Houdini, an Annotation Assistant for ESC/Java

Cormac Flanagan and K. Rustan M. Leino
25 August 2000
Compaq Systems Research Center
130 Lytton Ave., Palo Alto, CA 94301, U.S.A.
cormac.flanagan@compaq.com, rustan.leino@compaq.com

0 Introduction

The Extended Static Checker for Java (ESC/Java) is a tool for finding defects in Java programs [3, 12, 11]. It relies on the programmer to supply annotations describing program properties such as method preconditions, postconditions, and object invariants. These annotations allow ESC/Java to catch software defects using a method-local analysis. During this analysis, ESC/Java verifies that the annotations are consistent with the program, and it also uses the annotations to verify that each primitive operation (such as a dereference operation) will not raise a run-time exception (as might happen, for example, if a dereferenced pointer is null).

Other static checkers that follow this modular approach include conventional type checkers, which rely on type annotations to guide the type checking process, and `rccjava` [5], a static race condition checker, which relies on annotations describing the locking discipline.

A limitation of the modular checking approach is the burden on the programmer to supply annotations. Although programmers have grown accustomed to writing type annotations, they have been reluctant to provide additional annotations. In our experience, this reluctance has been the major obstacle to the adoption of ESC/Java. This annotation burden appears particularly pronounced when faced with the daunting task of applying ESC/Java to an existing (unannotated) code base.

To automate the annotation process, we have developed Houdini, an *annotation assistant* that infers suitable ESC/Java annotations for an unannotated program. Houdini reuses ESC/Java as a subroutine when inferring these annotation. Essentially, Houdini conjectures a large number of possible *candidate* annotations, and then uses ESC/Java to verify or refute each of these annotations.

This paper describes the design, implementation, and evaluation of Houdini. Our experience indicates that this approach is capable of inferring many useful annotations. These annotations significantly reduce the number of false alarms produced by ESC/Java (as compared with checking the original, unannotated program), and we have found that using Houdini significantly reduces the programmer time required to catch defects in unannotated programs.

The presentation of our results proceeds as follows. The following section starts by reviewing ESC/Java. Section 2 introduces the basic architecture of Houdini. Section 3 describes the heuristics for generating candidate annotations. Section 4 extends Houdini to handle libraries. Section 5 describes Houdini's user interface. Section 6 describes our experience using Houdini to catch defects in four test programs totalling 50,000 lines of code. Houdini is actually a third-generation annotation assistant; Section 7 outlines some prior approaches we have tried. Section 8 describes related work, and we conclude in Section 9.

1 Review of ESC/Java

ESC/Java is a programming tool for finding common programming errors in Java programs. It takes as input a Java program, possibly annotated with ESC/Java light-weight specifications, and produces as output a list of warnings of possible errors in the program. Because of its static and automatic nature, its use is reminiscent of that of a type checker. Under the hood, however, ESC/Java is powered by a more precise semantics engine and an automatic theorem prover.

ESC/Java performs modular checking: Every routine (method or constructor) is given a specification. ESC/Java checks that the implementation of each routine meets its specification, assuming that all routines called meet their specifications. The specification comes from user-supplied annotations. Note that ESC/Java does not trace into the code of a callee, even if the callee code is also given to the tool to be checked. By performing modular checking, ESC/Java can be applied to a class, or even a routine, at a time, without needing the entire program.

Figure 0 shows a simple Java class that demonstrates the use of typical ESC/Java annotations. The class implements an `n`-tuple of non-null values (of type `Object`), where `n` can be changed over the lifetime of the object. The constructor creates an empty tuple and the `put` method sets element `j` of the tuple to the given value `p`, extending the tuple size by 1 if `j == n`, and returning the previous value, if any, of this tuple element.

ESC/Java annotations are given as specially formatted Java comments. If

```
class Tuple {
  int n;
  //@ invariant 0 <= n;

  Tuple() { ... }              // constructor

  //@ requires 0 <= j;
  //@ requires j <= n;
  //@ requires p != null;
  //@ ensures j == n || \result != null;
  Object put(int j, Object p) { ... }

  ...

  Object a[];
  //@ invariant a != null;
  //@ invariant (\forall int i; 0 <= i && i < n
                                  ==> a[i] != null);
  //@ invariant n <= a.length;
}
```

Figure 0: Examples of typical ESC/Java annotations

the first character within the Java comment is an @-sign, ESC/Java parses the comment and expects to find legal annotations. The expressions occurring in ESC/Java annotations are mostly just side-effect free Java expression, but with some additions, including quantified expressions and some special keywords and functions.

