

Figure 2: Source code decision and its mapping to conditional branches.

Criteria #4: It must be observable by the conditional branches/instructions in the recorded program flow trace whether a source code condition was evaluated true or false. So far, the basic concept. The screenshot in figure 2 illustrates what has been described so far.

Observability Gaps and Their Causes

Practice has shown that criteria #2, #3 and #4 are not always fulfilled in every test scenario. When this is the case, Lauterbach speaks of observability gaps. Observability gap means that TRACE32 cannot detect whether a condition has been evaluated as true or false at a certain point in the program flow trace. In this case, no MC/DC coverage result can be displayed for the related decision. Here are the most likely causes of observability gaps and the countermeasures that need to be taken:

1. No dedicated compiler support

First of all, you should consider writing code coverage friendly code. Nested decisions or simple decisions in the assignment context, such as `return a==b`, may cause observability gaps. Here it is not guaranteed for every compiler that all (atomic) conditions are represented by a conditional branch/instruction on the object code level. This can be easily avoided by following some simple coding guidelines summarized by Lauterbach. However, if one cannot or will not modify the source code of a colleague or external provider, gaps are unavoidable. Criterion #2 is violated, but the observability gaps can be closed by targeted instrumentation. The second part of this article will provide details on this instrumentation mode.

On the other hand, the large number of core architectures and the associated diversity of compilers represents a challenge. An impressive number of cores offer the possibility to generate program flow trace. And there are a big number of compilers, especially for commonly used core architectures. The result is a large amount of possible core architecture/compiler pairings. There is no generic

heuristic for mapping source code decisions to conditional branches/instructions at object code level that generates an exact result for every possible pairing. In practice, TRACE32 has to tailor the mapping to the core architecture/compiler combination. Much, especially for common core/compiler combinations is already tailored.

For not yet supported core architecture/compiler pairings, for which the generic heuristic of TRACE32 does not provide an exact result, criterion #3 is not to be met. Lauterbach offers targeted instrumentation or even full instrumentation as a fallback for these cases.

2. Macros

A macro that is used in a decision can in itself contain decisions. The compiler expands all macros before compilation and handles the expanded statement as a single source block. During this step the source code locations of the decisions inside the macro are lost. In this case, criterion #3 is violated. A mapping of the inside-macro-decisions to the conditional branches/instructions is no longer possible. The resulting observability gap can be closed by targeted instrumentation.

GLOSSARY

while (!(Identity(a) >= -45) && Identity(b) && Identity(c) || d)

- A **CONDITION** (grey in the picture above) is a logical indivisible, atomic expression. It can only be "true" or "false".
- A **DECISION** (framed by turquoise rectangle) is a logical expression which can be composed of several (atomic) conditions separated by logical operators such as "or", "and", "not". It results in true or false.

