



University of Innsbruck

Institut of Computer Science

Distributed and Parallel Systems Group

## Insieme's Haskell-based Analysis Toolkit

Putting a HAT on Insieme

**Alexander Hirsch**



**Master Thesis**

**Supervisor:** Herbert Jordan PhD

30 April 2017



# Declaration

By my own signature I declare that I produced this work as the sole author, working independently, and that I did not use any sources and aids other than those referenced in the text. All passages borrowed from external sources, verbatim or by content, are explicitly identified as such.

Signed:.....Date: .....



# Abstract

This document provides an insight into the [Haskell-based Analysis Toolkit \(HAT\)](#), allowing developers and researchers to rapidly prototype and develop static program analyses. It is therefore composed of an introduction into the topic of static program analysis, an architectural overview of the framework, the specification of a variety of essential analyses and components, a tutorial for designing new analyses showcasing the framework, followed by an evaluation of the frameworks capabilities. The conclusion summarises the contributions of this thesis and provides an outlook on future work.

For the tutorial, and to showcase the framework, an array out-of-bounds analysis is constructed. This analysis is also utilised for evaluating the framework, by investigating its ability of implicitly expanding a program analysis language feature support. Furthermore we evaluate the frameworks performance in terms of execution time and memory requirements by analysing 63 622 different properties within a total of 320 example codes.

The complete code of this project is available on [GitHub](#)<sup>1</sup> embedded into the high-level analysis module of the Insieme compiler.

---

<sup>1</sup><https://github.com/insieme/insieme>



# Contents

<b>1. Introduction</b>	<b>1</b>
1.1. Attribution . . . . .	1
1.2. Motivation . . . . .	1
1.3. Objectives . . . . .	2
1.4. Picking Haskell . . . . .	2
1.5. Overview . . . . .	3
<b>2. Static Code Analysis</b>	<b>5</b>
2.1. Purpose . . . . .	5
2.2. Control-Flow Graph . . . . .	5
2.3. Data-Flow Analysis . . . . .	6
2.3.1. Lattice . . . . .	8
2.3.2. Reaching Definitions Example . . . . .	8
2.4. Constraint-Based Analysis . . . . .	10
2.4.1. Reaching Definitions Example (cont) . . . . .	11
2.5. Insieme CBA framework . . . . .	12
2.5.1. Property Space . . . . .	13
2.5.2. Analysis Variables . . . . .	13
2.5.3. Assignment . . . . .	14
2.5.4. Constraint Guards . . . . .	14
2.5.5. Constraint Generator . . . . .	14
2.5.6. Analysis . . . . .	14
2.5.7. Constraint Solver . . . . .	15
<b>3. Architecture</b>	<b>17</b>
3.1. Insieme . . . . .	17
3.2. INSPIRE . . . . .	18
3.2.1. Structure . . . . .	19
3.3. Extension Mechanism . . . . .	23
3.3.1. The Basic Language Extension . . . . .	24
3.3.2. Reference Extension . . . . .	26
3.3.3. Parallel Extension . . . . .	26
3.4. Semantics . . . . .	28
3.5. Haskell-based Analysis Toolkit . . . . .	33
<b>4. Framework Implementation</b>	<b>37</b>
4.1. Build Process . . . . .	37
4.1.1. Compiling HAT into a Library . . . . .	37
4.1.2. Linking the Shared Libraries . . . . .	38
4.2. Adapter . . . . .	39
4.2.1. Foreign Function Interface . . . . .	39
4.2.2. Haskell RunTime System . . . . .	40
4.2.3. Context . . . . .	41
4.3. IR Components . . . . .	41
4.3.1. Node Types . . . . .	41

## Contents

4.3.2.	Node Path . . . . .	42
4.3.3.	Binary Parser . . . . .	42
4.3.4.	Rose Tree – HAT’s INSPIRE Representation . . . . .	46
4.3.5.	Node Address . . . . .	46
4.3.6.	Visitors . . . . .	47
4.3.7.	Node References and Queries . . . . .	48
4.3.8.	Context . . . . .	48
4.4.	Utilities . . . . .	50
4.4.1.	Bound Sets . . . . .	50
4.4.2.	Arithmetic Formulas . . . . .	52
4.5.	Fixpoint Solver . . . . .	54
4.5.1.	Lattice . . . . .	54
4.5.2.	Variables . . . . .	55
4.5.3.	Assignments . . . . .	56
4.5.4.	Constraints . . . . .	56
4.5.5.	Solver . . . . .	57
4.6.	Modelling Arrays, Structs, and Unions . . . . .	58
4.7.	Generic Data-Flow Analysis . . . . .	60
4.7.1.	Operator Handler . . . . .	62
4.8.	Program Point . . . . .	62
4.8.1.	Program Point Analysis . . . . .	62
4.9.	Memory State . . . . .	64
4.9.1.	Reaching Definitions Analysis . . . . .	64
4.9.2.	Memory State Analysis . . . . .	65
4.10.	Specific Analyses . . . . .	66
4.10.1.	Identifier Analysis . . . . .	67
4.10.2.	Arithmetic Analysis . . . . .	68
4.10.3.	Boolean Analysis . . . . .	72
4.10.4.	Data Path Analysis . . . . .	77
4.10.5.	Reference Analysis . . . . .	77
4.10.6.	Callable Analysis . . . . .	79
4.11.	Helper Analyses . . . . .	81
4.11.1.	Reachability Analysis . . . . .	81
4.11.2.	Exit Point Analysis . . . . .	84
4.11.3.	Predecessor Analysis . . . . .	85
4.11.4.	Call Site Analysis . . . . .	86
<b>5.</b>	<b>Constructing a new Analysis</b> . . . . .	<b>87</b>
5.1.	Defining the Goals . . . . .	87
5.2.	Investigating INSPIRE . . . . .	89
5.3.	Designing the Analysis . . . . .	93
5.4.	Defining the Property Spaces . . . . .	94
5.5.	Defining the Interface . . . . .	94
5.6.	Setup Boilerplate . . . . .	95
5.7.	Write Test Cases . . . . .	96
5.8.	Implementation . . . . .	97
5.8.1.	Element Count Analysis . . . . .	98
5.8.2.	Getting the Array Index . . . . .	99
5.8.3.	Putting it Together . . . . .	100
5.9.	Summary . . . . .	100
<b>6.</b>	<b>Evaluation</b> . . . . .	<b>103</b>
6.1.	Qualitative Evaluation . . . . .	103
6.2.	Quantitative Evaluation . . . . .	105



<b>7. Conclusion</b>	<b>109</b>
7.1. Future work . . . . .	109
<b>A. INSPYER</b>	<b>111</b>
A.1. JSON Export . . . . .	112
A.2. User Interface . . . . .	112
A.3. Meta Data . . . . .	114
<b>B. Integration Test List</b>	<b>117</b>



# 1. Introduction

This document describes a framework for easing the implementation of data-flow analyses within the Insieme project. Note that throughout this document the terms framework, toolkit, and [Haskell-based Analysis Toolkit \(HAT\)](#) are used interchangeably.

The core contribution of this project is the implementation of a scalable, easy-to-use static program analysis framework, which can be used for rapid development and prototyping of data-flow analyses.

As this project is a joint effort between Herbert Jordan, Thomas Prokosch, and myself, a section dedicated to the attribution is provided next. It is followed by a motivation section, the objectives, and a justification for picking Haskell as platform for the framework. The introduction is concluded with a structural overview of the remaining document.

## 1.1. Attribution

The theoretical groundwork, as well as a first prototype has been created by Herbert Jordan, who also plays a major role in the current development of the framework. He is tasked with implementing the core components of the framework and overlooking the development process.

Thomas Prokosch helped with the initial technology assessment and prototyping components in Haskell before they are integrated into the framework.

The main tasks of the author of this thesis are:

- integration of the C++ / Haskell interface
- framework interface design
- [Intermediate Representation \(IR\)](#) data structure
- documentation and tutorial
- implementing utilities
- example analyses
- debugging tools

## 1.2. Motivation

Program analysis, in general, is a big topic, not only in compiler theory. Extracting certain properties automatically from an input program can be viable for multiple different use-cases (eg auto-tuning or code validation). When talking about compilers, many of their optimisations rely on some kind of [Data-Flow Analysis \(DFA\)](#) to identify potential optimisation options.

## 1. Introduction

The biggest nugget of motivation can be mined from the idea of having a generic **Constraint-Based Analysis (CBA)** framework at hand, which already provides a basic set of common analyses and can be extended rapidly as needed. The generic **DFA** analysis provided by the framework presented by this thesis handles (among others) declarations, parameter and return value passing, closures, support for composed data types, calls to unknown external functions, inter-procedural analyses, and memory locations.

By integrating it into a source-to-source compiler, and operating on a high-level intermediate representation, little boilerplate code is required to construct new analyses for various tasks. Prototypes for research as well as for development purposes can be efficiently implemented.

### 1.3. Objectives

The ultimate goal of this project is the realisation of a scalable implementation of the **CBA** framework developed by Jordan [10] exhibiting superior usability. The main objective of this thesis is to help lay the foundation for this framework and its integration with Insieme. To achieve the ultimate goal, further development will build upon this work in the future.

The main objective therefore consists of developing the core components of the framework. These are composed of a *fixpoint solver*, a smaller set of generic analyses, and a bigger set of basic (more specific) analyses. The intention behind each of these components and their relationship with each other is explained in Chapter 4.

The framework is a part of the Insieme project and therefore operators on the **Intermediate Representation** of the Insieme compiler – namely the **INSieme Parallel Intermediate REpresentation (INSPIRE)**. This **IR** is implemented as C++ data structure in the core of the compiler. However, as the framework is not realised using C++ (see below), a suitable representation of **INSPIRE** needs to be established. This also includes a way to transfer an **INSPIRE** program from the Insieme compiler to this representation. Furthermore, the framework should provide a way of relaying back analysis results to C++.

Also, it should be possible to use the framework as standalone, without the need for the Insieme compiler.

Last but not least, creating documentation of the framework and a manual for further analysis development is needed. This objective is reflected by this document.

### 1.4. Picking Haskell

The first prototype of the **CBA** framework has been implemented using C++. During its development drawbacks of using C++ manifested in the form of long compile times yielding a very slow development cycle. Due to this reason, a different programming language (platform) has been chosen for the framework.

Due to the mathematical nature of program analysis, a functional programming language is preferred for the implementation. The mathematical constructs defining analyses can be realised much simpler via the use of **Algebraic Data Types (ADTs)** offered by functional languages, compared to structs, classes, and records of imperative and object-oriented languages.

Haskell [3], popular amongst functional programming languages, features expressive syntax, a strong, static type system, and a growing community maintaining thousands of packages and libraries. It was therefore ultimately selected for this task. The average reader of this document may not be familiar with Haskell – or functional programming languages in general – and is therefore invited to pick it up and start their functional programming adventure. Lipovaca [14] composed an excellent book going by the name of *Learn You a Haskell for Great Good!: A Beginner's Guide*, endorsed by the Haskell community, and available online. Core concepts are explained with additional background information and useful analogies. The playful art style and pop culture references positively influence the learning experience. We also recommend the books *Real World Haskell* [17], *Haskell Design Patterns* [13], and *Haskell High Performance Programming* [19].

Other functional programming languages like OCaml [4], F# [2], Scala [5], and Erlang [1] were considered too, yet have been ruled out during an initial technology assessment phase. Since F# builds on the .NET framework it was quickly removed from the list of considerations. The overall footprint seems too large and no benefits of having the .NET framework on-board materialised. Both Scala and Erlang seem interesting as well as promising, although on further investigation Haskell promises a more advanced type system and people have reported that its easier applicable and better designed compared to Scala [16]. In contrast to Erlang, Haskell is statically typed – something we very much prefer as it makes it easier to catch bugs early in the development process. OCaml, in fact quite similar to Haskell, has been reconsidered multiple times during the prototyping phase. In the end, personal preference in the syntax and design of the standard library resulted in the selection of Haskell instead of OCaml.

For the remainder of this thesis, it is assumed that the reader has a basic understanding of the Haskell programming language. Nevertheless, details and more complex concepts are introduced when needed, along with additional references.

## 1.5. Overview

Chapter 2 provides a brief introduction to static code analysis, focusing on *flow-sensitive* analysis. The concept is communicated via an example utilising a more conventional DFA framework, followed by the constraint-based approach. This leads to a generic CBA framework, which is originally based on the one presented by Nielson et al. [15]. Additionally, customisations and adaptations towards integrating CBA capabilities into Insieme have been made by Jordan [10] to the framework, which not only seem promising in theory but also bear in mind real-world use-cases and practical applications – something not encountered that often in the field of program analysis.

Next, covered by Chapter 3, the overall architecture of Insieme's Haskell-based Analysis Toolkit (HAT) is depicted, delivering the big picture of the task at hand. This includes presenting the Insieme compiler followed by a detailed investigation of INSPIRE [11].

The implementation part, described in Chapter 4, reflects the documentation of the current state of the framework. The source code can be found on GitHub<sup>1</sup> embedded into the analysis module of the Insieme compiler. Presented code snippets have been simplified for clarity in some cases.

In Chapter 5 a new analysis is constructed showcasing the new framework. Furthermore, the newly implemented analysis is used to evaluate the framework in Chapter 6. Finally, Chapter 7 summarises the contributions of this work and provides an outlook on future development.

---

<sup>1</sup><https://github.com/insieme/insieme>



## 2. Static Code Analysis

This chapter provides a concise introduction to the topic of static program analysis with a focus on [Data-Flow Analysis \(DFA\)](#) and its extension, the *constraint-based approach*. It serves as a stepping stone to communicate the theoretical background for the following chapters.

It starts off with a few words about the general purpose of static code analysis and continues with the introduction of the [Control-Flow Graph \(CFG\)](#). After that the general concept of [DFA](#) is communicated together with a *reaching definitions* example. Finally, the constraint-based approach is tackled and applied to the same example for comparison. The chapter concludes with an overview of customisations made by Jordan [10] to the [Constraint-Based Analysis \(CBA\)](#) framework.

### 2.1. Purpose

Static code analysis is about deducing *properties* of a program without actually running the program. This happens typically in a compiler during the optimisation stage. Yet there are also other tools (referred to as static code analysers) which can run various checks on the code and extract properties from it. Checking for nullpointer dereferences, dead code, and estimating runtime complexity are just a few examples for checks, that can be conducted by an analysis. In case of a compiler, an optimisation commonly consists of an analysis identifying structures in the program suitable for optimisation, followed by a transformation of that structure into one that yields better performance<sup>1</sup>. Static code analysers are often used to identify bugs (eg race conditions, iterator invalidation, etc) not caught by the compiler. But they can also be useful to extract program properties for later use (eg auto-tuning).

As already mentioned, this introduction section focuses on [DFA](#), which is a *flow-sensitive* technique. Meaning, we take the control-flow of the input program into account at the time the analysis is performed. This leads to a more accurate result, but also requires a greater computational effort. *Flow-insensitive* analyses, on the other hand, are much faster. However, they yield less accurate results in the form of over-approximations. Since knowledge about flow-insensitive analysis is not required for this task, these kinds of techniques are not elaborated further.

There are, of course, different techniques how one could construct a [DFA](#). And while some of them provide clear benefits over others, they sometimes also suffer drawbacks. But what (nearly) all of them have in common is the utilisation of [CFGs](#). The following section introduces those.

### 2.2. Control-Flow Graph

The control flow of a program can be abstracted as a graph, in which each node represents a *basic block*. A basic block is a sequence of instructions without any jumps or jump targets (ie labels) in-

---

<sup>1</sup>Depending on the current objective, performance can refer to runtime, memory usage, power consumption, etc.

## 2. Static Code Analysis

between. It commonly begins with a jump target and ends with a jump instruction. Basic blocks are connected via directed edges, which represent potential execution paths. An edge that points to an already visited block during depth-first traversal is referred to as a *backward edge* and frequently introduced by a loop. Adding two additional empty blocks, the *entry block* and *exit block*, can help keeping algorithms operating on the graph simple. All control-flows start at the entry block and end at the exit block. [22]

What exactly constitutes a node of the CFG is not set in stone and varies depending on the application. While the textbook definition concerns itself with basic blocks, one could easily abstract away whole functions into a graph's nodes. Based on the application, attaching additional information to the edges can be helpful too.

Listing 1 displays a short C program with three functions. Figure 1 illustrates the corresponding CFG. Each function features a dedicated entry node annotated with the function's name and an exit (return) block. The example contains branching as well as loops.

```
1 | #include <stdio.h>
2 | #include <stdlib.h>
3 |
4 | void print_something(int times) {
5 |     for (int i = 0; i < times; ++i) {
6 |         puts("something");
7 |     }
8 | }
9 |
10 | void print_usage() {
11 |     puts("some usage");
12 | }
13 |
14 | int main(int argc, char* argv[]) {
15 |     if (argc < 2) {
16 |         print_usage();
17 |         return EXIT_FAILURE;
18 |     }
19 |     print_something(atoi(argv[1]));
20 |     return EXIT_SUCCESS;
21 | }
```

Listing 1: cfg.c

During a program analysis run certain properties are tracked across the nodes. These properties (eg reaching definitions) are recorded when a node is entered and left. These records are referenced by the labels of the related graph node together with either in or out.

For instance,  $REACH_{in}[S]$  and  $REACH_{out}[S]$  are the reaching definitions recorded at the entry point and exit point respectively of the graph node  $S$ . This notation is often shortened to  $in(S) / out(S)$  if the specific analysis can be derived from the current context [21].

## 2.3. Data-Flow Analysis

In this section the basis of [Data-Flow Analysis](#) is described. The description itself is keep very informal and is presented through an example.



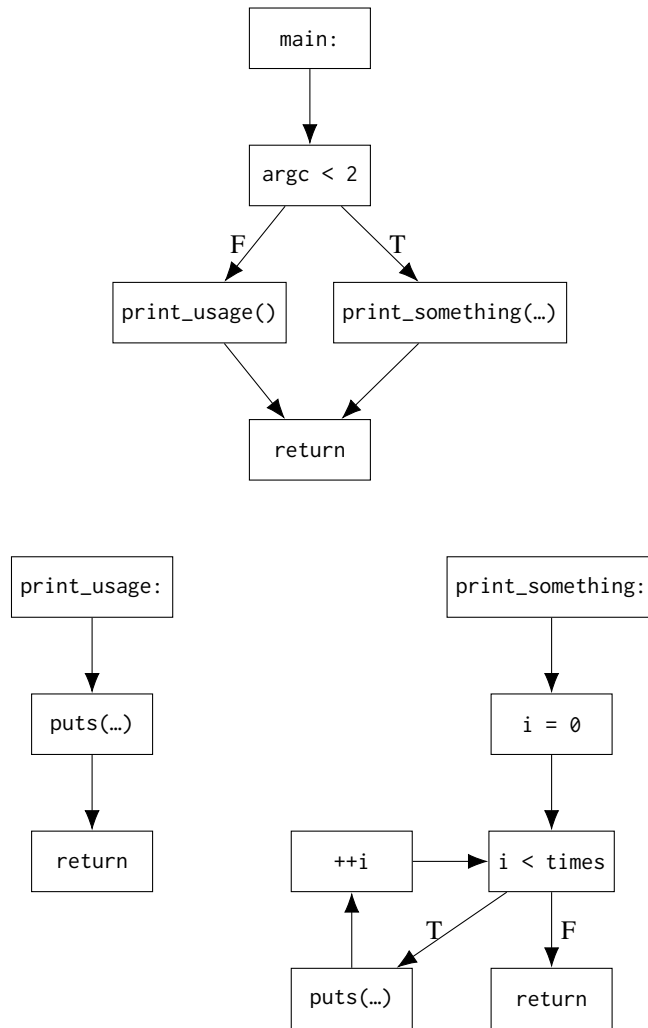


Figure 1.: A CFG corresponding to `cfg.c`.

## 2. Static Code Analysis

A **Data-Flow Analysis** is a static program analysis used to extract certain properties (eg *reaching definitions* or *live variables*) from a program. Traditionally, DFA assumes a given control-flow of the program – and while one could work on the **Abstract Syntax Tree (AST)** – using a CFG is often much simpler and yields similar outcome (with respect to accuracy regarding DFA) [7]. This technique can be cast into a framework as presented in the remainder of this document.

Furthermore, there are two ways an analysis can operate. It can either be a *must analysis* or a *may analysis*. A must analysis requires that a property must hold in all branches to be still considered after a meeting point (logical and). Alternatively, in a *may analysis* a property is considered after a meeting point if it holds in at least one branch (logical or). An example for a *must analysis* would be *available expressions* [7], while *reaching definitions* (see Section 2.3.2) is a *may analysis*.

Most DFA frameworks only yield sound solutions if the domain of the property space forms a *lattice*. Because of this, we now take a short detour and glance at the mathematical concept of a lattice.

### 2.3.1. Lattice

A complete lattice is an algebraic structure which consists of a partially ordered set in which every two elements share a unique upper bound (*join*) and a unique lower bound (*meet*). Furthermore there exists one element  $\perp$  that is smaller, and one element  $\top$  that is larger than every other element in the set. The purpose of the lattice, with respect to program analysis, is to model a domain of program properties. The *height* of a lattice is given by the longest path through partial order from greatest to least. Elements of the lattice represent flow values (in and out sets) while  $\perp$  and  $\top$  represent the *best-case* and *worst-case* information respectively. The *meet* operator  $\sqcap$  describes how two flow values are merged. [20, 6, 18]




Figure 2 shows an example application for a lattice where three variables  $i$ ,  $j$  and  $k$  are given. A flow value can exhibit any combination of them. Hence the  $\top$  and  $\perp$  value can be mapped to the set of all variables and  $\emptyset$  respectively. The height of the lattice, 4, which can be immediately deduced from the illustration and the ordering of the lattice ( $\sqsubseteq$ ) is mapped to  $\subseteq$ . This representation is known as a *Hasse Diagram* [23].

Figure 2.: A Hasse diagram of an example lattice.

There exists also the concept of a *semilattice*, which resembles a normal lattice, but only features *either* a least upper bound *or* greatest lower bound, known as *join-semilattice* or *meet-semilattice*, respectively.

The interested reader is encouraged to take a look at the more comprehensive introduction to lattices provided by Jovanovic [12].

### 2.3.2. Reaching Definitions Example

A definition of a variable at program point  $p_1$  *reaches* program point  $p_2$  if there is a control-flow path from  $p_1$  to  $p_2$  along which the variable defined by  $p_1$  is not redefined. Positive and negative examples are provided. [21]

$$\left| \begin{array}{l} p1: x = 42 \\ p2: y = 12 \\ p3 \end{array} \right.$$

The definition of  $x$  from  $p_1$  reaches the program point  $p_3$ .

$$\left| \begin{array}{l} p4: z = 24 \\ p5: z = 16 \\ p6 \end{array} \right.$$

The definition of  $z$  from  $p_4$  does not reach  $p_6$ . It is *killed* by the definition at  $p_5$ .

### Lattice

Before concerning ourselves with the DFA itself, we have to define the lattice of the property we would like to obtain. Let  $V$  be the set of all variables,  $S = 0, 1, \dots$  the set of all statements in the CFG and  $L = S \cup \{?\}$ . The lattice is then given by

$$(2^{V \times L}, \subseteq)$$

“?” is used to represent an input to the program. From this,  $\perp = \emptyset$  and  $\top = V \times L$  can be derived. Since the elements of  $V$  and  $L$  are ordered, we can also derive an order for  $V \times L$ . The join and meet operators can be mapped to  $\cup$  and  $\cap$  respectively.

### Data-Flow Analysis

The data-flow equations for calculating reaching definitions are: for all basic blocks  $b$  of a given CFG,

$$\begin{aligned} \text{in}(b) &= \bigcup_{p \in \text{pred}(b)} \text{out}(p) \\ \text{out}(b) &= \text{gen}(b) \cup (\text{in}(b) \setminus \text{kill}(b)) \end{aligned}$$

where  $\text{gen}(b)$  is the set of definitions contained in block  $b$ ,  $\text{kill}(b)$  the set of definitions *killed* in block  $b$  and  $\text{pred}(b)$  is the set of predecessors of  $b$ . *Killing* a definition means that it is no longer available in subsequent blocks, as illustrated in the initial example of this subsection.

Let us take the following program computing  $x!$  together with the corresponding CFG in Figure 3. Note that we only focus on the calculating part here and assign each statement its own block for simplicity. This example is based on the introduction provided by Nielson et al. [15, pp. 5–10].

Running a reaching definition analysis on this code results in the following twelve equations, two for each block.

$$\begin{array}{ll} \text{in}(1) = \{(x, ?)\} & \text{out}(1) = \{(y, 1)\} \cup (\text{in}(1) \setminus \{(y, l) \mid l \in L\}) \\ \text{in}(2) = \text{out}(1) & \text{out}(2) = \{(z, 2)\} \cup (\text{in}(2) \setminus \{(z, l) \mid l \in L\}) \\ \text{in}(3) = \text{out}(2) \cup \text{out}(5) & \text{out}(3) = \text{in}(3) \\ \text{in}(4) = \text{out}(3) & \text{out}(4) = \{(z, 4)\} \cup (\text{in}(4) \setminus \{(z, l) \mid l \in L\}) \\ \text{in}(5) = \text{out}(4) & \text{out}(5) = \{(y, 5)\} \cup (\text{in}(5) \setminus \{(y, l) \mid l \in L\}) \\ \text{in}(6) = \text{out}(3) & \text{out}(6) = \{(y, 6)\} \cup (\text{in}(6) \setminus \{(y, l) \mid l \in L\}) \end{array}$$

Each pair consists of a variable ( $x$ ,  $y$  or  $z$ ) and a label, where the label is either the number of the block or “?” if the variable is not initialised within this context. Uninitialised variables can be considered input to the program. When focusing on block 3 one can see that the information flowing into the

## 2. Static Code Analysis

```
1 | #include <stdio.h>
2 | #include <stdlib.h>
3 |
4 | int factorial(int x) {
5 |     int y = x;
6 |     int z = 1;
7 |     while (y > 1) {
8 |         z = z * y;
9 |         y = y - 1;
10 |    }
11 |    y = 0;
12 |    return z;
13 | }
14 |
15 | int main(int argc, char* argv[]) {
16 |     int x = atoi(argv[1]);
17 |     printf("%d! = %d\n", x, factorial(x));
18 |     return EXIT_SUCCESS;
19 | }
```

Listing 2: factorial.c, a simple factorial program.

block is the same as the information flowing out of the block since no variables are modified. At block 3 two different control-flows meet yielding a merge of facts.

One can observe a different result for the first equation when comparing it to the original source of this example. This change results from the way this example is presented. In our case the two variables  $y$  and  $z$  do not exist prior to executing the algorithm, hence the analysis does not need to consider them as input. [15, p. 7]

As stated by Nielson et al. [15, pp. 7–8], only a finite number of solutions exist for the derived system of equations. Furthermore each of these solutions is a *fixed point* and all of them are partially ordered. Most interestingly, the *least fixed point solution*, obtained this way, contains the fewest pairs of reaching definitions. Therefore this is the most accurate result, and the one we are interested in. It looks as follows:

$\text{in}(1) = \{(x, ?)\}$	$\text{out}(1) = \{(x, ?), (y, 1)\}$
$\text{in}(2) = \{(x, ?), (y, 1)\}$	$\text{out}(2) = \{(x, ?), (y, 1), (z, 2)\}$
$\text{in}(3) = \{(x, ?), (y, 1), (y, 5), (z, 2), (z, 4)\}$	$\text{out}(3) = \{(x, ?), (y, 1), (y, 5), (z, 2), (z, 4)\}$
$\text{in}(4) = \{(x, ?), (y, 1), (y, 5), (z, 2), (z, 4)\}$	$\text{out}(4) = \{(x, ?), (y, 1), (y, 5), (z, 4)\}$
$\text{in}(5) = \{(x, ?), (y, 1), (y, 5), (z, 4)\}$	$\text{out}(5) = \{(x, ?), (y, 5), (z, 4)\}$
$\text{in}(6) = \{(x, ?), (y, 1), (y, 5), (z, 2), (z, 4)\}$	$\text{out}(6) = \{(x, ?), (y, 6), (z, 2), (z, 4)\}$

## 2.4. Constraint-Based Analysis

Previously we dealt with the extraction of multiple equations from a CFG by traversing it, and then solving a system of equations to derive information (ie facts) about the program. The main idea behind the constraint-based approach is to directly extract these constraints and handle them explicitly, instead of implicitly handling equations given by the structure of the CFG. Due to the explicit handling of these constraints the two processes of generating and solving them can be separated, resulting in a *constraint generator* and *constraint solver*. [15, 10]




Figure 3.: A CFG corresponding to factorial.c.

We continue straight on with the example of the previous section to demonstrate the differences between the mentioned two approaches.

### 2.4.1. Reaching Definitions Example (cont)

Taking the constraint-based approach, the following list of constraints can be derived from the previous example. See Figure 3. Note that the format of constraint may be more complex (eg conditions), depending on the scenario. They may also be derived from other sources than the CFG (eg an AST).

$$\begin{array}{ll}
 \text{out}(1) \supseteq \text{in}(1) \setminus \{(y, l) \mid l \in L\} & \text{out}(1) \supseteq \{(y, 1)\} \\
 \text{out}(2) \supseteq \text{in}(2) \setminus \{(z, l) \mid l \in L\} & \text{out}(2) \supseteq \{(z, 2)\} \\
 \text{out}(3) \supseteq \text{in}(3) & \\
 \text{out}(4) \supseteq \text{in}(4) \setminus \{(z, l) \mid l \in L\} & \text{out}(4) \supseteq \{(z, 4)\} \\
 \text{out}(5) \supseteq \text{in}(5) \setminus \{(y, l) \mid l \in L\} & \text{out}(5) \supseteq \{(y, 5)\} \\
 \text{out}(6) \supseteq \text{in}(6) \setminus \{(y, l) \mid l \in L\} & \text{out}(6) \supseteq \{(y, 6)\}
 \end{array}$$

Upon close investigation of this list and the original program, a pattern emerges. Every time a variable is assigned, two constraints are generated. One for excluding (*killing*) all pairs generated by prior assignments of this variable, and one for adding (*generating*) a new pair corresponding to this assignment. Otherwise, if no variable is assigned, everything is simply passed through.

Note the similarities of this discovered pattern and the original dataflow equations:

$\text{out}(b) = \text{gen}(b) \cup (\text{in}(b) \setminus \text{kill}(b))$	Dataflow equation
$\text{out}(b) \supseteq \text{gen}(b)$	Constraints
$\text{out}(b) \supseteq \text{in}(b) \setminus \text{kill}(b)$	

Nevertheless, these constraints only tell us how information flows *through* each block, we are still

## 2. Static Code Analysis

missing how information flows *between* blocks. Therefore the following list is derived.

$$\begin{aligned} \text{in}(2) &\supseteq \text{out}(1) \\ \text{in}(3) &\supseteq \text{out}(2) & \text{in}(3) &\supseteq \text{out}(5) \\ \text{in}(5) &\supseteq \text{out}(4) \\ \text{in}(6) &\supseteq \text{out}(3) \end{aligned}$$

Again a pattern is immediately visible: each edge in the CFG results in a constraint connecting the exit set of a block with the entry set of its successor.

The last constraint we need relates to the input of our program.

$$\text{in}(1) \supseteq \{(x, ?)\}$$

Finally, the same solution obtained in the data-flow analysis approach is also valid for this system of constraints. In fact, it is also the least solution. [15, p. 10]

## 2.5. Insieme CBA framework

The constrained-based approach, as presented by Nielson et al. [15], resembles a generic framework suitable for various kinds of analysis. While we only focused on reaching definitions for now, other kinds (eg live variable analysis) can be realised in a similar fashion.

