

# Debugging with GDB, Valgrind, and Sanitizers

How to make C/C++ debugging bearable

Lars Quentin



# The Scientific Method of Debugging

## Steps

### 1 **Reproduce:**

- ▶ Find bad inputs

### 2 **Isolate:**

- ▶ Create **MINIMAL** version

### 3 **Hypothesize:**

- ▶ Try to think why it could happen

### 4 **Test:**

- ▶ Validate or discard afterwards

# The Scientific Method of Debugging

## Steps

- 1 Reproduce:**
  - ▶ Find bad inputs
- 2 Isolate:**
  - ▶ Create **MINIMAL** version
- 3 Hypothesize:**
  - ▶ Try to think why it could happen
- 4 Test:**
  - ▶ Validate or discard afterwards

## General Advice

- **Be Radical:**
  - ▶ Build an absolutely minimal version
- **No really...**
  - ▶ Can you go smaller?
- **Don't assume anything:**
  - ▶ (See JL-example)
- **Use all tools you can get:**
  - ▶ GCC is your friend
- **Disable all optimizations**

# Outline

- 1 Intro
- 2 GCC**
- 3 GDB
- 4 Sanitizers
- 5 Valgrind
- 6 Conclusion

# Offloading work to the Compiler

## 1 Enable debug symbols:

- ▶ -g: variable names, line numbers, ...
- ▶ -g3: More information, such as macro defs (native DWARF)
- ▶ -ggdb3: If you know you use gdb

# Offloading work to the Compiler

## 1 Enable debug symbols:

- ▶ -g: variable names, line numbers, ...
- ▶ -g3: More information, such as macro defs (native DWARF)
- ▶ -ggdb3: If you know you use gdb

## 2 Disable optimization

- ▶ -O0 if you want to be safe
- ▶ -Og if you trust gcc (I don't)

# Offloading work to the Compiler

## 1 Enable debug symbols:

- ▶ -g: variable names, line numbers, ...
- ▶ -g3: More information, such as macro defs (native DWARF)
- ▶ -ggdb3: If you know you use gdb

## 2 Disable optimization

- ▶ -O0 if you want to be safe
- ▶ -Og if you trust gcc (I don't)

## 3 Enjoy modern language: -std=...

# Offloading work to the Compiler

## 1 Enable debug symbols:

- ▶ -g: variable names, line numbers, ...
- ▶ -g3: More information, such as macro defs (native DWARF)
- ▶ -ggdb3: If you know you use gdb

## 2 Disable optimization

- ▶ -O0 if you want to be safe
- ▶ -Og if you trust gcc (I don't)

## 3 Enjoy modern language: -std=...

## 4 Enable STL checks (more later)

# Offloading work to the Compiler

## 1 Enable debug symbols:

- ▶ -g: variable names, line numbers, ...
- ▶ -g3: More information, such as macro defs (native DWARF)
- ▶ -ggdb3: If you know you use gdb

## 2 Disable optimization

- ▶ -O0 if you want to be safe
- ▶ -Og if you trust gcc (I don't)

## 3 Enjoy modern language: -std=...

## 4 Enable STL checks (more later)

## 5 Enable all warnings:

# Offloading work to the Compiler

## 1 Enable debug symbols:

- ▶ -g: variable names, line numbers, ...
- ▶ -g3: More information, such as macro defs (native DWARF)
- ▶ -ggdb3: If you know you use gdb

## 2 Disable optimization

- ▶ -O0 if you want to be safe
- ▶ -Og if you trust gcc (I don't)

## 3 Enjoy modern language: -std=...

## 4 Enable STL checks (more later)

## 5 Enable all warnings:

- ▶ Just -Wall... right?

## Enable all relevant warnings

- `-Wall`: Enables “all” common warnings

## Enable all relevant warnings

- `-Wall`: Enables “all” common warnings
- `-Wextra`: Even more warnings! ...but now we have all right?

