more or less standard fashion. However, it starts by introducing two polymorphic operators working in tandem that are somewhat unusual:

$$\mathrm{optional}\, X == \{x : \mathbb{F}\, X \mid \#x \leq 1\}$$
$$the[X] == \{x : X \bullet \{x\} \mapsto x\}$$

optional is used to indicate that a variable can hold a value or nil, as for example in:

Certificate
$id : CertificateId$
$validityPeriod : \mathbb{P}\, TIME$
$isValidatedBy : \mathrm{optional}\, KEYPART$

meaning that the certificate can be validated by some (asymmetric cryptographic) key or it cannot. When the key does exist and has to be retrieved from the certificate, then the *the* operator comes into play; for instance:

$$the\, cert.isValidatedBy = issuerCert.subjectPubK$$

Concerning the encoding of the Z specification in $\{log\}$, it is important to note that optional can be defined as follows:

$$\mathrm{optional}\, X == \{x : \mathbb{F}\, X \mid x = \emptyset \lor (\exists y : X \mid x = \{y\})\}$$

Likely, the team opted by the first definition in an attempt to keep the number of quantified formulas as low as possible. As we will show, introducing existentially quantified variables in $\{log\}$ is harmless, while introducing the cardinality operator is not.

Z schemas are heavily used to give structure to the main concepts formalized in the specification. For example, *Certificate* is used to define:

AttCertificate
$Certificate$
$baseCertId : CertificateId$
$tokenID : TOKENID$

which in turn is used to define:

PrivCert
$AttCertificate$
$role : PRIVILEGE$
$clearance : Clearance$

which is used to define a *Token*:

Token
$tokenID : TOKENID$

$idCert : IDCert$
$privCert : PrivCert$
$iandACert : IandACert$
$authCert : \mathrm{optional}\, AuthCert$

It can be said that the specification is divided into two parts: the real world peripherals interacting with the TIS and the TIS itself. The state of all real world entities is modeled with the *RealWorld* schema which is divided into two schemas, *TISControlledRealWorld* and *TISMonitoredRealWorld*. In total *RealWorld* comprises 11 state variables. In turn, the TIS state is modeled in schema *IDStation* which includes 12 schemas representing different subsystems. In total *IDStation* declares 36 state variables. The definition of *IDStation* follows the style of most Z specifications[1]:

IDStation
$\ldots$
$DoorLatchAlarm$
$\ldots$
$currentDisplay : DISPLAYMESSAGE$
$currentScreen : Screen$

$status \in \{gotFinger, waitingFinger, \ldots\} \Rightarrow$
$\quad ((\exists\, ValidToken \bullet goodT(\theta ValidToken) = \ldots)$
$\quad\quad \lor\ (\exists\, TokenWithValidAuth \bullet \ldots))$
$\ldots$
$currentScreen.screenStats = displayStats(\theta Stats)$

That is, it includes several schemas, declares two variables and 8 state invariants. Note that these state invariants are conjoined with those declared in some of the included schemas (e.g. *DoorLatchAlarm*). This implies a total of 13 state invariants. The reader can see two of them in the summary of the schema shown above. The first one gives an idea of the complexity of some of the predicates: it is a

1. We use ellipses to shorten the presentation.

quantified formula using Z's θ operator. The θ operator is one of the most complex logical operators of the Z notation.

Some of the state variables in *RealWorld* and *IDStation* have a non-enumerated free type, as for example:

$$KEYBOARD ::= noKB \mid badKB \mid keyedOps\langle\langle \ldots \rangle\rangle$$

which means the presence of structured infinite types.

The specification defines 25 major operations plus 3 that group some of these operations (e.g. *TISUserEntryOp* specifying the complete authentication process)[2]. All TIS operations are state transitions over *RealWorld* and *IDStation*. Many of these operations are assembled from simpler operation schemas by means of some non-trivial schema expressions; for instance:

$$TISArchiveLogOp \mathrel{\widehat=} StartArchiveLog \lor \ldots$$

$$\begin{aligned} &StartArchiveLog \mathrel{\widehat=} \\ &\quad (StartArchiveLogOK \fatsemi UpdateFloppy) \\ &\quad \lor StartArchiveLogWaitingFloppy \\ &\quad \lor [\, BadAdminLogout \mid \ldots \,] \end{aligned}$$

There is a particularly important operation schema, namely *TISProcessing*, specifying the overall processing activity of the TIS because it is used to state some important security properties as we will see. *TISProcessing* "calls" 20 TIS operations and updates the audit log:

$$\begin{aligned} &TISProcessing \mathrel{\widehat=} \\ &\quad (\cdots \lor TISUserEntryOp \lor \cdots \lor TISAdminOp \lor \ldots) \\ &\quad \land LogChange \end{aligned}$$

2.1 State Invariants

The specification follows the Z style concerning the encoding of state invariants. This means that state invariants are declared in the predicate part of the state schema (i.e., *IDStation*). Including the state invariants in the state schema implies that all operations trivially verify them. We exemplify this with a simple example not taken from the specification.

$$State \mathrel{\widehat=} [x : \mathbb{Z} \mid 0 \leq x]$$

$$Op \mathrel{\widehat=} [\Delta State \mid x' = x - 1]$$

Op cannot violate the state invariant ($0 \leq x$) because the state invariant is part of Op' definition:

$$\begin{aligned} &[\Delta State \mid x' = x - 1] \\ &\equiv [x, x' : \mathbb{Z} \mid 0 \leq x \land x' = x - 1 \land 0 \leq x'] \end{aligned}$$

So Op cannot be called when $x = 0$ because Op's predicate implies $0 < x$. This is called *implicit precondition* and raises due to the interaction between the state invariant and the operation predicate.

Another approach concerning state invariants is to prove that each operation verifies them. In this case the state invariant is not encoded in the state schema, preconditions are explicitly stated and a proof obligation is introduced:

$$State \mathrel{\widehat=} [x : \mathbb{Z}]$$

$$Inv \mathrel{\widehat=} [State \mid 0 \leq x]$$

$$Op \mathrel{\widehat=} [\Delta State \mid 0 < x \land x' = x - 1]$$

proof obligation $Inv \land Op \Rightarrow Inv'$

2. These are the operation schemas whose names begin with *TIS*.

This second approach has the advantage that all preconditions are explicit making the transition to the implementation simpler, but it has the disadvantage of having to discharge proof obligations.

In encoding the Z specification in $\{log\}$ we opted for this second approach because we want the $\{log\}$ prototype to be closer to an implementation (than the Z specification) and because $\{log\}$ can automatically discharge these proof obligations (see Section 5).

2.2 Security Properties

The team stated 6 security properties the specification must verify; one of them is not formalized, so we will not consider it. The team provided informal proofs of these properties. According to the documentation the most important of these properties is Property 1; we reproduce an excerpt of it here for the reader to have an idea of its complexity:

$$\begin{aligned} &\Delta IDStation;\ \Delta RealWorld \mid \\ &\quad TISOpThenUpdate \\ &\quad \land latch = locked \land latch' = unlocked \\ &\vdash \\ &\quad (\exists\, ValidToken \bullet goodT(\theta ValidToken) = \ldots \\ &\qquad \land UserTokenOKNoCurrCheck \\ &\qquad \land FingerOK) \\ &\quad \lor \\ &\quad (\exists\, TokenWithValidAuth \bullet goodT(\ldots) = \ldots \\ &\qquad \land UserTokenWithOKAuthCertNoCurrCheck) \\ &\quad \lor \\ &\quad (\exists\, ValidToken \bullet goodT(\theta ValidToken) = \ldots \\ &\qquad \land authCert \neq \emptyset \land (the\, authCert).role = guard) \end{aligned}$$

where:

$$TISOpThenUpdate \mathrel{\widehat=} TISProcessing \fatsemi TISUpdate$$

That is, the property must be proved for 20 of the TIS operations, followed by the *TISUpdate* operation.

It is important to remark that the above formula does not really correspond to the property the team would have wanted to prove—we call it the *intended property*. This is acknowledged by team in the technical documentation [15, page 11, 'Note on arguments']. The formula formalizing the intended property cannot be expressed in the Z notation. Indeed, that formula would have to predicate over (infinite) sequences of states instead of pairs of states as stated by the semantics of the Z notation [16], [17]. It is easy to observe that TIS is actually a reactive system perhaps making Z not the best notation to specify it.

We remark this point because we make an extra effort to overcome this limitation when proving the validity of Property 1 for the $\{log\}$ prototype (see Section 5.2.1). Hence, we end up producing not only a mechanized (automated) proof of Property 1 but also we produce proofs of several other properties that together provide much stronger arguments that the intended property actually holds.

