


ACCELERATING MISRA AUTOMOTIVE
SOFTWARE SAFETY COMPLIANCE WITH
STATIC APPLICATION SECURITY TESTING





INTRODUCTION

The MISRA C/C++ coding guidelines came about due to concern for safely using C and C++ in critical automotive systems. Since its inception in 1998, MISRA has become one of the most used coding standards in the automotive industry and has spread to other safety-critical devices in medical and industrial control. Static application security testing (SAST) solutions are needed to ensure software code quality, security and critical safety and enforce the standard, but not all tools are created equal. Sophisticated SAST solutions that provide support for the complex development process and perform more than simple syntax checking are desired to reduce risk, costs, and time to market.

SUPPORTING THE DEVELOPMENT PROCESS

SAST solutions support multiple aspects of the software development process, from early code development to post-launch debugging and forensics. Our whitepaper [Making Safety-Critical Software Development Affordable with Static Analysis](#) describes this in more detail. The recurring theme is that SAST plays a critical role in improving software quality, enforcing safe coding standards such as MISRA, and detecting defects and security vulnerabilities that are difficult to find during testing.

MISRA plays an important role in C/C++ development when applied to safety-critical automotive software. Enforcement of the standard is difficult to do manually, so SAST solutions are used as guidelines to enforce the standard as developers code. It is important to note, though, that there is a significant difference between advanced static analysis that supports MISRA rule checking and simpler tools that perform rule checking alone. GrammarTech's CodeSonar is capable of supporting MISRA coding enforcement but also providing valuable error detection and security vulnerability analysis that goes beyond simple code syntax checking.

CodeSonar is integrated with software development tools such as GitLab, GitHub, Jenkins, Jira, Eclipse, etc. These integrations allow seamless adoption of SAST into an existing development process. Source code quality and standards compliance can be checked right at the developer's desktop before checking into the build system. Defects and vulnerabilities can be automatically assigned for review and remediation. Audits can be done at any time and results distributed to the development team. This integration promotes efficiency and cost savings by enabling developers to fix issues early and avoid rework later in the process.

The MISRA guidelines themselves acknowledge the importance of tools in successfully using the guidelines:

...in its favor as a language is the fact that C is very mature, and consequently well-analysed and tried in practice. Therefore, its deficiencies are known and understood. Also, there is a large amount of tool support available commercially which can be used to statically check the C source code and warn the developer of the presence of many of the problematic aspects of the language

[MISRA-C:2004 Guidelines for the use of the C language in critical systems]

MISRA guidelines define a safer subset of C or C++ that should prevent many classes of errors. Following these guidelines does improve code quality and, in combination with modern development tools, rigorous testing and good software development practices should improve system safety and security as well (assuming the level of rigour remains the same throughout hardware and system development.) However, no coding standard is perfect and enforcing the rules isn't enough, by itself, to ensure better software quality, safety and security. Moreover, without tool support (as acknowledged by MISRA's authors) the standard doesn't stand on its own.

To really understand why SAST is crucial in producing safety-critical software for automotive applications, we need to first look at how defects get into programs, and second, how these defects are found and fixed. Defects are an unfortunate but unavoidable side effect of writing code. Research has shown that software defects are found at an average rate of one defect per ten lines of code. If developers are not scanning code and fixing these defects early on, finding and fixing those issues later in the software development life cycle (SDLC) will be labor intensive and costly.

But by the time the code is released, most commercial software will still have about one defect per 1,000 lines of code (KLOC). Open source code is a little better, at 0.68/KLOC. Cleanroom software, which combines formal methods of requirements and design with statistical usage testing, was found to have 0.1/KLOC. Considering that the software content of today's luxury car is well above 100 million lines of code, even in the absolute best-case scenario, there would still be at least 10,000 latent software defects and, in practice, there are likely more than 100,000. The defects that remain in the software are typically difficult to detect in post-development testing, and can result in poor interoperability with other subsystems or unchecked interfaces, unexpected behavior, outright failure or exploitable vulnerabilities.

Finding and fixing defects during post-development testing is extremely time consuming. When a tester finds an error or a failure, the root cause is in many cases unknown, forcing the developer to trace the problem back to the source. This means that the circumstances that created the failure must be reproduced, and the developer must study and understand the related code and the impact to other parts of the code base. The reality is that many bugs cannot be reproduced on demand and are never fixed.

The power of integrating a SAST solution into the development process is that it doesn't rely on test cases to find problems, nor does it require that an error or failure be reproduced. An advanced SAST solution can infer the runtime behavior of a program without actually running the program. Furthermore, when it identifies a problem, it also pinpoints the locations within the code and the interdependencies that created the failure. This makes the job of debugging far simpler. SAST does not eliminate the need for testing altogether, but it complements QA and testing activities procedures, providing an additional and critical validation technique. The reality is that in large and complex software systems, there are so many possible states, and such an astronomical number of possible paths of execution, that it is infeasible to exhaustively test them all. SAST, on the other hand, can explore these paths and state conditions in the aggregate, and is able to find problems that are missed in post-development code testing.



