

Assessment of the Java Programming Language for Use in High Integrity Systems*

Jagun Kwon, Andy Wellings, and Steve King
Department of Computer Science
University of York, UK
{jagun, andy, king}@cs.york.ac.uk

ABSTRACT

This paper sets a goal of investigating the use of Java in the development of high integrity systems. Based on previous studies, guidelines, and standards, we develop 23 criteria that are used for the following assessment of Java. A summary of the assessment is provided before we go on to review a few existing subsets of the language.

1. Introduction

Increasingly computers are being used in high integrity real-time systems; that is, systems where failure can cause loss of life, environmental harm, or significant financial penalties. Examples include space shuttles, nuclear power plants and medical instruments, and they typically have high development and maintenance costs due to the customised nature of their components. There are many general and sector-specific standards produced to assist in building such important systems, for example, U.S. DO178B, and MISRA guidelines.

Within high integrity systems, there has been a growing trend to use software, because it provides [25, 26, 38, 10]

- improved functionality
- increased flexibility in design and implementation
- reduced production cost
- enhanced management of complexity in application areas.

Over the recent years, Java has proved to be an appropriate vehicle for a diverse range of applications including web based intranets and embedded systems. Its relatively simple linguistic semantics, the adoption of well-understood approaches to managing software complexity, and support for concurrency seem to have contributed towards its popularity. Initially designed with embedded systems in mind, Java's main goal was to provide engineers with a reliable and cost-effective platform-independent environment. The burden of learning a new language is kept to the minimum for many existing C and C++ programmers, while helping them to discover errors earlier by means of strong type checking, array-bound checking, null-pointer checking, and so on [20]. Further, its support for concurrency, i.e. multi-threading and synchronisation mechanisms, together with the use of portable code (or *bytecode*) opens up a huge number of possibilities for many other applications, including high integrity systems.

However, despite all these valuable features, Java has been criticised for its unpredictable performance as well as some security concerns [4, 5, 3]. The automatic garbage collection and dynamic class loading mechanisms are often considered problematic, especially under time or performance-critical situations. Moreover, a number of security bugs in the Java virtual machine have been discovered since its first appearance, especially in the bytecode verifiers and Just-in-Time (JIT) compilers [19, 4]. These fears make the use of Java and its associated technology problematic for the development of high integrity systems.

Upon the realisation of such primary drawbacks of Java, many researchers have attempted to improve the situation, particularly in search of predictable real-time performance. For instance, *the Real-Time Specification for Java* [9, 8] and its reference implementation [46] have proved that Java can be a capable framework for concurrent real-time applications. The specification attempts to minimise any modification to the original language semantics and defines many additional classes that should be implemented in a supporting virtual machine. This, however, ironically leads to a language and run-time system that are complex to implement and have high overheads at run-time. Software produced in

* This work has been funded by the EPSRC under award number GR/M94113.
This paper is to appear in the April/2003 issue of the ACM SIGPLAN *Notices*.

that framework is also difficult to analyse with all the complex features, such as the asynchronous transfer of control (ATC) and dynamic class loading.

Bearing in mind the positive developments and drawbacks of Java, this paper investigates the use of the language in the development of high integrity systems. Based on the requirements gathered from several relevant standards and guidelines, we develop 23 criteria that are used for the following assessment of Java. A summary of the assessment is provided before we move on to review a few existing subsets of the language. Finally, a brief conclusion is drawn.

2. Requirements of Programming Language

A study by Bentley [7] summarises some of the well-known requirements of programming language for the development of high integrity systems including works by [14], [16], [48], [49] and [21]. It carries out an assessment on Java against all the requirements, producing a series of rationales. A subset of the language is also proposed, but unfortunately only sequential features are included. The outcome of the study is compatible to a large extent with our objective in this paper as the requirements are still of significant importance these days, and the chosen language is Java. Therefore we consider it as our starting point for a more complete and up-to-date assessment of the language.

In addition to Bentley's work, we have also considered more recent guidelines and standards, such as the Ada95 Trustworthiness study [15, 42, 43], ISO/IEC DTR 15942 [22], software guidelines from the U.S. Nuclear Regulatory Commission (NRC) [37], and a series of reports produced by the Motor Industry Software Reliability Association (MISRA) [27, 28-35].

3. Assessment Criteria

Since many of the requirements are redundant and even ambiguous in some cases, we inclusively categorise them into relevant assessment criteria along with appropriate references. However, it is important to note that this collection of criteria is neither complete¹ nor authoritative, but it attempts to amalgamate many different requirements into a balanced and informative framework for the assessment of programming languages. As in [21] we propose two levels of criteria, namely *Mandatory requirements* (Level 1) and *Desirable requirements* (Level 2).

3.1. Level 1 – Mandatory requirements

In Level 1 we identify compulsory requirements that a programming language must satisfy in order to be considered for use in implementing high integrity software. Appropriate justifications are made regarding each requirement. Readers are encouraged to refer to the references if in any doubt about rationales and specifics.

