

# Integrating Development Testing Practices: Extending Beyond Static Analysis

Development Testing techniques, such as pattern-based static code analysis, runtime error detection, unit testing, and flow analysis, are all valuable techniques for exposing defects. On its own, each technique can help you find specific types of errors. However, if you restrict yourself to applying just one or some of these techniques in isolation, you risk having defects that slip through the cracks.

A safer, more effective strategy is to use all of these complementary techniques in concert. This establishes a bulletproof framework that helps you find defects that are likely to evade specific techniques. It also creates an environment that helps you find functional problems, which can be the most critical and difficult to detect.

This paper will explain how Development Testing techniques such as pattern-based static code analysis, runtime error detection, unit testing, and flow analysis can be used together to find defects in an embedded C application. These techniques will be demonstrated using Parasoft C/C++test, an integrated solution for automating a broad range of best practices proven to improve C and C++ software development team productivity and software quality. The same techniques can be applied to Java (with Parasoft Jtest) and .NET code (with Parasoft dotTEST).

Automatically detecting memory corruption, deadlocks, and other defects is undoubtedly a vital activity for any development team, but the deadliest defects are functional errors. Functional errors often cannot be found automatically. We'll briefly discuss techniques for finding these defects at the conclusion of this paper.

## Introducing the Scenario

To provide a concrete example, we will introduce and demonstrate the recommended defect-finding strategies in the context of a scenario that we recently encountered: a simple sensor application that runs on an ARM board.


Assume that so far, we have created an application and uploaded it to the board. When we tried to run it, we did not see an expected output on the LCD screen.

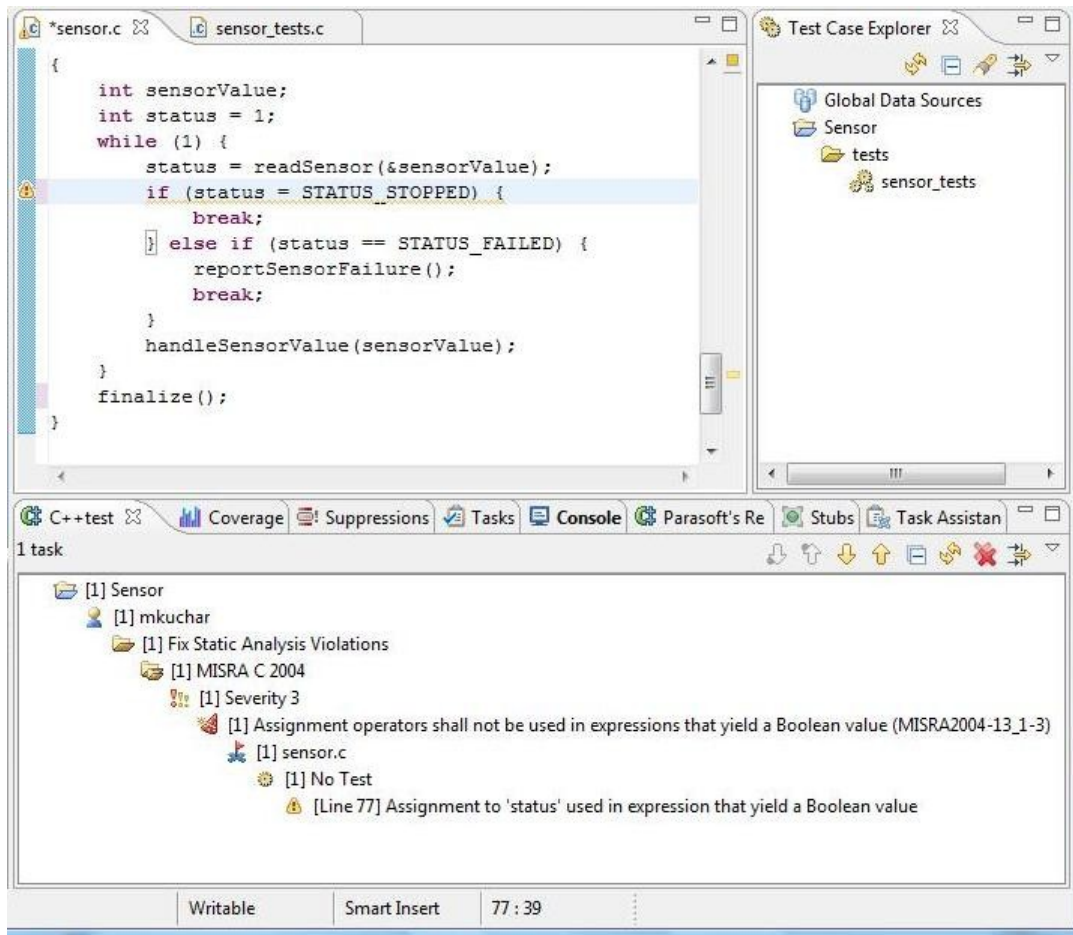
It's not working, but we're not sure why. We can try to debug it, but debugging on the target board is time-consuming and tedious. We would need to manually analyze the debugger results and try to determine the real problems on our own. Or, we might apply certain tools or techniques proven to pinpoint errors automatically.