A key contribution that advanced SAST solutions like CodeSonar provide to secure and safety-critical software development is the ability to find defects that have slipped through the traditional development techniques. As evidenced by high profile cases in the automotive industry, security and safety issues that make it to market and on the road are expensive to fix - orders of magnitude more expensive than during development. These defects can also be hazardous to drivers and put the automotive manufacturer at risk of financial and reputational damage.

Critical defects need to be detected early, even early enough to be caught during coding. Some examples of classes of defects that can be missed during development, even when good engineering practices, MISRA coding, and rigorous testing is done, are given below:

- **Concurrency defects** often occur randomly and only after a system has been integrated completely on the final hardware platform. Unlike other testing, GrammaTech's CodeSonar can reason about multithreaded/multitasking code behavior and detect dead locks, race conditions, and other types of concurrency errors.
- **Security vulnerabilities** are software defects that can be exploited to interfere with a system's behavior, create safety issues or expose critical data. Security is often overlooked in systems where safety is paramount. Advanced SAST solutions can detect security vulnerabilities in code that arise from malformed or tainted data outside expected values.
- **Tainted data analysis** helps to trace input data in the system to its use in the application and warn of any potential security vulnerabilities that arise. In today's hostile operating environment, it's foolhardy to assume that system input data is well-formed. Detecting these types of security issues is very difficult when data is passed across many functions. Automated SAST analysis can provide the full control and data path for tainted data, which allows for rapid remediation in the code base.
- **Complex inter-procedural defects** are difficult to detect, especially with unit and subsystem testing. CodeSonar does advanced inter-function (procedure) analysis of control and data flow of the entire scope of the program. Deep analysis decreases the rate of false positives (errors that are false) but also increases the rate of true positives (errors that are verified true). CodeSonar's analysis extends into executables, object files, and libraries.
- **Binary analysis** provides insight and error detection of compiled code as object files, libraries, and even executables. CodeSentry from GrammaTech performs binary software composition analysis to produce a software bill of materials (SBOM) that identifies the ingredient list of the software – third-party and open source components. Developers can use this SBOM to ensure that the all of the code (both binary and source) is up to the quality standard required for the project. CodeSentry can also automate a vulnerability report of these identified components to detect any hidden issues that could impact security and safety of the software.

MISRA RULE COVERAGE

When developers are required to adhere to coding standards, they must look for a solution that can help them find and fix violations as part of their workflow. One of the metrics they use to compare tools is the coverage (the proportion of rules that the tools claim to check), with a naïve strategy being to choose the tool that claims the highest coverage.

Unfortunately, the notion of coverage is not well-defined, and because there is no reliable source of information that can be used to compare coverage between tools, customers must trust vendors to reasonably interpret the term, and to report their coverage fairly.



WHILE SOME COVERAGE IS EASY,
COMPLEXITY IS MORE DIFFICULT

Some MISRA rules are simple enough that a straightforward code-syntax-checker can find all violations with no false positives. For such rules, coverage is easy—the tool can either find violations or not, with no ground in-between. In MISRA C 2012, such rules are labeled Decidable. If the violation can be detected by looking only at one compilation unit, the rule is also labeled Single Translation Unit. For example, rule 4.2 forbids the use of Trigraphs, and is labeled this way. Claiming coverage for this rule is believable and provable.

Less clear, however, is if a violation can only be reliably found if the tool needs to look at multiple compilation units at the same time (these are labeled System in the standard). For example, MISRA C 2012 rule 5.1's "External identifiers should be distinct" is certainly decidable, but the only way a tool can reliably find a violation is if it examines all compilation units and compares all such identifiers found in each.

If a tool claims to have full coverage of a rule with System scope, then it is only reasonable to believe that the tool is also capable of finding all the compilation units that contribute to the program. Over-approximation and under-approximation of the set can lead to both false positives and false negatives. Humans routinely get this wrong, so a user of a tool that does not offer an automatic way to determine the set is running the risk of getting incorrect results. The most effective tool is one that integrates tightly with the build system, as that is most often the most trustworthy source.

UNDECIDABILITY

In MISRA C 2012, some rules are labeled “Undecidable,” meaning that it is fundamentally impossible to have a method that can, in general, say for sure if a violation is present or not. Because of this property, the author of a checker must find a sweet spot that balances the risk of false positives with the risk of false negatives. Most of these rules require an analysis that is capable of reasoning about the execution of the program, so only the most sophisticated SAST solutions can be expected to do a good job. A good example is MISRA C 2012 rule 17.2, which forbids recursion, both direct and indirect (i.e., calls through function pointers).