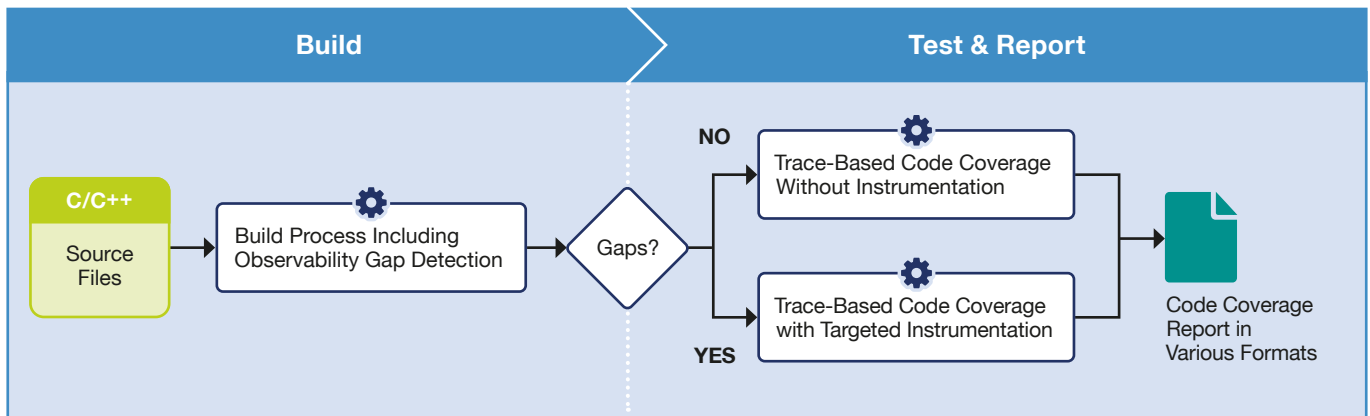


Figure 3: Two step workflow for TRACE32 multi-mode MC/DC coverage.

3. Highly-optimized code

Highly-optimized code is not recommended for trace-based code coverage analysis. For one, individual conditions may not be represented by conditional branches/instructions at the object code level. Criterion #2 is violated here. However, this can be remedied easily by targeted instrumentation. Highly optimized code is particularly challenging because it may not be possible to map the decisions exactly to the conditional branches/instructions. The violation of criterion #3 cannot be resolved in all cases.

Moderate optimization is recommended here. This is also advantageous because TRACE32 can display the results of the MC/DC coverage analysis clearly and in an intuitively and readable way.

4. Limitations of the trace protocol

The instruction set for a core architecture may contain conditional instructions. The compiler uses these to implement source code conditions at object code level. For trace-based code coverage to work, the trace protocol used must generate details about the execution of these conditional instructions. Unfortunately, this is not always the case. Currently there is no option that advises the compiler not to use conditional instruction. Observability gaps in program tracing are therefore inevitable. Criterion #4 is violated, but targeted instrumentation can be used to close the gaps.

5. Instruction set complexity

The challenges described in 1-4 are essentially the ones faced by cores with general-purpose RISC architecture. However, complex SoCs also contain coprocessors and special-purpose cores for which an instruction trace is generated. Examples are DSPs, configurable cores with user-defined instructions, timer IP and many more. Here, TRACE32 must always be specially adapted to the instruction set for an MC/DC coverage analysis. In this respect, it is always advisable to check with Lauterbach in good time.

Summarizing what has been described so far: a bit of source code instrumentation might be required to verify MC/DC by trace-based code coverage.

Workflow for Multi-Mode MC/DC Coverage

Since full instrumentation is only required in corner cases, we will concentrate on the two standard TRACE32 MC/DC coverage modes in the following.

- Trace-based code coverage without instrumentation
- Trace-based code coverage with targeted instrumentation

Figure 3 provides an overview of the workflow, which is organized into the two steps "Build Process" and "Test & Report".

Build Process

Figure 4 on the next page gives a detailed overview of the build process. The tasks of the build process are as follows:

1. Create a classic ELF file for the "Test & Report" step.
2. Create the .eca files that provide TRACE32 with the necessary decision details (as required by criterion #1).
3. Create an instrumented ELF file for the "Test & Report" step, in case that observability gaps have been detected.

All outputs of the build process required for the "Test & Report" step are drawn in turquoise and marked with a downward arrow in Figure 4. To generate the necessary outputs, the following TRACE32 products must be integrated into the build process.

- Clang-based command line tool t32cast; t32cast for C/C++ is free of charge and compiler independent.
- TRACE32 Instruction Set Simulator for the core architecture under test; the use of the TRACE32 Instruction Set Simulator requires a paid floating license.