3 The $\{log\}$ Constraint Solver

$\{log\}$ is a public available satisfiability solver and a set-based, constraint-based programming language implemented in Prolog [12]. $\{log\}$ implements a decision procedure for the theory of *hereditarily finite sets*, i.e., finitely

nested sets that are finite at each level of nesting [10]; a decision procedure for a very expressive fragment of the class of finite set relation algebras [11], [18]; a decision procedure for restricted intensional sets (RIS) [19], [20]; and uses Prolog's CLP(Q) to provide a decision procedure for integer linear arithmetic [21]. In $\{log\}$ sets and binary relations are first-class entities of the language. At the core of these decision procedures is set unification [22]. The set terms defined in all these three decision procedures can be combined in several ways: binary relations are hereditarily finite sets whose elements are ordered pairs and so set operators can take binary relations as arguments; RIS can be passed as arguments to set operators and freely combined with extensional sets. $\{log\}$ is an untyped formalism; variables are not declared; typing information can be encoded by means of constraints. Several in-depth empirical evaluations provide evidence that $\{log\}$ is able to solve non-trivial problems [11], [18], [19], [20], [23], in particular as an automated verifier of security properties [24]. Given that $\{log\}$ has been extensively described elsewhere, in this section we will show a few examples for the reader to understand how it works.

In $\{log\}$ set operators are encoded as constraints. For example: $\mathsf{un}(A, B, C)$ is a constraint interpreted as $C = A \cup B$. $\{log\}$ implements a wide range of set and relational operators covering most of those used in Z. For instance, in is a constraint interpreted as set membership (i.e., $\in$); $=$ is set equality; $\mathsf{dom}(F, D)$ corresponds to $\mathrm{dom}\, F = D$; $\mathsf{subset}(A, B)$ corresponds to $A \subseteq B$; $\mathsf{comp}(R, S, T)$ is interpreted as $T = R \fatsemi S$ (i.e., relational composition); and $\mathsf{apply}(F, X, Y)$ is equivalent to $\mathsf{pfun}(F) \;\&\; [X, Y]\ \mathsf{in}\ F$, where $\mathsf{pfun}(F)$ constrains F to be a (partial) function. Formulas in $\{log\}$ are conjunctions ($\&$) and disjunctions (or) of constraints; they must finish with a dot (as a Prolog query). Negation in $\{log\}$ is introduced by means of so-called *negated constraints*. For example $\mathsf{nun}(A, B, C)$ is interpreted as $C \neq A \cup B$ and nin corresponds to $\notin$—in general, a constraint beginning with 'n' identifies a negated constraint. For formulas to lay inside the decision procedures implemented in $\{log\}$, users must only use this form of negation.

The fact that set operators take a relational rather than a functional form makes it necessary to introduce variables to write compound expressions.

Example 1. The Z expression $x \in A \cap B \cap C$ is encoded as $X\ \mathsf{in}\ M \;\&\; \mathsf{inters}(A, B, W) \;\&\; \mathsf{inters}(W, C, M)$. □

Set terms can be of the following forms:

- A variable is a set term; variable names must start with an uppercase letter.
- $\{\}$ is the term interpreted as the empty set.
- $\{x/A\}$ is called *extensional set* and is interpreted as $\{x\} \cup A$; A must be a set term, x can be any term accepted by $\{log\}$ (basically, any Prolog uninterpreted symbol, integers, lists, ordered pairs, etc.).
- $\mathsf{ris}(X\ \mathsf{in}\ A, \phi)$ is called *restricted intensional set* (RIS) and is interpreted as $\{x : x \in A \wedge \phi\}$ where ϕ is any $\{log\}$ formula; A must be a set term and X is a bound variable local to the RIS. Actually, RIS have a more complex and expressive structure [19], [20].

Being a satisfiability solver, $\{log\}$ can be used as an automated theorem prover. To prove that formula ϕ is a theorem, $\{log\}$ has to be called to prove that $\neg\ \phi$ is unsatisfiable.

Example 2. We can prove that set union is commutative by asking $\{log\}$ to prove the following is unsatisfiable:

$$\mathsf{un}(A, B, C) \;\&\; \mathsf{un}(B, A, D) \;\&\; C\ \mathsf{neq}\ D.$$

As there are no sets satisfying this formula $\{log\}$ answers no. Note that the formula can also be written with the nun constraint: $\mathsf{un}(A, B, C) \;\&\; \mathsf{nun}(B, A, C)$. □

$\{log\}$ is also a programming language at the intersection of declarative programming, set programming [25] and constraint programming. $\{log\}$ programs can be structured by means of *clauses*—as in Prolog. A clause can be seen as a subroutine or procedure. Clauses can receive zero or more arguments. The only way a clause can return a value is by means of one or more of its arguments. Under certain conditions clauses behave as formulas. That is a $\{log\}$ clause can be seen as both a program and a formula. The following examples show the *formula-program duality* of $\{log\}$ code along with the notion of clause.

Example 3. If we want a program that updates function F in X with value Y provided X belongs to the domain of F and get an error otherwise, the $\{log\}$ code can be the following:

$$\begin{array}{l}
\mathsf{update}(F, X, Y, F_, Error)\ \text{:-} \\
\quad F = \{[X, V]/F1\} \;\&\; [X, V]\ \mathsf{nin}\ F1 \;\& \\
\quad F_ = \{[X, Y]/F1\} \;\& \\
\quad Error = ok \\
\quad \mathsf{or} \\
\quad \mathsf{comp}(\{[X, X]\}, F, \{\}) \;\& \\
\quad Error = err.
\end{array}$$

That is, update is a clause that receives F, X and Y and returns the modified F in $F_$ and the error code in $Error$—think of $F_$ as the value of F in the next state. As $\&$ and or are logical connectives and $=$ is logical equality, the order of the 'instructions' is irrelevant w.r.t. the functional result—although it can have an impact on the performance. Variable $F1$ is an existentially quantified variable representing the 'rest' of F w.r.t. to $[X, V]$. If $[X, V]$ does not belong to F then the unification between F and $\{[X, V]/F1\}$ will fail thus making update to execute the other branch.
Now we can call update by providing inputs and waiting for outputs:

$$\begin{array}{l}
\mathsf{update}(\{[setlog, 5], [hello, earth], [tokeneer, model]\}, \\
\qquad\qquad hello, world, G, E).
\end{array}$$

returns:

$$\begin{array}{l}
G = \{[hello, world], [setlog, 5], [tokeneer, model]\} \\
E = ok
\end{array}$$

□

Since update is also a formula we can prove properties true of it.

Example 4. If *Error* is equal to *err* then X does not belong to the domain of F. In order to prove this property we need to call $\{log\}$ on its negation:

$$\mathsf{update}(F, X, Y, F_, err) \;\; \& \;\; \mathsf{dom}(F, D) \;\; \& \;\; X \,\mathsf{in}\, D.$$

Then, $\{log\}$ answers no because the formula is unsatisfiable. □

As Example 3 shows, variables introduced in the clause body (e.g. $F1$) are existentially quantified variables. This is important because of the way existentially quantified formulas of the Z specification are encoded in $\{log\}$.

$\{log\}$ implements set unification but syntactic Prolog unification is still available as part of it. Prolog unification comes handy to encode some Z features and predicates.

Example 5. The 'dot' notation used in Z to access components of ordered pairs and variables of schema types, can be encoded by means of Prolog unification. For instance if $x : A \times B$, then the Z expression: $h = x.1$ can be encoded as $X = [X1, X2] \;\; \& \;\; H = X1$ or even as $X = [H, X2]$. Further, since $X2$ seems to be uninteresting yet another encoding is $X = [H, _]$. □

Example 6. A more elaborated example is the encoding of the Z predicate $x \in \mathrm{dom} f$ as $[X, _]\mathsf{in}F$. Indeed, $[X, _]\mathsf{in}F$ is readily rewritten as $F = \{[X, _]/F1\}$ which is interpreted as 'there is a pair in F whose first component is X', which in turn means that X belongs to the domain of F. A more direct encoding is $\mathsf{dom}(F, D) \;\; \& \;\; X \,\mathsf{in}\, D$ but it requires to compute the domain while the first encoding does not, meaning that the first encoding will, in general, yield more efficient $\{log\}$ code. Yet another encoding is $\mathsf{ncomp}(\{[X, X]\}, F, \{\})$ which uses the negated constraint of the comp constraint. That is, $\mathsf{ncomp}(\{[X, X]\}, F, \{\})$ is equivalent to $\mathsf{comp}(\{[X, X]\}, F, F1) \;\; \& \;\; F1 \,\mathsf{neq}\, \{\}$, which is a fourth encoding. □

As can be seen, frequently, unification introduces existentially quantified variables. In many circumstances these quantified formulas can be dealt with in a decidable manner [10], [11], [18], [19], [20]. A key aspect to preserve decidability is not to use logical negation but the negated constraints provided by $\{log\}$.