## Enable all relevant warnings

- `-Wall`: Enables “all” common warnings
- `-Wextra`: Even more warnings! ...but now we have all right?
- `-Wpedantic`: ISO compliance, relevant for -O2 Heisenbugs

## Enable all relevant warnings

- `-Wall`: Enables “all” common warnings
- `-Wextra`: Even more warnings! ...but now we have all right?
- `-Wpedantic`: ISO compliance, relevant for -O2 Heisenbugs
- `-Wshadow`: Local variable overshadows another (local, global) variable

## Enable all relevant warnings

- `-Wall`: Enables “all” common warnings
- `-Wextra`: Even more warnings! ...but now we have all right?
- `-Wpedantic`: ISO compliance, relevant for -O2 Heisenbugs
- `-Wshadow`: Local variable overshadows another (local, global) variable
- `-Wconversion`: Warns against conversions that might alter an value

## Enable all relevant warnings

- `-Wall`: Enables “all” common warnings
- `-Wextra`: Even more warnings! ...but now we have all right?
- `-Wpedantic`: ISO compliance, relevant for -O2 Heisenbugs
- `-Wshadow`: Local variable overshadows another (local, global) variable
- `-Wconversion`: Warns against conversions that might alter an value
- `-Wcast-align`: “Warn whenever a pointer is cast such that the required alignment of the target is increased. For example, warn if a `char *` is cast to an `int *` on machines where integers can only be accessed at two- or four-byte boundaries. ”

## Enable all relevant warnings

- `-Wall`: Enables “all” common warnings
- `-Wextra`: Even more warnings! ...but now we have all right?
- `-Wpedantic`: ISO compliance, relevant for -O2 Heisenbugs
- `-Wshadow`: Local variable overshadows another (local, global) variable
- `-Wconversion`: Warns against conversions that might alter an value
- `-Wcast-align`: “Warn whenever a pointer is cast such that the required alignment of the target is increased. For example, warn if a `char *` is cast to an `int *` on machines where integers can only be accessed at two- or four-byte boundaries. ”

### Notes

- `-Werror` is good for dev, but leave it off for others please.
- See <https://gcc.gnu.org/onlinedocs/gcc/Warning-Options.html>

# Asserts

- Remember: Asserts are **zero-cost in release mode!**
- Toggled via NDEBUG
- Use many asserts wherever you want ⇒ Never too many
- Assert the impossible (negative) not the expected (positive)
- Best Practise: No side effects!
  - ▶ **Bad:** `assert(p = malloc(10));`
  - ▶ **Good:** `void *p = malloc(10); assert(p != NULL);`
- Compile time checks via `static_assert`
  - ▶ `static_assert(sizeof(int) >= 4, "Weird embedded system");`

# STL guarantees

## Two possibilities

### 1 -D\_GLIBCXX\_ASSERTIONS

- ▶ Enables out of bounds checks
- ▶ Actually overwrites operator[]
- ▶ No risk in using it

### 2 -D\_GLIBCXX\_DEBUG

- ▶ Enables out of bounds **and iterator checks**
- ▶ **CHANGES ABI**
  - All your dependencies (Boost etc) also need -D\_GLIBCXX\_DEBUG
- ▶ You need to know what you're doing

# Pro-Tipp: printf-style debugging

Simplest implementation

```
1  template<class T>
2  void dbg_simple(T t) {
3  #ifndef NDEBUG
4      std::cerr << "[Debug] " << t << std::endl;
5  #endif
6  }
7
```

## Pro-Tipp: printf-style debugging

### Zero cost printf macros

```
1  #ifndef NDEBUG
2  void dbg(const auto &val,
3         const std::source_location loc = std::source_location::current()) {
4
5     // Get time struct
6     using namespace std::chrono;
7     auto t{system_clock::to_time_t(system_clock::now())};
8     struct tm tb; localtime_r(&t, &tb);
9
10    // Print out
11    std::cerr << "[" << std::put_time(&tb, "%H:%M:%S")
12    << "]" " << loc.function_name() << ":" << loc.line() << " | " << val << '\n';
13 }
14 #else
15 constexpr void dbg(const auto &) noexcept {}
16 #endif
```

# Pro-Tipp: printf-style debugging

C version: Only works for char \*

```
1  #ifndef NDEBUG
2  #include <stdio.h>
3  #include <time.h>
4
5  #define dbg_str(val) do { \
6      time_t _t = time(NULL); \
7      struct tm _tb; \
8      localtime_r(&_t, &_tb); \
9      fprintf(stderr, "[%02d:%02d:%02d] %s:%d | %s\n", \
10         _tb.tm_hour, _tb.tm_min, _tb.tm_sec, __func__, __LINE__, val); \
11 } while(0)
12 #else
13 #define dbg_str(val) ((void)0)
14 #endif
```