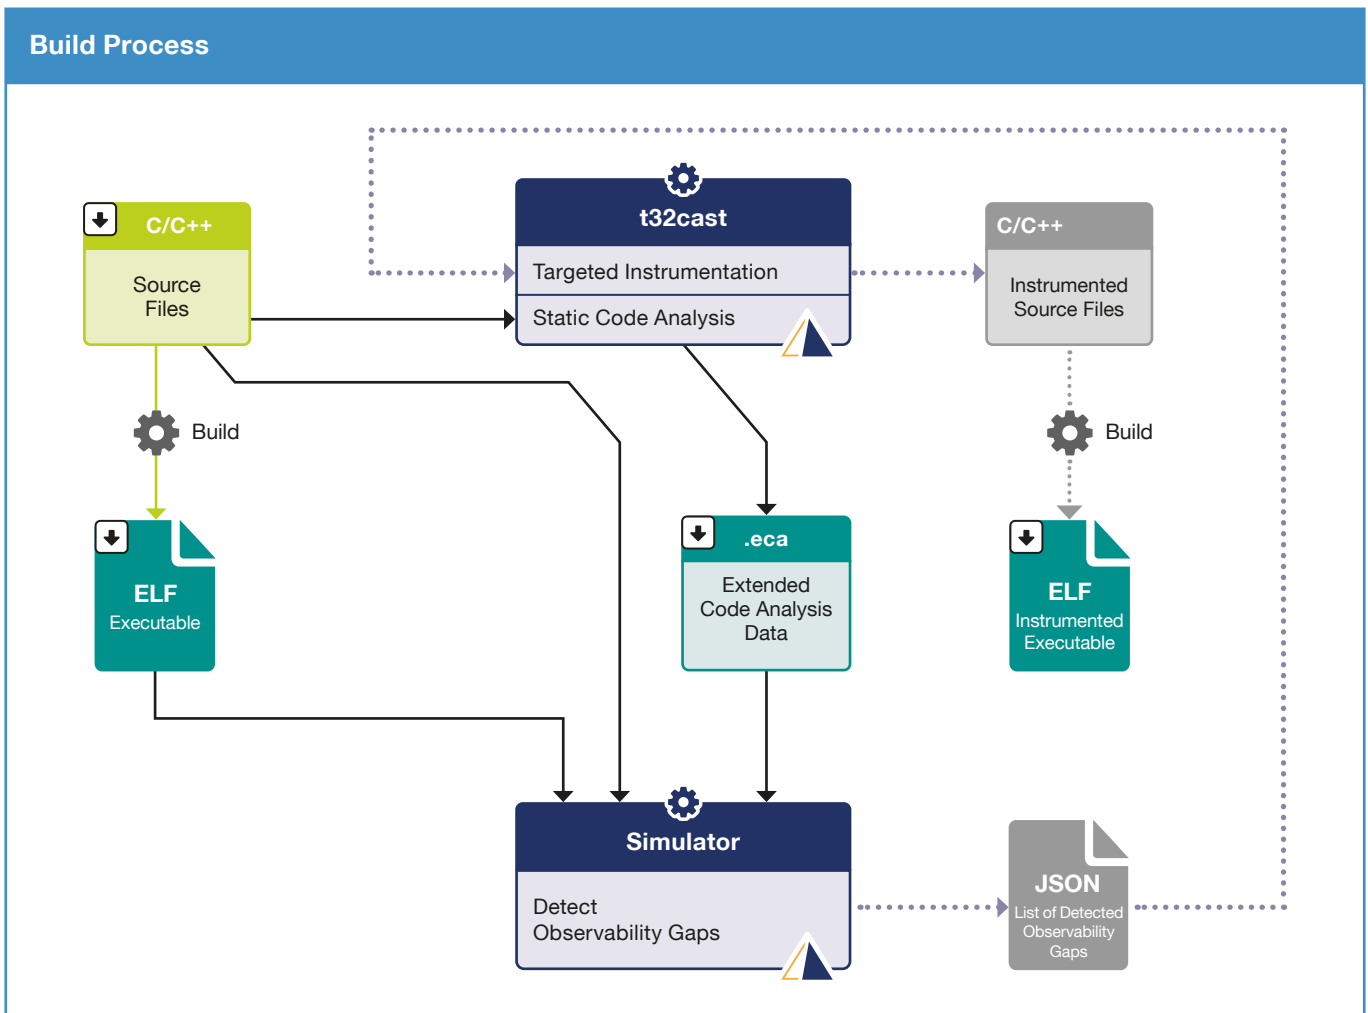


Figure 4: The build process for TRACE32 multi-mode code coverage.