Example 7. Concerning formula $x \in \mathrm{dom} f$ of Example 6, the encoding of its negation, $x \notin \mathrm{dom} f$, in {log} cannot be done simply with a formula such as $\forall Y : [X, Y] \,\mathsf{nin}\, F$ because this formula is outside the set of admissible {log} formulas. Conversely, the decidable way to handle this negation is either: $\mathsf{dom}(F, D) \;\; \& \;\; X \,\mathsf{nin}\, D$; or $\mathsf{comp}(\{[X, X]\}, F, \{\})$. □

The last example opens the issue of universal quantifiers in $\{log\}$. In $\{log\}$ universally quantified formulas are provided by mean of RIS. In effect, the introduction of RIS in $\{log\}$ allows for the definition of *restricted universal quantifiers* (RUQ). In general, if A is a set, then a RUQ is a formula of the following form:

$$\forall x \in A : \phi$$

It is easy to prove the following:

$$(\forall x \in A : \phi) \Leftrightarrow A \subseteq \{x : x \in A \wedge \phi\} \qquad (1)$$

Given that $\{x : x \in A \wedge \phi\}$ is the interpretation of $\mathsf{ris}(X \,\mathsf{in}\, A, \phi)$, the r.h.s. of (1) can be expressed as the $\{log\}$ formula:

$$\mathsf{subset}(A, \mathsf{ris}(X \,\mathsf{in}\, A, \phi))$$

In $\{log\}$ we have defined the foreach constraint to make RUQ easier to write:

$$\mathsf{foreach}(X \,\mathsf{in}\, A, \phi) \text{ :- } \mathsf{subset}(A, \mathsf{ris}(X \,\mathsf{in}\, A, \phi)).$$

We use these features to encode the Z specification in $\{log\}$, to automatically prove invariance lemmas and properties, and to provide a correct prototype of the Z specification in the form of a $\{log\}$ program.

4 Encoding the TIS Specification in $\{log\}$

In the introduction, we say 'encoding' and not 'implementing' the TIS specification due to the close resemblance between the Z and $\{log\}$ languages. That is, our point is that writing $\{log\}$ code from the specification is considerably more natural, evident and semantically equivalent than writing, say, SPARK code. Then, it looks like more as an encoding than as an implementation. In particular, $\{log\}$ implements all the logical, set and relational operators used in the TIS specification. Furthermore, these operators are not mere imperative implementations but real executable mathematical definitions. That is, they behave as logical or mathematical objects, as we have shown in Section 3. In other words, the $\{log\}$ prototype is a formula *quite* as the TIS specification is. We say 'quite' because of two reasons:

- The $\{log\}$ code can be used to *compute* results. Then, $\{log\}$ programmers may pay attention to implementation issues, such as efficiency.
- For different programming reasons (e.g. use of Prolog and set unification) we introduce some modifications in the $\{log\}$ code w.r.t. the specification.

This takes the discussion to another issue: the $\{log\}$ code is a *prototype* and not a *program*. Hence, users cannot expect the same computing efficiency from the $\{log\}$ code than from a typical imperative implementation. For instance, engineers can use the prototype to analyze functional scenarios but they cannot draw efficiency estimations. This section will make these points clear.

The $\{log\}$ encoding of the TIS specification as well as all the proof obligations can be found online: http://people.dmi.unipr.it/gianfranco.rossi/SETLOG/APPLICATIONS/tokeneer.zip.

In general, each Z schema is encoded as a $\{log\}$ clause. The variables declared in the schema become arguments of the clause. In many cases, if the schema declares variables through schema inclusion then the arguments of the clause are those schemas instead of the variables declared inside them. This somewhat preserves the structure of the specification. Schemas whose predicate part is empty are encoded as tuples of variables. We preserved the identifiers used in the specification as much as possible in the $\{log\}$ code. Recall that $\{log\}$ variables (constants) must start with an uppercase (lowercase) letter; while clauses can start only with a lowercase letter. Next state variables (e.g. x') are

encoded as {*log*} variables decorated with an underscore (e.g. *X_*)[3].

Figures 1 and 2 show parts of the encoding. In Figure 1 the specification is at the left and the corresponding encoding in {*log*}, at the right. We attempted to align each row of the specification with the corresponding row in {*log*}. In Figure 2 the specification is at the top and the encoding at the bottom. In general, the encoding in {*log*} is longer than the corresponding Z code, as is expected for any lower-level representation, but not that much (see Section 6).

Consider Figure 1. As can be seen, schema *CurrentToken* is encoded as clause `currentToken`. *CurrentToken* declares several variables through the inclusion of schema *ValidToken* which in turns declares those variables through the inclusion of schema *Token*. Then, in this case `currentToken` has two arguments: `Token`, corresponding to schema *Token*; and `Now`, corresponding to variable *now*. In order to make `Token` a valid one, `validToken(Token)` is conjoined. As expected, `validToken` is the encoding of schema *ValidToken* (not shown). In the specification *TIME* is a synonym for $\mathbb{N}$, so `Now` is constrained to be a non-negative integer by asserting `0 =< Now`. Given that *Token* is a schema declaring five variables it is encoded as a 5-tuple: `Token = [_,IDC,PC,IAC,_]` where the correspondence between Z and {*log*} variables is implemented by declaration order. Thus, `PC` corresponds to the third variable declared in *Token*, i.e., *privCert*. In `currentToken` only `IDC`, `PC` and `IAC` play some role, so the other two are hidden by putting underscores in their positions. The named variables can have any name because they are existentially quantified inside the clause. Besides, `Token = [_,IDC,PC,IAC,_]` forces the unification of the actual argument to be a 5-tuple; in case this is not true the clause will fail. The specification variables corresponding to `IDC`, `PC` and `IAC` are of schema types (e.g. *privCert* : *PrivCert*, see Section 2), so they are encoded as tuples, e.g. `PC = [_,PVP,_,_,_,_,_]`. In this particular case `PVP` corresponds to the Z expression *privCert.validityPeriod* (recall Example 5). Therefore, the predicate stated in *CurrentToken* is encoded as shown in Example 1.

Still in Figure 1, *UserTokenOK*'s predicate is based on a quantified schema expression using the θ operator. The type of *currentUserToken* is:

$$TOKENTRY ::= noT \mid badT \mid goodT\langle\langle Token\rangle\rangle$$

Then, $currentUserToken \in \mathrm{ran}\, goodT$ means that the current user token is indeed a *Token*. That is, the predicate is equivalent to $\exists t : Token \bullet currentUserToken = goodT(t)$. Avoiding quantified formulas is important but the encoding can introduce an existentially quantified variable without getting into troubles: `CurrentUserToken = goodT(_)`. Along the same lines, the predicate $(\exists CurrentToken \bullet \ldots)$ is easily encoded with `currentToken(Token,Now)`; that is, by asserting the existence of variables `Token` and `Now` satisfying `currentToken`. Then, $\theta ValidToken$ refers to the *Token* asserted in $(\exists CurrentToken \bullet \ldots)$. So, `goodT(Token) = CurrentUserToken` is asserted in {*log*}. This implies that `goodT(_) = goodT(Token)`, thus making `CurrentUserToken = goodT(_)` superfluous—actually, we included it to make the encoding more similar to the specification. Finally, the last three existentially quantified predicates are encoded along the same lines. That is, the encoding of schema *CertOK* is called on each certificate contained in the *Token*. Note that, for instance, $(\exists IDCert \ldots)$ corresponds to variable `IDC`; thus $\theta IDCert = idCert$ corresponds to `IDC = [[Id,VP,IVB],[_,_]]`; and `[Id,VP,IVB]` is the *Certificate* passed in as argument to `certOK`.

Now consider Figure 2. Schema *BioCheckRequired* specifies part of the *ValidateUserTokenOK* operation which in turn is part of the *TISValidateUserToken* operation. *BioCheckRequired* makes a 'delta' on *IDStation* and *RealWorld*; this is hidden in schema *UserEntryContext*. This is reflected in the encoding as clause `bioCheckRequired` waits for `IDStation` and `IDStation_`, meaning that the clause transitions from the former to the latter. As explained in Section 2, *IDStation* includes 12 schemas and declares two variables. Then, the encoding for *IDStation* is a 14-tuple respecting the declaration order given in the specification. As mentioned in Section 2.1, the invariants stated inside *IDStation* are moved out of it; see Section 5 for further details. Note that *IDStation*′ is encoded as `IDStation_`.

As *BioCheckRequired* states $\Xi UserToken$ the encoding forces `UserToken` and `UserToken_` to be the first elements of `IDStation` and `IDStation_`, respectively, while `UserToken_ = UserToken` is conjoined to the clause.

AddElementsToLog is an operation schema making a 'delta' on the auditing subsystem, *AudiLog*, and accessing the configuration subsystem, *Config*. In {*log*} this schema corresponds to clause `addElementsToLog`. The clause is called inside a `delay` predicate for efficiency reasons. After the first proof obligations were discharged it was evident that `addElementsToLog` was taking too long while not contributing to the proofs. Hence, we ask {*log*} to delay its execution as much as possible. *AddElementsToLog*'s predicate starts by asserting the existence of a finite set of audit records, *newElements*. The {*log*} encoding reflects that fact by declaring variable `_NewElements` inside the clause and passing it to `addElementsToLog`. This simplifies the encoding of some tricky parts of the specification.

