SBT-instrumentation: A Tool for Configurable Instrumentation of LLVM Bitcode

Martina Vitovská, Marek Chalupa, and Jan Strejček

Masaryk University, Brno, Czech Republic
`{xvitovs1, xchalup4, strejcek}@fi.muni.cz`

Abstract

The paper describes a member of the SYMBIOTIC toolbox called `sbt-instrumentation`, which is a tool for configurable instrumentation of LLVM bitcode. The tool enables a user to specify patterns of instructions and to define functions whose calls will be inserted before or after instructions that match the patterns. Moreover, the tool offers additional functionality. First, the instrumentation can be divided into phases in order to pass information acquired in an earlier phase to the later phases. Second, it can utilize results of some external static analysis by connecting it as a plugin. The `sbt-instrumentation` tool has been developed as the part of SYMBIOTIC responsible for inserting memory safety checks. However, its configurability opens the way to use it for many various purposes.

1 Introduction

Instrumentation is a technique widely used in the field of program analysis and verification that inserts an auxiliary code into the program that is being analyzed. A number of well-known tools use instrumentation to achieve their goals. For example, CLANG's AddressSanitizer [1] uses compile-time instrumentation to check for memory errors, JProfiler [2] uses instrumentation for profiling of Java applications, and CBMC [3] inserts code checking various defects like overflows or memory errors. However, the tools use mostly a single-purpose instrumentation that cannot be easily modified. In this paper, we present a tool `sbt-instrumentation` [4] that offers general instrumentation of LLVM bitcode [5] that is easily configurable according to the user's choice.

The tool was primarily designed for inserting memory safety checks that can detect invalid dereferences, invalid deallocations, and memory leaks. Roughly speaking, memory safety checking needs to insert two types of functions:

- tracking functions that build and maintain records about allocated memory blocks and

- checking functions that use these records to decide the validity of memory operations.

This initial application [6] brought several interesting and non-standard features to our tool. First, insertion of many memory safety checks can be avoided as a static pointer analysis can guarantee that many dereferences and deallocations are safe. Hence, `sbt-instrumentation` supports plugins performing pointer or other analyses and insertions of checks can be conditioned by an output of these plugins. Further, one needs to track only memory blocks relevant for inserted checks. Thus, the instrumentation first inserts checks and only then the necessary tracking functions are inserted. This motivated performing instrumentation in phases, where subsequent phases can use information gathered in the previous phases. The ability to use static analysis and phased instrumentation has led to a dramatic reduction of the number of inserted instructions by more than 85% in our experiments [6]. This reduction can make the instrumented code more efficient when executed or easier for subsequent analysis or verification.
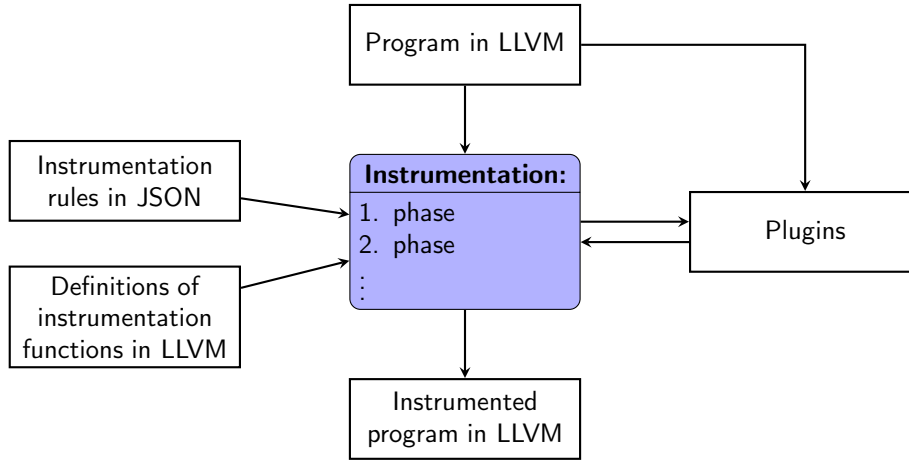
Figure 1: The scheme of configurable instrumentation.