The example in Figure 0 shows the `put` method to have some pre- and postconditions. The `requires` keyword declares a precondition and the `ensures` keyword declares a postcondition. The occurrences of `j` and `p` in these annotations refer to the method's parameters and `n` refers to the field declared in the class. The postcondition uses the special keyword `\result` to refer to the value returned by the method.

The example also shows several declarations of object invariants. ESC/Java checks that these are established by the constructor and maintained by other routines. (The details are actually more complicated, see the ESC/Java user's manual [11].) One of the invariant declarations uses a universal quantification and the ESC/Java implication operator `==>`.

Although ESC/Java annotations are mostly straightforward, inserting suitable annotations into a previously-unannotated program is somewhat time-consuming and tedious. Experience with ESC/Java and with a previous extended static checker for Modula-3 [1] indicates that a programmer can annotate an existing, unannotated program at the rate of 300 to 1000 lines per hour. The number of annotations inserted the into the program measures between about 6 and 10 percent of the total number of lines of code in the program.

2 Houdini architecture

Although ESC/Java works well on annotated programs, catching defects in legacy, unannotated programs using ESC/Java is a difficult process. It is possible to run ESC/Java on an unannotated program, but this produces an excessively large number of false alarms. Alternatively, one can manually insert appropriate annotations into the program, but this is a very time-consuming task for large programs.

Therefore, we would like to automate much of the annotation process by developing an *annotation assistant* that infers suitable ESC/Java annotations for a legacy, unannotated program. The following *Houdini* algorithm implements an annotation assistant. This algorithm leverages off ESC/Java's ability to perform precise procedure-local analysis.

Input: An unannotated program `P`

Output: ESC/Java warnings for an annotated version of `P`
Algorithm:
generate set of candidate annotations and insert into `P`;
repeat
  invoke ESC/Java to check `P`;
  remove any refuted candidate annotations from `P`;
until quiescence;
invoke ESC/Java to identify possible defects in `P`;

The first step in the algorithm is to generate a finite set of *candidate annotations*. This set is generated from the program text based on heuristics about what annotations might be useful in reasoning about the program's behavior. For example, since a common precondition in manually-annotated programs is that an argument of reference type is non-null, the candidate annotation set should include all preconditions of this form. Other useful heuristics for guessing candidate annotations are described in the following section.

Many of these guessed annotations will of course be incorrect. To identify these incorrect annotations, the Houdini algorithm invokes ESC/Java on the annotated program. Like any invocation of ESC/Java, this invocation may produce two kinds of warnings. The first kind concerns potential run-time errors, such as dereferencing the null pointer. These warnings are ignored by the Houdini algorithm.

The second kind of warning concerns invalid annotations. During the checking process, ESC/Java may discover that the property expressed by an annotation may not hold at a particular program point (for example, a method precondition may not hold at a call site of the method). The annotation assistant interprets such warnings as refuting incorrect guesses in the candidate annotation set, and removes these refuted annotations from the program text.

Since removing one annotation may cause subsequent annotations to become invalid, this check-and-refute cycle iterates until a fixpoint is reached. The resulting annotation set is a subset of the candidate set that is valid with respect to ESC/Java. Furthermore, this annotation set is the unique maximal such subset. For a proof of this property, and also a more efficient version of the basic algorithm presented here, we refer the interested reader to our companion paper [6].

After the check-and-refute loop terminates, the final step in the Houdini algorithm is to run ESC/Java one more time to identify potential run-time errors in the (now annotated) program. These warnings are then presented to the user, and are used as a starting point in identifying defects in the program.

3 Generating the candidate annotation set

The usefulness of the inferred annotations depends crucially on the initial candidate annotation set. Ideally, the candidate set should include all annotations that are likely to be useful in reasoning about the program's behavior. However, the candidate set should not be too large, because this would have a negative effect on the running time of the algorithm. Based on an inspection of a variety of hand-annotated programs and on our experience with ESC/Java and Houdini, we have developed the following heuristics for generating candidate annotations.

