Verifying a plaftorm for digital imaging: a multi-tool strategy[*]

Jónathan Heras[1], Gadea Mata[2], Ana Romero[2], Julio Rubio[2], and Rubén Sáenz[2]

[1] School of Computing, University of Dundee, UK
[2] Department of Mathematics and Computer Science, University of La Rioja, Spain
jonathanheras@computing.dundee.ac.uk, gadea.mata@unirioja.es,
ana.romero@unirioja.es, julio.rubio@unirioja.es,ruben.saenz@unirioja.es

Abstract. Fiji is a Java platform widely used by biologists and other experimental scientists to process digital images. In our research, made together with a biologists team, we use Fiji in some pre-processing steps before undertaking a homological digital processing of images. In a previous work, we have formalised the correctness of the programs which use homological techniques to analyse digital images. However, the verification of Fiji's pre-processing step was missed. In this paper, we present a *multi-tool* approach (based on the combination of Why/Krakatoa, Coq and ACL2) filling this gap.

1 Introduction

Fiji [27] is a Java platform widely used by biologists and other experimental scientists to process digital images. In our research, made together with a biologists team, we use Fiji in some pre-processing steps before undertaking a homological digital processing of images.

Due to the fact that the reliability of results is instrumental in biomedical research, we are working towards the certification of the programs that we use to analyse biomedical images – here, certification means verification assisted by computers. In a previous work, see [16,17], we have formalised two homological techniques to process biomedical images. However, in both cases, the verification of Fiji's pre-processing step was not undertaken.

Being a software built by means of plug-ins developed by several authors, Fiji is messy, very flexible (program pieces are used in some occasions with a completely different objective from the one they were designed), contains many redundancies and dead code, and so on. In summary, it is a big software system which has not been devised to be formally verified. So, this endeavour is challenging.

There are several approaches to verify Java code; for instance, proving the correctness of the associated Java bytecode, see [22]. In this paper, we use Krakatoa [11] to specify and prove the correctness of Fiji/Java programs. This experience allows us to evaluate both the verification of *production* Fiji/Java code, and the Krakatoa tool itself in an unprepared scenario.

Krakatoa uses some automated theorem provers (as Alt-Ergo [5] or CVC3 [3]) to discharge the proof obligations generated by means of the Why tool [11]. When a proof obligation cannot be solved by means of the automated provers, the corresponding statement is generated in Coq [9]. Then, the user can try to prove the missing property by interacting with this proof assistant.

In this picture, we add the ACL2 theorem prover [20]. ACL2 is an automated theorem prover but more powerful than others. In many aspects, working with ACL2 is more similar to interactive provers than to automated ones, see [20]. Instead of integrating ACL2 in the architecture of Why/Krakatoa, we have followed another path leaving untouched the Why/Krakatoa code. Our approach reuses a proposal presented in [2] to translate first-order Isabelle/HOL theories to ACL2 through an XML specification language called XLL [2]. We have enhanced our previous tools to translate Coq theories to the XLL language, and then apply the tools developed in [2] to obtain ACL2 files. In this way, we can use, unmodified, the Why/Krakatoa framework; the Coq statements are then translated (if needed) to ACL2, where an automated proof is tried; if it succeeds, Coq is only an intermediary specification step; otherwise, both ACL2 or Coq can be interactively used to complete the proof.

The organization of the paper is as follows. The used tools together with our general way of working are briefly presented in Section 2. Section 3 deals with a methodology to "tame" production Fiji code in such a way that it is acceptable for Why/Krakatoa – this method is general enough to be applied to any Java code. Section 4 describes an example of the kind of specification we faced. The role of ACL2, and the tools to interoperate between Coq and ACL2, are explained in Section 5. The exposition style along the paper tries to be clear (without much emphasis on formal aspects), driven by real examples extracted from our programming experience in Fiji; in the same vein, Section 6 contains a complete example illustrating the actual role of ACL2 in our setting. The paper ends with a conclusions section and the bibliography.

All the programs and examples presented throughout this paper are available at `http://www.computing.dundee.ac.uk/staff/jheras/vpdims/`.

2 Context, tools, method

2.1 Context

Fiji [27] is a Java program which can be described as a distribution of ImageJ [26]. These two programs help with the research in life sciences and biomedicine since they are used to process and analyse biomedical images. Fiji and ImageJ are open source projects and their functionality can be expanded by means of either

a macro scripting language or Java plug-ins. Among the Fiji/ImageJ plug-ins and macros, we can find functionality which allows us to binarise an image via different threshold algorithms, homogenise images through filters such as the "median filter" or obtain the maximum projection of a stack of images.