In *BioCheckRequired*, *status* is a variable declared in schema *Internal* (which is one of the 12 schemas included in *IDStation*). The encoding uses unification to access the corresponding variable, `Status`. Indeed, unification forces `Internal` to be the 12th component of `IDStation` and it is used again to force `Status` to be the first component of `Internal`. Then, we can state `Status = gotUserToken`. Note how the change of state of *status* is encoded by asserting `Status_ = waitingFinger`.

The predicate $\neg$ *UserTokenWithOKAuthCert* is encoded by calling the negation of `userTokenWithOK AuthCert`. As explained in Section 3 (see Example 7), in order to preserve decidability (general) logical negation should not be used. Instead, the negation has to be written in terms of {*log*} constraints. Given that Figure 1 includes the encoding of *UserTokenOK*, Figure 3 shows the encoding of $\neg$ *UserTokenOK* (in place of $\neg$ *UserTokenWithOKAuthCert*, which is nonetheless similar). There are several reasons for which the user token is not

3. The prime symbol is not allowed as part of a variable name in Prolog+{*log*}.

CurrentToken

ValidToken
now : *TIME*

$now \in idCert.validityPeriod$
$\quad \cap\, privCert.validityPeriod$
$\quad \cap\, iandACert.validityPeriod$

```
currentToken(Token,Now) :-
  Token = [_,IDC,PC,IAC,_] &
    IDC = [[_,VP,_],[_,_]] &
    PC = [_,PVP,_,_,_,_,_] &
    IAC = [_,IVP,_,_,_,_] &

  validToken(Token) &
  0 =< Now &

  Now in M2 &
    inters(VP,PVP,M1) &
    inters(M1,IVP,M2).
```

UserTokenOK

KeyStore
UserToken
currentTime : *TIME*

$currentUserToken \in \operatorname{ran} goodT$

$\exists\, CurrentToken \bullet$

$(goodT(\theta ValidToken) = currentUserToken$
$\wedge\ now = currentTime$
$\wedge\ (\exists\, IDCert \bullet \theta IDCert = idCert \wedge CertOK)$
$\wedge\ (\exists\, PrivCert \bullet \theta PrivCert = privCert \wedge CertOK)$
$\wedge\ (\exists\, IandACert \bullet$
$\qquad \theta IandACert = iandACert \wedge CertOK))$

```
userTokenOK(
    KeyStore,
    UserToken,CurrentTime,Now) :-
  0 =< CurrentTime &
  UserToken = [CurrentUserToken,_] &

  CurrentUserToken = goodT(_) &

  currentToken(Token,Now) &
    Token = [_,IDC,PC,IAC,_] &
      IDC = [[Id,VP,IVB],[_,_]] &
      PC = [PCId,PCVP,PCIVB,_,_,_,_] &
      IAC = [IAId,IAVP,IAIVB,_,_,_] &

    goodT(Token) = CurrentUserToken &
    Now = CurrentTime &
    certOK(KeyStore,[Id,VP,IVB]) &
    certOK(KeyStore,[PCId,PCVP,PCIVB]) &
    certOK(KeyStore,[IAId,IAVP,IAIVB]).
```

Fig. 1. Comparison between Z code and {*log*} code (part 1). Similarities are not only syntactic but semantic as well.

OK. The first one occurs when the current user token does not belong to the range of *goodT*. This implies that it must be either *noT* or *badT* (see *TOKENTRY* above). The encoding is `CurrentUserToken in {noT, badT}`. The other cases occur when there is a *token* : *Token* such that $currentUserToken = goodT(token)$. In this case, $(token, now)$ might not conform a current token (`not_currentToken(Token,Now)`); or *now* might not coincide with *currentTime* (`Now neq CurrentTime`); or any of the certificates stored in *token* is not OK (`not_certOK(KeyStore,[_,_,_])`).

To close this section, note that the {*log*} encoding of *ValidateUserTokenOK* turns out to be quite 'natural' (both, syntactically and semantically). This encoding becomes sort of a pattern by means of which most of the Z specification is encoded. Considering the size and complexity of the TIS Z specification, it can be argued that {*log*} can be used as an effecitive prototyping/programming language to encode many other Z specifications.

4.1 Potential Problems with the Encoding

In spite of the similarities between Z and {*log*}, in passing from the specification to the prototype several changes were introduced. We claim that, still, the prototype is a faithful representation of the specification as it is possible to prove that the former verifies essential properties of the latter—see Section 5. However, some of the changes might cause troubles when proving new properties. In this section we give a brief account of these differences.

Types are not always encoded. Z is a typed formalism, {*log*} is not. Typing information can be encoded using set membership and constraints such as pfun. We encoded the typing information needed to prove the properties discussed in Section 5. Some typing information is given as state invariants—as is nonetheless the case when, for instance, a set is implemented as a list. The encoding assumes the Z specification has been type-checked. Then, the encoding needs to check only the types of input values and the initial state. It would be possible to provide a type-checker for {*log*} programs based on special-purpose typing constraints.

BioCheckRequired

UserEntryContext

$\Xi UserToken$
$\Xi DoorLatchAlarm$
$\Xi Stats$

AddElementsToLog

$status = gotUserToken$
$userTokenPresence = present$

$\neg\ UserTokenWithOKAuthCert \land UserTokenOK$

$currentDisplay' = insertFinger$
$status' = waitingFinger$

$ValidateUserTokenOK \mathrel{\widehat=} BioCheckRequired \lor BioCheckNotRequired$

```
bioCheckRequired(IDStation,RealWorld,RealWorld_,IDStation_) :-
  IDStation = [UserToken,_,_,DoorLatchAlarm,_,_,Config,Stats,
               KeyStore,_,AuditLog,Internal,_,CurrentScreen] &
    UserToken = [_,UserTokenPresence] & DoorLatchAlarm = [CurrentTime,_,_,_,_,_] &
    Internal = [Status,_,_] & CurrentScreen = [ScreenStats,_,ScreenConfig] &
  RealWorld = [_,TISMonitoredRealWorld] &
    TISMonitoredRealWorld = [Now,_,_,_,_,_,_] &
  IDStation_ = [UserToken_,_,_,DoorLatchAlarm_,_,_,_,Stats_,
              _,_,AuditLog_,Internal_,CurrentDisplay_,CurrentScreen_] &
    Internal_ = [Status_,_,_] & CurrentScreen_ = [ScreenStats_,_,ScreenConfig_] &

  userEntryContext(IDStation,RealWorld,RealWorld_,IDStation_) &

  UserToken_ = UserToken &
  DoorLatchAlarm_ = DoorLatchAlarm &
  Stats_ = Stats &

  delay(addElementsToLog(Config,_NewElements,AuditLog,AuditLog_),false) &

  Status = gotUserToken &
  UserTokenPresence = present &

  delay(not_userTokenWithOKAuthCert(KeyStore,UserToken,CurrentTime),false) &
    userTokenOK(KeyStore,UserToken,CurrentTime,Now) &

  CurrentDisplay_ = insertFinger &
  Status_ = waitingFinger &

  ScreenStats_ = ScreenStats &
  ScreenConfig_ = ScreenConfig.

validateUserTokenOK(IDStation,RealWorld,RealWorld_,IDStation_) :-
  bioCheckRequired(IDStation,RealWorld,RealWorld_,IDStation_)
  or
  bioCheckNotRequired(IDStation,RealWorld,RealWorld_,IDStation_).
```

Fig. 2. Comparison between Z code and $\{log\}$ code (part 2). Similarities are not only syntactic but semantic as well.

```
not_userTokenOK(KeyStore,UserToken,
                CurrentTime,Now) :-
  UserToken = [CurrentUserToken,_] &

  (CurrentUserToken in {noT, badT}
   or
   CurrentUserToken = goodT(Token) &
    Token = [_,IDC,PC,IAC,_] &
      IDC = [[Id,VP,IVB],[_,_]] &
      PC = [PCId,PCVP,PCIVB,_,_,_,_] &
      IAC = [IAId,IAVP,IAIVB,_,_,_] &
   (not_currentToken(Token,Now)
    or
    Now neq CurrentTime
    or
    not_certOK(KeyStore,[Id,VP,IVB])
    or
    not_certOK(KeyStore,[PCId,PCVP,PCIVB])
    or
    not_certOK(KeyStore,[IAId,IAVP,IAIVB])
   )
  ).
```

Fig. 3. Encoding of $\neg$ *UserTokenOK*

Elements beyond {log}'s expressiveness. There are a few elements in the specification that are beyond the expressiveness of {*log*}. Schema types cannot be encoded as sets of records. In this particular specification this feature is used only once and can be circumvented. Free types declaring non-constant elements (e.g. *TOKENTRY* above) cannot be encoded as sets. However, it is possible to assert that a variable is of that type (e.g. `userTokenOK` and `not_userTokenOK`), and that all the elements of a set are of that type by means of the foreach constraint; not needed in this specification. {*log*} cannot express $\mathbb{N}$ nor $\mathbb{Z}$. It can express that a variable belongs to them and that all the elements of a set belong to them (this is used to encode the type of *sizeElement*). For this reason, we cannot encode the type of *authPeriod* and some values given in *InitConfig*. It is doubtful how these values can be implemented in any programming language as they entail to store a Cartesian product where one of the sets is $\mathbb{N}$.