For any field `f` declared in the program, we guess the following candidate invariants for `f`:

Type of `f`	Candidate invariants for `f`
integral type	`//@ invariant f` *`cmp expr`*`;`
reference type	`//@ invariant f != null;`

array type	`//@ invariant f != null;`
	`//@ invariant \nonnullelements(f);`
	`//@ invariant (\forall int i; 0 <= i && i <` *`expr`*
	`==> f[i] != null);`
	`//@ invariant f.length` *`cmp expr`*`;`

boolean	`//@ invariant f == true;` `//@ invariant f == false;`

Many of these candidate invariants are intended to help verify the absence of index-out-of-bounds errors. For each integral field `f` we guess several inequalities relating `f` to other integral fields and constants. The comparison operator *`cmp`* ranges over the six operators `<`, `<=`, `==`, `!=`, `>=`, and `>`, and *`expr`* is either an integral field declared earlier in the same class or an *interesting constant*. Interesting constants include the numbers -1, 0, 1, and also constant dimensions used in array allocation expressions (*e.g.*, `new int[4]`). For each field `f` of an array type, we also guess a number of inequalities regarding `f.length`. Although some of these inequalities are more useful than others, we include all of them for completeness.

Some of these guessed invariants are mutually inconsistent. For example, if a class declares an integral field `f` we will guess several invariants, including:

```
//@ invariant f < 0;
```

```
//@ invariant f >= 0;
```

Such inconsistent guesses do not cause a problem. When checking a constructor for the class, ESC/Java will refute at least one of these invariants, since the constructed instance cannot simultaneously satisfy both invariants.

We also guess candidates invariants that help verify the absence of null dereference errors. For each field `f` of a reference type, we guess `f != null`. For each field `f` of an array type, in addition to guessing the invariant `f != null`, we also guess the invariant `\nonnullelements(f)`, which states that each entry in the array is not null, and we guess an invariant that all entries in `f` up to ***`expr`*** (a field or an interesting constant) are not null. This last property is useful in reasoning about the behavior of stack-like data structures implemented using arrays.

We generate candidate preconditions and postconditions in a similar manner for each routine declared in the program. Candidate preconditions may include inequalities relating two argument variables, or relating an argument variable to a field declared in the same class. Candidate postconditions may relate the result variable `\result` to either argument variables or fields. In addition, we generate the candidate postcondition

```
//@ ensures \fresh(\result);
```

which states that the result of a method is a newly-allocated object, and hence not an alias of any existing object.

As an aid in identifying dead code, we generate the candidate annotation

```
//@ requires false;
```

for every routine in the program. An unrefuted `requires false` annotation indicates that the corresponding routine is never called.

For correctness reasons, we require that all applicable candidate annotations hold in the program's initial state. Hence, for the program entry point

```
public static void main(String args[]) { ... }
```

we only guess the following precondition, which is ensured by the Java run-time system:

```
//@ requires \nonnullelements(args);
```

4 Dealing with libraries

So far, we have described Houdini as a system that infers annotations based on an analysis of the entire program. However, the program may be linked with a library (or with several libraries) that we cannot analyze, either because we do not have source code for the library, or because the size of the library makes the analysis impractical.

If the library in question already includes ESC/Java annotations that specify its interface, or if we are willing to write such an interface specification, then it is straightforward to adapt the Houdini algorithm to analyze and annotate the remainder of the program with respect to this specification. In many cases, however, the size and complexity of the library makes writing an interface specification quite tedious. Hence, we would like to be able to infer annotations for a program even in the absence of ESC/Java specifications for all of its libraries.

Therefore, we extend Houdini so that it can analyze a program with respect to *guessed* specifications for these libraries. There are two main strategies for guessing library specifications. The first is to make *pessimistic* assumptions, for example, that all pointers returned by library methods may be null. Since many of these pointers will never be null, such pessimistic specifications lead to a large number of false alarms, and we have not found this approach cost-effective for static debugging.

An alternative strategy is to make *optimistic* assumptions about the behavior of libraries, for example, that all pointers returned by library methods will be non-null. Since some of these pointers may sometimes be null, this assumption is unsound, and may cause Houdini to miss certain run-time errors. However, in library clients, Houdini will still detect many other run-time errors, and since the optimistic specifications lead to many fewer false alarms, this appears to be a more cost-effective strategy for guessing library specifications.

