



## International Journal of Web Information Systems

Formal and semi-formal verification of a web voting system

Maximiliano Cristia Claudia Frydman

### Article information:

To cite this document:

Maximiliano Cristia Claudia Frydman , (2015), "Formal and semi-formal verification of a web voting system", International Journal of Web Information Systems, Vol. 11 Iss 2 pp. 183 - 204

Permanent link to this document:

<http://dx.doi.org/10.1108/IJWIS-11-2014-0042>

Downloaded on: 09 November 2016, At: 02:00 (PT)

References: this document contains references to 51 other documents.

To copy this document: [permissions@emeraldinsight.com](mailto:permissions@emeraldinsight.com)

The fulltext of this document has been downloaded 99 times since 2015\*

### Users who downloaded this article also downloaded:

(2015), "Development of mobile applications from existing Web-based enterprise systems", International Journal of Web Information Systems, Vol. 11 Iss 2 pp. 162-182 <http://dx.doi.org/10.1108/IJWIS-11-2014-0041>

(2015), "Web spam detection using trust and distrust-based ant colony optimization learning", International Journal of Web Information Systems, Vol. 11 Iss 2 pp. 142-161 <http://dx.doi.org/10.1108/IJWIS-12-2014-0047>

Access to this document was granted through an Emerald subscription provided by emerald-srm:563821 []

### For Authors

If you would like to write for this, or any other Emerald publication, then please use our Emerald for Authors service information about how to choose which publication to write for and submission guidelines are available for all. Please visit [www.emeraldinsight.com/authors](http://www.emeraldinsight.com/authors) for more information.

### About Emerald [www.emeraldinsight.com](http://www.emeraldinsight.com)

Emerald is a global publisher linking research and practice to the benefit of society. The company manages a portfolio of more than 290 journals and over 2,350 books and book series volumes, as well as providing an extensive range of online products and additional customer resources and services.

Emerald is both COUNTER 4 and TRANSFER compliant. The organization is a partner of the Committee on Publication Ethics (COPE) and also works with Portico and the LOCKSS initiative for digital archive preservation.

\*Related content and download information correct at time of download.

# Formal and semi-formal verification of a web voting system

Web voting  
system

183

Maximiliano Cristia  
*CIFASIS and UNR, Rosario, Argentina, and*  
Claudia Frydman  
*LSIS-CIFASIS, Marseille, France*

Received 25 November 2014  
Revised 9 December 2014  
15 December 2014  
Accepted 15 December 2014

## Abstract

**Purpose** – This paper aims to present the verification process conducted to assess the functional correctness of the voting system. Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET) is the most important research institution in Argentina. It depends directly from Argentina's President but its internal authorities are elected by around 8,000 research across the country. During 2011, the CONICET developed a Web voting system to replace the traditional mail-based process. In 2012 and 2014, CONICET conducted two Web election with no complaints from candidates and voters. Before moving the system into production, CONICET asked the authors to conduct a functional and security assessment of it.

**Design/methodology/approach** – This process is the result of integrating formal, semi-formal and informal verification activities from formal proof to code inspection and model-based testing.

**Findings** – Given the resources and time available, a reasonable level of confidence on the correctness of the application could be transmitted to senior management.

**Research limitations/implications** – A formal specification of the requirements must be developed.

**Originality/value** – Formal methods and semi-formal activities are seldom applied to Web applications.

**Keywords** Code review, Formal methods, Model-based testing, Z specifications

**Paper type** Research paper

## 1. Introduction

The Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET[1]) is the most important research institution in Argentina. Its president is nominated by Argentina's President, but some of its internal authorities, including some of the members of its Board of Directors, are elected by some of the more than 8,000 CONICET's researchers. Traditionally, the election process was carried out by mail (i.e. those allowed to vote had to send their votes by mail and, thus, the recount process was done manually). In 2011, the Board of Directors commissioned the development of an electronic voting system. Eventually, it was decided that CONICET's Systems Department would develop a Web-based system[2]. According to CONICET's authorities, its Web-based voting system is the first to be developed and used in Argentina.

As part of this important change, the Board of Directors issued a set of regulations as the legal support for the new form of the election process. These regulations set all the rules that voters and candidates have to follow to vote or be



voted. They take the form of a document written in natural language (Spanish) organized as a list of clauses.

The Board of Directors also required the Systems Department to subject the system to an external evaluation to determine whether it fulfilled functional, security and availability quality attributes, before moving it into production. Although the program is not large, it uses some complex technologies and a failure may have a high political impact. The System Department contacted the authors to perform the evaluation. One of the authors conducted the functional evaluation, and two security researchers performed the security and availability evaluation. This paper is an experience report on the verification activities conducted during the functional evaluation of the voting system. Note that the aim of the task was to assess the quality of an existing system and not to develop it from scratch.

Given the potential risk of a functional error, we decided to conduct an evaluation as formal as possible within the available time frame (one month). So, the first step was to write a formal specification of the set of regulations governing the election process issued by the Board of Directors (Section 2). Second, we formally proved that the specification verifies some properties as a way to have some confidence that the model is correct (Section 3). The third step consisted in inspecting the source code to find each pre- and post-condition of the model and annotating the source code with links to the corresponding predicates of the formal specification (Section 4). In the last step, we applied Fastest, a model-based testing (MBT) tool, to generate test cases (Section 5). In this way, the formal specification written in the first step was used in all the latter activities. As can be seen, the verification process combined formal, semi-formal and informal techniques, making it possible to finish the task within the schedule. In the rest of the paper, we further discuss these and other issues (Section 6) and survey some related papers (Section 7); our conclusions are in Section 8.

It is very important to remark that this project presents two important differences with respect to other verification efforts. First, in most other reports on the production of verified software, the development team can write a formal specification of the system and then develop the implementation they like the most, using many different formal techniques – formal refinement, program annotation combined with static analysis, etc. In this project, however, the implementation was already developed when its verification was requested – i.e. verification was an afterthought, as is always the case in the vast majority of industrial projects. Therefore, we were not allowed to generate the implementation we liked – mainly because of the available time – we had to inform about the correctness of the available implementation. Second, the implementation was developed by average programmers, and not, for instance, by people holding a PhD in software engineering, formal methods or formal verification. This is the case for almost all the software produced in the world.

On June 15, 2012, the first Web-based election of two members of the Board of Directors was carried out without any noticeable disruption or failure of the new election system. No voter and candidate complained about the election process or its outcome. The same happen after the 2014 election.

This paper is a reduced version of a technical report including the full  $Z$  specification and the  $Z/EVES$  proof scripts. It can be downloaded from [www.dropbox.com/s/c6bawikdrd608c9/votingSystem.tar.gz](http://www.dropbox.com/s/c6bawikdrd608c9/votingSystem.tar.gz). This paper is also an extended version of a

---

conference paper (Cristiá and Frydman, 2014a). It assumes that the reader has basic knowledge of set theory and predicate logic.

## 2. From requirements to formal specification

In this project, the functional requirements are the regulations set by CONICET for the election process. This document is essentially a legal document written in Spanish stipulating all the conditions for an election. It is divided in sections (for example, “On the Voters”, “On the Candidates to Be Elected”, etc.), and each section is organized as a list of articles or clauses. We derived a formal specification for the system from this document plus just a few questions to administrative personnel of CONICET (i.e. no intermediate representation of the requirements was developed).

We wrote a formal specification as the fundamental cornerstone of the verification process. We knew in advance that we could use it for many different verification activities. The Z notation (ISO, 2002; Spivey, 1992) was the chosen language because:

- we are fluent in it and some of its tools;
- Z is very good for this kind of requirements; and
- our research and tools on MBT (Cristiá *et al.*, 2010; Cristiá *et al.*, 2014; Cristiá *et al.*, 2011b) would be of great help (and, conversely, this would become a case study for our investigations).