The basic schema of `sbt-instrumentation` tool is depicted in Figure 1. Besides the LLVM bitcode to be instrumented, the tool needs to be supplied with two files created by a user: a file with definitions of so called *instrumentation functions* whose calls will be inserted into the code, and a JSON file with *instrumentation rules* that define how the LLVM bitcode should be instrumented with calls of instrumentation functions. In practice, the fact that our tool can insert just calls to instrumentation functions is not a restriction as these functions can contain arbitrary code.

The instrumentation proceeds in one or more phases, each phase defined by a set of instrumentation rules. In each phase, the tool goes through all instructions of the given LLVM bitcode and it looks for instructions matching any instrumentation rule of the current phase. If a match is found, conditions of the instrumentation rule are evaluated. This is where plugins can be queried. If conditions are satisfied, the rule is applied, i.e., a new code is inserted according to the rule and some information can be gathered for a later use. The result of the instrumentation process is again an LLVM bitcode.

2   Instrumentation Rules

Each instrumentation phase is defined by a set of instrumentation rules. An instrumentation rule consists of two parts saying *when* it should be applied and *what* its effect is. The first part of an instrumentation rule is specified by

- functions in which the rule is applied (typically `main` or all functions),

- a sequence of instructions that should be matched, and

- conditions under which the rule is applied.

The second part describes

- the instrumentation function call that should be inserted,

- where it should be inserted (before or after the matched sequence), and

- information-gathering effects of the rule, namely setting flags and remembering values in an auxiliary list.

Moreover, there are two other kinds of rules, namely rules for instrumentation of global variable declarations and rules instrumenting entry and exit points of functions.

In the rest of this section, we focus on conditions. The other parts of instrumentation rules are described in the following section.

2.1   Conditions

An instrumentation rule may contain conditions. To apply a rule, all conditions of the rule have to be satisfied.

Each condition consists of a *query* and a list of *expected results*. There are two kinds of queries, namely *flag queries* and *plugin queries*.

Flags are typically used to pass some information between instrumentation phases. They are set during application of instruction rules. Subsequently, a rule can be conditioned by a certain value of a specified flag. A flag query is simply the name of the flag. The query is satisfied if the current value of the flag is in the list of expected results. For example, memory safety instrumentation sets the flag `mallocPresent` to *"true"* when instrumenting a call to *malloc* in the first phase of the instrumentation. In the second phase, it inserts a check for memory leaks at the end of *main* only if the flag is set to *"true"*, i.e. if there is a call to *malloc* somewhere in the program.

Plugin queries are intended as an interface to static analyses. In general, a plugin query is formed by a keyword and some parameters. These parameters are values taken from the program being instrumented (typically the whole instruction matched by the rule or just an operand of this instruction). A query is satisfied if some plugin supporting this kind of query returns a string from the list of expected results. For example, a plugin query `canOverflow(op)`, where `op` is an instruction for an arithmetic operation with its operands, corresponds to the question "Can the result of `op` overflow?". Such a query can be answered by a plugin performing a range analysis. If the analysis says that the operation cannot overflow, the plugin answers *"false"* and it answers *"maybe"* otherwise. Besides the query `canOverflow(op)`, the range analysis plugin can also evaluate the query `canBeZero(value)` corresponding to the question whether `value` can be zero. Note that `value` is a term of the LLVM terminology and it can refer, for example, to a variable. We also support some queries evaluated by a pointer analysis including the following ones:

`isNull(pointer)` Does the given `pointer` point to NULL?

`isRemembered(value)` Was the given value or a pointer to the value remembered in the auxiliary list by some previously applied instrumentation rule? This information needs to be evaluated by a pointer analysis plugin because we ask not only about the value itself, but also about pointers to the value. We use this query in the configuration for checking memory safety to determine whether a check of a specific pointer was inserted in an earlier phase of the instrumentation.

`isValidPointer(addr, len)` Is a dereference of `len` bytes starting from the address `addr` a valid operation?

Plugins supporting these queries usually answer *"true"*, *"false"*, or *"maybe"*. If a plugin does not implement some query, it is not asked to evaluate it. The set of queries is not fixed and each

```
{
  "analyses": ["libRangeAnalysis.so"],
  "phases": [
      {
        "instructionsRules": [
            {
              "in": "*",
              "findInstructions": [
                  {
                    "returnValue": "*",
                    "instruction": "sdiv",
                    "operands": ["*","<t1>"]
                  }
              ],
              "conditions": [
                  {
                    "query": ["canBeZero","<t1>"],
                    "expectedResults": ["true"]
                  }
              ],
              "newInstruction": {
                    "instruction": "call",
                    "operands": ["<t1>","checkDivisionByZero"]
              },
              "where": "before"
            }
        ]
      }
  ]
}
```

Figure 2: Example of a JSON configuration file for inserting division-by-zero checks.

plugin can implement its own queries (without the need of changing the sbt-instrumentation source code).