Predicate outside the decision procedure. The universally quantified predicate in schema *ValidEnrol* (not included for brevity) lays outside the decision procedures implemented in {*log*}—and evidence suggests that this is a fundamental problem [11], [26]. The problem is that the property asserted for all elements in *issuerCerts* depends on *issuerCerts* itself. In other words, for each element in *issuerCerts* there must exist another element in it fulfilling a certain property. This creates a sort of recursive definition. However, the predicate *can* be encoded and it *can* be used for running the prototype. Problems could arise if {*log*} is asked to decide the satisfiability of a formula involving it, when *issuerCerts* remains variable. In that case {*log*} might enter an infinite loop. This is apparently not the case for the set of properties considered in the Tokeneer project.

Computationally hard predicate. The predicate $oldElements \cup auditLog' = auditLog \cup newElements$ in schema *AddElementsToLog* is computationally very hard when all the operands remain variable. We were able to circumvent this issue by enclosing its encoding inside a `delay` predicate.

As the empirical data shows (Section 5), these issues do not threaten the chances of using {*log*} as a verification tool for industrial-strength Z specifications.

5 A Verified Prototype

Now that we have the {log} program, we can use {*log*} to prove properties of it. We *automatically* prove two kinds of properties: invariance lemmas and security properties. In any case, recall that {*log*} can prove a theorem by proving that its negation is unsatisfiable. If the theorem is of the form $p \Rightarrow q$, then one should ask {*log*} to check if $p \wedge \neg q$ is unsatisfiable. In doing so consider that: *a)* the negation in $\neg q$ should be encoded as indicated in Sections 3 and 4; and *b)* if $p \wedge \neg q$ happens to be satisfiable, {*log*} will produce a finite representation of all the possible solutions (countermodels or counterexamples). This last feature comes handy to find out what are the possible causes for the formula not to be a theorem.

Section 6 provides quantitative figures about the verification work carried out with {*log*}.

5.1 Invariance Lemmas

As explained in Section 2.1, we move the state invariants included in schema *IDStation* out of it; this includes all the state invariants written in the schemas included in *IDStation*. In doing so, we must: *a)* add pre- or post-conditions to some operations; and *b)* prove that each TIS operation preserves each invariant.

Concerning *a)*, for instance, the encoding of *UnlockDoor* is augmented with the following conditions:

```
unlockDoor(DoorLatchAlarm,Config,
           DoorLatchAlarm_) :-
......
  (LatchUnlockDuration neq 0 &
   CurrentLatch_ = unlocked &
   DoorAlarm_ = silent
   or
   LatchUnlockDuration = 0 &
   CurrentLatch_ = locked &
   (AlarmSilentDuration neq 0 &
    DoorAlarm_ = silent
    or
    AlarmSilentDuration = 0 &
    (CurrentDoor = open &
     DoorAlarm_ = alarming
     or
     CurrentDoor = closed &
     DoorAlarm_ = silent
    )
   )
  ).
```

for *TISUnlockDoor* to preserve the invariant stated in schema *DoorLatchAlarm*:

$$currentLatch = locked \Leftrightarrow currentTime \geq latchTimeout$$
$$doorAlarm = alarming \Leftrightarrow$$
$$(currentDoor = open$$
$$\wedge\ currentLatch = locked$$
$$\wedge\ currentTime \geq alarmTimeout)$$

Note that the encoding of the invariant is closer to an implementation based on conditional statements than the Z invariant and that it takes into account other predicates included in *UnlockDoor* like $latchTimeout' = currentTime + latchUnlockDuration$. This would help in passing from the $\{log\}$ prototype to a definitive implementation.

Concerning *b)*, for each TIS operation we discharge a proof obligation of the form:

$$Invariant \wedge TISOperation \Rightarrow Invariant'$$

encoded as:

```
invariant(IDStation) &
tisOperation(IDStation,RealWorld,
             RealWorld_, IDStation_) &
not_invariant(IDStation_).
```

where `not_invariant` is the $\{log\}$ negation of `invariant`.

Example 8. The state invariant included in *IDStation*:

$$enclaveStatus \notin \{ notEnrolled, waitingEnrol, waitingEndEnrol \} \Rightarrow ownName \neq nil$$

is encoded as:

```
idStationInv07(IDStation) :-
  IDStation =
    [...,KeyStore,...,Internal,...] &
    KeyStore = [_,OwnName] &
    Internal = [_,EnclaveStatus,_] &
  (EnclaveStatus in {notEnrolled,
                     waitingEnrol,
                     waitingEndEnrol}
   or
   OwnName neq {}
  ).
```

while its negation is:

```
not_idStationInv07(IDStation) :-
  IDStation =
    [...,KeyStore,...,Internal,...] &
    KeyStore = [_,OwnName] &
    Internal = [_,EnclaveStatus,_] &
  EnclaveStatus nin {notEnrolled,
                     waitingEnrol,
                     waitingEndEnrol} &
  OwnName = {}.
```

□

5.2 Security Properties

All security properties formalized by the team—i.e., Property 1-4 and 6—are encoded in $\{log\}$. However, Property 2 cannot be proved by $\{log\}$ because it requires a decision procedure for integer intervals—which is, in fact, part of our current work. Properties 3, 4 and 6 are proved as formalized by the team. For instance, the encoding of Property 3 is the following[4]:

```
property3 :-
  idStationInv01(IDStation) &
  (tisEarlyUpdate(IDStation,RealWorld,
                  RealWorld_,IDStation_)
   or
   tisUpdate(IDStation,RealWorld,
             RealWorld_,IDStation_)
  ) &
......
  Latch_ = locked &
  CurrentDoor_ = open &
  CurrentTime_ >= AlarmTimeout &
  Alarm_ neq alarming.
```

5.2.1 Proving Property 1

As we have explained in Section 2.2, the formalization of Property 1 given by the team does not actually capture the intended property. This is acknowledged by the team in the technical documentation. The intended property requires to consider different execution sequences of the TIS operations. Each of these sequences takes the system from the initial state to a state where the property holds. The point made by the team is that the system can follow only those sequences.

Therefore, we make one more step by proving that the system can only execute those state sequences. Once the system arrives at the desired state, we prove that Property 1 holds. Hence, the proof of Property 1 involves (automatically) discharging 16 proof obligations—11 to prove the system can only execute certain state changes; 4 to prove that some properties hold at some of the traversed states; and 1 to prove Property 1. This proof strategy follows the informal proof made by the team [15, Section 3.2.1, pages 8-10].

Before giving details on the encoding of these proofs, it should be noted that in doing them we have to add $enclaveStatus \neq waitingStartAdminOp$ as an hypothesis to many of the auxiliary lemmas—otherwise they do not hold. This is so because, without that hypothesis, *TISShutdownOp* can be executed violating some desired properties. To the best of our knowledge, this extra hypothesis and the problems with *TISShutdownOp* are not mentioned in the documentation. Maybe this is obvious for the team and so they omitted it in the documentation.

Figures 4 and 5 depict the state sequences that lead the system to the state where Property 1 holds. Figure 6 shows the encoding of a lemma stating that if the system ever reaches $status = gotUserToken$ it is because the before state was $status = quiescent$—i.e., the first transition of Figure 4.

4. The ellipsis replace the unification between `IDStation` and the other arguments with tuples to have access to the variables.

That is, the $\{log\}$ code corresponds to the negation of the following formula:

$$\begin{array}{l} \Delta IDStation;\ \Delta RealWorld \mid \\ \quad TISOpThenUpdate \\ \quad \land\ enclaveStatus \neq waitingStartAdminOp \\ \quad \land\ status' = gotUserToken \land status' \neq status \\ \vdash \\ \quad status = quiescent \end{array}$$

where *TISOpThenUpdate* is the composition between the disjunction of all the TIS operations and the *TISUpdate* operation [15, page 5].

We prove one such lemma for each transition shown in Figures 4 and 5. In this way, the $\{log\}$ prototype is guaranteed to follow only those state sequences when it comes to the validity of Property 1.

Next, we prove that when the system reaches some of the states depicted in Figures 4 and 5 some properties hold. For example, when passing from *gotUserToken* to *waitingEntry*, the user token is checked for validity and the presence in the token of an authorization certificate is checked as well. Figure 7 shows the encoding of this lemma. As can be seen, we need `pfun(IssuerKey_)` as an hypothesis. In Z terms, this means that *issuerKey*′ is a partial function, which is the type given to the variable in the specification. This hypothesis is proved to be a state invariant and so it can be assumed without loss of generality.

Finally, Property 1 itself is proved by assuming the properties shown to be valid in the intermediate states.

Having discharged all the proof obligations concerning state invariants and those concerning security properties, the $\{log\}$ prototype is correct w.r.t. them. Therefore, it can be used as a correct implementation of the specification, at least on what concerns to the proven properties.