For libraries, we need to be careful not to guess contradictory annotations. In particular, if we guessed the contradictory postconditions

```
//@ ensures \result < 0;
//@ ensures \result >= 0;
```

for a library method, then ESC/Java would infer that the method never returns, and would not check code following a call to this method. Therefore, we only guess the following consistent postconditions for each library method:

Result type	Optimistic postconditions
integral type	`//@ ensures \result >= 0;`
reference type	`//@ ensures \result != null;`
array type	`//@ ensures \result != null;` `//@ ensures \nonnullelements(\result);`

We guess optimistic preconditions and invariants in a similar manner. The modified Houdini algorithm for dealing with libraries is as follows:

Input: An unannotated program `P`
A set of libraries `S` with specifications
A set of libraries `L` without specifications
Output: ESC/Java warnings for an annotated version of `P`
Algorithm:
generate and insert candidate annotations into `P`;
generate and insert optimistic annotations into `L`;
repeat
invoke ESC/Java to check `P` with respect to `L` and `S`;
remove any refuted candidate annotations from `P`;
remove any refuted optimistic annotations from `L`;
until quiescence;
invoke ESC/Java to identify possible defects in `P` with respect to `L` and `S`;

5 User Interface

We have described Houdini's output as consisting of an annotated version of the given program and the set of warnings that ESC/Java produces from this annotated program. To scrutinize Houdini's output, a user starts by going through the warnings. Unlike using ESC/Java, where a warning in one routine often points to a problem with that routine's implementation or specification, the warnings produced by Houdini are more often caused by some other part of the program. For example, suppose one of the warnings points out a possible null dereference of `t` in the following method:

```
char getFirstChar(String t) { return t.charAt(0); }
```

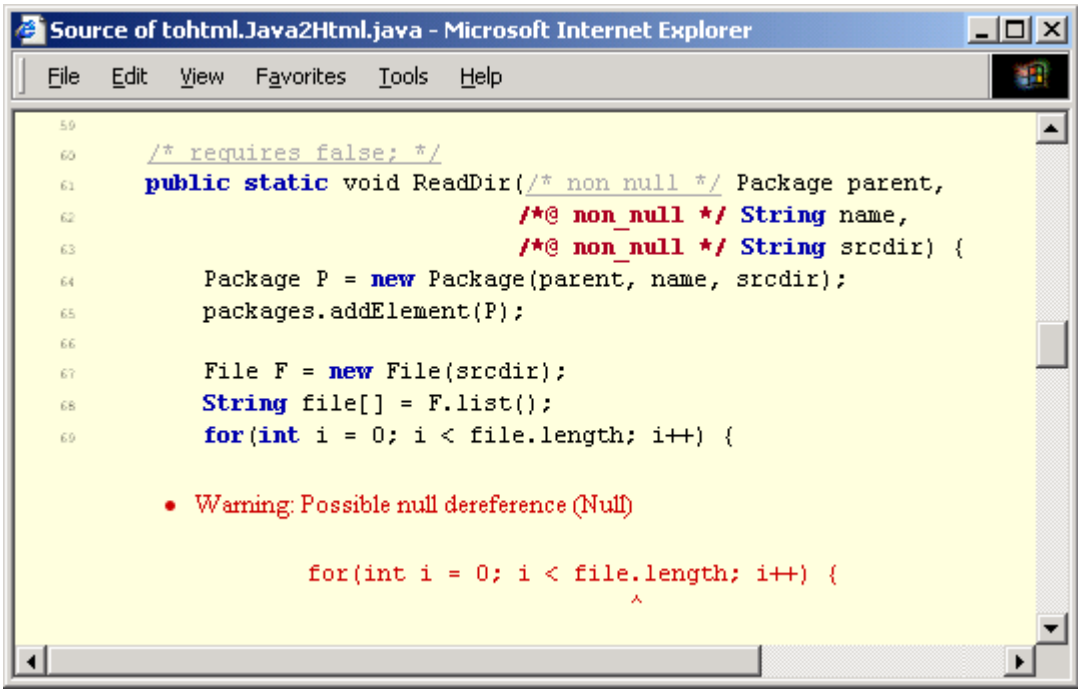

Figure 1: Screen shot showing the source view.

A user's first reaction might be: But I only intend `getFirstChar` to be called with non-null arguments! The ESC/Java user would then add the precondition `t != null`, which suppresses the spurious warning. But the Houdini user will instead ask: Why didn't Houdini infer this precondition?

In our experience with looking at Houdini output, we constantly asked questions such as these. To help answer these questions, we designed the tool to produce a log of the steps it takes and the results of those steps, including the ESC/Java output from each iteration. For each annotation removed, the log records the (first) ESC/Java warning that refuted the annotation.

We have built a simple but quite usable user interface in HTML, viewable and navigable through any web browser. The user interface starts by presenting a summary of the kinds of warnings that the final call to ESC/Java produces. This summary is followed by the actual list of warning messages with hyperlinks to the source view of the code at the location of the offending line.