3.1.1. Syntactical / Semantic Requirements

L.1.1.1	Type safety / Strong typing rules
References	[48], [16], [49], [21], [15], [22], [37], [33]
Rationale	Strongly typed languages help reduce errors in programs at compile-time. Moreover, type safety is often considered to be sufficient for ensuring the minimum nontrivial level of program safety, i.e. control flow safety, memory safety, and stack safety [24]. Thus it is strongly encouraged to use a type safe or strongly typed language, enhancing the integrity and security of software.
Specifics	Implicit type conversions must not be allowed. All data types should be statically analysable before program execution. Explicit type conversion rules should be clearly stated in the language standard or definition. There should be some ways to avoid access types or pointers.
Ratings	<ol style="list-style-type: none"> 1. Strongly typed / Statically analysable. 2. Strongly typed, but some types are analysable only at run-time, mainly due to the use of polymorphism in the language. 3. Not strongly typed and implicit type conversions are allowed.

¹ Some requirements or guidelines are deliberately missed out because they are either not relevant with respect to high integrity systems, or considered not reasonable in the context of modern programming languages. Examples include requirements on the use of a particular character set [USDoD1978], and improvements in wording or program presentation (of Ada83) [USDoD1990].

L.1.1.2	Side effects in expressions / Operator precedence levels / Initial values
References	[48], [37]
Rationale	Side effects in expressions can cause programs to behave in an ambiguous, or, possibly, unpredictable way, thus are not desirable. The precedence levels of all operators must be specified in the language definition; otherwise evaluation orders may vary from system to system.
Specifics	There should not be any time-dependent side effects in expressions. Operator precedence levels must clearly be defined in the standard. There should be no implicit initial values for variables.
Ratings	<ol style="list-style-type: none"> 1. All the above specifics are satisfied. 2. Not all the above specifics are satisfied, but there may be a subset of the language that meets the specifics. 3. The above specifics are not satisfied, and there is no reasonable way to improve the language.

L.1.1.3	Modularity / Structures
References	[48], [16], [21], [15], [37], [33]
Rationale	It must be straightforward to code and maintain programs in a high integrity programming language, so that the complexity of software becomes manageable. This is often achieved by means of visibility control (or scopes), functions, and objects in many modern languages, in which the integrity and security of software are generally improved.
Specifics	There should be sound mechanisms to structure and modularise program code both syntactically (in some form of determinable blocks or scopes) and semantically with clear interfaces. There should be no wild/unbounded jumps between different modules. Separate compilation of modules should be possible.
Ratings	<ol style="list-style-type: none"> 1. The language provides rich and precise means of structuring programs, and programs can be maintained in terms of modules or objects. 2. Such mechanisms are provided, but not cost-effective or efficient. 3. There is no reasonable approach.

L.1.1.4	Formal semantics / International standards
References	[48], [16], [49], [21], [15], [22], [33]
Rationale	A standardised language benefits the development of compilers and tools, and user training. Verification techniques can also be applied to a language with formally defined semantics.
Specifics	There should be a (international) standard definition of the language. There should be formally defined semantics of the language, or at least a subset of the language.
Ratings	<ol style="list-style-type: none"> 1. An internationally standardised formal definition exists. 2. The language or high integrity subset of it can be formally defined. 3. Unknown.

L.1.1.5	Well-understood
References	[16], [48], [21], [49], [33]
Rationale	A language with well-understood semantics and syntaxes will help to produce quality software, often cost-effectively.
Specifics	The language should be simple, well understood, easy to adopt, and easy to implement.
Ratings	<ol style="list-style-type: none"> 1. The language is well understood, and there are many trained developers and designers. 2. The language is well understood only by a limited number of people. 3. Unknown.

L.1.1.6	Support for domain specific or embedded applications
References	[48], [21]
Rationale	High integrity systems are often embedded systems that need to interface or control physical resources or (non-standard) peripheral devices. Therefore, a programming

	language designed with such applications in mind should be used.
Specifics	Robust mechanisms for controlling memory, I/O devices or other hardware are required.
Ratings	<ol style="list-style-type: none"> 1. The language naturally supports embedded applications. 2. There is a limited support, but external libraries or language extensions can be utilised. 3. No support provided or Unknown.

L.1.1.7	Concurrency / Parallel processing
References	[48], [21]
Rationale	Although concurrency is one of the main sources of complication in program analysis and verification (classified as only a desirable – not mandatory - feature in [21]), it is invaluable in modelling or capturing real-world problems. Thus, we believe this has to be an essential requirement for modern high integrity language.
Specifics	The following features should be included: Language-level support for multitasking or multithreading. Control over scheduling policy. Straightforward communication and synchronisation mechanism(s), plus facility to bound blocking.
Ratings	<ol style="list-style-type: none"> 1. All the above specifics are satisfied. 2. Only limited support is provided at the language-level, but external libraries or run-time systems can be utilised. 3. No reasonable support provided or Unknown.