5.3 The Prototype as a Verification Tool

The prototype can be used as an analysis tool taking advantage of its correctness. The most natural activity is to use it as an implementation to evaluate how the system behaves in specific functional scenarios.

Example 9. We may want to analyze different scenarios when the enclave's door is unlocked by asking $\{log\}$ to execute the following:

```
DLA = [5,locked,_,_,_,_] &
C = [10,4,_,_,_,_,_,_] &
doorLatchAlarmInv(DLA) & configInv(C,100) &
unlockDoor(DLA,C,DLA_).
```

where part of the first solution is:

```
DLA_ = [5,locked,unlocked,silent,9,19]
```

Observe that we give the before state (e.g. `DLA`), check whether it satisfies the invariants, and then run `unlockDoor` waiting for the next state (`DLA_`). Besides, note that the before state is only partially given (cf. the underscores in `DLA`) which allows to analyze several scenarios in one run. □

Being able to execute these scenarios in the presence of the customer can help to "check that this vital step [Z specification] has been achieved correctly" without needing a person "with full knowledge of the problem domain and a fluent reading knowledge of Z" [6]— the latter being not easy to find.

A more elaborated activity would be to use the $\{log\}$ prototype as part of a model-based testing method, e.g. [27], [28], [29]. Let's call $P_{\{log\}}$ and P_{SPARK} the $\{log\}$ prototype and the SPARK implementation corresponding to the Z specification, respectively. Given a test case generated from the specification, run it on both $P_{\{log\}}$ and P_{SPARK} and compare the outputs generated by each of them to decide whether or not the test case has uncovered an error in P_{SPARK}. The output generated by $P_{\{log\}}$ should be deemed correct as the prototype verifies all the stated properties. Hence, if the output generated by P_{SPARK} does not coincide it should be concluded that there is an error in it. Furthermore, $\{log\}$ itself can be used as a test case generator in the context of Z specifications [23].

Along the same line, $P_{\{log\}}$ can be used as a runtime or reference monitor of P_{SPARK}. Reference monitors have been proposed as a way to control the execution of secure systems, e.g. [30], [31]. In this scenario, $P_{\{log\}}$ and P_{SPARK} execute in parallel in such a way that every input sent to P_{SPARK} is also sent to $P_{\{log\}}$. If the outputs differ an alarm is fired. Clearly, the control is *ex post* (due to the, possibly sensible, different execution speeds) but in some circumstances this is much better than nothing. For example, in case a malfunction in P_{SPARK} allows an unauthorized access to the enclave, the late alarm fired by $P_{\{log\}}$ would allow for security personnel to correct the situation before is too late.

The intention behind these proposals and the mere fact of using $\{log\}$ is to show that $\{log\}$ can be useful, and not that it should replace existing tools or techniques used in development processes such as Altran Praxis' Correct by Construction. $\{log\}$ can fill some gaps when it comes to prototyping, automated proof and counterexample generation in the context of set-based specifications.

6 Quantitative Report

Besides proving all the lemmas described in Sections 5.1 and 5.2 we prove that:

1) Every disjunct of all the TIS operations is satisfiable. This verification is important because, as operations act in the antecedent of lemmas, if they are unsatisfiable then the lemma holds trivially. See for example Property 1 in Section 2.2.
2) The initial state satisfies all the state invariants. This is a standard verification step when state invariants are involved.
3) The conjunction of the negation of a schema and the schema itself is unsatisfiable, while those schemas are independently satisfiable. This applies only to those schemas whose negation appears either in the specification or as part of a lemma. These proofs are done to gain confidence in that each negation has been correctly written.

Table 1 summarizes key quantitative measures about the effort involved during this work. The first measured magnitude is effort (man-hours). It is difficult and can be misleading to estimate the effort needed to conduct the

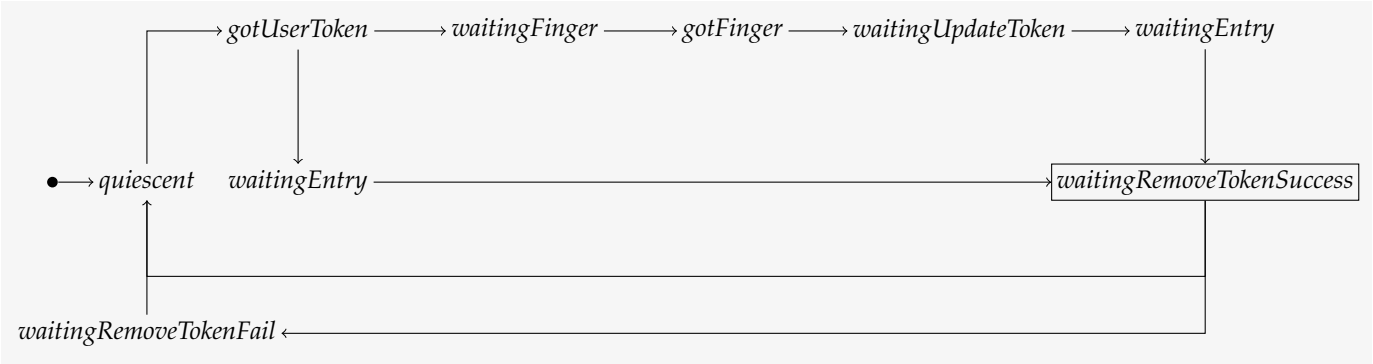

Fig. 4. Sequences of state changes leading to $status = waitingRemoveTokenSuccess$

Fig. 5. Sequences of state changes leading to $enclaveStatus = waitingStartAdminOp, rolePresent \neq \emptyset$

```
property1_01 :-
  IDStation = [_,_,_,DoorLatchAlarm,_,_,_,_,_,_,_,Internal,_,_] &
    Internal = [Status,EnclaveStatus,_] &
  EnclaveStatus = enclaveQuiescent &
  Status neq quiescent &
  tisOpThenUpdate(M,IDStation,RealWorld,RealWorld_,IDStation_) &
  IDStation_ = [_,_,_,_,_,_,_,_,_,_,_,Internal_,_,_] &
    Internal_ = [Status_,_,_] &
  Status neq Status_ & Status_ = gotUserToken.
```

Fig. 6. Encoding of a state change lemma

```
property1_09 :-
  IDStation = [_,_,_,_,_,_,_,_,_,_,_,Internal,_,_] &
    Internal = [Status,EnclaveStatus,_] &
  Status = gotUserToken & EnclaveStatus = enclaveQuiescent &
  tisOpThenUpdate(M,IDStation,RealWorld,RealWorld_,IDStation_) &
  IDStation_ = [UserToken_,_,_,_,_,_,_,_,KeyStore_,_,_,Internal_,_,_] &
    UserToken_ = [CurrentUserToken_,_] &
    Internal_ = [Status_,_,_] &
  Status_ = waitingEntry &
  (not_tokenWithValidAuth(CurrentUserToken_)
   or
   KeyStore_ = [IssuerKey_,_] &
   pfun(IssuerKey_) &
   not_userTokenWithOKAuthCertNoCurrencyCheck(KeyStore_,UserToken_,_)
  ).
```

Fig. 7. Encoding of a lemma stating an intermediate property

work carried out with {*log*} because it was done on an academic environment rather than a corporate one. Nevertheless, we think that an approximate number could help others. Our effort estimate for encoding the Z specification and discharging all the 523 proof obligations is between 40 and 120 man-hour. It depends on the experience of the engineer with Z, {*log*}, proof strategies and automated proof. Note that the effort reported by the team to write the Z specification and conduct informal proofs is 233.5 man-hour [7, Appendix B, page 71, item 3000]. A more elaborated tool environment could further reduce our estimate.

The next measured magnitude is size. As can be seen, the {*log*} prototype is around 30% larger than the Z specification measured in lines of code (LOC), but it is more than the double in terms of kilobytes. This is consistent with the fact that any implementation is expected to be more verbose than the specification. This verbosity can be observed in Figures 1 and 2. The {*log*} code for proof obligations (lemmas) is considerably larger than the prototype, but it means less than 10 LOC per lemma and it contains pretty-printing predicates to simplify proof execution. Much of the proof obligations code can be automatically generated—we did so with a few simple `bash` scripts but in an industrial environment it can be done much better.

The last aspect we measure concerns the number of proof obligations (523) and the time {*log*} spends in discharging them (851 seconds). As can be seen, most of the proof obligations corresponds to invariance lemmas as well as the computing time to discharge them (548 seconds). Most of the time is spent in proving the invariance of three TIS operations (*TISEnrolOp*, *TISValidateUserToken* and *TISPoll*). In average, each proof obligation runs in 1.6 seconds. There are at least two possible ways to reduce the computing time: parallelization and what we call *proof programming*. In effect, every lemma can be discharged independently of the others and so each of them can be run in a different thread. Prolog provides high-level parallelization predicates (e.g. `concurrent/4`) that can be easily used to considerably reduce the computing time.