Although the election process of CONICET shares many rules with that of any democratic nation, it has some peculiarities that distinguish it from other election systems. Just to mention some of the differences, consider the following:

- CONICET divides the country in geographical regions (GR).
- CONICET divides its researchers in so-called knowledge areas (KAs) (for instance, exact and natural sciences, social sciences and humanities, etc.).
- Each voter belongs to one and only one GR and KA.
- Each candidate nominates herself or himself for one and only one GR and KA.
- Each voter can vote up to three candidates of his KA but of different GR.
- Retired researchers may be allowed to vote.
- Not all active or retired researchers can vote or be candidates, there are many conditions they must verify for either one.
- Before a voter can be designated as a firm candidate, he or she has to be endorsed by at least 20 researchers allowed to vote.

The Z specification resulting from the requirements document is a rather standard Z model[3]. We added to the Z model a form of Jackson’s (1995) designations as a means of linking the formal model with the requirements. For instance, before introducing the set representing CONICET’s researchers, we provide the following semi-formal description (designation):

$x$  is an active or retired researcher of CONICET  $\approx x \in RSCH$   
so, then, we introduce the following  $Z$  paragraph:  
[ $RSCH$ ]

As can be seen, the left hand side of a designation is a natural language sentence over the vocabulary of the requirements, while the right hand side is a formal term of the model. Because this is not a technical document, we will not give all the designations corresponding to the portions of the specification included here; we prefer to be more informal in this regard.

### 2.1 A brief introduction to $Z$

$Z$  specifications take the form of state machines. That is, there is a state space and a set of state transitions defined over the state space. The state space is given by declaring a set of state variables, each ranging over a particular set of values. Then, the first step in a  $Z$  specification is to declare the sets or types of values for the state variables. These sets, as the rest of the elements of the specification, are taken from the requirements. Hence, we have defined the following basic sets:

- the set of knowledge areas, named  $KA$ ;
- the set of CONICET's researchers, named  $RSCH$ ; and
- the set of geographical regions in which the country is divided to organize the election, named  $GR$ .

Now, we can give the main components of the state of the system. We factored out the state of the system into two schemata. *Endorsements* record preliminary candidates and their endorsements, while *Votes* record firm candidates and voting data.

#### *Endorsements*

$$\begin{array}{l} \mathit{cand} : RSCH \rightarrow KA \\ \mathit{endors} : RSCH \rightarrow \mathbb{F} RSCH \end{array}$$

#### *Votes*

$$\begin{array}{l} \mathit{firmCand} : \mathbb{P} RSCH \\ \mathit{voted} : \mathbb{F} RSCH \\ \mathit{votes} : \text{seq}(\mathbb{F} RSCH) \end{array}$$

Think of these boxes, called schemas, as sets of elements. In  $Z$ :  $RSCH \rightarrow KA$  denotes the set of all partial functions from  $RSCH$  to  $KA$ ;  $\mathbb{F}RSCH$  is the set of all the finite sets whose elements belong to  $RSCH$ ;  $\mathbb{P}RSCH$  is the set of all the sets (finite or not) whose elements belong to  $RSCH$ ; and  $\text{seq}(\mathbb{F} RSCH)$  is the set of finite sequences (lists) whose elements belong to  $\mathbb{F}RSCH$ . In  $Z$ , partial functions are sets of ordered pairs. So, for instance, a partial function can be expressed as follows:  $\{a \mapsto 1, b \mapsto 2\}$ , where  $a \mapsto 1$  is a synonym for the more standard notation  $(a, 1)$ .

The meaning of state variables is as follows:

- if  $x \mapsto a \in \mathit{cand}$ , then  $x$  is a candidate willing to be endorsed for  $KA$   $a$ ;
- if  $x \in \mathit{endors}(c)$ , then voter  $x$  endorsed candidate  $c$ ;

- if  $x \in firmCand$ , then  $x$  is a firm candidate (i.e. he or she has been endorsed according to CONICET rules);
- if  $x \in voted$ , then  $x$  has already issued his or her vote; and
- $votes$  is the list of votes issued so far.

We used  $\mathbb{P}$  instead of  $\mathbb{F}$  for a variable whenever it was not necessary, a cardinality condition (for example, there is no cardinality condition over  $firmCand$ ). Note the particular type for  $votes$ : is not just a list of researchers but a list of finite sets of researchers, as voters can vote for up to three candidates. It would have been better to declare  $votes$  as bag (FRSCH) but Fastest (the MBT tool used for testing) does not currently supports bags.

The full state space of the election system is defined by schema  $ES$  (for election system) as follows:

$$ES = Endorsements \wedge Votes \wedge Today$$

where schema  $Today$  declares a single variable representing the current date.

Once we have the state space of the system, i.e. schema  $ES$ , we can specify the state transitions. Due to space restrictions, we will show only one of them – the rest is available in the technical report. State transitions are called operations, in  $Z$ . Each operation is given by a schema – as can be seen, in  $Z$  schemas are used for many purposes. Below we give the schema for the operation describing the conditions for a voter to issue his or her vote. The schema is called  $Vote$  and is given as the disjunction of eight schemas:

$Vote = =$

- (1)  $VoteOk$
- (2)  $\vee VoteWrongDate$
- (3)  $\vee ResearcherCannotVote$
- (4)  $\vee AlreadyVoted$
- (5)  $\vee VoteMoreThanThree$
- (6)  $\vee VoteNonCandidates$
- (7)  $\vee KACandidatesIsWrong$
- (8)  $\vee GRCandidatesIsWrong$

Think of such an expression as a nested conditional sentence of a programming language: if... then... else if ....  $VoteOk$  formalizes the situation when a voter successfully issues his or her vote; all the other schemas describe possible errors, and so the vote is not recorded.

In  $VoteOk$ , shown below,  $e?$  is the researcher issuing the vote and  $C?$  is his or her vote.

<i>VoteOk</i>	
$\Delta Votes$	
$\Xi Endorsements; \Xi Today$	
$e? : RSCH$	
$C? : \mathbb{F} RSCH$	
$rsp! : MSG$	
$bVote \leq today \leq fVote$	Pre – 1
$e? \in CanVote$	Pre – 2
$e? \notin voted$	Pre – 3
$\#C? \leq 3$	Pre – 4
$C? \subseteq firmCand$	Pre – 5
$\{myKA\ e?\} = myKA(\!C?)$	Pre – 6
$\#C? = \#(myGR(\!C?))$	Pre – 7
$voted' = voted \cup \{e?\}$	Pos – 1
$votes' = votes \wedge \langle C? \rangle$	Pos – 2
$firmCand' = firmCand$	

For now, do not pay too much attention to the expressions  $\Delta Votes$  and  $\Xi Endorsements; \Xi Today$ , we will explain them later. As we have said,  $Z$  operations are state transitions. Then, they link two states of the state space: the start or before state and the after state. Think of states as the initial and final contents of a table in a database. For example, if a method adds a record to a table it has changed the state of the table because the initial table is different from the final one.  $Z$  operations can also depend on input and output variables. Input variables are decorated with a question mark. So, for instance  $e?$  and  $C?$  are input variables. Think of input variables as the parameters waited by a method. Output variables are decorated with a bang symbol. So,  $rsp!$  is an output variable. Think of output variables as the values returned by a method.