What additional tasks do the two TRACE32 products handle in the build process?

1. t32cast analyzes the C/C++ source code. As a result of this static code analysis, an .eca file is generated for each source code file, which contains the required information about the decision structure.
2. The ELF file and the .eca files must be loaded into the TRACE32 Instruction Set Simulator to check for observability gaps. The result is saved in a JSON file.
3. If the JSON file is empty, the build process is complete.
4. If the JSON file is not empty, the source code must be instrumented at the locations where the observability gaps were detected. This is another task done with t32cast. The build process is completed here with the creation of an ELF file for the instrumented sources.

Before we proceed to the "Test & Report" step, a few words about the instrumentation performed by t32cast.

To ensure that TRACE32 can detect whether a condition evaluates true or false at any point in the recorded program flow, the source code is instrumented with the two instrumentation hooks t32__alpha() and t32__beta() for the detected observability gaps (see figure 6 on the last page). Both hooks are just function calls with an empty function body. TRACE32 evaluates these calls for MC/DC coverage analysis in addition to conditional branches/instructions.

The cost per instrumentation is very low because no additional interface to memory is required for TRACE32 targeted instrumentation. It is also convenient that no additional source code lines are generated, so the original source code can be used in the "Test & Report" step.

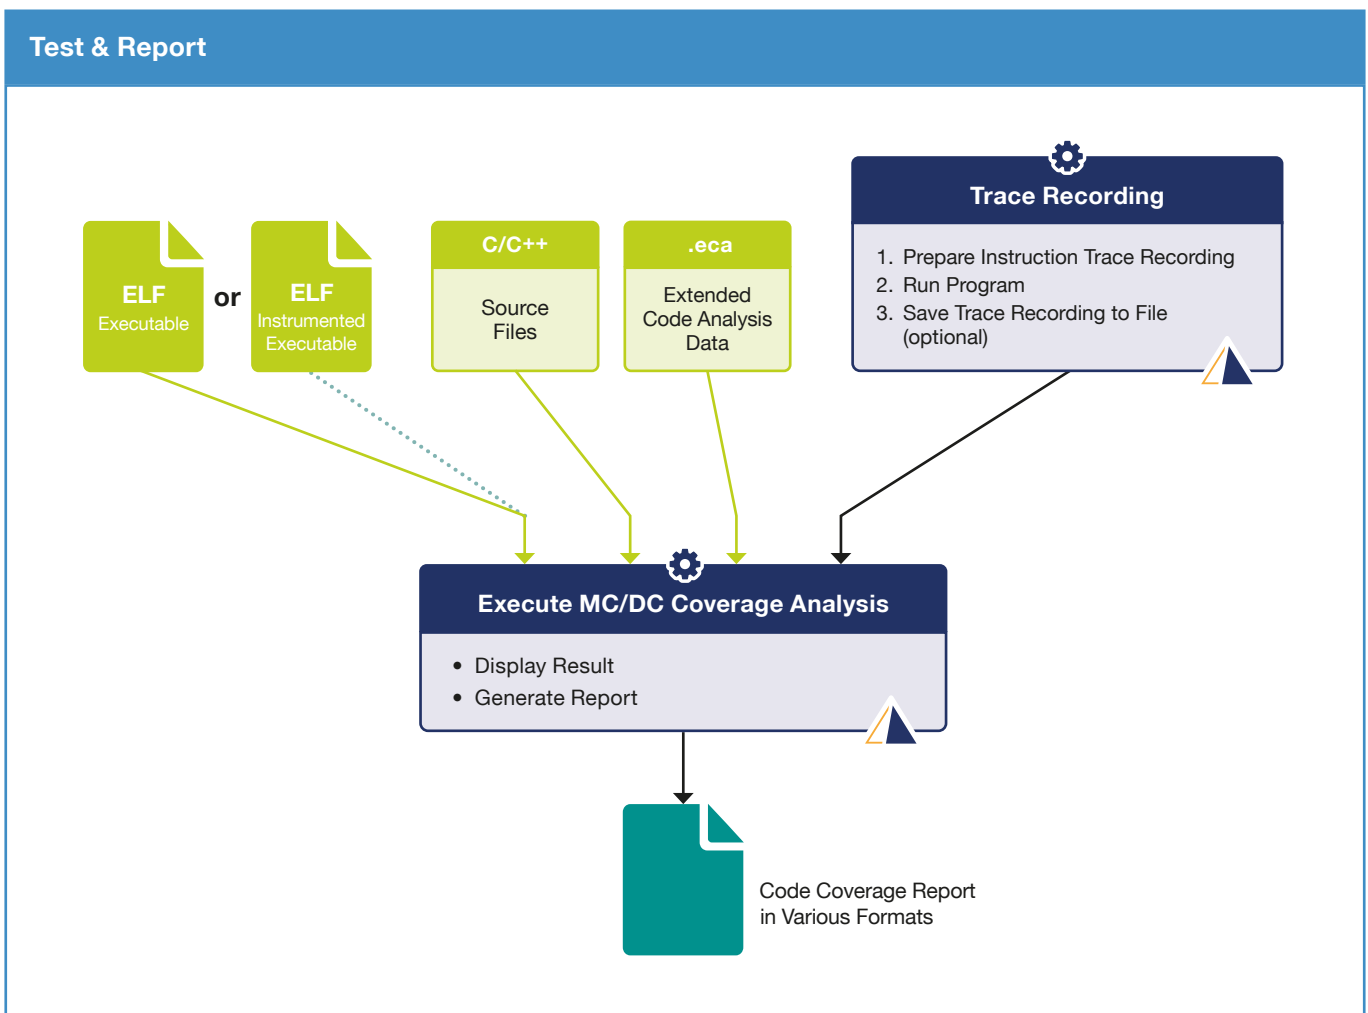


Figure 5: Test & Report for TRAC32 Multi-Mode Code Coverage.

