

Software certification – coding, code, and coders

Klaus Havelund and Gerard J. Holzmann
Laboratory for Reliable Software (LaRS)
Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive, Pasadena, California, 91109-8099
firstname.lastname@jpl.nasa.gov

ABSTRACT

We describe a certification approach for software development that has been adopted at our organization. JPL develops robotic spacecraft for the exploration of the solar system. The flight software that controls these spacecraft is considered to be *mission critical*. We argue that the goal of a software certification process cannot be the development of “*perfect*” software, i.e., software that can be formally proven to be correct under all imaginable and unimaginable circumstances. More realistically, the goal is to guarantee a software development process that is conducted by knowledgeable engineers, who follow generally accepted procedures to control known risks, while meeting agreed upon standards of workmanship. We target three specific issues that must be addressed in such a certification procedure: the coding process, the code that is developed, and the skills of the coders. The coding process is driven by standards (e.g., a coding standard) and tools. The code is mechanically checked against the standard with the help of state-of-the-art static source code analyzers. The coders, finally, are certified in on-site training courses that include formal exams.

Categories and Subject Descriptors

D.2.0 [Software Engineering]: General – standards.K.5.2. [Governmental Issues]: – Regulation. K.7.3: [The Computing Profession]: Testing, Certification, and Licensing.

General Terms

Design, Reliability, Standardization, Legal Aspects.

Keywords

Coding standards, code review, static source code analysis, logic model checking, runtime verification, safety-critical software.

1. INTRODUCTION

John Rushby once described the dilemma of current approaches to software verification or certification as follows: “*Because we cannot demonstrate how well we’ve done, we’ll show how hard we’ve tried.*” [1] The statement is apt. Few, if any, organizations feel confident enough about their software development processes that they are willing to give an absolute guarantee of its fitness for use or so much as the absence of preventable flaws in workmanship. As customers, and generally as users that have to rely on the safety and reliability of sometimes critically important software applications (e.g., as used by banks, car makers, or in medical devices), we are quite used to the opposite: we are routinely asked to sign disclaimers that hold the software makers invulnerable to flaws in workmanship and all possible damage that might be caused by it. This in itself is remarkable.

At some point, perhaps a few decades ago, we may have expected that standard market principles would solve this problem:

customers could have been expected to reject products that are delivered without warranties of fitness. But this is not what happened.

The driving principles that determine how software applications are developed and marketed give a significant advantage to the vendor who delivers a new service first, and merely commits to slowly improve while the product is in use, based on customer feedback. The customers, in this way, become part of what otherwise would be the testers, except this group of testers pays the vendor, instead of the reverse. As unsettling as this might be from a philosophical point of view, it works quite well for the vast majority of software products sold today.

A clear exception holds for the category of *safety-critical* software applications. Most will agree that different rules must apply here, since for these types of applications it cannot be considered acceptable for a vendor to decline all responsibility for the potential damage caused in return for a mere commitment to fix any problems not caught in the software development process until *after* they have manifested themselves to end-users. If we now look more carefully at *which* different rules are applied in these cases, we are in for a surprise. In many cases there are no software certification requirements, and those requirements that do exist can only be described as modest. Organizations are often only asked to show “how hard they’ve tried” and not that certain standards of workmanship are met.

As part of our research and work, we have inevitably gained experience with the analysis of many software products that are considered critical. At JPL this naturally includes the analysis of the control software for interplanetary spacecraft, but we have also been involved in a broad range of other types of safety-critical applications, including the investigation of specific aspects of automotive software (e.g., in the context of a study of the potential for sudden unintended acceleration of Toyota vehicles in 2010), medical device software, software used in the shutdown systems of nuclear power plants, railway signaling protocol software, etc.