Any  $Z$  operation is described by giving its pre- and post-conditions. Pre-conditions are predicates depending only on before-state and input variables. Before-state variables are the variables declared in schemas  $Votes$  and  $Endorsements$ . Post-conditions are predicates that can depend on any kind of variable. Usually, post-conditions depend on before-state and after-state variables. After-state variables are the before-state variables but decorated with a prime symbol. So, for example,  $voted'$  is an after-state variable. Before-state and after-state variables are made accessible to  $Vote$  by the expressions  $\Delta Votes$  and  $\Xi Endorsements$ . Indeed,  $\Delta Votes$  means “ $Vote$  may change the value of some of the state variables declared in  $Votes$ ”, while  $\Xi Endorsements$  means “ $Vote$  cannot change the value of the state variables declared in  $Endorsements$ ”.

Now, let us take a closer look to  $Vote$ . Recall that the vote,  $C?$ , may contain up to three candidates, so it is defined as a finite set of researchers. Then, we include the pre-condition  $\#C? \leq 3$  to enforce that at most three candidates can be voted.  $bVote$  and  $fVote$  are two “dates”, representing the initial and final date of the voting period.  $today$  is a state variable declared in schema  $Today$ , which represents the current date. So the precondition  $bVote \leq today \leq fVote$  says that a vote can be issued only of the current

date is between the initial and final dates of the voting period. *CanVote* is the set of CONICET's researchers that are allowed to vote. It is defined as follows:

$$Can\ Vote == \left\{ i: RSCH \mid myKA\ i \in eKA \wedge (i \in ret \Rightarrow i \in hiredRet) \wedge i \in inIntranet \right\}$$

where *myKA* is a function that yields the *KA* of each researcher; *eKA* is the set of *KA*'s for which the current election is open; *ret* is the set of retired researchers; *hiredRet* is the set of retired researchers that has been hired by CONICET's to continue with their duties; and *inIntranet* is the set of researchers that have an user account in CONICET's Intranet.

Two interesting conditions of this election system are that voters can vote only candidates of their same *KA* ( $\{myKA\ e?\} = myKA(|C?|)$ ), but voters must vote candidates all from different *GR* ( $\#C? = \#(myGR(|C?|))$ ). In *Z*, ( $|_|\_$ ) is the relational image of a set under a function; and  $\#$  is the cardinality operator.

Under all these conditions, the system shall add the voter to the set of people who has already voted and the vote to the logical ballot box. This is expressed by the post-conditions:

$$\begin{aligned} voted' &= voted \cup \{e?\} \\ votes' &= votes \cap \langle C? \rangle \end{aligned}$$

where  $\cup$  is the classical union set operator and  $\cap$  concatenates two sequences.

Recall that *Vote* is defined as the disjunction of eight schemas where *VoteOk* describes the conditions when a vote is successfully issued and all the others represent error conditions. Below, we include one of these "error" schemas:

$$AlreadyVoted == [\exists ES; e?: RSCH \mid e? \in voted \text{ Pos } -3]$$

Note that this schema includes the expression  $\exists ES$  meaning that it will not change the value of any of the variables declared in *ES*. So *AlreadyVoted* is saying that if the person willing to issue a vote, (*e?*) has already voted ( $e? \in voted$ ) then nothing will change.

All the operations defined in the model share the same structure (i.e. a schema defined as the disjunction of some schemas one of which represents the successful case and the rest represent erroneous situations).

## 2.2 State invariants

A state invariant is a predicate, depending on state variables, that is true of every state of the state space. As can be seen, we have not included state invariants in the state schemas as is customary in *Z*. Rather, we write them in a separate schema as follows[4]:



*Invariants**ES* $cand \subseteq CanBeCandidates$  $dom\ endors = dom\ cand$  $\bigcup(\text{ran}\ endors) \subseteq CanVote$  $\forall i : dom\ endors \bullet endors\ i \neq \emptyset \Rightarrow \{myKA\ i\} = myKA(\{endors\ i\})$  $\forall i, j : dom\ endors \mid i \neq j \bullet endors\ i \cap endors\ j = \emptyset$  $today \leq fEndors \Rightarrow firmCand = \emptyset$  $fEndors < today \Rightarrow firmCand = dom(endors \triangleright \{A : \text{ran}\ endors \mid 20 \leq \#A\})$  $today < bVote \Rightarrow voted = \emptyset \wedge votes = \langle \rangle$  $\bigcup(\text{ran}\ votes) \subseteq firmCand$  $\#votes = \#voted$  $voted \subseteq CanVote$ 

and then we include a proof obligation such as the following for each operation in the model:

**theorem** VotePreservesInvariants

$$Invariants \wedge Vote \Rightarrow Invariants'$$

where *Vote* is the schema defined in the previous section and *Invariants'* is the same than *Invariants* but where all the state variables are decorated with a prime. Then if theorem VotePreservesInvariants is proved, it means that if the state invariants hold before executing *Vote*, then they will hold right after its execution. In other words, *vote* preserves the invariants.

If invariants are codified in this way then all preconditions must be explicit (Jacky, 1996; Saaltink, 1999; Spivey, 1992) and Fastest becomes more efficient. In a sense, this way of writing invariants changes “calculate the full precondition of an operation” by “discharge a proof obligation” (Cristiá *et al.*, 2014; Section 2.7 for a few more details).

Some of the predicates in *Invariants* are the formalization of CONICET rules. For example, see the seventh predicate in *Invariants*. CONICET established that after the end of the period reserved for endorsements, only candidates who have got at least 20 endorsements become firm candidates. In the model, *fEndors* represents the final date of the endorsements period. In  $Z$ , *dom* is the domain of a function; *ran* is the range of a function; and if *f* is a function and *A* is a set, then  $f \triangleright A$  (called range restriction) is the function obtained from *f* by removing all the ordered pair whose second component belongs to *A*. Therefore:

$$fEndors < today \Rightarrow firmCand = dom(endors \triangleright \{A : \text{ran}\ endors \mid 20 \leq \#A\})$$

That is, if the current date (*today*) is after the end of the endorsement period (*fEndors*), then the set of firm candidates (*firmCand*) must be the domain of *endors* restricted to those who got at least 20 endorsements. Recall that *endors* is a state variable declared in schema *Endorsements* as  $RSCH \rightarrow \mathbb{F}RSCH$ , and whose meaning is that *endors*(*e*) is the set of researchers who endorsed *e* for the current election.

Other invariants arise from the model itself. For instance, the domains of *cand* and *endors* must always be the same because even a candidate who got no endorsements is represented as having an empty set of endorsements.

### 3. Proving properties of the specification

The Z specification was verified under the Z/EVES system (Saaltink, 1997b). Therefore, we divided the task in two steps:

- (1) discharging all the proof obligations automatically generated by Z/EVES; and
- (2) proving that all the operations preserve all the invariants.

Z/EVES generates a proof obligation, called domain check, every time a partial function is applied to an argument[5]. The proof obligation asks to prove that the argument belongs to the domain of the function. For example, the following is part of the proof obligation generated by Z/EVES for the *Vote* operation:

**theorem** axiom Vote\$domainCheck  
 $\Delta Votes \wedge \exists Endorsements \wedge \exists Time$   
 $\wedge e? \in RSCH \wedge C? \in FRSCH$   
 $\Rightarrow (iVote \leq today \wedge \dots \wedge \{myKA\ e?\} = myKA(|C?|))$   
 $\Rightarrow C? \in \text{dom } \# \wedge myGR(|C?|) \in \text{dom } \#$   
 .....

where  $C? \in \text{dom } \#$  is trivial to prove, as the domain of  $\#$  is  $\mathbb{F}X$  for any type  $X$ ; but  $miRE(|C?|) \in \text{dom } \#$  is more difficult because it requires to prove that the relational image of a finite set through a function is a finite set. We proved this and other theorems involving the cardinality operator with the help of the extension to the Z mathematical toolkit (ZMT) (Saaltink, 1997a) proposed by Freitas (2004).