At this point, we can start crossing our fingers as we try to debug the application with the debugger. Or, we can try to apply an automated testing strategy in order to peel errors out of the code. If it's still not working after we try the automated techniques, we can then go to the debugger as a last resort.

## Pattern-Based Static Code Analysis

Assuming that we don't want to take the debugging route unless it's absolutely necessary, we'll start by running pattern-based static code analysis, which detects a problem:



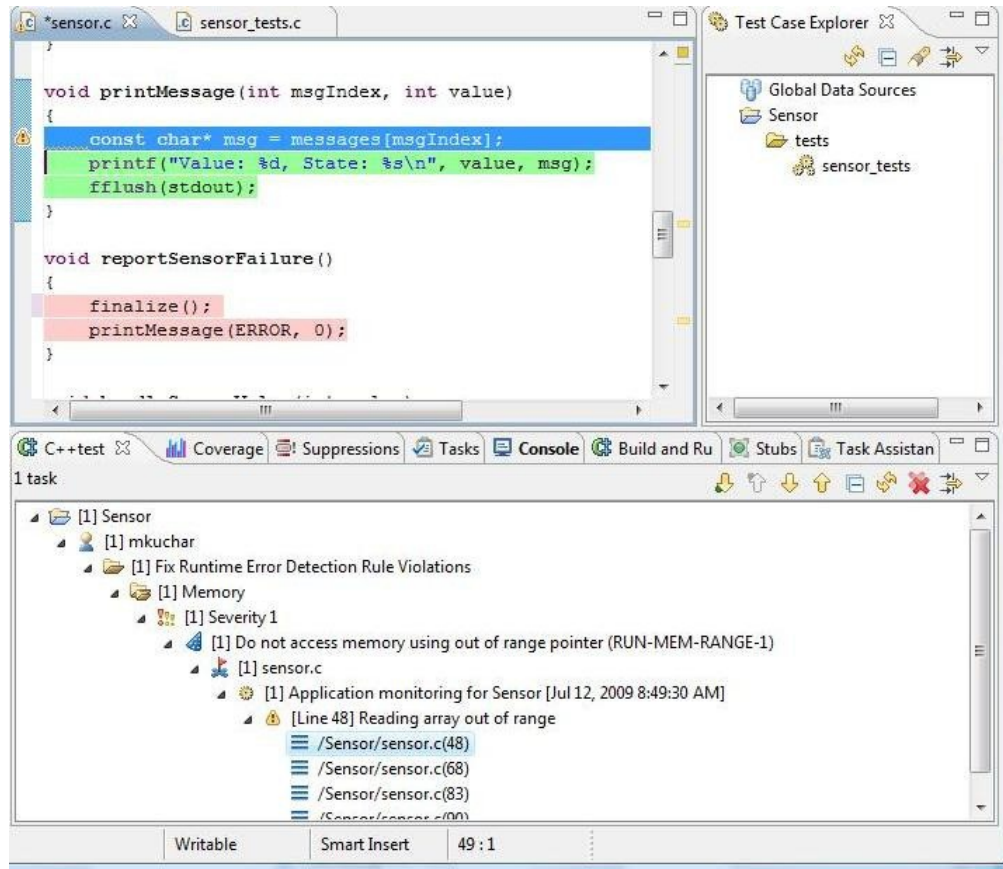


This is a violation of a MISRA rule that says that there is something suspicious with this assignment operator. Indeed, our intention was not to use an assignment operator, but rather a comparison operator. So, we fix this problem and rerun the program.