In the source code view (shown in Figure 1), the user interface displays all of the candidate annotations guessed by Houdini. A refuted annotation is grayed out and hyperlinks to the source line where ESC/Java issued the warning that refuted the annotation. We also insert the warning messages into the code in the source view.

Thus, in the example above, the Houdini user would look at the preconditions that Houdini guessed initially for `getFirstChar`. These would include a grayed-out precondition `t != null`, and clicking on this refuted annotation would bring the user to a call site where, as far as ESC/Java could tell, the actual parameter of `t` may be null. This may lead the user to understand whether this is a real problem with the program or a spurious warning, or this may just be one of a number of similar steps required to get to the source of the problem.

6 Experience

We have applied Houdini to tens of thousands of lines of unannotated programs. Here we report on the four programs that have been subjected to many test runs of various versions of the tool. They are:

- Java2Html [10], a 500-line program that turns Java programs into color-coded HTML pages,
- WebSampler, a 2,000-line program that performs statistical samplings of trace files generated by the web crawler Mercator [9],

warning type	Java2Html		WebSampler		PachyClient		Cobalt	
	before	after	before	after	before	after	before	after
null	41	2	85	13	952	396	3114	417
array bounds	3	2	110	18	50	25	142	52
type cast	7	7	4	3	111		269	
other	5	0	52	16	123	23	519	87
all	56	11	251	50	1236	444	4044	556

Table 0: A breakdown of the warnings produced by ESC/Java before and after applying Houdini. [We happened not to have measured "after" for type casts for PachyClient and Cobalt. We will take a new measurement for the final version of this paper.]

- PachyClient, the 11,000-line graphical user interface of the web-based email program Pachyderm [13], and
- "Cobalt", a proprietary 36,000-line program.

These programs had been tested and used, in some cases extensively, before Houdini was applied to them.

Table 0 shows a breakdown of the various warnings produced by ESC/Java before and after running Houdini.

Java2Html was something of a lucky draw: of the 4 non-cast warnings reported, all 4 indicated real errors in the source code. For example, the program will crash if a line in the input exceeds 1024 characters or if the system call `File.list()` returns null (Figure 1).

Of the warnings produced on WebSampler, we have found 3 real errors. One of these was part of a class that was borrowed from the web crawler Mercator: the method

```
int InputStream.read(byte[] b);
```

is supposed to read characters into the given array `b`, returning the number of characters read or returning -1 if the input stream is at end-of-file. Consequently, if the length of `b` is 0, the method can only return 0 or -1, but which of these should be returned if the stream is at end-of-file? The `InputStream` documentation says -1, but the Mercator subclass of `InputStream` returned 0.

Houdini also issued several warnings pointing out places where WebSampler assumed its input to have a particular format. For example, in some situations,

WebSampler assumes the next line of input to be at least 4 characters long and to consist only of characters that can be interpreted as a decimal number. The warnings pointing out these infelicities were considered spurious, since WebSampler is only intended to work on well-formed input.

We have only scrutinized about a dozen of the PachyClient warnings. Nevertheless, two of these pointed out infelicities that compelled the author of the code to make code changes. Technically, one can argue that these infelicities were not errors, but they did make the code overly brittle. The author changed the code to guard against possible future failures.

One hundred of the 556 warnings of Cobalt were scrutinized and revealed 3 real errors in the code. This is the only one of the four programs for which we used optimistic library assumptions.

For the first three programs, we did not measure the user time required to scrutinize the warnings, but in the case of Cobalt, this time was measured to be 9 hours. Toward the end of this time, the scrutiny proceeded at a higher pace than in the beginning, partly because of getting more familiar with the tool's output and partly because of repeated spurious warnings.

Despite the precision of ESC/Java, Houdini still produces many false alarms. One cause of false alarms is that Houdini may fail to guess the right annotations for the given program. In particular, Houdini does not guess disjunctions, such as in the following postcondition of the method `put` from Figure 0:

```
//@ ensures j == n || \result != null
```

Another cause of false alarms is the incompleteness of ESC/Java, comprising the incompleteness of the underlying theorem prover, the incompleteness of ESC/Java's axiomatization of Java's operators (like the full semantics of bitwise-and), and the incompleteness of ESC/Java's light-weight annotation language. An ESC/Java user would know to insert `nowarn`, `assume`, and `axiom` annotations to make up for any such incompleteness (see the ESC/Java user's manual [11]), but Houdini does not employ such annotations.