3   Configuration

In this section, we describe the basic structure of a JSON file defining instrumentation rules. We illustrate the structure with the configuration for inserting division-by-zero checks presented in Figure 2.

A JSON file with instrumentation rules may contain the following fields, where only the field phases is mandatory.

analyses	List of paths to plugins.
flags	List of flags that can be set during instrumentation.
phases	List of instrumentation phases. Each phase contains a list of `instructionsRules` and/or `globalVariablesRules`. The phases are processed in the order given by their position in the list.

Instrumentation rules may use *configuration variables*. Syntactically, these variables are enclosed by the < and > characters. Configuration variables serve to store some parts of matched instructions (e.g., their operands) for use in conditions, inserted function calls, and to be stored in the auxiliary list.

3.1 `instructionsRules`

Each element of `instructionsRules` is a JSON object described by several fields:

in Name of a function, in which this rule should be considered for application. For example, `"in":  main` means that this rule will be applied only in the *main* function. It can also be set to `"*"` meaning that it should be taken into consideration in all functions (see Figure 2).

findInstructions Sequence of instructions we are searching for. For each instruction in the sequence, we need to fill in a field `instruction` that specifies the name of the instruction to be matched. The field `returnValue` allows to remember the return value of the instruction in a given configuration variable. It can be omitted or set to `"*"` if the return value is not needed. Finally, the field `operands` enables either to match the operands or to remember them in configuration variables. We can also optionally fill in the field `getTypeSize`. This field can be used only with `load`, `store` or `alloca` instructions and it stores the size of the type of the value that is being loaded, stored, or allocated to the given configuration variable.

An example of a `findInstructions` field can be found in Figure 2. We are searching for a sequence of instructions of length one, namely for a `sdiv` instruction. In this case, we do not care about the return value, so the `returnValue` field is set to `"*"`. The second operand of `sdiv` will be stored in the configuration variable `<t1>` (the first operand will be skipped as it is set to `"*"`).

conditions An optional field specifying a list of conditions that have to be satisfied in order to apply the rule (see Section 2.1). A condition consists of the fields `query` and `expectedResults`. The `query` is a list where the first element is the name of a query and other elements are parameters passed to the query. The `expectedResults` is a list of expected results of the query.

In Figure 2, the `conditions` list contains one condition that has to be satisfied for the rule to be applied. It is a plugin query that will be satisfied if the plugin loaded from the shared object `libRangeAnalysis.so` will answer *"true"* to the query `canBeZero` with the value of `<t1>` passed as a parameter. The following code shows a condition with a flag query that is satisfied if the flag `mallocPresent` is set to *"true"*.

```
"conditions": [
    {
        "query":["mallocPresent"],
        "expectedResult":["true"]
    }
]
```

newInstruction A new instruction that is to be inserted. It contains two mandatory fields: `instruction` that specifies the name of the new instruction (for now, only `call` instructions are supported), and `operands` of the instruction. The last operand is the name of the function that should be called. For example, the instrumentation rule described in Figure 2 will insert a call to the function `checkDivisionByZero` with the argument stored in `<t1>`.

where The location of insertion. It can be *"before"* or *"after"* the found sequence of instructions. Alternatively, it can have the value *"entry"* or *"return"* saying that the code should be inserted at the entry point or before every `return` instruction of the current function, respectively. If the value is *"entry"* or *"return"*, the rule does not need to contain the `findInstructions` field (as it is ignored).

setFlags This optional field describes the list of pairs `[flag, string]` that sets each `flag` to the corresponding `string` if the rule was applied. The following code sets the flag `loadFlag` to *"true"* and the flag `testFlag` to *"false"*.

```
"setFlags":[["loadFlag","true"],["testFlag","false"]]
```

remember An optional field specifying the name of a configuration variable whose value will be stored in the auxiliary list if the rule is applied.

