Lecture 09

Optimization, debugging, profiling, testing.

Advanced Programming - SISSA, UniTS, 2024-2025

Pasquale Claudio Africa

25 Nov 2024

Outline

1. Optimization

2. Debugging

3. Profiling

4. Testing

Optimization

Code optimization

Code optimization is the process of enhancing a program's performance, efficiency, and resource utilization without changing its functionality. It involves improving execution speed, reducing memory usage, and enhancing overall system responsiveness.

Optimization techniques

- **Compiler optimizations:** Utilize compiler features to automatically enhance code during compilation.

- **Algorithmic optimization:** Improve the efficiency of algorithms and data structures.

- **Manual refactoring:** Restructure code for better readability, maintainability, and performance.

- **Profiling and analysis:** Use profiling tools to identify and optimize performance bottlenecks.

Optimization options

The compiler enhances performance by optimizing CPU register usage, expression refactoring, and pre-computing constants.

- Disable optimization during debugging.

- Pass the `-O{n}` (`n={0,1,2,s,3}`) flag to the compiler to control optimization level, with `-Os` for space optimization and `-O3` for maximum optimization. Here a detailed list of optimizations enabled with each flag.

- Defining the `-DNDEBUG` preprocessor variable, standard assertion are ignored, resulting in faster code.

Loop unrolling

It is beneficial to unroll small loops. For example, transform:

```cpp
for (int i = 0; i < n; ++i) {
  for(int k = 0; k < 3; ++k) {
    a[k] += b[k] * c[i];
  }
}
```

to:

```cpp
for (int i = 0; i < n; ++i) {
  a[0] += b[0] * c[i];
  a[1] += b[1] * c[i];
  a[2] += b[2] * c[i];
}
```

Compiler may unroll loops with `-funroll-loops`, but better performance isn't guaranteed.

Prefetching constant values

Prefetch constant values inside the loop for further optimization:

```cpp
for (int i = 0; i < n; ++i) {
  auto x = c[i];
  a[0] += b[0] * x;
  a[1] += b[1] * x;
  a[2] += b[2] * x;
}
```

Avoid `if` inside nested loops

`if` statements, especially in nested loops, can be costly. Consider these improvements:

```cpp
for(int i = 0; i < 10000; ++i) {
  for (int j = 1; j < 10; ++j) {
    if(c[i] > 0)
      a[i][j] = 0;
    else
      a[i][j] = 1;
  }
}
// Better:
for(int i = 0; i < 10000; ++i)
  if(c[i] > 0)
    for(int j = 0; j < 10; ++j) {
      a[i][j] = 0;
  else
    for(int j = 0; j < 10; ++j) {
      a[i][j] = 1;
}
```

Sum of a vector: two strategies compared

```cpp
double sum1(double *data, const size_t &size) {
    double sum{0};
    for (size_t j = 0; j < size; ++j)
        sum += data[j];
    return sum;
}

double sum2(double *data, const size_t &size) {
    double sum{0}, sum1{0}, sum2{0}, sum3{0};
    size_t j;
    for (j = 0; j < (size - 3); j += 4) {
        sum += data[j + 0];
        sum1 += data[j + 1];
        sum2 += data[j + 2];
        sum3 += data[j + 3];
    }
    for (; j < size; ++j)
        sum += data[j];
    sum += sum1 + sum2 + sum3;
    return sum;
}
```

Which one is faster, `sum1` or `sum2`?

The number of floating point operations is the same in both cases!

The answer is not straightforward: it depends on the computer's architecture.

On my laptop (Intel(R) Core(TM) Ultra 7 155H CPU @ 4.80GHz), `sum2` **is approximately 10 times faster than** `sum1` with `size = 1e9` ! (see `examples/unrolling/unrolling.cpp`).