In the frame of the ForMath European project [1], one of the tasks is devoted to the topological aspects of digital image processing. The objective of that consists in formalising enough mathematics to verify programs in the area of biomedical imaging. In collaboration with the biologists team directed by Miguel Morales, two plug-ins for Fiji have been developed (SynapCountJ [24] and NeuronPersistentJ [23]); these programs are devoted to analyse the effects of some drugs on the neuronal structure. At the end of such analysis, some homological processing is needed (standard homology groups in SynapCountJ and persistent homology in NeuronPersistentJ). As explained in the introduction, we have verified these last steps [16, 17]. But all the pre-processing steps, based on already-built Fiji plug-ins and tools, kept unverified. This is the gap we try to fill now, by using the facilities presented in the sequel.

2.2 Tools

Why/Krakatoa: Specifying and verifying Java code. The Why/Krakatoa tools [11] are an environment for proving the correctness of Java programs annotated with JML [7] specifications which have been successfully applied in different context, see [4]. The environment involves three distinct components: the Krakatoa tool, which reads the annotated Java files and produces a representation of the semantics of the Java program into Why's input language; the Why tool, which computes proof obligations (POs) for a core imperative language annotated with pre- and post-conditions, and finally several automated theorem provers which are included in the environment and are used to prove the POs. When some PO cannot be solved by means of the automated provers, corresponding statements are automatically generated in Coq [9], so that the user can then try to prove the missing properties in this interactive theorem prover. The POs generation is based on a Weakest Precondition calculus and the validity of all generated POs implies the soundness of the code with respect to the given specification. The Why/Krakatoa tools are available as open source software at `http://krakatoa.lri.fr`.

Coq and ACL2: Interactive theorem proving. Coq [9] is an interactive proof assistant for constructive higher-order logic based on the Calculus of Inductive Construction. This system provides a formal language to write mathematical definitions, executable algorithms and theorems together with an environment for semi-interactive development of machine-checked proofs. Coq has been successfully used in the formalisation of relevant mathematical results; for instance, the recently proven Feit-Thompson Theorem [13].

ACL2 [20] is a programming language, a first order logic and an automated theorem prover. Thus, the system constitutes an environment in which algo-

rithms can be defined and executed, and their properties can be formally specified and proved with the assistance of a mechanical theorem prover. ACL2 has elements of both interactive and automated provers. ACL2 is automatic in the sense that once started on a problem, it proceeds without human assistance. However, non-trivial results are not usually proved in the first attempt, and the user has to lead the prover to a successful proof providing a set of lemmas, inspired by the failed proof generated by ACL2. This system has been used for a variety of important formal methods projects of industrial and commercial interest [15] and for implementing large proofs in mathematics.

2.3 Method

In this section, we present the method that we have applied to verify Fiji code. This process can be split into the following steps.

1. Transforming Fiji code into compilable Krakatoa code.
2. Specifying Java programs.
3. Applying the Why tool.
4. If all the proof obligations are discharged automatically by the provers integrated in Krakatoa, stop; the verification has ended.
5. Otherwise, study the failed attempts, and consider if they are under-specified; if it is the case, go again to step (2).
6. Otherwise, consider the Coq expressions of the still-non-proven statements and transform them to ACL2.
7. If all the statements are automatically proved in ACL2, stop; the verification has ended.
8. Otherwise, by inspecting the failed ACL2 proofs, decide if other specifications are needed (go to item (2)); if it is not the case, decide if the missing proofs should be carried out in Coq or ACL2.

The first step is the most sensitive one, because it is the only point where informal (or, rather, semi-formal) methods are needed. Thus, some unsafe, and manual, code transformation can be required. To minimize this drawback, we apply two strategies:

– First, only well-known transformations are applied; for instance, we eliminate inheritance by "flattening" out the code, but without touching the real behaviour of methods.
– Second, the equivalence between the original code and the transformed one is systematically tested.

Both points together increase the reliability of our approach; a more detailed description of the transformations needed in step (1) are explained in Section 3. Step (2) is quite well-understood, and some remarks about this step are provided in Section 4. Steps (3)-(6) are mechanized in Krakatoa. The role of ACL2 (steps (6)-(8)) is explained in Section 5 and, by means of an example, in Section 6.

3 Transforming Fiji-Java to Krakatoa-Java

In its current state, the Why/Krakatoa system does not support the complete Java programming language and has some limitations. In order to make a Fiji Java program compilable by Krakatoa we have to take several steps.