There is improvement: some output is displayed on the LCD. However, the application crashes with an access violation. Again, we have a choice to make. We could try to use the debugger, or we can continue applying automated error detection techniques. Since we know from experience that automated error detection is very effective at finding memory corruptions such as the one we seem to be experiencing, we decide to try runtime error detection.

## Runtime Error Detection on the Complete Application

To perform runtime error detection, we have C/C++test instrument the application. This instrumentation is so lightweight that it is suitable for running on the target board. After uploading and running the instrumented application, then downloading results, the following error is reported:



This indicates reading an array out of range at line 48. Obviously, the `msgIndex` variable must have had a value that was outside the bounds of the array. If we go up the stack trace, we see that we came here with this print message with a value that was out of range (because we put an improper condition for it before calling function `printMessage()`). We can fix this by taking away unnecessary conditions (`value <= 20`).

```
void handleSensorValue(int value)
{
    initialize();
    int index = -1;
    if (value >= 0 && value <= 10) {
        index = VALUE_LOW;
    }
    else if ((value > 10) && (value <= 20)) {
        index = VALUE_HIGH;
    }
    printMessage(index, value);
}
```

When we rerun the application, no more memory errors are reported. After the application is uploaded to the board, it seems to work as expected. But there should still be concern. We found one instance of a memory overwrite in the code paths that we exercised, but how can we be sure that there are no more memory overwrites in the code that we did not exercise?

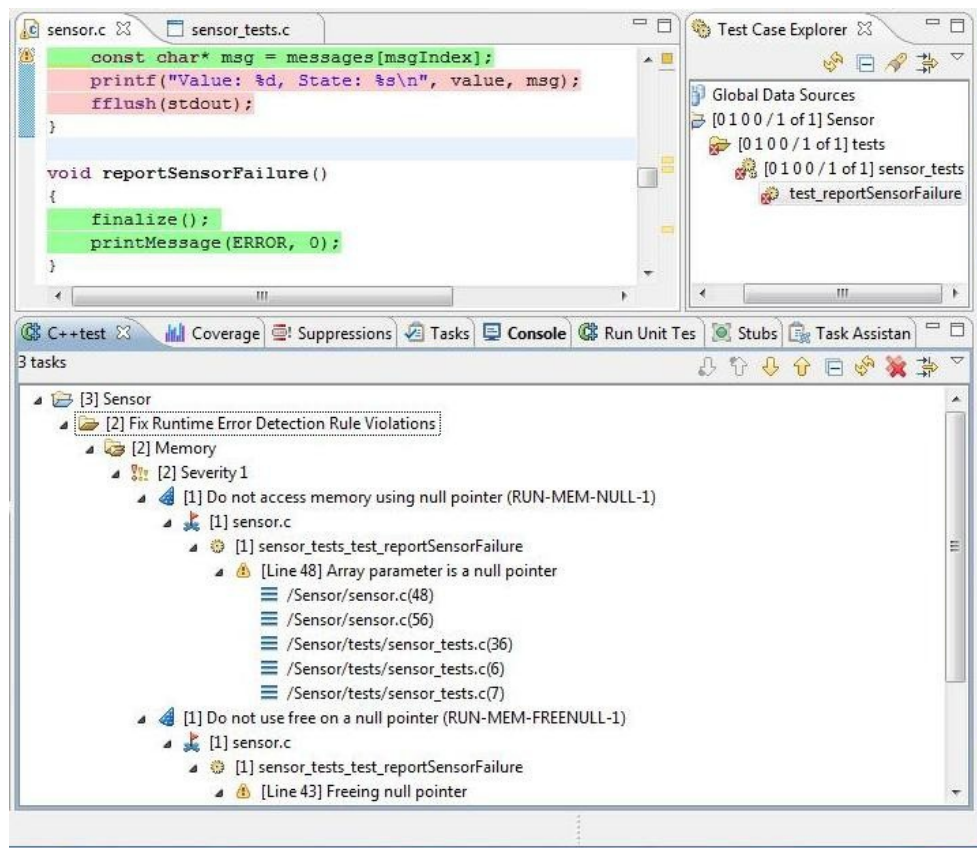
If we look at the coverage analysis, we see that the following function has not been exercised at all: `reportSensorFailure()`. We need to test this function by creating a unit test that will call this function.

## Unit Testing with Runtime Error Detection

The general process for unit testing with runtime error detection is:

1. Create a test case skeleton using C/C++-test's test case wizard
2. Fill in some test code
3. Run this test case—exercising just this one previously-untested function—with runtime error detection enabled.