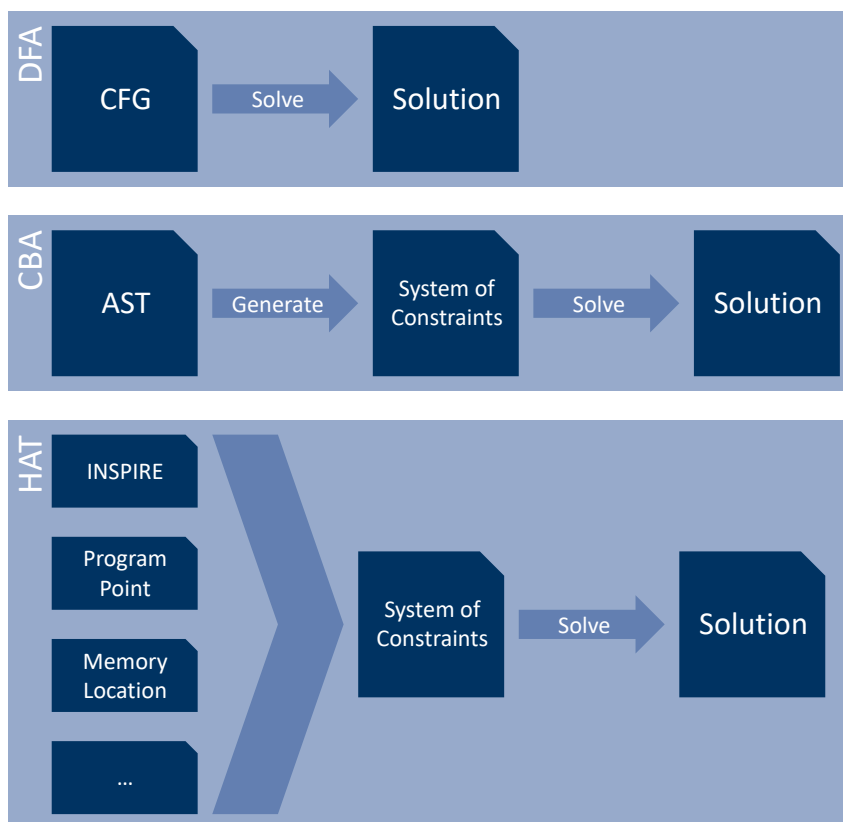


Figure 4.: Illustrating the difference between DFA, CBA and HAT.

In Figure 4 the core difference between DFA, CBA and our Haskell-based Analysis Toolkit (HAT) are illustrated. In DFA a system of equations is implicitly derived from the CFG and immediately solved to obtain a valid solution for an analysis. In CBA the generation and solving of constraints is replaced by two stages: first generating a system of constraints and solving the system afterwards. Both stages are clearly separated and the second is initiated *after* the first one is completed. HAT, on the other hand, relaxes the strict design of two separated stages. This is made possible by the introduction of a lazy constraint solver [10].

The lazy constraint solver allows the toolkit to generate constraints on the fly, as needed, and to inject them into the system of constraints during an analysis run. Furthermore, HAT is not limited to using an (annotated) CFG. As already hinted in the introduction to this chapter, not all analyses operate on a CFG. Instead, they could directly use an AST or their own custom data structure.

Using the INSieme Parallel Intermediate REpresentation (INSPIRE) together with additional sources of information, including but not limited to Program Points and Memory Locations, analyses can be more flexible and can output more accurate results than by just using a CFG. In general, the framework allows multiple different structures to be associated with *analysis variables*. Also, the extraction of a CFG from an AST can be skipped. In particular, this design enables the unification of control- and data-flow problems, facilitating the computation of the actual (inter-procedural) flow of control as part of the DFA.

The design of the modified CBA framework, including lots of details and examples, can be found in [10, pp. 198–289]. Certain parts have been cherry-picked and summarised in the remainder of this section. Section 3.2 is dedicated to INSPIRE.

### 2.5.1. Property Space

A property space consists of a domain, representing all possible values a property may exhibit and a *combination operator*. The operator is used to merge the results of different control-flow paths considered by the analysis. [10, p. 187]

Using the reaching definitions example, presented previously, the property space would be given by

$$(2^D, \bigcup_D)$$

where

$$\begin{aligned} D &= V \times L \\ \bigcup_D : 2^D &\rightarrow D \\ \{d_1, d_2, \dots, d_N\} &\mapsto \bigcup_{1 \leq i \leq N} d_i \end{aligned}$$

Each property space is inducing a lattice  $(D, \subseteq)$  by defining  $a \subseteq b \iff b = \bigcup\{a, b\}$ . Thus,  $\perp = \bigcup \emptyset$  and  $\top = \bigcup D$ .

### 2.5.2. Analysis Variables

Analysis variables are given by a pair connecting the type of analysis (eg arithmetic value) and some id associated with an element of an underlying structure. For example, given the labelled expression

## 2. Static Code Analysis

$[[7]^1 + [4]^2]^3$  where the labels are written in superscript, the variable  $A(3)$  refers to the *arithmetic value* of the expression labelled 3. This would lead to the following set of constraints.

$$\begin{aligned} \{7\} &\subseteq A(1) \\ \{4\} &\subseteq A(2) \\ \{x + y \mid x \in A(1), y \in A(2)\} &\subseteq A(3) \end{aligned}$$

### 2.5.3. Assignment

An assignment maps each analysis variable to a value of the corresponding property space. At the beginning of an analysis each analysis variable is initialised with the  $\perp$  value of the respective property space. The solver (see below) modifies this mapping until all constraints are satisfied. The final assignment corresponds to the (least) solution of the system of constraints.

### 2.5.4. Constraint Guards

A detail omitted in the previous example is the option of having *conditional constraints*. These are, in other words, *guarded* constraints which are only considered if their *guard* evaluates to true. Taking one of the previous constraints we encountered out of context and attaching a random guard, just for the sake of it, gives us the following expression

$$x = 4 \wedge y < 2 \implies \text{in}(4) \setminus \{(z, l) \mid l \in L\} \supseteq \text{out}(4)$$

, where ' $x = 4 \wedge y < 2$ ' is the guard and ' $\text{in}(4) \setminus \{(z, l) \mid l \in L\} \supseteq \text{out}(4)$ ' is the constraint. Also, the term on the right of the  $\supseteq$  operator must be a single variable – the constraint or target variable.

Constraints without an explicit guard can be viewed as one with a guard that is always true.

### 2.5.5. Constraint Generator

A constraint generator constructs new analysis variables and constraints. The constraints are attached to the variable and may depend on other analysis variables. Hence connections (ie dependencies) between variables are established. For new analyses, a constraint generator must be provided. It hooks into the framework and is used for the corresponding analysis. The framework already comes with some analyses and therefore their related property spaces and constraint generators. They can be used as a starting point for new analyses.

### 2.5.6. Analysis

An analysis combines a given property space with a constraint generator. It (and its variables) are distinguished from other analyses (and their variables) by an identifier. Already mentioned examples have been  $A$ ,  $\text{REACH}_{\text{in}}$ , and  $\text{REACH}_{\text{out}}$ .



### 2.5.7. Constraint Solver

As presented by Nielson et al. [15], the solver was required to have a *complete* set of constraints from the get go. The modifications made by Jordan [10] include a *lazy-solver* capable of incorporating constraints during the solving process. Furthermore, this addition enables analyses to be more flexible since they can take temporary results into account – partial solutions obtained by the solver can influence the generation process of new constraints.

Additionally, in the original framework constraint guards were required to be *monotone* predicates. With the new *local restart* feature this restriction can be relaxed. This enables the analysis to make assumptions. This can lead to more accurate results, if the assumptions turn out to be valid. But results obtained from an assumption that turns out to be invalid need to be *forgotten* through resetting a set of interdependent analysis variables. This mechanism works fine as long as reset-events do not repeatedly trigger each other<sup>1</sup>, see [10, pp. 216] for more information.

The constraint solver itself is a fixpoint solver which utilises worklists. The algorithm initialises the assignment with the property spaces'  $\perp$  values and updates it step by step until all constraints are satisfied. Section 4.5.5 covers its internal working.

---

<sup>1</sup>This problem does not occur if the constraint set is *stratifiable*.



## 3. Architecture

This chapter provides a grand overview of the system’s architecture. It starts with a brief outline of the Insieme project and continues with investigating the **INSieme Parallel Intermediate REpresentation (INSPIRE)**. Next, each module of the toolkit is described shortly by summarising its purpose.

### 3.1. Insieme

The Insieme project is split into two parts, the Insieme compiler and the Insieme Runtime. Both are designed in such a way that they can also be used separately. Figure 5 displays the typical pipeline where the compiler is combined with the runtime library. Insieme’s mission statement follows. [8]




Figure 5.: The Insieme compiler architecture.

Parallel computing systems have become omnipresent in recent years through the penetration of multi-core processors in all IT markets, ranging from small scale embedded systems to large scale supercomputers. These systems have a profound effect on software development in science as most applications are not designed to exploit multiple cores. The complexity in developing and optimizing parallel programs will rise sharply in the future, as many-core computing systems become highly heterogeneous in nature, integrating general purpose cores with accelerator cores. Modern and in particular future parallel computing systems will be so complex that it appears to be impossible for any human programmer to effectively parallelize and optimize programs across architectures.

The main goal of the *Insieme project* of the *University of Innsbruck* is to research ways of automatically optimizing parallel programs for homogeneous and heterogeneous multi-core architectures and to provide a source-to-source compiler that offers such capabilities to the user.

Insieme’s Runtime maintains a set of worker threads to concurrently processing tasks. It knows about the system architecture and can take records of previously run programs for scheduling and optimisation decisions into account. We won’t concern ourselves with the runtime any further since no interaction with it is required for this work.

### 3. Architecture

Insieme's compiler is a source-to-source compiler taking C/C++ code as input and generating standard C++11 code as output (optionally with OpenCL<sup>1</sup>). Along with this, the frontend accepts the AllScale API<sup>2</sup>, OpenMP<sup>3</sup> and Cilk<sup>4</sup> for parallelism. An important aspect of Insieme is that it is built as a framework, not just a single compiler with a fixed set of options. This design decision keeps Insieme's architecture modular and prevents interference between modules.

Despite resembling a framework, the design is still similar to other compilers. It consists of a frontend, a backend and some complex parts in-between them. The frontend's purpose is to validate and transform input code into **INSPIRE**. The original input code is no longer needed after transformation. The **INSPIRE** program is all that is needed for subsequent steps and is capable of modelling parallel constructs and control-flow. The complex parts are made up of the core module (mainly responsible for managing **INSPIRE**), the analysis module and the transformation module. The analysis module is relevant for identifying code optimisation candidates and feature extraction. It takes a program represented in **INSPIRE** as input and, depending on the configuration, runs various analyses on it. The information derived in this step can further be used in the transformation module to realise optimisations. The optimised program can then be fed to the backend, which is responsible for producing target code.

## 3.2. INSPIRE

In order to use the Insieme infrastructure and create analyses with the new toolkit, a certain understanding of Insieme's **Intermediate Representation (IR)** is required. This section communicates the most important aspects of **INSPIRE**. The information provided here comes in handy, when implementing an analysis.

**INSPIRE** has multiple different representations. Inside the compiler it is handled as a **Direct Acyclic Graph (DAG)**, but a text representation resembling a functional programming language is available as well, for development and testing purposes. *Pretty printer* and *parser* are part of the infrastructure allowing conversions between the **DAG** and text representation, while a binary dumper as well as a **JavaScript Object Notation (JSON)**<sup>5</sup> dumper are provided for exporting **INSPIRE**. These exports can then be used by other tools like the **Haskell-based Analysis Toolkit (HAT)** or **INSPYER** (see Appendix A).

The structure of **INSPIRE** is presented next.

---

<sup>1</sup><https://www.khronos.org/openc1>

<sup>2</sup><http://www.allscale.eu>

<sup>3</sup><http://www.openmp.org>

<sup>4</sup><https://www.cilkplus.org>

<sup>5</sup><http://www.json.org>

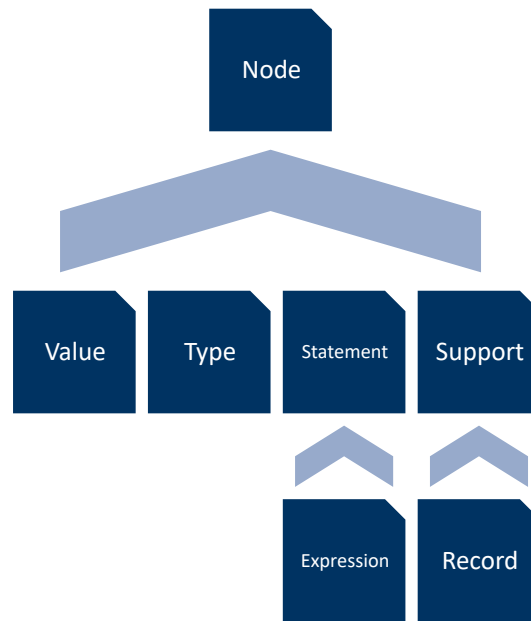


Figure 6.: Hierarchy composed of node categories.

### 3.2.1. Structure

Each node belongs to a category which are hierarchically organised, as can be seen in Figure 6. On top we have the base class `Node` from which the six categories `Value`, `Type`, `Statement`, `Program` and `Support` are derived. Furthermore `Expression` is derived from `Statement` and `Record` is derived from `Support`.

In this subsection the *complete* structure of *INSPIRE* is given, yet only selected node types, relevant for the remainder of this document, are described. For further information see [11] and [10, pp. 37–172].

#### Notation

The block below is just an example used to communicate the notation used in this section. Categories are always written in italic, while the node type defined is written in bold. The types stated after a node definition are the fields (children) of that node. The following block defines the node categories *Example*. It contains four different types of nodes, `BreakStmt`, `CompoundStmt`, `ForStmt` and `DeclarationStmt`. `BreakStmt` has no children and is therefore a terminal. `CompoundStmt` contains a list (annotated by the square brackets) of *Statements*. `ForStmt` holds a `DeclarationStmt`, two *Expressions* followed, by a `CompoundStmt`. Lastly, the `DeclarationStmt` consists of a `Declaration` and a `Variable`. Note that member field indices are currently the main method of addressing sub-structures in *HAT*.

```

Example ::= BreakStmt
           | CompoundStmt [Statement]
           | ForStmt DeclarationStmt Expression Expression CompoundStmt
           | DeclarationStmt Declaration Variable
  
```

### 3. Architecture

The actual structure of [INSPIRE](#) follows.

#### Nodes

```
Node ::= Value  
        | Type  
        | Statement  
        | Support  
        | Program [Expression]
```

This category builds the root of the hierarchy and only defines one node type itself. Program represents an entire program with its entry points.

#### Value Nodes

```
Value ::= BoolValue bool  
          | CharValue char  
          | IntValue int  
          | UIntValue unsigned  
          | StringValue std::string
```

These nodes are simple wrappers around values of the stated C++ type.

#### Type Nodes

```
Type ::= FunctionType Types Type UIntValue Types  
        | GenericType StringValue Parents Types  
        | GenericTypeVariable StringValue Types  
        | NumericType Expression  
        | TagType TagTypeReference TagTypeDefinition  
        | TagTypeReference StringValue  
        | TupleType [Type]  
        | TypeVariable StringValue  
        | VariadicGenericTypeVariable StringValue Types  
        | VariadicTypeVariable StringValue
```

GenericTypes and GenericTypeVariables are used for abstract operators (see Section 3.3) and, hence mainly for language extensions. TagTypes, on the other hand, are for records / structs. Yet, these nodes are of little relevance to the program analysis since the semantics are mainly modelled by the language extensions. Also note that the framework assumes a valid input program, therefore no type checking is required by the framework.

## Statement Nodes

```

Statement ::= Expression
           | BreakStmt
           | ContinueStmt
           | GotoStmt StringValue
           | LabelStmt StringValue
           | CompoundStmt [Statement]
           | DeclarationStmt Declaration Variable
           | IfStmt Expression CompoundStmt CompoundStmt
           | WhileStmt Expression CompoundStmt
           | ForStmt DeclarationStmt Expression Expression CompoundStmt
           | MarkerStmt UIntValue Statement
           | ReturnStmt Declaration
           | SwitchStmt Expression SwitchCases CompoundStmt
           | ThrowStmt Expression
           | TryCatchStmt CompoundStmt [CatchClause]

```

The provided statements are equivalent to their C/C++ counterparts, just note the ForStmt which is actually a range-based for-loop. The Expression inside its Declaration marks the start, while the other two Expressions inside the ForStmt are the end and step, respectively. The MarkerStmt wraps another Statement connecting it to a unique identifier which can be used later on. This is especially helpful for testing and debugging purposes.

## Expression Nodes

```

Expression ::= BindExpr FunctionType Parameters CallExpr
              | CallExpr Type Expression [Declaration]
              | CastExpr Type Expression
              | InitExpr GenericType Expression Expressions
              | JobExpr GenericType Expression Expression
              | LambdaExpr FunctionType LambdaReference LambdaDefinition
              | LambdaReference FunctionType StringValue
              | Literal Type StringValue
              | MarkerExpr Type UIntValue Expression
              | TupleExpr TupleType Expressions
              | Variable Type UIntValue

```

Like Statement nodes, Expression nodes are akin to C/C++. Similar to the statements, expressions can be marked too using MarkerExpr. BindExpr enables the construction of closures, while JobExpr is used for parallel constructs.

### 3. Architecture

#### Support Nodes

```
Support ::= Record
| CatchClause Variable CompoundStmt
| Declaration Type Expression
| Expressions [Expression]
| Field StringValue Type
| Fields [Field]
| Lambda FunctionType Parameters CompoundStmt
| LambdaBinding LambdaReference Lambda
| LambdaDefinition [LambdaBinding]
| MemberFunction StringValue BoolValue Expression
| MemberFunctions [MemberFunction]
| Parameters [Variable]
| Parent BoolValue UIntValue Type
| Parents [Parent]
| PureVirtualMemberFunction StringValue FunctionType
| PureVirtualMemberFunctions [PureVirtualMemberFunction]
| SwitchCase Literal CompoundStmt
| SwitchCases [SwitchCase]
| TagTypeBinding TagTypeReference Record
| TagTypeDefinition [TagTypeBinding]
| Types [Type]
```

These nodes are utilised for building up types, expressions and statements. As the name of this node category hints, the purpose of these nodes is to keep the remaining structure of [INSPIRE](#) less complex and easier to work with.

#### Record Nodes

```
Record ::= Struct StringValue Fields Expressions Expressions BoolValue
MemberFunctions PureVirtualMemberFunctions Parents
| Union StringValue Fields Expressions Expressions BoolValue
MemberFunctions PureVirtualMemberFunctions
```

Struct is used to model C++ classes, including virtual functions and inheritance. The Parents node holds a list of parent (ie base) classes plus information about the type of inheritance (public, protected, or private). A Struct is composed of (in order): name, fields, constructors, destructors, destructor virtual flag, member functions, and pure virtual member functions. Since unions can not inherit structure, Union does not feature Parents.



### 3.3. Extension Mechanism

**INSPIRE** supports *language extensions*, similar to (abstract) libraries, which allow the structured integration of complex functionality (eg parallel directives). Insieme itself already uses this mechanism to model the semantics of C++ programs.

Language extensions do not add additional node types, instead they build *abstract data types* (and operators) by combining the previously (Section 3.2) defined node types (core language). An abstract data type (not to be confused with **Algebraic Data Type**) is defined by its semantics and handled as a black box. For example, the abstract data type `int<4>` represents a 4 byte wide integer, yet we do not state *how* this integer is represented in its binary form.

Abstract operators enable us to work with abstract data types, despite not having their inner workings defined. For instance, the abstract operator `int_add` is defined semantically as adding two integers together, however we do not state how exactly this addition is to be performed.

By combining abstract operators we are able to craft new, *derived* operators having their semantics defined programmatically. For instance, given the abstract integer operation modulo (%) and equality (==), we are able to define a derived operator calculating the *greatest common divisor* of two integers. The definition combines aforementioned abstract operators with constructs of the **INSPIRE** core language (eg `IfStmt`, `ReturnStmt`, etc). In **INSPIRE** these derived operators are realised as lambdas.

Each abstract data type and operator is available in **INSPIRE** under its own name and operators are called like regular functions using a `CallExpr`. A more detailed description of the extension mechanism can be found in [10, pp. 119–141].

Extensions are defined in the `lang` subdirectory of the Insieme core module via the use of preprocessor macros. The following list displays the different options available for building abstract data types and operators.

**Abstract / Composed Type** A new, generic abstract data type can be created using the `LANG_EXT_TYPE` macro. The two arguments are the new abstract data type's handle in C++ and the *specification* in **INSPIRE**. The following code snippet creates the boolean and 4 byte wide integer types.

```
1 | LANG_EXT_TYPE(Bool, "bool")
2 | LANG_EXT_TYPE(Int4, "int<4>")
```

Composed types are created by combining type constructors, like tuple, struct, or function type with already declared data types. The following code snippet composes the type `int4_pair` using two 4 byte wide integers.

```
| TYPE_ALIAS("int4_pair", "(int<4>, int<4>)")
```

**Abstract / Derived Operators** An abstract operator is declared using the `LANG_EXT_LITERAL` macro with an id (C++ handle), name, and its respective type. The integer modulo operation and equality are provided as examples.

```
1 | LANG_EXT_LITERAL(IntAdd, "int_mod", "(int<'a>, int<'a>) -> int<'a>")
2 | LANG_EXT_LITERAL(IntEq, "int_eq", "(int<'a>, int<'a>) -> bool")
```

Derived operators are defined using `LANG_EXT_DERIVED`. The name, under which they are available in **INSPIRE** is derived from their C++ handle (first argument) by converting it from *Camel-Case* to *snail\_case* – in this case `int_gcd`.

### 3. Architecture

```
1 | LANG_EXT_DERIVED(IntGcd, "(a : int<'a>, b : int<'b>) -> int<'a> {"
2 | "   if (int_eq(a, 0)) {"
3 | "       return b;"
4 | "   } else {"
5 | "       return int_gcd(int_mod(b, a), a);"
6 | "   }"
7 | "}")
```

**Constants** Constants are actually a special case of abstract operators. They are abstract operators which's type is not a function type. We therefore use the same macro as for creating abstract operators. Here the two boolean constants true and false are declared.

```
1 | LANG_EXT_LITERAL(True, "true", "bool")
2 | LANG_EXT_LITERAL(False, "false", "bool")
```

Constants are used directly in [INSPIRE](#) using a Literal node – no CallExpr is needed.

**Operators and Analyses** The operators of language extensions represent semantics which need to be handled by an analysis. For abstract operators, the analysis designer *has to* provide an OperatorHandler (Section 4.7.1) to interpret their semantics accordingly. This is optional for derived operators, as their semantics are encoded in their definitions.

#### 3.3.1. The Basic Language Extension

The list of node types of [INSPIRE](#) does not contain primitive types and operations like integers or boolean, nor does it feature basic operators like + or ||. These are defined in the *basic language extension* using a similar set of macros as previously introduced by this section.

The content of this language extension can further be divided into logical sections, which are presented next. The implementation of the content covered in this section can be found in:

File: code/core/include/insieme/core/lang/inspire\_api/basic.def

#### Identifiers

The *basic language extension* defines identifiers as a distinct type of the [INSPIRE](#) language. There do not exist any operators for this type.

```
| TYPE(Identifier, "identifier")
```

This abstract data type is used by the identifier analysis in Section 4.10.1 to determine the value of expressions representing identifiers.

#### Arithmetic Types and Operations

This part of the basic language extension defines various numeric types and operations. Here we find the generic integer type int<'a> as well as its specialisations (eg int<4>, int<16>, etc). All primitive

operations (eg `int_add`, `uint_lshift`, etc) are defined as abstract operators. Table 1 shows the five basic integer operations.

Operator / Type	Description
<code>int_add</code> ( <code>int&lt;'a&gt;</code> , <code>int&lt;'a&gt;</code> ) -> <code>int&lt;'a&gt;</code>	Integer addition.
<code>int_sub</code> ( <code>int&lt;'a&gt;</code> , <code>int&lt;'a&gt;</code> ) -> <code>int&lt;'a&gt;</code>	Integer subtraction.
<code>int_mul</code> ( <code>int&lt;'a&gt;</code> , <code>int&lt;'a&gt;</code> ) -> <code>int&lt;'a&gt;</code>	Integer multiplication.
<code>int_div</code> ( <code>int&lt;'a&gt;</code> , <code>int&lt;'a&gt;</code> ) -> <code>int&lt;'a&gt;</code>	Integer division.
<code>int_mod</code> ( <code>int&lt;'a&gt;</code> , <code>int&lt;'a&gt;</code> ) -> <code>int&lt;'a&gt;</code>	Integer modulo.

Table 1.: Abstract operators defined by the *basic language extension* for integer operations.

The semantics implied by these operations needs to be handled by the corresponding analysis (in this case the arithmetic analysis, described in Section 4.10.2).

### Boolean Types and Operations

The boolean type with the two literals `true` and `false` is defined as follows:

```

1 | TYPE(Bool, "bool")
2 | LITERAL(True, "true", "bool")
3 | LITERAL(False, "false", "bool")

```

The three boolean operations *and*, *or*, and *not* are implemented as derived operators modelling short-circuit evaluation. They are summarised by Table 2

Operator / Type	Description
<code>bool_not</code> ( <code>bool</code> ) -> <code>bool</code>	Boolean <i>not</i> .
<code>bool_and</code> ( <code>bool</code> , <code>()=&gt;bool</code> ) -> <code>bool</code>	Boolean <i>and</i> supporting short-circuit evaluation.
<code>bool_or</code> ( <code>bool</code> , <code>()=&gt;bool</code> ) -> <code>bool</code>	Boolean <i>or</i> supporting short-circuit evaluation..

Table 2.: Derived operators defined by the *basic language extension* for boolean operations.

Numeric comparisons also fall in this part and are defined as abstract operators for the types `char`, `int<'a>`, `uint<'a>` and `real<'a>`. The following table (Table 3) summarises them for the type `int<'a>`.

### 3. Architecture

Operator / Type	Description
<code>int_eq</code> ( <code>int&lt;'a&gt;</code> , <code>int&lt;'a&gt;</code> ) -> <code>bool</code>	Integer equality comparison.
<code>int_ne</code> ( <code>int&lt;'a&gt;</code> , <code>int&lt;'a&gt;</code> ) -> <code>bool</code>	Integer inequality comparison.
<code>int_lt</code> ( <code>int&lt;'a&gt;</code> , <code>int&lt;'a&gt;</code> ) -> <code>bool</code>	Integer less than comparison.
<code>int_gt</code> ( <code>int&lt;'a&gt;</code> , <code>int&lt;'a&gt;</code> ) -> <code>bool</code>	Integer greater than comparison.
<code>int_le</code> ( <code>int&lt;'a&gt;</code> , <code>int&lt;'a&gt;</code> ) -> <code>bool</code>	Integer less than or equal comparison.
<code>int_ge</code> ( <code>int&lt;'a&gt;</code> , <code>int&lt;'a&gt;</code> ) -> <code>bool</code>	Integer greater than or equal comparison.

Table 3.: Abstract operators defined by the *basic language extension* for integer comparison.

#### 3.3.2. Reference Extension

In **INSPIRE**, references are the only means to address memory locations and memory locations are used for modelling mutual state.

Implicit and explicit C/C++ references are modelled by this extension. Implicit references are translated to *plain* references in **INSPIRE**. Explicit references, as available in C++, are mapped to `cpp_ref` or `cpp_rref` depending on whether we are dealing with an ordinary, l-value, or r-value reference. In addition to the reference kind, it can be stated if the referenced variable is `const`, `volatile` or both.

Several operators are provided for memory management. For allocation the generic operator `ref_alloc` is provided returning a reference to a memory location allocated to fit an element of the requested type. This function can be combined with memory location information to specify whether the data is allocated on the heap or stack. `ref_delete` allows memory to be freed after it is no longer needed.

Accessing the data stored in a memory location addressed by a reference for reading is done by using `ref_deref` and `ref_assign` for writing. The following two tables provide an overview of common abstract (Table 4) and derived (Table 5) operators defined by this language extension.

File: `code/core/include/insieme/core/lang/reference.h`

#### 3.3.3. Parallel Extension

Up until this point we have not really talked about the parallel constructs available in **INSPIRE**. This extension covers the relevant primitives. At the most fine grained level **INSPIRE** provides parallelism by allowing multiple threads to be spawned at once, resulting in a *thread group* and joining them later on. The thread group is always tasked with a specific *job*. Basic identification queries like getting the index of the current thread in a group and getting the overall group size are available. Common parallel directives like barriers, mutexes and atomics are provided as well.

Operator / Type	Description
<code>ref_alloc</code> ( <code>type&lt;'a&gt;</code> , <code>memloc</code> ) -> <code>ref&lt;'a,f,f&gt;</code>	Allocates memory for an object of given type at a given memory location.
<code>ref_decl</code> ( <code>type&lt;ref&lt;'a,'c,'v,'k&gt;&gt;</code> ) -> <code>ref&lt;'a,'c,'v,'k&gt;</code>	References memory allocated in a surrounding declaration context.
<code>ref_delete</code> ( <code>ref&lt;'a,f,'v&gt;</code> ) -> <code>unit</code>	Frees memory of the given reference.
<code>ref_deref</code> ( <code>ref&lt;'a,'c,'v,'k&gt;</code> ) -> <code>'a</code>	Used to obtain the data stored in the memory location linked to the given reference.
<code>ref_assign</code> ( <code>ref&lt;'a,'c,'v,'k&gt;</code> , <code>'a</code> ) -> <code>unit</code>	Used to update the value stored in the memory location linked to the given reference.
<code>ref_reinterpret</code> ( <code>ref&lt;'a,'c,'v,'k&gt;</code> , <code>type&lt;'b&gt;</code> ) -> <code>ref&lt;'b,'c,'v,'k&gt;</code>	A reinterpret cast altering the actual interpretation of the referenced memory cell.
<code>ref_narrow</code> ( <code>ref&lt;'a,'c,'v,'k&gt;</code> , <code>datapath&lt;'a,'b&gt;</code> ) -> <code>ref&lt;'b,'c,'v,'k&gt;</code>	Obtain a reference to a sub-object within a referenced object.
<code>ref_expand</code> ( <code>ref&lt;'b,'c,'v,'k&gt;</code> , <code>datapath&lt;'a,'b&gt;</code> ) -> <code>ref&lt;'a,'c,'v,'k&gt;</code>	The inverse operation to <code>ref_narrow</code> .
<code>ref_null</code> ( <code>type&lt;'a&gt;</code> , <code>type&lt;'a&gt;</code> , <code>type&lt;'v&gt;</code> ) -> <code>ref&lt;'a,'c,'v,plain&gt;</code>	Creates a <i>null-reference</i> pointing to no memory location.
<code>ref_cast</code> ( <code>ref&lt;'a,'c,'v,'k&gt;</code> , <code>type&lt;'nc&gt;</code> , <code>type&lt;'nv&gt;</code> , <code>type&lt;'nk&gt;</code> ) -> <code>ref&lt;'a,'nc,'nv,'nk&gt;</code>	A simple reference cast merely altering the view on the otherwise untouched memory location.

Table 4.: Abstract operators defined by the reference language extension.

Operator / Type	Description
<code>ref_new</code> ( <code>type&lt;'a&gt;</code> ) -> <code>ref&lt;'a,f,f&gt;</code>	Allocates memory for an object of given type on the heap by using <code>ref_alloc</code> .
<code>ref_array_element</code> ( <code>ref&lt;array&lt;'a,'s&gt;,'c,'v,plain, int&lt;8&gt;&gt;</code> ) -> <code>ref&lt;'a,'c,'v,plain&gt;</code>	Provides access to an element in an array by using <code>ref_narrow</code> .
<code>ref_member_access</code> ( <code>ref&lt;'a,'c,'v,'k&gt;</code> , <code>identifier</code> , <code>type&lt;'b&gt;</code> ) -> <code>ref&lt;'b,'c,'v,plain&gt;</code>	Provides access to an element of a struct / union by using <code>ref_narrow</code> .
<code>ref_scalar_to_array</code> ( <code>ref&lt;'a,'c,'v,plain&gt;</code> ) -> <code>ref&lt;array&lt;'a&gt;,'c,'v,plain&gt;</code>	A reference-navigation operator providing an array view on a scalar using <code>ref_expand</code> .

Table 5.: Derived operators defined by the reference language extension.

### 3. Architecture

This extensions will by investigated further as it is not relevant to this work.

File: code/core/include/insieme/core/lang/parallel.h

## 3.4. Semantics

In this section we look at some example input codes, and compare them to the corresponding **INSPIRE** representations.

### Empty Program

Below you can see the C++ input program on the left, the corresponding **INSPIRE** pretty printer output on the right and the simplified tree structure below. Omissions are indicated by an ellipsis.

```
1 | int main() {
2 |     return 0;
3 | }

1 | decl IMP_main : () -> int<4>;
2 | // Inspire Program
3 | int<4> function IMP_main (){
4 |     return 0;
5 | }
```

```
1 | (Program |
2 |   (LambdaExpr |
3 |     (FunctionType | ... )
4 |     (LambdaReference |
5 |       (FunctionType | ... )
6 |       (StringValue "IMP_main"))
7 |     (LambdaDefinition |
8 |       (LambdaBinding |
9 |         (LambdaReference | ... )
10 |        (Lambda |
11 |          (FunctionType | ... )
12 |          (Parameters )
13 |          (CompoundStmt |
14 |            (ReturnStmt | ... ))))))))
```

Listing 3: inspire\_semantics/empty.{cpp,ir,tree}

The pretty printer output starts off with a declaration of the function `IMP_main`, here **IMP** stands for **Insieme Mangling Prefix**. It is followed by the definition of `IMP_main`. `() -> int<4>` is the type of `IMP_main` and states that no parameters are taken and an integer of 4 bytes is returned.