The second step in verifying the specification consisted in proving that all operations preserve the state invariants described in schema *Invariants* shown in page 8. We did this by proving a theorem like *VotePreservesInvariants* shown at page 8. However, proving such a theorem can be cumbersome because some parts of the proof can be quite difficult, and there are many cases to consider, as each operation includes many schemata and there are many invariants. But, on the other hand, many of these cases are trivial to prove. Hence, we analyzed which invariants would be, in principle, non-trivial to prove for a given operation, and thus, we defined a theorem for each of these cases. For example we have:

**theorem** EndorseOkPI5  
 $\text{dom } cand = \text{dom } endors$   
 $\wedge (\forall i, j: \text{dom } endors | i \neq j \bullet endors\ i \cap endors\ j = \emptyset)$   
 $\wedge Endorse\ Ok$   
 $\Rightarrow (\forall i, j: \text{dom } endors' | i \neq j \bullet endors'\ i \cap endors'\ j = \emptyset)$

This theorem is an intermediate theorem for the operation, representing a voter endorsing a preliminar candidate. Note that in this case, we prove that only the schema corresponding to the successful case (*EndorseOk*) preserves only the fifth invariant (see

the fifth predicate in schema *Invariants* at page 8). There are three more theorems like this one for the *Endorse* operation. Once all these theorems are proved, the main theorem (i.e. that establishing that *Endorse* preserves all the invariants) can be proved easily by iterating over all the schemata that define the operation and all the invariants. A similar approach was applied for all the operations.

Although the ZMT includes a number of theorems about all the mathematical theories supported by the Z notation and the Z/EVES system uses them for proofs, from time to time, it is necessary to prove some new result about mathematics. In this project, we needed to prove five of these theorems. For example, we proved that the relational image of a finite set through a partial function is a finite set:

**theorem** grule finiteRelimgIsFinite [X, Y]  
 $\forall f: X \mapsto Y; A: \mathbb{F}X \cdot f(|A|) \in \mathbb{F}Y$

This theorem was necessary to prove two domain checks. As an another example, we can mention the following theorem which is a particular form of the *applyInRanPfun* theorem of the ZMT and was used in the proofs, confirming that operation *PublishCandidates* preserves the fourth and fifth invariants[6]:

**theorem** applyInRanSinglePfun [X, Y]  
 $\forall A: \mathbb{P}X; B: \mathbb{P}Y \cdot \forall f: A \mapsto B \cdot \forall x: \text{dom } f; y: Y \cdot \text{ran } f = \{y\} \Rightarrow fx = y$

All the proofs performed during this step gave us a reasonable confidence that the specification is a faithful formalization of the requirements document; hence, the specification can be used as the guide for verification.

#### 4. Specification-guided code inspection

Having confidence on the correctness of the specification is of no help to users if it is not used as a means to gain confidence on the correctness of the implementation. Our approach to gain confidence on the correctness of the implementation was based on two verifications activities:

- (1) inspecting the code to see if it refines the specification; and
- (2) run some test cases derived from the specification.

Clearly, the first activity is static, while the second is dynamic; thus, we made use of two fundamental verification strategies. In this section, we explain how we used the specification to inspect the source code of the application.

Before showing how we inspected the code, we will give some insight on the implementation of the voting system. As we have said in the introduction, the application was developed by CONICET's System Department. They decided to program it as a Grails application. Grails ([SpringSource, 2013a](#)) is a Web application framework for the Java Virtual Machine which, in turn, takes advantage of the Groovy programming language ([SpringSource, 2013b](#)). That is, the application is an object-oriented program written in a combination of high-level programming languages (Java and Groovy) and complex frameworks (Grails). We can neither show nor make public the full application for confidentiality issues, but we will show excerpts of it.

We used the specification as the guide to conduct the inspection of the source code. Therefore, the first step was to identify the data structures and variables that refine the state variables declared in schemata *Endorsements* and *Votes* (see page 4) and the input variables declared in the operation schemata like *Vote*. This is documented in a table similar to Table I. Note that building such a table might not be trivial because the specification was written from the requirements document and not from the code. For example, as is said in the requirements document, the system stores the name of each candidate plus his or her CV and photo, but we decided to abstract away these details in the specification. As another example, the requirements document, and thus the specification, does not mention data encryption at all, but it is used in the implementation.

We will comment on the first three rows of Table I. The implementation of variable *firmCand* is a table in a database called *BallotPaper* which holds data about the candidates that voters may vote and a variable in memory, *papers*, which, at some point, is assigned with the contents of that table through a Grails' mechanism. In turn, variable *voted* is implemented by saving the system date when the voter votes in his or her database record. Voters are held in a table named *Elector*, one of whose fields is the identification of the voter in the system, and the other is the date when the voter has voted; this table acts as the electoral roll. *votes* is refined by a table in the same database named *EncryptedVote*. The table's name reflects the fact that, due to security considerations, developers decided to encrypt each vote and then present it in the table. Each of the elements stored in *EncryptedVote* is the result of encrypting an instance of class *Vote*. In turn, *Vote* stores not only the name of candidates but also their GR, although that could be traced through the data model.

Once we had clear how state variables were implemented, we identified all the major operations of the specification. This involves finding the relation between the signature of operations at the specification level and subroutines in the implementation, and checking whether they match. We will analyze the case of a person issuing his or her vote – i.e. *Vote*. The code dealing with voting is divided in several files two of which, *VoteController.groovy* and *VoteService.groovy* are the main ones. Control flow starts in *VoteController.groovy*. There the system presents to the user all the candidates, he/she can select some of them and press a button. Then, a method, called *issue()*, in *VoteService.groovy* is called, and if everything is

Specification	Implementation	Comments
<i>firmCand</i>	<i>BallotPaper</i> , <i>papers</i>	<i>BallotPaper</i> is a table in the database; <i>papers</i> is a variable in memory.
<i>voted</i>	<i>Elector</i>	Is a table in the database. If date is non-empty, elector has voted.
<i>votes</i>	<i>EncryptedVote</i>	Is a table in the database. Each record is an encrypted instance of vote.
<i>C?</i>	<i>Vote</i>	Includes also the election and GR of candidates.

**Table I.**  
Implementation variables refining specification variables

right, the vote is saved in the database. If one looks at the first method executed in `VoteController.groovy`, then its interface mismatches that of *Vote* in the specification because it does not wait any input parameters. The refinement of the interface of *Vote* is found in the signature of `issue()` which is called from `VoteController.groovy`.

The last step during the code inspection was to check whether all the pre- and post-conditions are correctly implemented. Due to the differences in the structure of the implementation and the specification and those introduced by the (implicit or *de facto*) refinement made by programmers, this step is perhaps the most difficult to do, although the more likely to uncover errors. We will continue with the analysis of the operation *Vote*. First, we labeled each pre- and post-condition in the specification as can be seen in schema *VoteOk* on page 2.1. These identifiers can latter be used to annotate the implementation as a means to link both representations. Second, we used the information gathered in the two previous steps to focus the inspection on specific implementations units. For example, to inspect the implementation of *Vote*, we read code in `VoteController.groovy` and `VoteService.groovy`, and we looked for the refinement of variables *voted*, *firmCand*, *votes*, etc. Every time we found the implementation of a pre- or post-condition, we annotated the program as follows:

```
class VoteController{
    VoteService voteService
    def beforeInterceptor = {
        def userPrincipal = (AttributePrincipal)
            RCH.currentRequestAttributes().request.userPrincipal
        def elector =
            Elector.findById(userPrincipal.attributes.usnum)
        //Begin Vote::Pre-3, Vote::Pos-3
        if (elector.voteDate) {
            render 'You already voted.'
            return false
        }
        //End Vote::Pre-3, Vote::Pos-3
    }
}
```

Note the comments before and after the conditional structure. Pre – 3 is one of the preconditions of schema *VoteOk* ( $e \notin voted$ ); Pos – 3 is a post-condition present in schema *AlreadyVoted*. Observe how the precondition was implemented: instead of having a table or file storing all the persons who have issued their votes so far, programmers decided to augment the electoral roll with an extra column that if empty means that the corresponding person has not voted and, otherwise, it stores the date when the person issued his or her vote. As can be seen, pre- and post-conditions may be annotated together because the “else” branch of conditional sentences is not always present.