The trouble is, claims of coverage are often made that ignore whether a tool is good or bad at finding violations. If a tool can only find the most obvious and superficial instances of violations, is it reasonable for it to claim that it has coverage of that rule? The other side of that coin is interesting to consider too — if a tool finds all violations, but also reports so many false positives that is impractical to inspect them all, is it fair to say that it has coverage?



COVERAGE BREADTH

The final aspect of rule coverage that makes it complicated is that coding standards are usually quite loosely defined, whereas SAST solutions must have a precise definition of the properties that they are looking for. Consequently, it is common for a checker to detect a property that is either a superset or a subset of what the rule requires.

For example, let's consider CodeSonar's coverage of MISRA C 2012 Rule 2.2: “There shall be no dead code.” For the purposes of this rule, dead code is code that is executed, but whose removal cannot affect program behavior. CodeSonar has an Unused Value checker, which finds places where a variable is assigned a value that is never subsequently used. All such places violate the MISRA rule, but there are other ways in which the rule can be violated that are not detected by this checker. Thus, the Unused Value checker covers only a subset of what the rule specifies, and other CodeSonar checkers fill in the gaps.

In some cases, the rule and the checker are not in a strict subset/superset relationship. They may overlap a lot or a little, or the checker may detect a property that is not a direct violation but is very likely to lead to a violation of the rule.

In CodeSonar, the policy is to claim coverage only if there is a large overlap between what the rule specifies and what our checker will find, and where the checker does not yield warnings that would be reasonably judged to be false positives for that rule (notwithstanding that they may be true positives otherwise).

MISRA C 2012, RULE 1.3

There is one MISRA C 2012 rule in particular for which the issue of claimed versus real coverage is acute. One should be wary of a SAST tool claiming coverage of MISRA C 2012 Rule 1.3: “There shall be no occurrence of undefined or critical unspecified behavior.” This rule is so broad that an additional 10-page Appendix in the standard outlines possible scenarios to avoid. This in turn references the C standards: those for C90/99 enumerate 230 instances of undefined behavior (65 of these are not covered by any other MISRA rule), and 51 instances of critical unspecified behavior (of which 17 are not covered by any other MISRA rule). Furthermore, the rationale for Directive 4.1 adds: “the presence of a run-time error indicates a violation of Rule 1.3.”

Consequently, Rule 1.3 specifies an enormous amount of forbidden behavior, including null pointer dereferences, buffer overruns, use of uninitialized memory, data races, use after free errors, and many of the other hazards of programming in C.

The problems with rule coverage claims should be clear – although this one rule (of the 143 in the standard) constitutes only 0.7% of the standard, it covers maybe 50% of the truly unpleasant kinds of failures that can befall a C program.

Furthermore, if a tool is to claim coverage of Rule 1.3, it should have checkers that have a good-sized intersection with all of those undesirable behaviors. If a tool can only find a tiny percentage of them, then it is unreasonable for it to claim coverage. In contrast, advanced SAST solutions, such as CodeSonar, are explicitly designed to find the kind of run-time errors that constitute violations of Rule 1.3. Analysis that make it possible for them to find such defects with reasonable precision must be whole-program, path-sensitive, aware of hazardous information-flows, and capable of reasoning about concurrently-executing threads.

ACCELERATING TIME TO CERTIFICATION

Obvious advantages of SAST solutions include risk-reduction and time-reduction from finding and fixing defects and vulnerabilities in early stages of development. The cost savings of finding critical defects is significant compared to finding and fixing these bugs in system integration or worse, when products are in service. Additionally, CodeSonar provides automated documentation to support testing and quality/robustness evidence. Much of the manpower used in safety certifications is documentation and evidence production.

Automation, and specifically SAST analysis reduces this burden significantly. As a TUV SUD certified ISO 26262 solution, CodeSonar provides assurances to developers that it can be integrated into a safety-critical development project without further certification requirements – something that adds costs and risk otherwise.

CONCLUSION

The tools needed to support successful safety-critical projects require more than simple source analysis and MISRA rule checking. Enterprise-level development projects require sophisticated solutions that support and enhance the full software development life cycle, including integrating with other development automation tools. The ability to go beyond MISRA rule enforcement and prevent critical defects and vulnerabilities from leaking through the development process pays big dividends in cost and risk reduction.

For more information or to request a CodeSentry demo,
contact GrammaTech or visit www.grammatech.com

U.S Sales: 888-695-2668

International Sales: +1-607-273-7340

Email: sales@grammatech.com