L1.2. Application of verification techniques / Predictability

L.1.2.1	Functional predictability
References	[21], [15], [22], [37], [33]
Rationale	High integrity software must be proven to be predictable in terms of its functional behaviours.
Specifics	All or most of the following analysis techniques should be applicable. Control flow analysis Data flow analysis Information flow analysis Symbolic execution Formal code verification
Ratings	<ol style="list-style-type: none"> 1. All techniques in the above specifics or feasible alternatives can be utilised. 2. Not all techniques can be utilised due to the complex features of the language, but sub-setting the language may improve such analyses. 3. Unknown or there is no cost-effective way of utilising such analysis techniques.

L.1.2.2	Temporal predictability / Timing analysis
References	[21], [15], [37], [33]
Rationale	In addition to the functional predictability, timely behaviours of such software and systems must also be guaranteed.
Specifics	Worst Case Execution Time (WCET) of each process must be obtainable, so that schedulability analysis can be performed.
Ratings	<ol style="list-style-type: none"> 1. Tightly bounded execution time(s) can be obtained. 2. Loosely bounded execution time(s) can be obtained. 3. Unpredictable or there is no known way to obtain WCET.

L.1.2.3	Resource usage analysis
References	[16], [21], [15], [37], [33]
Rationale	It is important to identify what resources are needed and how they are utilised, so that errors such as stack overflow may not occur, and system implementations may be kept economical.
Specifics	The following properties should be analysable. Memory (or heap) usage Stack usage

	Any other resources to be utilised in the application area.
Ratings	<ol style="list-style-type: none"> 1. Exact prediction of the above specifics is possible. 2. Worst-case analysis is possible, but not practical. 3. Unpredictable.

L1.3. Language Processors / Run-time environment / Tools

L.1.3.1	Certified language translators / Run-time environments
References	[48], [21], [22], [33]
Rationale	There must be a high level of assurance in language processors, especially compilers.
Specifics	A formally certified compiler by an authoritative or trusted body should be used. Low-level code should be traceable in accordance with source code. Run-time environments should also be certified if used.
Ratings	<ol style="list-style-type: none"> 1. There are one or more certified language translators, and they are formally proven to be flawless. 2. Language translators may contain several known errors or malfunctions that are well documented, but they will not affect the development of high integrity software. 3. Unknown.

L.1.3.2	Run-time support / Environment issues
References	[48], [21], [49], [15], [22], [37]
Rationale	Libraries (or any additional code) or run-time support may make it complex to perform some analyses, such as WCET and control flow analyses. Hence, all such additional code should be predictable and analysable in terms of safety and timeliness. Minimising implementation dependencies is also encouraged.
Specifics	All the behaviours of additional code should be well understood. All timing information of the underlying run-time system and libraries should be known and accurate.
Ratings	<ol style="list-style-type: none"> 1. There is concrete information on the functional and temporal behaviours of all libraries and run-time system. 2. Only worst-case analysis is possible. 3. Unknown.

3.2. Level 2 – Desirable Requirements

The requirements at this level are not immediately necessary but beneficial in that they help produce more efficient, comprehensible, and structured systems. Note that ratings are not provided at this level because they are meant to be additional desirables.

L2.1. Syntactical / Semantic Requirements

L.2.1.1	Exception Handling / Failure behaviour
References	[48], [16], [21], [15], [37], [33]
Rationale	Handling errors while a high integrity system is operating is sometimes seen as undesirable on account of additional overheads and unpredictable behaviours. However, if any sort of error can occur, then the system should gracefully degrade, or recover after some corrections.
Specifics	Robust and analysable run-time error detection and handling mechanism should exist. Failure behaviours should be programmable.

L.2.1.2	Model of Mathematics
References	[16], [48]
Rationale	As often required in some high integrity systems, the language should have a rigorous model of maths defined in the language standard.
Specifics	A model of both integer and floating point arithmetic should be defined within the language standard. Procedures for checking if operational arithmetic at run-time is correct should exist.

L.2.1.3	Support for User documentation
References	[48], [49], [37]
Rationale	Languages that allow user comments will undoubtedly improve program readability and maintainability. Some language processors may make use of annotations to detect subtle logical errors in programs or to obtain extra information.
Specifics	There should be some way of commenting programmer's intentions within source code.

L.2.1.4	Support for a range of static types including subtypes and enumeration types
References	[48], [49], [22], [33]
Rationale	It is easier to perform any analyses or checks on static types than on dynamic types. Enumeration types with a limited number of values also help reduce errors.
Specifics	None.

L.2.1.5	Coding style guidelines
References	[21], [37], [33]
Rationale	Coding style guidelines may help reduce the gap between well-established Software Engineering principles and the actual practice of programming in a particular language.
Specifics	None.

L.2.1.6	Support for abstraction and information hiding
References	[22], [21], [49], [33]
Rationale	Employing abstraction or information hiding techniques (e.g. object orientation) can greatly decrease software complexity. Thus they are beneficial in program design, development, and maintenance.
Specifics	None.

L.2.1.7	Assertion checking
References	[48]
Rationale	It may sometimes be desirable to check for user specified assertions before or while programs are executing.
Specifics	None.

