Project N°: FP6-015905

Project Acronym: MOBIUS

Project Title: Mobility, Ubiquity and Security

Instrument: Integrated Project

Priority 2: Information Society Technologies

Future and Emerging Technologies

Deliverable 3.5

Intermediate Report on Program Verification Environment and Annotation Generation

Due date of deliverable: 2008-02-29 (T0+30)

Actual submission date: 2008-03-26

Start date of the project: 1 September 2005 Duration: 48 months

Organisation name of lead contractor for this deliverable: UCD

Revision: Creation

Project co-funded by the European Commission in the Sixth Framework Programme (2002-2006)		
Dissemination level		
PU	Public	✓
PP	Restricted to other programme participants (including Commission Services)	
RE	Restricted to a group specified by the consortium (including Commission Services)	
CO	Confidential, only for members of the consortium (including Commission Services)	

Contributions

Site	Contributed to chapter
ETH	2
INRIA	2
RU	2
UCD	1 2 3 4 5 6
WU	2

This document has been written by Dermot Cochran (UCD), Werner Dietl (ETH), Radu Grigore (UCD), Marieke Huisman (INRIA), Joe Kiniry (UCD), Erik Poll (RU), and Aleksy Schubert (WU).

Executive Summary:

Intermediate Report on Program Verification Environment and Annotation Generation

This document describes the results of tasks 3.6 and 3.7 of the MOBIUS project. The focus of these tasks is the design and development of an integrated program verification environment with annotation generation facilities.

Task 3.6 develops a program verification environment for the interactive verification of Java programs that integrates the techniques developed in WP3. Such a program environment is needed to tackle the realistic case studies in WP5, and, more generally, is crucial if we are to seriously evaluate the MOBIUS PCC infrastructure.

The MOBIUS verification environment supports the development of specifications, programs and proofs of their security properties. It will support both source code and bytecode level verification, and be able to produce the output necessary for the generation of PCC certificates.

The result of this task in the first two years of the project has been three releases (called the alpha, beta, and gamma releases) of the platform, as well as dozens of independent releases of its various subsystems. These initial releases support reasoning about source code and bytecode and focus on single-threaded Java only. It also generates proof obligations for a single proof generation tool. The environment also interfaces to several automated theorem provers and interactive proof assistants.

Task 3.7 focus on generating JML annotations for existing source code and bytecode. Automata or state machines are a convenient and accessible specification formalism for many properties. Techniques are being investigated to generate JML annotations from a simple automaton that specifies security properties. If a program that is monitored by the automaton does not go “wrong,” then the generated annotations will not be violated during any execution of the program. The algorithm and its correctness proof are formalised with a theorem prover. The generated annotations are then propagated, to assist static verification. A simple version of the propagation algorithm has been implemented.

Contents

Chapter 1

Overview of the Program Verification Environment

Many large software engineering projects have functional specifications that are thousands of pages long. These are usually written in a natural language, such as English, and with semi-formal notation such as diagrams and tables. There are often subtle ambiguities or internal contradictions within the software requirements, which are often not discovered until very late in the software development life cycle, simply because the project becomes too big for any one developer or analyst to fully understand and remember all the relevant details at the same time.

Popular high-level ”modeling” languages in current use (e.g., UML) are not amenable to formal analysis because of a lack of semantics and interoperable tool support. Even so, high-level representations are compelling and useful in modern software engineering practice and software engineers regularly struggle with using these inadequate languages, even for mission critical applications.

The Mobius PVE is a customized version of the Eclipse platform that integrates a set of tools that support software development using state-of-the-art formal methods and software engineering techniques. The Mobius PVE supports the specification and verification of JML-annotated Java programs at the source code level as well as at the bytecode level using logic- and type-based techniques and the Proof-Carrying Code paradigm. The properties of particular interest include security (e.g., data is secret and never leaked) and resource guarantee properties (e.g., this method will never use more than this much memory or that much CPU) as well as total functional correctness. The Mobius PVE is thus, in a sense, a prototype realization of the U.K. Grand Challenge 6 Verifying Compiler, specialized to the domain of Java programs [HM05].

The Mobius PVE integrates theoretical and technical best-practices, leveraging the hard work of many research groups within, as well as outside of, the MOBIUS grant consortium. Many Java software engineers use the Eclipse development environment, together with plugins such as FindBugs and PMD [Ecl]. The Mobius PVE is bundled with Eclipse and these well-known third-party plugins so that Java developers are presented with a pre-configured development environment without the need to manually install each plugin.

For example, a Java developer might use the Mobius PVE to derive a JML specification from a high-level functional specification. The JML subsystem of MOBIUS might then be used to typecheck the JML specification, and, in a complementary fashion, the ESC/Java2 used look for inconsistencies in the specification. A variety of other pre-configured subsystems like CheckStyle, FindBugs, PMD, and others are executed asynchronously while the developer writes, modifies, and evolves their software architecture. Feedback about the completeness and quality of the system's documentation, formal specification, and program code are always at the developer's fingertips. Once the Java code has been verified against the JML specification, both statically using ESC/Java2 and dynamically using hand-written and automatically generated unit tests using JML's runtime assertion checking, then the developer has very high confidence that the requirements have been implemented correctly. At this stage, one performs automatic and interactive verification using the MOBIUS DirectVCGen and Prover Editor, fully verifying functional and non-functional properties.

Chapter 2

Subsystems of the MOBIUS Program Verification Environment

The Mobius PVE is built on top of the Eclipse platform. The primary units of software within Eclipse are known as features, which are in turn composed of one or more plugins. We generally call these two kinds of units subsystems in Mobius PVE documentation.

The following sections summarize all of these subsystems at a high level. More information on each subsystem is available via the MOBIUS Trac server, within publications related to specific tools cited appropriately, and from the original development team, in the case of externally developed tools.

The Mobius PVE is a customized version of Eclipse that contains over forty subsystems dedicated to reliable, dependable software development. Over half of these subsystems were implemented by MOBIUS partners.

Development on some of these subsystems (e.g., the static checker ESC/Java2) began before the project started, but their development continued within the project as their capabilities directly complemented the goals of the Mobius PVE. Other subsystems (e.g., the Umbra bytecode subsystem) were developed entirely within and for the MOBIUS project and the Mobius PVE.

This chapter contains a brief summary of each subsystem in the Mobius PVE. There are two forms of subsystems: those (co-)developed by MOBIUS partners (known as "MOBIUS subsystems"), and those developed outside of MOBIUS (known as "non-MOBIUS subsystems"); section 2.1 discusses the former, and section 2.2 the latter. Subsystems developed within MOBIUS are often discussed in more detail unless a non-MOBIUS subsystem has some special interaction with a MOBIUS one.

Some subsystems are exposed to the verification application developer via a programmatic API. Others are full-fledged Eclipse features and/or plugins. Many subsystems, even those that are plugins, also have rich command-line interfaces and are executed as normal processes.

In general, all MOBIUS subsystems written in Java (the majority of them) were developed with various version of the Mobius PVE itself. Moreover, some of them have full formal specifications and have been partially or fully validated (e.g., extensively unit tested) and/or verified (e.g.,, ESC/Java2 reports no errors) with the Mobius PVE. When a subsystem has one of these properties, it is highlighted in the text.

As discussed in section 4.2, the source code, documentation, version control logs, and bug and feature tracking database of all subsystems is available online via our Trac-based

CDE located at http://mobius.ucd.ie/. All subsystems developed within MOBIUS are Open Source and released under the GPLv2 or a BSD-style license.

2.1 Subsystems Developed within MOBIUS

This section summarizes in some detail all of the subsystems that ship as part of the current "gamma" (3rd) release of the Mobius PVE. Most of these subsystems are still under development and are evolving to match the needs of Mobius PVE users and MOBIUS partners and their research and industry collaborators.

Information about these components (their description, status, location in the source repository, etc.) is also summarized in the MOBIUS Trac wiki via the Components wiki page.

Many MOBIUS subsystems are developed by MOBIUS partners working in collaboration with non-MOBIUS researchers throughout North America and Europe.

Many subsystems have interdependencies, many of which grew organically over several years of development, as one plugin will depend upon the availability of another. In fact, teasing out, documenting, and formalizing these dependencies in plugin descriptors has been a main organizational effort in Mobius PVE development.

Broadly speaking, subsystems are organized into three dimensions. The first dimension is organized by language:

- Java
- Java Modeling Language (JML)
- Java Bytecode
- Bytecode Modeling Language (BML)
- Coq

The second, by action:

- Static Analyses
- Abstract Syntax Trees (AST)

The third, by interface:

- Application Programmer Interfaces
- Backends
- Interactive Plugins

Some subsystems lie along multiple facets of these dimensions. For example, many subsytems are interactive, but also have APIs. Others, tie together every facet. For example, in the case of ESC/Java2, it reasons about Java source code, bytecode, and JML, contains several different ASTs, provides several APIs, uses many backends, and finally, is also an interactive plugin.

2.1.1 Java Modeling Language Subsystems

Several different subsystems are included that lex, parse, represent, process, and reason about the Java Modeling Language (JML), as JML is core to program specification and verification within MOBIUS [LPC+05]. The inclusion of several subsystems with overlapping functionality is, in nearly all instances, more a historical artifact than a conscious developer choice.

The JML4 Plugin

ESC/Java2 is currently being maintained to support new verification functionality, but its compiler frontend has yet to support Java 5 source code.

JML4 is a customization of the Eclipse subsystem that support Java program development, the Java Developers Toolkit (org.eclipse.jdt). JML4 provides runtime assertion checking and has extension points which are used by other subsystems, like a customized bridge that connects ESC/Java2 to JML4 [CJK07].

JML4 currently enhances Eclipse versions 3.3 to 3.4 with:

- scanning and parsing of nullity modifiers,
- enforcement of the JML non-null type system, both statically and at runtime, and
- the ability to read and make use of the extensive JML API library specifications.

This subset of features was chosen so as to exercise some of the basic capabilities that any JML extension to Eclipse would need to support.

Also, the functionality added by the chosen subset of features is useful in its own right, somewhat independent of other JML features; i.e., the capabilities form a natural extension to the existing embryonic Eclipse support for nullity analysis.

JML4 development continues, led by Concordia University, and mainly supported by an NSF Infrastructure grant in the U.S.A. A JML4 “bootcamp” took place in Florida in the first quarter of 2008. At that event developers from all over the world, including MOBIUS-funded researchers, were introduced to the design and development of the JML4 plugin.

It is expected that the first public release of JML4 will take place in 2008 and will support all of JML level 0 and level 1. These technical issues (the level of JML support and rich Eclipse integration) as well as social aspects (the size and activity level of the JML community working on/with JML4) means that JML4 is a strong candidate as a frontend for future MOBIUS tools.

The JML2 Eclipse Plugin

This plugin provides a simple integration of the common JML tools into Eclipse 3.3. It is the main interactive plugin used by MOBIUS developers within the Mobius PVE at this point in time.

The following features are provided:

- simple installation of the core JML tools, including
- executing the JML checker,
- executing the JML runtime assertion checker (RAC),

- executing a program that was compiled with RAC, and
- configuration dialogs for JML checker and RAC.

The current version of the plugin is available from an ETHZ website and will be integrated into the main MOBIUS releases website and Trac in the coming year.

2.1.2 Extended Static Checking Subsystems

ESC/Java2 has been under development for over a decade. The original ESC/Java was developed by a team at Digital Equipment Corporation's Systems Research Center in the mid-1990s. After that team disbanded and DEC-SRC was shut down (after Compaq bought Digital and then Hewlett-Packard bought Compaq), Hewlett-Packard released the source code of ESC/Java for research and educational purposes.

At that time, a team at Radboud University Nijmegen, led by Joe Kiniry, restarted research and development on Extended Static Checking for Java, and have since then released over fifteen versions of ESC/Java2. ESC/Java2 shares only a few subsystems with the original ESC/Java.

This section summarizes the core subsystems of the current release of ESC/Java2 (2.0b4) included with the Mobius PVE that remain from the original ESC/Java. All other extensions to ESC/Java2 that have taken place over the past five years are described in other sections, as they have all been incorporated into the Mobius PVE itself. These extensions include the new VC ASTs, the Automated Prover API, the the new Fx7 and Z3 backends, universes typechecking, and support for higher-order extended static checking, which is not discussed in this deliverable.