1. Delete annotations. Krakatoa JML annotations will be placed between `\*@` and `@*\`. Therefore, we need to remove other Java Annotations preceded by `@`.
2. Move the classes that are referenced in the file that we want to compile into the directory *whyInstallationDir/java_api/*. For example, the class *RankFilters* uses the class *java.awt.Rectangle*; therefore, we need to create the folder *awt* inside the *java* directory that already exists, and put the file *Rectangle.java* into it. Moreover, we can remove the body of the methods because only the headers and the fields of the classes will be taken into consideration. We must iterate this process over the classes that we add. The files that we add into the *java_api* directory can contain `import`, `extends` and `implements` clauses although the file that we want to compile cannot do it – Krakatoa does not support these mechanisms. This is a tough process: for instance, to make use of the class *Rectangle*, we need to add fifteen classes.
3. Reproduce the behaviour of the class that we want to compile. Considering that we are not able to use `extends` and `implements` clauses, we need to move the code from the upper classes into the one that we want to compile in order to have the same behaviour. For instance, the class *BinaryProcessor* extends from *ByteProcessor* and inside its constructor it calls the constructor of *ByteProcessor*; to solve this problem we need to copy the body of the super constructor at the beginning of the constructor of the class *BinaryProcessor*. If we find the use of interfaces, we can ignore them and remove the `implements` clause because the code will be implemented in the class that makes use of the interface.
4. Remove `import` clauses. We need to delete them from the file that we want to compile and change the places where the corresponding classes appear with the full path codes. If for example we are trying to use the class *Rectangle* as we have explained in Step 2, we need to replace it by *java.awt.Rectangle*.
5. Owing to package declarations are forbidden, we need to remove them with the purpose of halting "*unknown identifier packageName*" errors.
6. Rebuild native methods. The Java programming language allows the use of *native* methods, which are written in C or C++ and might be specific to a hardware and operating system platform. For example, many of the methods in the class *Math* (which perform basic numeric operations such as the elementary exponential, logarithm, square root, and trigonometric functions) simply call the equivalent method included in a different class named *StrictMath* for their implementation, and then the code in *StrictMath* of these methods is just a *native* call. Since native methods are not written in Java, they cannot be specified and verified in Krakatoa. Therefore, if our Fiji program uses some native methods, it will be necessary to rewrite them

with our own code. See in Section 6 our implementation (and specification) of the native method `sqrt` computing the square root of a number of type double, based on Newton's algorithm.

7. Add a clause in *if-else* structures in order to remove "*Uncaught exception: Invalid_argument("equal: abstract value")*". We can find an example in the method `filterEdge` of the class *MedianFilter* where we have to replace the last *else...* clause by *else if(true)....*

8. Remove debugging useless references. We have mentioned in a previous step that we can only use certain static methods that we have manually added to the Why core code and therefore we can remove some debugging instructions like `System.out.println(...)`. We can find the usage of standard output printing statement in the method `write` of the class *IJ*.

9. Modify the declaration of some variables to avoid syntax errors. There can be some compilation errors with the definition of some floats and double values that match the pattern *<number>f* or *<number>d*. We can see an example in the line 180 of the file *RankFilters.java*; we have to transform the code: `float f = 50f;` into `float f = 50`.

10. Change the way that Maximum and Minimum float numbers are written. Those two special numbers are located in the file *Float.java* and there are widely used to avoid overflow errors, but they generate an error due to the *eP* exponent. To stop having errors with expressions like `0x1.fffffeP+127d` we need to convert it into `3.4028235e+38f`.

4 Specifying programs for digital imaging

As already said in Section 2.2, Fiji and ImageJ are open source projects and many different people from many different teams (some of them not being computer scientists) are involved in the development of the different Fiji Java plug-ins. This implies that the code of these programs is in general not suitable for its formal verification and a deep previous transformation process, following the steps explained in Section 3, is necessary before introducing the Java programs into the Why/Krakatoa system. Even after this initial transformation, Fiji programs usually remain complex and their specification in Krakatoa is not a direct process. In this section we present some examples of Fiji methods that we have specified in JML trying to show the difficulties we have faced.

Once that a Fiji Java program has been adapted, following the ideas of Section 3, and is accepted by the Why/Krakatoa application, the following step in order to certify its correctness consists in specifying its behaviour (that is, its precondition and its postcondition) by writing annotations in the Java Modelling Language (JML) [7] . The precondition of a method must be a proposition introduced by the keyword `requires` which is supposed to hold in the pre-state, that is, when the method is called. The postcondition is introduced by the keyword `ensures`, and must be satisfied in the post-state, that is, when the method returns normally. The notation `\result` denotes the returned value. To differentiate the value of a variable in the pre- and post- states, we can use the keyword `\old` for the pre-state.

Let us begin by showing a simple example. The following Fiji method, included in the class *Rectangle*, translates an object by given horizontal and vertical increments `dx` and `dy`.

```
/*@ ensures x == \old(x) + dx && y == \old(y) + dy;
  @*/
public void translate(final double dx, final double dy) {
    this.x += dx; this.y += dy;
}
```

The postcondition expresses that the field `x` is modified by incrementing it by `dx`, and the field `y` is increased by `dy`. In this case no precondition is given since all values of `dx` and `dy` are valid, and the keyword `\result` does not appear because the returned type is `void`.

Using this JML specification, Why/Krakatoa generates several lemmas (*Proof Obligations*) which express the correctness of the program. In this simple case, the proof obligations are elementary and they can be easily discharged by the automated theorem provers Alt-Ergo [5] and CVC3 [3], which are included in the environment. The proofs of these lemmas guarantee the correctness of the Fiji method `translate` with respect to the given specification.