Proof programming refers to the application of some {*log*} control predicates (e.g. `delay`) to impose some order in the execution of constraints. As we have shown in Section 5, enclosing a constraint in a `delay` predicate can speed up the execution of the prototype. However, if the delayed constraint is important to discharge a particular proof obligation the proof will take longer compared to a goal where that constraint is not delayed. Hence, the best strategy is to delay a constraint for a proof while not to delay it for another. This requires some *proof programming*. In general proof programming has to be applied only to some constraints. It is an area that deserves to be further explored.

The numbers reported in Table 1 provide evidence about the usefulness of {*log*} as a verification tool for Z specifications.

6.1 Platform where the verification was executed

The verification of the {*log*} prototype of the TIS specification was performed on a Latitude E7470 (06DC) with a 4 core Intel(R) Core™ i7-6600U CPU at 2.60GHz with 8 Gb of main memory, running Linux Ubuntu 18.04.4 (LTS) 64-bit with kernel 4.15.0-106-generic. {*log*} 4.9.6-21c over SWI-Prolog (multi-threaded, 64 bits, version 7.6.4) was used during the experiments.

The execution time of each collection of proof obligations is given by `T` in the following {*log*} formula:

```
prolog_call(get_time(Ti)) &
<collection proof obligations>
prolog_call(get_time(Te)) &
prolog_call(T is Te - Ti).
```

where `prolog_call` is a {*log*} facility to access the Prolog interpreter.

The {*log*} code used in this work can be downloaded from http://people.dmi.unipr.it/gianfranco.rossi/SETLOG/APPLICATIONS/tokeneer.zip along with instructions to set up the environment.

7 Other works about the Tokeneer Project

Besides members of the team in charge of the Tokeneer project, other researchers have used it as a case study, a benchmark or just as an industrial-scale problem. We will comment on the most relevant ones w.r.t. our work.

Rivera et al. [32] undertake the Tokeneer project in Event-B. They use the Rodin toolset for discharging proof obligations and the EventB2Java code generator to create a Java program from the Event-B model. The TIS model consists of an abstract machine and 6 refinements. The full development resulted in 334 proof obligations, of which more than 90% are discharged automatically using Rodin. These proof obligations should coincide with a subset of those proved with {*log*}. Rivera and his colleagues prove Properties 1, 2 and 3 with Rodin, although it is not clear whether or not they are automatically discharged. As can be seen, {*log*} is able to automatically discharge 100% of the proof obligations. Besides, {*log*} produces a prototype while the Event-B approach needs to apply EventB2Java. Likely, the resulting Java program will be more efficient than the {*log*} prototype, but the former is *only* a program while the latter is a program *and* a formula. Although it is not exactly the same, there is a sort of Java version of {*log*} called JSetL [33], [34].

Answer Set Programming, which is close to constraint programming, has been used to generate counterexamples for false and unprovable verification conditions (VC) of the Tokeneer project [35]. These VC correspond to those generated during the verification of the SPARK implementation. As we have shown, {*log*} returns a finite representation of all the solutions of any satisfiable formula. These solutions are counterexamples when the intention is to discharge a VC—that is, the VC is negated, submitted to {*log*} and it returns a solution (counterexample) meaning that the VC is not a theorem. Along the same lines, in a technical report, Jackson and Passmore apply an SMT solving-based tool to prove SPARK VC of the Tokeneer project [36]. The tool calls CVC 3, Yices, Z3 and Simplify. Roughly, the tool proves more than 90% of the VC. Tokeneer has also been used as a case study for the formal verification framework Echo [37]. Echo uses PVS as a theorem prover and SPARK as

TABLE 1
Summary of Quantitative Report

Effort		
From 40 to 120 man-hour (prototype + proofs)		
Size		
	KLOC	KByte
Z (LaTeX)	2	52
Prototype ({*log*})	2.6	108
Lemmas ({*log*})	5.1	276
Proof Obligations (Lemmas)		
Collection	Number	Time (s)
Negations	93	4
Invariance	325	533
TISEnrolOp		107
TISValidateUserToken		114
TISPoll		281
Satisfiability of initial state	12	0.1
Satisfiability of operations	74	0.9
Security properties	19	313
Totals	523	851

programming language. According to the paper "in 90% of the cases, the PVS theorem prover could not prove the implication lemmas completely automatically".

Although, the VC are at the SPARK level and the proof obligations discharged by {*log*} and Rodin (cf. Rivera et al. above) are at the specification level, there should be a clear relation between them as the SPARK program should implement the specification. Hence, {*log*} looks promising as a VC verifier.

Abdelhalim et al. [38] apply CSP to formalize fUML activity diagrams and FDR as a model checker to the Tokeneer specification. Specifically the authors found several deadlock scenarios in the form of counterexamples generated by FDR.

The work authored by Moy and Wallenburg [39] is interesting because they find problems in Tokeneer, although it was formally verified. Moy and Wallenburg's goal is to find out why these problems were not found when the system was verified and to propose verification activities that could have detected these problems. Specifically, the authors propose to complement formal verification with static analysis and code reviews. {*log*} might be considered as part of the toolbox proposed by Moy and Wallenburg as it is at the intersection of several programming and verification paradigms. For instance, it can be used to perform proofs and to run functional scenarios.

Woodcock at al. [29] apply an assertion-guided model-based robustness testing method to the Tokeneer project. Robustness testing checks that a system can handle unexpected user input or software failures. They use a model of the system for code generation (the Z original specification) and a separate model for test case generation (an Alloy model); these models are independently produced from the requirements. The test case specifications are fed into the Alloy Analyzer, and test cases are automatically generated as counterexamples. This allowed the authors to detect nine anomalous behaviors. {*log*} can be used in place of Alloy and its analyzer. In fact the Alloy Analyzer does not implement a decision procedure for sets and binary relations, as the one provided by {*log*}. Then, the Alloy Analyzer might fail in finding a counterexample while {*log*} might not.

8 Conclusions

We have encoded the Z specification of the Tokeneer project in {*log*}. This encoding can be used as a functional prototype. Then, we used {*log*} to automatically proved hundreds of proof obligations over the prototype itself. That is, we took advantage of the formula-program duality of {*log*} code to produce a verified prototype w.r.t. the proven properties. In this way, {*log*} is used as a programming language and a verification engine using the same and only representation of the system. The case study provides evidence that {*log*} can be helpful in analyzing set-based formal specifications like those written in the Z formal notation.

The most interesting future work is to apply {*log*} to discharge the verification conditions generated during the verification of the SPARK implementation.

References

[1] X. Leroy, "Formal verification of a realistic compiler," *Commun. ACM*, vol. 52, no. 7, pp. 107–115, 2009. [Online]. Available: http://doi.acm.org/10.1145/1538788.1538814

[2] T. C. Murray, D. Matichuk, M. Brassil, P. Gammie, T. Bourke, S. Seefried, C. Lewis, X. Gao, and G. Klein, "sel4: From general purpose to a proof of information flow enforcement," in *2013 IEEE Symposium on Security and Privacy, SP 2013, Berkeley, CA, USA, May 19-22, 2013*. IEEE Computer Society, 2013, pp. 415–429. [Online]. Available: https://doi.org/10.1109/SP.2013.35

[3] S. King, J. Hammond, R. Chapman, and A. Pryor, "Is proof more cost-effective than testing?" *IEEE Trans. Software Eng.*, vol. 26, no. 8, pp. 675–686, 2000. [Online]. Available: https://doi.org/10.1109/32.879807

[4] A. Hall and R. Chapman, "Correctness by construction: Developing a commercial secure system," *IEEE Software*, vol. 19, no. 1, pp. 18–25, 2002. [Online]. Available: https://doi.org/10.1109/52.976937

[5] Common Criteria Recognition Arrangement, "Common criteria for information technology security evaluation, part 1: Introduction and general model, version 3.1. release 5," Tech. Rep., 2017. [Online]. Available: https://www.commoncriteriaportal.org/files/ccfiles/CCPART1V3.1R5.pdf

[6] J. Barnes, R. Chapman, R. Johnson, J. Widmaier, D. Cooper, and B. Everett, "Engineering the Tokeneer enclave protection software," in *Proceedings of the IEEE International Symposium on Secure Software Engineering*. IEEE, 2006.

[7] D. Cooper and J. Barnes, "Tokeneer ID Station EAL5 demonstrator: Summary report," Altran Praxis, Tech. Rep., 2008. [Online]. Available: https://www.adacore.com/uploads/downloads/Tokeneer_Report.pdf

[8] J. M. Spivey, *The Z notation: a reference manual*. Hertfordshire, UK, UK: Prentice Hall International (UK) Ltd., 1992.

[9] B. Potter, J. Sinclair, and D. Till, *An introduction to formal specification and Z*. Prentice Hall PTR Upper Saddle River, NJ, USA, 1996.

[10] A. Dovier, C. Piazza, E. Pontelli, and G. Rossi, "Sets and constraint logic programming," *ACM Trans. Program. Lang. Syst.*, vol. 22, no. 5, pp. 861–931, 2000.