ESCJava2

The Extended Static Checker for Java version 2 (ESC/Java2) is a programming tool that attempts to find common run-time errors in JML-annotated Java programs by static analysis of the program code and its formal annotations. Users control the amount and kinds of checking that ESC/Java2 performs by annotating their programs with specially formatted comments called pragmas.

ESC/Java2 is a large-scale subsystem within the Mobius PVE. Its architecture relies upon several of the other subsystems described in this document including the original DEC-SRC Java Frontend and the Simplify theorem prover, both of which are described in this section, and several other new MOBIUS subsystems, most of which are described in later sections of this chapter.

In a nutshell, ESC/Java2 translates JML-annotated Java source code into verification conditions that are passed to an automated theorem prover. Responses from the prover are translated back into informational messages useful to a normal programmer.

To implement these features, ESC/Java2 uses a combination of Javafe to lex and parse Java sourcecode, an internal bytecode subsystem to parse Java bytecode, and its own JML lexer and parser, all of which produce a JML-annotated Java AST. This representation is then translated into a semantically equivalent guarded command, dynamic single assignment-based representation, which is then used to generate a verification condition with either a weakest precondition or a strongest postcondition VC generator. This VC is represented in one of two forms, the original DEC-SRC VC AST, or the new MOBIUS sorted VC AST. This VC is passed to an automated theorem prover via one of two mechanisms, either the original DEC-SRC Simplify-specific interface, or the new MOBIUS automated prover backend. Likewise, responses from the automated prover flow back through these same interfaces.

In summary, ESC/Java2 uses more MOBIUS subsystems than any other component, except the Mobius PVE itself. It represents the core system used to perform research on, and evolve existing, generic verification functionality, as well as exercise and extend various other Mobius PVE subsystems.

Javafe: The Original DEC-SRC Java Frontend

Javafe is the original DEC-SRC Java Frontend. It lexes and parses Java 1.5 bytecode and Java 1.4 source code. It is used by several other MOBIUS subsystems including ESC/Java2, RCC, Houdini, and the universes typechecker. It was originally developed by DEC-SRC but has been improved and extended over the past five years in a variety of ways.

Some MOBIUS subsystems under development use Javafe as their frontend, as it is a reliable, well-understood subsystem. Its main drawbacks are (1) that it only supports pre-Java 1.5 source code, thus does not support parameterized classes, enumerations, foreach loops, etc., and (2) its license is non-standard and not a recognized Free/Open Source license.

Simplify

The Simplify automated theorem prover was written by Greg Nelson and others at DEC-SRC around fifteen years ago. It is implemented in Modula-III. Until very recently it was still considered the premier automated theorem prover for software verification tools. It supports reasoning about a decidable fragment of first-order logic that includes a theory of arrays, uninterpreted function symbols, and linear arithmetic.

The only recent development on Simplify that has taken place is the MOBIUS port of it to Mac OS X, building new releases with new Modula-III compilers, as well as some simple bug fixes. There are no plans on doing any development on Simplify, as the world of first-order theorem provers has changed quite a bit over the past decade or so. Instead, we are working on integrating ESC/Java2 with new-generation SMT-LIB provers, as discussed later in this chapter.

Simplify has been wrapped in the new MOBIUS prover backend, discussed in section 2.1.8, thus is accessible in a simple, transparent fashion like many other new generation automated provers.

2.1.3 Abstract Syntax Tree Subsystems

Six AST subsystems exist within the Mobius PVE covering Java, JML, Java bytecode, guarded-commands, BoogiePL, and verification conditions. These subsystems are not described here in detail as it is beyond the scope of this report.

2.1.4 Java Bytecode Subsystems

One of the end goals of the MOBIUS project is the production of PCC certificates for Java. To accomplish this goal, reading, writing, manipulating, specifying, and reasoning about Java bytecode is mandatory. This section discusses the initial subsystems developed for the Mobius PVE to support these objectives.

Bico

Bicolano [Pic06] is a Coq library that formalizes the memory model of Java, as well as the bytecode instructions and their semantics. To translate Java classfiles into this Coq formalization, a tool called Bico has been developed within MOBIUS. This is joint work of INRIA and Warsaw University [mC07a].

For each classfile, Bico generates a type, signature, and implementation formalized in Coq. This generation is modular. Bico also generates summary files to simplify the use of the definitions.

The main limitation of Bico is due to the size of the Java system library. The classes contained in this library are usually not generated, because the resolving of the dependencies would make Bico generate the Coq formalisation for all the system libraries (more than 200 classes). This is not a problem by itself, but loading the generated files into Coq uses a large amount of time and memory. Therefore, Bico is most useful when applied on user-written class files, where this problem tends not to arise.

The Umbra Bytecode Editor

Umbra is an editor for bytecode files which supports editing of BML specifications. This editor can manipulate bytecode files in a textual form. It can also add, change, and remove instruction mnemonics so that the actual functionality encoded in the class file methods changes.

In addition to the modification of bytecode programs, Umbra supports viewing, adding, deleting and modifying specifications written in the Bytecode Modeling Language (BML), a counterpart of JML at the bytecode level.

A separate part of the Umbra editor is a library called BMLlib. BMLlib parses fragments of BML specifications, prints them out in the aforementioned textual format, and reads and writes BML specifications in class files. More details are presented in [SCB^{+}08].

2.1.5 Application Programmer Interfaces Subsystems

Two subsystems developed for MOBIUS are exclusively used via APIs at this time. These subsystems focus on two critical activities in program verification: interacting with automated provers and making assertions and logging data about a system as it executes.

The Automated Prover API

The MOBIUS Automated Prover API provides a generic programming interface to automated theorem provers. The initial version of this API was designed by MOBIUS partners working in tandem with leaders in the SMT-LIB community. The intention is to have some future variant of this API be the standard API for all automated theorem provers, at least within the SMT-LIB community.

The API was developed by performing an analysis of the core operations provides by all major automated provers: prover initialization and shutdown, hinting to a prover about resource utilization bounds (primarily, time and space), declaration of new theories (via axiomatic definitions), making and retracting assumptions (as part of, e.g., a background predicate of a theory), and checking the validity of formulae.

The initial version of this API was designed with a formal specification written in JML, and example implementations were provided in Java, C, and OCaml. Communication with theorem provers is supported via pipes/sockets, native methods, and an XML-RPC-based messaging interface. This API is now used by ESC/Java2 and the CVC3 prover, and interfaces with the Z3 and Simplify theorem provers as well [BT07].

The Logging API

The MOBIUS Logging subsystem (or MobLog for short) is an API used by programmers to make assertions, and emit log messages, about a software system. At its core, it combines Java and JML's assert statements with functionality akin to that found in the Java logging framework (via the java.util.logging package) and the Apache Foundation's log4j logging service.

The original design for the Logging subsystem predates not only the MOBIUS project, but the existence of Java's logging package and log4j [Kin98]. But, even though this framework is about a decade old, it still has several features not found in log4j or similar frameworks including first-class logging contexts, integrated statistics collection, logging categories, and more.

Within this project, it has been updated and extended for incorporation into the Mobius PVE for several reasons.

First, as it was designed and written with formal specifications (initially written in Jass, now in JML), it represents a case study in verification and thus will be used to test various other MOBIUS subsystems. All classes are annotated with lightweight JML assertions, and a high-level BON [WN95] specification is also included with the release. Moreover, the entire software system has been updated and analyzed within the Mobius PVE itself, passing all static checkers, including ESC/Java2.

Additionally, the MobLog framework supports complex concurrent and distributed systems' development, like that found in the Mobius PVE. Tracing the execution of, and understanding the behavior of, multiple concurrent threads requires a thread-aware logging framework.

Finally, its availability under MOBIUS control provides an opportunity to unify the several disparate logging subsystems found within other MOBIUS components. Thus, it is our aim that all Mobius PVE logging, tracing, and statistics-gathering is performed via this single unified subsystem.

2.1.6 Interactive Plugins Subsystems

In addition, three subsystems developed by MOBIUS partners are exclusively used as interactive Eclipse plugins. These subsystems focus on code and specification folding, extended static checking, and interacting with higher-order theorem provers.

The JML Folding Editor

Folding editors allow one to manually or automatically show or hide various structured pieces of a document—for example, hiding the contents of a section while only keeping the section title, which is similar to hiding the implementation of a method while keeping its signature. The process of hiding and showing a usually larger or complex construct behind a concise visual representation is called “folding”.

Basic folding editors permit one to manually fold and unfold constructs like source code documentation and implementation structures like method and loop bodies. Advanced folding editors provide additional features like saving and restoring the state of folding when a file is closed and then opened again, or automatically folding a set of related constructs. An example of the latter is the "short" or "contract" form supported by IDEs like EiffelStudio, whereby the implementations of all features (fields and methods) of a class are folded and hidden from the developer—only natural language documentation and contracts remain in view.

The JML Folding Editor, known under the development code name as "Grok," is an Eclipse plugin that has built-in support for folding JML specifications in an intelligent fashion. It is based upon the Open Source Coffee-Bytes folding plugin.

Manual manipulation of folding is supported, as is customizable defaults in folding behavior. For example, when a file is initially opened, should loop bodies be folded or unfolded? Perhaps only all private method bodies should be folded?

The MOBIUS extensions to this plugin focus on support for folding, in an intelligent fashion, JML formal specifications. For example, groups of adjacent related specifications written in line comments are folded as a unit. Furthermore, the state of folding within a file is saved when the Mylyn subsystem subsection 2.2.1 is in use. Finally, some basic notions of fisheye views [Fur86] have been included in the subsystem.

Our next steps in the research work focus on further evolving and experimenting with fisheye notions of Degree of Interest in the context of software systems with formal specifications. Given the numerous syntactic and semantic relationships that exist between code constructs and specification constructs, and observing the ways in which verification-centric software development and verification takes place, there are many fertile opportunities for good computer-human interaction research.

ESC/Java2

As discussed in subsection section 2.1.2, the Extended Static Checker for Java version 2 (ESC/Java2) is a programming tool that attempts to find common run-time errors in JML-annotated Java programs by static analysis of the program code and its formal annotations. Users control the checking that ESC/Java2 performs by annotating their programs with specially formatted comments called pragmas.

ESC/Java2 is available in two forms:

- ESC/Java2 as a command-line tool with a simple Swing GUI front-end, or
- ESC/Java2 as an Eclipse 3.3 (Europa) plugin.

This latter, plugin-based version of this subsystem has several features that deeply integrate it into the MOBIUS PVE.

With respect to environment configuration, its preferences system is integrated into the Mobius PVE at the IDE-level and at the project-level. The former supports configuration at the systems-level, the latter settings specific to a project that are saved with the project.

Focusing on use, the integration with development is transparent. Each time a file is saved the extended static checker is run in the background asynchronously. Thus, from the developer point of view, warnings and errors appear as regular Eclipse markers on lines of code and as run underlines within program code and specifications. Likewise, all developer feedback appears within the appropriate Eclipse “Warning” and “Error” panes. The extended static checker is also executed manually on a variety of different constructs (on a package, a class, a method, etc.) with a single mouse action.

The MOBIUS Prover Editor

The MOBIUS Prover Editor is an integrated frontend for interactive theorem provers, currently the Coq prover, that is used to prove the generated verification conditions interactively. Originally developed for JACK, it has now been integrated in the MOBIUS environment. See [BBC$^+$07b] for more information.

The basic idea behind the design of the MOBIUS Prover Editor is that reading and writing large-scale formal specifications is much the same as reading and writing large-scale program source [Kin04]. Common programming interface idioms like outline views, syntax highlighting, (possibly type-aware) identifier completion, file and construct templates, etc. are naturally adopted and adapted to support development of, and reasoning about, formal models expressed in higher-order provers.

Additionally, the human-computer interface between the user and the prover via the Mobius PVE must not only be simple, but most importantly, familiar to traditional interactive prover users. Thus, most standard key bindings for the Emacs-based front-end are available from within the MOBIUS Prover Editor.

While the MOBIUS Prover Editor primarily supports Coq, development is underway to also support the PVS theorem prover, porting the improvements previously identified by Kiniry and Owre [KO03] and reemphasizing the Prover Editor’s prover-generic nature.

2.1.7 Static Analysis Subsystems

The Mobius PVE contains over a half dozen static analysis subsystems. Some of these static checker were wholly developed within the MOBIUS project (e.g., ESC/Java2, as described in subsection 2.1.2), some that were developed by others, but evolved and extended within the project, and still others that were adopted and configured to complement the core analyses of MOBIUS. This section describes the second kind of static analysis subsystems—tools that MOBIUS has adopted and extended for our own purposes.