NASA’s shuttle software [2] is often mentioned as an example of how critical software systems can reach a high level of safety and reliability. This software indeed has an exemplary track record of having a low residual fault density rate over the approximately three decades of use. Like any other human design, it is, of course, not completely free from defects, nor can it be expected to be. One could well say that the first principle adopted in the design of any system that is meant to be reliable is the recognition that no single system component can be perfect: every part has breaking points, some known and some unknown. Reliability and safety, therefore should be treated as system properties, not component properties. (And to complete the argument, in almost all cases of interest the software is merely one component in a larger system

that includes also hardware and human operators as essential elements.)

The software used for commercial airplanes, similarly, has an enviable track record for reliability. Again, the track record is not for perfection, because in any sufficiently complex system there are always residual defects that are discovered only after a system is delivered and goes into operation. The goal for certified software, therefore, cannot be to put a process in place that guarantees correctness under all circumstances – the goal is to produce a safe and reliable system that is built by competent, well-trained developers, following a process that controls risks and meets the evolving standard of skilled workmanship. In one sentence here, then we touch on three separate targets for a software certification process: the code, the coder, and the coding process that is followed. The certification process that is followed in the aerospace industry (e.g., for software used to control commercial airplanes), targets primarily the coding process, e.g. with standards such as DO-178B [3]. There are no strict requirements here for the use of specific verification tools, or for the certification of software developers themselves. The target is only to secure that due diligence was used in the development process itself.

2. CERTIFICATION PROCESS

At JPL, in the development of the control software for interplanetary spacecraft, such as the Mars Exploration Rovers [4], we have adopted a different process. The intent of this process is to subject not just the coding process, but also the code, and the coders, and to some extent the software managers as well, to some form of certification.

2.1 The Coding Process

For flight code, JPL has adopted an Institutional Coding Standard [5], which it requires compliance in all newly developed mission-critical software written at JPL.¹ The coding standard deliberately captures only risk-related rules for which compliance can be verified mechanically. Other than most other coding standards, then, this standard has real teeth: except in rare cases, non-compliance is not an option for our flight software developers. Because all rules in the standard are specifically risk related (i.e., we can often point at a mission anomaly or mission loss that was caused by the violation of the underlying principle), approval for non-compliance is also rarely requested or granted. An example of a risk-related rule in this coding standard is the abolition of all dynamic memory use and of recursive code. Some of the motivation for the rules can be traced to the Power of Ten rules, described in [6]. JPL further imposes fairly strict requirements on the code review process some of which is detailed in [7].

2.2 The Coders

Starting in 2010, JPL adopted a new procedure for the certification of flight software developers. The procedure itself is still subject to some revision, but once fully operational no software developer will be able to touch flight code (develop, manage, or modify) without having successfully completed a JPL specific Flight Software Certification course. The course consists of three modules, focusing on (a) computing science principles, (b) JPL software development standard processes, and (c) software risk and software vulnerabilities. Each module takes two

full days of instruction, for a total of six days for all three modules combined. Each module ends with an exam that must be completed with a passing grade. At the time of writing, the first twenty software developers have successfully completed this course, and have received their certificates, others have not passed and will have to take the course, and the exams, again. New classes are held several times a year, until all software developers have been certified. At that point, we will likely add refresher courses for those who are already certified, in addition to the basic certification course itself, to keep pace with continuing developments in this field. The certification course intends to certify that all developers of critical code are familiar with basic computing science theory, and standard algorithms, are intimately familiar with the risks inherent in the use of the programming languages that are typically used for flight code, and understand not just the letter but also the rationale for the coding standard that they are expected to follow. The certification course also introduces developers to the tools (e.g., static analyzers, code review tools, and unit test tools) that they will be using in flight software development.

Perhaps as an aside, JPL has also instituted an (as yet non-required) course for senior management. Senior management normally has deep experience with spacecraft and mission design, but less so with software design principles. To date, most of JPL's senior management has taken and completed this course. The course is repeated once a year, and is by invitation only.