L2.2. Language Processors / Run-time environment / Tools

L.2.2.1	Certified (static/dynamic) analysis tools
References	[21], [22], [33]
Rationale	In order to gain more confidence in high integrity software it is imperative to use certified analysis tools, which may check for errors, such as, race conditions and deadlocks.
Specifics	None.

L.2.2.2	Interface to other languages
References	[48]
Rationale	There are some situations where a program written in a high-level language needs to interact with existing libraries or other low-level routines that are written in different languages. In such cases there should be a means of interfacing our program with such routines.
Specifics	None.

L.2.2.3	Code optimisation
References	[48]
Rationale	It is always advantageous to improve the efficiency of programs by means of optimising them. However, optimisation should not alter the semantics of correct programs, nor compromise the application of analysis techniques.
Specifics	None.

L.2.2.4	Code portability
References	[21], [37]
Rationale	Since there exists a diverse range of code-executing platforms, it is often considered beneficial to have a portable program representation, so that all necessary analyses may be applied once for all.
Specifics	None.

4. Assessment of the Java programming language

In this section, Java is assessed against the criteria developed earlier. A summary is provided at the end.

4.1. Assessment of Java against Level 1

L.1.1. Syntactical / Semantic Requirements

L.1.1.1. Type safety / Strong typing rules

Java is a strongly typed language. For all primitive types, implicit type conversions are not allowed (all possible conversions are stated in the language specification), and programs are analysable before running them. But for dynamic reference types, it is not always straightforward to statically analyse code, but is generally possible only at run-time because of the use of, for example, inherited interfaces and local classes within different scopes.

Rating: 2. Strongly typed, but some types are analysable only at run-time, mainly due to the use of polymorphism in the language.

L.1.1.2. Side effects in expressions / Operator precedence levels / Initial values

Side effects can occur in Java if expressions contain embedded assignments, sub-operators, and method invocations. Many side effects, however, can be eliminated via the use of a code checker or analyser, and a subset of Java. Operator precedence levels are defined in the specification [20], but the large number is at times seen undesirable as it becomes more difficult for programmers to learn [7, 48]. All types in Java have default initial values, but compilers issue warnings if any variables are used before initialisation. It should also be noted that some returned values of a method can be *quietly discarded* without any warning [20], i.e. when there is no assignment expression for a method call that returns a value.

Rating: 2. Not all the above specifics are satisfied, but there may be a subset of the language that meets the specifics.

L.1.1.3. Modularity / Structures

In Java, programs are organised as objects that normally consists of visible and non-visible data fields and methods. Abstraction and encapsulation mechanisms are also provided through classes and interfaces, and packages (into which related classes are organised) also enhance modularity and structure of software. In addition to this, the language includes various means of controlling program flows, including the exception-handling mechanism. Separate compilation is always possible.

Rating: 1. The language provides rich and precise means of structuring programs, and programs can be maintained in terms of modules or objects.

L.1.1.4. Formal semantics / International standards

There are no stable standards for Java although the language specification [20] serves as an informal standard for the time being. There are some formal semantics of Java, for example, in Action semantics [50, 12], in Denotational Semantics [2], and in other BNF-like notations [1], most of which are based on parts of the language. Drossopoulou and Eisenbach [18] have also defined a series of subsets of Java and proved their type soundness.

Rating: 2. The language or a high integrity subset of it can be formally defined.

L.1.1.5. Well-understood

Java is a familiar programming language to many existing C/C++ programmers, which means that no extensive training is usually required and there may well be many trained engineers. In addition, some of the problematic features in C/C++ (such as pointer operations) are removed, which all results in a dramatic increase in productivity. However, the excessive number of APIs and other additional mechanisms can be hard to master.

Rating: 1. The language is well understood, and there are many trained developers and designers.

L.1.1.6. Support for domain specific or embedded applications

One of the main application areas for which Java was first developed was embedded systems. In pure Java, however, it is not possible to control underlying hardware without appropriate native methods implemented in different languages. Even then, it is still difficult to implement systems with rigorous safety and real-time requirements, thanks mostly to the overheads incurred by the garbage collection mechanism, and virtual machines *per se*. There has been much research on scheduling the garbage collector and improving the efficiency of code transformation, even though it has not proven particularly effective so far. In the recent years, the Real-Time Specification for Java [8] and Real-Time Core Extensions [23] have been defined, so that real-time applications will certainly benefit from reference implementations of such specifications.

Rating: 2. There is a limited support, but external libraries or language extensions can be utilised.

L.1.1.7. Concurrency / Parallel processing

Java supports concurrent execution of multiple threads, as well as some key synchronisation mechanisms, for example, the monitor and synchronized blocks/methods. Programmers can also allocate a priority to threads, which nevertheless is not of any significant value, as they have no control over scheduling mechanisms implemented in the virtual machine and underlying kernel. Recently, two of the specifications for real-time Java, i.e. one from Sun Microsystems [8] and the other J Consortium [23], define various features that real-time systems require, especially with regard to scheduling, memory management, synchronisation, time, and exceptions.

Rating: 1. All the above specifics are satisfied.