# Outline

- 1 Intro
- 2 GCC
- 3 GDB**
- 4 Sanitizers
- 5 Valgrind
- 6 Conclusion

# GDB: The GNU Debugger

## What is it?

- The de-facto debugger on Linux
- Made for C/C++, can do a lot more
- **Advantages over printf:**
  - ▶ Look into crashes (segfaults)
  - ▶ Step Over/Into/Out
  - ▶ Call Stack
  - ▶ Non-Invasive

## Rules for good use

- 1 Compile with `-ggdb3 -O0`
- 2 Compile without sanitizers
- 3 Realize Heisenbugs can exist
  - ▶ They obviously change timings
- 4 As Han Solo said:  
“Good luck... you’re gonna need it”

# HELP I HAVE A SEGFAULT (80% of all use cases)

## What to do

- 1 Deep breath
- 2 Run the program inside gdb
- 3 Let it crash (SIGSEGV)
- 4 Look at the *backtrace* how you got there
- 5 Look at the local state to find the condition
- 6 If unclear: Real debugging begins

# Livedemo: HELP I HAVE A SEGFAULT (80% of all use cases)

What to do

```
1  # Compile at least unoptimized with debug symbols
2  gcc -O0 -ggdb3 ./my_program.c -o my_program
3  gdb ./my_program
4
5  (gdb) run <args> # Run the program (with optional args)
6  # wait for SIGSEGV...
7  (gdb) bt # Print the call stack
8  (gdb) bt full # Print the call stack WITH local variables
```

# GDB TUI Mode

- GDB has a bad rep because people think its cryptic
- But: It has a Text User Interface (TUI) mode!

How to use it

- **At start:** `gdb -tui ./my_program`
- **On the fly:** Press `Ctrl-X` followed by a

Note: TUIs and `stdout`

- **Problem:** `stdout` can break TUI mode
- **Solution:** Force redraw via `Ctrl-L`

## GDB: Most important commands for stepping

- `break main.cc:42` (Set a breakpoint)
- `next (n)` (Step over functions)
- `step (s)` (Step into functions)
- `finish (fin)` (Step out of functions)
- `advance main.cc:42` (Continue until you are there)
- `print var (p)` (Show value of variable)
- `display var` (Always show value of variable)

# Life Hack: gdb - dashboard

- Python based gdb config with dashboard capabilities
- Looks good, nooby friendly
- Single new command: dashboard
- Even syntax highlighting! (with pysegments)
- [github.com/cyrus-and/gdb-dashboard](https://github.com/cyrus-and/gdb-dashboard)

```
Assembly
0x0005555555514a save_to_memory+1 mov %rsp,%rbp
0x0005555555514d save_to_memory+4 mov %edi,-0x14(%rbp)
0x00055555555150 save_to_memory+7 movq $0x0,-0x8(%rbp)
0x00055555555158 save_to_memory+15 mov -0x8(%rbp),%rax
0x0005555555515c save_to_memory+19 mov -0x14(%rbp),%edx
0x0005555555515f save_to_memory+22 mov %edx,(%rax)
0x00055555555161 save_to_memory+24 nop
0x00055555555162 save_to_memory+25 pop %rbp
0x00055555555163 save_to_memory+26 ret

Breakpoints
Expressions
History
Memory
Registers
rax 0x0000000000000000 rbx 0x0007ffffffe248 rcx 0x2c49d3daf6b8660
rdx 0x0000000000000014 rsi 0x000555555592a8 rdi 0x000000000000014
rbp 0x0007ffffffe0e8 rsp 0x0007ffffffe0e8 r8 0x0000000000000000
r9 0x0000000000000000 r10 0x0000000000000000 r11 0x000000000000202
r12 0x0000000000000000 r13 0x0007ffffffe258 r14 0x0007ffff7ff0000
r15 0x000555555557dd8 rip 0x0005555555515f eflags [ PF IF RF ]
cs 0x00000033 ss 0x0000002b ds 0x00000000
es 0x00000000 fs 0x00000000 gs 0x00000000
fs_base 0x0007fff7dab740 gs_base 0x0000000000000000

Source
1 #include <stdio.h>
2
3 void save_to_memory(int value) {
4     int *ptr = NULL;
5     *ptr = value;
6 }
7
8 void process_results(int total) {
9     printf("sth sth processing\n");
10 }

Stack
[0] from 0x0005555555515f in save_to_memory+22 at test.c:5
[1] from 0x00055555555168 in process_results+38 at test.c:10
[2] from 0x000555555551a5 in main+74 at test.c:19

Threads
[1] id 13560 name a.out from 0x0005555555515f in save_to_memory+22 at test.c:5

Variables
arg value = 20
loc ptr = 0x0: Cannot access memory at address 0x0
```