2.3 The Code

The code, finally, is rightfully subject to the strictest requirements. Flight code, e.g. for the MSL mission [4], is checked nightly for compliance with the JPL coding standard [5], and subjected to rigorous tests with four separate state-of-the-art static source code analysis tools [7] (at the time of writing this includes the commercial tools coverity, codesonar, and semmlle, and the research tool uno). The warnings generated by each of these tools is combined with the output of mission-specific checkers that secure compliance with naming conventions, coding style, etc.. In addition, all warnings, if any (there should be none), from the standard C compiler, used in pedantic mode with all warnings enabled are included in the results that are provided to the software developers as part of the standard 'scrub' interface [7]. The developers are required to close out all reports before a formal code review is initiated. In peer code reviews, an additional source of input is provided by designated peer code reviewers, and added to the 'scrub' results.

Separately, key parts of the software design are also checked for correctness and compliance with higher level design requirements with the help of logic model checkers, such as Spin [8]. Training in the use of logic model checkers is tacitly provided via (optional) graduate-level courses taught by members of the JPL Laboratory for Reliable Software in the Computer Science Department at the California Institute of Technology. Approximately ten JPL employees outside the Laboratory for Reliable Software have so far taken and passed this course, and have become proficient in the use of logic model checkers on the analysis and verification flight software.

3. REGULATORY PROCESS

As noted in the introduction, members of our team have been involved in a broad range of software analysis applications, targeting not only aerospace but also safety-critical software used in automobiles, medical devices, and in the shutdown systems of nuclear power plants. It is perhaps noteworthy that at present there

¹ Most flight software, by a significant margin, is traditionally written in the C programming language, and therefore the JPL coding standard targets this language.

do not appear to be any strict regulatory requirements on the development of these critical types of software applications, neither on the code or the coders, on the organization that employs the coders, or on the processes that are followed in the coding process.

In the automotive industry there is reasonable consensus on at least one set of coding guidelines: the one developed by the organization MiraLtd, and known as the MISRA-C Coding Guidelines [9]. Curiously, although many developing organizations have publically expressed support for these guidelines, there is no requirement (or verification) that they actually comply with them.

Compliance with any reasonable standard, e.g., [5,6,9], can make it significantly simpler to analyze code for potential anomalies, and to revise, and maintain it longer term. Much the same is true in the medical device industry, where the FDA does not require compliance with any specific coding standard or software development process, and goes no further than to *recommend* the use of state-of-the-art static source code analyzers as part of software development process, without actually requiring evidence that this is done. Similarly, the Nuclear Regulatory Commission has issued no comparable requirements for any software used in the shutdown systems of future nuclear power plants, nor does it seem to have plans to do so, as a key part of the licensing process.

We believe that in each of these cases the lack of requirements on software development is an omission that should be corrected. Where not following generally accepted principles for safe software development could be regarded as a lack of workmanship on the part of the developer or developing organization, with the potential effect of contributing to preventable software failure, inadequate regulation for safety-

critical software systems that we all rely on could well be regarded as a failure of the regulatory process.

4. ACKNOWLEDGMENT

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

5. REFERENCES

- [1] Rushby, J. Verified Software Systems – the certification perspective. <http://www.csl.sri.com/users/rushby/vsr-roadmap-ccert-apr06.pdf>.
- [2] Fishman, Charles, They Write the Right Stuff, <http://www.fastcompany.com/magazine/06/writestuff.html>.
- [3] DO-178B, Software considerations in airborne systems and equipment, , <http://en.wikipedia.org/wiki/DO-178B>.
- [4] Mars Science Laboratory mission (MSL), http://www.nasa.gov/mission_pages/msl/index.html.
- [5] JPL Institutional Coding Standard for the C Programming Language, http://lars-lab.jpl.nasa.gov/JPL_Coding_Standard_ext.pdf
- [6] The “Power of Ten” Coding Rules, <http://spinroot.com/p10/>
- [7] G.J. Holzmann, Scrub: a tool for code reviews, Innovations in Systems and Software Engineering, Vol. 6, No. 4, 2010, pp. 311-318.
- [8] G.J. Holzmann, The Spin Model Checker – Primer and Reference Manual, Addison-Wesley, 2004.
- [9] MISRA-C:2004, Guidelines for the use of the C language in critical systems, Mira Ltd.