Why? The Streaming SIMD Extensions (SSE2) instruction set of the CPU allows for parallelization at the microcode level. It's a super-scalar architecture with multiple instruction pipelines to execute several instructions concurrently during a clock cycle. The code of `test2` better exploits this capability.

Take-home message: Counting operations doesn't necessarily reflect performance. Compiler optimizers can transform `sum1` into `sum2` automatically. Sometimes, giving it a hand is beneficial.

Cache friendliness

Efficiency often depends on how variables are accessed in memory. Access variables contiguously for cache pre-fetching effectiveness. For example, if `mat` is a dynamic matrix organized **row-wise**:

```cpp
// Not cache-friendly, inefficient.
for (j = 0; j < n_cols; ++j) {
  for (i = 0; i < n_rows; ++i) {
    a += mat(i, j);
  }
}

// Cache-friendly, thus more efficient.
for (i = 0; i < n_rows; ++i) {
  for (j = 0; j < n_cols; ++j) {
    a += mat(i, j);
  }
}
```

Debugging

Static analysis vs. debugging (1/2)

Static analysis

- **Nature:** Examines code without executing it.

- **Purpose:** Identifies potential issues and coding standards violations.

- **Tools:** Code linters, security scanners, and complexity analyzers.

- **Integration:** Often part of development workflows or continuous integration.

Debugging

- **Nature:** Inspects and troubleshoots code during runtime.

- **Purpose:** Locates and resolves bugs, runtime errors, and unexpected behavior.

- **Tools:** Debuggers with features like breakpoints and variable inspection.

- **Integration:** Interactive process during development or post-runtime.

Static analysis vs. debugging (2/2)

Key differences

- **Timing:** Static analysis is pre-runtime; debugging is during or post-runtime.

- **Focus:** Static analysis emphasizes code quality; debugging resolves runtime issues.

- **Use cases:** Static analysis is proactive; debugging is reactive.

- **Automation:** Static analysis tools can be automated; debugging is more interactive.

- **Complementarity:** Both are complementary, with static analysis preventing issues and debugging addressing runtime problems.

Static analysis

Static analysis tools analyze source code by inspecting it for potential issues, vulnerabilities, or adherence to coding standards. Common ones include:

- cppcheck
- cpplint
- clang-tidy

Some of the checks they perform:

- Automatic variable checking.
- Bounds checking for array overruns.
- Unused functions, variable initialization and memory duplication.
- Invalid usage of Standard Template Library functions and idioms.
- Missing `#include`s.
- Memory or resource leaks, performance errors.

Other useful tools

- Cling is an interactive C++ interpreter, built on LLVM and Clang. It's part of the ROOT project at CERN and can be integrated into a Jupyter workspace (see here). While experimental, an interpreter aids in code prototyping.

- Compiler Explorer to check how code translates into assembly language.

- C++ Insights allows viewing source code through a compiler's eyes.

Debuggers

Debuggers are software tools that enable developers to inspect, analyze, and troubleshoot code during the development process. They provide a set of features for identifying and fixing errors in programs.

Key features

- **Breakpoints:** Pauses program execution at specified points to inspect variables and code.
- **Variable inspection:** Allows developers to examine the values of variables during runtime.
- **Step-through execution:** Enables line-by-line execution for precise debugging.
- **Call stack analysis:** Displays the sequence of function calls leading to the current point in code.

Debugging

During code development, debugging allows step-by-step execution. To use a debugger, compile with `-g` (which implies no optimization). `-g` adds information for locating source lines in machine code.

Two debugging types:

- **Static debugging:** Analyze core dump if code aborts.

- **Dynamic debugging:** Execute through a debugger, breaking at points to examine variables.

Two common debuggers are `gdb` and `lldb`. See, e.g.:

- Debugging with GDB
- GDB cheat sheet
- GDB tutorial

Debugging levels

Debugging can be at different levels, and using `-g` `-O` together is allowed. `-g` tells the compiler to provide extra information for the debugger. However, line-by-line debugging reliability decreases with optimization. `-g` implies `-O0` by default.