L.1.2. Application of verification techniques / Predictability

L.1.2.1. Functional predictability

Due to the recent development of sophisticated analysis algorithms and tools it is now possible, to some extent, to analyse Java programs in terms of control and data flow. Nevertheless, some complex features of Java, such as the exception handling mechanism and monitors, are still not considered, or at least are immaturely handled. Formal verification is even harder for Java as there is no complete formal semantics. However, a constant progress is made in this area, and especially *Model-checking* technology is proving strong in the verification of Java programs. For example, the Java Pathfinder 2 [11] developed by NASA can detect race conditions, deadlocks, and violations of user-specified assertions.

Rating: 2. Not all techniques can be utilised due to the complex features of the language, but sub-setting the language may improve such analyses.

L.1.2.2. Temporal predictability / Timing analysis

It is well known that with all the sometimes-superfluous features like the garbage collector and virtual machine support, it is hard to obtain tight execution-time bounds for Java threads, and such timing analyses are all dependent on eventual target architectures and base operating systems (if utilised). Some techniques, however, have been suggested (e.g. [6, 41]), and the release of the specifications for real-time Java [8] will certainly improve the current situation.

Rating: 2. Loosely bounded execution time(s) can be obtained.

L.1.2.3. Resource usage analysis

On account of the presence of the background garbage collector, it is generally difficult to predict how much memory space will be in use at a given moment in time, or even deducing the worst case can become impractical (and dependent on which garbage collection algorithms are employed). However, subsets of Java or of the Real-Time Specification for Java [8], such as [40] in which garbage collection is excluded, will ease this sort of analysis.

Rating: 2. Worst-case analysis is possible, but not practical.

L.1.3. Language Processors / Run-time environment / Tools

L.1.3.1. Certified language translators / Run-time environments

To the best of our knowledge, Java compiler and virtual machine validation is still an on-going research work. Whereas it may never be possible to formally exploit and validate such complex software, some attempts have been made to conduct conformity assessment of Java or Java-like language processors to the language specification and industry standards, for example see [39]. Reported errors are reasonably well documented and updated.

Rating: 2. Language translators may contain several known errors or malfunctions that are well documented, but they will not affect the development of high integrity software.

L.1.3.2. Run-time support / Environment issues

It is not easy to perform analyses on additional code, i.e. that of variable run-time systems, APIs, native methods, unless a sound standard for such program entities is developed.

Rating: 3. Unknown.

4.2. Assessment of Java against Level 2

L.2.1. Syntactical / Semantic Requirements

L.2.1.1. Exception Handling / Failure behaviour

Java has a wide variety of predefined exception classes, and programmers are also allowed to define customised (checked) exceptions and program's behaviours. Uncaught exceptions, i.e. unchecked exceptions or errors, can become problematic as they may simply result in the system halting.

L.2.1.2. Model of Mathematics

Java provides a rich set of integer and floating point data types, and the `java.math` package can be used to assist in more rigorous mathematical applications. While the utilisation of the standard IEEE 754 arithmetic semantics is seen as universally beneficial in terms of compatibility, it is occasionally not desirable as it hinders the utilisation of advanced hardware, for example, built-in co-processors [7].

L.2.1.3. Support for User documentation

Java provides two ways of commenting source code. Furthermore, there is a facility for automatically generating on-line documentation of user classes, i.e. *javadoc* tool.

L.2.1.4. Support for a range of static types including subtypes and enumeration types

Subtypes and enumeration types are not supported in Java, but may possibly be emulated with additional overheads.

L.2.1.5. Coding style guidelines

There are coding style documents available at the web site of Sun Microsystems [45]. However, none of them specifically addresses high integrity or real-time applications.

L.2.1.6. Support for abstraction and information hiding

As an object oriented language, Java offers abstraction by means of the abstract class type and interface, where no implementation details are allowed. Information hiding is also naturally supported.

L.2.1.7. Assertion checking

There was not any specific language construct for assertion checking in the previous releases of Java. However, in the new release of Java, i.e. J2SE 1.4, a simple assertion facility is provided, so that programs can be checked against assertions. See <http://www.jcp.org/jsr/detail/41.jsp> for more information.

L.2.2. Language Processors / Run-time environment / Tools

L.2.2.1. Certified (static/dynamic) analysis tools

A large number of analysis tools have been developed to assist in debugging Java programs, but most of them are not certified by reliable bodies or standards. However, as mentioned above, tools such as Java PathFinder 2 (from NASA) and the Extended Static Checker for Java (from Compaq) appear to be successful in detecting many known errors.

L.2.2.2. Interface to other languages

Java cannot directly interface to programs written in other languages. But, it is possible to invoke native methods, mostly written in C, of the run-time environment. This will result in poor portability.

L.2.2.3. Code optimisation

Most of the available optimisation techniques are not applied until Java programs reach their target or virtual machine for security reasons. Different quality of code or performance may be generated depending on how code is processed, i.e. bytecode can be interpreted, compiled *Just-in-Time*, or compiled *Ahead-of-Time*. It is complex to statically analyse optimised native code in relation to high-level bytecode. Optimisation will also make the complexity of compiler and tool validation more difficult.