With C/C++-test, this entire operation takes just seconds. The results show that the function is now covered, but new errors are reported:



Our test case uncovered more memory-related errors. We have a clear problem with memory initialization (null pointers) when our failure handler is being called. After further analysis, it's clear that we mixed an order of calls in `reportSensorValue()`. The `finalize()` function is being called before `printMessage()`, but `finalize()` actually frees memory used by `printMessage()`.

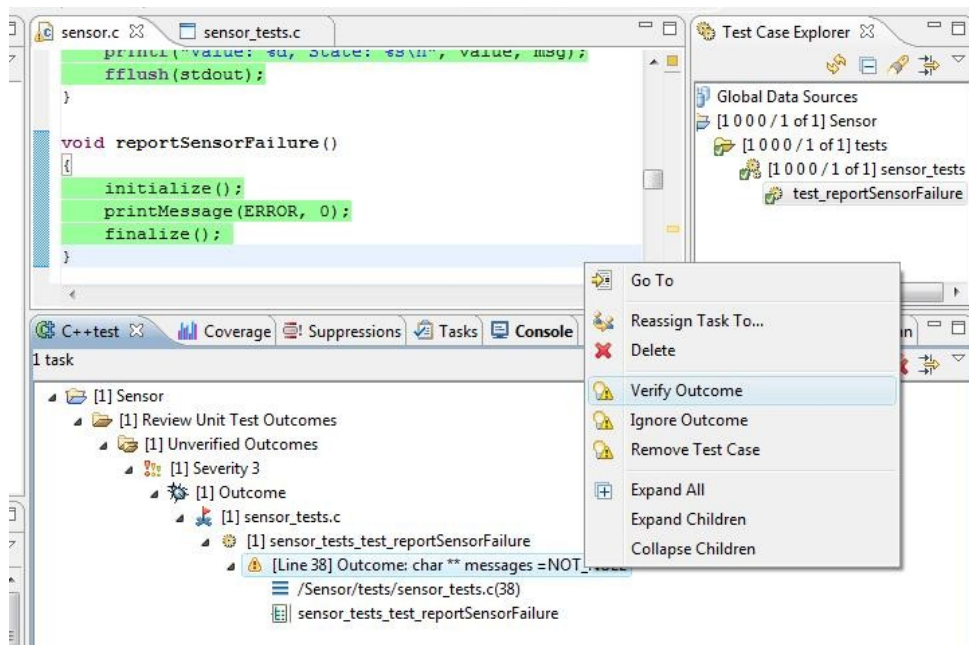
```
void finalize()
{
    if (messages) {
        free(messages[0]); free(messages[1]); free(messages[2]);
    }
    free(messages);
}
```

We fix this order then rerun the test case one more time.

That resolves one of the errors reported. Now, let's look at the second problem reported: `AccessViolationException` in the print message. This occurs because these table messages are not initialized. To resolve this, we call the `initialize()` function before printing the message. The repaired function looks as follows:

```
void reportSensorFailure()
{
    initialize(); printMessage(ERROR, 0); finalize();
}
```

When we rerun the test, only one test case task, which is not really an error is reported, was not validated. All we need to do here is verify the outcome in order to convert this test into a regression test. C/C++test will do this for us automatically by creating an appropriate assertion.