In the tree structure we see that the root node is a `Program` containing a `LambdaExpr`. Every function in the input code is translated to a lambda in **INSPIRE** (unless inlined). The lambda is composed of `LambdaExpr`, `LambdaReference`, `LambdaDefinition`, `LambdaBinding` and finally `Lambda` nodes. The `Lambda` node holds the `CompoundStmt` of the `main` function with its single `ReturnStmt`. Note that there are no `Declarations` in the actual **IR** structure. The `Declaration` seen in the pretty print is a presentation artefact. So are comments.

### Variable Declaration

Next we define a single integer variable and observe the resulting **INSPIRE** output.

```

1 | int main() {
2 |     int i = 42;
3 |     return 0;
4 | }

1 | decl IMP_main : () -> int<4>;
2 | // Inspire Program
3 | int<4> function IMP_main (){
4 |     var ref<int<4>,f,f,plain> v1 = 42;
5 |     return 0;
6 | }

1 | (Program | ...
2 |   (CompoundStmt |
3 |     (DeclarationStmt |
4 |       (Declaration |
5 |         (GenericType | ... )
6 |         (Literal |
7 |           (GenericType | ... )
8 |           (StringValue "42"))))
9 |       (Variable |
10 |        (GenericType | ... )
11 |        (UIntValue 1)))
12 |     (ReturnStmt | ... )...))

```

Listing 4: inspire\_semantics/decl.{cpp,ir,tree}

The original variable `i` gets translated to `v1` of type `ref<int<4>,f,f,plain>`, meaning it is a *plain*<sup>1</sup> reference to an integer. The reference is neither `const` nor `volatile`.

The `DeclarationStmt` creates an uninitialised memory location for a 4 byte integer on the stack. It is automatically destroyed when the reference goes out of scope. The value 42 is used to initialise the memory location created by the `Declaration`, which ends up being referenced by `v1`. All variables are indexed by natural numbers. In our case it receives an id of 1.

### Basic Arithmetic

Now we add some basic arithmetic by multiplying two integer variables.

<sup>1</sup>A reference might either be `plain`, `cpp_ref`, or `cpp_rref` depending on its C++ semantic.

### 3. Architecture

```

1 | int main() {
2 |     int a = 2;
3 |     int b = 3;
4 |     int c = a * b;
5 |     return 0;
6 | }

1 | decl IMP_main : () -> int<4>;
2 | // Inspire Program
3 | int<4> function IMP_main (){
4 |     var ref<int<4>,f,f,plain> v1 = 2;
5 |     var ref<int<4>,f,f,plain> v2 = 3;
6 |     var ref<int<4>,f,f,plain> v3 = (*v1) * (*v2);
7 |     return 0;
8 | }

1 | (Program | ...
2 |   (CompoundStmt |
3 |     (DeclarationStmt | ... )
4 |     (DeclarationStmt | ... )
5 |     (DeclarationStmt |
6 |       (Declaration |
7 |         (GenericType | ... )
8 |         (CallExpr |
9 |           (GenericType | ... )
10 |          (Literal |
11 |            (FunctionType | ... )
12 |            (StringValue "int_mul")))
13 |        (Declaration |
14 |          (GenericType | ... )
15 |          (CallExpr |
16 |            (GenericType | ... )
17 |            (Literal |
18 |              (FunctionType | ... )
19 |              (StringValue "ref_deref")))
20 |        (Declaration |
21 |          (GenericType | ... )
22 |          (Variable |
23 |            (GenericType | ... )
24 |            (UIntValue 1))))))
25 |   (Declaration |
26 |     (GenericType | ... )
27 |     (CallExpr |
28 |       (GenericType | ... )
29 |       (Literal |
30 |         (FunctionType | ... )
31 |         (StringValue "ref_deref")))
32 |   (Declaration |
33 |     (GenericType | ... )
34 |     (Variable |
35 |       (GenericType | ... )
36 |       (UIntValue 2))))))
37 |   (Variable |
38 |     (GenericType | ... )
39 |     (UIntValue 3)))
40 |   (ReturnStmt | ... ))...)

```

Listing 5: inspire\_semantics/mul.{cpp,ir,tree}

Again, the variables  $a$  and  $b$  are translated to  $v_1$  and  $v_2$  respectively, referencing memory locations storing their corresponding integer value. When multiplying the referenced value, we need to dereference first using the dereference operator  $*$ . In C/C++ this *variable-dereferencing* is implicit, while in **INSPIRE** those steps are explicit.

Multiplication and dereferencing are part of the arithmetic and reference extension. They are provided as operators and can be called via `CallExpr`.

In the `CompoundStmt` (line 2 in `mul.tree`) we see two `DeclarationStmt`s on top. They contain the same structure as in the previous example. The third `DeclarationStmt`, on the other hand, holds two `CallExpr`s (lines 15 and 27 in `mul.tree`) for dereferencing the variables  $v_1$  and  $v_2$  and another `CallExpr` (line 8 in `mul.tree`) for multiplying the two referenced integer values.



## Function Call

Let us investigate how parameters and return values are passed. We therefore call a function `square` passing in an integer parameter and receive the squared value as return.

```

1 | int square(int x) {
2 |     return x * x;
3 | }
4 |
5 | int main() {
6 |     int i = square(3);
7 |     return 0;
8 | }

1 | decl IMP_main : () -> int<4>;
2 | decl IMP_square : (int<4>) -> int<4>;
3 | def IMP_square = function (v1 : ref<int<4>,f,f,plain>)
4 |     -> int<4> {
5 |         return (*v1) * (*v1);
6 |     };
7 | // Inspire Program
8 | int<4> function IMP_main (){
9 |     var ref<int<4>,f,f,plain> v2 = IMP_square(3);
10 |     return 0;
11 | }

1 | (Program | ...
2 |   (CompoundStmt |
3 |     (DeclarationStmt |
4 |       (Declaration |
5 |         (GenericType | ... )
6 |         (CallExpr |
7 |           (GenericType | ... )
8 |           (LambdaExpr |
9 |             (FunctionType | ... )
10 |            (LambdaReference | ... )
11 |            (LambdaDefinition |
12 |              (LambdaBinding |
13 |                (LambdaReference | ... )
14 |                (Lambda |
15 |                  (FunctionType | ... )
16 |                  (Parameters |
17 |                    (Variable |
18 |                      (GenericType | ... )
19 |                      (UIntValue 1)))
20 |                  (CompoundStmt |
21 |                    (ReturnStmt |
22 |                      (Declaration | ... ))))))))
23 |         (Declaration |
24 |           (GenericType | ... )
25 |           (Literal |
26 |             (GenericType | ... )
27 |             (StringValue "3")))))
28 |       (Variable |
29 |         (GenericType | ... )
30 |         (UIntValue 2)))
31 |     (ReturnStmt | ... )...

```

Listing 6: `inspire_semantics/call.{cpp,ir,tree}`

In addition to the `IMP_main` declaration we now also have a declaration of `IMP_square` receiving and returning an `int<4>`. In the function definition the variable `v1` references the passed-in value. Note that since `C/C++` is implicitly allocating memory for parameters passed by value, the parameter type is actually `ref<int<4>,f,f,plain>` instead of `int<4>`.

In the [Abstract Syntax Tree \(AST\)](#), the `CallExpr` (line 6 in `call.tree`) directly contains the `LambdaExpr` (line 8 in `call.tree`) holding the translated `IMP_square` function plus the argument to the function, in this case the literal `3` (line 25 in `call.tree`). The list of `Declarations` (line 23 in `call.tree`) in `CallExpr` reflects the arguments of the call. The `IMP_square` function's body is composed of a `CompoundStmt` (line 20 in `call.tree`) with a single `ReturnStmt` (line 21 in `call.tree`). The squaring of values happens inside the `Declaration` (line 22 in `call.tree`) like shown in the previous example.

### 3. Architecture

#### For Loop

This example is given to illustrate the use of ForStmt.

```
1 | int main() {
2 |     int sum = 0;
3 |     for (int i = 0; i < 10; i++) {
4 |         sum += i;
5 |     }
6 |     return 0;
7 | }

1 | decl IMP_main : () -> int<4>;
2 | // Inspire Program
3 | int<4> function IMP_main (){
4 |     var ref<int<4>,f,f,plain> v1 = 0;
5 |     {
6 |         for( int<4> v20 = 0 .. 10 : 1) {
7 |             comp_assign_add(v1, v20);
8 |         }
9 |     }
10 |     return 0;
11 | }
```

```
1 | (Program | ...
2 |   (CompoundStmt |
3 |     (DeclarationStmt |
4 |       (Declaration | ... )
5 |       (Variable |
6 |         (GenericType | ... )
7 |         (UIntValue 1)))
8 |     (CompoundStmt |
9 |       (ForStmt |
10 |        (DeclarationStmt |
11 |          (Declaration |
12 |            (GenericType | ... )
13 |            (Literal |
14 |              (GenericType | ... )
15 |              (StringValue "0"))))
16 |          (Variable |
17 |            (GenericType | ... )
18 |            (UIntValue 20)))
19 |        (Literal |
20 |          (GenericType | ... )
21 |          (StringValue "10"))
22 |        (Literal |
23 |          (GenericType | ... )
24 |          (StringValue "1"))
25 |        (CompoundStmt |
26 |          (CallExpr |
27 |            (GenericType | ... )
28 |            (LambdaExpr |
29 |              (FunctionType | ... )
30 |              (LambdaReference |
31 |                (FunctionType | ... )
32 |                (StringValue "comp_assign_add")))
33 |              (LambdaDefinition | ... ))
34 |            (Declaration |
35 |              (GenericType | ... )
36 |              (Variable |
37 |                (GenericType | ... )
38 |                (UIntValue 1)))
39 |            (Declaration |
40 |              (GenericType | ... )
41 |              (Variable |
42 |                (GenericType | ... )
43 |                (UIntValue 20))))))
44 |      (ReturnStmt | ... ))...)
```

Listing 7: inspire\_semantics/for.{cpp,ir,tree}

The ForStmt models a range-based for-loop with its beginning, end, and step size. The compound assignment operator += inside the for-loop is translated to the comp\_assign\_add derived operator.

For simplicity the Insieme frontend is translating every for-loop into a while-loop in a first step, before

attempting to restore for-loops qualifying as range-based loops.

When the original for-loop was translated to a WhileStmt, it was placed in a new CompoundStmt (line 8 in `for.tree`) to account for the scoping of the index variable. This CompoundStmt remains even after the WhileStmt has been converted to a ForStmt. The ForStmt consists of a DeclarationStmt (line 10 in `for.tree`) initialising the loop index, two literals (lines 19 and 22 in `for.tree`) for marking the end and stepsize, and finally a CompoundStmt for the translated body of the original for-loop (line 25 in `for.tree`. Here it contains the call to the `comp_assign_add` operator (line 26 in `for.tree`).

The derived operator `comp_assign_add` is a lambda, including an implementation defining its semantic by combining other primitives.

### 3.5. Haskell-based Analysis Toolkit

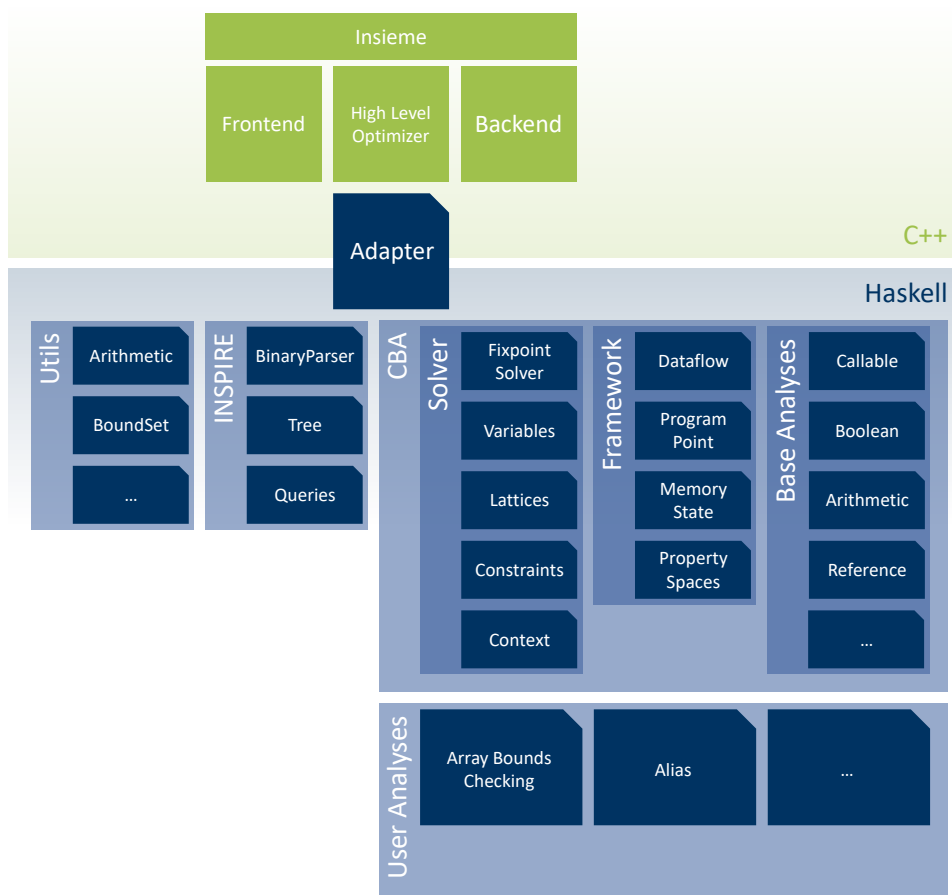


Figure 7.: Simplified architecture of HAT, interfacing with Insieme.

In Figure 7 we see a slightly simplified design of HAT. At the top we can see Insieme with its basic structure and beneath it the Adapter module attaching the toolkit to the High Level Optimiser. The INSPIRE module is responsible for holding an input program and providing queries to specific parts of the underlying data structure. The analysis module, here titled CBA, can be divided into three

### 3. Architecture

submodules. Note that the architecture may evolve over time and certain parts of the module layout may be revisited in the future. This document describes the state of the toolkit at the time of writing (February 2017).

**Adapter** The adapter is the only module of the toolkit existing in both, the C++ and Haskell domain of Insieme. It therefore bridges these two worlds via a *foreign function interface*. Data, relevant for the analysis, can be passed in both directions. Typically, an **INSPIRE** representation of an input program is transferred to the toolkit plus the information what analysis should be run on which part of the program. After the analysis has been completed the result may be transferred back to the C++ domain. Depending on the chosen analysis, the result can have many different formats (see analyses presented in Section 4.10).

**Utilities** The two interesting submodules in Utils are Arithmetic and BoundSet. Arithmetic provides a data structure for simple arithmetic formulas, which can be combined using basic arithmetic operations. BoundSet is also a data structure, similar to the regular `Data.Set` found in Haskell's `containers` package<sup>1</sup>. The main difference is that you can set an upper limit on how many elements the set can hold. Upon adding another element, that is not already a member of the set, the set turns into `Universe` representing every possible value of the domain. This mechanism is required for some property spaces to ensure corresponding analyses to terminate.

**INSPIRE** The input program, represented in **INSPIRE**, is transferred to the toolkit through a binary export. The **INSPIRE** module maintains a `BinaryParser` which takes the binary encoding as input and creates a representation of it, in the form of a *rose tree*, in Haskell. The related concepts of *node paths* and *node addresses* are translated as well and a collection of queries eases working with the tree.

A rose tree is defined as a tree where each node can have a variable number of children. The most basic definition looks like this (in Haskell):

```
| data Tree a = Node a [Tree a]
```

**Framework** During an analysis the toolkit constructs additional data structures which expose critical information for analyses. These data structures include **INSPIRE**, Program Points, and Memory State information but can be extended as needed. The framework is responsible for elevating a user defined analysis to declarations, parameter and return value passing, closures, calls to unknown external functions, inter-procedural analyses and memory locations automatically. This is illustrated in Chapter 5 and achieved via the use of base analyses.

**Base Analyses** The generic **Data-Flow Analysis (DFA)**, part of the framework, not only utilises multiple of these specialised analyses, but the Base Analyses themselves are specialisations of the generic **DFA**. Therefore the Base Analyses and Framework module depend on each other. These Base Analyses are necessary to correctly deduce the control flow of a program. For instance, the Arithmetic and Boolean value analysis are important for branches and loops, while `Callable` is vital for inter-procedural analyses.

---

<sup>1</sup><https://hackage.haskell.org/package/containers>

**User Analyses** Last but not least, user defined analyses build on top of the framework. While the illustration shows them combined into a single module, this is not necessary. How complex these analyses are is left to the user and we will see how Array Bounds Checking is realised later on in Chapter 5. Note that the framework imposes only few restrictions on what the property space can be. The downside of this feature is that the adapter needs to be extended in some cases. Otherwise the result cannot be used outside the Haskell domain.



## 4. Framework Implementation

Now, that our stage has been properly decorated with the required background knowledge and a grand overview, we can dive into the actual implementation of the framework.

First we take a look at the build infrastructure and the adapter used for communication between Haskell and C++. The modelling of the *INSieme Parallel Intermediate REpresentation (INSPIRE)* in Haskell follows. The fixpoint solver; together with modelling of arrays, structures, and unions; program points; and memory state is presented next. The final part investigates the base analyses required and provided by the framework.

### 4.1. Build Process

First, we examine the build process. Remember that the Adapter module exists in both, the Haskell and C/C++ domain, as it bridges the gap between them. Because of this, we do not solely rely on the Haskell Stack<sup>1</sup> build tool but also on CMake<sup>2</sup>, utilised by Insieme. For Insieme's fully automated build process, CMake invokes Haskell Stack with the required parameters to build the Haskell modules of the toolkit.

#### 4.1.1. Compiling HAT into a Library

Haskell Stack is not only designed for automating the build process of Haskell packages, but also serves as a package manager to download, build and install third-party packages. Packages for this mechanism are hosted on Stackage<sup>3</sup> in different versions. Furthermore Haskell Stack also sets up a suitable compiler for the chosen *release*. On Stackage a release, like `1ts-7.0`, groups together a set of packages, a specific version for each of these packages, and a specific version of the *Glasgow Haskell Compiler (GHC)*.

In our case we are using the release `1ts-6.17` which brings `GHC-7.10.3` along. This release contains almost all packages we need, except `attoparsec-binary`<sup>4</sup> and `directory`<sup>5</sup>. Fortunately, Haskell Stack allows us to specify these requirements in the corresponding config file (`stack.yaml`) where we also state which release we are using. When specifying these extra dependencies, we also need to supply a specific version for each of them. In our case the entries are `attoparsec-binary-0.2` and `directory-1.2.6.3`.

Haskell Stack downloads all required packages and the Haskell compiler into a directory referred to as Stack Root. By default this is a folder named `.stack` inside the home-directory of the user invoking

---

<sup>1</sup><https://www.haskellstack.org>

<sup>2</sup><https://cmake.org>

<sup>3</sup><https://www.stackage.org>

<sup>4</sup><https://www.stackage.org/package/attoparsec-binary>

<sup>5</sup><https://www.stackage.org/package/directory>

## 4. Framework Implementation

the Haskell Stack binary `stack`. This default setting can be overwritten by setting the `STACK_ROOT` environment variable when executing `stack`. The version of `stack` we are currently using (1.0.4) turned out to be quite picky about the ownership of the directory containing the Stack Root folder. If the directory belongs to a different user than the one invoking `stack`, it complains – even though the user has sufficient permissions to create the Stack Root folder.

A pre-compiled binary for Haskell Stack is available for common distributions and allows one to use it without having any Haskell related stuff installed. The only dependencies are the [GNU Multiple Precision \(GMP\)](#) and `zlib` libraries, which are widely available for any Linux based platform.

### 4.1.2. Linking the Shared Libraries

Stack itself is capable of building a shared library out of a Haskell project. Typically each third-party package is built as a shared library and so is the [Haskell-based Analysis Toolkit \(HAT\)](#). The library resulting from HAT, named `libHSinsieme-hat-0.2.so`, depends on the shared libraries of the third-party packages. Furthermore, neither the shared library of HAT nor any of the other third-parties contain the Haskell runtime environment. We therefore need to link the shared library `libHSrts-ghc7.10.3.so` containing the runtime explicitly when building Insieme.

During this process we discovered that Stack, at least in the version used, strips symbols, which are required for the linking process, from a shared library. It is therefore required to patch Haskell Stack, and hence remove this behaviour, since there was no option provided to prevent this from happening. The patch can be viewed in Listing 8 and shows that the only change needed is to pass the `--disable-library-stripping` flag to the Cabal<sup>1</sup> library. The Cabal library is responsible for the compilation process and used by Stack. Luckily Insieme comes with a dependencies installer capable of building third-party software from source and injecting patches.

```
1  --- a/src/Stack/Types/Build.hs
2  +++ b/src/Stack/Types/Build.hs
3  @@ -606,6 +606,7 @@ configureOptsNoDir :: EnvConfig
4      -> [String]
5  configureOptsNoDir econfig bco deps wanted isLocal package = concat
6      [ depOptions
7  +   , ["--disable-library-stripping"]
8     , ["--enable-library-profiling" | boptsLibProfile bopts || boptsExeProfile bopts]
9     , ["--enable-executable-profiling" | boptsExeProfile bopts && isLocal]
10    , ["--enable-split-objs" | boptsSplitObjs bopts]
```

Listing 8: Patch for Haskell Stack preventing stripping of symbols from shared libraries.

A different issue, mainly arising from the way CMake operates, is *finding* the resulting shared library. Files generated during the build process, as well as the output files, are placed inside a folder named `.stack-work` inside the project directory. In our case the wanted shared library is located at

```
.stack-work/dist/x86_64-linux/Cabal-1.22.5.0/build/libHSinsieme-hat-0.2-7Yep1HX5wt11qxdm
3bLxLA-ghc7.10.3.so.
```

Yet, the `.stack-work` directory does not exist prior to the build process and certain meta information is encoded in the path. Among architecture (x86\_64), Cabal library version (1.22.5.0), and GHC

---

<sup>1</sup><https://www.haskell.org/cabal>



version (7.10.3), a so-called *package key*<sup>1</sup> (7Yep1HX5wt11qxdm3bLxLA) is present. This package key is a hash of the package name, package version, and its dependencies. It is therefore very likely to change during development. CMake requires one to explicitly state the name of a shared library for linking plus its location, which must be known during the CMake run – before building the library. In our case, since the file name is likely to change, a workaround has been put into place. A script named `copyLibs.hs` is executed by CMake after invoking `stack`, which finds the shared library inside the `.stack-work` directory using a `glob`<sup>2</sup> and copies it to CMake’s build directory of HAT. Version and package key are stripped resulting in the file name `libHSinsieme-hat.so`.

Now linking this shared library can be stated explicitly in the CMake configuration, but the linker proofs to be to smart for our con. The original file name (containing version and package key) is present inside the shared library’s meta data<sup>3</sup> which is used by the linker. The linker now complains that it cannot find the file with this name – so we need another workaround for the workaround. This is achieved by copying the library to the build directory not only once, but twice. Once with the version and package key stripped from the file name so we can state it in the CMake configuration and once without anything stripped so the linker finds the correct file. This is to be found in the following file of the Insieme code base:

File: `code/analysis/src/cba/haskell/insieme-hat/CMakeLists.txt`

## 4.2. Adapter

The Adapter is responsible for attaching HAT to the Insieme compiler. It therefore consists of Haskell, as well as, C/C++ code. Through the use of a **Foreign Function Interface (FFI)** data can be transferred between the two domains. This works by exporting Haskell and C++ function in such a way that they can be linked and called like regular C functions. It follows that the names are not mangled and the platform’s default *C calling convention*<sup>4</sup> is used for these functions.

### 4.2.1. Foreign Function Interface

Both languages provide a simple mechanism to export functions for FFI. Two examples, one for C++ (left) and one for Haskell (right) are given.

<pre> 1   #include &lt;iostream&gt; 2   3   extern "C" { 4       void hello_world(void) { 5           std::cout &lt;&lt; "Hello World" 6               &lt;&lt; std::endl; 7       } 8   }</pre>	<pre> 1   {-# LANGUAGE ForeignFunctionInterface #-} 2   3   foreign export ccall "hello_world" 4       helloWorld :: IO () 5   6   helloWorld = putStrLn "Hello World"</pre>
--	--

In C++ the directive `extern "C"` is all that is required to export the function correctly. In the Haskell domain we first have to declare that we are using the FFI language extension in the first line. Next, in addition to a regular function definition (line 6) we declare that a function should be exported using C calling convention (`ccall`) and with the symbol `hello_world` in line 3. The following line

<sup>1</sup><https://ghc.haskell.org/trac/ghc/wiki/Commentary/Packages/Concepts>

<sup>2</sup>[https://en.wikipedia.org/wiki/Glob\\_\(programming\)](https://en.wikipedia.org/wiki/Glob_(programming))

<sup>3</sup>see Executable and Linkable Format (ELF)

<sup>4</sup>[https://en.wikipedia.org/wiki/Calling\\_convention](https://en.wikipedia.org/wiki/Calling_convention)

#### 4. Framework Implementation

(4) belongs to that **FFI** declaration and states which Haskell function is to be exported, including the type signature of that function. Note that all functions exported via the **FFI** live in the `IO` monad.

That alone is quite nice, but not very useful. We also require to *import* a function exposed via the **FFI**. The following code does exactly that.

```
1 | extern "C" {
2 |     void hello_world(void);
3 | }
4 | {-# LANGUAGE ForeignFunctionInterface #-}
5 | foreign import ccall "hello_world"
6 |     helloWorld :: IO ()
```

In C++ we just need to wrap the function declaration inside `extern "C"`. For Haskell instead of `export` we now use `import` followed by stating C calling convention and the name of the symbol to import. The following line (4) states under which name that function is made available in Haskell and which type signature it has.

While this example does not illustrate data being transferred between the two domains, it shows the basic structure of building an **FFI**. Primitives<sup>1</sup> can be passed back and forth just like regular function parameters / return values. Haskell's `Foreign.Ptr` and `Foreign.C.Types` modules contain the relevant types for sending and receiving data over the **FFI**. Additionally the module `Foreign.C.String` is dedicated to transferring C strings.

In this context we have two different types of pointers, depending on where the target pointed to is located. If the pointer refers to an object living in the C++ domain we use the generic type `Ptr a` in Haskell. When exposing a pointer to a Haskell object<sup>2</sup> we use `StablePtr a`. The function

```
| newStablePtr :: a -> IO (StablePtr a)
```

gives us a pointer to a Haskell object and marks the object as *stable*, meaning it will not be moved or deallocated by the garbage collector. Similar to dynamic memory management, we later have to *un-stable* the object by calling `freeStablePtr`. Note that the pointers can be passed back and forth between the two domains as needed. To dereference a pointer in Haskell one can either use `deRefStablePtr` or `peek` depending on the type.

Most of the time we do not use `peek` directly. Since working with arrays and strings is quite common, Haskell serves us some convenience functions like `peekArray` and `withArrayLen` for reading, and `pokeArray` for writing.

Despite the fact that *marshalling*<sup>3</sup> can be done in Haskell, we do not use it in the framework. **INSPIRE**, being the most complex data structure transferred from C++ to Haskell, is exported to its binary format, which is then parsed into a rose tree (see Section 3.5). To pass complex data structures from Haskell to C++, we first recreate the data structure in C++ and export its constructor(s) via **FFI**. An example of this is provided in the `Arithmetic` module (Section 4.4.2).

#### 4.2.2. Haskell RunTime System

In order to use Haskell in a C++ project, like *Insieme*, we not only have to link together all the correct parts and use a correct **FFI**. We also have to initiate (and tear down) the **Haskell RunTime System** (**Haskell RTS**). The two functions `hs_init` and `hs_exit` fulfil this functionality, where the initialise

<sup>1</sup>integers, doubles and pointers

<sup>2</sup>By *Haskell object* in C++ we actually mean an instantiated **Algebraic Data Type** (ADT)

<sup>3</sup>[https://en.wikipedia.org/wiki/Marshalling\\_\(computer\\_science\)](https://en.wikipedia.org/wiki/Marshalling_(computer_science))

function optionally takes arguments which are then passed to the Haskell RTS. Since GHC does no longer support calling `hs_exit` from a non-static object's destructor, the Haskell RTS is initialised and terminated by a static object. No arguments are passed to the runtime.

```

1 | extern "C" {
2 |     void hs_init(int*, const char**[]);
3 |     void hs_exit(void);
4 | }
5 |
6 | class Runtime {
7 | public:
8 |     Runtime() { hs_init(nullptr, nullptr); }
9 |     ~Runtime() { hs_exit(); }
10| } rt;

```

File: code/analysis/src/cba/haskell/runtime.cpp

### 4.2.3. Context

For easier interaction between the two domains a *context* is established. This *context* consists of a Haskell object and a C++ object which are linked to each other by pointers. We will see its benefits when talking about the [Intermediate Representation \(IR\)](#) components in [Section 4.3](#).

On the Haskell end, the *context* is used to retain a received [INSPIRE](#) program and the solver's state. Holding on to the solver's state is important to reuse already computed information for new analysis runs. When using the toolkit without the Insieme compiler, this *context* can still be used by setting the reference to its corresponding C++ object to `nullptr`. By allowing for this *dummy context* we do not require an additional interface for the toolkit because it is used without Insieme.

## 4.3. IR Components

Next, we take a look at the components relevant to model [INSPIRE](#). The bottom layer is composed of the node types and a rose tree. Node addresses, visitors and queries build upon this. An [INSPIRE](#) program is converted to its binary representation (using Insieme's binary dumper) and passed to the toolkit via the [FFI](#). The binary dump's format can be inferred manually from the Insieme source code<sup>1</sup>. Relevant parts are covered in the presentation of the binary parser.

The binary dump, representing an [INSPIRE](#) program also comes with a list that states operators used by that program together with a node address pointing to the function describing it.

### 4.3.1. Node Types

The full set of node types is defined in the `ir_nodes.def` file of the Insieme core. It is processed via the use of the C pre-processor to construct an enum holding all node types, which is then used by Insieme. Luckily we also have the option to process such a file in Haskell using `c2hs`<sup>2</sup>. With

<sup>1</sup>code/core/src/dump/binary\_dump.cpp

<sup>2</sup><https://github.com/haskell/c2hs>

#### 4. Framework Implementation

this tool, we can get a full list of available node types in exactly the same order as the Insieme core does. Maintaining the order is important since the binary dump simply uses the index in this enum for communicating the node type. Yet we do not use this enum directly. What we seek is something like this:

```
1 | data NodeType = BoolValue   Bool
2 |                 | CharValue  Char
3 |                 | IntValue   Int
4 |                 | UIntValue  Int
5 |                 | StringValue String
6 |                 | CompoundStmt
7 |                 | IfStmt
8 |                 | WhileStmt
9 |                 | CompoundStmt
10 | --                 | ...
```

This data structure is a list of *all* possible node types, where the *value nodes* have their value already attached. With this approach we run into two problems. First, we cannot derive this exact type definition from `ir_nodes.def` using `c2hs`. Second, we cannot let the compiler derive an enum instance for this data structure. To circumvent these two problems we have to resort to [Template Haskell \(TH\)](#). [TH](#) offers a way to modify the Haskell syntax tree prior to compilation – it is basically a pre-process.

First the `ir_nodes.def` file gets included filling a C enum that is then converted to Haskell using `c2hs`. All entries of this enum are prefixed with `NT_`. Based on this enum [TH](#) is used to define a data structure `NodeType` similar to the enum and two functions `fromNodeType` and `toNodeType`. `NodeType` is composed of all node types (without the `NT_` prefix) and value nodes have their corresponding value attached to them. This is exactly the data structure we want.

File: `code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Inspire/NodeType.chs`

#### 4.3.2. Node Path

```
| type NodePath = [Int]
```

A *node path* is a list of child indices describing a path along the tree starting at some node. It is commonly used to reference a node in [INSPIRE](#). Each index tells you which child to go to next. You have reached the target upon depleting the list.