Race Condition Checker

The Race Condition Checker (RCC) was originally developed by researchers at the Digital Systems Research Lab (DEC-SRC) in the late 1990s [FF00]. Its source code was released to the general public along with the source for Simplify, Javafe, ESC/Java, and Houdini.

RCC’s purpose is to (conservatively) statically check for the absence of race conditions via analysis of (type) annotated Java source code. As RCC’s analysis is sound, if an annotated Java program passes RCC, then it is guaranteed to have no race conditions.

As the Mobius PVE is meant to eventually support reasoning about concurrent Java systems, determining whether a Java module is race-free is a critical first step in verifying any other properties. Unfortunately, RCC’s source code release barely compiled, and was completely unmaintained.

Consequently, MOBIUS developers have “revived” RCC and released RCC version 2. This updated version of RCC incorporates the evolution of the Java language semantics, as expressed within the third revision of the JLS [GJSB05] and as realized in Java version 5, particularly with respect to the the Java Memory Model [Con04]. Moreover, several critical flaws in the original RCC type system and type checker were identified and corrected in this release.

Houdini

Like the Race Condition Checker (RCC), Houdini was also originally developed by researchers at the Digital Systems Research Lab (DEC-SRC) in the late 1990s [FL00]. Its source code was released to the general public along with the source for Simplify, Javafe, ESC/Java, and RCC.

Houdini automatically generates specifications written in JML for a Java module by statically analyzing the structure of program code. It has a number of built-in heuristics and performs some lightweight syntactic and semantic reasoning about method bodies to "guess" potential method preconditions, for example. These guesses are then fed to ESC/Java2 to see whether or not they are valid. If a specification seems to be valid (recall that ESC/Java is neither sound nor complete), then the specification is retained, otherwise it is discarded.

Houdini generates reports of its behavior in HTML. Houdini is also capable of taking a snapshot of its state for increasing the reliability of long runs (as running Houdini on a large code base can take hours, or even days). Finally, Houdini can run in a client-server mode, utilizing a (potentially larger, and more powerful) remote compute server.

2.1.8 Backend Subsystems

Several new "backend" subsystems have been developed with MOBIUS resources. These systems are "backend" in the sense that they are used by many other subsystems and are deeply embedded in other tools, thus they are mainly interfaced via APIs and are heavily reused.

These subsystems include automated theorem provers, verification condition generators, specification generators, and intermediate representations for reasoning. All of these subsystems are used within the Mobius PVE, many for several purposes.

Simplify Backend

As discussed in section 2.1.2, Simplify is an automated theorem prover mainly used for software and hardware verification. Like several other automated provers, the MOBIUS Simplify backend bridges Simplify to the MOBIUS Automated Prover API, discussed in section 2.1.5.

Communication with Simplify is accomplished exclusively via piped data streams at this time. Additionally, if the prover terminates abnormally or becomes unavailable, it is restarted automatically. Finally, terminating a validity check via a timeout is also supported.

Fx7

Fx7 is an automated theorem prover implemented in the Nemerle programming language [ML06] by Michal Moskal, a visiting MOBIUS PhD student in 2007. Fx7 supports the same set of theories as Simplify and Z3, as well as a similar set of concrete syntaxes as Z3.

As Nemerle is implemented in Microsoft's .NET platform, the Mono runtime must be installed on non-Microsoft Windows Mobius PVE host operating systems. Currently the Fx7 prover is a hand-installed extension to the Mobius PVE. It will be included in a future release of the Mobius PVE, complete with a pre-packaged runtime for user convenience.

Fx7 has been “bridged” into the MOBIUS Automated Prover API in a fashion akin to the other provers discussed in this section. Piped datastreams are used to communicate with the prover, and starting and stopping the prover and controlling its resources use is provided via the API.

In addition, because Fx7 is exclusively developed by MOBIUS-supported researchers, including contributors from UCD, various experiments in automated proving and verification are easy to carry out. In particular, Fx7 generates proofs of valid formulae and includes two proof checkers which run on embedded devices like PDAs [GM07]. These capabilities represent the first concrete experiences with proof generation and checking the Mobius PVE, albeit for extended static checking with ESC/Java2 rather than full functional verification.

Z3

The Z3 automated theorem prover is developed by Microsoft Research in Redmond [?]. It and Simplify are the primary theorem provers of the Boogie reasoning framework. Z3 is released in binary-only form. The most recent releases are only built for the Microsoft Windows operating system. An older, out-of-date release was built for Linux to conform to the SMT-LIB competition’s rules.

Z3 is a very powerful modern prover, thus it has been “bridged” into the MOBIUS Prover API as well. It supports the Simplify and SMT-LIB concrete syntaxes, as well as its own.

Whereas Z3 only runs on the Microsoft Windows platform, the Mobius PVE runs on all modern platforms, including Apple’s OS X and UNIX variants like Linux, Solaris, etc. We wish to support as many modern theorem provers as possible within the Mobius PVE, but supporting Z3 on non-Windows systems is a non-trivial task.

Recently Z3 has been tested running within the a Wine runtime [win08] on Intel-based systems running Linux. This mechanism (spawning Z3 within Wine, communication via piped datastreams, etc.) is now built into the MOBIUS Automated Prover API.

In theory, such a mechanism will work on all Intel-based Wine runtimes, like those available for other UNIX variants and OS X. In addition, we plan on experimenting with executing Z3 within a virtualized Microsoft Windows installation (e.g., using Parallels or VMware, both of which are used for Mobius PVE development and testing) on OS X on Intel.

2.1.9 Verification Condition Generation

To perform software verification a verification condition (VC) must be generated in some fashion. The Mobius PVE contains two primary VC generators (or VCGen for short) in its gamma release. A new VCgen, included in the FreeBoogie subsystem, will be included in a later release.

Verification condition generation is essentially an algorithm that operates upon some input that describes a program, and produces a logical verification condition describes some desired properties of that program. Thus, each VCGen takes some AST as input, and produces some AST as output. The ASTs discussed in subsection 2.1.3 are used for these VCGen subsystems.

The ESC/Java2 Verification Condition Generator

As discussed in subsection 2.1.2, ESC/Java2 produces a VC by processing one of two input ASTs, either the GC or the DSA form, and producing one of two VC ASTs, either the original DEC-SRC ESC/Java ad hoc VC AST, or the MOBIUS sorted AST.

VCgen is accomplished using either a weakest precondition or a strongest postcondition calculus. These calculi are discussed in detail in [CHK05].

The original ESC/Java VCGen focused on produced VCs describing common programming errors like null pointer exceptions and indexing an array out of bounds. The updated VCGen found in ESC/Java2 also produces VCs describing proper use of constructs found in JML that did not exist in the original pragma language of ESC/Java like model fields.

The background theory of the original VCGen is expressed in the concrete syntax of Simplify, within the built-in theory of Simplify. The background theory of the experimental ESC/Java2 VCGen is expressed in several concrete syntaxes in several logics, including those of Coq, PVS, Simplify, and the SMT format.

Because of this multi-prover nature, the new VCGen is prover-independent, insofar as it generates a VC in the generic representation of the MOBIUS VC AST, which is, in turn, pretty-printed in a theory- and prover-aware fashion.

The MOBIUS Direct Verification Condition Generator

The MOBIUS DirectVCGen is a tool built upon Bico, Bicolano, the ESC/Java2 parser, as well as Mobius PVE's prover backend. It is built around two verification condition generators, one over source code and one over bytecode.

The MOBIUS DirectVCGen is called “direct” because, contrary to the other verification condition generators of the Mobius PVE, it uses a simple weakest precondition calculi without any use of an intermediate language. The source VCGen is written in Java and the bytecode VCGen is written in Coq. It is part of the proof transforming compiler scheme, as explained in the deliverable 4.3 [mC07b].

One of the main goals of this tool is to translate the proofs for the verification conditions generated from the source code to proofs for the verification conditions generated from the bytecode. Therefore, both VCGen subsystems have to be reasonably similar, to allow this transformation. Both weakest precondition calculi have exactly the same set of rules. The only thing that differs between the two calculi is their way of handling variables. Of course, bytecode variables can be easily programatically deduced from the source code ones.

The background predicates for both verification conditions are found in the same theory, that of Bicolano. The tool Bico is used to translate a program into the Coq formalization. The backend for the VCGen over source code is a specific Coq backend that uses Bicolano as its memory model.

The FreeBoogie Verification Condition Generator

The FreeBoogie VCGen is meant to replicate and extend the VCGen found in the original Boogie tool from Microsoft Research. In particular, it will support seperation logic operators natively in an extended version of BoogiePL. As this VCGen is still under development, it will not be described here in detail.

The Universes Typecheckers

The universes typesystems [MR07a] are used to describe ownership, containment, and aliasing properties of Java programs. The version of the universes typesystem described in Muller et al. [MR07b] is implemented in two subsystems in the Mobius PVE: in the core JML tool suite found in the JML2 Eclipse plugin, and in ESC/Java2.

Inferring and Propagating Annotations

Verifying a program typically requires many of program annotations. Error messages produced by a program analysis tool, which indicate failed verification attempts, can be used to automate the generation and propagation of program annotations in code.

This approach is pursued in the CANAPA tool [CFJ$^+$06]: CANAPA uses the error messages of ESC/Java2 to infer and propagate non-null annotations. CANAPA comes with an Eclipse plug-in.

In an attempt to further reduce the burden of annotation writing further, we have implemented an algorithm dedicated to the generation of “obvious” minimal preconditions to avoid not only null-pointer, but also array-out-of-bounds exceptions (see [BBC$^+$07b]). This algorithm re-uses the implementation of the wp-calculus. A possibility to improve the algorithm is to apply a simple analysis to identify possible class invariants.

Specification Generation

An alternative way to alleviate the burden of annotating code is to provide some higher-level specification formalism, which makes developing specifications less work. These high-level specifications are then automatically translated into to more basic program annotations.

The earlier work on security requirements reported in Deliverable 1.2 [mC06] suggested that so-called midlet navigation graphs [Cré06] could be such a higher-level specification formalism. Such graphs are essentially a form of security automata, and the existing industry standard for testing MIDP applications [Ini06] already prescribes the use of a similar, albeit completely informal, notion of “flow diagram.”

Security automata are a convenient way to describe security policies. Traditionally they are used to monitor an application at runtime, and to interrupt it as soon as it violates the security policy. However, in our setting of PCC, we would like to use security automata as JML specifications, and verify adherence to them statically. The JML annotations effectively inline the monitor in the application.

The translation of security automata into program annotations is not trivial, and there is the danger that the semantics of the annotations is not equivalent to the effect of runtime monitoring.

The translation from security automata to program annotations is defined in several steps:

1. complete the security automata, adding a special trap state that should not be reached (whereas partial automata, describing only the allowed transitions are sufficient for monitoring);

2. generate special set-annotations at the level of method specifications, capturing the behavior of the security automata, together with an invariant stating that the error state should not be reached; and

3. inline the annotations into the method body.

We have developed a generic program semantics that is instantiated for monitoring and run-time annotation checking. We prove that every translation step preserves the program behavior. In particular this means that monitoring will not find a security violation iff run-time annotation checking will not find any annotation violation (provided that earlier existing annotations are not violated either).

A full description of the translation, its formalisation and the correctness proof are available in [?]. The correctness proofs revealed several subtleties that have to be considered in the algorithm, and necessary assumptions about program behavior.

Using an annotation propagation algorithm (as in [PBB+04]) the generated annotations give rise to pre- and postcondition annotations in the whole application, that can be verified statically.

In the simplest form of security automata, each edge corresponds to some unique event, typically some method call, but for convenient specifications one can allow richer languages for decorating edges. In [PR08] we define suitable notions of refinement for MIDP navigation graphs, and propose a translation of these to JML annotations.

We have also adapted AutoJML, a tool we developed earlier in Nijmegen [HO03] to generate JML annotations from automata-like specifications, to cope with richer forms of automata. In [PS07] this has been used to generate JML annotations to verify that the MIDP-SSH implementation is correct with respect to a formal specification of the SSH protocol. Here our formal specification probably goes into much more detail than one would typically give in a security automaton.