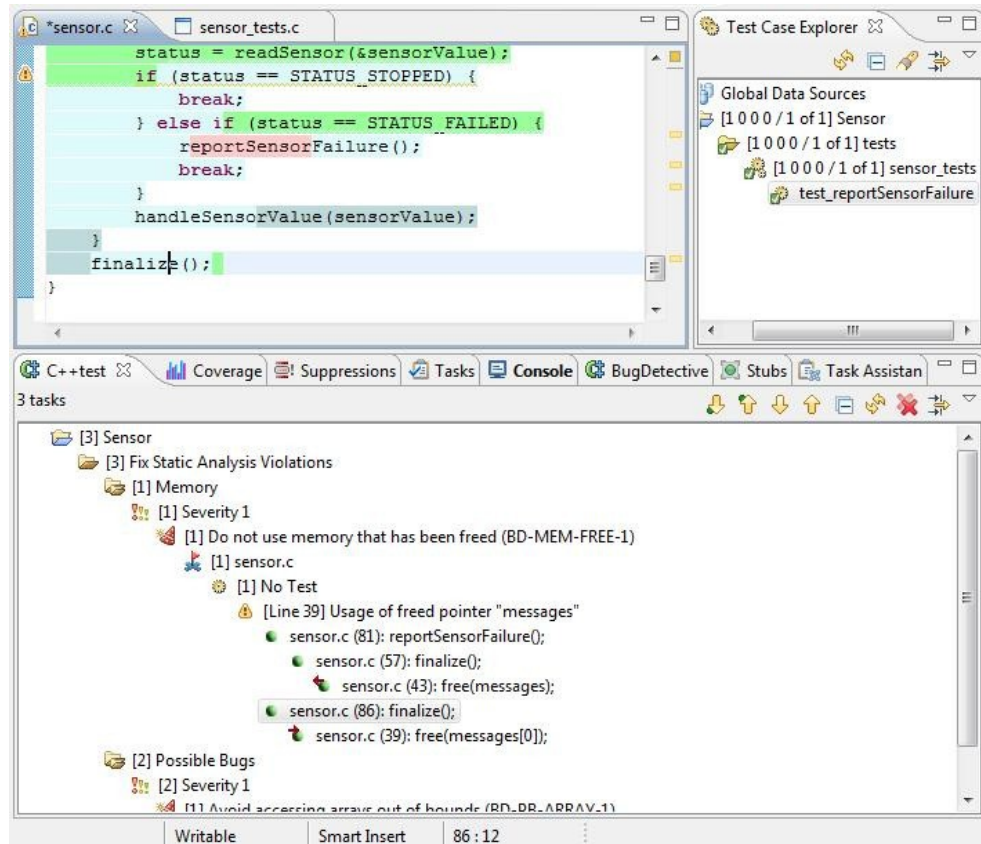
Next, we run the entire application again. The coverage analysis shows that almost the entire application was covered, and the results show that no memory error problems occurred.

But we are not quite finished, yet. Even though we ran the entire application and created unit tests for an uncovered function, there are still some paths that are not covered. We can continue with unit test creation, but it would take some time to cover all of the paths in the application. Alternatively, we could simulate those paths with flow analysis.

## Flow-Based Static Analysis

The flow analysis test configurations simulate different paths through the system and check if there

are potential problems in those paths. The following issues are reported:



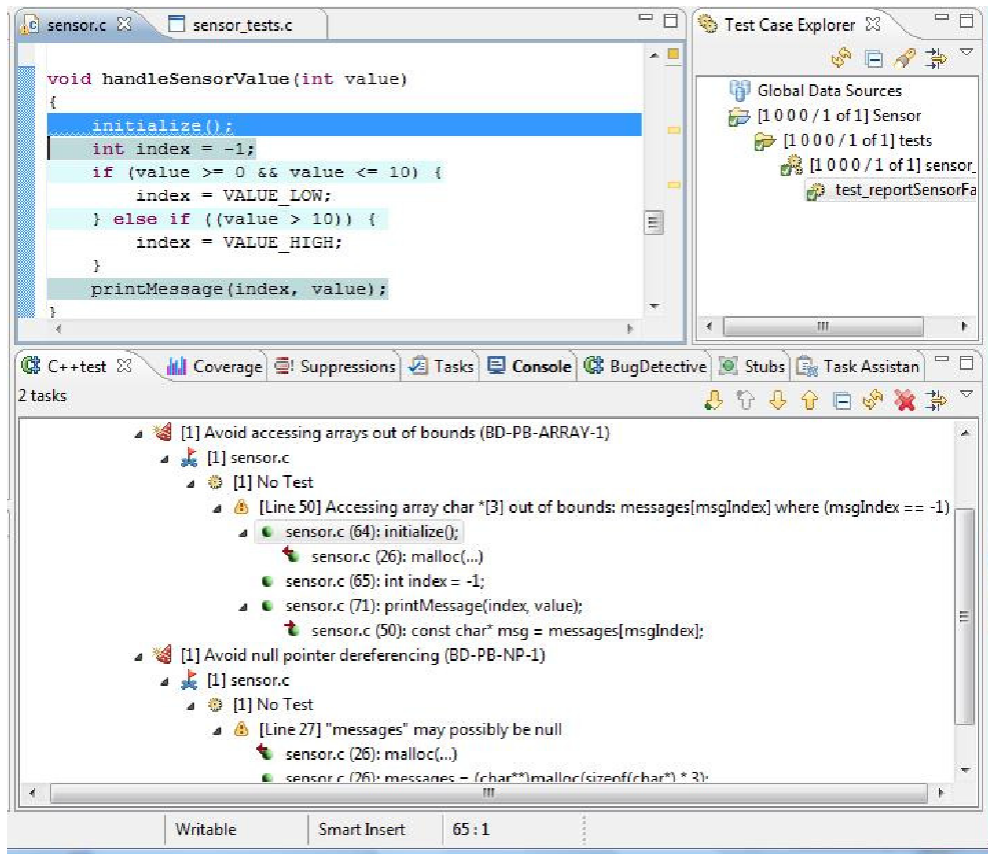
A potential path—one that was not covered—is discovered in which a double free in the `finalize()` function is possible. The `reportSensorValue()` function calls `finalize()`, then the `finalize()` calls `free()`. Also, `finalize()` is called again in the `mainLoop()`. We could fix this by making `finalize()` more intelligent, as shown below:

```

void finalize()
{
    if (messages) {
        free(messages[0]); free(messages[1]); free(messages[2]);
        free(messages); messages = 0;
    }
}