File: `code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Inspire/NodeAddress.hs`

#### 4.3.3. Binary Parser

The binary parser of the `Inspire` module builds upon `Attoparsec`<sup>1</sup>, which is a library for constructing combinator parsers. Together with the `attoparsec-binary`<sup>2</sup> package it becomes quite easy to build a binary parser for the [INSPIRE](#) export. The binary dump is mainly composed of three differently

---

<sup>1</sup><https://hackage.haskell.org/package/attoparsec>

<sup>2</sup><https://hackage.haskell.org/package/attoparsec-binary>

sized integer types and strings. We therefore define ourselves the following parser combinators.

```

1 | anyInt8 :: Parser Int
2 | anyInt8 = fromIntegral <$> anyWord8
3
4 | anyInt16 :: Parser Int
5 | anyInt16 = fromIntegral <$> anyWord16le
6
7 | anyInt32 :: Parser Int
8 | anyInt32 = fromIntegral <$> anyWord32le
9
10 | parseString :: Parser String
11 | parseString = liftM BS8.unpack $ take =<< anyInt32

```

Integers larger than one byte are dumped in little endian<sup>1</sup>, while strings are always prefixed with their length and do not contain a terminator.

The header of the binary dump starts with the magic number `0x494e5350495245` and a list of *converters*. Like strings, lists are prefixed with the number of elements. Converters are only relevant when using annotations attached to the **INSPIRE** program and since we do not use them, they are discarded by the main parser. The header is parsed with this set of combinators.

```

1 | parseHeader :: Parser [String]
2 | parseHeader = parseMagicNr *> parseConverters
3
4 | parseMagicNr :: Parser Word64
5 | parseMagicNr = word64le 0x494e5350495245
6 |                 <|> fail "wrong magic number"
7
8 | parseConverter :: Parser String
9 | parseConverter = parseString
10
11 | parseConverters :: Parser [String]
12 | parseConverters = do
13 |     n <- anyInt32
14 |     count n parseConverter

```

The main part of the binary dump – the encoding of the **IR** nodes – is a list of nodes, where each node starts with an integer encoding its type. If the node is a value node, the value follows, otherwise a list of integers follows telling us the node's children. The integers are list indices referring to other nodes in the same list. Each node can be followed by a list of annotations which is discarded by our parser. We first parse the whole list of nodes from the dump to create an integer map of, so called, `DumpNodes`, where the key is the corresponding index in the list. A `DumpNode` is a node where child nodes are referenced through indices, instead of actual nodes. The following code parses a single

---

<sup>1</sup><https://en.wikipedia.org/wiki/Endianness>

#### 4. Framework Implementation

DumpNode.

```
1 data DumpNode = DumpNode IR.NodeType [Int]
2
3 parseDumpNode :: Parser DumpNode
4 parseDumpNode = do
5     -- get node type
6     t <- toEnum <$> anyInt16
7
8     -- create corresponding node
9     n <- case t of
10         -- value nodes
11         IR.NT_BoolValue   -> DumpNode <$> IR.BoolValue <$> (/=0) <$> anyInt8   <*> pure []
12         IR.NT_CharValue   -> DumpNode <$> IR.CharValue <$> chr <$> anyInt8   <*> pure []
13         IR.NT_IntValue    -> DumpNode <$> IR.IntValue <$> anyInt32   <*> pure []
14         IR.NT_UIntValue   -> DumpNode <$> IR.UIntValue <$> anyInt32   <*> pure []
15         IR.NT_StringValue -> DumpNode <$> IR.StringValue <$> parseString <*> pure []
16
17         -- other nodes
18         _                 -> do
19             c <- anyInt32
20             is <- count c anyInt32
21             return $ DumpNode (IR.fromNodeType t) (fromIntegral <$> is)
22
23     -- skip annotations
24     a <- anyInt32
25     count a anyWord64le
26
27     return $ n
```

After this phase is completed, the DumpNodes are gradually converted into a single rose tree. The first node, the one with index zero, marks the root of the tree. Nodes may be referenced by the indices multiple times, yet it is important to construct a node only once and reuse already constructed ones. Otherwise we end up with multiple instances of the same node in memory, wasting not only a big amount of space, but also a lot of time when checking for equality later on. This concept is referred to as *node sharing* and, while Haskell does not offer us a way to handle references explicitly, we can deal with them by carefully using the constructor and already defined nodes. Each constructed node is temporarily stored using the State monad and a map during the construction of the rose tree.

```
1 connectDumpNodes :: IntMap.IntMap DumpNode -> IR.Tree
2 connectDumpNodes dumpNodes = evalState (go 0) IntMap.empty
3 where
4     go :: Int -> State (IntMap.IntMap IR.Tree) IR.Tree
5     go index = do
6         memo <- get
7         case IntMap.lookup index memo of
8             Just n -> return n
9             Nothing -> do
10                 let (DumpNode irnode is) = dumpNodes IntMap.! index
11                     children <- mapM go is
12                     let n = IR.mkNode index irnode children []
13                         modify (IntMap.insert index n)
14                 return n
```

Now that we have the rose tree representing the received **INSPIRE** program we still need to do one slight modification. We go over each node one more time and attach information about operators (ie

### 4.3. IR Components

language extensions). This step is done by the function `markBuiltins` in line 15 (next code snippet). As already mentioned, after the binary dump of the *INSPIRE* program, a list of node paths addressing operators is provided by the binary dump.

We now have all the parts to construct the complete binary parser. The function `parseBinaryDump` combines what we have established in this section. It takes the whole binary dump as input and returns either the constructed rose tree or an error message.

```
1 | parseBinaryDump :: BS.ByteString -> Either String IR.Tree
2 | parseBinaryDump = parseOnly $ do
3 |     -- parse components
4 |     parseHeader
5 |     n         <- anyInt32
6 |     dumpNodes <- IntMap.fromList <$> zip [0..] <$> count n parseDumpNode
7 |     m         <- anyInt32
8 |     dumpBuiltins <- Map.fromList <$> count m parseBuiltin
9 |     l         <- anyInt32
10 |    paths      <- count l parseNodePath
11 |
12 |    -- connect components
13 |    let root    = connectDumpNodes dumpNodes
14 |        builtins = resolve root <$> dumpBuiltins
15 |        ir      = markBuiltins root builtins
16 |
17 |    return ir
18 |
19 | where
20 |     resolve :: IR.Tree -> [Int] -> IR.Tree
21 |     resolve node [] = node
22 |     resolve node (x:xs) = resolve (IR.getChildren node !! x) xs
```

File: `code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Inspire/BinaryParser.hs`

Let us now take a final look at the adapter to see how this parser is used.

```
1 | foreign export ccall "hat_initialize_context"
2 | initializeContext :: Ctx.CContext -> CString -> CSize -> IO (StablePtr Ctx.Context)
3 |
4 | initializeContext context_c dump_c size_c = do
5 |     dump <- BS8.packCStringLen (dump_c, fromIntegral size_c)
6 |     let Right ir = BinPar.parseBinaryDump dump
7 |         newStablePtr $ Ctx.mkContext context_c ir
```

Upon setting up a context, the Haskell side receives a reference to the corresponding C++ object (`Ctx.CContext`) and the binary dump in form of a byte array (plus its length). This array is *packed* – necessary so the parser can work with it – and then parsed. The result is our rose tree. This allows us to setup an analysis context in Haskell.

File: `code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Adapter.hs`

## 4. Framework Implementation

### 4.3.4. Rose Tree – HAT’s INSPIRE Representation

We start this subsection off with a look at the definition of HAT’s IR tree.

```
1 | data Tree = Tree { getID      :: Int
2 |                   , getNodeType :: NodeType
3 |                   , getChildren :: [Tree]
4 |                   , builtinTags :: [String]
5 |                   }
6 |
7 | instance Eq Tree where
8 |     (==) = (==) `on` getID
9 |
10 | instance Ord Tree where
11 |     compare = compare `on` getID
```

Each instance of this data structure represents one node and we refer to the whole tree by the root node. Furthermore, each node is uniquely identified by an id which eases the definition of ordering and equality (lines 7–11). From the binary parser we simply use the list index of a node as id. We can see each node has a node type, which in case of being a value node also contains the value. Next we have the list of children, followed by a List of tags marking subtrees as being operators.

This data structure corresponds to Insieme’s `NodePtr` in C++.

File: `code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Inspire.hs`

### 4.3.5. Node Address

A node address is basically defined as a node path plus a root node. Since we are working a lot with node addresses during an analysis we attach more information to these instances for faster access to relevant data. The definition follows.

```
1 | data NodeAddress = NodeAddress { getPathReversed  :: NodePath
2 |                               , getNode          :: IR.Tree
3 |                               , getParent         :: Maybe NodeAddress
4 |                               , getRoot          :: IR.Tree
5 |                               , getAbsoluteRootPath :: NodePath
6 |                               }
```

The first thing to notice is that the node path is actually stored in reverse. The main reason for this is to speed up the common case of creating node addresses of children, given their parent.

Next we have three references, one to the addressed node, an optional one pointing to the parent node address (`Nothing` in case of being the root node address) and a reference to the root.

The *absolute node path* was introduced some time later during development. Together with the functions `append` and `crop` one can create node addresses with roots different than the one of the original tree.

For transferring node addresses between C++ and Haskell the following two functions are exposed



by the adapter.

```

1 | foreign export ccall "hat_mk_node_address"
2 |   mkNodeAddress :: StablePtr Ctx.Context -> Ptr CSize -> CSize
3 |                 -> IO (StablePtr Addr.NodeAddress)
4 |
5 | mkNodeAddress ctx_hs path_c length_c = do
6 |   ctx <- deRefStablePtr ctx_hs
7 |   path <- peekArray (fromIntegral length_c) path_c
8 |   newStablePtr $ Addr.mkNodeAddress (fromIntegral <$> path) (Ctx.getTree ctx)
9 |
10 | foreign export ccall "hat_node_path_poke"
11 |   nodePathPoke :: StablePtr Addr.NodeAddress -> Ptr CSize -> IO ()
12 |
13 | nodePathPoke addr_hs path_c = do
14 |   addr <- deRefStablePtr addr_hs
15 |   pokeArray path_c $ fromIntegral <$> Addr.getAbsolutePath addr

```

`mkNodeAddress` receives the node path as an array, its length, and a reference to the current context. From the context the root node is taken and used for constructing the node address together with the node path. A stable pointer to the newly constructed node address is returned. `nodePathPoke` takes such a stable pointer, dereferences it and places the indices of the node path in the provided array.

This data structure corresponds to Insieme's `NodeAddress` in C++.

File: `code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Inspire/NodeAddress.hs`

### 4.3.6. Visitors

Since we have defined our own tree structure for **INSPIRE** we also have to provide some basic visitor functionality. For performance critical reasons one can stop the visiting process at a controlled depth. This is referred to as *pruning*. The following visitors are provided.

`foldTree` Folds the whole tree. No pruning takes place. (fold is the Haskell equivalent of a reduction)

`foldTreePrune` Same as `foldTree`, but allows you to prune at any point.

`foldAddress` Folds the subtree starting at a given address. No pruning takes place.

`foldAddressPrune` Same as `foldAddress`, but allows you to prune at any point.

`collectAll` Returns a list of node addresses to all nodes satisfying a given predicate of the `NodeType`. No pruning takes place.

`collectAllPrune` Same as `collectAll`, but allows you to prune at any point.

`foldTree` and `foldTreePrune` operate on nodes while `foldAddress` and `foldAddressPrune` operate on node addresses.

File: `code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Inspire/Visit.hs`

## 4. Framework Implementation

### 4.3.7. Node References and Queries

To facilitate using the tree IR structure and node addresses a list of useful queries are implemented. These include, for example, checking if the given node is an operator or getting its type. The type-class `NodeReference` has been defined in order to allow nodes of the tree and node addresses to use the same queries, hence both of them implement this type class. The list of queries can be extended later on. Note that there exist a few<sup>1</sup> queries which can only be used with node addresses. This is due to the fact that a node address allows you to look at the parent of the referred node, which is not possible by using only plain nodes.

Furthermore a function `findDecl` is provided, which takes a node address pointing to a variable and tries to find its declaration.

File: `code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Inspire/Query.hs`

### 4.3.8. Context

Up until now we have already heard that a context is established in both domains and keeps track of the solver's state. We now take a look at the context's definition in C++ and how it is setup and used.

---

<sup>1</sup>Currently `isLoopIterator`, `hasEnclosingStatement`, `isEntryPoint` and `isEntryPointParameter`.

The important parts are displayed below.

```

1 | typedef void* StablePtr;
2 | typedef void* HaskellNodeAddress;
3 |
4 | class Context {
5 |
6 |     StablePtr context_hs;
7 |
8 |     core::NodePtr root;
9 |
10 |     std::map<core::NodeAddress, HaskellNodeAddress> addresses;
11 |
12 | public:
13 |
14 |     Context();
15 |     explicit Context(const core::NodePtr& root);
16 |     ~Context();
17 |
18 |     // move semantics
19 |     Context(const Context& other) = delete;
20 |     Context(Context&& other) = default;
21 |
22 |     Context& operator=(const Context& other) = delete;
23 |     Context& operator=(Context&& other) = default;
24 |
25 |     // -- haskell engine specific requirements --
26 |
27 |     StablePtr getHaskellContext() const;
28 |     void setHaskellContext(StablePtr);
29 |
30 |     void setRoot(const core::NodePtr&);
31 |     core::NodePtr getRoot() const;
32 |
33 |     HaskellNodeAddress resolveNodeAddress(const core::NodeAddress& addr);
34 |     core::NodeAddress resolveNodeAddress(const HaskellNodeAddress& addr);
35 |
36 | };

```

The context holds a stable pointer to its counter part in the Haskell domain and a specific **INSPIRE** program is associated with it. Furthermore it caches already transferred node addresses. Since the context in C++ is responsible for freeing (ie *un-stabling* Haskell objects) it has a custom destructor and no copy semantics. Node addresses can be resolved in both directions using `resolveNodeAddress`. To use it, just create an instance of the class, passing the root of an **INSPIRE** program to the constructor. Functions initiating an analysis take the context as their first argument, which is then used to convert **IR** information. This allows us to isolate the **FFI** code from the rest of Insieme. Two function declarations, providing the user-interface to HAT-based analyses, are given as examples. Their signatures do not contain any Haskell related parts.

```

1 | bool maybeTrue(Context& ctxt, const ExpressionAddress& expr);
2 |
3 | ArithmeticSet getArithmeticValue(Context& ctxt, const ExpressionAddress& expr);

```

First the context needs to be established as described. Afterwards both functions can be used directly. What exactly these two functions do is covered later in the Base Analyses part (Sections 4.10.2 and 4.10.3).

## 4. Framework Implementation

File: code/analysis/include/insieme/analysis/cba/haskell/context.h

### 4.4. Utilities

Before continuing with investigating the actual framework we take a look at some utilities used by the framework. The two most notable are bound sets and arithmetic formulas.

#### 4.4.1. Bound Sets

A *bound set* is similar to a regular set, apart from the fact that it is associated with an upper bound. When the number of elements within the set exceeds this upper bound, the set turns into *universe* representing every possible value of the domain. Of course, using this set as property spaces for program analyses likely yields over approximations as soon as the set turns into *universe*. Yet it is necessary to maintain termination of the analyses (thus decidability). Note that the bound can be adjusted even during an analysis run to maximise flexibility, but adjusting it changes the type of the bound set. Its definition:

```
| data BoundSet bb a = BoundSet (Data.Set.Set a) | Universe
```

The functions provided are quite similar to the ones from `Data.Set`, as can be seen by the export

list.

```

1 | module Insieme.Utils.BoundsSet (
2 |   BoundsSet(Universe),
3 |   empty, singleton,
4 |   size, null, isUniverse, member,
5 |   fromSet, fromList, toSet, toList,
6 |   insert, applyOrDefault, filter, map, lift2,
7 |   union, intersection, cartProduct,
8 |   -- ...
9 | ) where

```

Bounds are defined in the following way.

```

1 | class Typeable b => IsBound b where
2 |   bound :: p b a -> Int
3 |
4 | getBound :: IsBound bb => BoundsSet bb a -> Int
5 | getBound bs = bound bs
6 |
7 | changeBound :: (IsBound bb, IsBound bb', Ord a)
8 |              => BoundsSet bb a -> BoundsSet bb' a
9 | changeBound Universe = Universe
10 | changeBound (BoundsSet s) = fromSet s
11 |
12 | data Bound10 = Bound10
13 | instance IsBound Bound10 where
14 |   bound _ = 10
15 |
16 | data Bound100 = Bound100
17 | instance IsBound Bound100 where
18 |   bound _ = 100

```

The type-class requires the implementation of `bound` which takes a *bounded object* as input and returns the *capacity* associated with this bound. Two bounds, `Bound10` and `Bound100`, are already defined. More bounds can be added as needed. Having the bound encoded into the type of the bounded container prevents accidental combining of differently bound containers. A function `changeBound` can be used to explicitly convert between them.

Although we took `Data.Set` as a reference, bound sets are not interoperable with it. Yet, sometimes an *unbound* set is required (or helpful) and it would be convenient to use the same interface. We therefore define ourselves the `Unbound` bound.

```

1 | data Unbound = Unbound
2 |
3 | instance IsBound Unbound where
4 |   bound _ = -1
5 |
6 | type UnboundSet a = BoundsSet Unbound a
7 |
8 | toUnboundSet :: (IsBound bb, Ord a) => BoundsSet bb a -> UnboundSet a
9 | toUnboundSet = changeBound

```

Some analyses yield a result in the form of a bound set (eg arithmetic value). Therefore a similar data structure exists at the C++ side. With it, results wrapped in a bound set can be transferred across the two domains as long as the element type is transferable.

## 4. Framework Implementation

Along with this, convenience functions are provided, among them `map`, `lift2`, `cartProduct`, and `applyOrElseDefault`.

File: `code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Utils/BoundSet.hs`

### 4.4.2. Arithmetic Formulas

The purpose of the Arithmetic module is to model arithmetic formulas. For now only basic integer operations are implemented. On both sides – C++ and Haskell – a data structure is defined to model these formulas. Here is the definition of the Haskell part.

```
1  -- v...variable type
2  -- c...constant type
3
4  -- Something like x2.
5  data Factor c v = Factor { base    :: v
6                        , exponent :: c }
7
8  -- Something like x2y2.
9  data Product c v = Product { factors :: [Factor c v] }
10
11 -- Somethign like -2x2y3.
12 data Term c v = Term { coeff  :: c
13                    , product :: Product c v }
14
15 -- Something like -2x2y3 + 5z3 + 2.
16 data Formula c v = Formula { terms  :: [Term c v] }
```

The types of variables `v` and coefficient / exponent `c` are generic. Node addresses are commonly used as variables to reference a variable (or expression) in the [INSPIRE](#) program. Using integer types from `Foreign.C.Types` instead of `Integer` allows more accurate modelling of overflow behaviour.

Two formulas can be added, subtracted and multiplied. Division by a constant is supported, but may only be accurate if the remainder is zero. For now this model seemed adequate, but can be replaced in the future by one more accurate.

In addition to the representation of arithmetic formulas, we bring a different definition of *numeric ordering* along.

```
1  data NumOrdering = NumLT | NumEQ | NumGT | Sometimes
2
3  class NumOrd a where
4      numCompare :: a -> a -> NumOrdering
5
6  fromOrdering :: Ordering -> NumOrdering
7  fromOrdering LT = NumLT
8  fromOrdering EQ = NumEQ
9  fromOrdering GT = NumGT
```

A formula can either always be greater than, less than, or equal to another formula; or the ordering depends on the value of variables (`Sometimes`). These four different outcomes are modelled here.

For transferring an arithmetic formula from Haskell to C++ we have to reconstruct the formula part by

part. For this purpose the corresponding constructors are wrapped and made available to the Adapter via the FFI as seen below.

```

1 | foreign import ccall "hat_mk_arithmetic_formula"
2 |   arithmeticFormula :: Ptr (Ptr CArithmeticTerm) -> CSize -> IO (Ptr CArithmeticFormula)
3 |
4 | foreign import ccall "hat_mk_arithmetic_product"
5 |   arithmeticProduct :: Ptr (Ptr CArithmeticFactor) -> CSize -> IO (Ptr CArithmeticProduct)
6 |
7 | foreign import ccall "hat_mk_arithmetic_term"
8 |   arithmeticTerm :: Ptr CArithmeticProduct -> CULong -> IO (Ptr CArithmeticTerm)
9 |
10 | foreign import ccall "hat_mk_arithmetic_factor"
11 |   arithmeticFactor :: Ptr CArithmeticValue -> CInt -> IO (Ptr CArithmeticFactor)
12 |
13 | foreign import ccall "hat_mk_arithmetic_value"
14 |   arithmeticValue :: Ctx.CContext -> Ptr CSize -> CSize -> IO (Ptr CArithmeticValue)

```

The following function does all the heavy lifting on the Haskell side.

```

1 | passFormula :: Integral c
2 |             => Ctx.CContext
3 |             -> Ar.Formula c Addr.NodeAddress
4 |             -> IO (Ptr CArithmeticFormula)
5 |
6 | passFormula ctx_c formula_hs = do
7 |   terms_c <- forM (Ar.terms formula_hs) passTerm
8 |   withArrayLen' terms_c arithmeticFormula
9 |   where
10 |    passTerm :: Integral c => Ar.Term c Addr.NodeAddress -> IO (Ptr CArithmeticTerm)
11 |    passTerm term_hs = do
12 |      product_c <- passProduct (Ar.product term_hs)
13 |      arithmeticTerm product_c (fromIntegral $ Ar.coeff term_hs)
14 |
15 |    passProduct :: Integral c => Ar.Product c Addr.NodeAddress -> IO (Ptr CArithmeticProduct)
16 |    passProduct product_hs = do
17 |      factors_c <- forM (Ar.factors product_hs) passFactor
18 |      withArrayLen' factors_c arithmeticProduct
19 |
20 |    passFactor :: Integral c => Ar.Factor c Addr.NodeAddress -> IO (Ptr CArithmeticFactor)
21 |    passFactor factor_hs = do
22 |      value_c <- passValue (Ar.base factor_hs)
23 |      arithmeticFactor value_c (fromIntegral $ Ar.exponent factor_hs)
24 |
25 |    passValue :: Addr.NodeAddress -> IO (Ptr CArithmeticValue)
26 |    passValue addr_hs = withArrayLen' (fromIntegral <$> Addr.getAbsolutePath addr_hs)
27 |                               (arithmeticValue ctx_c)
28 |
29 |    withArrayLen' :: Storable a => [a] -> (Ptr a -> CSize -> IO b) -> IO b
30 |    withArrayLen' xs f = withArrayLen xs (\s a -> f a (fromIntegral s))

```

The formula is reconstructed from the bottom up, where value here refers to node addresses and are the variables of the formula.

File: code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Utils/Arithmetic.hs

## 4. Framework Implementation

### 4.5. Fixpoint Solver

This section describes the solver itself together with the `Lattice` type-class and used data structures for Variables, Assignments, and Constraints.

#### 4.5.1. Lattice

For implementation, two different names for the (complete) lattice and join-semi-lattice are used – `ExtLattice` and `Lattice` respectively.

```
1 | class (Eq v, Show v, Typeable v, NFData v) => Lattice v where
2 |     join :: [v] -> v
3 |     join [] = bot
4 |     join xs = foldr1 merge xs
5 |
6 |     merge :: v -> v -> v
7 |     merge a b = join [a,b]
8 |
9 |     bot :: v
10 |     bot = join []
11 |
12 |     less :: v -> v -> Bool
13 |     less a b = (a `merge` b) == b
14 |
15 | class (Lattice v) => ExtLattice v where
16 |     top :: v
```

`join` describes how a list of values is combined into a single one. `merge` is just the binary version of `join` as can be seen by the default implementation. `less` determines whether one element of the lattice is less<sup>1</sup> than another element. This type-class must be instantiated by the property space of each analyses.

---

<sup>1</sup>  $a < b \iff a$  is a more accurate value for an analyses than  $b$



## 4.5.2. Variables

Analysis variables are described briefly in Section 2.5.2. An analysis variable is associated with a list of constraints restricting potential Assignments and a unique identifier. The definition follows.

```

1  data Identifier = Identifier { analysis :: AnalysisIdentifier
2                                , idValue  :: IdentifierValue
3                                , idHash   :: Int  }
4
5  data AnalysisIdentifier = AnalysisIdentifier { aidToken :: TypeRep
6                                                  , aidName  :: String
7                                                  , aidHash   :: Int  }
8
9  data IdentifierValue = IDV_Expression NodeAddress
10                       | IDV_ProgramPoint ProgramPoint
11                       | IDV_MemoryStatePoint MemoryStatePoint
12                       | IDV_Other BS.ByteString
13
14  deriving (Eq,Ord,Show)
15
16  data Var = Var { index  :: Identifier
17                 , constraints :: [Constraint]
18                 , bottom  :: Dynamic
19                 , valuePrint :: Assignment -> String }
20
21  instance Eq Var where
22      (==) a b = (index a) == (index b)
23
24  instance Ord Var where
25      compare a b = compare (index a) (index b)
26
27
28  newtype TypedVar a = TypedVar Var
29                      deriving (Show,Eq,Ord)
30
31  mkVariable :: (Lattice a) => Identifier -> [Constraint] -> a -> TypedVar a

```

Note that most of the complexity of these definitions comes from the necessary performance optimizations – in particular for Variable comparison.

The identifier used for analysis variables contains an identifier for the corresponding type of analysis (AnalysisIdentifier) and a *value* (NodeAddress, ProgramPoint, MemoryStatePoint, or byte string) followed by a hash. The Identifier’s hash is used to speed up equality checks and takes the hash of AnalysisIdentifier into account.

The bottom value of an analysis variable depends on the analysis’ property space. We therefore use Dynamic as type. We could use a generic type variable for this, but this would make it difficult to directly work with and combine variables of different analyses. See the related Haskell Wiki page for more information about using heterogeneous collections<sup>1</sup>.

The structure TypedVar is used to statically connect a variable with its related property space – where possible. This is used in the Assignment. A Var is created and wrapped by mkVariable given an Identifier, a list of Constraints, and the bottom value of the corresponding property space. Only TypedVars may be created.

<sup>1</sup>[https://wiki.haskell.org/Heterogenous\\_collections](https://wiki.haskell.org/Heterogenous_collections)

## 4. Framework Implementation

### 4.5.3. Assignments

An Assignment is defined as a mapping from analysis variables to values of the related property space. The interface definition is given below.

```
1 | newtype Assignment = Assignment -- ...
2 |
3 | get :: (Typeable a) => Assignment -> TypedVar a -> a
4 |
5 | set :: (Typeable a, NFData a) => Assignment -> TypedVar a -> a -> Assignment
```

get returns the value associated to the given TypedVar in the Assignment. The use of TypedVar allows us to return the correct type of the value. set modifies an Assignment by associating a given TypedVar with the given value.

### 4.5.4. Constraints

```
1 | data Event = None
2 |             | Increment
3 |             | Reset
4 |
5 | data Constraint = Constraint
6 |   { dependingOn      :: Assignment -> [Var]
7 |     , update         :: (Assignment -> (Assignment,Event))
8 |     , updateWithoutReset :: (Assignment -> (Assignment,Event))
9 |     , target         :: Var }
```

A Constraint is always associated with the one variable which's value it is constraining, here referred to as target. Depending on an Assignment a list of variables can be deduced on which the Constraint depends on.

A Constraint provides us with the two functions update and updateWithoutReset which mutate a given Assignment. The update function also gives us information about what happened in the form of an Event. The three possible outcomes are.

None The Assignment did not change, we can continue without any additional work.

Increment The Assignment changed during the update monotonously, all depending variables need to be reconsidered.

Reset The current Assignment changed in a non-monotonous way, hence a *local restart* needs to be triggered.

The function updateWithoutReset is used to prevent local restarts from triggering each other by forcing monotonous updates. How exactly these three events are handled can be seen in Section 4.5.5.

### Dependencies between Variables

It is essential that the dependencies between variables is correctly established for all analyses. The framework's functions expect this to be done correctly by the analysis designer. This means that for a given Assignment, all analysis variables used to calculate the value of the current analysis variable

must be present in the dependency list of the current variable. The correct use is illustrated in the following code snippet, given two analysis variables `var1` and `var2`, and an `Identifier` for the new analysis variable (lines 1–3).

```

1 | var1 :: Solver.TypedVar Int
2 | var2 :: Solver.TypedVar Int
3 | varId :: Identifier
4 |
5 | con :: Constraint
6 | con = createConstraint dep val var
7 |
8 | var :: Solver.TypedVar Int
9 | var = mkVariable varId [con] Solver.bot
10 |
11 | dep :: Assignment -> [Var]
12 | dep _ = [Solver.toVar var1, Solver.toVar var2]
13 |
14 | val :: Assignment -> Int
15 | val a = Solver.get a var1 + Solver.get a var2

```

In lines 6 and 9, a constraint and analysis variable are generated. Since the value (`val`) of the constraint depends on the values yielded by `var1` and `var2`, the new analysis variable `var` depends on `var1` and `var2`. This is stated in lines 6 and 12 where `dep` contains all analysis variable `var` is depending on, which is used to create the constraint.

### 4.5.5. Solver

Before talking about the solving process itself, we model the state of the solver. `SolverState`, as seen below, captures the current `Assignment` and all processed analysis variables (and thus their constraints) using the already introduced `VariableIndex` structure (Section 4.5.2).

```

1 | SolverState = SolverState { assignment    :: Assignment
2 |                               , variableIndex :: VariableIndex }

```

The `solve` function takes an initial `SolverState` and a list of desired analysis variables as input, yielding a new `SolverState` containing the modified `Assignment`, containing the least fixpoint solution for the passed analysis variables.

```

| solve :: SolverState -> [Var] -> SolverState

```

Behind the scene `solve` calls `solveStep` which contains the main solver logic.

```

| solveStep :: SolverState -> Dependencies -> [IndexedVar] -> SolverState

```

`solveStep` takes as input a `SolverState`, the `Dependencies` between variables and a *worklist* in form of a list of `IndexedVar`. The result is a new `SolverState`. The function takes one `IndexedVar` after another from the worklist (third parameter) and processes the constraints associated with this variable.

To process a constraint, the update function of this constraint is called with the current assignment, yielding a new assignment and an `Event` (Section 4.5.4). In case the `Event` is `None`, the processing simply continues since the `Assignment` did not change. If it is `Increment` the constraints depending on the target variable of the current constraint need to be re-evaluated and the related variables are therefore

#### 4. Framework Implementation

added to the worklist. The old Assignment has been updated by the constraint's update. On Reset the function behaves similarly as on Increment, however the transitive closure of all dependent variables is reset to their respective bottom values – unless the current target value is indirectly depending on itself.

At each step, newly generated analysis variables are added to the worklist whenever encountered. The VariableIndex, as well as the Dependencies, are updated accordingly.

File: code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Analysis/Solver.hs

### 4.6. Modelling Arrays, Structs, and Unions

This section describes how arrays, structs, and unions are modelled in the current state of the framework. This is a bit tricky since in C/C++ (and thus in [INSPIRE](#)) all three structures can be combined with each other, multiple times. For example having an array of structs or a struct embedded in a struct embedded in a union, and so on.

The two main building blocks we need are DataPath and FieldIndex. A FieldIndex is used to reference a field in case of a struct or union, or an element (index) in case of an array. DataPath allows us to reference elements of multi-level structures using a list of FieldIndex. Both of these building blocks are implemented in a generic way to be easily extensible in the future. Thus FieldIndex is a type-class for which SimpleFieldIndex is our main implementation.

```
1 | class (Eq v, Ord v, Show v, Typeable v, NFData v) => FieldIndex v where
2 |     -- ...
3 |
4 | data SimpleFieldIndex = Field String
5 |                       | Index Int
6 |                       | UnknownIndex
7 |
8 | data Symbol = Constant NodeAddress
9 |             | Variable NodeAddress
10 |
11 | instance FieldIndex SimpleFieldIndex where
12 |     -- ...
13 |
14 | type SymbolicFormula = Ar.Formula CInt Symbol
```

The data structure Symbol is used together with CInt to specialise the generic arithmetic formula data structure (see Section 4.10.2). The result, SymbolicFormula, is used when accessing an element of an array.

```
1 | data DataPath i = Root
2 |                 | Narrow [i]
3 |                 | Expand [i]
4 |                 | Invalid
```

DataPath is used to access the underlying structure of *structured* memory locations, where Root refers to the whole structure and Narrow to a path starting at the root pointing to a specific sub-structure. Expand is the inverse to Narrow and is used to navigate from a sub-structure to an enclosing structure [10, pp. 133–134].