Unfortunately, this is not the general situation because, as already said, Fiji code has not been designed for its formal verification and can be very complicated; so, in most cases, Krakatoa is not able to prove the validity of a program from the given precondition and postcondition. In order to formally verify a Fiji method, it is usually necessary to include annotations in the intermediate points of the program. These annotations, introduced by the keyword `assert`, must hold at the corresponding program point. For loop constructs (while, for, etc), we must give an *inductive invariant*, introduced by the keyword `loop_invariant`, which is a proposition which must hold at the loop entry and be preserved by any iteration of the loop body. One can also indicate a `loop_variant`, which must be an expression of type integer, which remains non-negative and decreases at each loop iteration, assuring in this way the termination of the loop. It is also possible to declare new logical functions, lemmas and predicates, and to define *ghost variables* which allow one to monitor the program execution.

Let us consider the following Fiji method included in the class *RankFilters*. It implements Hoare's find algorithm (also known as *quickselect*) for computing the nth lowest number in part of an unsorted array, generalizing in this way the computation of the median element. This method appears in the implementation of the"median filter", a process very common in digital imaging which is used in order to achieve greater homogeneity in an image and provide continuity, obtaining in this way a good binarization of the image.

```
/*@ requires buf!=null && 1<= bufLength <= buf.length && 0<=n <bufLength;
  @ ensures Permut{Old,Here}(buf,0,bufLength-1)
  @    && (\forall integer k; (0<=k<=n-1 ==> buf[k]<=buf[n])
  @            && (n+1<=k<=bufLength-1 ==> buf[k]>=buf[n]))
  @    && \result==buf[n] ;
```

```
    @*/
public final static float findNthLowestNumber
                (float[] buf, int bufLength, int n) {
    int i,j;
    int l=0;
    int m=bufLength-1;
    float med=buf[n];
    float dum ;
    while (l<m) {
        i=l ;
        j=m ;
        do {
            while (buf[i]<med) i++ ;
            while (med<buf[j]) j-- ;
            dum=buf[j];
            buf[j]=buf[i];
            buf[i]=dum;
            i++ ; j-- ;
        } while ((j>=n) && (i<=n)) ;
        if (j<n) l=i ;
        if (n<i) m=j ;
        med=buf[n] ;
    }
    return med ;
}
```

Given an array `buf` and two integers `bufLength` and `n`, the Fiji method `findNthLowestNumber` returns the $(n+1)$-th lowest number in the first `bufLength` components of `buf`. The precondition expresses that `buf` is not null, `bufLength` must be an integer between 1 and the length of `buf`, and `n` is an integer between 0 and `bufLength` -1. The definition of the postcondition includes the use of the predicate `Permut`, a predefined predicate, which expresses that when the method returns the (modified) `bufLength` first components of the array `buf` must be a permutation of the initial ones. The array has been reordered such that the components $0, \ldots, n-1$ are smaller than or equal to the component n, and the elements at positions $n+1, \ldots, $ `bufLength` -1 are greater than or equal to that in n. The returned value must be equal to `buf[n]`, which is therefore the $(n+1)$-th lowest number in the first `bufLength` components of `buf`.

In order to prove the correctness of this program, we have included different JML annotations in the Java code. First of all, loop invariants must be given for all `while` and `do` structures appearing in the code. Difficulties have been found in order to deduce the adequate properties for invariants which must be strong enough to imply the program (and other loops) postconditions; automated techniques like discovery of loop invariants [18] will be used in the future. We show as an example the loop invariant (and variant) for the exterior while, which is given by the following properties:

```
/*@ loop_invariant
```

```
@ 0<=l<=n+1 && n-1<=m<=bufLength-1 && l<=m+2
@ && (\forall integer k1 k2; (0<=k1<=n && m+1<=k2<=bufLength-1)
@           ==> buf[k1]<=buf[k2])
@ && (\forall integer k1 k2; (0<=k1<=l-1 && n<=k2<=bufLength-1)
@           ==> buf[k1]<=buf[k2])
@ && Permut{Pre,Here}(buf,0,buf.length-1) && med==buf[n]
@ && ((l<m)==> ((l<=n)&&(m>=n)));
@ loop_variant m - l+2;
@*/
```

To help the automated provers to verify the program and prove the generated proof obligations it is also necessary to introduce several assertions in some intermediate points of the program and to use ghost variables which allow the system to deduce that the loop variant decreases.

Our final specification of this method includes 78 lines of JML annotations (for only 24 Java code lines). Krakatoa/Why produces 175 proof obligations expressing the validity of the program. The automated theorem prover Alt-Ergo is able to demonstrate all of them, although in some cases more than a minute (in an ordinary computer) is needed; another prover included in Krakatoa, CVC3, is, on the contrary, only capable of proving 171. The proofs of the lemmas obtained by means of Alt-Ergo certify the correctness of the method with respect to the given specification.