```

After running flow analysis again, only two problems are reported:



First, we may be accessing a table with `index -1` because the integral `index` is set initially to `-1`, and a possible path exists through the `if` statement that does not set this integral to the correct value before calling `printMessage()`. Runtime analysis did not lead to such a path, and it might be that such path would never be taken in the field. This is one argument for using a range of defect prevention techniques: flow analysis shows potential paths, not necessarily paths that will be taken during actual application execution. It's better to be safe than sorry, so we fix this potential error easily by removing the unnecessary condition (`value >= 0`).

```

void handleSensorValue(int value)
{
    initialize();
    int index = -1;
    if (value <= 10) {
        index = VALUE_LOW;
    } else {
        index = VALUE_HIGH;
    }
    printMessage(index, value);
}

```

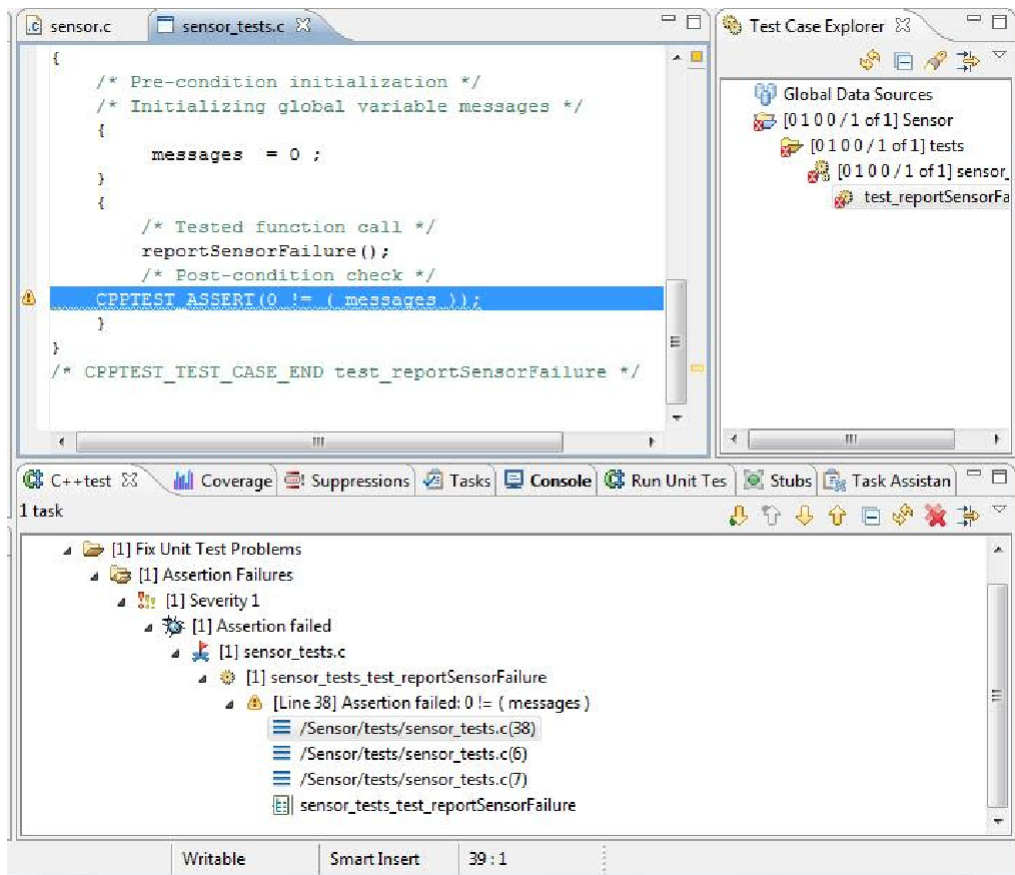
We'll use a similar method to fix the final error reported. After rerunning flow analysis, no more issues are reported.