L.2.2.4. Code portability

Following the “*write once and run everywhere*” motto, Java has become a truly portable programming language for most of the well-known platforms. In addition, Java chips with an integrated virtual machine and processor also start to appear. However, a problem can arise when non-standard processors or operating systems are utilised, where the burden of developing a new virtual machine is left to the system developer.

4.3. Summary of Assessment

Most of the Level 1 criteria are not, or loosely met by Java. Below is a summarising classification of the strengths and weaknesses identified above.

Strengths

Java is a strongly typed object-oriented language that provides an excellent means of modularising and structuring programs (L.1.1.3), and is well understood (L.1.1.5). It also supports concurrent execution of multiple threads as well as some key synchronisation mechanisms (L.1.1.7).

Weaknesses

Reference types are not generally amenable to static checking (L.1.1.1). Furthermore, side effects can occur in expressions, and some returned values may be quietly discarded (L.1.1.2). There are several formal definitions and semantics of Java, but they are predominantly concerned with parts of the language (L.1.1.4). Embedded applications, in which hardware control is essential, can only be supported by means of native methods (L.1.1.6), although implementations of the specifications for Real-Time Java are expected to solve this problem.

It is not straightforward to apply various analysis techniques directly to Java due to some of its complex features (L.1.2.1), but there appear to be some evolving analysis tools. Timing analysis is also difficult to perform on Java code (L.1.2.2), as is resource usage analysis (L.1.2.3).

There is no formally validated Java compiler and virtual machine, but conformity-checking tools do exist (L.1.3.1). It is also complex to perform any analyses on additional code, such as that of APIs and run-time systems (L.1.3.2).

Regarding Level 2 requirements, the following strengths and weaknesses have been identified.

Strengths

- Integrated exception handling mechanism (L.2.1.1)
- Rich set of integer and floating-point data types, and java.math package (L.2.1.2)
- Support for user documentation (L.2.1.3)
- General coding style guidelines (L.2.1.5)
- Support for abstraction and information hiding (L.2.1.6)
- Code portability (L.2.2.4)

Weaknesses

- Overhead and complexity of exception handling mechanism (L.2.1.1)
- The utilisation of the standard IEEE 754 arithmetic semantics can be overhead, and no exception is generated for particular operations (L.2.1.2)
- No support for subtypes and enumeration types (L.2.1.4)
- No coding guidelines for high integrity applications (L.2.1.5)
- No assertion checking facility (L.2.1.8)
- Shortage of certified analysis tools (L.2.2.1)
- Difficulty in interfacing to programs written in other languages (L.2.2.2)
- Complexity of analysing optimised code (L.2.2.3)

5. Review of Subsets

Burns et al. [13] suggest that a restricted programming model or profile can help produce efficient and predictable systems by removing language features with high overheads, and complex and erroneous semantics. Along these lines, there have been a few subsets or profiles for Java suggested in the literature².

² In fact, there are subsets of Java defined for other purposes than for use in high integrity systems. For example, in [Drossopoulou+1999] the authors define a series of subsets in order to prove the type soundness of them.

Bentley [7] defines a sequential subset of Java after assessing the language. The subset consists of 21 rules that are effectively derived from [21], [36] and his assessment. All the rules are categorised into six groups, most of which are concerned with the use of some problematic features of the language itself. However, while this subset will undoubtedly help produce analysable and predictable *sequential* programs, it can be criticised for its restriction on multithreading, one of Java's inherent elements. Without the language-level support for multithreading and all the associated synchronisation mechanisms, Java may not be considered as a great evolution from its predecessors. In addition to this, the subset also fails to address issues on the object-oriented programming model of the language, as well as real-time issues.

Puschner and Wellings [40] suggest a *Ravenscar*-like profile for the Real-Time Specification for Java [8], and it is in fact a predecessor of the work presented in this paper. The profile is primarily focused on leaving out complex features of the RTSJ. However, little attention is paid to Java's sequential language constructs (unlike [7]) and object-orientation features that can be problematic in performing various static analyses. Furthermore, the profile is not consistent with the current version of the RTSJ. Kwon et.al. [52], however, extend Puschner's work to produce a more complete and up-to-date version of the profile. It is based on important guidelines and standards, such as the software guidelines from the U.S. Nuclear Regulatory Commission [37].

A sub-committee has been formed within the Real-Time Java Working Group of the J Consortium to produce a high integrity profile based on the Real-Time Core Extensions [23]. The profile has not publicly been released yet, but according to Dobbing [17] it will resemble the Ravenscar profile for Ada95 [13]. It consists of four main themes: partitioning, memory management, concurrency, and error recovery, respectively. Like the one proposed in [40], this profile is mainly focused on sub-setting the Real-Time Core Extensions [23], but does not address issues on the use of problematic language constructs and object-orientation features of Java.

6. Conclusions

We have reviewed important requirements of programming language for the development of high integrity software, and defined 23 assessment criteria derived from the requirements. The criteria are divided into two groups, namely, *Mandatory requirements* (Level 1) and *Desirable requirements* (Level 2). Appropriate references and rationale for each criterion are given, and suitable ratings are also provided for the Level 1 requirements.