FreeBoogie

FreeBoogie aims to be a backend for verification tools such as Spec$^{\#}$ and future versions of ESC/Java and similar tools. It takes as input a BoogiePL (Boogie Programming Language) file, generates verification conditions, calls provers such as Simplify and Z3, and reports the errors in a way that relates to the BoogiePL source. It is a work in progress.

The most important part missing is the verification condition generation. The tool can already be used to parse, typecheck, and build symbol tables and flow graphs for BoogiePL code. The AST data structures are generated from a short description of the BoogiePL abstract grammar by a tool included in FreeBoogie. The backend communicates with Simplify and with Z3 running in Simplify mode. (There are only minor differences in the way the two provers communicate.) The design allows adding other provers easily, in particular there is support for sort-checking that should facilitate adding sorted provers.

The API is available at http://radu.ucd.ie/freeboogie/doc/. The source is available via the MOBIUS Trac at https://mobius.ucd.ie/browser/src/freeboogie/trunk/FreeBoogie.

The Spec$^{\#}$ system from Microsoft Research generates BoogiePL from the Spec$^{\#}$ language, which is a variant of Microsoft's C$^{\#}$ programming language with annotations. Another way of obtaining BoogiePL from a higher level language is via JACK [BBC+07a], which translates Java with JML annotations into a variant of BML, and B2BPL, which was developed at ETH under the MOBIUS project and translates BML into BoogiePL. In fact, one design goal of BoogiePL was to be used as an intermediate language by multiple verification tools, thereby reducing the implementation effort needed to support multiple high-level languages and multiple static analyzes at the same time.

An initial release of the FreeBoogie subsystem, which supports parsing, name resolution, flow graphs, and typechecking, took place at T0+18.

This concludes a summary of Mobius PVE subsystems developed under the MOBIUS project.

2.2 Subsystems Used within MOBIUS but Developed Elsewhere

Currently eight major "external" subsystems are included in the Mobius PVE. All of these subsystems have programmatic interfaces available to Mobius PVE developers and are wrapped in Eclipse plugins for Mobius PVE users. Many of them are also "tuned" for verification, as described in the following pages.

These particular tools have been chosen for inclusion in the Mobius PVE because they have been very useful to MOBIUS partners in verification-centric software engineering over the past several years [KZ08]. As new tools appear in the Eclipse community, they are objective evaluated by MOBIUS partner UCD for inclusion in future releases of the Mobius PVE.

2.2.1 Mylyn

Mylyn is a "Task-Focused Interface" plugin developed originally by researchers at the University of British Columbia, and is now a full-fledged Open Source project recognized by the Eclipse Foundation. Mylyn makes tasks first-class entities within Eclipse and supports offline, asynchronous interations with the MOBIUS Trac bug and feature tracker from within Eclipse, and thus from within the Mobius PVE.

Additionally, when a task is "active", Mylyn watches the Mobius PVE user work and figures out what constructs (files, classes, methods, theories, definitions, lemmas, etc.) are important with respect to that particular task. This means that only those constructs relevant to the current work, be it developing new software or performing a verification, are visible at any point in time. When a task is labeled as "inactive", Mylyn automatically snapshots the current state of the Mobius PVE and puts everything away so the user is returned to a clean initial desktop. Furthermore, this context can be saved as a file or attached to a Trac wiki page or ticket so that others may return to the same state as the original Mobius PVE user.

2.2.2 Version Control

MOBIUS developers primarily use two version control systems: CVS and Subversion. Eclipse comes with a CVS plugin by default. Two primary Subversion plugins are available today: Subclipse and Subversive. After evaluating both options, we chose to ship Subclipse with the Mobius PVE, as it is more reliable and has a more appropriate feature set for dependable systems development. No special configuration of our version control plugins has taken place.

The Eclipse Mobius PVE developer workspace that we provide via the MOBIUS Trac server is pre-configured and points (anonymously) to all of the version control repositories relevant to MOBIUS work, including those hosted by MOBIUS partners, and those hosted by external sites like SourceForge and our North American colleagues.

The student/instructional/demonstration Mobius PVE workspace that we provide does not point to any version control repositories.

2.2.3 CheckStyle

CheckStyle is a plugin that performs source code style checking. Style checking involves statically analyzing both documentation and source code against simple, syntactic rules. At this time CheckStyle does not check JML annotations. Thus, beyond providing example coding styles codified for checking with this plugin, no special configuration of the CheckStyle plugin has taken place.

As described in chapter 4, a major part of the MOBIUS process is ensuring that source code is of a certain quality, measured in terms of the number of errors and warnings produced by several static checkers working in tandem. The first of these checkers is CheckStyle.

The Mobius PVE comes with several example coding styles: a basic style that promotes readable code, a moderate style developed by MOBIUS partner Warsaw University that focuses on dependable code, and a rigorous style developed by MOBIUS partner University College Dublin that focuses on verifiable code. Additionally, all non-basic styles inherit from the basic style, thus all checks in the basic style are also contained in the moderate and rigorous styles.

Style checkers work hand-in-hand with code generators. In Eclipse, there are three basic kinds of code generators: code templates and skeletons, the results of refactoring, and writing new code "by-hand" while using Eclipse's built-in code formatter.

Consequently, each of these mechanisms for generating code must be "tuned" to match the expected structures of the tool that checks code. As such, the Mobius PVE comes with appropriate code templates and a code formatting specifications to match each coding style. Refactoring need not be "tuned" as refactored code is automatically styled according to the active code formatting style.

In summary, the Mobius PVE assists developers in writing clear, readable, maintainable documentation and program code by both generating code to a particular per-project style and checking that code conforms to said style.

The next level of rigor in static checking uses not just syntactic checks, but some lightweight semantic checks as well. The two main plugins used in the Mobius PVE for such are FindBugs and PMD, covered in the following two sections.

2.2.4 FindBugs

Static checkers like FindBugs check for programming problems beyond syntactic code style checking. FindBugs is a configurable bytecode-level checker that analyzes code for common programming errors. FindBugs does not use a formal semantics or a theorem prover; its checks are entirely based upon the structure of Java bytecode.

In the Mobius PVE, we have "tuned" FindBugs to perform those checks that complement and support writing code which is verifiable using MOBIUS verification tools. In other words, if a program passes the Mobius PVE pre-configured static checkers, it is more amenable to formal verification.

Each check that is enabled in FindBugs has been evaluated and objectively analyzed to determine if it should be enabled or disabled for verification-centric programming. This evaluation is ongoing as new checks become available as FindBugs evolves.

2.2.5 PMD

Likewise, we have "tuned" the static checker PMD in a similar fashion. PMD analyzes source code, rather than bytecode. Thus, it is not usable in all circumstances, as sometimes source is not available due to licensing, legal, or intellectual property-related reasons.

Evaluation and maintaining the configuration of PMD in the Mobius PVE mirrors that of FindBugs—as the plugin evolves, so does the set of checks enabled by default to support verification-centric development.

2.2.6 Metrics

A variety of plugins support measuring a suite of standard software engineering metrics in program code. Years of experience with automatic and interactive verification have taught Mobius PVE developers that, code which has a specific range of values for a set of well-understood metrics are indicative of program code that is easier to verify. Therefore, providing a plugin that performs metric static analysis is useful to Mobius PVE users.

Two metrics-related plugins ship with the current "gamma" version of the Mobius PVE: NCSS and Metrics.

NCSS

JavaNCSS is a simple program that performs some very basic analysis of program source code. In particular, JavaNCSS counts the Non-Commenting Source Statements (NCSS) of Java source at the module and method level, as well as summarizing a variety of other tertiary data.

The NCSS plugin included with the Mobius PVE provides an extremely simple interface to triggering JavaNCSS from within Eclipse. Analysis is only supported at the project level, and the output of said analysis is provided within a read-only window containing the raw text output of JavaNCSS.

As we gain more experience with and "tune" the Metrics plugin discussed in the next section, it is expected that future versions of the Mobius PVE will not include the NCSS plugin.

Metrics

The Metrics plugin, in addition to performing all of the actions that NCSS does, also is capable of analyzing source code for well-known metrics found in the software engineering literature. Also, metrics at the class and method level are summarized in a more useful contextual structured fashion, rather than in a plain text window. For example, when a given code construct is selected with the mouse, the Metrics reporting window automatically updates to reflect the metrics of that construct.

MOBIUS partners have not gained much hands-on experience in some of these more advanced kinds of metrics, like various kinds of code coupling analyses. Thus, there is little concrete evidence about the verifiability of code with respect to such metrics. In fact, as there is little scientific published data on such, the integration of verification and metrics tools in the Mobius PVE provides the foundation for new research in such.

2.2.7 EclEmma

The last external plugin integrated into the Mobius PVE focuses on code coverage analysis. The code coverage plugin chosen for use within the Mobius PVE is called EclEmma.

Emma is a code coverage analysis tool that, during program execution (typically during unit testing), automatically instruments program bytecode to gather information about various kinds of standard code coverage measures, including branch and statement coverage. Emma reports such analysis back to the programmer through a series of reports.

EclEmma wraps Emma into an Eclipse plugin with a rich user interface. Coverage data is summarized to the user in a simple graphical fashion where program source is highlighted in various ways to indicate if a given construct has been executed during a given test run. For example, if a particular method was never executed during unit testing, its body is highlighted in red.

No customization has yet taken place for EclEmma. The main customization requested thus far is the ability to perform code coverage analysis during runtime assertion checking. Highlighting the program source code and executable specifications (in particular, method contracts and invariants) executed in a given test run will be particularly useful to Mobius PVE users. Such data will help users understand the completeness of their specifications with respect to their program code, as well as the completeness of their specifications with respect to their unit tests.

Additionally, given EclEmma performs aggregate coverage analysis, e.g., a given class's executable statements are 90% covered by the current unit tests. Such data will give a concrete measure that can be related to other aggregate measure like those reported by the variety of static analysis tools included in the Mobius PVE like FindBugs and ESC/Java2. A verification "dashboard" is planned that succinctly, visually, interactively summarizes all of this information for the programmer and a development team.

This concludes a summary of Mobius PVE subsystems developed by parties outside of the MOBIUS project.

Furthermore, this concludes a summary of all Mobius PVE subsystems.

Chapter 3

Using the MOBIUS Program Verification Environment

As the Mobius PVE is simply an enriched version of Eclipse, Eclipse users will find the Mobius PVE quite familiar. This chapter provides a high-level developer walk-through of the configuration and use of the Mobius PVE.

3.1 General Walk-through of Using the Mobius PVE

The first step in using the Mobius PVE is to obtain and install it from the MOBIUS website.

3.1.1 Downloading

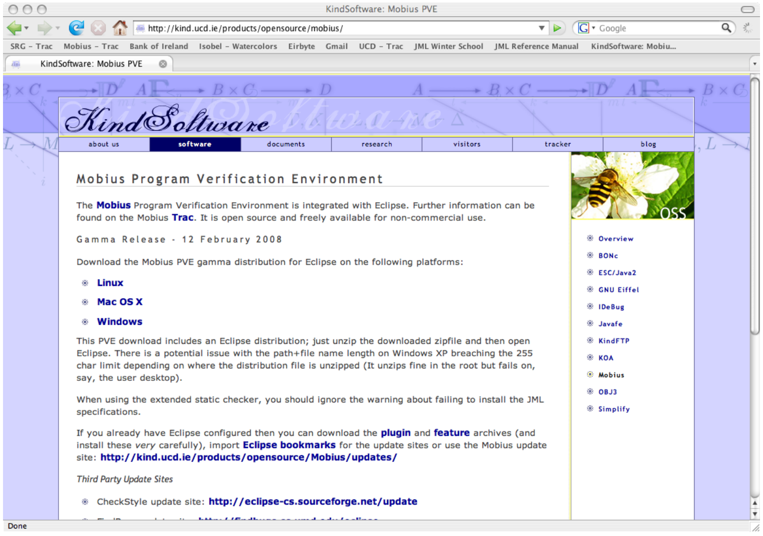

Mobius Program Verification Environment

The **Mobius** Program Verification Environment is integrated with Eclipse. Further information can be found on the Mobius **Trac**. It is open source and freely available for non-commercial use.

Gamma Release - 12 February 2008

Download the Mobius PVE gamma distribution for Eclipse on the following platforms:

- **Linux**
- **Mac OS X**
- **Windows**

This PVE download includes an Eclipse distribution; just unzip the downloaded zipfile and then open Eclipse. There is a potential issue with the path+file name length on Windows XP breaching the 255 char limit depending on where the distribution file is unzipped (It unzips fine in the root but fails on, say, the user desktop).