We now combine these building blocks with the `ComposedValue` type-class.

```

1 | -- c...composed value type, for modelling the value of composed bojects
2 | --   like arrays, structs, ...
3 | -- i...field index type
4 | -- v...leaf value type to be covered
5 |
6 | class (Solver.ExtLattice c, FieldIndex i, Solver.ExtLattice v)
7 |     => ComposedValue c i v | c -> i v where
8 |
9 | toComposed :: v -> c
10 | toValue    :: c -> v
11 |
12 | setElement :: DataPath i -> c -> c -> c
13 | getElement :: DataPath i -> c -> c
14 |
15 | composeElements :: [(i,c)] -> c
16 |
17 | top :: c

```

Note the functional dependency<sup>1</sup> `c -> i v` stating that `i` and `v` are uniquely determined by `c`. Take a look at [9] for more information on the topic of functional dependencies. An instance is provided for our use in the framework.

```

1 | data ValueTree i a = Leaf a
2 |                   | Node (Data.Map.Map i (Tree i a))
3 |                   | Empty
4 |                   | Inconsistent
5 |
6 | instance (FieldIndex i, Solver.ExtLattice a) => ComposedValue (Tree i a) i a where
7 |     -- ...

```

Each level in this `ValueTree` data structure corresponds to one level in the structured memory. Let us look at an example.

<pre> 1   int a[2][3] = { 2       {1, 2, 3}, 3       {4, 5, 6} 4   }; </pre>	<pre> 1   let a = composeElements [ 2       (Index 0, composeElements [ (Index 0, toComposed 1), 3                                 (Index 1, toComposed 2), 4                                 (Index 2, toComposed 3) ]), 5       (Index 1, composeElements [ (Index 0, toComposed 4), 6                                 (Index 1, toComposed 5), 7                                 (Index 2, toComposed 6) ])] </pre>
--	--

On the left we see the original structure from an input program, on the right the corresponding `ValueTree` instance. Since we are representing a 2D array, the `ValueTree` instance also has two levels and `Index` is used for field indexing (see definition of the `SimpleFieldIndex` type in Section 4.6). One more example using a struct follows.

<pre> 1   struct vector3 { 2       int x; 3       int y; 4       int z; 5   }; 6   7   struct vector3 u = {7, 8, 9}; </pre>	<pre> 1   let u = composeElements [ (Field "x", toComposed 7), 2                             (Field "y", toComposed 8), 3                             (Field "z", toComposed 9) ] </pre>
---	--

<sup>1</sup>[https://wiki.haskell.org/Functional\\_dependencies](https://wiki.haskell.org/Functional_dependencies)

## 4. Framework Implementation

Since we are using a struct instead of `Index` we use `Field` together with the corresponding field name as string.

File: `code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Analysis/Entities/FieldIndex.hs`

File: `code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Analysis/Framework/PropertySpace/ComposedValue.hs`

File: `code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Analysis/Framework/PropertySpace/ValueTree.hs`

## 4.7. Generic Data-Flow Analysis

The generic **Data-Flow Analysis (DFA)** is implemented by a combination of data structures, more specific analyses, and by the main function `dataflowValue`. Some of the data structures and more specific analyses are covered in Section 4.5.5. The remaining analyses are covered afterwards and grouped under Section 4.10.

To easily manage the more specific analyses the `DataFlowAnalysis` data structure represents a specific instance of a **DFA** and is composed of the following fields.

```
1 | -- a...analysis type
2 | -- v...property space
3 |
4 | data DataFlowAnalysis a v = DataFlowAnalysis
5 |   { analysis           :: a,
6 |     , analysisIdentifier :: Solver.AnalysisIdentifier,
7 |     , variableGenerator :: NodeAddress -> Solver.TypedVar v,
8 |     , topValue          :: v,
9 |     , entryPointParameterHandler :: NodeAddress -> Solver.TypedVar v,
10 |    , initialValueHandler  :: NodeAddress -> v,
11 |    , initValueHandler    :: v }
```

The `analysis` field is equivalent to the `aidToken` field in `AnalysisIdentifier`. It serves the purpose of distinguishing different analyses by their type. Yet, we also have to maintain the `AnalysisIdentifier` itself.

The `variableGenerator` is used for hooking the variable / constraint generator of the specific analysis into the generic framework. It follows the top value of the corresponding property space. As a safe fallback, the top value is used by the **DFA** whenever something unsupported is encountered.

The `entryPointParameterHandler` is used for covering parameters of program entry points – commonly `argc`, `argv`.

`initialValueHandler` and `initValueHandler` specify the respective value for memory locations – either upon an initialisation or program start. The `initialValueHandler` is consulted for dynamic memory locations, while `initValueHandler` deals with static memory locations.

While the data constructor `DataFlowAnalysis` is exported, a simplified version is also provided.

```

1 | mkDataFlowAnalysis :: (Typeable a, Solver.ExtLattice v)
2 |                   => a
3 |                   -> String
4 |                   -> (NodeAddress -> Solver.TypedVar v)
5 |                   -> DataFlowAnalysis a v

```

In this function the `analysisIdentifier` is populated with an `AnalysisIdentifier` constructed from the given analysis token and string. The third parameter is used as the variable / constraint generator. The respective top value can be deduced from the type and is outputted as result by all handlers, thus `entryPointParameterHandler`, `initialValueHandler`, and `initValueHandler` are set to yield the property space's top value.

An instance of `DataFlowAnalysis` can then be used with the main function of generic DFA `dataflowValue`. Its signature looks as follows.

```

1 | dataflowValue :: (ComposedValue.ComposedValue a i v, Typeable d)
2 |              => NodeAddress
3 |              -> DataFlowAnalysis d a
4 |              -> [OperatorHandler a]
5 |              -> Solver.TypedVar a

```

The first argument specifies the node for which a variable, representing the dataflow value, needs to be computed. Next an instance of `DataFlowAnalysis` is required to indicate the specific nature of the analysis. A list of `OperatorHandlers` (see Section 4.7.1) follows. The result is, of course, an analysis variable corresponding to the specific analysis. Note that `dataflowValue` can be specialised to a variable / constraint provider by fixing the second and third parameter.

Figure 8 illustrates how dataflow values are aggregated for the expression  $1 + 2 \times 3$ . As can be seen, the aggregation happens bottom-up.

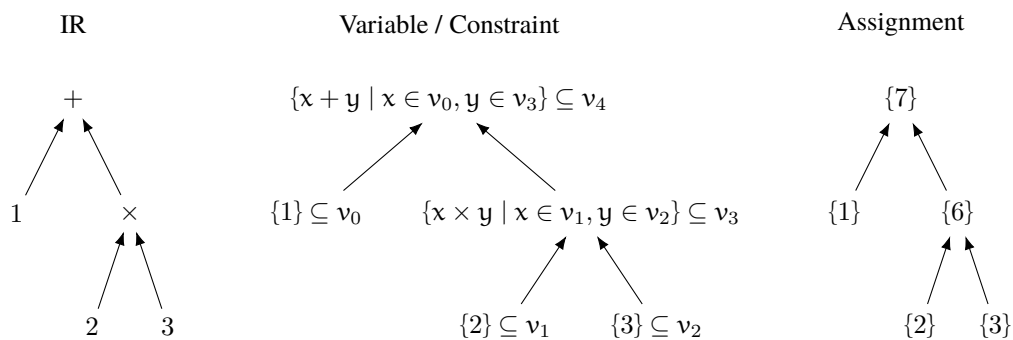


Figure 8.: Illustrating dataflow value aggregation of `dataflowValue`.

File: code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Analysis/Framework/Dataflow.hs

## 4. Framework Implementation

### 4.7.1. Operator Handler

Before continuing on, the concept of `OperatorHandler` is presented. They provide a simple means to hook more specific logic into the generic `DFA` (or `ProgramPoint` analysis as we will see later).

This is used to interpret the semantics of abstract operators (and optionally derived operators) (see Section 3.3). The definition is straight forward and looks like this.

```
1 | data OperatorHandler a = OperatorHandler { covers    :: NodeAddress -> Bool,
2 |                                           , dependsOn :: Assignment -> [Var],
3 |                                           , getValue  :: Assignment -> a }
```

`covers` is used by the generic `DFA` to check whether this operator handler is to be used for the interpretation of a call to a given function symbol / operator. The other two functions house the logic of the handler itself.

## 4.8. Program Point

In the next two subsections we focus on the two interconnected concepts of program points and memory state. Both concepts have their dedicated generic variable / constraint generator which are utilised by the generic `DFA`.

The definition of program points follows.

```
1 | data Phase = Pre | Internal | Post
2 |
3 | data ProgramPoint = ProgramPoint NodeAddress Phase
```

A `ProgramPoint` models a specific point in the execution of an `INSPIRE` program. For this purpose a `NodeAddress` referencing a specific node is used together with a `Phase`. The `Phase` describes at which phase of the execution the referenced node is referenced. It can be either

`Pre` before the referred node is processed;

`Internal` while the referred node is processed (only relevant for interpreting operators); or

`Post` after the referred node is processed.

### 4.8.1. Program Point Analysis

The purpose of the program point analysis is to derive properties from a given `ProgramPoint`. The generic variable / constraint generator function `programPointValue` of this analysis has the following



signature.

```

1 | programPointValue :: (Solver.ExtLattice a)
2 |                   => ProgramPoint
3 |                   -> (ProgramPoint -> Solver.Identifier)
4 |                   -> (ProgramPoint -> Solver.TypedVar a)
5 |                   -> [OperatorHandler a]
6 |                   -> Solver.TypedVar a

```

The result is a variable representing the requested information, for which the given `ProgramPoint` (first argument) represents a state value. The second argument is simply a variable id generator used to generate (unique) identifiers for analysis variables. The third argument is the analysis for which `programPointValue` is the base. The list of `OperatorHandlers` is used to intercept and interpret certain operators.

`programPointValue` utilises the predecessor analysis (covered in Section 4.11.3). The following code example and figure illustrate how the **INSPIRE** tree is traversed by `programPointValue`.

```

1 | // pseudo code
2 | {
3 |   a = 6;
4 |   b = 7;
5 |   f(a, b);
6 | }

```




Figure 9.: Illustrating **INSPIRE** backward traversal by `programPointValue`.

Each node in Figure 9 has two grey bullets, one to its left, the other to its right. They represent the `ProgramPoint`'s Phases, Pre or Post respectively. These points are traversed in backward direction by the program point analysis – thus in reverse execution order. This concept is the basic building block for modelling the backwards analyses in the analysis framework.

Like the data-flow function, this is a generic base for analyses – not an analysis itself.

**Data-Flow vs Program Point** The purpose of a data-flow analysis is to determine the value(s) a given expression may exhibit. In contrast to this, a program point analysis determines the state(s) of some (shared) structure at some point of program execution.

## 4.9. Memory State

A memory location is defined by a `NodeAddress` referencing its creation point. Furthermore, we can combine `MemoryLocation` with `ProgramPoint` to define ourselves `MemoryStatePoint`. A `MemoryStatePoint` references the state of a memory location at a given `ProgramPoint`.

```
1 | data MemoryLocation = MemoryLocation NodeAddress
2 |
3 | data MemoryStatePoint = MemoryStatePoint ProgramPoint MemoryLocation
```

In order to find the state of a `MemoryLocation` at a given `ProgramPoint`, a reaching definitions analysis has to be conducted.

### 4.9.1. Reaching Definitions Analysis

A reaching definitions analysis is capable of identifying all definitions of a given memory location reaching the `ProgramPoint` of interest. We therefore now look at the definition of the reaching definitions analysis' property space.

```
1 | data Definition = Initial
2 |                 | Creation
3 |                 | Declaration NodeAddress
4 |                 | MaterializingCall NodeAddress
5 |                 | Assignment NodeAddress
6 |                 | Initialization NodeAddress
7 |
8 | type Definitions = BSet.UnboundSet Definition
9 |
10 | instance Lattice Definitions where
11 |     bot = BSet.empty
12 |     merge = BSet.union
13 |
14 | instance ExtLattice Definitions where
15 |     top = BSet.Universe
```

Definition models the different ways a memory location in `INSPIRE` can be initialised.

**Initial** The definition is obtained at program startup.

**Creation** The variable is left undefined at its point of allocation.

**Declaration** The definition is conducted by an assignment triggered through a *materializing*<sup>a</sup> declaration. The node address refers to that Declaration.

**MaterializingCall** The definition is conducted by an assignment triggered through a materializing call. The node address refers to that CallExpr.

**Assignment** The definition is conducted by an assignment. The node address refers to that assignment.

**Initialization** The definition is conducted by an initialisation expression. The node address refers to that initialisation.

---

<sup>a</sup>A temporary (r-value) is written to a memory location so that it can be referenced – constructor and destructor are managed accordingly – see [https://clang.llvm.org/doxygen/classclang\\_1\\_1MaterializeTemporaryExpr.html#details](https://clang.llvm.org/doxygen/classclang_1_1MaterializeTemporaryExpr.html#details).

The signature of `reachingDefinitions` looks as follows:

```
| reachingDefinitions :: MemoryStatePoint -> Solver.TypedVar Definitions
```

This analysis is used by the memory state analysis.

## 4.9.2. Memory State Analysis

Finding the state of a `MemoryLocation` at a given `ProgramPoint` can now be achieved by the memory state analysis, which's variable / constraint generator builds on top of the reaching definitions analysis.

The analysis variable constructed by `reachingDefinitions`, for a given `MemoryState ms`, can be passed to the Solver's `get` function together with an `Assignment a` to yield the `Definitions`.

```
| reachingDefVar = reachingDefinitions ms
| reachingDefVal = Solver.get a reachingDefVar
```

This is almost the information we want. From the `Definitions` we need to deduce the actual value the memory location exhibits. This is done by the following step.

```
1 | definingValueVars a = BSet.applyOrElse [] (concat . (map go) . BSet.toList)
2 |                       $ reachingDefVal a
3 |   where
4 |     go (Declaration      addr) = [variableGenerator analysis $ goDown 1 addr]
5 |     go (MaterializingCall addr) = [variableGenerator analysis $ addr]
6 |     go (Assignment      addr) = [definedValue addr ml analysis]
7 |     go (Initialization  addr) = [definedValue addr ml analysis]
8 |     go _                 = []
```

Here, `definedValue` models the underlying memory structure of the given `MemoryLocation ml` using an instance of the `ComposedValue` type-class. Note that neither for `Initial` nor for `Creation` we can deduce a value more accurate than the worst-case  $\top$  element (line 8).

With this last piece in place we can define the memory state analysis' variable / constraint generator

## 4. Framework Implementation

function memoryStateValue.

```
1 memoryStateValue (MemoryStatePoint _ ml@(MemoryLocation loc)) analysis = var
2   where
3     var = Solver.mkVariable varId [con] Solver.bot
4     con = Solver.createConstraint dep val var
5
6     dep a = (Solver.toVar reachingDefVar) : (
7         if BSet.isUniverse defs || BSet.member Creation defs
8         then []
9         else (map Solver.toVar $ definingValueVars a)
10    )
11   where
12     defs = reachingDefVal a
13
14   val a = case () of
15     _ | BSet.isUniverse defs      -> Solver.top
16     | BSet.member Initial defs    -> Solver.merge init value
17     | otherwise                  -> value
18   where
19     init = initialValueHandler analysis loc
20
21     value = if BSet.member Creation $ defs
22             then ComposedValue.top
23             else Solver.join $ map (Solver.get a) (definingValueVars a)
24
25     defs = reachingDefVal a
26
27   varId = -- ..
```

The `initialValueHandler` is used in line 16 if the set of definitions `defs` contains an `Initial`. The `initialValueHandler`'s result is merged with deduced values. Constraints (and dependencies) are created accordingly.

### 4.10. Specific Analyses

Now, after covering the generic `DFA`, we look at more *specific analyses*. As already mentioned, these specialised analyses are used by the generic `DFA`. While the generic `DFA` covers the `INSPIRE` core language, the specific analyses cover individual language extensions. This relationship is illustrated in Figure 10.

As each language extension typically defines a new type and operators, each matching analysis defines an abstract value space for the values of expressions composed by the language extension. Furthermore, it models the semantics of the associated operators on top of these property spaces. Thereby we have the cases:

- Operators are abstract literals, hence an `OperatorHandler` has to be defined.
- Operators are derived functions, for which their interpretation is implicitly covered by interpreting their definitions – or an `OperatorHandler` may be defined, if desired.

Another important aspect to note is that the property spaces of analyses are not unique. There are infinitely many different ones, differing in their correctness, accuracy and evaluation performance.

The designer of an analyses has to chose the right one for a given task.

Also note that not all language extensions are clearly separated. In Insieme's current state, the *basic language extension* combines some frequently used language constructs like arithmetic and boolean.

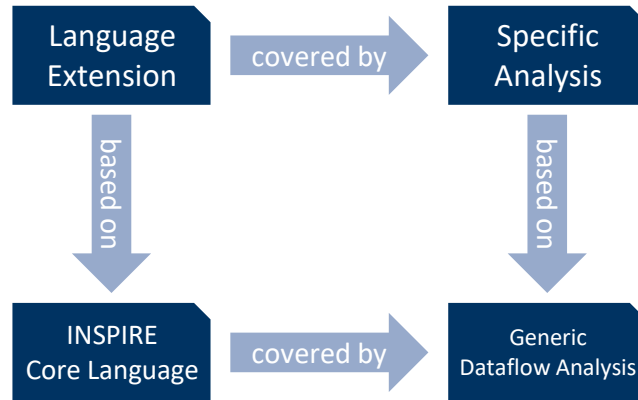


Figure 10.: Connection between language a extension, INSPIRE core language, a specific analysis and the generic DFA.

#### 4.10.1. Identifier Analysis

The purpose of the identifier analysis is to determine the value of expressions representing identifiers. Thus expressions of type `Identifier` (see Section 3.3.1).

This analysis is used by the data path analysis (see Section 4.10.4) to track the field names of structs and unions. Its property space is defined as follows.

```

1 | data Identifier = Identifier String
2 |   deriving (Eq,Ord,Generic,NFData)
3 |
4 | instance Show Identifier where
5 |   show (Identifier s) = s
6 |
7 | type IdentifierSet = BSet.UnboundSet Identifier
8 |
9 | instance Solver.Lattice IdentifierSet where
10 |   bot  = BSet.empty
11 |   merge = BSet.union
12 |
13 | instance Solver.ExtLattice IdentifierSet where
14 |   top  = BSet.Universe
  
```

The bottom value of the property space represents no identifiers, while the top element represents all identifiers – hence `BSet.empty` and `BSet.Universe` are used respectively. When two data-flows merge, the resulting data-flow should contain all identifiers from the two branches, therefore `BSet.union` is used as merge operator. Also note that the identifier analysis is a *may analysis*.

For the analysis to be usable by the framework, a unique identifier is required. Together with this

## 4. Framework Implementation

identifier we provide the constraint generator function `identifierValue`.

```
1 | data IdentifierAnalysis = IdentifierAnalysis
2 |   deriving (Typeable)
3 |
4 | identifierValue :: NodeAddress -> Solver.TypedVar (ValueTree SimpleFieldIndex IdentifierSet)
5 | identifierValue addr = case getNode addr of
6 |
7 |   IR.Node IR.Literal [_, IR.Node (IR.StringValue x) _] ->
8 |     Solver.mkVariable (idGen addr) [] (compose $ BSet.singleton (Identifier x))
9 |
10 |   _ -> dataflowValue addr analysis []
11 |
12 | where
13 |   analysis = mkDataFlowAnalysis IdentifierAnalysis "I" identifierValue
14 |   idGen     = mkVarIdentifier analysis
15 |   compose  = ComposedValue.toComposed
```

In line 13, a `DataFlowAnalysis` is instantiated using the unique identifier, a string "I" for easy identification of the associated analysis variables, and `identifierValue`, holding the main logic of the identifier analysis.

The interpretation of most `IR` constructs is handled by the generic `DFA`. We only need to cover the nodes relevant to us. In our case a `Literal`, where the second child is a `StringValue`<sup>1</sup>. An analysis variable is created, using `mkVariable` (line 8) and a newly generated identifier (line 14). No constraints are associated with the variable and the value associated with the variable is a (bound) set containing the identifier, wrapped up in a `ValueTree`.

One can also see that that no `OperatorHandlers` are used when invoking `dataflowValue` in line 10. This means that all information derived from `Literal` nodes is passed along the data-flow without any additional modifications.

**Connection Between Generic DFA and Specialised Analysis** In line 10 we can see the interaction between the generic `DFA` and a specialised analysis. The user would simply invoke `identifierValue` and pass the result to the solver. If the node targeted is not a `Literal`, the input program falls back to the generic `DFA`, which calls our specialised analysis function `identifierValue` for the relevant nodes. The results are merged according to the `Lattice` instantiation.

File: `code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Analysis/Identifier.hs`

### 4.10.2. Arithmetic Analysis

Section 4.4.2 has already presented the concept and implementation of arithmetic formulas. Their main use appears in the arithmetic analysis. Its task is to derive the arithmetic values a given expression may exhibit.

As mentioned in Section 3.3.1, arithmetic types and operators are defined in the *basic language extension*. The arithmetic analysis thus covers this part of the *basic language extension*.

---

<sup>1</sup>Note that this is an simplification that does not cause harm since *all* literals are covered, where only those of the type 'identifier' would be required.

The type alias `SymbolicFormula`, previously introduced in Section 4.6, is used together with `BoundSet` to compose the arithmetic property space.

```

1 | type SymbolicFormulaSet b = BSet.BoundSet b SymbolicFormula
2 |
3 | instance BSet.IsBound b => Solver.Lattice (SymbolicFormulaSet b) where
4 |     bot   = BSet.empty
5 |     merge = BSet.union
6 |
7 | instance BSet.IsBound b => Solver.ExtLattice (SymbolicFormulaSet b) where
8 |     top   = BSet.Universe

```

Similar to the identifiers analysis, the arithmetic analysis is a *may analysis*, therefore `BSet.empty`, `BSet.Universe`, and `BSet.union` are used for  $\perp$ ,  $\top$ , and the merge operator respectively.

An integer literal in the input program is translated to *INSPIRE* as a `Literal` containing a `StringValue` node. As the string contained needs to be parsed, a utility `parseInt` is provided. `parseInt` understands the different C/C++ integer literal formats (eg `0xf`, `072`, `1u1`, `-24`).

```

1 | data CInt = CInt32 Int32 -- ^ Represents @int@
2 |           | CInt64 Int64 -- ^ Represents @long@
3 |           | CUInt32 Word32 -- ^ Represents @unsigned int@
4 |           | CUInt64 Word64 -- ^ Represents @unsigned long@
5 |     deriving (Eq,Ord,Generic,NFData)
6 |
7 | instance Show CInt where -- ...
8 | instance Num CInt where -- ...
9 | instance Enum CInt where -- ...
10 | instance Real CInt where -- ...
11 | instance Integral CInt where -- ...
12 |
13 | parseInt :: String -> Maybe CInt

```

Before looking at the body of `arithmeticValue` we need to setup the variable / constraint generator implementing the arithmetic analysis, to be used by the solver.

```

1 | data ArithmeticAnalysis = ArithmeticAnalysis
2 |     deriving (Typeable)
3 |
4 | analysis = (mkDataFlowAnalysis ArithmeticAnalysis "A" arithmeticValue) {
5 |
6 |     initialValueHandler = \a -> compose $ BSet.singleton
7 |                                     $ Ar.mkVar
8 |                                     $ Constant (Addr.getNode a) a,
9 |
10 |     initValueHandler = compose $ BSet.singleton $ Ar.zero
11 | }
12 |
13 | idGen = mkVarIdentifier analysis

```

The default initial value handler yield the top element of the property space. For the arithmetic analysis this default is not desired, hence we override it in lines 6–10. For static memory locations (handled by `initValueHandler`) we can assume the memory location’s value to be initialised to 0. For dynamic memory locations we handle the location as a variable, which enables us to reason about it in more detail than just assuming its content to be  $\top$ . For instance, we can deduce that the value of the memory location subtracted from itself yields 0 as result.

#### 4. Framework Implementation

Now we are ready to look at `arithmeticValue`.

```
1 arithmeticValue :: Addr.NodeAddress
2   -> Solver.TypedVar (ValueTree SimpleFieldIndex
3                       (SymbolicFormulaSet BSet.Bound10))
4
5 arithmeticValue addr = case Addr.getNode addr of
6
7   IR.Node IR.Literal [t, IR.Node (IR.StringValue v) _] | isIntType t ->
8     case parseInt v of
9       Just cint -> Solver.mkVariable (idGen addr) [] (compose $ BSet.singleton
10                                     $ Ar.mkConst cint)
11
12       Nothing   -> Solver.mkVariable (idGen addr) [] (compose $ BSet.singleton
13                                     $ Ar.mkVar c)
14
15     where
16       c = Constant (Addr.getNode addr) addr
17
18   IR.Node IR.Variable _ | Addr.isLoopIterator addr ->
19     Solver.mkVariable (idGen addr) [] (compose $ BSet.Universe)
20
21   IR.Node IR.Variable (t:_) | isIntType t && isFreeVariable addr ->
22     Solver.mkVariable (idGen addr) [] (compose $ BSet.singleton
23                                     $ Ar.mkVar
24                                     $ Variable (Addr.getNode addr) addr)
25
26   IR.Node IR.CastExpr (t:_) | isIntType t -> var
27     where
28       var = Solver.mkVariable (idGen addr) [con] Solver.bot
29       con = Solver.forward (arithmeticValue $ Addr.goDown 1 addr) var
30
31   _ -> dataflowValue addr analysis ops
32
33   where
34     -- ...
```

In case the given node address points to a `Literal` of type integer (line 7), the integer parser utility is consulted (line 8). If the parse was successful, the result is wrapped in a (constant) arithmetic formula, a (bound) set, and a `ValueTree` (lines 9 and 10). Otherwise the given node is handled as a mathematical constant and wrapped similarly (lines 12–15).

If we encounter a variable, which is a loop-iterator (line 17), we immediately return `Universe` as the range is commonly too big to consider all possible values in the analysis. Free variables of type integer are handled explicitly by `arithmeticValue` (line 20). Such an explicitly handled variable is wrapped in an arithmetic formula, a (bound) set, and a `ValueTree`. Other variables are covered by the default dataflow analysis.

A `CastExpr` of (target) type integer simply forwards the `arithmeticValue` of its argument. For all other nodes, the generic `DFA` is used, hooking the defined operator handlers and this analysis into the framework. The `DFA` also covers declarations, parameter and return value passing, closures, calls to unknown external functions, inter-procedural analyses, and memory locations.

Since all arithmetic operators are abstract, `OperatorHandlers` need to be defined, however their implementation is straightforward. The encountered operator is simply mapped to the corresponding



arithmetic operation.

```

1 | ops = [add, mul, sub, div, mod]
2 |
3 | add = OperatorHandler cov dep (val Ar.addFormula)
4 |   where
5 |     cov a = any (isBuiltin a) ["int_add", "uint_add", "gen_add"]

```

The result of each integer operation depends on the arithmetic value of its operands. Thus we invoke `arithmeticValue` on the left- and right-hand side (lines 1 and 2).

```

1 | lhs = arithmeticValue $ Addr.goDown 2 addr
2 | rhs = arithmeticValue $ Addr.goDown 3 addr
3 |
4 | val op a = compose $ (BSet.lift2 op) (extract $ Solver.get a lhs)
5 |                                     (extract $ Solver.get a rhs)
6 |
7 | dep _ = Solver.toVar <$> [lhs, rhs]
8 |
9 | compose = ComposedValue.toComposed
10 | extract = ComposedValue.toValue

```

For division, we first need to check if it is safe to divide using the two operands (see Section 4.4.2). If possible, the division is applied, otherwise `Universe` is used as result.

```

1 | div = OperatorHandler cov dep val
2 |   where
3 |     cov a = any (isBuiltin a) [ "int_div", "uint_div", "gen_div" ]
4 |     val a = compose $ tryDiv (extract $ Solver.get a lhs) (extract $ Solver.get a rhs)
5 |
6 |     tryDiv x y = if not (BSet.isUniverse prod) && all (uncurry Ar.canDivide)
7 |                 (BSet.toList prod)
8 |                 then BSet.map (uncurry Ar.divFormula) prod
9 |                 else BSet.Universe
10 |   where
11 |     prod = BSet.cartProduct x y

```

We have already described how an arithmetic formula is passed from Haskell back to C++ in Section 4.4.2. We therefore now look at making the arithmetic analysis available to C++.

```

1 | foreign export ccall "hat_arithmetic_value"
2 |   arithValue :: StablePtr Ctx.Context
3 |               -> StablePtr Addr.NodeAddress
4 |               -> IO (Ptr CArithmeticSet)
5 |
6 | arithValue ctx_hs expr_hs = do
7 |   ctx <- deRefStablePtr ctx_hs
8 |   expr <- deRefStablePtr expr_hs
9 |   let (res,ns) = Solver.resolve (Ctx.getSolverState ctx) (Arith.arithmeticValue expr)
10 |      results = ComposedValue.toValue res
11 |      let ctx_c = Ctx.getCContext ctx
12 |          ctx_nhs <- newStablePtr $ ctx { Ctx.getSolverState = ns }
13 |          updateContext ctx_c ctx_nhs
14 |      passFormulaSet ctx_c $ BSet.map (fmap SymbolicFormula.getAddr) results

```

`arithValue`, exported as `hat_arithmetic_value`, receives a `Context` and an expression. The arithmetic

#### 4. Framework Implementation

analysis is invoked in line 9. The initial state of the solver is taken from the received Context. A new solver state `ns`, as well as the result of the analysis `res` is returned by `Solver.resolve`. The Context is updated with the new solver state in line 13, while the result is passed to C++ using `passFormulaSet`.

As the arithmetic analysis yields a (bound) set of arithmetic formulas, we pass this whole set via the FFI. Since we already know how a single formula is passed, we can easily build upon this knowledge and define ourselves `passFormulaSet`.

```
1 | foreign import ccall "hat_mk_arithmetic_set"
2 |   arithmeticSet :: Ptr (Ptr CArithmeticFormula) -> CInt -> IO (Ptr CArithmeticSet)
3 |
4 | passFormulaSet :: Integral c
5 |                 => Ctx.CContext
6 |                 -> BSet.BoundSet bb (Ar.Formula c Addr.NodeAddress)
7 |                 -> IO (Ptr CArithmeticSet)
8 |
9 | passFormulaSet _ BSet.Universe = arithmeticSet nullPtr (-1)
10 | passFormulaSet ctx_c bs = do
11 |   formulas <- mapM (passFormula ctx_c) (BSet.toList bs)
12 |   withArrayLen' formulas arithmeticSet
13 |   where
14 |     withArrayLen' :: Storable a => [a] -> (Ptr a -> CInt -> IO b) -> IO b
15 |     withArrayLen' xs f = withArrayLen xs (\s a -> f a (fromIntegral s))
```

The imported function `hat_mk_arithmetic_set` simply takes an array of arithmetic formulas (C++) and combines them into an `ArithmeticSet`, the C++ equivalent to `SymbolicFormulaSet`. A negative length (second argument) indicates that the set is `Universe`.

`passFormulaSet` passes every formula of the arithmetic set from Haskell to C++ first, and uses `arithmeticSet` to combine them to an `ArithmeticSet` (C++).

The arithmetic analysis is provided in C++:

```
1 | extern "C" {
2 |   ArithmeticSet* hat_arithmetic_value(StablePtr ctx, const HaskellNodeAddress expr_hs);
3 | }
4 |
5 | ArithmeticSet getArithmeticValue(Context& ctxt, const ExpressionAddress& expr) {
6 |   auto expr_hs = ctxt.resolveNodeAddress(expr);
7 |   ArithmeticSet* res_ptr = hat_arithmetic_value(ctxt.getHaskellContext(), expr_hs);
8 |   ArithmeticSet res(std::move(*res_ptr));
9 |   delete res_ptr;
10 |   return res;
11 | }
```

File: `code/analysis/src/cba/haskell/arithmetic_analysis.cpp`

File: `code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Adapter.hs`

File: `code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Analysis/Arithmetic.hs`

#### 4.10.3. Boolean Analysis

Next, the boolean analysis is presented. Its task is to describe the truth value an expression might have during execution. For the property space we define the possible outcomes, together with the

Lattice instance.

```

1 | data Result = AlwaysTrue
2 |     | AlwaysFalse
3 |     | Both
4 |     | Neither
5 | deriving (Eq,Enum,Show,Generic,NFData)
6 |
7 | instance Solver.Lattice Result where
8 |     bot = Neither
9 |
10 |     merge Neither x = x
11 |     merge x Neither = x
12 |     merge x y | x == y = x
13 |     merge _ _ = Both
14 |
15 | instance Solver.ExtLattice Result where
16 |     top = Both

```

If the analysis can determine that the boolean value of an expression is constant (ie is not input dependent), the result is either `AlwaysTrue` or `AlwaysFalse`. If the boolean value is not constant, the result is `Both`. `Neither` is the initial state, until more accurate information regarding the analysed expressions could be determined. Also, it is the value of expressions which are never evaluated (dead code) due to the implicit dead-code detection of our framework.