In this particular example, the automated theorem provers integrated in Krakatoa are enough to discharge all the proof obligations. In other cases, some properties are not proven, and then one should try to prove them using *interactive* theorem provers, as Coq. In this architecture, we also introduce the ACL2 theorem prover, as explained in the next section.

5 The role of ACL2

In this section, we present the role played by ACL2 in our infrastructure to verify the correctness of Java programs. The Why platform relies on automated provers, such as Alt-Ergo or CVC3, and interactive provers, such as Coq or PVS, to discharge proof obligations; however, it does not consider the ACL2 theorem prover to that aim. We believe that the use of ACL2 can help in the proof verification process. The reason is twofold.

– The scope of automated provers is smaller than the one of ACL2; therefore, ACL2 can prove some of the proof obligations which cannot be discharged by automated provers.
– Moreover, interactive provers lack automation; then, ACL2 can automatically discharge proof obligations which would require user interaction in interactive provers.

We have developed *Coq2ACL2*, a Proof General extension, which integrates ACL2 in our infrastructure to verify Java programs; in particular, we work with

ACL2(r) a variant of ACL2 which supports the real numbers [12] – the formalisation of real analysis in theorem provers is an outstanding topic, see [6]. Coq2ACL2 features three main functions:

F1. it transforms Coq statements generated by Why to ACL2;
F2. it automatically sends the ACL2 statements to ACL2; and
F3. it displays the proof attempt generated by ACL2.

If all the statements are proved in ACL2; then, the verification process is ended. Otherwise, the statements must be manually proved either in Coq or ACL2.

The major challenge in the development of Coq2ACL2 was the transformation of Coq statements to ACL2. There is a considerable number of proposals documented in the literature related to the area of theorem proving interoperability. We have not enough space here to do a thorough review, but we can classify the translations between proof assistants in two groups: *deep* [8, 14, 19] and *shallow* [10, 21, 25].

In our work, we took advantage of a previous shallow development presented in [2], where a framework called *I2EA* to import Isabelle/HOL theories into ACL2 was introduced. The approach followed in [2] can be summarized as follows. Due to the different nature of Isabelle/HOL and ACL2, it is not feasible to replay proofs that have been recorded in Isabelle/HOL within ACL2. Nevertheless, Isabelle/HOL statements dealing with first order expressions can be transformed to ACL2; and then, they can be used as a schema to guide the proof in ACL2.

A key component in the framework presented in [2] was an XML-based specification language called *XLL* (that stands for Xmall Logical Language). XLL was developed to act as an intermediate language to port Isabelle/HOL theories to both ACL2 and an Ecore model (given by UML class definitions and OCL restrictions) – the translation to Ecore serves as a general purpose formal specification of the theory carried out. The transformations among the different languages are done by means of XSLT and some Java programs. We have integrated the Coq system into the I2EA framework as can be seen in Figure 1; in this way, we can reuse both the XLL language and some of the XSLT files developed in [2] to transform (first-order like) Coq statements to ACL2.

In particular, functionality **F1** of Coq2ACL2 can be split into two steps:

1. given a Coq statement, Coq2ACL2 transforms it to an XLL file using a Common Lisp translator program; then,
2. the XLL file is transformed to ACL2 using an XSLT file previously developed in [2].

In this way, ACL2 has been integrated into our environment to verify Java programs. As we will see in the following section, this has meant an improvement to automatically discharge proof obligations.

Fig. 1. (Reduced) Architecture of the I2EA framework integrating Coq.

6 The method in action: a complete example

In our work, we deal with images acquired by microscopy techniques from biological samples. These samples have volume and the object of interest is not always in the same plane. For this reason, it is necessary to obtain different planes from the same sample to get more information. This means that several images are acquired in the same XY plane at different levels of Z. To work with this stack of images, it is often necessary to make their *maximum projection*. To this aim, Fiji provides several methods such as maximum intensity or standard deviation to obtain the maximum projection of a set of images.

In this section, we consider the Fiji code for computing the maximum projection of a set of images based on the standard deviation, which uses in particular the method `calculateStdDev` located in the class *ImageStatistics*.

```
double calculateStdDev(double n, double sum, double sum2) {
    double stdDev = 0.0;
    if (n>0.0) {
        stdDev = (n*sum2-sum*sum)/n;
        if (stdDev>0.0)
            stdDev = Math.sqrt(stdDev/(n-1.0));
        else
            stdDev = 0.0;
    } else
    stdDev = 0.0;
}
```