Test & Report

For the test, the C/C++ sources, the corresponding .eca data and the corresponding ELF file must be loaded into the TRAC32 Debugger/Instruction Set Simulator, which executes the MC/DC coverage analysis (see figure 5). There are two options for trace recording:

1. The trace recording and the execution of the MC/DC coverage analysis are performed as a sequential task. The MC/DC coverage analysis is executed directly for the trace data recorded in the debugger.
2. Trace recording and the execution of the MC/DC coverage analysis are separate tasks, performed by two different test teams. In this case, the trace data must be saved in a file after recording and reloaded into TRAC32 for the MC/DC coverage analysis at a later point in time.

Recommended Workflow for Safety-Related Projects

As described above, trace-based MC/DC requires reduced optimization and may even require some code instrumentation. It is crucial that the embedded software built specifically for the coverage test purpose behaves in exactly the same way as the production software that will ultimately control the embedded system. Therefore, it is necessary to test both software variants side by side for the entire test life cycle. Figure 7 on the last page shows the testing workflow recommended for safety related projects.

Outlook

TRACE32 Multi-Mode Code Coverage allows us to achieve our goal of supporting MC/DC coverage for a wide range of core architectures, trace protocols and compilers. In many cases, you will get by without instrumentation at all. The targeted instrumentation that may be necessary requires never more than 10% additional memory for the whole object code and has a minimal time overhead.