#### 4. Framework Implementation

Again, we have to setup an identifier for the analysis and provide the main logic of the analysis.

```
1 | data BooleanAnalysis = BooleanAnalysis
2 |   deriving (Typeable)
3 |
4 | booleanValue :: NodeAddress -> Solver.TypedVar (ValueTree SimpleFieldIndex Result)
5 | booleanValue addr =
6 |   case () of _
7 |     | isBuiltin addr "true" -> Solver.mkVariable (idGen addr) []
8 |                                     $ compose AlwaysTrue
9 |     | isBuiltin addr "false" -> Solver.mkVariable (idGen addr) []
10 |                                     $ compose AlwaysFalse
11 |     | otherwise              -> dataflowValue addr analysis ops
12 |
13 | where
14 |   compose = ComposedValue.toComposed
15 |   extract = ComposedValue.toValue
16 |
17 |   analysis = mkDataFlowAnalysis BooleanAnalysis "B" booleanValue
18 |   idGen = mkVarIdentifier analysis
19 |
20 |   ops = [ lt, le, eq, ne, ge, gt ]
21 |
22 |   le = OperatorHandler cov dep (val cmp)
23 |   where
24 |     cov a = any (isBuiltin a) ["int_le", "uint_le"]
25 |     cmp x y = case numCompare x y of
26 |       NumEQ    -> AlwaysTrue
27 |       NumLT    -> AlwaysTrue
28 |       Sometimes -> Both
29 |       _        -> AlwaysFalse
30 |
31 |   -- lt, eq, ne, ge and gt following the same pattern as le
32 |
33 |   lhs = arithmeticValue $ goDown 2 addr
34 |   rhs = arithmeticValue $ goDown 3 addr
35 |
36 |   dep a = Solver.toVar <$> [lhs, rhs]
37 |
38 |   val op a = combine (extract $ Solver.get a lhs) (extract $ Solver.get a rhs)
39 |   where
40 |     combine BSet.Universe _ = compose Both
41 |     combine _ BSet.Universe = compose Both
42 |     combine x y = compose . Solver.join $ [ u `op` v | u <- BSet.toList x
43 |                                             , v <- BSet.toList y ]
```

Initially, we check if a given node is equivalent to the operators true or false. If so, we wrap AlwaysTrue or AlwaysFalse in a ValueTree and return an analysis variable with it as initial value. Otherwise we resort to the generic DFA.

The OperatorHandlers lt, le, eq, ne, ge, and gt are setup to cover their respective comparators in lines 21–30. Furthermore, we have to state the dependency (lines 32–35) between the left and right hand side’s arithmetic value and the current expression.

arithmeticValue, part of the arithmetic analysis (see Section 4.10.2), yields a (bound) set of arithmetic formulas representing the arithmetic values of an expression. The OperatorHandler defined in line 21 covers the operators int\_le and uint\_le, which are used in INSPIRE for the (unsigned) integer comparison *less than or equal*. The remaining OperatorHandlers are defined similarly. Together

with `val`, defined in line 37, it

- extracts the (bound) sets of arithmetic formulas for the left and right hand sides;
- compares each element of the lhs's set with each element of the rhs's set, using `numCompare` (see Section 4.4.2); and
- maps the result of each comparison to the corresponding analysis outcome, defined by the `OperatorHandler`;

It is important not that the logical operations `bool_not`, `bool_and`, and `bool_or` are defined in the *basic language extension* using `IfStmts` and closures to correctly model short-circuit-evaluation of C/C++. We do not need to handle this explicitly thanks to the framework.

This analysis is made available to Insieme via FFI.

```

1 | foreign export ccall "hat_check_boolean"
2 |   checkBoolean :: StablePtr Ctx.Context -> StablePtr Addr.NodeAddress -> IO CInt
3 |
4 | checkBoolean ctx_hs expr_hs = do
5 |   ctx <- deRefStablePtr ctx_hs
6 |   expr <- deRefStablePtr expr_hs
7 |   let (res,ns) = Solver.resolve (Ctx.getSolverState ctx) $ booleanValue expr
8 |       let ctx_c = Ctx.getCContext ctx
9 |           ctx_nhs <- newStablePtr $ ctx { Ctx.getSolverState = ns }
10 |           updateContext ctx_c ctx_nhs
11 |           evaluate $ fromIntegral $ fromEnum $ ComposedValue.toValue res

```

The structure of `checkBoolean` is quite similar to `arithValue`, defined Section 4.10.2. It is sufficient to use a `CInt` for the boolean analysis' result as it is just an enumeration. The conversion happens by combining `fromIntegral`, `fromEnum`, and `ComposedValue.toValue`.

#### 4. Framework Implementation

In C++, `hat_check_boolean` is imported and wrapped.

```
1 | extern "C" {
2 |     int hat_check_boolean(StablePtr ctxt, const HaskellNodeAddress expr_hs);
3 | }
4 |
5 | enum class BooleanAnalysisResult : int {
6 |     AlwaysTrue,
7 |     AlwaysFalse,
8 |     Both,
9 |     Neither
10| };
11|
12| BooleanAnalysisResult checkBoolean(Context& ctxt, const ExpressionAddress& expr) {
13|     auto expr_hs = ctxt.resolveNodeAddress(expr);
14|     int res_hs = hat_check_boolean(ctxt.getHaskellContext(), expr_hs)
15|     auto res = static_cast<BooleanAnalysisResult>(res_hs);
16|     if(res == BooleanAnalysisResult::Neither) {
17|         std::vector<std::string> msgs{"Boolean Analysis Error"};
18|         throw AnalysisFailure(msgs);
19|     }
20|     return res;
21| }
```

The result of `hat_check_boolean` is received as integer and mapped to the corresponding enum value in line 15. The following four convenience functions are provided.

```
1 | bool isTrue(Context& ctxt, const ExpressionAddress& expr) {
2 |     return checkBoolean(ctxt, expr) == BooleanAnalysisResult::AlwaysTrue;
3 | }
4 |
5 | bool isFalse(Context& ctxt, const ExpressionAddress& expr) {
6 |     return checkBoolean(ctxt, expr) == BooleanAnalysisResult::AlwaysFalse;
7 | }
8 |
9 | bool maybeTrue(Context& ctxt, const ExpressionAddress& expr) {
10|     auto res = checkBoolean(ctxt, expr);
11|     return res == BooleanAnalysisResult::AlwaysTrue
12|         || res == BooleanAnalysisResult::Both;
13| }
14|
15| bool maybeFalse(Context& ctxt, const ExpressionAddress& expr) {
16|     auto res = checkBoolean(ctxt, expr);
17|     return res == BooleanAnalysisResult::AlwaysFalse
18|         || res == BooleanAnalysisResult::Both;
19| }
```

Note that for clarity `BooleanAnalysisResult` has been defined in C++ *and* Haskell. In reality, a C enum is defined in a header file which is imported in C++ and Haskell (via `c2hs`). This way they share the exact same definition and are thus implicitly consistent.

File: `code/analysis/src/cba/haskell/boolean_analysis.cpp`

File: `code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Adapter.hs`

File: `code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Analysis/Boolean.chs`

#### 4.10.4. Data Path Analysis

DataPaths have been covered in Section 4.6. The framework also contains an analysis for deducing the (bound) set of possible DataPath values for a given expression. This analysis is used by the reference analysis (see Section 4.10.5).

As we can use the DataPath entity directly we can focus immediately on the lattice instance.

```

1 | type DataPathSet i = BSet.UnboundSet (DataPath i)
2 |
3 | instance (FieldIndex i) => Solver.Lattice (DataPathSet i) where
4 |     bot   = BSet.empty
5 |     merge = BSet.union
6 |
7 | instance (FieldIndex i) => Solver.ExtLattice (DataPathSet i) where
8 |     top = BSet.Universe

```

Note that the data path analysis is, again, a *may analysis*.

dataPathValue does not need to handle specific types of nodes, it is only concerned with operators as it covers the datapath language extension<sup>1</sup>.

We therefore only have to define the relevant OperatorHandlers. The following operators are handled:

dp\_root This operator indicates the root of the object, hence DataPath.root is wrapped and returned.

dp\_member The identifier analysis is used to deduce the target field name, which is then used to modify the DataPath. One DataPath.step is added using the field name.

dp\_element Causes a recursion to find the *nested path* within and uses the arithmetic analysis to find the index used when calling the operator. Both results are then merged.

dp\_component Same as dp\_element.

Note that this analysis is currently not accessible from C++.

File: code/analysis/src/cba/haske11/insieme-hat/src/Insieme/Analysis/DataPath.hs

#### 4.10.5. Reference Analysis

Upon encountering a reference during an analysis, we are often interested in the location this reference points to. The reference analysis is tasked with answering this question. Note that the reference language extension is covered by Section 3.3.2.

In the common case, a reference can be defined by the creation point of the (root) object and a DataPath to reference a sub-object within. The other two edge cases include NULL references and uninitialised

<sup>1</sup>File: code/core/include/insieme/core/lang/datapath.h

#### 4. Framework Implementation

references. This composes our property space.

```
1 | data Reference i = Reference { creationPoint :: NodeAddress
2 |                               , dataPath     :: DP.DataPath i }
3 |                               | NullReference
4 |                               | UninitializedReference
5 |   deriving (Eq,Ord,Show,Generic,NFData)
6 |
7 | type ReferenceSet i = BSet.UnboundSet (Reference i)
8 |
9 | instance (FieldIndex i) => Lattice (ReferenceSet i) where
10 |   bot  = BSet.empty
11 |   merge = BSet.union
12 |
13 | instance (FieldIndex i) => ExtLattice (ReferenceSet i) where
14 |   top = BSet.singleton UninitializedReference
```

Again we have a *may analysis*.

Next we take a look at the corresponding constraint generator function.

```
1 | data ReferenceAnalysis = ReferenceAnalysis
2 |   deriving (Typeable)
3 |
4 | referenceValue :: (FieldIndex i)
5 |                => NodeAddress
6 |                -> TypedVar (ValueTree i (ReferenceSet i))
7 |
8 | referenceValue addr = case getNodeType addr of
9 |   IR.Literal ->
10 |     mkVariable (idGen addr) [] $ compose
11 |               $ BSet.singleton
12 |               $ Reference (crop addr) DP.root
13 |
14 |   IR.Declaration | isMaterializingDeclaration (Addr.getNode addr) ->
15 |     mkVariable (idGen addr) [] $ compose
16 |               $ BSet.singleton
17 |               $ Reference addr DP.root
18 |
19 |   IR.CallExpr | isMaterializingCall (Addr.getNode addr) ->
20 |     mkVariable (idGen addr) [] $ compose
21 |               $ BSet.singleton
22 |               $ Reference addr DP.root
23 |
24 |   _ -> dataflowValue addr analysis opsHandler
25 | where
26 |   analysis = (mkDataFlowAnalysis ReferenceAnalysis "R" referenceValue){
27 |     entryPointParameterHandler = epParamHandler,
28 |     initValueHandler = compose $ BSet.singleton $ NullReference }
29 |   idGen = mkVarIdentifier analysis
30 |   compose = ComposedValue.toComposed
31 |   -- ...
```

As can be seen, the main node types we are interested in are Literals, Declarations and CallExprs where the latter two are only relevant when they are materializing. This is due to the fact that these are the only occurrences where memory is allocated.



The operators defined by the reference language extension are covered by the `OperatorHandlers`. The following list describes how each of them are handled. Note that we do not have to handle derived operators as their interpretation is covered by their definition.

`ref_alloc` The reference points to a newly allocated object, we can therefore use the current node address as creation point and `DataPath.root` to model the resulting reference.

`ref_decl` The target reference can be modelled by using the enclosing `Declaration` of this node as creation point. Again we use `DataPath.root` to refer to the whole object.

`ref_null` Yields a `NullReference`.

`ref_narrow` Since this operator is used to obtain a reference to a sub-object within a referenced object, we not only have to initiate another reference analysis for the referenced object, but also combine the `DataPath(s)` of that analysis' results with the `DataPath` provided to the invocation of `ref_narrow`.

`ref_expand` Handled similarly to `ref_narrow`, just expanding instead of narrowing the `DataPath` of references.

`ref_cast` Causes a recursive call of the reference analysis to the referenced object.

`ref_reinterpret` Same as `ref_cast`.

File: `code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Analysis/Reference.hs`

File: `code/core/include/insieme/core/lang/reference.h`

### 4.10.6. Callable Analysis

The callable analysis is a special case of specific analyses as it does not cover a language extension. Instead it covers function types, which are part of the [INSPIRE](#) core language, together with the related node types `Literal`, `Lambda`, and `BindExpr`.

It is used to identify the target function of a call expression. The three different callable types are

#### 4. Framework Implementation

Lambda, Literal, and Closure. We therefore have the following property space:

```
1 data Callable = Lambda NodeAddress
2               | Literal NodeAddress
3               | Closure NodeAddress
4   deriving (Eq,Ord,Generic,NFData)
5
6 type CallableSet = BSet.UnboundSet Callable
7
8 instance Solver.Lattice CallableSet where
9   bot  = BSet.empty
10  merge = BSet.union
11
12 instance Solver.ExtLattice CallableSet where
13   top  = BSet.Universe
```

The analysis is straight forward in that it handles the related nodes.

```
1 data CallableAnalysis = CallableAnalysis
2   deriving (Typeable)
3
4 callableValue :: NodeAddress -> Solver.TypedVar (ValueTree SimpleFieldIndex CallableSet)
5 callableValue addr = case getNodeType addr of
6   IR.LambdaExpr ->
7     Solver.mkVariable (idGen addr) [] (compose $ BSet.singleton
8                                             $ Lambda (fromJust $ getLambda addr))
9
10  IR.LambdaReference ->
11    Solver.mkVariable (idGen addr) [] (compose $ getCallables4Ref addr)
12
13  IR.BindExpr ->
14    Solver.mkVariable (idGen addr) [] (compose $ BSet.singleton (Closure addr))
15
16  IR.Literal ->
17    Solver.mkVariable (idGen addr) [] (compose $ BSet.singleton (Literal addr))
18
19  _ -> dataflowValue addr analysis []
20 where
21   analysis = mkDataFlowAnalysis CallableAnalysis "C" callableValue
22   idGen    = mkVarIdentifier analysis
23   compose  = ComposedValue.toComposed
24   getCallables4Ref = -- ...
```

The only special case is encountered when dealing with a `LambdaReference`. In this case the tree is searched upwards from the `LambdaReference` for a matching `LambdaDefinition`. If the search was successful the `Lambda` contained in the `LambdaDefinition` is used, otherwise the property space's  $\top$  value is returned. This is achieved by `getCallables4Ref`. This is indicating that any unknown target function may be represented by the targeted expression.

File: code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Analysis/Callable.hs

## 4.11. Helper Analyses

The analyses described in this section are different from the previously described specific analyses in that they do not cover a certain language extension (or **INSPIRE** core language construct) and are not based on the generic **DFA**. However, they are supporting the definition of the generic **DFA** as well as the specific analyses.

### 4.11.1. Reachability Analysis

The reachability analysis is divided into two parts, `reachableIn` and `reachableOut`. Despite both of them being distinct analyses with their own respective analysis variables, they work together in tandem to answer the question of reachability. In detail, `reachableIn` indicates whether the beginning of the given node is reachable. In contrast to this, `reachableOut` indicates whether the end of a given node is reachable during some execution trace.

For instance, the beginning of a statement inside a `CompoundStmt` is reachable iff the end of the previous statement is reachable. If there is no previous statement within the `CompoundStmt`, the reachability depends on whether the beginning of the `CompoundStmt` is reachable.

The property space definitions for both analyses are trivial:

```

1 | newtype Reachable = Reachable Bool
2 |   deriving (Eq,Show,Generic,NFData)
3 |
4 | instance Solver.Lattice Reachable where
5 |   bot = Reachable False
6 |   merge (Reachable a) (Reachable b) = Reachable $ a || b

```

Note that this definition for the property space is sufficient and does not require to be an `ExtLattice` as this analysis is not a **DFA**.

It follows the definition of the analysis identifier token:

```

1 | data ReachableInAnalysis = ReachableInAnalysis
2 |   deriving (Typeable)
3 |
4 | reachableInAnalysis :: Solver.AnalysisIdentifier
5 | reachableInAnalysis = Solver.mkAnalysisIdentifier ReachableInAnalysis "R[in]"
6 |
7 | data ReachableOutAnalysis = ReachableOutAnalysis
8 |   deriving (Typeable)
9 |
10 | reachableOutAnalysis :: Solver.AnalysisIdentifier
11 | reachableOutAnalysis = Solver.mkAnalysisIdentifier ReachableOutAnalysis "R[out]"

```

#### Reachable In

We now look at the definitions of the constraint generators `reachableIn` and `reachableOut`. As both of them are too long to be displayed easily in this thesis as a whole, they are broken up into their

#### 4. Framework Implementation

logical parts.

```
1 reachableIn :: NodeAddress -> Solver.TypedVar Reachable
2 reachableIn a | isRoot a = Solver.mkVariable (reachableInIdGen a) [] (Reachable True)
3 reachableIn a = case getNodeType parent of
4     -- ...
5     where
6         parent = fromJust $ getParent a
7         idGen = Solver.mkIdentifierFromExpression reachableInAnalysis a
8         compose = ComposedValue.toComposed
```

We define the beginning of the root node of **INSPIRE** to be always reachable reflecting the premise that the analysed code fragment is indeed invoked. For a given node, not being the root, we need to distinguish a couple of cases, depending on the type of the parent node. All of the following code snippets (until reaching the definition of `reachableOut`) are replacing line 4.

```
| IR.Lambda -> Solver.mkVariable (idGen a) [] (Reachable True)
```

Considering Lambdas to always be reachable is an over-approximation needed to reduce runtime complexity. Otherwise the necessity for a context sensitive reachability analysis would arise.

```
1 IR.CompoundStmt -> var
2     where
3         n = getIndex a
4         var = Solver.mkVariable (idGen a) [con] Solver.bot
5         con = if n == 0
6             then Solver.forward (reachableIn parent) var
7             else Solver.forward (reachableOut $ goDown (n-1) parent) var
```

As already mentioned by the introduction to this subsection, reachability inside a `CompoundStmt` depends on the previous node inside the `CompoundStmt`, or if there is no previous node, the reachability of the `CompoundStmt` itself.

```
1 IR.IfStmt -> var
2     where
3         var = Solver.mkVariable (idGen a) [con] Solver.bot
4         con = case getIndex a of
5             0 -> Solver.forward (reachableIn parent) var
6
7             1 -> Solver.forwardIf (compose Boolean.AlwaysTrue)
8                                 (Boolean.booleanValue $ goDown 0 parent)
9                                 (reachableOut $ goDown 0 parent)
10                                var
11
12            2 -> Solver.forwardIf (compose Boolean.AlwaysFalse)
13                                (Boolean.booleanValue $ goDown 0 parent)
14                                (reachableOut $ goDown 0 parent)
15                                var
16
17            _ -> error "index out of bound"
```

If the target node is the *then* or *else* branch of an `IfStmt`, the reachability depends on the boolean value of the condition. Therefore we need to conduct a boolean analysis as can be seen in lines 8 and 13.

The condition is considered reachable if the IfStmt itself is reachable.

```

1 | IR.WhileStmt -> var
2 |   where
3 |     var = Solver.mkVariable (idGen a) [con] Solver.bot
4 |     con = case getIndex a of
5 |         0 -> Solver.forward (reachableIn parent) var
6 |
7 |         1 -> Solver.forwardIf (compose Boolean.AlwaysTrue)
8 |                               (Boolean.booleanValue $ goDown 0 parent)
9 |                               (reachableOut $ goDown 0 parent)
10 |                               var
11 |
12 |     _ -> error "index out of bound"

```

Regarding reachability, handling a WhileStmt is identical to handling an IfStmt except that there is no else branch to worry about.

```

1 | _ -> var
2 |   where
3 |     var = Solver.mkVariable (idGen a) [con] Solver.bot
4 |     con = Solver.forward (reachableIn parent) var

```

For all other nodes, the reachability of the target is equal to the reachability of the parent node.

## Reachable Out

We now continue with the definition of reachableOut.

```

1 | reachableOut :: NodeAddress -> Solver.TypedVar Reachable
2 | reachableOut a = case getNodeType a of
3 |     -- ...
4 |   where
5 |     idGen a = Solver.mkIdentifierFromExpression reachableOutAnalysis a

```

For the main part we require a case distinction. The following code snippets replace line 3.

```

1 | IR.ReturnStmt -> Solver.mkVariable (idGen a) [] Solver.bot
2 |
3 | IR.ContinueStmt -> Solver.mkVariable (idGen a) [] Solver.bot
4 |
5 | IR.BreakStmt -> Solver.mkVariable (idGen a) [] Solver.bot

```

The three simple cases are ReturnStmt, ContinueStmt, and BreakStmt. As they directly break the control flow, we can say that the end of these statements is never reachable, hence the use of Solver.bot,

## 4. Framework Implementation

corresponding to `Reachable False`.

```
1 | IR.CompoundStmt -> var
2 |   where
3 |     var = Solver.mkVariable (idGen a) [cnt] Solver.bot
4 |     cnt = if numChildren a == 0
5 |           then Solver.forward (reachableIn a) var
6 |           else Solver.forward (reachableOut $ goDown (numChildren a - 1) a) var
```

The end of a `CompoundStmt` is reachable iff the end of the last statement in it is reachable. Note that (for simplicity) we ignore the possibility of exceptions being raised on destructors, leading to an over-approximation in the result.

```
1 | IR.IfStmt -> var
2 |   where
3 |     var = Solver.mkVariable (idGen a) [t,e] Solver.bot
4 |     t = Solver.forward (reachableOut $ goDown 1 a) var
5 |     e = Solver.forward (reachableOut $ goDown 2 a) var
```

In case of an `IfStmt`, the end is reachable, iff the end of one of the branches is reachable.

```
1 | IR.WhileStmt -> var
2 |   where
3 |     var = Solver.mkVariable (idGen a) [cnt] Solver.bot
4 |     cnt = Solver.forwardIf (ComposedValue.toComposed Boolean.AlwaysFalse)
5 |                           (Boolean.booleanValue $ goDown 0 a)
6 |                           (reachableOut $ goDown 0 a)
7 |     var
```

The end of a `WhileStmt` is reachable iff the condition can evaluate to false. Again, we therefore require the boolean analysis.

```
1 | _ -> var
2 |   where
3 |     var = Solver.mkVariable (idGen a) [con] Solver.bot
4 |     con = Solver.forward (reachableIn a) var
```

All other nodes follow the same schema where the end is reachable iff the beginning is reachable.

File: `code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Analysis/Reachable.hs`

### 4.11.2. Exit Point Analysis

The exit point analysis tries to determine the possible exit points of a functions. These are typically `ReturnStmts` and the end of the function body. Reachability analysis (Section 4.11.1) is used to determine whether the control flow can reach a potential exit point, or not.

The definition of the corresponding property space is trivial:

```

1 | newtype ExitPoint = ExitPoint NodeAddress
2 |   deriving (Eq,Ord,Generic,NFData)
3 |
4 | type ExitPointSet = BSet.UnboundSet ExitPoint
5 |
6 | instance Solver.Lattice ExitPointSet where
7 |     bot    = BSet.empty
8 |     merge  = BSet.union

```

Similar to the reachability analysis (previous section), the exit point analysis is not a *DFA* and therefore does not require the instantiation of an `ExtLattice`.

Only `BindExpr` and `Lambda` are of interest to this analysis. All other node types yield an empty set of `ExitPoints`. For a `BindExpr` the exit point is simply the nested `CallExpr`. For a given `Lambda`, all `ReturnStmts` are collected and checked for reachability. All reachable exit points are joined in an `ExitPointSet`.

File: `code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Analysis/ExitPoint.hs`

### 4.11.3. Predecessor Analysis

The traversal of `ProgramPoints` has already been illustrated in Figure 9. The predecessor analysis models this traversal accordingly. It is split into three main parts: the first part handles traversal from a `Pre ProgramPoint` to its previous `Post ProgramPoint`, the second deals with `Internal ProgramPoints`, and the last one is concerned with the predecessor of `Post ProgramPoints`.

For a given `ProgramPoint` the list of predecessors is computed and represented as constraints in our framework. The property space simply consists of a list of `ProgramPoints` and concatenation.

```

1 | type PredecessorList = [ProgramPoint]
2 |
3 | instance Solver.Lattice PredecessorList where
4 |     join [] = []
5 |     join xs = foldr1 (++) xs

```

For the predecessor analysis we do not use a (bound) set, since for each node, there only exist a limited number of possible predecessors.

The definition of the corresponding analysis identifier and the signature of the analysis' constraint generator function follows.

```

1 | data PredecessorAnalysis = PredecessorAnalysis
2 |   deriving (Typeable)
3 |
4 | predecessorAnalysis :: Solver.AnalysisIdentifier
5 | predecessorAnalysis = Solver.mkAnalysisIdentifier PredecessorAnalysis "pred_of"
6 |
7 | predecessor :: ProgramPoint -> Solver.TypedVar PredecessorList

```

The three parts mentioned can be distinguished by pattern matching on the `ProgramPoint` parameter of the function. The base case for this traversal is encountered when asking for the predecessors of

#### 4. Framework Implementation

the root's Pre ProgramPoint. Since there is no execution step before this program point, the resulting analysis variable is bound to the empty list.

```
1 | predecessor p@(ProgramPoint a Pre) | isRoot a = Solver.mkVariable (idGen p) [] []
2 | predecessor p@(ProgramPoint a Pre) = case getNodeType parent of
3 |     -- handle part 1
4 |
5 | predecessor p@(ProgramPoint a Internal) = case getNodeType a of
6 |     -- handle part 2
7 |
8 | predecessor p@(ProgramPoint a Post) = case getNodeType a of
9 |     -- handle part 3
10 |
11 |
12 | idGen :: ProgramPoint -> Solver.Identifier
13 | idGen pp = Solver.mkIdentifierFromProgramPoint predecessorAnalysis pp
```

The second part, concerned with Internal ProgramPoints, utilises the exit point Analysis (see Section 4.11.2) and callable analysis (see Section 4.10.6) to link to exit points of potential target functions.

An Internal ProgramPoint is returned by the third part only for CallExprs. All other node types either yield a Pre ProgramPoint if they are a basic expression, like Variable or Literal, or the corresponding Post ProgramPoint otherwise.

File: code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Analysis/Predecessor.hs

#### 4.11.4. Call Site Analysis

The call site analysis identifies all possible call sites for a given Lambda or BindExpr. We can use node addresses to reference a call site, the lattice definition follows:

```
1 | newtype CallSite = CallSite NodeAddress
2 |   deriving (Eq,Ord,Generic,NFData)
3 |
4 | instance Show CallSite where
5 |   show (CallSite na) = "Call@" ++ (prettyShow na)
6 |
7 | type CallSiteSet = BSet.UnboundSet CallSite
8 |
9 | instance Solver.Lattice CallSiteSet where
10 |   bot  = BSet.empty
11 |   merge = BSet.union
```

Call sites can only be derived for Lambda and BindExpr nodes, and only when it is not the root node. In the simple case all call sites are statically determined. Otherwise, we have to utilise the callable analysis (see Section 4.10.6) to identify all potential call sites.

File: code/analysis/src/cba/haskell/insieme-hat/src/Insieme/Analysis/CallSite.hs



## 5. Constructing a new Analysis

In this chapter we construct a new analysis using the established framework. Our goal is to implement an *array out-of-bounds* analysis. We go through the construction, implementation, and integration into INSIEME step by step:

1. Defining the goals of the new analysis.
2. Investigating the relevant constructs and language extensions of the [INSieme Parallel Intermediate REpresentation \(INSPIRE\)](#).
3. Deriving from the gathered information what components / additional analyses are required to fulfil our goal.
4. Defining the property spaces of the analyses to implement.
5. Defining the interface for communication between INSIEME and the [Haskell-based Analysis Toolkit \(HAT\)](#).
6. Setting up the required boilerplate code.
7. Writing a handful of test cases to allow for checking the implementation during development.
8. Implementing the required analyses.
9. Summarising what files have been added / modified by the process.

The implemented analysis is used to evaluate the current state of the framework.

### 5.1. Defining the Goals

The goal is to create an analysis which can identify out-of-bounds accesses to arrays. For simplicity we will only concern ourselves with accessing elements of an array beyond its upper bound. Also covering the lower bound will be a trivial extension and is left as an exercise to the reader.

## 5. Constructing a new Analysis

The following code fragment will illustrate the essentials:

```
1  { // Case 1
2      int a[21];
3      a[42];
4  }
5
6  { // case 2
7      int a[21];
8      a[20];
9  }
10
11 { // Case 3
12     int a[21];
13     int *ap = a + 10;
14     ap[20];
15 }
16
17 { // Case 4
18     auto a = new int[21];
19     a[42];
20     delete[] a;
21 }
22
23 { // Case 5
24     int a[21];
25     int b;
26     a[b];
27 }
28
29 { // Case 6
30     int a[21];
31     int b = a[42];
32 }
```

**Case 1** Accessing an element beyond the array's upper bound. The analysis returns `IsOutOfBounds` for `a[42]`.

**Case 2** Accessing an element within the array's upper bound. The analysis returns `IsNotOutOfBounds` for `a[20]`.

**Case 3** Same as Case 1, yet the array is accessed indirectly via a pointer offset to the base of the array. Same result as in Case 1.

**Case 4** Same as Case 1, yet the array is allocated using the `new` keyword. Same result as in Case 1.

**Case 5** Accessing the array with an undefined index, the analysis returns `MaybeOutOfBounds`.

**Case 6** Same as Case 1, but involves an assignment.

This marks the base for the following sections. Note that the framework will cover more complicated cases (eg where a pointer to an array is passed to a function and the function accesses the underlying array) automatically thanks to the generic [Data-Flow Analysis \(DFA\)](#).

## 5.2. Investigating *INSPIRE*

The different cases defined in the previous section are used as input code the Insieme compiler. The analysis runs on the resulting *INSPIRE* program, constructed by the frontend of the compiler. Using `insiemecc` we can feed each of the cases to the compiler and retrieve the corresponding **Intermediate Representation (IR)** in various formats. `case1.cpp` is the input to the compiler, while `case1.ir`, `case1.tree`, and `case1.json` compose the output we are interested in. The following command is used to obtain the files.

```
| insiemecc --dump-tree case1.tree --dump-json case1.json --dump-ir case1.ir case1.cpp
```

`case1.ir` contains the pretty-printed text representation of the corresponding *INSPIRE* program. `case1.tree` can be viewed with a regular text editor to investigate the node structure of this *INSPIRE* program. Furthermore `case1.json` also contains the node structure, but in a format that can be viewed using *INSPYER* (see Appendix A), which provides a more interactive viewing experience.

Let us now take a look at the first case.

## 5. Constructing a new Analysis

```

1 | int main(void) {
2 |     int a[21];
3 |     a[42];
4 | }
1 | decl IMP_main : () -> int<4>;
2 | // Inspire Program
3 | int<4> function IMP_main (){
4 |     var ref<array<int<4>,21>,f,f,plain> v1 =
5 |         ref_decl(type_lit(ref<array<int<4>,21>,f,f,plain>));
6 |     *ptr_subscript(ptr_from_array(v1), 42);
7 |     return 0;
8 | }

1 | (Program | ...
2 |   (CompoundStmt |
3 |     (DeclarationStmt |
4 |       (Declaration |
5 |         (GenericType | ... )
6 |         (CallExpr |
7 |           (GenericType |
8 |             (StringValue "ref")
9 |             (Parents )
10 |            (Types |
11 |              (GenericType |
12 |                (StringValue "array")
13 |                (Parents )
14 |                (Types |
15 |                  (GenericType | ... )
16 |                  (NumericType |
17 |                    (Literal |
18 |                      (GenericType | ... )
19 |                      (StringValue "21"))))
20 |                  ... )))
21 |            (Literal |
22 |              (FunctionType | ... )
23 |              (StringValue "ref_decl"))
24 |            (Declaration | ... )))
25 |     (Variable | ... ))
26 |   (CallExpr |
27 |     (GenericType | ... )
28 |     (Literal |
29 |       (FunctionType | ... )
30 |       (StringValue "ref_deref"))
31 |     (Declaration |
32 |       (GenericType | ... )
33 |       (CallExpr |
34 |         (GenericType | ... )
35 |         (LambdaExpr |
36 |           (FunctionType | ... )
37 |           (LambdaReference |
38 |             (FunctionType | ... )
39 |             (StringValue "ptr_subscript"))
40 |           (LambdaDefinition | ... ))
41 |         (Declaration | ... )
42 |         (Declaration | ... )))
43 |     (ReturnStmt | ... )...