The Java programming language and its associated environments are then assessed against the two levels of criteria, and we conclude that Java is a good general language, yet not appropriate as a whole for the development of high integrity systems that require rigorous and predictable language features, compilation systems, and tools. However, Java may be able to qualify as a suitable vehicle in the future with the help of sub-setting the language (e.g. [52]) and future developments of formal mechanisms, although most of the currently proposed subsets do not address all the necessary areas required for high-integrity real-time systems. There is perhaps some movement towards a standardisation through the Java 2 Platform Micro Edition (J2ME) that introduces profiles for resource constrained mobile devices. One could devise a profile for high-integrity real-time systems.

7. References

- [1] J. Alves-Foss and D. Frincke, *Formal Grammar for Java*, in LNCS 1523 Formal Syntax and semantics of Java (ed. J. Alves-Foss), Springer-Verlag, Berlin, 1999.
- [2] J. Alves-Foss and F. S. Lam, *Dynamic Denotational Semantics of Java*, in LNCS 1523 Formal Syntax and semantics of Java (ed. J. Alves-Foss), Springer-Verlag, Berlin, 1999.
- [3] W. Amme, N. Dalton, M. Franz, and J. Von Ronne, *SafeTSA: A Type Safe and Referentially Secure Mobile-Code Representation Based on Static Single Assignment Form*, Accepted for the 2001 ACM SIGPLAN Conference on Programming Language Design and Implementation 2001.
- [4] Andrew W. Appel, *Protection against untrusted code: The JIT compiler security hole, and what you can do about it*, <http://www-106.ibm.com/developerworks/library/untrusted-code/>, as of January 2001.
- [5] A. Azevedo, A. Nicolau, and J. Hummel, *Java Annotation-Aware Just-In-Time (AJIT) Compilation System*, ACM 1999 Java Grande Conference, 1999.
- [6] I. Bate, G. Bernat, G. Murphy, P. Puschner, *Low-level analysis of a portable WCET analysis framework*, 6th IEEE Real-Time Computing Systems and Applications (RTCSA), 2000.
- [7] S. Bentley, *The Utilisation of the Java Language in Safety Critical System Development*, MSc dissertation, Department of Computer Science, University of York, 1999.
- [8] G. Bollella, et al, *The Real-Time Specification for Java*, Addison-Wesley, 2000.

- [9] G. Bollella and J. Gosling, *The Real-Time Specification for Java*, IEEE Computer, Vol. 33, No. 6, June 2000.
- [10] J. P. Bowen and M. G. Hinchey, *High-Integrity System Specification and Design*, Springer-Verlag London, 1999.
- [11] G. Brat, K. Havelund, S. Park, and W. Visser, *Java PathFinder – Second Generation of a Java Model Checker*, In Proceedings of Post-CAV Workshop on Advances in Verification, Chicago, July 2000.
- [12] D. F. Brown and D. A. Watt, *JAS: a Java Action Semantics*, in Proc. of 2nd International Workshop on Action Semantics (ed. Mosses, P.D., and Watt, D.A.), BRICS NS-99-3, University of Aarhus, Denmark, 1999.
- [13] A. Burns, B. Dobbing, and G. Romanski, *The Ravenscar Tasking Profile for High Integrity Real-Time Programs*, In L. Asplund, editor, Proceedings of Ada-Europe 98, LNCS, Vol. 1411, pages 263-275, Berlin Heidelberg, Germany, Springer-Verlag 1998.
- [14] B. A. Carré, T. J. Jennings, F. J. Maclennan, P. F. Farrow, and J. R. Garnsworthy, *SPARK – The SPADE Ada Kernel*, 3rd ed, Program Validation Limited, 1990.
- [15] D. Craigen, M. Saaltink and S. Michell, *Ada 95 Trustworthiness Study; A Framework for Analysis*, ORA Canada, 29 November 1995.
- [16] W. J. Cullyer, S. J. Goodenough, and B. A. Wichmann, *The Choice of Computer Languages for use in Safety-Critical Systems*, Software Engineering Journal, March 1991.
- [17] B. Dobbing, *The Ravenscar Profile for High-Integrity Java Programs?*, ACM Ada Letters, Vol. 21, Issue. 1, March 2001.
- [18] S. Drossopoulou and S. Eisenbach, *Describing the Semantics of Java and Proving Type Soundness*, in LNCS 1523 Formal Syntax and semantics of Java (ed. J. Alves-Foss), Springer-Verlag, Berlin, 1999.
- [19] Li Gong, *Inside Java™ 2 Platform Security: Architecture, API Design, and Implementation*, Addison-Wesley, 1999.
- [20] J. Gosling, B. Joy, G. Steele, and G. Bracha, *The Java Language Specification*, 2nd Edition, Addison Wesley, 2000.
- [21] A. Hutcheon, B. Jepson, D. Jordan, and I. Wand, *A Study of High Integrity Ada: Language Review*, Technical Report SLS31c/73-1-D, Version 2, York Software Engineering, University of York, July 1992.
- [22] *Programming Languages – Guide for the Use of the Ada Programming Language in High Integrity Systems*, ISO/IEC DTR 15942, ISO/IEC WG9, 1999.
- [23] J Consortium, *International J Consortium Specification: Real-Time Core Extensions*, Revision 1.0.14, www.j-consortium.org, September 2000.
- [24] D. Kozen, *Language Based Security*, Technical Report TR99-1751, Cornell University, 1999.
- [25] N. G. Leveson, *Software Safety: Why, What, and How*, Computing Surveys, Vol. 18, No. 2, ACM, June 1986.
- [26] N. G. Leveson, *Software Safety in Embedded Computer Systems*, Communications of the ACM, Vol. 34, No. 2, February 1991.
- [27] The Motor Industry Software Reliability Association, *Development Guidelines for Vehicle Based Software*, ISBN 0952415607, MIRA Ltd., November 1994.
- [28] The Motor Industry Software Reliability Association, *Report 1: Diagnostics and Integrated Vehicle Systems*, MIRA Ltd., February 1995.
- [29] The Motor Industry Software Reliability Association, *Report 2: Integrity*, MIRA Ltd., February 1995.
- [30] The Motor Industry Software Reliability Association, *Report 3: Noise, EMC and Real-Time*, MIRA Ltd., February 1995.
- [31] The Motor Industry Software Reliability Association, *Report 4: Software in Control Systems*, MIRA Ltd., February 1995.
- [32] The Motor Industry Software Reliability Association, *Report 5: Software Metrics*, MIRA Ltd., February 1995.
- [33] The Motor Industry Software Reliability Association, *Report 6: Verification and Validation*, MIRA Ltd., February 1995.
- [34] The Motor Industry Software Reliability Association, *Report 7: Subcontracting of Automotive Software*, MIRA Ltd., February 1995.
- [35] The Motor Industry Software Reliability Association, *Report 8: Human Factors in Software Development*, MIRA Ltd., February 1995.
- [36] The Motor Industry Software Reliability Association, *Guidelines for the use of the C language in vehicle based software*, The Motor Industry Research Association, 1998.