When using the extended static checker, you should ignore the warning about failing to install the JML specifications.

If you already have Eclipse configured then you can download the **plugin** and **feature** archives (and install these *very* carefully), import **Eclipse bookmarks** for the update sites or use the Mobius update site: **http://kind.ucd.ie/products/opensource/Mobius/updates/**

Third Party Update Sites

- CheckStyle update site: **http://eclipse-cs.sourceforge.net/update**

Figure 3.1: The Mobius PVE Download Site at KindSoftware.

The Mobius PVE is available from the MOBIUS Trac server http://mobius.ucd.ie/ and via the KindSoftware research group website (the research group of the MOBIUS UCD partner), as seen in Figure 3.1. Pre-configured Eclipse builds are available for all main platforms including Microsoft Windows, Apple's Mac OS X, and Linux. Download the appropriate archive; it can be unpacked anywhere on the local disk.

Sample Workspace A sample Eclipse workspace is also available from the aforementioned sites. This workspace includes example projects and documentation about the Mobius PVE.

3.1.2 Configuring a Java Project for Verification

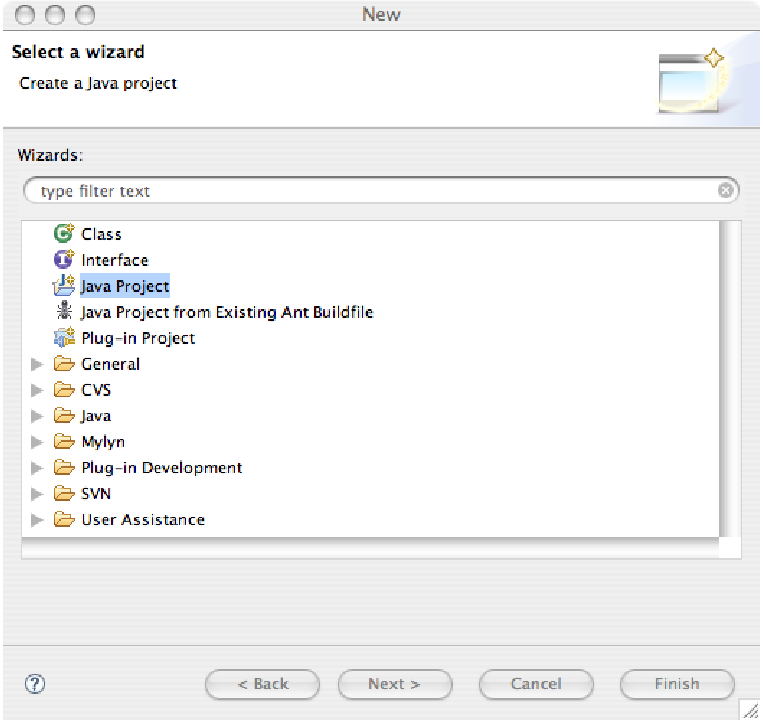

Select a wizard

Create a Java project

Figure 3.2: Creating a new Java Eclipse project from scratch.

Each main subsystems included in the current "gamma" release of the Mobius PVE is configured separately. First one must create a new Java project, as seen in Figure 3.2 or import an existing Java project. One may also simply use an existing Eclipse workspace with the Mobius PVE.

Coding Standards

The first subsystem that one should configure is the CheckStyle coding standard checker and the complementary Eclipse code formatter. The CheckStyle plugin is configured via the Project Properties dialog, available either via the Project menu or via a right mouse action on the project name in the resources browser.

The Mobius PVE comes pre-configured with the MOBIUS coding standard built-in, as seen in Figure 3.3.

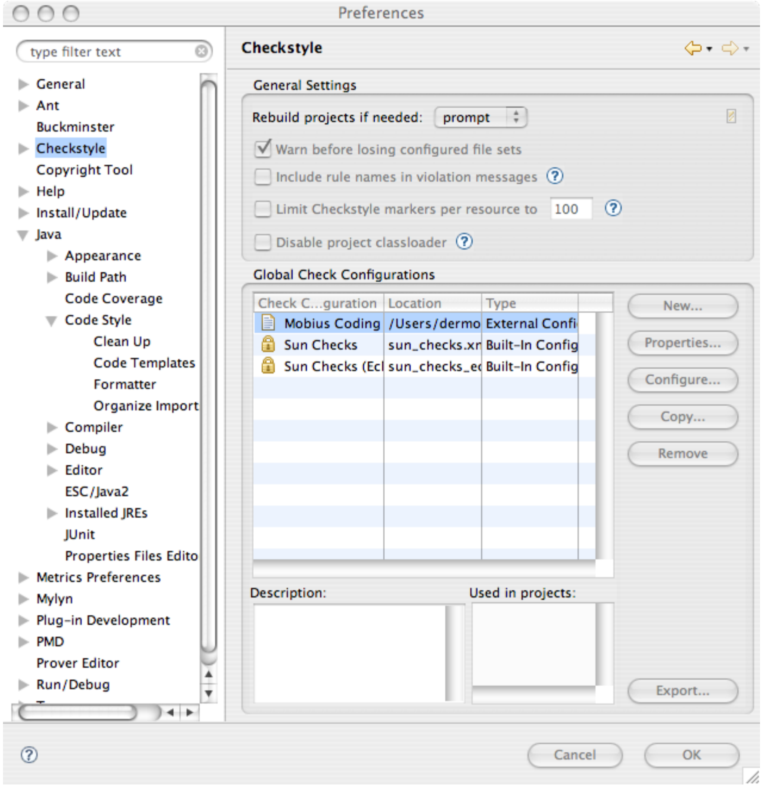

Checkstyle

General Settings

Rebuild projects if needed: prompt

☑ Warn before losing configured file sets

□ Include rule names in violation messages

□ Limit Checkstyle markers per resource to 100

□ Disable project classloader

Global Check Configurations

Check C...guration	Location	Type
Mobius Coding	/Users/dermo	External Confi
Sun Checks	sun_checks.xr	Built-In Config
Sun Checks (Ecl	sun_checks_ec	Built-In Config

Description:

Used in projects:

Figure 3.3: The default CheckStyle configuration.

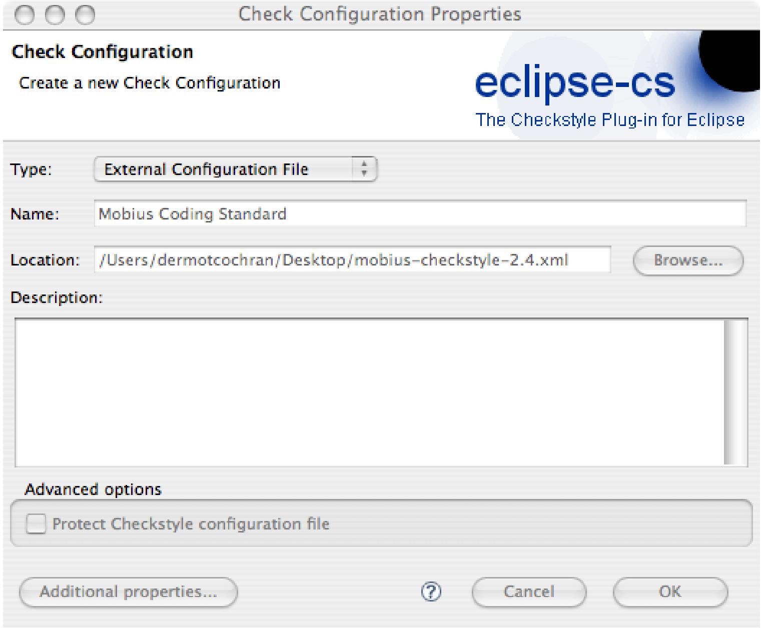

Check Configuration Properties
eclipse-cs
The Checkstyle Plug-in for Eclipse
Browse...
Additional properties...
Cancel
OK

Check Configuration

Create a new Check Configuration

Type: External Configuration File

Name: Mobius Coding Standard

Location: /Users/dermotcochran/Desktop/mobius-checkstyle-2.4.xml

Description:

Advanced options

□ Protect Checkstyle configuration file

Figure 3.4: Importing a new coding standard.

If one wishes to further configure the default coding standard, either the configuration can be tweaked via the plugin's preferences panel, or one can duplicate and edit the external MOBIUS coding standard configuration file, available from the MOBIUS Trac server. Figure 3.4 illustrates this possibility.

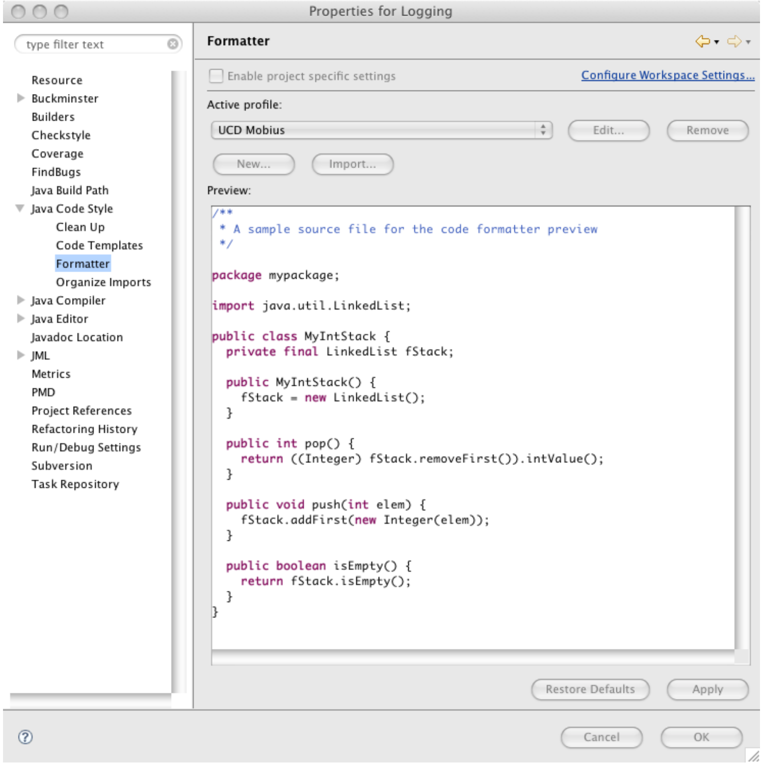

Formatter

□ Enable project specific settings

Active profile: UCD Mobius

Preview:

```
/**
 * A sample source file for the code formatter preview
 */

package mypackage;

import java.util.LinkedList;

public class MyIntStack {
  private final LinkedList fStack;

  public MyIntStack() {
    fStack = new LinkedList();
  }

  public int pop() {
    return ((Integer) fStack.removeFirst()).intValue();
  }

  public void push(int elem) {
    fStack.addFirst(new Integer(elem));
  }

  public boolean isEmpty() {
    return fStack.isEmpty();
  }
}
```

Figure 3.5: Configuring the Eclipse Code Formatter.

Configuring the Eclipse code formatter takes place in a similar fashion. The Mobius PVE is pre-configured with several code formatters. Figure 3.5 shows the example UCD code formatter configuration that is included with the sample workspace.

Configuring and Using the JML Plugin

The JML plugin is configured on a per-project basis. The JML typechecker, compiler, and runtime assertion checker are configured independently. Figure 3.6 shows the preferences panel for configuring the JML typechecker.

JML typechecking or compilation is currently accomplished via a menu action, as seen in Figure 3.7. In a future version of the Mobius PVE, autobuilders will be intelligently triggered instead.

Configuring and Using Static Checkers

Several static checkers are included in the Mobius PVE. First, configuring and executing the main static checker, ESC/Java2, is summarized.