```

Listing 9: case1.{cpp,ir,tree}

As can be seen in case1.ir and case1.tree, the array is constructed using the operator `ref_decl` and accessed via the `ptr_subscript` operator. Note that the pointer subscript operation is *not* prohibited, however, dereferencing its result *is*. We therefore target calls to `ref_deref` with our analysis.

Furthermore, it is important to note that the array is never handled directly, but via a reference. Hence our analysis is concerned with the type `ref<array<'a>>`. This also means that (at least one of the components) is very similar to the reference analysis discussed earlier in Section 4.10.5. It follows that a look at the reference language extension is essential. Furthermore, as the array is also handled as pointer, taking a look at the pointer language extension is helpful too.

The relevant parts can be found in the header files `insieme/core/lang/reference.h` and `insieme/core/pointer.h`. First thing to note here is that the pointer language extension builds on top of the

reference extension.

For the creation process of the array, `ref_decl` is used. As this operator is *not* derived, we have to handle it explicitly using an `OperatorHandler`. See the declaration in `reference.h`

```
| LANG_EXT_LITERAL(RefDecl, "ref_decl", "(type<ref<'a','c','v','k>>) -> ref<'a','c','v','k>")
```

This size of the array (21) can be inferred from its argument `type_lit(ref<array<int<4>, 21>, f, f, plain>)`.

For the access we need to investigate the pointer language extension and immediately notice the following type alias

```
| TYPE_ALIAS("ptr<'a','c','v>", "( ref<array<'a','c','v>, int<8> )");
```

which indicates that a pointer in *INSPIRE* is handled as a pair of array reference and offset. Its use can be seen when inspecting the definition of the derived operator `ptr_subscript`.

```
1 | LANG_EXT_DERIVED(PtrSubscript, R"(
2 |     (p : ptr<'a','c','v>, i : int<8>) -> ref<'a','c','v> {
3 |         return p.0[p.1 + i];
4 |     }
5 | )")
```

In line 3, the offset `i`, used for the `ptr_subscript` call, is simply added to the second part of the pair `p` before the actual subscript happens.

As this function requires a pointer as argument, the original array is converted to a pointer by the frontend using the derived operator `ptr_from_array`, which in turn uses the abstract operator `ref_reinterpret`. `ref_reinterpret` is a reinterpret cast altering the actual interpretation of the referenced memory call.

```
1 | // from pointer.h
2 | LANG_EXT_DERIVED(PtrFromArray, R"(
3 |     (r : ref<array<'a','s>,'c','v>) -> ptr<'a','c','v> {
4 |         return ( ref_reinterpret(r,type_lit(array<'a,inf>)), 0l );
5 |     }
6 | )")
7 |
8 | // from reference.h
9 | LANG_EXT_LITERAL(RefReinterpret, "ref_reinterpret", "(ref<'a','c','v','k>, type<'b>)"
10 |     " -> ref<'b','c','v','k>")
```

The structure of Case 2 is identical to this one, the only difference is the literal for accessing the array element. We therefore now look at Case 3.

## 5. Constructing a new Analysis

```

1 | int main(void) {
2 |     int a[21];
3 |     int *ap = a + 10;
4 |     ap[20];
5 | }

```

```

1 | decl IMP_main : () -> int<4>;
2 | // Inspire Program
3 | int<4> function IMP_main (){
4 |     var ref<array<int<4>,21>,f,f,plain> v1 =
5 |         ref_decl(type_lit(ref<array<int<4>,21>,f,f,plain>));
6 |
7 |     var ref<ptr<int<4>>,f,f,plain> v2 =
8 |         ptr_add(ptr_from_array(v1), 10);
9 |
10 |    *ptr_subscript(*v2, 20);
11 |
12 |    return 0;
13 | }

```

Listing 10: case3.{cpp,ir}

In this **INSPIRE** program, the array is accessed via the pointer  $v_2$  which is already offset by 10 elements. This offset is created by using the derived operator `ptr_add`. Its definition follows:

```

1 | LANG_EXT_DERIVED(PtrAdd, R"(
2 |     (p : ptr<'a','c','v', i : int<8>) -> ptr<'a','c','v'> {
3 |         return ( p.0, p.1 + i );
4 |     }
5 | )")

```

As only this derived operator is used, no additional changes to the analysis are necessary compared to Case 1 and 2.

Case 4 uses a different mechanism for allocating the array.

```

1 | int main(void) {
2 |     auto a = new int[21];
3 |     a[42];
4 |     delete[] a;
5 | }

```

```

1 | decl IMP_main : () -> int<4>;
2 | // Inspire Program
3 | int<4> function IMP_main (){
4 |     var ref<ptr<int<4>>,f,f,plain> v1 =
5 |         ptr_from_array(
6 |             // type
7 |             <ref<array<int<4>,21>,f,f,plain>>
8 |
9 |             // memory location
10 |            (ref_new(type_lit(array<int<4>,21>)))
11 |
12 |            // initialisation
13 |            {}
14 |        );
15 |
16 |    *ptr_subscript(*v1, 42);
17 |
18 |    ref_delete(ptr_to_array(*v1));
19 |
20 |    return 0;
21 | }

```

Listing 11: case4.{cpp,ir}

The syntax from line 6 to line 13 corresponds to an `InitExpr`, where the type is written between `< >`,

the memory location between ( ), and the initialisation between { }. In this case the type is straight forward, `ref_new` is used to allocate the array on the heap, and the array is not initialised.

We therefore have to look at the operator `ref_new` and see that it is derived and uses `ref_alloc` which is an abstract operator. This abstract operator needs to be handled by the analysis using an `OperatorHandler`.

```

1 | LANG_EXT_DERIVED(RefNew, R"(
2 |     (t : type<'a>) -> ref<'a,f,f> {
3 |         return ref_alloc(t, mem_loc_heap );
4 |     }
5 | )")
6 |
7 | LANG_EXT_LITERAL(RefAlloc, "ref_alloc", "(type<'a>, memloc) -> ref<'a,f,f>")

```

Case 5 does not introduce any new constructs, thus we can continue our investigation process.

In Case 6, the call to `ref_deref` we are interested in is wrapped in a `Declaration` which in turn is wrapped in a `DeclarationStmt` to model the assignment. Since nothing else changed, we continue with a summary.

To summarise our findings we put together a table listing the relevant operators, shortly describing each of them.

Operator	Type	Description
<code>ptr_add</code>	<i>derived</i>	Offsets a given pointer with a given offset.
<code>ptr_from_array</code>	<i>derived</i>	Casts a given pointer to an array.
<code>ptr_subscript</code>	<i>derived</i>	Returns a reference to a specific array element.
<code>ref_alloc</code>	<i>abstract</i>	Allocate an object of the given type at a specific memory location (stack / heap), returning a reference to it.
<code>ref_decl</code>	<i>abstract</i>	Declares an object of the given type and returning a reference to it.
<code>ref_deref</code>	<i>abstract</i>	Obtain the data stored in the memory location referenced by the given reference.
<code>ref_new</code>	<i>derived</i>	Allocates an object of the given type at the heap, returning a reference to it.
<code>ref_reinterpret</code>	<i>abstract</i>	Alters the actual interpretation of the referenced memory cell.

### 5.3. Designing the Analysis

Our design of this new analysis is based on the investigations done in the previous section. As has already been mentioned, the type we are concerned with the most is `ref<array<'a>>`. It follows that the operators of the reference language extensions are relevant to this analysis.

## 5. Constructing a new Analysis

We decide to split the analysis into two parts. One for identifying the index (ie offset) used to access the array, one for deriving the size (ie number of elements) of the array.

The first part can be accomplished fully by the reference analysis (see Section 4.10.5). The key to this is by investigating the corresponding `DataPath` derived by the reference analysis.

For the second part we have create a new specialised `DFA`, tracking the number of elements of an allocated array. We will therefore refer to this analysis as *element count analysis*.

A call to our (yet to construct) function `outOfBounds` will trigger both analyses and compare the results. If an out-of-bounds access is noticed (ie an access where the offset is equal or greater to the array's element count), `IsOutOfBounds` is returned. If one (or both) of the triggered analyses could not determine an accurate enough result, `MaybeOutOfBounds` is returned. Otherwise we can safely say that no out-of-bounds (according to our set goals) occurs returning `IsNotOutOfBounds`.

Next we define the property spaces of our two *sub*-analyses.

## 5.4. Defining the Property Spaces

This is straightforward as we have already covered the reference analysis in Section 4.10.5 and `DataPaths` in Section 4.10.4. Using the established reference analysis on the offset, we will get back a set of `References`, where `References` are defined as follows:

```
1 | data Reference i = Reference { creationPoint :: Location
2 |                               , dataPath     :: DP.DataPath i }
3 |                               | NullReference
4 |                               | UninitializedReference
5 | deriving (Eq,Ord,Show,Generic,NFData)
```

Where we are interested in the `DataPath` of a `Reference`. The type parameter `i` will be a `SimpleFieldIndex` (see Section 4.6). From the `DataPath` we can infer the index used for accessing the array.

The element count analysis provides us a `SymbolicFormulaSet` like the arithmetic analysis (see Section 4.10.2). This set contains `SymbolicFormulas` modelling the element count of the array. We reuse the `Lattice` and `ExtLattice` instance definitions of `SymbolicFormula`.

## 5.5. Defining the Interface

In this case, the interface is trivial and can be derived from the design process. The possible outcomes are gathered in a single `Algebraic Data Type (ADT)`.

```
1 | OutOfBoundsResult = MaybeOutOfBounds
2 |                   | IsNotOutOfBounds
3 |                   | IsOutOfBounds
4 | deriving (Eq,Ord,Enum,Show,Generic,NFData)
```

We could have used the same approach as for the boolean analysis (see Section 4.10.3), by defining this enumeration in a dedicated header file, which is then used in Haskell and C++ to ensure that the



mapping is always correct. Yet, for simplicity, we chose not to and define an identical enumeration in C++. The definition goes into a new header file, dedicated to the out-of-bounds analysis:

```

1 | // insieme/analysis/cba/haskell/out_of_bounds_analysis.h
2 |
3 | // namespace insieme::analysis::cba::haskell
4 | enum class OutOfBoundsResult : int {
5 |     MaybeOutOfBounds,
6 |     IsNotOutOfBounds,
7 |     IsOutOfBounds,
8 | };

```

Furthermore we add the prototype of the analysis function to the same header file.

```

1 | // insieme/analysis/cba/haskell/out_of_bounds_analysis.h
2 |
3 | // namespace insieme::analysis::cba::haskell
4 | OutOfBoundsResult getOutOfBounds(Context& ctxt, const CallExprAddress& expr);

```

The Haskell function representing the out-of-bounds analysis has the following signature and goes into a new file `OutOfBounds.hs`:

```

| outOfBounds :: SolverState -> NodeAddress -> (OutOfBoundsResult, SolverState)

```

As the interface for C++ as well as Haskell has now been established, the next step is to write the necessary boiler plate code in the adapter.

## 5.6. Setup Boilerplate

The first thing we take care of is the exporting of the Haskell `outOfBounds` function. As our result is an enumeration we can simply pass an integer representing the corresponding enum value via the

## 5. Constructing a new Analysis

Foreign Function Interface (FFI). The following code is added to `Adapter.hs`.

```
1 | foreign export ccall "hat_out_of_bounds"
2 |   outOfBounds :: StablePtr Ctx.Context -> StablePtr Addr.NodeAddress -> IO CInt
3 |
4 | outOfBounds ctx_hs expr_hs = do
5 |     ctx <- deRefStablePtr ctx_hs
6 |     expr <- deRefStablePtr expr_hs
7 |     let (result,ns) = OOB.outOfBounds (Ctx.getSolverState ctx) expr
8 |         let ctx_c = Ctx.getCContext ctx
9 |             ctx_nhs <- newStablePtr $ ctx { Ctx.getSolverState = ns }
10 |             updateContext ctx_c ctx_nhs
11 |     return $ fromIntegral $ fromEnum result
```

The export is available under the symbol `hat_out_of_bounds` which we import in C++ next.

```
1 | extern "C" {
2 |     namespace hat = insieme::analysis::cba::haskell;
3 |     int hat_out_of_bounds(hat::StablePtr ctx, const hat::HaskellNodeAddress expr_hs);
4 | }
```

This import and the wrapper around it (next code snippet) both go into a new source file `out_of_bounds_analysis.cpp`.

```
1 | // namespace insieme::analysis::cba::haskell
2 | OutOfBoundsResult getOutOfBounds(Context& ctxt, const core::CallExprAddress& call) {
3 |     const auto& refext = call.getNodeManager()
4 |         .getLangExtension<core::lang::ReferenceExtension>();
5 |
6 |     if (!refext.isCallOfRefDeref(call)) {
7 |         return OutOfBoundsResult::IsNotOutOfBounds;
8 |     }
9 |
10 |     auto call_hs = ctxt.resolveNodeAddress(call);
11 |     int res = hat_out_of_bounds(ctxt.getHaskellContext(), call_hs);
12 |     return static_cast<OutOfBoundsResult>(res);
13 | }
```

As the analysis is to be only invoked on calls to the `ref_deref` operator, we import the reference language extension in lines 3 and 4 and ensure that the argument `call` is indeed a call to `ref_deref` in lines 6–8.

Note that we have to add the new Haskell module (defined by `OutOfBounds.hs`) to the list of modules in `insieme-hat.cabal`.

## 5.7. Write Test Cases

Before actually implementing the analysis, setting up a few test cases is a good idea. This eases the development process as we can immediately verify the correctness of our implementation. Since we have already defined six cases the analysis should cover, we can morph these cases into unit tests.

Insieme uses the Google Test Framework<sup>1</sup> (aka gtest) for most of its infrastructure testing – therefore we also use gtest and integrate the tests into the infrastructure.

Important to note here is that we are not using the original input code for our tests, but the corresponding **INSPIRE** programs. Otherwise we would have a strong coupling between the analysis module and Insieme’s frontend.

Each test case for the out-of-bounds analysis goes into the file `out_of_bounds_analysis_test.cc` and will be structured as follows.

```

1 | TEST(OutOfBounds, Basic) {
2 |     NodeManager mgr;
3 |     IRBuilder builder(mgr);
4 |     Context ctx;
5 |
6 |     auto stmt = builder.parseStmt(
7 |         "{"
8 |         "   var ref<array<int<4>, 21>> a = ref_decl(type_lit(ref<array<int<4>, 21>>));"
9 |         "   *ptr_subscript(ptr_from_array(a), 42);"
10 |        "}"
11 |    ).as<CompoundStmtPtr>();
12 |    auto call = CompoundStmtAddress(stmt)[1].as<core::CallExprAddress>();
13 |
14 |    ASSERT_EQ(OutOfBoundsResult::IsOutOfBounds, getOutOfBounds(ctx, call));
15 | }

```

In line 1, `OutOfBounds` is the group this test case belongs to and `Basic` is its (unique) name. The argument to the `parseStmt` call in line 6 is the **INSPIRE** program for this test. The `CallExpr` of interest is extracted from it in line 12 and passed to the out-of-bounds analysis as input in line 14. We expect the result to be `IsOutOfBounds`.

Note that the extracted `CallExpr` (input to the analysis) is a call to `ref_deref` as the problem of out-of-bound access occurs only upon dereferencing.

The file `out_of_bounds_analysis_test.cc` is placed inside the test sub-directory of the analysis module. Insieme’s build infrastructure will automatically create an executable allowing us to run all test cases of the out-of-bounds analysis in a single go. We are also able to run only a subset of the test cases by using the `--gtest_filter` option provided by gtest.

## 5.8. Implementation

The big picture of this analysis has been communicated in Section 5.3. Before starting with implementing the element count analysis we define ourselves the following helper functions.

```

1 | maybeToBool :: Maybe Bool -> Bool
2 | maybeToBool = Data.Foldable.or
3 |
4 | goesDown :: [Int] -> NodeAddress -> NodeAddress
5 | goesDown l = foldr1 (.) $ goDown <$> reverse l

```

<sup>1</sup><https://github.com/google/googletest>

## 5. Constructing a new Analysis

### 5.8.1. Element Count Analysis

The element count analysis is a specialised DFA and therefore requires some boilerplate code:

```
1 | data ElementCountAnalysis = ElementCountAnalysis
2 |   deriving (Typeable)
3 |
4 | elementCountAnalysis :: DataFlowAnalysis ElementCountAnalysis
5 |                       (ValueTree SimpleFieldIndex
6 |                       (SymbolicFormulaSet BSet.Bound10))
7 |
8 | elementCountAnalysis = (mkDataFlowAnalysis ElementCountAnalysis "EC" elementCount)
9 |
10 | elementCount :: NodeAddress -> Solver.TypedVar (ValueTree SimpleFieldIndex
11 |                                                  (SymbolicFormulaSet BSet.Bound10))
12 |
13 | elementCount addr = dataflowValue addr elementCountAnalysis ops
14 |   where -- ...
```

This analysis covers the operators of the reference language extension and therefore uses `OperatorHandlers` (`ops`) to model their semantics. The simplest case is encountered when dealing with a `ref_null` as the element count is zero.

```
1 | refNull = OperatorHandler cov dep val
2 |   where
3 |     cov a = isBuiltin a "ref_null"
4 |     dep _ = []
5 |     val a = toComposed $ BSet.singleton $ Ar.zero
```

Next, the operators `ref_cast`, `ref_reinterpret`, `ref_narrow`, and `ref_expand` simply forward the element count of their (first) argument as their semantics do not yield any modifications.

```
1 | noChange = OperatorHandler cov dep val
2 |   where
3 |     cov a = any (isBuiltin a) ["ref_cast", "ref_reinterpret", "ref_narrow", "ref_expand"]
4 |     dep _ = [Solver.toVar baseRefVar]
5 |     val a = Solver.get a baseRefVar
6 |
7 |     baseRefVar = elementCount $ goDown 1 $ goDown 2 addr
```

`ref_decl` and `ref_new` are used for the creation of arrays as we have observed at the beginning of this chapter. From them we can extract the size of the array.

```
1 | creation = OperatorHandler cov dep val
2 |   where
3 |     cov a = any (isBuiltin a) ["ref_decl", "ref_new"] && isRefArray
4 |     dep _ = Solver.toVar arraySize
5 |     val a = Solver.get a arraySize
6 |
7 |     arraySize = arithmeticValue $ goesDown [0,2,0,2,1,0] addr
8 |
9 |     isRefArray = maybeToBool $ isArrayType <$> (getReferencedType $ goDown 0 addr)
```

Note the check in line 3 as we only cover creations of arrays. From the call to `ref_decl` or `ref_new` we invoke the arithmetic analysis on the node representing the size of the array and forward its result.

The location of the array size, here the node path  $[0, 2, 0, 2, 1, 0]$  can be inferred from Listing 9.

We also support scalars, their element count will be one.

```

1 | scalar = OperatorHandler cov dep val
2 |   where
3 |     cov a = any (isBuiltin a) ["ref_decl", "ref_new"] && not isRefArray
4 |     dep _ = []
5 |     val _ = toComposed $ BSet.singleton $ Ar.one

```

This completes our list of operator handlers.

```
ops = [refNull, noChange, creation, scalar]
```

## 5.8.2. Getting the Array Index

The element count analysis provides us with a `SymbolicFormulaSet`. In order to (easily) check for out-of-bounds accesses, the array index part should provide its result in a similar data structure. From the reference analysis we get a `ReferenceSet`, where the containing `References` provide the `DataPath` from which the index can be extracted. Therefore we now construct the parts responsible for bringing the result into the wanted data structure

```

1 | toDataPath :: Ref.Reference i -> Maybe (DP.DataPath i)
2 | toDataPath (Ref.Reference _ dp) = Just dp
3 | toDataPath _                    = Nothing
4 |
5 | toPath :: DP.DataPath i -> Maybe [i]
6 | toPath dp | DP.isInvalid dp = Nothing
7 | toPath dp                    = Just $ DP.getPath dp
8 |
9 | toFormula :: [SimpleFieldIndex] -> Maybe SymbolicFormula
10 | toFormula []                  = Just Ar.zero
11 | toFormula (Index i:_)        = Just $ Ar.mkConst $ fromIntegral i
12 | toFormula _                  = Nothing

```

Note that each of the functions returns a `Maybe` as there exists the possibility for the conversion to fail. This would happen, for instance, if the received `ReferenceSet` contains a `NullReference`.

All that is left to do, is to combine these parts using Kleisli composition, convert the results to `ArrayAccess`, and adjust the bound; yielding:

```

1 | data ArrayAccess = ArrayAccess SymbolicFormula
2 |                   | InvalidArrayAccess
3 |   deriving (Eq,Ord,Show,Generic,NFData)
4 |
5 | convertArrayIndex :: Ref.ReferenceSet SimpleFieldIndex -> BSet.UnboundSet ArrayAccess
6 | convertArrayIndex = BSet.changeBound
7 |                   . BSet.map (maybe InvalidArrayAccess ArrayAccess)
8 |                   . BSet.map (toDataPath >=> toPath >=> toFormula)

```

## 5. Constructing a new Analysis

### 5.8.3. Putting it Together

Now that we have the array size and the index used to access it in a usable format, we can put the two parts together, finally implementing `outOfBounds`. `arraySize` and `arrayIndex` is processed as follows.

```
1 | (arrayIndex',ns) = Solver.resolve init
2 |     $ Ref.referenceValue
3 |     $ goDown 1 $ goDown 2 addr
4 |
5 | arrayIndex :: BSet.UnboundSet ArrayAccess
6 | arrayIndex = convertArrayIndex $ toValue arrayIndex'
7 |
8 | (arraySize',ns') = Solver.resolve ns
9 |     $ elementCount
10 |     $ goDown 1 $ goDown 2 addr
11 |
12 | arraySize :: SymbolicFormulaSet BSet.Unbound
13 | arraySize = BSet.changeBound $ toValue arraySize'
```

First the `arrayIndex` is derived using the reference analysis and an initial `SolverState` (`init`). The resulting `SolverState` (`ns`) is then used with the new element count analysis to derive `arraySize`.

Next, each element of `arrayIndex` is compared with each element of `arraySize` and checked for out-of-bounds access.

```
1 | isOutOfBounds :: ArrayAccess -> SymbolicFormula -> Bool
2 | isOutOfBounds (ArrayAccess i) s = Ar.numCompare i s /= Ar.NumLT
3 | isOutOfBounds InvalidArrayAccess _ = True
4 |
5 | oobs = BSet.toList $ BSet.lift2 isOutOfBounds arrayIndex arraySize
```

The result of the analysis is determined by whether `oobs` contains at least a single `True`.

```
1 | outOfBounds :: SolverState -> NodeAddress -> (OutOfBoundsResult,SolverState)
2 | outOfBounds init addr = (result,ns')
3 |   where
4 |     result = case () of
5 |       _ | BSet.isUniverse arrayIndex -> MaybeOutOfBounds
6 |         | BSet.isUniverse arraySize -> MaybeOutOfBounds
7 |         | or oobs -> IsOutOfBounds
8 |         | otherwise -> IsNotOutOfBounds
```

## 5.9. Summary

To summarise the construction process a table is presented providing an overview of what files have been added / modified together with a short description.

File	State	Description
<code>Adapter.hs</code>	<i>modified</i>	Now also contains Haskell export declarations for analysis function.
<code>OutOfBounds.hs</code>	<i>added</i>	Defines the out-of-bounds analysis and its property space.
<code>insieme-hat.cabal</code>	<i>modified</i>	Contains a list of all Haskell modules of <a href="#">HAT</a> .
<code>out_of_bounds_analysis.cpp</code>	<i>added</i>	Wraps the imported out-of-bounds analysis function for C++.
<code>out_of_bounds_analysis.h</code>	<i>added</i>	Defines the analysis result and function prototype in C++.
<code>out_of_bounds_analysis_test.cc</code>	<i>added</i>	Contains a unit test suite for the out-of-bounds analysis.

The `adapter` and `insieme-hat.cabal` are the only files that have been modified, all other files are simply added to the code base. Note that it is not mandatory to have all [FFI](#) export declarations located in the adapter. We could also have stated those directives in `OutOfBounds.hs`, yet we decided against it in order to keep all [FFI](#) export declarations (and their relevant logic) in one single (Haskell) module.





## 6. Evaluation

The evaluation of the toolkit is done by using the newly constructed analysis from Chapter 5. This process is split into two parts, one focusing on the *qualitative* features, the other focusing on the *quantitative*.

### 6.1. Qualitative Evaluation

The qualitative evaluation focuses on the ability of the toolkit to lift an analysis to more complex input programs than explicitly covered by the analysis implementation. Non of the explicitly developed analyses components is dealing with function calls. Yet, as we have used the mechanisms provided by the framework, no modifications are needed to make the analysis inter-procedural. To show this, two additional test cases are constructed:

```
1 | def fun = function (arr : ref<ptr<int<4>>>) -> ptr<int<4>> {
2 |     return ptr_add(*arr, 20);
3 | };
4 |
5 | {
6 |     var ref<array<int<4>,21>> arr = ref_decl(type_lit(ref<array<int<4>,21>>));
7 |     var ref<ptr<int<4>>> ap = fun(ptr_from_array(arr));
8 |     *ptr_subscript(*ap, 10);
9 | }
```

Listing 12: case7.ir

```
1 | def fun = function (arr : ref<ptr<int<4>>>) -> ptr<int<4>> {
2 |     return ptr_add(*arr, 10);
3 | };
4 |
5 | {
6 |     var ref<array<int<4>,21>> arr = ref_decl(type_lit(ref<array<int<4>,21>>));
7 |     var ref<ptr<int<4>>> ap = fun(ptr_from_array(arr));
8 |     *ptr_subscript(*ap, 10);
9 | }
```

Listing 13: case8.ir

Note the different offset in line 2 (both cases).

For the `ref_deref` in line 8 (the first `*` operator in both cases), `IsOutOfBounds` must be returned by the analysis in Case 7, and `IsNotOutOfBounds` in Case 8.

## 6. Evaluation

We also provide two more test cases running the analysis on a scalar, instead of an array.

```
1 | {  
2 |   var ref<int<4>> a = ref_decl(type_lit(ref<int<4>>));  
3 |   var ref<ptr<int<4>>> ap = ptr_from_ref(a);  
4 |   *ptr_subscript(*ap, 42);  
5 | }
```

Listing 14: case9.ir

```
1 | {  
2 |   var ref<int<4>> a = ref_decl(type_lit(ref<int<4>>));  
3 |   var ref<ptr<int<4>>> ap = ptr_from_ref(a);  
4 |   *ptr_subscript(*ap, 0);  
5 | }
```

Listing 15: case10.ir

Case 9 should yield `IsOutOfBounds` while Case 10 should yield `IsNotOutOfBounds`.

As with the initial six cases in Section 5.7, Cases 7–10 are converted to unit tests and executed.

```
> ./ut_analysis_cba_haskell_out_of_bounds_analysis_test --gtest_filter=*AcrossFunctionCall*
Running main() from gtest_main.cc
Note: Google Test filter = *AcrossFunctionCall*
[=====] Running 2 tests from 1 test case.
[-----] Global test environment set-up.
[-----] 2 tests from OutOfBounds
[ RUN      ] OutOfBounds.AcrossFunctionCall
[       OK ] OutOfBounds.AcrossFunctionCall (472 ms)
[ RUN      ] OutOfBounds.AcrossFunctionCallNotOutOfBounds
[       OK ] OutOfBounds.AcrossFunctionCallNotOutOfBounds (286 ms)
[-----] 2 tests from OutOfBounds (758 ms total)

[-----] Global test environment tear-down
[=====] 2 tests from 1 test case ran. (759 ms total)
[ PASSED  ] 2 tests.
```

```
> ./ut_analysis_cba_haskell_out_of_bounds_analysis_test --gtest_filter=*Scalar*
Running main() from gtest_main.cc
Note: Google Test filter = *Scalar*
[=====] Running 2 tests from 1 test case.
[-----] Global test environment set-up.
[-----] 2 tests from OutOfBounds
[ RUN      ] OutOfBounds.Scalar
[       OK ] OutOfBounds.Scalar (456 ms)
[ RUN      ] OutOfBounds.ScalarNotOutOfBounds
[       OK ] OutOfBounds.ScalarNotOutOfBounds (281 ms)
[-----] 2 tests from OutOfBounds (737 ms total)

[-----] Global test environment tear-down
[=====] 2 tests from 1 test case ran. (737 ms total)
[ PASSED  ] 2 tests.
```

The new test cases pass, evincing that the framework indeed transparently integrates support for language features not explicitly covered.

## 6.2. Quantitative Evaluation

For the quantitative evaluation we utilise Insieme’s integration tests and run the out-of-bounds analysis on every `ref_deref` contained. For this, we do not care about the actual result, but about the runtime. This evaluates whether the analysis and framework provide sufficient enough performance.

As the integration tests are input files to the compiler, they are only available as C/C++ sources. We therefore chose to create another driver (`haskell_dumper`) dedicated to reading C/C++ input files and dumping the binary representation of the corresponding **INSIEME Parallel Intermediate REpresentation (INSPIRE)** program. Next, this dump is read by another binary (`insieme-hat`) which then collects all node address of `ref_deref` calls and runs the out-of-bounds analysis on them. To maximise modularity, the part for finding target nodes (in this case calls to `ref_deref`) and running the

## 6. Evaluation

analysis is kept concise in dedicated functions.

```
1 | findTargets :: NodeAddress -> [NodeAddress] -> [NodeAddress]
2 | findTargets addr xs = case getNode addr of
3 |     IR.Node IR.CallExpr _ | Q.isBuiltin (goDown 1 addr) "ref_deref" -> addr : xs
4 |     _ -> xs
5 |
6 | analysis :: NodeAddress -> State SolverState OutOfBoundsResult
7 | analysis addr = do
8 |     state <- get
9 |     let (res, state') = outOfBounds state addr
10 |     put state'
11 |     return res
```

The binary dump is read from stdin and each analysis run uses the SolverState outputted by the previous run.

```
1 | main :: IO ()
2 | main = do
3 |     -- parse binary dump (valid input expected)
4 |     dump <- .getContents
5 |     let Right ir = parseBinaryDump dump
6 |
7 |     let targets = foldTreePrune findTargets Q.isType ir
8 |     let res = evalState (sequence $ analysis <$> targets) initState
9 |
10 |     start <- getCurrentTime
11 |     evaluate $ force res
12 |     end <- getCurrentTime
13 |
14 |     let ms = round $ (*1000) $ toRational $ diffUTCTime end start :: Integer
15 |     printf "%d ms" ms
```

Note that in line 11 `res` is evaluated to normal form using `force` from the `deepseq` package. This is necessary since Haskell features lazy-evaluation and `evaluate` will only evaluate the argument to weak head normal form<sup>1</sup> (ie not fully). `force` on the other hand evaluates the argument completely.

The evaluation is performed on an Intel(R)Xeon(R)CPUE5-2699v3@2.30GHz with 256 GB of memory.

Overall 320 integration tests have been used, containing about  $3.81 \times 10^8$  nodes (not shared) altogether, from which 63 662 are calls to `ref_deref`. To minimise measurement errors, each integration test is processed by `insieme-hat` 10 times from which the arithmetic mean of the outputted execution times is recorded. The result is displayed in Figure 11. It shows that the majority of the test cases can be completed in under 30 ms. The full list of integration tests (with the average execution time of the 10 runs) is provided in Appendix B.

---

<sup>1</sup>See [https://wiki.haskell.org/Weak\\_head\\_normal\\_form](https://wiki.haskell.org/Weak_head_normal_form) for more details.