[11] M. Cristiá and G. Rossi, "Solving quantifier-free first-order constraints over finite sets and binary relations," *J. Autom. Reasoning*, vol. 64, no. 2, pp. 295–330, 2020. [Online]. Available: https://doi.org/10.1007/s10817-019-09520-4

[12] G. Rossi, "{*log*}," http://people.dmi.unipr.it/gianfranco.rossi/setlog.Home.html, 2008. [Online]. Available: http://people.dmi.unipr.it/gianfranco.rossi/setlog.Home.html

[13] A. Praxis". (2008) Tokneer software and project documents. [Online]. Available: http://www.adacore.com/uploads/downloads/tokeneer.zip

[14] J. Barnes, "Tokeneer ID Station: Formal specification," Altran Praxis, Tech. Rep., 2008, find it as /tokeneer/docs/41_2_Formal_Functional_Specification/41_2.pdf inside tokeneer.zip. [Online]. Available: http://www.adacore.com/uploads/downloads/tokeneer.zip

[15] D. Cooper, "Tokeneer ID Station: Security properties," Altran Praxis, Tech. Rep., 2008, find it as /tokeneer/docs/40_4_Security_Properties/40_4.pdf inside tokeneer.zip. [Online]. Available: http://www.adacore.com/uploads/downloads/tokeneer.zip

[16] A. Evans, "An improved recipe for specifying reactive systems in Z," in *ZUM '97: The Z Formal Specification Notation, 10th International Conference of Z Users, Reading, UK, April 3-4, 1997, Proceedings*, ser. Lecture Notes in Computer Science, J. P. Bowen, M. G. Hinchey, and D. Till, Eds., vol. 1212. Springer, 1997, pp. 275–294. [Online]. Available: https://doi.org/10.1007/BFb0027293

[17] L. Lamport, "TLZ," in *Z User Workshop, Cambridge, UK, 29-30 June 1994, Proceedings*, ser. Workshops in Computing, J. P. Bowen and J. A. Hall, Eds. Springer/BCS, 1994, pp. 267–268. [Online]. Available: https://doi.org/10.1007/978-1-4471-3452-7_15

[18] M. Cristiá and G. Rossi, "A set solver for finite set relation algebra," in *Relational and Algebraic Methods in Computer Science - 17th International Conference, RAMiCS 2018, Groningen, The Netherlands, October 29 - November 1, 2018, Proceedings*, ser. Lecture Notes in Computer Science, J. Desharnais, W. Guttmann, and S. Joosten, Eds., vol. 11194. Springer, 2018, pp. 333–349. [Online]. Available: https://doi.org/10.1007/978-3-030-02149-8_20

[19] ——, "A decision procedure for restricted intensional sets," in *Automated Deduction - CADE 26 - 26th International Conference on Automated Deduction, Gothenburg, Sweden, August 6-11, 2017, Proceedings*, ser. Lecture Notes in Computer Science, L. de Moura, Ed., vol. 10395. Springer, 2017, pp. 185–201. [Online]. Available: https://doi.org/10.1007/978-3-319-63046-5_12

[20] ——, "Automated reasoning with restricted intensional sets," *CoRR*, vol. abs/1910.09118, 2019. [Online]. Available: http://arxiv.org/abs/1910.09118

[21] C. Holzbaur, F. Menezes, and P. Barahona, "Defeasibility in CLP(Q) through generalized slack variables," in *CP*, ser. Lecture Notes in Computer Science, E. C. Freuder, Ed., vol. 1118. Springer, 1996, pp. 209–223.

[22] A. Dovier, E. Pontelli, and G. Rossi, "Set unification," *Theory Pract. Log. Program.*, vol. 6, no. 6, pp. 645–701, 2006.

[23] M. Cristiá, G. Rossi, and C. S. Frydman, "{log} as a test case generator for the Test Template Framework," in *SEFM*, ser. Lecture Notes in Computer Science, R. M. Hierons, M. G. Merayo, and M. Bravetti, Eds., vol. 8137. Springer, 2013, pp. 229–243.

[24] M. Cristiá and G. Rossi, "Automated proof of Bell–LaPadula security properties," *J. Autom. Reasoning*, vol. n/a, no. n/a, jul 2020. [Online]. Available: https://link.springer.com/article/10.1007/s10817-020-09577-6

[25] J. T. Schwartz, R. B. K. Dewar, E. Dubinsky, and E. Schonberg, *Programming with Sets - An Introduction to SETL*, ser. Texts and Monographs in Computer Science. Springer, 1986. [Online]. Available: http://dx.doi.org/10.1007/978-1-4613-9575-1

[26] H. Andréka, S. R. Givant, and I. Németi, *Decision problems for equational theories of relation algebras*. American Mathematical Soc., 1997, vol. 604.

[27] P. Stocks and D. Carrington, "A Framework for Specification-Based Testing," *IEEE Transactions on Software Engineering*, vol. 22, no. 11, pp. 777–793, Nov. 1996.

[28] M. Cristiá, P. Albertengo, C. S. Frydman, B. Plüss, and P. Rodríguez Monetti, "Tool support for the Test Template Framework," *Softw. Test., Verif. Reliab.*, vol. 24, no. 1, pp. 3–37, 2014.

[29] J. Woodcock, E. G. Aydal, and R. Chapman, "The tokeneer experiments," in *Reflections on the Work of C. A. R. Hoare*, A. W. Roscoe, C. B. Jones, and K. R. Wood, Eds. Springer, 2010, pp. 405–430. [Online]. Available: https://doi.org/10.1007/978-1-84882-912-1_17

[30] G. Betarte, J. D. Campo, C. Luna, and A. Romano, "Formal analysis of Android's permission-based security model,," *Sci. Ann. Comput. Sci.*, vol. 26, no. 1, pp. 27–68, 2016. [Online]. Available: https://doi.org/10.7561/SACS.2016.1.27

[31] C. Luna, G. Betarte, J. D. Campo, C. Sanz, M. Cristiá, and F. Gorostiaga, "A formal approach for the verification of the permission-based security model of Android," *CLEI Electron. J.*, vol. 21, no. 2, 2018. [Online]. Available: https://doi.org/10.19153/cleiej.21.2.3

[32] V. Rivera, S. Bhattacharya, and N. Cataño, "Undertaking the tokeneer challenge in event-b," in *Proceedings of the 4th FME Workshop on Formal Methods in Software Engineering, FormaliSE@ICSE 2016, Austin, Texas, USA, May 15, 2016*. ACM, 2016, pp. 8–14. [Online]. Available: https://doi.org/10.1145/2897667.2897671

[33] M. Cristiá and G. Rossi, "Programming in Java with restricted intensional sets," in *Proceedings of the 3rd International Workshop on Sets and Tools co-located with the 6th International ABZ Conference, SETS@ABZ 2018, Southamptom, UK, June 5, 2018.*, ser. CEUR Workshop Proceedings, M. Cristiá, D. Delahaye, and C. Dubois, Eds., vol. 2199. CEUR-WS.org, 2018, pp. 17–31. [Online]. Available: http://ceur-ws.org/Vol-2199/paper2.pdf

[34] G. Rossi and F. Bergenti, "Nondeterministic programming in Java with JSetL," *Fundam. Inform.*, vol. 140, no. 3-4, pp. 393–412, 2015. [Online]. Available: https://doi.org/10.3233/FI-2015-1260

[35] F. Schanda and M. Brain, "Using answer set programming in the development of verified software," in *Technical Communications of the 28th International Conference on Logic Programming, ICLP 2012, September 4-8, 2012, Budapest, Hungary*, ser. LIPIcs, A. Dovier and V. S. Costa, Eds., vol. 17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2012, pp. 72–85. [Online]. Available: https://doi.org/10.4230/LIPIcs.ICLP.2012.72

[36] P. B. Jackson and G. O. Passmore, "Proving SPARK verification conditions with SMT solvers," University of Edinburgh, Tech. Rep., 2009. [Online]. Available: http://homepages.inf.ed.ac.uk/pbj/papers/vct-dec09-draft.pdf

[37] X. Yin and J. C. Knight, "Formal verification of large software systems," in *Second NASA Formal Methods Symposium - NFM 2010, Washington D.C., USA, April 13-15, 2010. Proceedings*, ser. NASA Conference Proceedings, C. A. Muñoz, Ed., vol. NASA/CP-2010-216215, 2010, pp. 192–201.

[38] I. Abdelhalim, J. Sharp, S. A. Schneider, and H. Treharne, "Formal verification of tokeneer behaviours modelled in fuml using CSP," in *Formal Methods and Software Engineering - 12th International Conference on Formal Engineering Methods, ICFEM 2010, Shanghai, China, November 17-19, 2010. Proceedings*, ser. Lecture Notes in Computer Science, J. S. Dong and H. Zhu, Eds., vol. 6447. Springer, 2010, pp. 371–387. [Online]. Available: https://doi.org/10.1007/978-3-642-16901-4_25

[39] Y. Moy and A. Wallenburg, "Tokeneer: Beyond formal program

verification," *Embedded Real Time Software and Systems*, vol. 24, 2010.

Maximiliano Cristiá to be completed.

Gianfranco Rossi to be completed.