After scrutinizing a serious warning, a user manually fixes the error in the program. If instead the user determines the warning to be spurious, a prudent thing to do is to convey to Houdini the missing pieces of information. For example, if the cause of a spurious warning is that Houdini didn't guess some annotation, possibly with a disjunction like the one shown above, then the user can manually insert the missing annotation into the program.

Thus, Houdini users are likely to manually insert ESC/Java annotations into the program, just like ESC/Java users would. However, Houdini users insert many

fewer annotations. We predict that leaving such a small number of annotations in the code will be acceptable to most programmers.

These manually-inserted annotations give future runs of Houdini more information about the program, which may cause Houdini to refute fewer candidate annotations and to generate fewer spurious warnings. Thus, users are likely to rerun Houdini after adding some manual annotations. This iterative process can be quite cost effective, because even small manual interventions can prevent what otherwise might have resulted in a cascade of spurious refutations.

A consequence of this iterative process is that the Houdini running time must not be too long. Since Houdini, like all other static debuggers, competes in practice with software testing, it seems reasonable that a dozen iterations may be done over the course of a couple of weeks. This means that each run of Houdini must be fast enough to complete overnight, say, finishing within 16 hours.

The version of Houdini reported on in this paper does not meet the overnight challenge: the running time on the 36,000-line Cobalt program was 62 hours. We remain optimistic, however, for several reasons. First, some preliminary experiments with algorithmic improvements seem promising. Second, measurements of the work performed during each of Houdini's iterations suggest the operations of Houdini to be parallelizable. And third, we have built a prototype of a dynamic refuter, which creates an instrument version of the program that, when run, records which candidate annotations are violated (i.e., refuted) during the execution. This significantly reduces the number of candidate annotations that are left to be refuted by ESC/Java.

Since a run of Houdini may take many hours, an important aspect of Houdini's usability is that it is restartable. The system periodically writes a snapshot of its current state to disk, so that if that particular run is abruptly terminated (for example, by a power failure), it can later be restarted at the most recent snapshot.

Finally, we give some measurements that give a preliminary idea of the effectiveness of the various candidate annotations described in section 3. Table 1 shows a breakdown of the kinds of annotations inserted into the Cobalt program and how many of these are acutally valid annotations. Is this table, the count of valid annotations includes many annotations that hold hold in all Java programs (for example, $-1 <= \texttt{a.length}$) and annotations that are subsumed by other annotations (for example, `x != -1` is subsumed by `x > 0`); it remains for future work to avoid such valid but redundant annotations.

Type of annotation	Preconditions		Postconditions		Invariants		Total		
	guessed	valid	guessed	valid	guessed	valid	guessed	valid	%
`==`	2130	380	985	175	435	59	3550	614	17
`!=`	2130	745	985	346	435	165	3550	1256	35
`<`	2130	560	985	268	435	105	3550	933	26
`>`	2130	669	985	354	435	151	3550	821	23
`<=`	2130	668	985	312	435	157	3550	1174	33
`>=`	2130	537	985	202	435	82	3550	1137	23
`f != null`	509	469	229	180	983	708	1721	1357	79
`nonnullelems`	54	44	21	13	36	23	111	80	72
`\forall...`	841	228	260	95	125	74	1226	397	32
`f`	47	13	51	12	39	3	137	34	25
`!f`	47	17	51	13	39	4	137	28	20
`\fresh`	0	0	229	69	0	0	229	69	30
`false`	780	130	0	0	0	0	780	130	17
exact type	37	7	11	4	14	8	62	19	31
Total	15095	4467	6762	2043	3850	1539	25707	8049	31

Table 1: Numbers of candidate annotations generated by the heuristics of Section 3, and number and percentage of these annotations that are valid.

7 Other Annotation Assistants

We also experimented with two other ESC/Java annotation assistants that were less successful than Houdini.

The first annotation assistant started with an unannotated program and iteratively added annotations, rather than like Houdini guessing initially a candidate set of annotations and iteratively refuting these. To make this first annotation assistant work, ESC/Java was modified to output a suggestion with each warning, whenever possible. For example, if a warning points to the dereference of a formal parameter `p` that is not changed by the routine body being checked, the suggestion is to add a precondition `p != null`. If ESC/Java makes no suggestion, Houdini uses the default suggestion "put a `nowarn` on the offending line". Each suggestion has the property that by following the suggestion and rerunning ESC/Java, the warning will be suppressed (but other, new warnings may be generated). This requirement ensures the termination of the annotation assistant, which iteratively ran ESC/Java and followed the suggestions until ESC/Java produced no more suggestions.