## Regression Testing

To ensure that everything is still working, let's re-run the entire analysis. First, we run the application with runtime error detection, and everything seems fine. Then, we run unit testing with runtime error



detection, and a task is reported:



Our unit test detected a change in the behavior of the `reportSensorFailure()` function. This was caused by our modifications in `finalize()`—a change that we made in order to correct one of the previously-reported issues. This task alerts us to that change, and reminds us that we need to review the test case and then either correct the code or update the test case to indicate that this new behavior is actually the expected behavior. After looking at the code, we discover that the latter is true and we update the assertion correct condition.

```
/* CPPTEST_TEST_CASE_BEGIN test_reportSensorFailure */
/* CPPTEST_TEST_CASE_CONTEXT void reportSensorFailure(void) */

void sensor_tests_test_reportSensorFailure()
{
    /* Pre-condition initialization */
    /* Initializing global variable messages */
    {
        messages    = 0 ;
    }
    {
        /* Tested function call */ reportSensorFailure();
        /* Post-condition check */

        CPPTEST_ASSERT(0 == ( messages ));
    }
}
/* CPPTEST_TEST_CASE_END test_reportSensorFailure */
```

As a final sanity check, we run the entire application on its own—building it in the IDE without any runtime error detection. The results confirm that it is working as expected.

## Example Recap

We had a problem with our application not running as expected, and we had to decide between two approaches to resolving this: running in the debugger, or applying automated error detection techniques. If we decided to run code through the debugger, we would have observed that some variable was habitually being assigned the same value. We would have had to deduct from this that the problem was actually caused by an assignment operator being used instead of comparison.

Static analysis found this logical error for us automatically. This type of error could not have been found by runtime memory analysis because it has nothing to do with memory. It probably would not be found by flow analysis either because flow analysis traverses the execution rather than validate whether conditions are proper.

After we fixed this problem, the application ran, but it still had memory problems. Memory problems are very difficult to see under a debugger; when you are in a debugger, you don't really remember the sizes of memory. Automated tools do, however. So, to find these memory problems, we instrumented the entire application, and ran it with runtime memory analysis. This told us exactly what chunk of memory was being overwritten.

However, upon reviewing the coverage analysis results, we learned that some of the code was not covered while testing on the target. Getting this coverage information was simple since we had it tracked automatically, but if we were using a debugger, we would have had to try to figure out exactly how much of the application we verified. This is typically done by jotting notes down on paper and trying to correlate everything manually.

Once the tool alerted us to this uncovered code, we decided to leverage unit testing to add additional execution coverage to our testing efforts. Indeed, this revealed yet another problem. During normal testing on the target, covering those functions may be almost impossible because they might be hardware error-handling routines—or something else that is only executed under very rare conditions. This can be extremely important for safety critical applications. Imagine that there is a software error in code that should handle a velocity sensor problem in an airplane: instead of flagging a single device as non-working, we have a system corrupt. Creating a unit test case to cover such an execution path is very often the only way to effectively test it.

Next, we cleaned those problems and also created a regression test case by verifying the outcome (as one of the reported tasks guided us to do). Then, we ran flow analysis to penetrate paths that were not executed during testing on the target—even with the unit test. Before this, we had nearly 100% line coverage, but we did not have that level of path coverage. Flow-based static analysis uncovered some potential problems. They didn't actually happen and they might have never happened. They would surface only under conditions that were not yet met during actual execution and might never be met in real life situations. However, there's no guarantee that as the code base evolves, the application won't end up in those paths.

Just to be safe, we fixed the reported problem to eliminate the risk of it ever impacting actual application execution. While modifying the code, we also introduced a regression, which was immediately detected when we re-ran unit testing. Among these automated error detection methods, regression testing is unique in its ability to detect functional changes and verify that code modifications do not introduce functional errors or unexpected side effects. Finally, we fixed the regression, retested the code, and it all seems fine.