Debugging levels and special optimization options linked to debugging:

- `-g0` : No debugging information.
- `-g1` : Minimal information for backtraces.
- **The default debugging level is 2.**
- `-g3` : Extra information, including macro definitions.
- `-Og` : Special optimization option. Enables optimizations without interfering with debugging.

Main commands of `gdb` / `lldb`

- `run` : Run the program.

- `break` : Set a breakpoint at a line/function.

- `where` : Show location and backtrace.

- `print` : Display variable/expression value.

- `list n` : Show lines around line n.

- `next` : Go to the next instruction, proceeding through subroutines.

- `step` : Go to the next instruction, entering called functions.

- `continue` : Continue executing.

- `backtrace` : Print memory stack after program aborts.

- `quit` : Exit the debugger.

- `help` : Displays help information.

Other debugging tools

`valgrind` , a suite of tools for debugging and profiling. It can find memory leaks, unassigned variables, or check memory usage:

Find memory leaks:

```
valgrind --tool=memcheck --leak-check=yes --log-file=file.log executable
```

Check memory usage:

```
valgrind --tool=massif --massif-out-file=massif.out --demangle=yes executable
ms_print massif.out > massif.txt
```

`massif.txt` indicates memory usage during the program execution.

Profiling

Profilers

A **profiler** in software development is a tool or set of tools designed to analyze the runtime behavior and performance of a computer program. It provides detailed information about resource utilization, execution times, and function calls during the program's execution.

Key objectives

- **Performance analysis:** Profilers offer insights into how much time a program spends in different functions, helping identify performance bottlenecks.

- **Resource usage:** They measure memory consumption, CPU utilization, and other system resources, aiding in optimizing resource-intensive operations.

- **Function call tracing:** Profilers track the sequence of function calls, enabling developers to understand the flow of execution.

gprof

`gprof` is the GCC simple profiler. In order to use it, compile the code with the `-pg` option at both the compilation and linking stages.

When executing the code, it generates a file called `gmon.out`, which is then utilized by the profiler:

```
gprof --demangle executable > file.txt
```

Then `file.txt` will contain valuable information about the program execution.

Main options of `gprof`

`gprof` offers a range of options. The main ones are:

- `--annotated-source[=symspec]` : Prints annotated source code. If `symspec` is specified, print output only for matching symbols.

- `-I dirs` : List of directories to search for source files.

- `--graph[=symspec]` : Prints the call graph analysis.

- `--demangle` : Demangles mangled names (essential for C++ programs).

- `--display-unused-functions` : As the name says.

- `--line` : Line-by-line profiling (but maybe better use gcov).

`callgrind`

`callgrind` is a tool of `valgrind` that you may call, for instance, as:

```
valgrind --tool=callgrind --callgrind-out-file=grind.out --dump-line=yes ./myprog
```

Compile the program with `-g` and **optimization activated**. The option `--dump-lines` is used for line-by-line profiling.

Afterward, post-process the binary file `grind.out`, e.g., using `kcachegrind`:

```
kcachegrind grind.out
```

It opens a graphical interface.

Other profilers

There are alternative profilers, some useful in a parallel environment:

- perf : Lightweight CPU profiling.

- gperftools : Formerly Google Performance Tools.

- TAU (Tuning and Analysis Utilities) : Profiling and tracing toolkit for parallel programs.

- Scalasca : Performance analysis for parallel applications on distributed memory systems.

Testing

Verification vs. validation

Verification: **Ensuring correct implementation**

Conducted during development, tests **individual components** separately. Specific tests demonstrate correct functionality, covering the code and checking for memory leaks.

Validation: **Confirming desired behavior**

Performed on the **final code**. Assesses if the code produces the intended outcome - convergence, reasonable results, and expected computational complexity.

Types of testing

- **Unit testing:** Testing individual components (functions, methods, or classes) to ensure each behaves as expected. It focuses on a specific piece of code in isolation.