Although many of the suggestions were good, this annotation assistant had two severe limitations. First, it didn't produce enough suggestions. Not only did new heuristics become increasingly more complicated, but because of the requirement of only suggesting measures that are sure to suppress the warning, the heuristics for when to make a suggestion frequently had side conditions that were not met (such as "...and the parameter is not assigned to by the body"), and thus the default suggestion would be used. Second, the annotation assistant produced incorrect suggestions. The cascading effects of incorrect annotations based on such suggestions generated too much clutter and caused many more spurious warnings to be generated in the end.

The second annotation assistant used a whole-program set-based analysis [4, 8] to identify which variables and fields could never be null and inserted corresponding annotations into the program. These annotations were useful in verifying many dereference operations. However, the inferred annotations did not include numeric inequalities (which are necessary for verifying the absence of array bounds errors), and did not include properties such as

```
//@ invariant (\forall int i; 0 <= i && i < expr
                            ==> f[i] != null);
```

(which are sometimes necessary for verifying dereference operations).

An advantage of the Houdini algorithm is its generality: it is not closely dependent on the underlying checker, and can be used to infer annotations for a variety of modular static checkers. We have successfully ported Houdini to a second checker, the race condition checker for Java [5]. This checker extends Java's type system with additional checks that verify the absence of race conditions. The checker relies on additional type annotations to describe aspects of the locking discipline, for example, the protecting lock of each field. Adapting the Houdini algorithm to guess many such type annotations was straightforward, and, as with ESC/Java, we have found the annotations inferred by the system to be useful.

8 Related work

Predicate abstraction is a technique for analyzing an infinite state system, given a set of predicates over the state space of the system [7]. This technique uses ideas from model checking to find a boolean combination of the predicates that holds in all reachable states, and that holds in a minimum of unreachable states.

The Houdini algorithm can be viewed as a variant of predicate abstraction in which each candidate annotation corresponds to a predicate. Interestingly, the Houdini algorithm does not consider arbitrary boolean formulae over these predicates; it only considers conjunctions of predicates. This restriction means that Houdini cannot infer disjunctions or implications of predicates. For example, given the two predicates

```
j == n
\result != null
```

Houdini could not infer the property

```
j == n || \result != null
```

This restriction reduces the maximum number of iterations of the algorithm from exponential to linear in the number of predicates; however, there is a corresponding, though small, increase in the number of false alarms produced by the system.

Abstract interpretation [0] is a standard framework for developing and describing program analyses. We can view the Houdini algorithm as an abstract interpretation, where the abstract state space is the power set lattice over the candidate annotations and the checker is used to compute the abstract transition relation. As usual, the choice of the abstract state space controls the conservative approximations performed by the analysis. In our approach, it is easy to tune

these approximations by choosing the set of candidate annotations appropriately, provided that this set remains finite and that the annotations are understood by the checker.

An interesting aspect of our approach is that the checker can use arbitrary techniques (for example, weakest preconditions in the case of ESC) for performing procedure-local analysis. If these local analysis techniques allow the checker to reason about sets of intermediate states that cannot be precisely characterized using the abstract state space, then an annotation assistant may yield more precise results than a conventional abstract interpretation that exclusively uses abstract states to represent sets of concrete states.

PREfix is a static programming tool that warns about possible errors in C and C++ code [14]. There are no annotations involved in using PREfix, which is mostly an advantage. We find the presence of annotations in Houdini's output, including the refuted annotations, helpful when scrutinizing the tool's output. Annotations also provide a general and convenient way for users to supply the tool with missing, uninferred facts. The technology underlying PREfix is different from the precise semantics engine and automatic theorem prover that underly ESC/Java, but perhaps the differences are bigger than they need to be.

Daikon is a system that uses an empirical approach to find probable invariants [2]. These invariants are found by creating an instrumented version of the program that records a trace of intermediate program states; running the instrumented program on a test suite; and then analyzing the generated traces off-line to determine properties that hold throughout all runs. Given a suitably complete test suite, the inferred properties are likely to be true program invariants.

9 Conclusions

This paper describes a technique for building an annotation assistant for a modular static checker. The annotation assistant reuses the checker as a subroutine; it works by guessing a large number of annotations and using the checker to verify or refute each guessed annotation.