The inputs are `n` (the number of data to be considered), `sum` (the sum of all considered values; in our case, these values will obtained from the pixels in an image) and `sum2` (the sum of the squares of the data values). The method `calculateStdDev` computes the standard deviation from these inputs and assigns it to the field `stdDev`. The specification of this method is given by the following JML annotation.

```
/*@ requires ((n==1.0)==> sum2==sum*sum) && ((n<=0.0) || (n>=1.0)) ;
  @ behaviour negative_n :
  @   assumes  n<=0.0 || (n>0.0 && (n*sum2-sum*sum)/n <=0.0);
  @   ensures stdDev == 0.0;
  @ behaviour normal_behaviour :
  @   assumes n>=1.0 && ((n*sum2-sum*sum)/n  > 0.0);
  @   ensures is_sqrt(stdDev,(double)((n*sum2-sum*sum)/n/(n-1.0)));
@*/
```

The precondition, introduced by the keyword `requires`, expresses that in the case $n = 1$ (that is, there is only one element in the data) the inputs `sum` and `sum2` must satisfy $sum2 = sum * sum$. Moreover we must require that `n` is less than or equal to 0 or greater than or equal to 1 to avoid the possible values in the interval $(0, 1)$; for `n` in this interval one has $n - 1 < 0$ and then it is not possible to apply the square root function to the given argument $stdDev/(n - 1.0)$. This fact has not been taken into account by the author of the Fiji program because in all real applications the method will be called with `n` being a natural number; however, to formalise the method we must specify this particular situation in the precondition. For the postcondition we distinguish two different behaviours: if `n` is non-positive or `sum` and `sum2` are such that $n * sum2 - sum * sum < 0$, the field `stdDev` is assigned to 0; otherwise, the standard deviation formula is applied and the result is assigned to the field `stdDev`. The predicate `is_sqrt` is previously defined.

For the proof of correctness of the method `calculateStdDev` in Krakatoa, it is necessary to specify (and verify) the method `sqrt`. The problem here, as already explained in Section 3, is that the method `sqrt` of the class *Math* simply calls the equivalent method in the class *StrictMath*, and the code in *StrictMath* of the method `sqrt` is just a native call and might be implemented differently on different Java platforms. In order to give a JML specification of the method `sqrt` is necessary then to rewrite it with our own code. The documentation of *StrictMath* states "*To help ensure portability of Java programs, the definitions of some of the numeric functions in this package require that they produce the same results as certain published algorithms. These algorithms are available from the well-known network library netlib as the package "Freely Distributable Math Library", fdlibm*". In the case of the square root, one of these *recommended* algorithms is Newton's method; based on it, we have implemented and specified in JML the computation of the square root of a given (non-negative) input of type double.

```
/*@ requires c>=0 && epsi > 0 ;
  @ ensures \result >=0 &&  (\result*\result>=c)
  @    && \result*\result - c < epsi ;
  @*/
public double sqrt(double c, double epsi){
    double t;
    if (c>1) t= c;
        else t=1.1;
    /*@ loop_invariant
```

```
      @ (t >= 0) && (t*t> c) ;
      @*/
    while (t* t - c  >= epsi) {
        t = (c/t + t) / 2.0;
    }
    return t;
}

/*@ requires c>=0  ;
  @ ensures (\result >=0) && (\result*\result>=c)
  @ && (\result*\result - c < 1.2E-7);
  @*/
public double sqrt(double c){
    double eps=1.2E-7;
    return sqrt(c,eps);
}
```

The first method computes the square root of a double x with a given precision epsi; the second one calls the previous method with a precision less than $1.2E-7$. Using JUnit, we have run one million tests between $1E9$ and $1E-9$ to show that the results of our method sqrt have similar precision to those obtained by the *original* method Math.sqrt. Here, we applied the "first test, then verify" approach – intensive testing can be really useful to find bugs (and can save us time) before starting the verification process.

From the given JML specification for the Fiji method calculateStdDev and our sqrt method, Why/Krakatoa produces 52 proof obligations, 9 of them corresponding to lemmas that we have introduced and which are used in order to prove the correctness of the programs. Alt-Ergo is able to prove 50 of these proof obligations, but two of the lemmas that we have defined remain unsolved. CVC3 on the contrary only proves 44 proof obligations.

The two lemmas that Alt-Ergo (and CVC3) are not able to prove are the following ones:

```
/*@ lemma double_div_pos :
  @ \forall double x y; x>0 && y > 0 ==> x / y > 0;
  @*/
/*@ lemma double_div_zero :
  @ \forall double x y; x==0.0 && y > 0 ==> x / y == 0.0;
  @*/
```

In order to discharge these two proof obligations, we can manually prove their associated Coq expressions.

```
Lemma double_div_zero : (forall (x_0_0:R), (forall (y_0:R),
  ((eq x_0_0 (0)%R) /\ (Rgt y_0 (0)%R) -> (eq (Rdiv x_0_0 y_0) (0)%R)))).

Lemma double_div_pos : (forall (x_13:R), (forall (y:R),
  ((Rgt x_13 (0)%R) /\ (Rgt y (0)%R) -> (Rgt (Rdiv x_13 y) (0)%R)))).
```