- **Integration testing:** Verifying that different components/modules of the software work together as intended. It deals with interactions between different parts of the system.

- **Regression testing:** Ensuring recent code changes do not adversely affect existing functionalities. It involves re-running previous tests on the modified codebase to catch unintended side effects.

Importance of testing

- **Early detection of bugs:** Testing allows early detection and fixing of bugs, reducing the cost and time required for debugging later in the development process.

- **Code reliability:** Testing ensures the code behaves as expected and provides reliable results under different conditions.

- **Documentation:** Test cases serve as documentation for how different parts of the code are expected to work/to be used. They help other developers understand the intended behavior of functions and classes.

Unit testing in C++

In C++, unit testing often uses frameworks like Google Test , Catch2 , or CTest itself (from the CMake ecosystem).

Here's a simple example using `gtest` :

```cpp
#include "mylibrary.h"
#include "gtest/gtest.h"

TEST(MyLibrary, AddTwoNumbers) {
  EXPECT_EQ(add(2, 3), 5);
}

int main(int argc, char **argv) {
  ::testing::InitGoogleTest(&argc, argv);
  return RUN_ALL_TESTS();
}
```

In this example, we test the `add` function from the `mylibrary` module.

Test-Driven Development (TDD)

TDD is a software development approach where tests are written before the actual code. The cycle is writing a test, implementing the code to pass the test, and refactoring.

Advantages

TDD encourages modular and testable code, ensuring all parts of the codebase are covered by tests. It also starts by thinking at how code should be used (bottom-up strategy), possibly guiding the design of the interface exposed.

Process

1. Write a test defining a function or improvements succinctly.
2. Run the test to ensure it fails, showing it doesn't pass.
3. Write the simplest code to make the test pass.
4. Run the test and verify it passes.
5. Refactor the code for better structure or performance.

Continuous Integration (CI) and testing

- **CI:** Frequently integrating code changes into a shared repository. Automated builds and tests ensure new changes don't break existing functionalities.

- **Benefits:**

  - Early detection of integration issues.

  - Regular validation of code against the test suite.

  - Confidence in the stability of the codebase.

- **Popular CI Tools:**

  - Jenkins

  - Travis CI

  - GitHub Actions

  - GitLab CI

Coverage

Code coverage is a metric used in software testing to measure the extent to which source code is executed during the testing process. It provides insights into which parts of the codebase have been exercised by the test suite and which parts remain untested.

Key concepts

- **Lines of code**: Code coverage is often expressed as a percentage of lines of code that have been executed by tests. The goal is to have as close to 100% coverage as possible.

- **Branches** and **paths**: In addition to lines, code coverage can also consider branches and execution paths within the code. This provides a more detailed analysis of the code's behavior.

Coverage with gcov

`GCC` supports program coverage with `gcov`. Compile with `-g -fprofile-arcs -ftest-coverage` and no optimization. For shared objects with `dlopen`, add the option `-Wl,--dynamic-list-data`.

Run the code, producing `gcda` and `gcno` files. Use `gcov` utility:

```
gcov [options] source_file_to_examine [or executable]
```

Text files with code and execution counts for each line are created.

Main options of gcov

`gcov` offers various options:

- `--demangled-names`: Demangle names, useful for C++.
- `--function-summaries`: Output summaries for each function.
- `--branch-probabilities`: Write branch frequencies to the output file.

`lcov` `and` `genhtml`: nice graphical tools for `gcov`

The `gcov` output is verbose. With `lcov` and `genhtml`, you get a graphical view:

Compile with `gcov` rules, then:

```
lcov --capture --directory project_dir --output-file cov.info
genhtml cov.info --output-directory html
```

`project_dir` is the directory with `gcda` and `gcno` files. In the `html` directory, open `index.html` in your browser.

➡️ **Introduction to Python.**