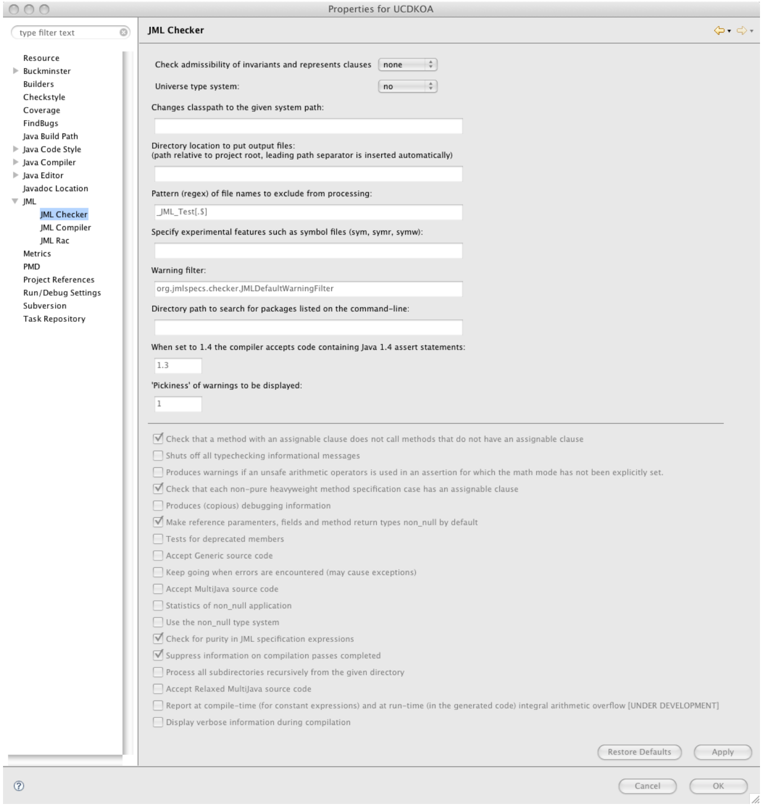

JML Checker

Check admissibility of invariants and represents clauses: none

Universe type system: no

Changes classpath to the given system path:

Directory location to put output files:
(path relative to project root, leading path separator is inserted automatically)

Pattern (regex) of file names to exclude from processing: _JML_Test[.$]

Specify experimental features such as symbol files (sym, symr, symw):

Warning filter: org.jmlspecs.checker.JMLDefaultWarningFilter

Directory path to search for packages listed on the command-line:

When set to 1.4 the compiler accepts code containing Java 1.4 assert statements: 1.3

'Pickiness' of warnings to be displayed: 1

- ☑ Check that a method with an assignable clause does not call methods that do not have an assignable clause
- ☐ Shuts off all typechecking informational messages
- ☐ Produces warnings if an unsafe arithmetic operators is used in an assertion for which the math mode has not been explicitly set.
- ☑ Check that each non-pure heavyweight method specification case has an assignable clause
- ☐ Produces (copious) debugging information
- ☑ Make reference paramenters, fields and method return types non_null by default
- ☐ Tests for deprecated members
- ☐ Accept Generic source code
- ☐ Keep going when errors are encountered (may cause exceptions)
- ☐ Accept MultiJava source code
- ☐ Statistics of non_null application
- ☐ Use the non_null type system
- ☑ Check for purity in JML specification expressions
- ☑ Suppress information on compilation passes completed
- ☐ Process all subdirectories recursively from the given directory
- ☐ Accept Relaxed MultiJava source code
- ☐ Report at compile-time (for constant expressions) and at run-time (in the generated code) integral arithmetic overflow [UNDER DEVELOPMENT]
- ☐ Display verbose information during compilation

Figure 3.6: Configuring the JML Checkers.

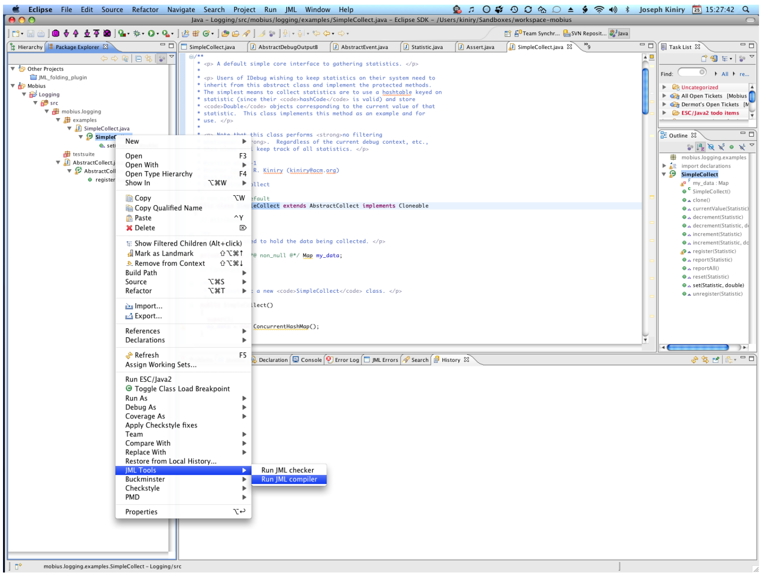

Figure 3.7: Executing the JML Compiler.

ESC/Java2 The main static checker included with the Mobius PVE at this time is ESC/Java2. ESC/Java2 is configured via menu actions and preferences panels.

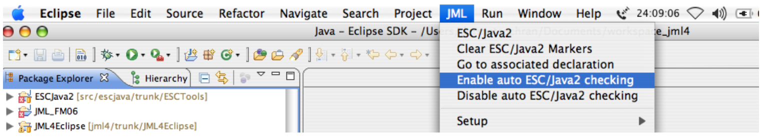

Figure 3.8: Enabling the ESC/Java Nature.

To initially configure the plugin, one must enable the ESC/Java "nature" (a set of configurations that ensures the plugin is properly configured for use) and "builder" (an action that is triggered at specified times, typically to build a system or to perform some analysis). The JML menu contains item to perform initial setup of the plugin (Figure 3.8).

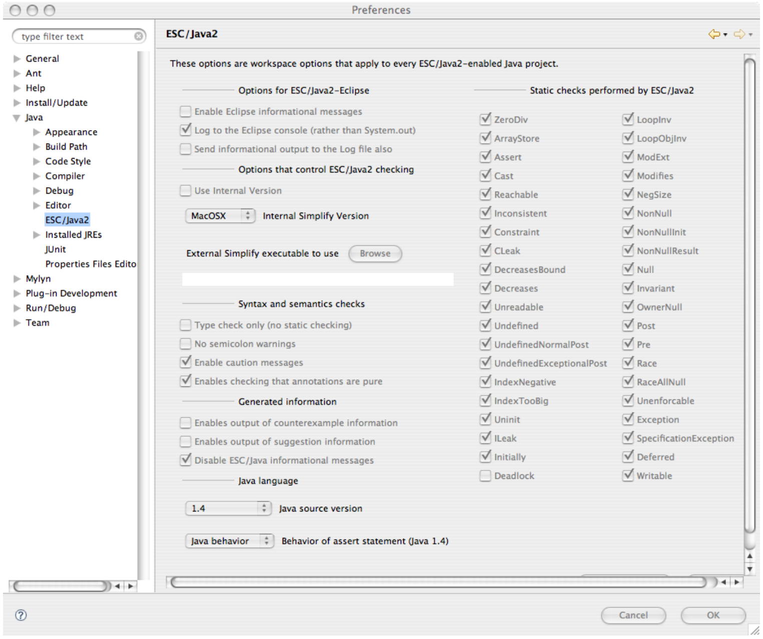

ESC/Java2

These options are workspace options that apply to every ESC/Java2-enabled Java project.

Options for ESC/Java2-Eclipse

□ Enable Eclipse informational messages
☑ Log to the Eclipse console (rather than System.out)
□ Send informational output to the Log file also

Options that control ESC/Java2 checking

□ Use Internal Version

MacOSX Internal Simplify Version

External Simplify executable to use Browse

Syntax and semantics checks

□ Type check only (no static checking)
□ No semicolon warnings
☑ Enable caution messages
☑ Enables checking that annotations are pure

Generated information

□ Enables output of counterexample information
□ Enables output of suggestion information
☑ Disable ESC/Java informational messages

Java language

1.4 Java source version

Java behavior Behavior of assert statement (Java 1.4)

Static checks performed by ESC/Java2

☑ ZeroDiv
☑ ArrayStore
☑ Assert
☑ Cast
☑ Reachable
☑ Inconsistent
☑ Constraint
☑ CLeak
☑ DecreasesBound
☑ Decreases
☑ Unreadable
☑ Undefined
☑ UndefinedNormalPost
☑ UndefinedExceptionalPost
☑ IndexNegative
☑ IndexTooBig
☑ Uninit
☑ ILeak
☑ Initially
□ Deadlock
☑ LoopInv
☑ LoopObjInv
☑ ModExt
☑ Modifies
☑ NegSize
☑ NonNull
☑ NonNullInit
☑ NonNullResult
☑ Null
☑ Invariant
☑ OwnerNull
☑ Post
☑ Pre
☑ Race
☑ RaceAllNull
☑ Unenforcable
☑ Exception
☑ SpecificationException
☑ Deferred
☑ Writable

Figure 3.9: Configuring ESC/Java2.

Project-specific configuration of ESC/Java2 is accomplished via a Project Properties panel, as seen in Figure 3.9. Via this control panel, one controls all of the standard warnings for ESC/Java2, the use of various theorem provers, and programming language variants, including JavaCard (see Figure 3.10).

The ESC/Java builder can be enabled to run each and every time a Java source file is saved. This is accomplished via another item in the JML menu, as seen in Figure 3.11.

Finally, ESC/Java2 can be manually executed on any project subunit (file, package, etc.) via a right mouse action on the appropriate resource (Figure 3.12).

Other Static Checkers Configuring plugins like PMD, FindBugs, Mylyn, EclEmma and Metrics all happen in the standard fashion. Decide which, if any, of the non-MOBIUS static check-

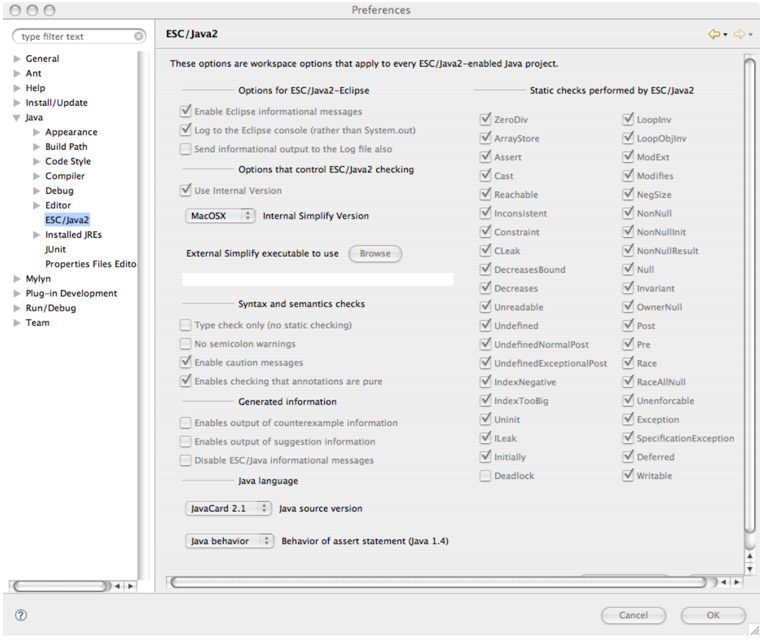

ESC/Java2

These options are workspace options that apply to every ESC/Java2-enabled Java project.

Options for ESC/Java2-Eclipse

- ☑ Enable Eclipse informational messages
- ☑ Log to the Eclipse console (rather than System.out)
- □ Send informational output to the Log file also

Options that control ESC/Java2 checking

- ☑ Use Internal Version
- MacOSX — Internal Simplify Version
- External Simplify executable to use [Browse]

Syntax and semantics checks

- □ Type check only (no static checking)
- □ No semicolon warnings
- ☑ Enable caution messages
- ☑ Enables checking that annotations are pure

Generated information

- □ Enables output of counterexample information
- □ Enables output of suggestion information
- □ Disable ESC/Java informational messages

Java language

- JavaCard 2.1 — Java source version
- Java behavior — Behavior of assert statement (Java 1.4)

Static checks performed by ESC/Java2

- ☑ ZeroDiv
- ☑ ArrayStore
- ☑ Assert
- ☑ Cast
- ☑ Reachable
- ☑ Inconsistent
- ☑ Constraint
- ☑ CLeak
- ☑ DecreasesBound
- ☑ Decreases
- ☑ Unreadable
- ☑ Undefined
- ☑ UndefinedNormalPost
- ☑ UndefinedExceptionalPost
- ☑ IndexNegative
- ☑ IndexTooBig
- ☑ Uninit
- ☑ ILeak
- ☑ Initially
- □ Deadlock
- ☑ LoopInv
- ☑ LoopObjInv
- ☑ ModExt
- ☑ Modifies
- ☑ NegSize
- ☑ NonNull
- ☑ NonNullInit
- ☑ NonNullResult
- ☑ Null
- ☑ Invariant
- ☑ OwnerNull
- ☑ Post
- ☑ Pre
- ☑ Race
- ☑ RaceAllNull
- ☑ Unenforcable
- ☑ Exception
- ☑ SpecificationException
- ☑ Deferred
- ☑ Writable

Figure 3.10: Configuring ESC/Java2 for use with JavaCard.