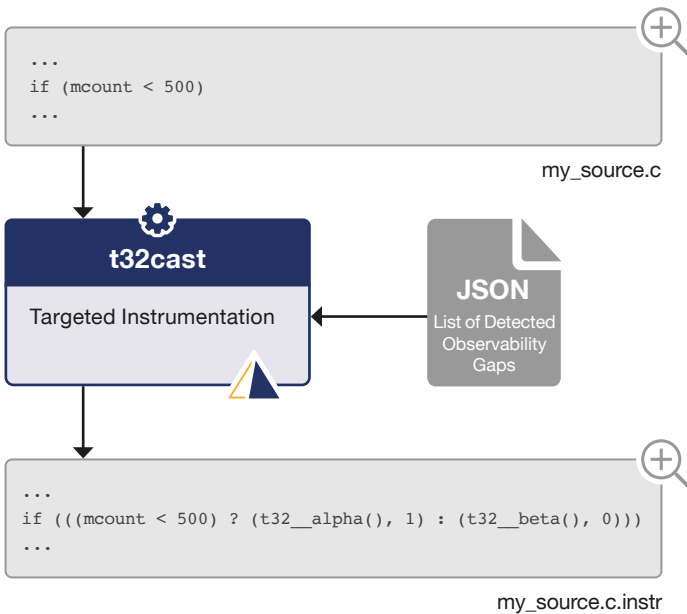


Figure 6: Instrumentation example for the multi-mode code coverage.

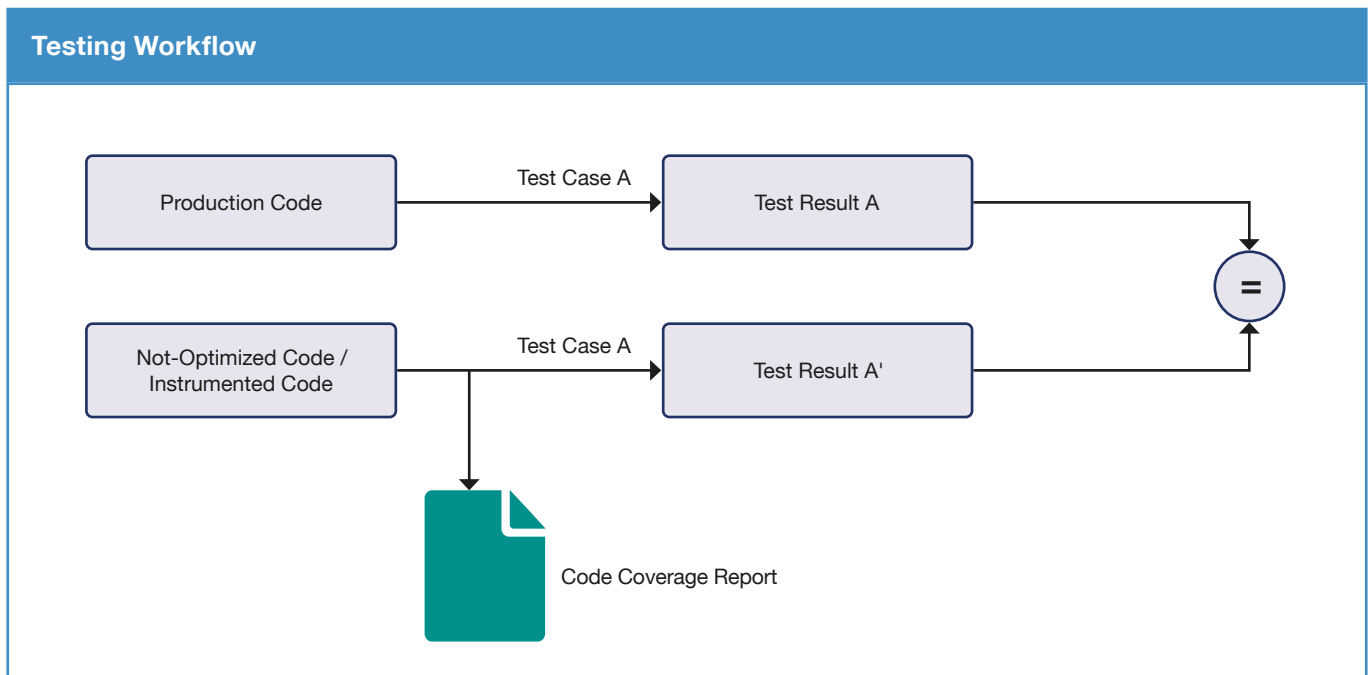


Figure 7: Testing workflow for safety-related projects.