Both lemmas can be proven in Coq in less than 4 lines, but, of course, it is necessary some experience working with Coq. Therefore, it makes sense to delegate those proofs to ACL2. Coq2ACL2 translates the Coq lemmas to the following ACL2 ones. ACL2 can prove both lemmas without any user interaction (a screenshot of the proof of one of this lemmas in ACL2 is shown in Figure 2).

```
(defthm double_div_zero
 (implies (and (realp x_0_0) (realp y_0) (and (equal x_0_0 0) (> y_0 0)))
          (equal (/ x_0_0 y_0) 0)))

(defthm double_div_pos
 (implies (and (realp x_13) (realp y) (and (> x_13 0) (> y 0)))
          (> (/ x_13 y) 0))
```
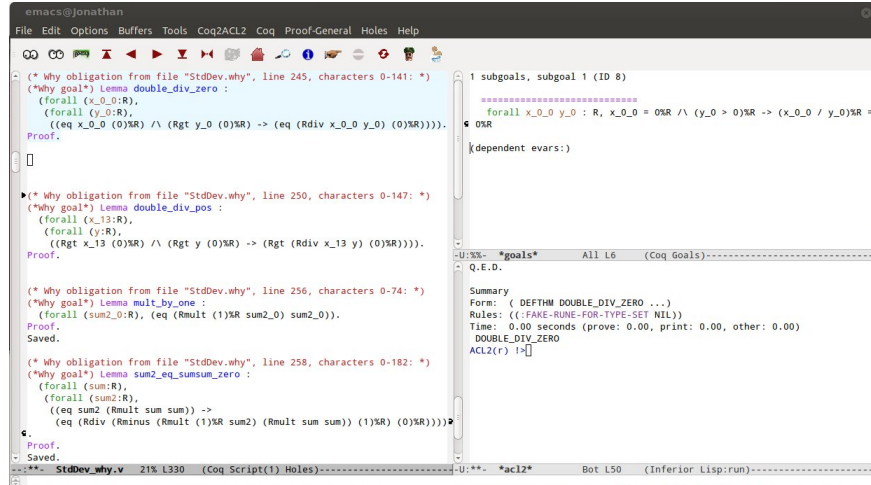

Fig. 2. Proof General with Coq2ACL2 extension. The Coq2ACL2 extension consists of the Coq2ACL2 menu and the right-most button of the toolbar. Left: the Coq file generated by the Why tool. Top Right: current state of the Coq proof. Bottom Right: ACL2 proof of the lemma.

7 Conclusions and further work

This paper reports an experience to verify actual Java code, as generated by different-skilled programmers, in a multi-programmer tool called Fiji. As one could suspect, the task is challenging and, in some sense, the objectives are impossible to accomplish, at least in their full extent – after our experiments, we have found that the Fiji system is *unsound*, but the errors are minor (e.g. a

variable declared as a real number but which should be declared as an integer) and can be easily corrected.

Nevertheless, we defend the interest of this kind of experimental work. It is useful to evaluate the degree of maturity of the verification tools (Krakatoa, in our case). In addition, by a careful examination of the code really needed for a concrete application, it is possible to isolate the relevant parts of the code, and then it is possible to achieve a complete formalisation. Several examples in our text showed this feature, see Section 4.

In addition to Krakatoa, several theorem provers (Coq and ACL2) have been used to discharge some proof obligations that were not automatically proved by Krakatoa. To this aim, it has been necessary the integration of several tools, and our approach can be considered as semi-formal: we keep transformations as simple as possible, and substantiate the process by systematic testing.

As a further interest of our work, we have reused a previous interoperability-proposal [2], between Isabelle and ACL2, to get an integration of ACL2 (through a partial mapping from Coq to ACL2), without touching the Krakatoa kernel.

Future work includes several improvements in our method. Starting from the beginning, the transformation from real Java code to Krakatoa one could be automated (Section 3 can be understood as a list of requirements to this aim). Then, a formal study of this transformation could be undertaken to increase the reliability of our method. In addition, we can try to automatically reconstruct ACL2 proofs in Coq.

As for applications, more verification is needed to obtain a certified version of, for instance, the SynapCountJ plug-in [24]. The preliminary results presented in this paper allow us to be reasonably optimistic with respect to the feasibility of this objective.