In `VoteService.groovy`, the method `record()` receives the ID of the voter and the encrypted vote and, among other things, checks again if the user can vote and, if not, it writes the system date in the record of the electoral roll corresponding to the user:

```
private record(String userId, byte[] encryptedVote) {  
.....  
//Begin Vote::Pre-3, Vote::Pos-3  
Elector elector = Elector.findById(userId)  
if (elector.voteDate)  
    throw new RuntimeException("User already voted.")  
//End Vote::Pre-3, Vote::Pos-3  
.....  
//Begin Vote::Pos-1  
elector.voteDate = new Date()  
//End Vote::Pos-1  
}
```

The reader can see that we repeated the annotation in this piece of code, and we added an annotation regarding the first post-condition of schema *VoteOk*.

Clearly, the mere presence of a sentence implementing a given condition in the specification does not guarantee the correctness of the implementation. Correctness depends also on the sentences before and after the one that has been annotated. However, once the implementation has been inspected, evaluated as correct and annotated in this way, a convenient IDE can assist the development team during maintenance because the tool can bring specific predicates of the specification into attention, can show all the pieces of code implementing a given condition, etc. For example, note that we annotated the beginning of Pre – 3 in *VoteController* after the value of *elector* is set (see the *def* instruction right above the annotation). Clearly, if this variable is incorrectly assigned, the precondition will be incorrectly implemented. Therefore, if during maintenance, the assignment sentence is changed or a new assignment to *elector* is added, a convenient IDE can detect that this variable is used in a critical precondition and warn the programmer.

## 5. Generating test cases from the specification

The *Z* specification was also used to generate test cases to exercise the implementation. Testing the implementation is important even after code inspection because many third-party components with which the application interacts may fail. This is true in general but particularly important for the application being considered because it is implemented over and interacts with very complex components like the JVM, Groovy, Grails, MySQL, etc.

Test cases were generated by following a model-based testing method known as Test Template Framework (TTF) and by using Fastest, a tool that semi-automates the TTF. Given that the TTF and Fastest have been extensively described (Cristiá and Plüss, 2010; Cristiá *et al.*, 2011a; Cristiá *et al.*, 2014; Cristiá *et al.*, 2011b; Cristiá and Frydman, 2012; Stocks and Carrington, 1996), here we will show, by means of a running example, how we applied them to generate test cases. The TTF and Fastest are used for unit and functional testing.

Assume the specification has been loaded in Fastest. Consider schema *Vote* (shown at page 5) and its input space. The input space of an operation is a *Z* schema declaring all the input and state variables of the operation and not restricted by a predicate. First, we partition the input space of *Vote* by applying the so-called testing tactics. For example, we can apply Disjunctive Normal Form (DNF) in which case Fastest generates nine test

conditions, including the following three (where  $Vote^{IS}$  is the schema representing the input space of  $Vote$ ):

$Vote_1^{DNF}$ $Vote^{IS}$ <hr/> $bVote \leq today \leq fVote$ $e? \in CanVote$ $e? \notin voted$ $0 \leq \#C?$ $\#C? \leq 3$ $C? \subseteq firmCand$ $\{myKA\ e?\} = myKA(\!C?)$ $\#C? = \#(myGR(\!C?))$	$Vote_3^{DNF}$ $Vote^{IS}$ <hr/> $fVote < today$
	$Vote_5^{DNF}$ $Vote^{IS}$ <hr/> $e? \in voted$

Note that these schemata include only input and (unprimed) state variables. DNF guarantees that the main situations described in  $Vote$  are going to be tested. For instance, a user successfully issuing the vote ( $Vote_1^{DNF}$ ); a user trying to vote after the election ( $Vote_3^{DNF}$ ); and a user trying to vote more than once ( $Vote_5^{DNF}$ ). Any of these test conditions can be further partitioned by applying other testing tactics. For example, we can partition  $Vote_1^{DNF}$  by applying a standard partition to the operator  $\cup$  in  $voted' = voted \cup \{e?\}$ , thus yielding the following test conditions among others:

$Vote_1^{SP}$ $Vote_1^{DNF}$ <hr/> $hanVotado = \{\}$ $\{e?\} = \{\}$	$Vote_4^{SP}$ $Vote_1^{DNF}$ <hr/> $hanVotado \neq \{\}$ $\{e?\} \neq \{\}$ $hanVotado \cap \{e?\} = \{\}$
$Vote_2^{SP}$ $Vote_1^{DNF}$ <hr/> $hanVotado = \{\}$ $\{e?\} \neq \{\}$	$Vote_7^{SP}$ $Vote_1^{DNF}$ <hr/> $hanVotado \neq \{\}$ $\{e?\} \neq \{\}$ $\{e?\} = hanVotado$

As can be seen, some of the conditions are unsatisfiable ( $Vote_1^{SP}$ , for instance), but Fastest implements an algorithm that can eliminate many of them (Cristiá *et al.*, 2010). Note that these test conditions include the conditions of  $Vote_1^{DNF}$  because this schema is included in the others. Therefore, for example, a test case derived from  $Vote_2^{SP}$  will exercise the system when a user is allowed to vote, issues a valid vote and is the first person to vote; while  $Vote_4^{SP}$  will do the same but there should be another user who has already voted.

More testing tactics can be applied to further partition the input space of the operation. For instance, another standard partition can be applied to  $\cap$ , in  $votes' = votes \cap \langle C? \rangle$ , to partition all the satisfiable children of  $Vote_1^{DNF}$ . Test conditions can be arranged in a so-called testing tree by following the schema inclusions. Partitioning is

automatically done by Fastest while the elimination of unsatisfiable test conditions is semi-automatic.

Once the engineer is done with partitioning, he or she can generate test cases from the leaves of the testing tree. In the TTF, generating a test case means finding constants that satisfy the predicate of the test condition. In Fastest, this is done by calling a satisfiability algorithm[7]. As with the elimination of unsatisfiable test conditions, this algorithm is necessarily incomplete and thus semi-automatic. However, according to our experiments, in average, the tool automatically finds test cases for at least 80 per cent of the satisfiable test conditions; the rest must be calculated by the user. Below we show two typical test cases generated by Fastest (*rsch0* is a constant of type *RSCH* identifying a particular researcher):

$Vote_2^{TC}$
$Vote_2^{SP}$
<pre> today = 5 e? = rsch5 C? = {rsch0} firmCand = {rsch0} voted = {} votes = {} </pre>

$Vote_4^{TC}$
$Vote_4^{SP}$
<pre> today = 5 e? = rsch5 C? = {rsch0} firmCand = {rsch0} voted = {rsch0} votes = {{rsch0}} </pre>

These test cases, however, are Z terms that cannot be executed by the application. Usually, these test cases are called abstract test cases given that they are at the same level of abstraction of the specification. Hence, we gave developers precise instructions on how to write JUnit (Saff, 2014) test cases from the Z test cases, and we translated some of them as examples. For the translation, we used some of the information gathered during the code inspection activity. For example, we used the data structures we identified as the refinement of the state and input variables. That is, each activity provides useful data for the other activities.