We have used to this technique to develop an annotation assistant, called Houdini, for ESC/Java. Our experience with Houdini indicates that this approach shows promise. Houdini is capable of inferring a large number of useful annotations, which significantly reduce the number of false alarms produced by ESC/Java (as compared with checking an unannotated program). As expected, these inferred annotations reduce the programmer time required to annotate a program, which

is important since this annotation overhead appears to be the major barrier to the cost-effective use of advanced modular checkers like ESC/Java.

In our experience, a natural strategy for using Houdini is to maintain a code base with a few manually-inserted annotations that Houdini cannot infer, and to rely on Houdini to infer additional annotations whenever the code needs to be checked. Thus, we expect that Houdini will be considered by users as a static checker in itself, and not just as an annotation assistant for the underlying checker ESC/Java.

A number of issues remain for future work, including refining the heuristics for the generation of the candidate annotations, improving the performance of Houdini (significant progress has already been made in this direction), and enhancing the user interface. However, the system developed to date has already proven capable of catching defects in a variety of large, real-world programs.

Acknowledgements

We are grateful to our colleagues who contributed in various ways to the Houdini project: Yuan Yu suggested and helped us develop the dynamic refuter. He also helped scrutinize the Houdini output on WebSampler. Steve Freund ported Houdini to `rccjava`. Raymie Stata was part of creating the annotation assistant that used set-based analysis. Jim Saxe helped think about how to get the underlying theorem prover to work well with Houdini's demands. Lawrence Markosian and Dr. Maggie Johnson of Reasoning, Inc., Mountain View, CA scrutinized the Houdini output on Cobalt.

References

[0] P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Conference Record of the Fourth Annual ACM Symposium on Principles of Programming Languages*, pages 238–252, January 1977.

[1] David L. Detlefs, K. Rustan M. Leino, Greg Nelson, and James B. Saxe. Extended static checking. Research Report 159, Compaq Systems Research Center, December 1998. Available from `research.compaq.com/SRC/publications/src-rr.html`.

[2] Michael D. Ernst, Adam Czeisler, William G. Griswold, and David Notkin. Quickly detecting relevant program invariants. In *Proceedings of the 22nd International Conference on Software Engineering (ICSE 2000), Limerick, Ireland*, June 2000.

[3] Extended Static Checking home page, Compaq Systems Research Center. On the Web at `research.compaq.com/SRC/esc/Esc.html`.

[4] Cormac Flanagan. *Effective Static Debugging via Componential Set-Based Analysis*. PhD thesis, Rice University, Houston, Texas, May 1997.

[5] Cormac Flanagan and Steven N. Freund. Type-based race detection for Java. In *Proceedings of the 2000 ACM SIGPLAN conference on Programming Design and Implementation (PLDI)*, pages 219–232, 2000.

[6] Cormac Flanagan, Rajeev Joshi, and K. Rustan M. Leino. Annotation inference for modular checkers. *Information Processing Letters*, 2000. To appear.

[7] S. Graf and H. Saidi. Construction of abstract state graphs with PVS. In O. Grumberg, editor, *CAV 97: Computer Aided Verification*, Lecture Notes in Computer Science 1254, pages 72–83. Springer-Verlag, 1997.

[8] Nevin Heintze. Set-based analysis of ML programs. In *Proceedings of the ACM Conference on Lisp and Functional Programming*, pages 306–317, 1994.

[9] Allan Heydon and Marc A. Najork. Mercator: A scalable, extensible web crawler. *World Wide Web*, 2(4):219–229, December 1999.

[10] Java2html, Compaq Systems Research Center. On the Web at `research.compaq.com/SRC/software/`.

[11] K. Rustan M. Leino, Greg Nelson, and James B. Saxe. ESC/Java user's manual. Technical Note 2000-002, Compaq Systems Research Center, September 2000. To appear.

[12] K. Rustan M. Leino, James B. Saxe, and Raymie Stata. Checking Java programs via guarded commands. In Bart Jacobs, Gary T. Leavens, Peter Müller, and Arnd Poetzsch-Heffter, editors, *Formal Techniques for Java Programs*, Technical Report 251. Fernuniversität Hagen, May 1999. Also available as Technical Note 1999-002, Compaq Systems Research Center, from `research.compaq.com/SRC/publications/src-tn.html`.

[13] The Pachyderm email system, Compaq Systems Research Center. On the Web at `research.compaq.com/SRC/pachyderm/`, 1997.

[14] Prefix. Intrinsa, Mountain View, CA. 1999.