References

1. ForMath: Formalisation of Mathematics, European Project. `http://wiki.portal.chalmers.se/cse/pmwiki.php/ForMath/ForMath`.
2. J. Aransay et al. A report on an experiment in porting formal theories from Isabelle/HOL to Ecore and ACL2. Technical report, 2012. `http://wiki.portal.chalmers.se/cse/uploads/ForMath/isabelle_acl2_report`.
3. C. Barrett and C. Tinelli. CVC3. In *19th International Conference on Computer Aided Verification (CAV'07)*, volume 4590 of *LNCS*, pages 298–302, 2007.
4. G. Barthe, D. Pointcheval, and S. Zanella-Béguelin. Verified Security of Redundancy-Free Encryption from Rabin and RSA. In *Proceedings 19th ACM Conference on Computer and Communications Security (CCS'12)*, pages 724–735, 2012.
5. F. Bobot, S. Conchon, E. Contejean, M. Iguernelala, S. Lescuyer, and A. Mebsout. The Alt-Ergo automated theorem prover, 2008. `http://alt-ergo.lri.fr/`.
6. S. Boldo, C. Lelay, and G. Melquiond. Formalization of Real Analysis: A Survey of Proof Assistants and Libraries. Technical report, 2013. `http://hal.inria.fr/hal-00806920`.
7. L. Burdy et al. An overview of JML tools and applications. *International Journal on Software Tools for Technology Transf*, 7(3):212–232, 2005.

8. M. Codescu et al. Towards Logical Frameworks in the Heterogeneous Tool Set Hets. In *Post-Proceedings 20th International Workshop on Recent Trends in Algebraic Development Techniques (WADT'10)*, volume 7137 of *LNCS*, pages 139–159, 2012.

9. Coq development team. The Coq Proof Assistant, version 8.4. Technical report, 2012. `http://coq.inria.fr/`.

10. E. Denney. A Prototype Proof Translator from HOL to Coq. In *13th International Conference on Theorem Proving in Higher Order Logics (TPHOLs'00)*, volume 1869 of *LNCS*, pages 108–125, 2000.

11. J. Filliâtre and C. Marché. The Why/Krakatoa/Caduceus Platform for Deductive Program verification. In *19th International Conference on Computer Aided Verification (CAV'07)*, volume 4590 of *LNCS*, pages 173–177, 2007.

12. R. Gamboa and M. Kaufmann. Non-Standard Analysis in ACL2. *Journal of Automated Reasoning*, 27(4):323–351, 2001.

13. G. Gonthier et al. A Machine-Checked Proof of the Odd Order Theorem. In *Proceedings 4th Conference on Interactive Theorem Proving (ITP'13)*, LNCS, 2013.

14. M. J. C. Gordon, M. Kaufmann, and S. Ray. The Right Tools for the Job: Correctness of Cone of Influence Reduction Proved Using ACL2 and HOL4. *Journal of Automated Reasoning*, 47(1):1–16, 2011.

15. D. Hardin, editor. *Design and Verification of Microprocessor Systems for High-Assurance Applications*. Springer, 2010.

16. J. Heras, T. Coquand, A. Mörtberg, and V. Siles. Computing Persistent Homology within Coq/SSReflect. *To appear in ACM Transactions on Computational Logic*, 2013.

17. J. Heras, M. Poza, and J. Rubio. Verifying an Algorithm Computing Discrete Vector Fields for Digital Imaging. In *AISC/MKM/Calculemus (CICM'12)*, volume 7362 of *LNCS*, pages 216–230, 2012.

18. A. Ireland and J. Stark. On the automatic discovery of loop invariants, 1997.

19. M. Jacquel, K. Berkani, D. Delahaye, and C. Dubois. Verifying B Proof Rules Using Deep Embedding and Automated Theorem Proving. In *Proceedings 9th International Conference on Software Engineering and Formal Methods (SEFM'11)*, volume 7041 of *LNCS*, pages 253–268, 2011.

20. M. Kaufmann and J S. Moore. ACL2 version 6.0, 2012. `http://www.cs.utexas.edu/users/moore/acl2/`.

21. C. Keller and B. Werner. Importing HOL Light into Coq. In *Proceedings 1st International Conference on Interactive Theorem Proving (ITP'11)*, volume 6172 of *LNCS*, pages 307–322, 2011.

22. H. Liu and J S. Moore. Java Program Verification via a JVM Deep Embedding in ACL2. In *Proceedings 17th International Conference on Theorem Proving in Higher Order Logics (TPHOLs'04)*, volume 3223 of *LNCS*, pages 184–200, 2004.

23. G. Mata. NeuronPersistentJ. `http://imagejdocu.tudor.lu/doku.php?id=plugin:utilities:neuronpersistentj:start`.

24. G. Mata. SynapCountJ. `http://imagejdocu.tudor.lu/doku.php?id=plugin:utilities:synapsescountj:start`.

25. S. Obua and S. Skalberg. Importing HOL into Isabelle/HOL. In *3rd International Joint Conference on Automated Reasoning (IJCAR'06)*, volume 4130 of *LNCS*, pages 298–302, 2006.

26. W. S. Rasband. ImageJ: Image Processing and Analysis in Java. Technical report, U. S. National Institutes of Health, Bethesda, Maryland, USA, 1997–2012.

27. J. Schindelin et al. Fiji: an open-source platform for biological-image analysis. *Nature Methods*, 9(7):676–682, 2012.