It is important to remark that, although we had to write the JUnit test cases by hand, Fastest and the TTF were very helpful because without them, we also would have to design the test cases. Test case design is the most time-consuming and error-prone activity of testing (Jones and Bonsignour, 2011). JUnit users need to design the test cases by hand and then write them in JUnit. Furthermore, according to (Cristiá *et al.*, 2011b), we proposed FTCRL, a method to refine abstract test cases generated by Fastest. In other words, FTCRL would semi-automate the translation from Fastest test cases into



test cases similar to JUnit test cases. However, the implementation of the method is not ready yet.

## 6. Discussion

The Z specification presented here is about 450 lines of Z code (L<sup>A</sup>T<sub>E</sub>X mark-up) in a 20-page document. The implementation has approximately 2,575 lines of Grails and Java code. We proved 31 theorems, 7 of which were discharged automatically. Nine of the theorems are domain checks automatically generated by Z/EVES; 5 are theorems about mathematical properties; 6 are the main theorems (i.e. those proving that each operation preserves all the invariants); and 11 are auxiliary theorems about the specification – i.e. each of them proves that the main schema of each operation preserves one invariant. The proof scripts total 2,045 lines of commands. We generated only 68 test cases, mainly because we were running out of time and because we knew developers will not have time to run more. However, we estimate that Fastest could generate around 200 test cases. The test cases that were generated cover the main functional alternatives of all the operations.

The specification was written in 20 man-hours and its verification took around 80 man-hours. Code inspection was performed in 32 man-hours, although it could have been less provided we have had a deeper knowledge of the implementation technologies. Test case generation required to translate the specification from Spivey's *de facto* standard into the ISO standard, although, in this case, it was very simple. Generating the test cases took less than 8 man-hours, mainly because the specification contains many axiomatic descriptions whose values are underspecified and must be fully specified before test case generation. Then, our work totaled around 140 man-hours. We do not have figures about the time consumed during test case refinement and execution because it was performed by developers.

During the verification of the specification, four minor errors were found. Three of them concerned with declaring some variables as  $\mathbb{P}$  instead of  $\mathbb{F}$ . The remaining one was a missing precondition in one operation. The code inspection revealed that only two operations of the model were implemented (*Vote* and *Recount*). That is, all the information regarding preliminary and firm candidates (i.e. their endorsements, KA and GR) and the electoral roll is loaded manually from different sources. We are not sure whether top management was aware of this fact. In the implemented operations, code inspection revealed no errors. Performing a code inspection before engaging in testing can make testing very cost-effective, as many errors are discovered during the inspection. Furthermore, the code inspection yields a documented project (specification and implementation), making it easier to introduce changes and bug-fixes. As far as we have been told, test cases were run and all the errors found were corrected.

We know that there have been verification efforts producing verified software that involve larger specifications and implementations than the one presented here (see Section 7). However, our project presents crucial differences with many, if not all, of them. These other projects are performed entirely, including the programming stage, by experts on formal verification – many of them are either PhD students within first-class research groups or people already holding a PhD. Furthermore, in all these projects, the implementation is developed once the specification has been written and proved correct. Indeed, many formal techniques have been thought to be applied under the assumption of a correct specification. In many of these projects, the creators of the techniques being

applied are those who run them. In contrast, the project presented here involves the verification of a program that was developed by average programmers before the verification activities were even thought. This means, for example, that we were not allowed to choose the programming language. In other words, verification was an afterthought and the implementation was not done by researchers, as it is in almost all the software produced nowadays. It is true, however, that the verification was performed by researchers. In summary, we think that the project presented here is closer to the way the software industry works, making it appealing to evaluate the applicability of formal methods. It remains as a challenge to run a larger project but in a similar context.

## 7. Related work

The Z notation, and its extensions, is being used for formal specification since the early eighties. It has been used to specify a wide range of systems; we will mention just a few of the latest specifications to give an idea of the kind of requirements formalized with Z. Perhaps one of the most praised Z specifications in recent years is the Tokeneer abstract specification (Barnes *et al.*, 2006). This project was an experiment performed by the NSA to prove that formal methods are cost-effective in real-world software. Altran Praxis from UK was finally hired for the job. They wrote a Z specification that was latter verified to be correct by proving some security properties. Also a low-level design was written in Z and the SPARK code was formally verified and reviewed. The net result was that only two errors were discovered after delivery although more recent studies show that more errors exist (Moy and Wallenburg, 2010). Altran Praxis has used these technologies in many projects of different application domains.

Cristiá *et al.* (2011a) have used Z for modeling aerospace software such as satellite communication protocols, part of a launching vehicle control software and the ECSS-E-70-41A aerospace standard. In all these projects, the TTF and Fastest were applied to generate test cases.

Frydman and her colleagues have combined Z with DEVS (Zeigler *et al.*, 2001) for the validation of discrete event systems via simulation and formal methods (Sqali *et al.*, 2009; Trojet *et al.*, 2009). Object-Z has been used in the formalization of the Web Service Modeling Ontology for the Semantic Web (Wang *et al.*, 2012) and also Z and Z/EVES were applied in the same domain (Khan *et al.*, 2012). The main purpose of these works was to provide a precise and unambiguous specification for concepts that have traditionally been informal.

Security systems have also been the focus of Z specifications. For example, Haidar *et al.* (2009) combined Z and CSP (Roscoe, 1997) to provide a formal specification for the Audited Credential Delegation architecture which would help virtual organization in managing the identities of their users. Security is often a critical aspect of some systems, but there are systems that are critical in their own. Gomes and Oliveira (2009) have written a Z specification for a cardiac pacing system which is one of the challenges proposed as part of the Verification Grand Challenge. This specification comprises 4,000 lines of Z. As a last example, we can mention the specification of the safety properties of a railway interlocking system (Zafar *et al.*, 2012), which is one of the traditional targets of formal specification due to the potential damages a failure may cause.

Z/EVES has been the proof assistant used in many projects where Z was the specification language. Khan *et al.* (2012) and 32 use Z/EVES to discharge the proof

obligations automatically generated by the tool itself. However, some have used Z/EVES to prove properties of the specification as a mean to gain confidence on its correctness. For example, [Sun \*et al.\* \(2005\)](#) used this tool to prove 40 theorems about feature modeling. [Amálio \*et al.\* \(2006\)](#) used Z/EVES to prove properties of FTL, a formal language that allows the description of formal templates written in any formal language, in particular templates for the Z notation) and for the UML modeling notation ([Amálio \*et al.\*, 2004](#)). [Dong \*et al.\* \(2002\)](#) explore the benefits of using Z/EVES to detect inconsistencies in semantic Web services, although it seems that they did not prove a reasonable amount of properties. [Yuan \*et al.\* \(2006\)](#) used Z/EVES to prove some security properties (separation of duty, in this case) of a state-based, role-based access control (RBAC) model. [Freitas and Woodcock \(2008\)](#) have extensively used Z/EVES for proving properties of complex systems such as the Mondex Electronic Purse ([Woodcock and Freitas, 2006](#)) and a POSIX file store ([Freitas \*et al.\*, 2007](#)), contributing to the verified software repository. [Cristiá \*et al.\* \(2010\)](#) have used Z/EVES to prove that the so-called pruning rules are sound and so they can be used to eliminate inconsistent test conditions from testing trees.