3.2   `globalVariablesRules`

`globalVariablesRules` can be used to instrument declarations of global variables. The original motivation was in memory safety checking, where we needed to insert functions tracking memory blocks corresponding to global variables. Each element of `globalVariablesRules` is a JSON object with the following fields:

findGlobals Contains a mandatory field `globalVariable` that stores the address of the global variable to the given configuration variable, and an optional field `getTypeSize` that gets the size of the type of the global variable. In the following example, the field `findGlobals` stores the address of a global variable in `<t1>` and stores the size of its type in `<t2>`.

```
"findGlobals": {
    "globalVariable": "<t1>",
    "getTypeSize": "<t2>"
}
```

conditions The same as in `instructionsRule`.

newInstruction The same as in `instructionsRule`.

in The name of the function, at the beginning of which the new instruction will be inserted. This time, the field cannot be set to `"*"`.

4   Implementation and Usage

The `sbt-instrumentation` tool can be run by executing the compiled binary `sbt-instr`:

<div align="center"><code>./sbt-instr <config> <input> <definitions> <output></code></div>

where the arguments have the following meaning:

- `<config>` is a file with instrumentation rules in JSON as described in Section 3,

- `<input>` is the LLVM bitcode to be instrumented,

- `<definitions>` is the LLVM bitcode with definitions of instrumentation functions, and

- `<output>` is the name of the output file.

As we primarily use the instrumentation for programs written in the C programming language, there is a script `sbt-instr-c` that works directly with C programs:

<div align="center"><code>./sbt-instr-c <config> <input> [-output=<output>]</code></div>

The arguments have the same meaning as for `sbt-instr`, but now the `<input>` is a program in LLVM or C. The script compiles the program automatically to LLVM if needed and the final instrumented LLVM bitcode is stored to `<output>` file (if given) or into `out.bc`. Note that there is no argument for definitions of instrumentation functions. The path to the file with definitions is specified by a new top-level field `file` in the JSON configuration. The definitions of instrumentation functions can be given in LLVM or in C (in which case they are automatically compiled into LLVM).

The `sbt-instrumentation` tool is distributed as a part of the toolbox called Symbiotic [7] and comes with two predefined configurations used for program analysis, namely a configuration for checking memory safety [6] briefly described in Section 1 and a configuration for checking signed integer overflows. The latter configuration inserts a check before every binary operation over signed integers that may potentially overflow based on the results of a range analysis.

`sbt-instrumentation` tool together with the predefined configurations for checking memory safety and integer overflows can be found at

<div align="center">https://github.com/staticafi/sbt-instrumentation</div>

under Apache 2.0 license. The tool is implemented in the C++ programming language. It uses an open-source parser for the JSON format JsonCpp[1]. The `sbt-instrumentation` repository also contains the script `sbt-instr-c` written in Python.

To compile and run the tool, it is necessary to have CMake (minimal version 2.8.8) and the LLVM (minimal version 3.9.1) together with Clang installed. To use the predefined configurations for checking memory safety or integer overflows, it is also needed to install the `dg` library[2] for pointer analysis or the `ra` library[3] for range analysis, respectively.

---

[1] https://github.com/open-source-parsers/jsoncpp
[2] https://github.com/mchalupa/dg
[3] https://github.com/xvitovs1/ra

5   Related Work

Even though most of the tools for a program analysis implement their own instrumentation aimed directly at their purposes, there are a few papers about general-purpose instrumentation. However, none of these tools supports instrumentation of LLVM bitcode and they do not provide the full functionality of our tool, in particular, the interface to static analyses.

ATOM [8] is a tool for a binary instrumentation that allows a user to describe how the code will be instrumented and to define the code that will be inserted. The user has to implement their own instrumentation routine with the help of ATOM's C-language user interface. This routine defines the way of iterating over instructions and when to insert new code. FIT [9] is a binary instrumentation tool that is backwards compatible with ATOM, but it supports more different architectures.

Geimer et al. [10] propose an approach for general entry/exit instrumentation and explore the basic constructs that are needed for a configurable instrumentation routine. They evaluate the proposed approach by implementing a prototype based on the instrumentation tool from the TAU performance-analysis framework [11] which currently works with C, C++, and Fortran programs.