# Outline

- 1 Intro
- 2 GCC
- 3 GDB
- 4 Sanitizers**
- 5 Valgrind
- 6 Conclusion

# What are Sanitizers?

- Compiler features that add extra checks (instrumentation) at compile time
  - ▶ Runtime analysis, no static analysis!
  - ▶ Very fast compared to valgrind (ca 2x slowdown)
- Originally developed by Google for LLVM/Clang
  - ▶ Also ported to GCC later
- Compile rules:
  - ▶ -g3 for debug symbols
  - ▶ -O0 for easier debugging
  - ▶ **-fno-omit-frame-pointer** for error traces!
- **Recommendation:** (ASan + UBSan) when developing

# AddressSanitizer (ASan) and LeakSanitizer (LSan)

- **Compiler flag:** `-fsanitize=address`
- **What it does:** Tracks allocations and puts redzones around variables to detect illegal access
- **Examples:**
  - ▶ Heap out-of-bounds (r/w after `malloc`-region)
  - ▶ Stack out-of-bounds (r/w past a local array)
  - ▶ Use-after-free
  - ▶ Double-free
- **LSan included:** LeakSanitizer is automatically enabled as well.
  - ▶ Checks at exit if everything was freed (or thus leaked)

# ASAN (LSan) Live Example (Who can see the error)

A good C++ compile command

```
1  #include<stdlib.h>
2  #include<stdio.h>
3
4  int main() {
5      size_t N=10;
6      printf("Allocating array of %zu integers...\n", N);
7      int *arr = malloc(N * sizeof(int));
8
9      for (int i = 0; i <= N; ++i)
10         arr[i] = i * 2;
11
12     for (int i = 0; i <= N; ++i)
13         printf("%d ", arr[i]);
14     printf("\n");
15 }
```

# UndefinedBehaviorSanitizer (UBSan)

- **Compiler flag:** `-fsanitize=undefined`
- **What it does:** Detects undefined behavior at RUNTIME (i.e. input dependent)
- **Why its needed:** Compilers are allowed to assume no UB happens
  - ▶ Otherwise: You *might* lose correctness (and might is even worse to debug!)
- **Examples:**
  - ▶ Signed integer overflow
  - ▶ Division by zero
  - ▶ Unaligned pointer aliasing
  - ▶ Shifting out of bounds (`1u >> 32`)

## Other sanitizers (use carefully)

### ThreadSanitizer (TSan)

- `-fsanitize=thread`
- Concurrent data accesses without locking
- Deadlocks
- Problems:
  - ▶ Up to 10x slower
  - ▶ Cannot be combined with ASan+LSan+UBSan

## Other sanitizers (use carefully)

### ThreadSanitizer (TSan)

- `-fsanitize=thread`
- Concurrent data accesses without locking
- Deadlocks
- Problems:
  - ▶ Up to 10x slower
  - ▶ Cannot be combined with ASan+LSan+UBSan

### MemorySanitizer (MSan)

- `-fsanitize=memory`
- Reading uninitialized variables
- if-branching based on it
- **PROBLEMS:**
  - ▶ Requires ALL linked libraries to be recompiled. Including `libc++`, your MPI, ...
  - ▶ Very easy to get wrong, especially on clusters.
- If you really need it: Try a static analyzer first...