Program annotation has a long tradition in the formal methods community and in other fields of programming; we will review some representative works. [Cataño and Huisman \(2002\)](#) annotate an application with ESC/Java ([Flanagan \*et al.\*, 2013](#)) in such a way that a formal functional specification is provided. The authors say that application developers might be more inclined to write formal specifications if specifications are written in a language closer to the programming language being used. While this might indeed be true, a disadvantage of this approach is that the specification becomes less abstract. This work reports a serious impact on the quality and documentation of the project. Developers at Altran Praxis annotate their SPARK programs with data and information flow clauses that are latter analyzed by the SPARK Examiner ([King \*et al.\*, 2000](#)). They also annotate programs with pre- and post-conditions to perform a functional verification. In this case, the SPARK Examiner generates proof obligations that, in general, cannot be discharged automatically. This work is also interesting because some of the SPARK annotations are derived from the Z specification by following some naming conventions and by running simple type translations. Note, however, that the broader context of the project is quite different: in the Altran Praxis case, they developed the specification and the implementation allowing them to select the latter; in our case, we could not select the implementation because it was already given – for instance, we could not decide to implement the voting system in SPARK. There are many other works investigating different aspects of program annotation but they are not closely related with the ideas presented here ([Bartetzko \*et al.\*, 2001](#); [Frade and Pinto, 2011](#); [Jacobs, 2004](#); [Marché \*et al.\*, 2004](#)).

Some properties of larger programs than the one presented here were automatically proved correct by means of many static analysis techniques. For example, [Berdine \*et al.\* \(2011\)](#) describe a tool (Slayer) that can automatically prove memory safety of industrial systems such as Windows device drivers. Another advanced tool that was applied to large, safety-critical, embedded, real-time software is presented by [Blanchet \*et al.\* \(2002\)](#). They acknowledge that their tool works for a restricted class of programs and properties. VeriFast and Frama-C are another two static analysis tools that have been applied to large programs ([Hartig \*et al.\*, 2010](#); [Philippaerts \*et al.\*, 2014](#)). Although all

these tools represent remarkable achievements in their field, the proven properties are not functional or do not fully cover a functional verification of the implementation.

We could not find works presenting the combined application of all the techniques shown in this paper in the same project, that is: formal specification, formal verification of the specification, program annotation with respect to the specification as the basis for code inspection, and model-based testing as a complementary verification activity.

## 8. Conclusions

During the verification of the voting system reported in this paper, we applied four techniques ranging from formal to informal ones. Given the resources and time available, we were able to transmit to senior management a reasonable level of confidence on the correctness of the application. The first election was carried out without any noticeable failure.

In our opinion, the most valuable contribution of this report is that these techniques were applied to an industrial project under very realistic conditions. That is, a project where average developers implemented a program whose verification was scheduled once it was finished. When this is the case, verification has a very low budget and tight schedule. In this context, these techniques proved to be effective and efficient. Moreover, the separation between a standard team of developers and a group of researchers in charge of the verification may be a good strategy in many projects. However, this setting poses some challenges to formal methods. The main reason is that the verification team has practically no influence on the implementation technologies. In turn, this implies that many formal techniques must be adapted or lightened.

We believe that the combination between a formal specification, a code inspection guided by the specification and model-based testing (as was done in this project) can be the basis of a verification methodology for mission critical applications whose verification is requested once the implementation is finished.

## Notes

1. In English: National Scientific and Technical Research Council.
2. CONICET's Systems Department is not a research unit.
3. In the project, we used both Spivey's Z for Z/EVES and standard Z (ISO, 2002) for Fastest; in the paper we use only the latter.
4. Just take a look at it for now; later we give some details about it.
5. Actually, Z/EVES also generates proof obligations for definite descriptions ( $\mu$ -terms), but we did not use them.
6. *PublishCandidates* is the first step during the election process in which CONICET publish the list of candidates looking for endorsements.
7. The algorithm currently implemented by Fastest has been described by [Cristiá et al. \(2014\)](#) but is being replaced by one based on the *{log}* (setlog) tool (see [Cristiá et al., 2012](#), for details).

## References

- Amálio, N., Stepney, S. and Polack, F. (2004), "Formal proof from UML models", in Davies, J., Schulte, W. and Barnett, M. (Eds), 6th International Conference on Formal Engineering Methods, Seattle, WA, 8-12 November, pp. 418-433.

- Amálio, N., Stepney, S. and Polack, F. (2006), "A formal template language enabling metaproof", in Misra, J., Nipkow, T. and Sekerinski, E. (Eds), *14th International Symposium on Formal Methods, Hamilton, 21-27 August*, pp. 252-267.
- Barnes, J., Chapman, R., Johnson, R., Widmaier, J., Cooper, D. and Everett, B. (2006), "Engineering the Tokeneer enclave protection software", *Proceedings of the IEEE International Symposium on Secure Software Engineering, Washington, DC*.
- Bartetzko, D., Fischer, C., Möller, M. and Wehrheim, H. (2001), "Jass – java with assertions", *Electronic Notes in Theoretical Computer Science*, Vol. 55 No. 2, pp. 103-117.
- Berdine, J., Cook, B. and Ishtiaq, S. (2011), "Slayer: memory safety for systems-level code", in Gopalakrishnan, G. and Qadeer, S. (Eds), 23rd International Conference, Snowbird, UT, 14-20 July, pp. 178-183.
- Blanchet, B., Cousot, P., Cousot, R., Feret, J., Mauborgne, L., Miné, A., Monniaux, D. and Rival, X. (2002), *The Essence of Computation*, Springer-Verlag, New York, NY, pp. 85-108.
- Cataño, N. and Huisman, M. (2002), "Formal specification and static checking of gemplus' electronic purge using ESC/Java", *Proceedings of the International Symposium of Formal Methods Europe on Formal Methods - Getting IT Right, London*, pp. 272-289.
- Flanagan, C., Leino, K.R.M., Lillibridge, M., Nelson, G., Saxe, J.B. and Stata, R. (2013), *Extended Static Checking for Java*, Compaq Systems Research Center, Palo Alto, CA.
- Cristiá, M., Albertengo, P., Frydman, C.S., Plüss, B. and Monetti, P.R. (2011a), "Applying the test template framework to aerospace software", in Rash, J.L. and Rouff, C. (Eds), *34th IEEE Software Engineering Workshop (SEW)*, Limerick, pp. 128-137.
- Cristiá, M., Albertengo, P. and Rodríguez Monetti, P. (2010), "Pruning testing trees in the test template framework by detecting mathematical contradictions", in Fiadeiro, J.L. and Gnesi, S. (Eds), 8th IEEE International Conference on Software Engineering and Formal Methods (SEFM), Pisa, pp. 268-277.
- Cristiá, M. and Frydman, C.S. (2014), "A functional verification of a web voting system", in Murgante, B., Misra, S., Rocha, A.M.A.C., Torre, C.M., Rocha, J.G., Falcão, M.I., Taniar, D., Apduhan, B.O. and Gervasi, O. (Eds), 14th International Conference, Guimarães, 30 June-3 July, pp. 640-655.
- Cristiá, M. and Plüss, B. (2010), "Generating natural language descriptions of Z test cases", in Kelleher, J.D., Namee, B.M., van der Sluis, I., Belz, A., Gatt, A. and Koller, A. (Eds), The International Natural Language Generation Conference, Dublin, pp. 173-177.
- Cristiá, M., Albertengo, P., Frydman, C.S., Plüss, B. and Rodríguez Monetti, P. (2014), "Tool support for the test template framework", *Software Testing, Verification and Reliability*, Vol. 24 No. 1, pp. 3-37.
- Cristiá, M., Rossi, G. and Frydman, C. (2012), "{log} as a test case generator for the test template framework", Technical report, Dipartimento di Matematica, Università di Parma, Parma.
- Cristiá, M., Hollmann, D., Albertengo, P., Frydman, C.S. and Monetti, P.R. (2011b), "A language for test case refinement in the Test Template Framework", in Qin, S. and Qiu, Z. (Eds), 13th International Conference on Formal Engineering Methods, Durham, 26-28 October, pp. 601-616.
- Cristiá, M. and Frydman, C.S. (2012), "Extending the Test Template Framework to deal with axiomatic descriptions, quantifiers and set comprehensions", in Derrick, J., Fitzgerald, J.A., Gnesi, S., Khurshid, S., Leuschel, M., Reeves, S. and Riccobene, E. (Eds), Third International Conference, ABZ 2012, Pisa, 18-21 June, pp. 280-293.