Stephenson et al. [12] designed a flexible instrumentation framework called SASSI for NVIDIA GPUs. The motivation was to allow injecting user-defined code into GPU programs instead of choosing from predefined profiling functions. The user therefore can specify *what* will be inserted and *where* it will be inserted.

6   Conclusion

We have introduced `sbt-instrumentation`, a tool for general-purpose configurable instrumentation of LLVM bitcode. The instrumentation process can be divided into multiple phases. Every phase can gather information about what has been already instrumented and this information can be used by subsequent phases. Furthermore, the tool provides an extensible API for static analyses in the form of plugins. Injecting new code can be then conditioned by the results of external static analyses.

We described the workflow of the tool and the format for configuration of the instrumentation process. `sbt-instrumentation` was integrated into the Symbiotic toolbox where it plays a crucial role for memory safety checking.

In the future, we plan to develop more configurations and try to plug in new static analyses. We also want to add support for replacing parts of the matched code.

References

[1] Konstantin Serebryany, Derek Bruening, Alexander Potapenko, and Dmitry Vyukov. AddressSanitizer: A Fast Address Sanity Checker. In *Proceedings of the 2012 USENIX Conference on Annual Technical Conference*, USENIX ATC'12, pages 28–28. USENIX Association, 2012.

[2] Java Profiler - JProfiler. https://www.ej-technologies.com/products/jprofiler/overview.html. Accessed: 2018-07-24.

[3] Daniel Kroening and Michael Tautschnig. CBMC - C bounded model checker - (competition contribution). In *Tools and Algorithms for the Construction and Analysis of Systems - 20th International Conference, TACAS 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014. Proceedings*, volume 8413 of *Lecture Notes in Computer Science*, pages 389–391. Springer, 2014.

[4] Martina Vitovská. Instrumentation of LLVM IR. Master's thesis, Masaryk University, Faculty of Informatics, Brno, 2018.

[5] The LLVM compiler infrastructure. http://llvm.org, 2017.

[6] Marek Chalupa, Jan Strejček, and Martina Vitovská. Joint forces for memory safety checking. In *Model Checking Software - 25th International Symposium, SPIN 2018, Malaga, Spain, June 20-22, 2018, Proceedings*, volume 10869 of *Lecture Notes in Computer Science*, pages 115–132. Springer, 2018.

[7] Marek Chalupa, Martina Vitovská, and Jan Strejček. Symbiotic 5: Boosted instrumentation (competition contribution). In *Tools and Algorithms for the Construction and Analysis of Systems - 24th International Conference, TACAS 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018, Proceedings, Part II*, volume 10806 of *Lecture Notes in Computer Science*, pages 442–446. Springer, 2018.

[8] Amitabh Srivastava and Alan Eustace. ATOM - A system for building customized program analysis tools. In *Proceedings of the ACM SIGPLAN'94 Conference on Programming Language Design and Implementation (PLDI), Orlando, Florida, USA, June 20-24, 1994*, pages 196–205. ACM, 1994.

[9] Bruno De Bus, Dominique Chanet, Bjorn De Sutter, Ludo Van Put, and Koen De Bosschere. The design and implementation of FIT: a flexible instrumentation toolkit. In *Proceedings of the 2004 ACM SIGPLAN-SIGSOFT Workshop on Program Analysis For Software Tools and Engineering, PASTE'04, Washington, DC, USA, June 7-8, 2004*, pages 29–34. ACM, 2004.

[10] Markus Geimer, Sameer Shende, Allen D. Malony, and Felix Wolf. A generic and configurable source-code instrumentation component. In *Computational Science - ICCS 2009, 9th International Conference, Baton Rouge, LA, USA, May 25-27, 2009, Proceedings, Part II*, volume 5545 of *Lecture Notes in Computer Science*, pages 696–705. Springer, 2009.

[11] Sameer Shende and Allen D. Malony. The Tau Parallel Performance System. *IJHPCA*, 20(2):287–311, 2006.

[12] Mark Stephenson, Siva Kumar Sastry Hari, Yunsup Lee, Eiman Ebrahimi, Daniel R. Johnson, David W. Nellans, Mike O'Connor, and Stephen W. Keckler. Flexible software profiling of GPU architectures. In *Proceedings of the 42nd Annual International Symposium on Computer Architecture, Portland, OR, USA, June 13-17, 2015*, pages 185–197. ACM, 2015.