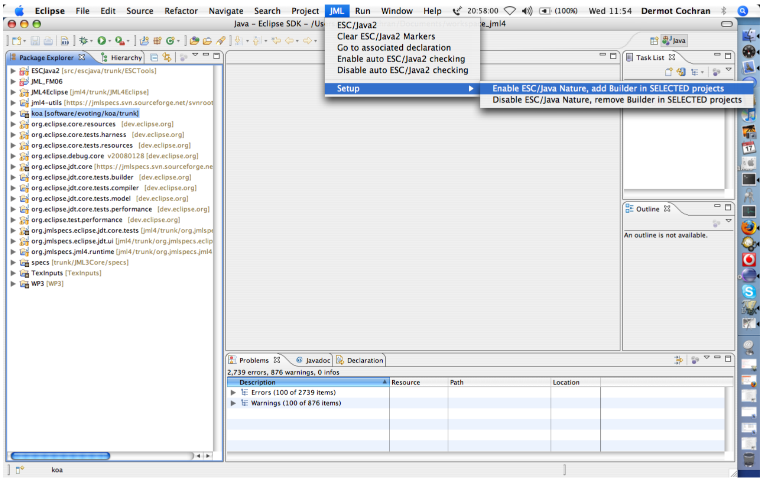
Eclipse File Edit Source Refactor Navigate Search Project JML Run Window Help
Wed 11:54 Dermot Cochran
Java - Eclipse SDK
ESC/Java2
Clear ESC/Java2 Markers
Go to associated declaration
Enable auto ESC/Java2 checking
Disable auto ESC/Java2 checking
Setup
Enable ESC/Java Nature, add Builder in SELECTED projects
Disable ESC/Java Nature, remove Builder in SELECTED projects
Package Explorer
Hierarchy
ESCJava2 [src/escjava/trunk/ESCTools]
JML_FM06
JML4Eclipse [jml4/trunk/JML4Eclipse]
jml4-utils
koa [software/evoting/koa/trunk]
org.eclipse.core.resources
org.eclipse.core.tests.harness
org.eclipse.core.tests.resources
org.eclipse.debug.core
org.eclipse.jdt.core
org.eclipse.jdt.core.tests.builder
org.eclipse.jdt.core.tests.compiler
org.eclipse.jdt.core.tests.model
org.eclipse.jdt.core.tests.performance
org.eclipse.test.performance
org.jmlspecs.eclipse.jdt.core.tests
org.jmlspecs.eclipse.jdt.ui
org.jmlspecs.jml4.runtime
specs [trunk/JML3Core/specs]
TexInputs [TexInputs]
WP3 [WP3]
Task List
Outline
An outline is not available.
Problems
Javadoc
Declaration
2,739 errors, 876 warnings, 0 infos
Description
Resource
Path
Location
Errors (100 of 2739 items)
Warnings (100 of 876 items)
koa

Figure 3.11: Enabling automatic execution of ESC/Java2.

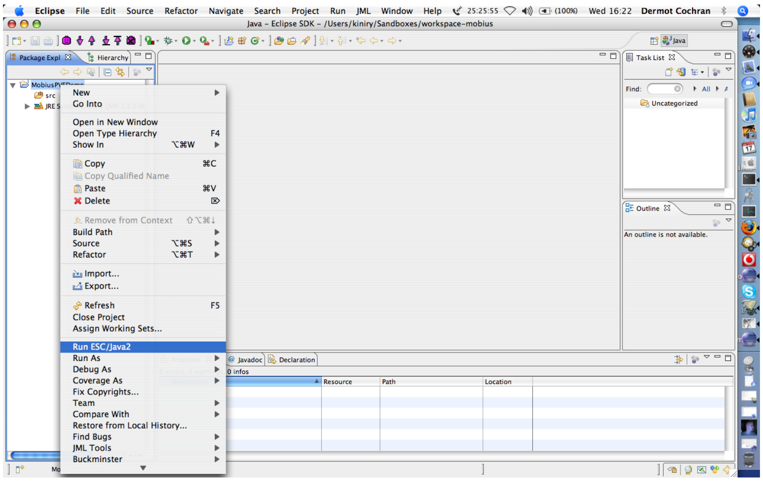

Figure 3.12: Manually Executing ESC/Java2.

ers should be enabled as auto-builders. (An auto-builder is a builder that is automatically triggered when a specific action takes place, e.g., when a file is saved.)

As all plugins are pre-configured to compliment verification with MOBIUS subsystems, their preference settings should, in general, should be untouched.

Configuring Other Subsystems

Several other subsystems (the JML Folding subsystem, the MOBIUS Coq editor, the Umbra bytecode viewer, etc.) are configured in similar ways, but are not discussed here for brevity's sake

3.1.3 Performing Verification

Performing verification within the Mobius PVE is integrated directly into the standard practices of software development.

Executing MOBIUS and non-MOBIUS Static Checkers

All static checkers can be executed manually with a right mouse button action and/or a menu operation, typically at multiple levels of granularity (e.g., check a method, a class, a file, a package, etc.). Some static checkers like CheckStyle, PMD, FindBugs, and ESC/Java2 can also be configured as autobuilders, which are triggered each time a resource changes.

Generating a Verification Condition

As discussed in subsection 2.1.9, verification conditions are generated by three different subsystems shipped with the Mobius PVE at the moment.

VC generation from ESC/Java2 is built-in and VCs are only retained if special switches, currently unavailable via the Eclipse ESC/Java2 plugin, are used.

VC generation with the MOBIUS DirectVCGen is also only controlled via command-line tools at this time.

Finally, as discussed in the aforementioned chapter, development of the VC generator for FreeBoogie is still underway.

Executing the MOBIUS Prover Editor

The MOBIUS Prover Editor is an Eclipse “perspective”. It is automatically loaded and initialized anytime a Coq theory file or proof script file is loaded within the Mobius PVE. No special configuration is necessary. If the prover needs to be reset, a small reinitialize button is available in the toolbar.

Viewing and Manipulating BML-annotated Java Bytecode

Finally, viewing and editing Java bytecode with Umbra is accomplished by simply opening a bytecode file. While viewing bytecode, a variety of buttons are available that permit one to refresh the view, jump to source corresponding to the currently selected bytecode fragment, etc.

This completes our brief summary of installation, configuration, and use of the Mobius PVE.

Chapter 4

The MOBIUS Software Development Process

Software developed within the MOBIUS project takes place using a variety of tools and processes. As the Mobius PVE has been under development during the first two years of the project, MOBIUS developers are now able to use the prototypeMobius PVE to verify existing and ongoing work. Likewise, software development processes change to accommodate new verification tools and related technologies as they become available. Most developers have used, to date, tools and processes which are related to the intent of the Mobius PVE, e.g. older versions of the JML tools or standard open source static checkers.

The core elements of any medium-to-large scale software development effort, a coding standard, a collaborative development environment, and a version control repository, were put in place very early in MOBIUS development. This chapter briefly summarizes these resources and the current and intended use of each.

4.1 The MOBIUS Coding Standards

The MOBIUS coding standard is derived from the UCD coding standard, originally developed by Joe Kiniry while at Caltech in 1996, then updated while he worked at Radboud University Nijmegen from 2002–2004. This standard has influenced the programming standards of many other research groups and organizations, including Sun Microsystems.

This coding standard stipulates the syntactic style of Java source code including identifier naming, code layout, documentation style and completeness, etc. Additionally, guidelines on appropriate values of a variety of metrics measured by subsystems like those described in subsection 2.2.6, are provided.

The base MOBIUS coding standard is codified in as a CheckStyle style and an Eclipse code formatter configuration. Variant coding styles, specific to various projects, groups, development teams, and sites, are derived from this base standard. They are also codified using CheckStyle and Eclipse. Example “moderate” and “rigorous” standards are provided with the Mobius PVE, as discussed in subsection 2.2.6.

While not all source for all subsystems has been checked against a specific coding standard as of yet, the majority of MOBIUS-specific Mobius PVE code has been checked against one of three different standards. As Mobius PVE development continues, and the MOBIUS development process evolves and becomes more codified within the Mobius PVE, new and existing subsystems are checked against new and existing styles.

4.2 The MOBIUS Collaborative Development Environments

Several web-based Collaborative Development Environments (CDEs) have been used for Mobius PVE development thus far.

Tools developed prior to the MOBIUS project's genesis were hosted at open CDEs sites like SourceForge and private CDEs like the SoS Bugzilla and the UCD GForge.

The initial CDEs used for MOBIUS coordination and some development effort were GForges hosted by INRIA (located at http://mobius.inria.fr) and UCD (located at http://sort.ucd.ie/).

To better coordinate development on the Mobius PVE, a single CDE has been configured at UCD, as it leads the main task relating to Mobius PVE development (Task 3.6). A Trac server, securely hosted at UCD, now hosts nearly all of Mobius PVE development. Existing collaborative artifacts like bug and feature request tickets, patches, repository contents, and wiki pages have been migrated from the various older CDEs mentioned above into this Trac server.

4.2.1 The MOBIUS Trac

The MOBIUS Trac is an integrated project management system, which includes a wiki for technical documentation, a Subversion server for source code (and other artifacts) control, and a ticket management system for defect reports and feature requests. It is integrated with Eclipse using the Mylyn plugin, as described in subsection 2.2.1, which is distributed as part of the Mobius PVE.

The MOBIUS Trac is accessed via the URL https://mobius.ucd.ie/.

4.3 The MOBIUS Version Control Repository

Like CDEs, a variety of version control repositories have been used over the past several years as well. Consolidation of these repositories is ongoing. The MOBIUS Subversion repository is hosted as part of the Trac discussed in the last section.

Consolidation and reorganization of the Subversion repository is still ongoing. For example, the jars and and libs directories are being combined and reorganized.

A Mobius PVE development mailing list called mobius-tools also exists. It is currently hosted at INRIA but will be moved to UCD as well for better integration with the MOBIUS Trac server.

This concludes this summary of the MOBIUS software development process and related servers and tools.

Chapter 5

Gamma Release Manifest

The “gamma” release of the Mobius PVE is based upon the Eclipse Europa (3.3.x) IDE for Java Developers.

5.1 Features and Plugins Included

The following MOBIUS plugins and features are included in the “gamma” release. Subsystems in bold italics are developed wholly by MOBIUS partners; subsystems in bold are developed partially by MOBIUS partners; subsystems in italics are maintained by MOBIUS partners.

- Coq Editor plugin 0.9.4
- JML2 Eclipse plugin 1.0.4
- ESC/Java2 feature 2.0.7
- Javafe plugin 2.0.8
- Simplify feature 1.5.6
- Prover Editor Feature 0.9.7
- Umbra plugin 1.0.2

and the following third-party plugins and features:

- BCEL feature 1.0.0
- Buckminster 0.2.0
- Checkstyle feature 4.3.3
- EclEmma Java Code Coverage feature 1.3.0
- FindBugs feature 1.2.1
- Metrics feature 1.3.6
- Mylyn feature 2.1.0

- PMD feature 3.2.4
- Subclipse feature 1.2.4

Additionally, the following subsystems are included in, or can be added to, the Mobius PVE, but not wrapped directly in a feature or a plugin:

- The Fx7 theorem prover from Michal Moskal
- Microsoft Research's Z3 theorem prover

5.2 Sample Workspace

The sample workspace contains the following projects:

- ESCJava2 plugin demo
- ETAPS Tutorial demo
- Examples

5.3 Platforms

The gamma release has been tested on the following platforms:

- Mac OS X 10.4 Tiger
- Mac OS X 10.5 Leopard
- Linux - Fedora Core
- Windows XP
- Windows Vista

Chapter 6

Publications Relating to the MOBIUS Program Verification Environment

The following publications relate directly to subsystems and functionality of the Mobius PVE, or case studies performed with the Mobius PVE, which are authored by MOBIUS partners.