- [37] H. Hetcht, M. Hecht, S. Graff, et al, *Review Guidelines for Software Languages for Use in Nuclear Power Plant Systems*, NUREG/CR-6463, U.S. Nuclear Regulatory Commission, 1997, also available at <http://fermi.sohar.com/J1030/index.htm>, last accessed in January 2002.
- [38] D. L. Parnas, A. J. van Schouwen, and S. P. Kwan, *Evaluation of Safety-Critical Software*, Communications of the ACM, Vol. 33, No. 6, June 1990.
- [39] *JETS, A Perennial Validation Suite for the Java language*, http://www.peren.com/pages/jets_set.htm, last accessed in December 2001.
- [40] P. Puschner and A. J. Wellings, *A Profile for High-Integrity Real-Time Java Programs*, Proceedings of ISORC 2001.
- [41] P. Puschner, G. Bernat, *WCET Analysis of Reusable Portable Code*, Proceedings of the 13th Euromicro International Conference on Real-Time Systems, 2001.
- [42] M. Saaltink and S. Michell, *Ada 95 Trustworthiness Study; Analysis of Ada 95 for Critical Systems*, V2.0, ORA Canada, 27 March 1997.
- [43] M. Saaltink and S. Michell, *Ada 95 Trustworthiness Study; Guidance on the Use of Ada95 in the Development of High Integrity Systems*, V2.0, ORA Canada, 27 March 1997.
- [44] N. Storey, *Safety-Critical Computer Systems*, Addison Wesley Longman 1996.
- [45] Sun Microsystems, *Code Conventions for the Java Programming Language*, available at <http://java.sun.com/docs/codeconv/html/CodeConvTOC.doc.html>, written in April 1999, last accessed in December 2001.
- [46] TimeSys™, *Products and Services: Real-Time Java*, available at <http://www.timesys.com/rtj/index.html>, last accessed in January 2002.
- [47] U.K. Ministry of Defence, *The Procurement of Safety Critical Software in Defence Equipment*, INTERIM Defence Standard 00-55 (PART 1: REQUIREMENTS)/Issue 1, 5 April 1991.
- [48] U.S. Department of Defence, *Requirements for High Order Computer Programming Languages "STEELMAN"*, U.S. Department of Defence, 1978.
- [49] U.S. Department of Defence, *Ada 9X Requirements*, Office of the Under Secretary for Defence Applications, Washington, D.C., December 1990.
- [50] D. A. Watt and D. F. Brown, *Formalising the Dynamic Semantics of Java*, In Proceedings of the Third International Workshop on Action Semantics (AS2000), Recife, Brazil, May 2000.
- [51] D. A. Wheeler, *Ada, C, C++, and Java vs. the Steelman*, ACM Ada Letters, Vol. 17, Issue 4, July 1997.
- [52] J. Kwon, A. Wellings and S. King, *Ravenscar-Java Profile: A High Integrity Profile for Real-Time Java*, Proceedings of the ACM Java Grand-ISCOPE conference 2002, Seattle, WA, USA, November 2002