- Dong, J.S., Sun, J. and Wang, H.H. (2002), "Z approach to semantic web", in George, C. and Miao, H. (Eds), 4th International Conference on Formal Engineering Methods, Shanghai, 21-25 October, pp. 156-167.
- Frade, M.J. and Pinto, J.S. (2011), "Verification conditions for source-level imperative programs", *Computer Science Review*, Vol. 5 No. 3, pp. 252-277.
- Freitas, L. (2004), "Z/Eves extended Z toolkit", Technical Report, University of York, Heslington.
- Freitas, L., Fu, Z. and Woodcock, J. (2007), "Posix file store in z/eves: an experiment in the verified software repository", 12th IEEE International Conference on Engineering Complex Computer Systems, Auckland, pp. 3-14.
- Freitas, L. and Woodcock, J. (2008), "Mechanising mondex with z/eves", *Formal Aspects of Computing*, Vol. 20, pp. 117-139.
- Gomes, A.O. and Oliveira, M.V. (2009), "Formal specification of a cardiac pacing system", *Proceedings of the 2nd World Congress on Formal Methods, Berlin*, pp. 692-707.
- Haidar, A.N., Coveney, P.V., Abdallah, A.E., Ryan, P.Y.A., Beckles, B., Brooke, J.M. and Jones, M.A.S. (2009), "Formal modelling of a usable identity management solution for virtual organizations", in Bryans, J. and Fitzgerald, J.S. (Eds), *Formal Aspects of Virtual Organisations 2009 (FAVO2009)*, *Electronic Proceedings in Theoretical Computer Science (EPTCS)*, Eindhoven, pp. 41-50.
- Hartig, K., Gerlach, J., Soto, J. and Busse, J. (2010), "Formal specification and automated verification of safety-critical requirements of a railway vehicle with frama-c/Jessie", in Schnieder, E. and Tarnai, G. (Eds), *Formal Methods for Automation and Safety in Railway and Automotive Systems*, Springer, pp. 145-153.
- ISO (2002), "Information technology – Z formal specification notation – syntax, type system and semantics", Technical Report ISO/IEC 13568, International Organization for Standardization.
- Jackson, M. (1995), *Software Requirements & Specifications: A Lexicon of Practice, Principles and Prejudices*, ACM Press/Addison-Wesley Publishing, New York, NY.
- Jacky, J. (1996), *The Way of Z: Practical Programming with Formal Methods*, Cambridge University Press, New York, NY.
- Jacobs, B. (2004), "Weakest pre-condition reasoning for java programs with jml annotations", *The Journal of Logic and Algebraic Programming*, Vol. 58 Nos 1/2, pp. 61-88.
- Jones, C. and Bonsignour, O. (2011), *The Economics of Software Quality*, Addison-Wesley.
- Khan, S.A., Hashmi, A.A., Alhumaidan, F. and Zafar, N.A. (2012), "Semantic web specification using Z-notation", *Life Science Journal*, Vol. 9 No. 4.
- King, S., Hammond, J., Chapman, R. and Pryor, A. (2000), "Is proof more cost-effective than testing?", *IEEE Transactions on Software Engineering*, Vol. 26 No. 8, pp. 675-686.
- Marché, C., Paulin-Mohring, C. and Urbain, X. (2004), "The krakatoa tool for certification of java/javacard programs annotated in jml", *Journal of Logic and Algebraic Programming*, Vol. 58 Nos 1/2, pp. 89-106.
- Moy, Y. and Wallenburg, A. (2010), "Tokeneer: beyond formal program verification", *Proceeding 5th International Congress on Embedded Real Time Software and Systems (ERTS'10)*, Toulouse.
- Philippaerts, P., Mühlberg, J.T., Penninckx, W., Smans, J., Jacobs, B. and Piessens, F. (2014), "Software verification with VeriFast: industrial case studies", *Science of Computer Programming*, Vol. 82 No. 1, pp. 77-97.

- Roscoe, A.W. (1997), *The Theory and Practice of Concurrency*, Prentice Hall PTR, Upper Saddle River, NJ.
- Saaltink, M. (1999), *The Z/EVES 2.0 User's Guide*, Ora Canada.
- Saaltink, M. (1997a), "The Z/EVES mathematical toolkit version 2.2 for Z/EVES version 1.5", Technical Report, ORA Canada.
- Saaltink, M. (1997b), "The Z/EVES system", in Bowen, J.P., Hinchey, M.G. and Till, D. (Eds), 10th International Conference of Z Users, Reading, 3-4 April, pp. 72-85.
- Saff, D. (2014), "JUnit.org – resources for test driven development", available at: <http://junit.org/>
- Spivey, J.M. (1992), *The Z Notation: A Reference Manual*, Prentice Hall International, Hertfordshire.
- SpringSource (2013a), *Grails – The Search is Over*, SpringSource.
- SpringSource (2013b), *Groovy – A Dynamic Language for the Java Platform*, SpringSource.
- Sqali, M., Trojet, W., Torres, L. and Frydman, C. (2009), "Combining interaction and state based modelling to validate system specification via simulation and formal methods", Winter Simulation Conference (WSC 2009), Poster Session, Austin, TX.
- Stocks, P. and Carrington, D. (1996), "A framework for specification-based testing", *IEEE Transactions on Software Engineering*, Vol. 22 No. 11, pp. 777-793.
- Sun, J., Zhang, H. and Wang, H. (2005), "Formal semantics and verification for feature modeling", *ICECCS '05: Proceedings of the 10th IEEE International Conference on Engineering of Complex Computer Systems, Washington, DC*, pp. 303-312.
- Trojet, W., Sqali, M., Frydman, C., Torres, L. and el-amine Hamri, M. (2009), "Validating the global behaviour of a system described with scenarios using GDEVS and Z", *21th European Modeling and Simulation Symposium (EMSSO'9)*, Tenerife - Canary Islands.
- Wang, H.H., Gibbins, N., Payne, T.R. and Redavid, D. (2012), "A formal model of the semantic web service ontology (WSMO)", *Information Systems*, Vol. 37 No. 1, pp. 33-60.
- Woodcock, J. and Freitas, L. (2006), "Z/eves and the mondex electronic purse", in Barkaoui, K., Cavalcanti, A. and Cerone, A. (Eds), *Theoretical Aspects of Computing - ICTAC (2006)*, Springer Berlin Heidelberg, pp. 15-34.
- Yuan, C., He, Y., He, J. and Zhou, Z. (2006), "A verifiable formal specification for rbac model with constraints of separation of duty", in Lipmaa, H., Yung, M. and Lin, D. (Eds), *Information Security and Cryptology*, Springer Berlin Heidelberg, pp. 196-210.
- Zafar, N.A., Khan, S.A. and Araki, K. (2012), "Towards the safety properties of moving block railway interlocking system", *International Journal of Innovative Computing, Information and Control*, Vol. 8 No. 8.
- Zeigler, B.P., Kim, T.G. and Praehofer, H. (2000), *Theory of Modeling and Simulation*, Academic Press, Orlando, FL.

### Corresponding author

Maximiliano Cristia can be contacted at: [cristia@cifasis-conicet.gov.ar](mailto:cristia@cifasis-conicet.gov.ar)

---

For instructions on how to order reprints of this article, please visit our website:

[www.emeraldgroupublishing.com/licensing/reprints.htm](http://www.emeraldgroupublishing.com/licensing/reprints.htm)

Or contact us for further details: [permissions@emeraldinsight.com](mailto:permissions@emeraldinsight.com)