As you can see, all of the testing methods we applied—pattern-based static code analysis, memory analysis, unit testing, flow analysis, and regression testing—are not in competition with one another, but rather complement one another. Used together, they are an amazingly powerful tool that

provides an unparalleled level of automated error detection. In fact, such a comprehensive and integrated approach is promoted by VDC Research, who recently recommended consistent application of a broad set of complementary defect detection/prevention practices—in the context of an overarching standardized process—as a "must have" for improving software development quality and predictability.<sup>1</sup>

## Final Thoughts: Code Review and Regression Testing

In sum, by automatically finding many defects related to memory and other coding errors, we were able to get the application up and running successfully. However, it's important to remember that the deadliest defects are actually functional errors: instances where the application is not working according to specification. Unfortunately, these errors are much more difficult to find.

One of the best ways to find to find such defects is through peer code reviews. With at least one other person inspecting the code and thinking about it in context of the requirements, you gain a very good assessment of whether the code is really doing what it's supposed to.

Another helpful strategy is to create a regression test suite that frames the code, enabling you to verify that it continues to adhere to specification. In the sample scenario described above, unit testing was used to force execution of code that was not covered by application-level runtime error detection: it framed the current functionality of the application, then later, as we modified the code, it alerted us to a functionality problem. In fact, such unit test cases should be created much earlier: ideally, as you are implementing the functionality of your application. This way, you achieve higher coverage and build a much stronger "safety net" for catching critical changes in functionality.

Parasoft assists with both of these tasks: from automating and managing the peer code review workflow, to helping the team establish, continuously run, and maintain an effective regression test suite.

## About Parasoft's Development Testing Platform

Parasoft Development Testing Platform (DTP) enables Continuous Testing. Leveraging policies, DTP consistently applies software quality practices across teams and throughout the SDLC. It enables your quality efforts to shift left—delivering a platform for automated defect prevention and the uniform measurement of risk. DTP helps organizations:

- Leverage policies to align business expectations with development activities
- Prevent software defects and eliminate rework—reducing technical debt
- Focus development efforts on quality tasks that have the most impact
- Comply with internal, industry, or government standards
- Integrate security best practices into application development
- Leverage multivariate analysis to discover application hotspots that harbor hidden defects

Parasoft's Development Testing platform covers C, C++, Java, and .NET code.

## About Parasoft C/C++test

Parasoft C/C++test is an integrated Development Testing solution for automating a broad range of best practices proven to improve software development team productivity and software quality. C/C++test enables coding policy enforcement, static analysis, runtime error detection, automated peer code review, and unit and component testing to provide teams a practical way to ensure that

---

<sup>1</sup> <http://alm.parasoft.com/embedded-software-vdc-report/>

their C and C++ code works as expected. C/C++test can be used both on the desktop under common development IDEs, as well as in batch processes via command line interface for regression testing. It also integrates with Parasoft's reporting system, which provides interactive Web-based dashboards with drill-down capability, allowing teams to track project status and trends based on C/C++test results and other key process metrics.

C/C++test reduces the time, effort, and cost of testing embedded systems applications by enabling extensive testing on the host and streamlining validation on the target. As code is built on the host, an automated framework enables developers to start testing and improving code quality before the target hardware is available. This significantly reduces the amount of time required for testing on the target. The test suite built on the host can be reused to validate software functionality on the simulator or the actual target.

## About Parasoft

Parasoft researches and develops software solutions that help organizations deliver defect-free software efficiently. To combat the risk of software failure while accelerating the SDLC, Parasoft offers a Development Testing Platform and Continuous Testing Platform. Parasoft's enterprise and embedded development solutions are the industry's most comprehensive—including static analysis, unit testing, requirements traceability, coverage analysis, API testing, dev/test environment management, service virtualization and more. The majority of Fortune 500 companies rely on Parasoft in order to produce top-quality software consistently and efficiently as they pursue agile, lean, DevOps, compliance, and safety-critical development initiatives.

## Contacting Parasoft

### Headquarters

101 E. Huntington Drive, 2nd Floor  
Monrovia, CA 91016  
Toll Free: (888) 305-0041  
Tel: (626) 305-0041  
Email: [info@parasoft.com](mailto:info@parasoft.com)

### Global Offices

Visit [www.parasoft.com/contact](http://www.parasoft.com/contact) for contacting Parasoft in EMEA, APAC, and LATAM.