# Outline

- 1 Intro
- 2 GCC
- 3 GDB
- 4 Sanitizers
- 5 Valgrind**
- 6 Conclusion

# Valgrind

## What is Valgrind

- Dynamic analysis tool
- Spiritual predecessor to sanitizers
- Synthetic CPU that takes *any* compiled binary, adds instrumentation, and runs it

## Valgrind tools (selection)

- **Memcheck:** Memory error detector
- **Cachegrind:** Cache and branch-prediction profiler
- **Callgrind:** Call-graph generating cache profiler
  - ▶ **KCachegrind** GUI for analysis

# Valgrind Memcheck

- Memcheck is the default tool in Valgrind.
- It detects memory-management problems
- Roughly equivalent to ASan+LSan+**MSan**
  - ▶ Works without recompilation! (opposed to MSan)
- Can detect if your program
  - ▶ Leaks memory (with stack traces where it was allocated)
  - ▶ Missuses uninitialized values
  - ▶ Accesses memory incorrectly (such as alignment or out of bounds)
  - ▶ Does illegal frees of heap blocks (double free, stack frees, free on new allocs)
  - ▶ Does illegal reads and writes (for example an overlapping memcpy)
- **NOTE:** Can be flaky with high optimization levels.

# Comparison Memcheck Sanitizers

## ■ Advantages Valgrind

- ▶ No recompilation required
- ▶ Good MSan alternative
- ▶ Allows blackbox debugging (into any third party libraries)
- ▶ No compiler support required

## ■ Disadvantages Valgrind

- ▶ **Very** slow (20-30x overhead, order of magnitude higher than sanitizers)
- ▶ High memory overhead (since it tracks everything)
- ▶ Can get confused when doing vectorization (-O3)

## ■ Personal recommendation:

Use Sanitizers when dev, Valgrind when you have to (you'll know)

# Memcheck: Example

Use of uninitialized values in syscalls

```
1 $ cat > uninit_syscall.c <<'EOF'
2 #include <stdlib.h>
3 #include <unistd.h>
4 int main( void )
5 {
6     char* arr = malloc(10);
7     int* arr2 = malloc(sizeof(int));
8     write( 1 /* stdout */, arr, 10 );
9     exit(arr2[0]);
10 }
11 EOF
```

# Memcheck: Example

## Use of uninitialized values in syscalls

```
1 $ gcc -o uninit_syscall uninit_syscall.c -Wall -g
2 $ module load valgrind
3 $ valgrind ./uninstall_syscall
4 ...
5 Syscall param write(buf) points to uninitialised byte(s)
6   at 0x25A48723: __write_nocancel (in /lib/tls/libc-2.3.3.so)
7   by 0x259AFAD3: __libc_start_main (in /lib/tls/libc-2.3.3.so)
8 Address 0x25AB8028 is 0 bytes inside a block of size 10 alloc'd
9   at 0x259852B0: malloc (vg_replace_malloc.c:130)
10  by 0x80483F1: main (a.c:5)
11 Syscall param exit(error_code) contains uninitialised byte(s)
12 ...
```

# Outline

- 1 Intro
- 2 GCC
- 3 GDB
- 4 Sanitizers
- 5 Valgrind
- 6 Conclusion**

# How to compile safely (For Sanitizers/General)

A good C++ compile command

```
1 g++ -Wall -Wextra -Wpedantic \  
2     -Wshadow -Wconversion -Wcast-align \  
3     -ggdb3 -O0 -std=c++23 \  
4     -fsanitize=address -fsanitize=undefined \  
5     -fno-omit-frame-pointer \  
6     -D_GLIBCXX_ASSERTIONS \  
7     ./src/main.cc -o main
```

# How to compile safely (For gdb/valgrind)

A good C++ compile command

```
1  # NO SANITIZERS
2  g++ -Wall -Wextra -Wpedantic \
3      -Wshadow -Wconversion -Wcast-align \
4      -ggdb3 -O0 -std=c++23 \
5      -fno-omit-frame-pointer \
6      -D_GLIBCXX_ASSERTIONS \
7      ./src/main.cc -o main
```

## Conclusion: What to use when

- **Logic Errors:** GDB or printf-debugging
- **Crashes (segfaults):** GDB or ASAN+UBSan
- **Memleaks:** Valgrind or LSan
- **Uninitialized Memory:** Valgrind
- **Data Races/Unintentional Indeterminism:** TSan
- **Problem with closed-source lib:** Valgrind