- [BBC^{+}07a] G. Barthe, L. Burdy, J. Charles, B. Grégoire, M. Huisman, J.-L. Lanet, M. Pavlova, and A. Requet. JACK: a tool for validation of security and behaviour of Java applications. In Formal Methods for Components and Objects, Lecture Notes in Computer Science. Springer-Verlag, 2007. To appear.

- [CFJ^{+}06] M. Cielecki, J. Fulara, K. Jakubczyk, . Jancewicz, J. Chrzszcz, A. Schubert, and . Kamiski. Propagation of JML non-null annotations in Java programs. In Principles and Practices of Programming in Java, 2006.

- [?] P. Czarnik and A. Schubert. Extending operational semantics of the Java bytecode. In Trustworthy Global Computing, Lecture Notes in Computer Science. Springer-Verlag, 2007.

- [CHK05] P. Chalin, C. Hurlin, and J. R. Kiniry. Integrating static checking and interactive verification: Supporting multiple theories and provers in verification. In Verified Software: Tools, Technologies, and Experiences (VSTTE), Zurich, Switzerland, 2005.

- [GM07] R. Grigore and M. Moskal. Edit and verify. In Workshop on First-Order Theorem Proving, September 2007.

- [Jan07] M. Janota. Assertion-based loop invariant generation. In Workshop on Invariant Generation, Hagenberg, Austria, 2007. Workshop at CALCULEMUS 2007.

- [JGM07] M. Janota, R. Grigore, and M. Moskal. Reachability analysis for annotated code. In Specification and Verification of Component-Based Systems, a workshop at ESEC/FSE, 2007.

- [KMD06] J. R. Kiniry, A. Morkan, and B. Denby. Soundness and completeness warnings in ESC/Java2. In Specification and Verification of Component-Based Systems, 2006.

- [Kin07] J. R. Kiniry. Formally counting electronic votes (but still only trusting paper). In IEEE International Conference on Engineering of Complex Computer Systems, Auckland, New Zealand, 2007.

- [?] J.R. Kiniry and F. Fairmichael. Semantic properties for lightweight specification in knowledgeable development environments, 2007. In preparation.

- [KCT07] J.R. Kiniry, D. Cochran, and P. Tierney. A verification-centric realization of e-voting. In International Workshop on Electronic Voting Technologies, Boston, Massachusetts, 2007. Workshop at USENIX/ACCURATE 2007.

- [?] J.R. Kiniry, A. Morkan, D. Cochran, F. Fairmichael, P. Chalin, M. Oostdijk, and E. Hubbers. The KOA remote voting system: A summary of work to date. In Trustworthy Global Computing, Lucca, Italy, 2006.

- [KZ08] J.R. Kiniry and D.M. Zimmerman. Secret ninja formal methods. In Formal Methods, 2008. In press.

- [LM07] H. Lehner and P. Müller. Formal translation of bytecode into BoogiePL. In M. Huisman and F. Spoto, editors, Bytecode Semantics, Verification, Analysis and Transformation, Electronic Notes in Theoretical Computer Science, 2007.

- [MK07] M. Moskal, J. opuszaski, and J. Kiniry. E-matching for fun and profit. In The 5th International Workshop on Satisfiability Modulo Theories (SMT), Berlin, Germany, July 2007.

- [PS07] E. Poll and A. Schubert. Verifying an implementation of SSH. In Workshop on Issues in the Theory of Security, pages 164–177. IFIP WG1.7, 2007.

- [SC06] A. Schubert and J. Chrzszcz. ESC/Java2 as a tool to ensure security in the source code of Java applications. In Software Engineering Techniques: Design for Quality, IFIP, Warsaw, 2006. Springer-Verlag.

- [SCB$^+$08] A. Schubert, J. Chrzszcz, T. Batkiewicz, J. Paszek, and W. Ws. Technical aspects of class specification in the byte code of java language. In To be published in Bytecode'08 proceedings, Electronic Notes in Theoretical Computer Science. Elsevier, 2008.

Bibliography

[BBC$^+$07a] G. Barthe, L. Burdy, J. Charles, B. Grégoire, M. Huisman, J.-L. Lanet, M. Pavlova, and A. Requet. JACK: A tool for validation of security and behaviour of Java applications. In Formal Methods for Components and Objects: Revised Lectures from the 5th International Symposium FMCO 2006, number 4709 in Lecture Notes in Computer Science, pages 152–174. Springer-Verlag, 2007.

[BBC$^+$07b] Gilles Barthe, Lennart Beringer, Pierre Crégut, Benjamin Grégoire, Martin Hofmann, Peter Müller, Erik Poll, Germán Puebla, Ian Stark, and Eric Vétillard. MOBIUS: Mobility, ubiquity, security — objectives and progress report. In Trustworthy Global Computing: Revised Selected Papers from the Second Symposium TGC 2006, number 4661 in Lecture Notes in Computer Science. Springer-Verlag, 2007.

[BT07] C. Barrett and C. Tinelli. CVC3. In Computer Aided Verification, volume 4590 of Lecture Notes in Computer Science, pages 298–302. Springer, 2007.

[CFJ$^+$06] M. Cielecki, J. Fulara, K. Jakubczyk, . Jancewicz, J. Chrzszcz, A. Schubert, and . Kamiski. Propagation of JML non-null annotations in Java programs. In Principles and Practices of Programming in Java, 2006.

[CHK05] P. Chalin, C. Hurlin, and J. R. Kiniry. Integrating static checking and interactive verification: Supporting multiple theories and provers in verification. In Verified Software: Tools, Technologies, and Experiences (VSTTE), 2005.

[CJK07] P. Chalin, P. R. James, and G. Karabotsos. The architecture of JML4, a proposed integrated verification environment for JML. Technical report, Dependable Software Research Group, Concordia University, May 2007. ENCS-CSE-TR 2007-006.

[Con04] The JSR 133 Consortium. The Java memory model causality test cases, 2004. http://www.cs.umd.edu/˜pugh/java/memoryModel/CausalityTestCases.html.

[Cré06] Pierre Crégut. Midlet navigation graphs. Unpublished draft, 2006.

[Ecl] Eclipse website. www. See http://www.eclipse.org/.

[FF00] C. Flanagan and S. N. Freund. Type-based race detection for Java. In Programming Languages Design and Implementation, pages 219–232, New York, NY, USA, 2000. ACM Press.

[FL00] C. Flanagan and K.R.M. Leino. Houdini, an annotation assistant for ESC/Java. Technical Report 2000-003, DEC-SRC, December 2000.

[Fur86] G.W. Furnas. Generalized fisheye views. ACM SIGCHI Bulletin, 17(4):16–23, April 1986.

[GJSB05] J. Gosling, B. Joy, G. Steele, and G. Bracha. The Java Language Specification, Third Edition, chapter 17. Sun Microsystems, 2005. http://java.sun.com/docs/books/jls/.

[GM07] R. Grigore and M. Moskal. Edit and verify. In Workshop on First-Order Theorem Proving, September 2007.

[HM05] C. A. R. Hoare and J. Misra. Verified software: Theories, tools, experiments vision of a grand challenge project. A companion paper for VSTTE 2005, held in Züich, Switzerland., July 2005.

[HO03] E. Hubbers and M. Oostdijk. Generating JML specifications from UML state diagrams. In Forum on specification & Design Languages, pages 263–273. University of Frankfurt, 2003. Proceedings appeared as CD-Rom with ISSN 1636-9874.

[Ini06] Unified Testing Initiative. Unified testing criteria for Java technology-based applications for mobile devices. Technical report, Sun Microsystems, Motorola, Nokia, Siemens, Sony Ericsson, May 2006. Version 2.1.

[Jan07] M. Janota. Assertion-based loop invariant generation. In Workshop on Invariant Generation, 2007. Workshop at CALCULEMUS 2007.

[JGM07] M. Janota, R. Grigore, and M. Moskal. Reachability analysis for annotated code. In Specification and Verification of Component-Based Systems, Dubrovnik, Croatia, 2007. Workshop at ESEC/FSE 2007.

[KCT07] J.R. Kiniry, D. Cochran, and P. Tierney. A verification-centric realization of e-voting. In International Workshop on Electronic Voting Technologies, Boston, Massachusetts, 2007. Workshop at USENIX/ACCURATE 2007.

[Kin98] J. R. Kiniry. IDebug: An advanced debugging framework for Java. Technical Report CaltechCSTR:1998.cs-tr-98-16, California Institute of Technology, 1998.

[Kin04] J.R. Kiniry. Formalizing the user's context to support user interfaces for integrated mathematical environments. Electronic Notes in Theoretical Computer Science, 103, 2004.

[Kin07] J. R. Kiniry. Formally counting electronic votes (but still only trusting paper). In IEEE International Conference on Engineering of Complex Computer Systems, Auckland, New Zealand, 2007.

[KMD06] J. R. Kiniry, A. Morkan, and B. Denby. Soundness and completeness warnings in ESC/Java2. In Specification and Verification of Component-Based Systems, 2006.

[KO03] J.R. Kiniry and S. Owre. Improving the PVS user interface. In User Interfaces for Theorem Proving, September 2003.

[KZ08] J.R. Kiniry and D.M. Zimmerman. Secret ninja formal methods. In Formal Methods, 2008.

[LM07] H. Lehner and P. Müller. Formal translation of bytecode into BoogiePL. In M. Huisman and F. Spoto, editors, Bytecode Semantics, Verification, Analysis and Transformation, Electronic Notes in Theoretical Computer Science, 2007.

[LPC$^+$05] G. T. Leavens, E. Poll, C. Clifton, Y. Cheon, C. Ruby, D. Cok, and J. Kiniry. JML Reference Manual, July 2005. In Progress. Department of Computer Science, Iowa State University. Available from http://www.jmlspecs.org.

[mC06] MOBIUS Consortium. Deliverable 1.2: Framework-specific and application-specific security requirements, 2006. Available online from http://mobius.inria.fr.

[mC07a] MOBIUS Consortium. Deliverable 4.2: Certificates, 2007. Available online from http://mobius.inria.fr.

[mC07b] MOBIUS Consortium. Deliverable 4.3: Intermediate report on proof-transforming compiler, 2007. Available online from http://mobius.inria.fr.

[ML06] M. Moskal and J. Lopuszaski. Fast quantifier reasoning with lazy proof explication. http://nemerle.org/ malakith/smt/smt-tr-1.pdf, 2006. Available via http://nemerle.org/.

[MK07] M. Moskal, J. opuszaski, and J. Kiniry. E-matching for fun and profit. In The 5th International Workshop on Satisfiability Modulo Theories (SMT), Berlin, Germany, July 2007.

[MR07a] P. Müller and A. Rudich. Formalization of ownership transfer in universe types. Technical Report 556, ETH Zurich, 2007.

[MR07b] P. Müller and A. Rudich. Ownership transfer in universe types. In ACM Conference on Object-Oriented Programming Systems, Languages, and Applications, 2007. To appear.

[PBB$^+$04] M. Pavlova, G. Barthe, L. Burdy, M. Huisman, and J.-L. Lanet. Enforcing high-level security properties for applets. In P. Paradinas and J.-J. Quisquater, editors, Smart Card Research and Advanced Application. Kluwer Academic Publishing, 2004.

[Pic06] D. Pichardie. Bicolano – Byte Code Language in Coq. http://mobius.inria.fr/ bicolano. Summary appears in [?], 2006.

[PR08] E. Poll and J. N. Ravelo. Generalised navigation graphs, their refinement, and their use for midp applications, 2008. Manuscript. 48 pages.

[PS07] E. Poll and A. Schubert. Verifying an implementation of SSH. In Workshop on Issues in the Theory of Security, pages 164–177. IFIP WG1.7, 2007.

[SC06] A. Schubert and J. Chrzszcz. ESC/Java2 as a tool to ensure security in the source code of Java applications. In Software Engineering Techniques: Design for Quality, IFIP, Warsaw, 2006. Springer-Verlag.

[SCB+08] A. Schubert, J. Chrzszcz, T. Batkiewicz, J. Paszek, and W. Ws. Technical aspects of class specification in the byte code of java language. In To be published in Bytecode'08 proceedings, Electronic Notes in Theoretical Computer Science. Elsevier, 2008.

[win08] The Wine user guide. Available via http://www.winehq.org/, 2008.

[WN95] K. Waldén and J.-M. Nerson. Seamless Object-Oriented Software Architecture - Analysis and Design of Reliable Systems. Prentice Hall, 1995.