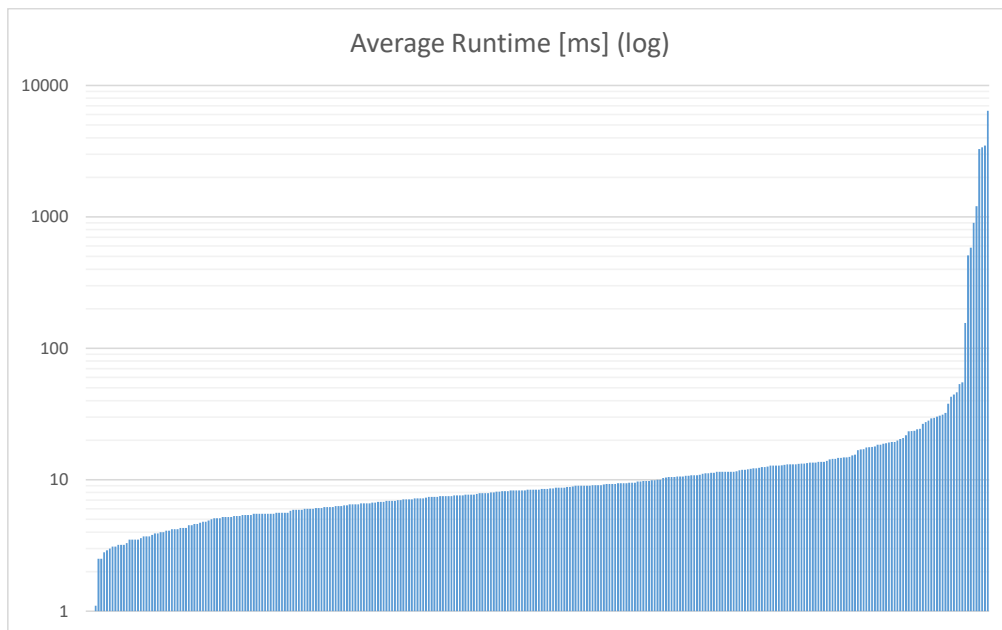


Figure 11.: Runtime of the 320 test cases.

Furthermore, memory consumption has been recorded using GNU time 1.7. Figure 12 displays the recorded Maximum Resident Set Size. The majority of test cases do not consume more than 100 MB of memory.

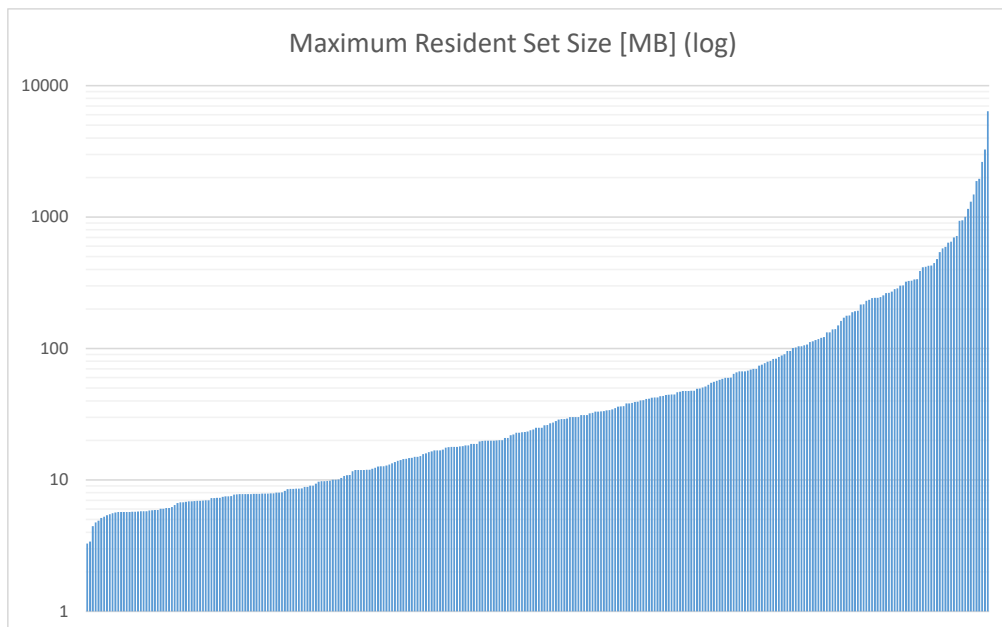


Figure 12.: Memory consumption of the 320 test cases.



## 7. Conclusion

In this thesis a framework for implementing various static program analyses has been presented. It utilises a **Constraint-Based Analysis (CBA)** framework, which is originally derived from the **Data-Flow Analysis (DFA)** framework presented by Nielson et al. [15] and has further been customised by Jordan [10]. This toolkit, created using Haskell, integrates seamlessly into the Insieme compiler, but can also be used as standalone.

It enables developers and researchers to rapidly develop and use specialised static program analyses, for various different kinds of applications (eg auto-tuning or identifying optimisation candidates in a compilation process). In Chapter 5 we have shown that crafting a new analysis only requires a small amount of boiler plate code and some moderate implementation effort. Furthermore we have shown that the framework elevates the capabilities of newly realised analyses while maintaining decent performance.

### 7.1. Future work

While the framework's current state already allows one to use it for various tasks, certain language features of the **INSieme Parallel Intermediate REpresentation (INSPIRE)** are not yet supported. As stated in Section 1.3, this merely builds the foundation upon which development will iterate on in future.

Regarding documentation, the core concept of the framework itself, as well as analyses implemented using it, are quite straight forward to grasp. Yet, further documentation of certain edge cases would be helpful. However, the main challenge encountered by new developers is understanding the different node types, constructs, and language extensions of **INSPIRE**. This issue can be addressed by providing more in-depth documentation on language extensions and their underlying semantics, as well as their connections to other language extensions and the **INSPIRE** core language.

On the framework side, creating a catalogue of convenience functions (mainly queries) for all the different **INSPIRE** constructs would further ease the creation and prototype process of analysis.

Unit-testing an analysis is not that much of a hassle when using the framework together with Insieme. For standalone purposes, however more tooling would be required to be efficient at writing tests.

And last, but certainly not least, debugging techniques and tools are needed. At the current state of development, debugging mainly resorts to using Haskell's `Debug.Trace` and dumping the computed `Assignment`, render it via `graphviz`<sup>1</sup>. A notable improvement would be to have a dedicated viewer for the `Assignment`, in a similar fashion as **INSPYER** to **INSPIRE**.

---

<sup>1</sup><http://www.graphviz.org/>





# A. INSPYER

The INSPYER tool was created to visualise a program represented in Insieme’s *Intermediate Representation (IR)*, also known as the *INSieme Parallel Intermediate REpresentation (INSPIRE)*, in an interactive way. Using the text representation of *INSPIRE*, as seen in Section 3.4, is helpful for inspecting an *INSPIRE* program. Yet it is a bit cumbersome to work with – even when using an editor, like Vim<sup>1</sup>, which supports code folding, jumping to corresponding closing brackets, and allows bookmarking locations. The main idea is to have a *Graphical User Interface (GUI)* application which allows one to collapse subtrees which are not interesting and display useful information for each node in one line (ie row). Navigation functions like *goto* and *search* are available as well.

INSPYER is created as server-less web application using only HTML<sup>2</sup>, CSS<sup>3</sup>, and JavaScript. This allows for it to be run locally using a browser or even host it online<sup>4</sup> – in our case we are using GitHub Pages<sup>5</sup>. Bootstrap<sup>6</sup>, an HTML, CSS, and JavaScript framework is used to create the user interface in little time and to provide a pleasing look and feel. A screenshot of INSPYER can be seen in Figure 13 where Insieme’s FFT<sup>7</sup> integration test has been loaded.

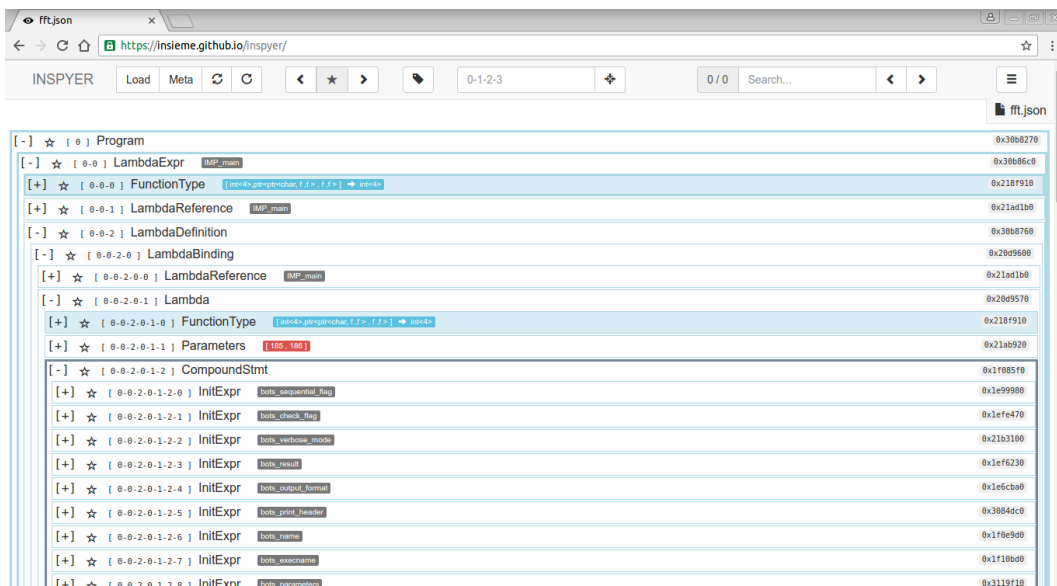


Figure 13.: Screenshot of INSPYER visualising the FFT integration test.

<sup>1</sup><http://www.vim.org>

<sup>2</sup>HyperText Markup Language

<sup>3</sup>Cascading Style Sheets

<sup>4</sup><https://insieme.github.io/inspyer>

<sup>5</sup><https://pages.github.com>

<sup>6</sup><http://getbootstrap.com>

<sup>7</sup>Fast Fourier Transformation

## A.1. JSON Export

An **INSPIRE** program is exported using Insieme’s *JavaScript Object Notation (JSON) dumper*. The format is straight forward and similar to the binary dump. An example is given in Listing 16. The **JSON** object is required to have a field named `root` which identifies the root node. Each **IR** node is attached to the object using a unique key, for simplicity the original memory address is used. The information dumped for one node consists of the node type, its children and (in case its a value node) the value. One can already see in the example, that the list of children is just a list of the memory locations of the child nodes. The memory location serves as a unique identifier for each node. This representation produces node sharing, similar to the binary dump.

```

1 | {
2 |   "root" : "0x30b8270",
3 |   "0x30b8270" : {
4 |     "Kind" : "Program",
5 |     "Children" : ["0x30b86c0"]
6 |   },
7 |   "0x30b86c0" : {
8 |     "Kind" : "LambdaExpr",
9 |     "Children" : ["0x218f910", "0x21ad1b0", "0x30b8760"]
10 |  },
11 |   "0x1f588f0" : {
12 |     "Kind" : "StringValue",
13 |     "Value" : "int",
14 |     "Children" : []
15 |  },
16 |   ...
17 | }

```

Listing 16: Example of an **INSPIRE** JSON dump.

The user loads the **JSON** dump either by selecting it via the corresponding dialogue (using the button labelled *Load* located in the top bar) or by dragging and dropping the file into the main area of **INSPYER**. Doing this will trigger the JavaScript logic, loading the file using `FileReader` and `JSON.parse`. The **JSON** dump is then interpreted. A class `Node`, dedicated for maintaining and working with nodes, has been implemented. The tree visualisation is generated incrementally – upon expanding a node, its children will be created. This keeps the DOM<sup>1</sup> small and simple, which is important to provide adequate performance – even small input programs can yield a big **INSPIRE** program. Having more than 1000 nodes displayed simultaneously results in an unresponsive (ie laggy) user experience. This limitation is inherited by the browser.

## A.2. User Interface

The tree can be navigated either by using the mouse or the keyboard. Clicking on a node will cause it to expand (or collapse if it is already expanded). When using the keyboard, the currently *selected* node is outlined by a thicker, darker border. Hitting  will do the same to the selected node as clicking on it. Other nodes can be selected by using the cursor keys to move the selection.

---

<sup>1</sup>Document Object Model

↑ and ↓ will change the selection on the current level of the tree. ← moves the selection to the parent node. → selects the first child of the currently selected node, expanding it if necessary.

Each line, representing a node, consists of (from left to right):

- an indicator showing whether the node is expanded or collapsed;
- a bookmark button;
- the corresponding node address;
- the node type;
- its **INSPIRE** type<sup>1</sup> (eg `int<4>`);
- its value<sup>1</sup> (eg `9` or `"Hello"`);
- its variable id<sup>1</sup>; and
- the unique identifier in form of the original memory location.

The controls (navigation bar in Figure 13) of INSPYER allow you to (again from left to right):

- load an **INSPIRE** program from a **JSON** dump;
- load meta information from a *meta file* (see Appendix A.3);
- reload the **JSON** dump and meta file (keeping bookmarks and search results);
- reload the **JSON** dump and meta file (clearing bookmarks and search results);
- move the selection back and forth between bookmarked nodes;
- show / hide node addresses – useful deep inside the tree, where node addresses can get unpleasantly long;
- directly jump to a node given a node address;
- search the tree using a **Regular Expression (Regex)**, jumping back and forth between results; and
- show a modal window containing information about hotkeys.

The modal window can be extended easily to hold more, infrequently used functions. There are currently no such functions implemented, but may be needed in the future. An example for this would be statistics about the currently loaded dump.

The search functionality is implemented using web workers<sup>2</sup>. This allows us to search the tree using a dedicated thread – the **GUI** remains responsive even if it takes multiple seconds to search the whole tree. The set of search results will be updated periodically so the user can inspect results matching the query before the search process is completed.

The **Fast Fourier Transform (FFT)** example contains 34 254 nodes and can be fully searched in about 15 seconds – ‘InitExpr’ was used as search string, yielding 25 hits. Note that the number of hits has a huge performance impact as each hit results in a new JavaScript object, which also has to be transferred from the web worker to the main thread. Searching the **FFT** example for ‘StringValue’ yields over 2 500 000 results before the browser kills the main thread (and web worker). This took about two minutes to happen. These tests have been run on an Intel(R) Core(TM) i3-3120M CPU @ 2.50GHz with 16 GB DDR3 Memory @ 1600Hz in Chromium 51.

The web worker API requires an external script file for the thread, but loading the script later on is not possible from a clients filesystem due to security concerns. A workaround is to have the script embedded into the page and create a **Binary Long Object (blob)** from it. This **blob** will then be passed to the worker. Therefore the logic of the search web worker is found in `index.html` at the

<sup>1</sup>If available.

<sup>2</sup>[https://developer.mozilla.org/en-US/docs/Web/API/Web\\_Workers\\_API/Using\\_web\\_workers](https://developer.mozilla.org/en-US/docs/Web/API/Web_Workers_API/Using_web_workers)

## A. INSPYER

bottom, inside a script tag with type `javascript/worker`. It is important to set the type to something different from `text/javascript` to prevent the browser from executing this code right away.

### A.3. Meta Data

The *meta file* is a **JSON** formatted file that is associated with one **JSON** dump of an **INSPIRE** program and holds additional information. Neither what kind of information can be contained nor where that information comes from is fixed. The meta file can:

- bookmark specific nodes, the first bookmark will be jumped to on load;
- expand specific nodes;
- highlight specific nodes;
- attach labels to specific nodes; and
- attach multiline information (including markup), shown when a node is expanded.

As can be seen in Listing 17, node addresses are used to reference the specific nodes.

```
1 | {
2 |   "bookmarks": [
3 |     "0-0-2-0-1-1",
4 |     "0-0-2-0-1-2"
5 |   ],
6 |   "expands": [
7 |     "0-0-2-0-1-2-1-0",
8 |     "0-0-2-0-1-2-6-0",
9 |     "0-0-2-0-1-2-12-0",
10 |    "0-0-2-0-1-2-15-0"
11 |  ],
12 |   "labels": {
13 |     "0-0-2-0-1-2-1": "some information",
14 |     "0-0-2-0-1-2-3": "some other information"
15 |  },
16 |   "highlights": [
17 |     "0-0-2-0-1-2-4",
18 |     "0-0-2-0-1-2-5",
19 |     "0-0-2-0-1-2-6",
20 |     "0-0-2-0-1-2-7",
21 |     "0-0-2-0-1-2-8"
22 |  ],
23 |   "bodies": {
24 |     "0-0-2-0-1-2-6": "Some <b>additional</b> information text"
25 |  }
26 | }
```

Listing 17: An example meta file.

Insieme maintains a helper class for generating meta files. It is currently used in the analysis stress test and the module responsible for semantic checks. The class declaration is seen in Listing 18. One can manually instantiate this class, save the information and export it afterwards. This, on the other hand, requires making the instance available at every point where information needs to be added – and at the location where it should be exported. Because of this, and since there usually only exists

one INSPIRE program per run of the Insieme compiler, one static instance of the class and functions wrapping the default instance's functionality are provided. Listing 19 displays these provided utility functions.

```

1 | class MetaGenerator {
2 |     private:
3 |         NodePtr root;
4 |         std::set<NodeAddress> bookmarks;
5 |         std::set<NodeAddress> expands;
6 |         std::set<NodeAddress> highlights;
7 |         std::map<NodeAddress, std::string> labels;
8 |         std::map<NodeAddress, std::string> bodies;
9 |
10 |         void checkRoot(const NodePtr root);
11 |
12 |     public:
13 |         explicit MetaGenerator(const NodePtr root);
14 |         void addBookmark(const NodeAddress addr);
15 |         void addExpand(const NodeAddress addr);
16 |         void addHighlight(const NodeAddress addr);
17 |         void addLabel(const NodeAddress addr, const std::string label);
18 |         void addBody(const NodeAddress addr, const std::string body);
19 |         void dump(std::ostream& out);
20 | };

```

Listing 18: MetaGenerator utility class declaration.

```

1 | void addBookmark(const NodeAddress addr);
2 |
3 | void addExpand(const NodeAddress addr);
4 |
5 | void addHighlight(const NodeAddress addr);
6 |
7 | void addLabel(const NodeAddress addr, const std::string label);
8 |
9 | void addBody(const NodeAddress addr, const std::string body);
10 |
11 | void dumpTree(std::ostream& out, const NodePtr root);
12 |
13 | void dumpMeta(std::ostream& out);

```

Listing 19: Declarations of functions wrapping the static, default instance of MetaGenerator.



## B. Integration Test List

The following list contains all 320 integration tests used for the quantitative evaluation (see Section 6.2). Each has been run 10 times and the execution time of the out-of-bounds analysis, as well as the overall memory consumption has been recorded. The number in parenthesis states the average time consumed by all 10 runs, followed by the memory consumption.

```
cilk/fib (14.8 ms, 5780 kB)
ocl/ocl_mat_mult (1203.7 ms, 140 736 kB)
ocl/ocl_mat_trans (582.2 ms, 69 052 kB)
ocl/ocl_stencil_3d (899.8 ms, 83 396 kB)
omp/c/atomic (32.2 ms, 7300 kB)
omp/c/autocollapse (155.6 ms, 19 896 kB)
omp/c/backward (509.7 ms, 56 768 kB)
omp/c/barrier (29.6 ms, 6872 kB)
omp/c/barrier_numthreads (26.6 ms, 6868 kB)
omp/c/barrier_simple (37.9 ms, 7880 kB)
omp/c/bots/fib/fib_if (3495.4 ms, 426 680 kB)
omp/c/bots/fib/fib_manual (3389.3 ms, 428 728 kB)
omp/c/bots/nqueens/nqueens_if (3283.2 ms, 418 484 kB)
omp/c/bots/nqueens/nqueens_manual (6425.0 ms, 650 944 kB)
omp/c/capture_be (7.6 ms, 8580 kB)
omp/c/critical (7.2 ms, 6816 kB)
omp/c/datasharing (6.6 ms, 6916 kB)
omp/c/delannoy (13.1 ms, 27 016 kB)
omp/c/dijkstra (24.4 ms, 178 612 kB)
omp/c/even_odd_task (10.7 ms, 20 056 kB)
omp/c/flush (6.5 ms, 7892 kB)
omp/c/func_pointer (8.2 ms, 7508 kB)
omp/c/functions (6.5 ms, 9816 kB)
omp/c/global_firstprivate (6.6 ms, 7772 kB)
omp/c/heated_plate (19.9 ms, 101 760 kB)
omp/c/jacobi (21.8 ms, 79 260 kB)
omp/c/locks (11.2 ms, 12 716 kB)
omp/c/matrix_mul_dyn_1D_omp (14.9 ms, 39 348 kB)
omp/c/matrix_mul_dyn_2D (18.5 ms, 120 760 kB)
omp/c/matrix_mul_static (13.7 ms, 30 116 kB)
omp/c/meta_info_test (3.5 ms, 5644 kB)
omp/c/nested_loops (13.9 ms, 43 444 kB)
omp/c/odd_even_sort (14.8 ms, 40 376 kB)
omp/c/parallel_merge_stress (5.0 ms, 5788 kB)
omp/c/pendulum (53.5 ms, 577 208 kB)
omp/c/private_iterator (10.5 ms, 11 660 kB)
omp/c/privatized_global (7.7 ms, 13 972 kB)
omp/c/qap2 (30.8 ms, 230 856 kB)
omp/c/rec_global (6.1 ms, 5884 kB)
omp/c/reduction (8.3 ms, 13 696 kB)
omp/c/scr/loopsA (30.2 ms, 323 248 kB)
omp/c/scr/loopsB (29.3 ms, 328 364 kB)
omp/c/scr/loopsC (31.3 ms, 335 544 kB)
omp/c/scr/loopsD (28.2 ms, 327 348 kB)
omp/c/scr/lu_reduction (12.5 ms, 36 244 kB)
omp/c/scr/pi (7.2 ms, 8852 kB)
omp/c/set_num_threads (5.4 ms, 7428 kB)
omp/c/simple_pfor (9.7 ms, 8856 kB)
omp/c/single (7.2 ms, 7836 kB)
omp/c/stream/stream_d (23.6 ms, 100 788 kB)
omp/c/void_ptr_capturing (5.2 ms, 12 376 kB)
omp/c/volatile (5.9 ms, 8612 kB)
omp/c/wavefront (10.9 ms, 38 288 kB)
omp/c/wtime (7.5 ms, 8532 kB)
seq/c/anon_enum (5.1 ms, 7312 kB)
seq/c/anonymous (12.3 ms, 42 380 kB)
seq/c/anyref (11.3 ms, 40 324 kB)
seq/c/args (4.9 ms, 6932 kB)
seq/c/arithmetic (7.1 ms, 13 108 kB)
seq/c/assert (11.5 ms, 30 104 kB)
seq/c/assign (5.5 ms, 10 892 kB)
seq/c/assign_expr (11.1 ms, 17 828 kB)
seq/c/backward (12.8 ms, 70 084 kB)
seq/c/bind (4.5 ms, 9056 kB)
seq/c/bots/fib (44.6 ms, 301 744 kB)
seq/c/bots/nqueens (46.3 ms, 479 684 kB)
seq/c/bugs/anon_multi_tu (6.9 ms, 11 900 kB)
seq/c/bugs/basename_bug (3.6 ms, 5692 kB)
seq/c/bugs/char_ops (2.8 ms, 6748 kB)
seq/c/bugs/extern_array (7.4 ms, 20 088 kB)
seq/c/bugs/for_loops (5.2 ms, 7904 kB)
seq/c/bugs/function_ptr (18.8 ms, 105 880 kB)
seq/c/bugs/global_func_ptr (7.4 ms, 12 156 kB)
seq/c/bugs/global_rec_struct (10.4 ms, 47 520 kB)
seq/c/bugs/global_rec_struct_typedef (10.6 ms, 47 512 kB)
seq/c/bugs/hex (4.3 ms, 8264 kB)
seq/c/bugs/long_long_replace (6.8 ms, 11 892 kB)
seq/c/bugs/multiple_tu_struct_decl (9.1 ms, 20 876 kB)
seq/c/bugs/null_function_ptr (12.1 ms, 57 744 kB)
seq/c/bugs/ptrptr_bug (7.8 ms, 19 868 kB)
seq/c/bugs/rec_data_struct (5.4 ms, 49 360 kB)
seq/c/bugs/rec_func_param_type (13.1 ms, 66 980 kB)
seq/c/bugs/rec_func_ptr (3.5 ms, 8580 kB)
seq/c/bugs/rec_struct_typedef (8.6 ms, 47 488 kB)
seq/c/bugs/scalar_pointer (9.8 ms, 17 808 kB)
seq/c/bugs/static_ptr (6.0 ms, 6236 kB)
seq/c/bugs/string_escape_chars (4.0 ms, 7744 kB)
seq/c/bugs/switch_cases (4.7 ms, 7816 kB)
seq/c/bugs/void_ptr (8.7 ms, 17 812 kB)
seq/c/c/comma (1.1 ms, 5240 kB)
seq/c/c/long_long (4.6 ms, 8000 kB)
seq/c/capture_global (8.3 ms, 12 708 kB)
seq/c/comma (7.1 ms, 14 448 kB)
seq/c/compound_operators (10.6 ms, 29 096 kB)
seq/c/conditional_strings (4.8 ms, 9784 kB)
seq/c/delannoy (9.3 ms, 28 816 kB)
seq/c/demo/loop (5.5 ms, 7872 kB)
seq/c/demo/loop_parallel (5.9 ms, 7864 kB)
seq/c/enum (17.6 ms, 23 420 kB)
seq/c/even_odd (5.8 ms, 24 948 kB)
seq/c/expressions (6.6 ms, 10 044 kB)
seq/c/ext_multi (3.2 ms, 5696 kB)
seq/c/extended_for_test (13.0 ms, 16 784 kB)
seq/c/extern_global (1.0 ms, 4456 kB)
seq/c/external (3.1 ms, 5696 kB)
seq/c/external_type (6.4 ms, 8012 kB)
seq/c/external_value (4.2 ms, 5720 kB)
seq/c/file_io (8.4 ms, 27 280 kB)
seq/c/for_loops (19.4 ms, 217 524 kB)
seq/c/func_pointer (15.5 ms, 65 980 kB)
seq/c/func_pointer_simple (7.5 ms, 11 968 kB)
seq/c/gettimeofday (8.3 ms, 12 880 kB)
seq/c/global_data (6.2 ms, 9380 kB)
seq/c/global_init (1.0 ms, 3284 kB)
seq/c/global_multi (6.3 ms, 7288 kB)
seq/c/global_nested (6.2 ms, 9888 kB)
seq/c/health_struct (10.5 ms, 77 460 kB)
seq/c/hello_world (3.1 ms, 5584 kB)
```

## B. Integration Test List

seq/c/indirection (8.0 ms, 19 888 kB)  
seq/c/initialization (6.5 ms, 18 780 kB)  
seq/c/instrumentation\_region (5.5 ms, 7844 kB)  
seq/c/jacobi (13.3 ms, 66 988 kB)  
seq/c/literals (5.1 ms, 7804 kB)  
seq/c/loop\_transform (42.8 ms, 446 120 kB)  
seq/c/loops (17.9 ms, 55 732 kB)  
seq/c/lu\_decomposition (13.1 ms, 58 820 kB)  
seq/c/matrix\_mul\_dyn\_1d (13.4 ms, 35 240 kB)  
seq/c/matrix\_mul\_dyn\_2d (17.7 ms, 90 552 kB)  
seq/c/matrix\_mul\_dynamic\_size (9.3 ms, 32 180 kB)  
seq/c/matrix\_mul\_recursive (11.5 ms, 177 600 kB)  
seq/c/matrix\_mul\_static (9.7 ms, 30 124 kB)  
seq/c/matrix\_transpose (8.0 ms, 11 880 kB)  
seq/c/mem\_alloc (10.6 ms, 19 644 kB)  
seq/c/minMax (4.2 ms, 6996 kB)  
seq/c/multi\_tu (3.9 ms, 5720 kB)  
seq/c/multi\_tu\_mutual (6.1 ms, 23 924 kB)  
seq/c/mutual\_recursion\_global (7.9 ms, 51 396 kB)  
seq/c/n\_body (27.4 ms, 254 356 kB)  
seq/c/param\_passing (8.4 ms, 16 800 kB)  
seq/c/pendulum (55.2 ms, 639 676 kB)  
seq/c/pointer\_arith (11.9 ms, 44 700 kB)  
seq/c/pointers (11.8 ms, 42 372 kB)  
seq/c/ptr\_assign (9.5 ms, 47 792 kB)  
seq/c/pragma\_match (1.0 ms, 3404 kB)  
seq/c/pthread/single\_thread (9.0 ms, 10 340 kB)  
seq/c/ptr\_offset (12.2 ms, 50 600 kB)  
seq/c/qap (23.4 ms, 171 464 kB)  
seq/c/quicksort (14.4 ms, 116 156 kB)  
seq/c/rec\_nasty (5.1 ms, 7852 kB)  
seq/c/rec\_struct (8.4 ms, 33 180 kB)  
seq/c/recursions (7.2 ms, 44 400 kB)  
seq/c/recursions2 (8.2 ms, 18 824 kB)  
seq/c/ref\_passing (6.2 ms, 10 064 kB)  
seq/c/ref\_vector\_array\_pointer (14.4 ms, 33 164 kB)  
seq/c/sigrt (6.7 ms, 11 896 kB)  
seq/c/simple\_jacobi (14.7 ms, 64 180 kB)  
seq/c/simple\_kernel (5.3 ms, 8032 kB)  
seq/c/simple\_recursion (5.9 ms, 10 852 kB)  
seq/c/sizeof (3.9 ms, 6444 kB)  
seq/c/stack (17.1 ms, 74 160 kB)  
seq/c/static\_init (7.5 ms, 12 676 kB)  
seq/c/stencil3d (18.5 ms, 112 072 kB)  
seq/c/structs (7.5 ms, 11 932 kB)  
seq/c/tiling/dsyk (12.0 ms, 41 400 kB)  
seq/c/tiling/dtrmm (9.0 ms, 31 152 kB)  
seq/c/transpose (23.5 ms, 246 220 kB)  
seq/c/trunc (4.3 ms, 6964 kB)  
seq/c/union (8.8 ms, 33 144 kB)  
seq/c/variable\_arrays (10.7 ms, 20 868 kB)  
seq/c/variable\_sized\_struct (19.2 ms, 191 932 kB)  
seq/c/vector\_partial\_init (9.2 ms, 31 092 kB)  
seq/c/vectors (8.6 ms, 19 832 kB)  
seq/c/void\_ptr (7.0 ms, 18 284 kB)  
seq/cpp/adv\_op\_over1 (24.2 ms, 2 630 244 kB)  
seq/cpp/argument\_types (9.4 ms, 287 128 kB)  
seq/cpp/assign\_expr (11.5 ms, 25 000 kB)  
seq/cpp/base\_of (7.7 ms, 13 412 kB)  
seq/cpp/boost/regex (12.8 ms, 43 640 kB)  
seq/cpp/boost/shared\_ptr (14.7 ms, 6 384 616 kB)  
seq/cpp/bugs/allocate\_matrix (8.7 ms, 16 760 kB)  
seq/cpp/bugs/boost/boost\_array (2.9 ms, 5136 kB)  
seq/cpp/bugs/boost/boost\_regex\_search (13.7 ms, 88 696 kB)  
seq/cpp/bugs/boost/def\_initialization (12.9 ms, 1 155 720 kB)  
seq/cpp/bugs/boost/global\_intercepted (8.5 ms, 16 992 kB)  
seq/cpp/bugs/case\_default (9.9 ms, 41 588 kB)  
seq/cpp/bugs/const\_obj\_init (12.8 ms, 122 212 kB)  
seq/cpp/bugs/const\_obj\_param (11.5 ms, 389 496 kB)  
seq/cpp/bugs/const\_ref\_arg (12.6 ms, 937 568 kB)  
seq/cpp/bugs/ctor\_base\_ptr (11.2 ms, 1 958 492 kB)  
seq/cpp/bugs/ctor\_inherit (9.4 ms, 414 352 kB)  
seq/cpp/bugs/ctor\_ptr\_overload (8.7 ms, 95 904 kB)  
seq/cpp/bugs/duplicated\_ctor (7.1 ms, 44 732 kB)  
seq/cpp/bugs/empty\_pod (7.5 ms, 32 372 kB)  
seq/cpp/bugs/error\_handler (12.8 ms, 14 172 kB)  
seq/cpp/bugs/ewc\_capturing (6.2 ms, 9056 kB)  
seq/cpp/bugs/expr\_with\_cleanups (11.6 ms, 54 900 kB)  
seq/cpp/bugs/expression\_with\_cleanups (6.6 ms, 699 028 kB)  
seq/cpp/bugs/gettimeofday (7.4 ms, 8528 kB)  
seq/cpp/bugs/global\_init (9.5 ms, 46 944 kB)  
seq/cpp/bugs/globals\_dtor (9.1 ms, 47 652 kB)  
seq/cpp/bugs/init\_ptr (9.4 ms, 52 864 kB)  
seq/cpp/bugs/init\_struct (9.4 ms, 18 060 kB)  
seq/cpp/bugs/intlit (3.2 ms, 5712 kB)  
seq/cpp/bugs/just\_decl\_template (5.5 ms, 6124 kB)  
seq/cpp/bugs/long\_long\_thing (11.9 ms, 14 996 kB)  
seq/cpp/bugs/loops (8.5 ms, 31 164 kB)  
seq/cpp/bugs/malformed\_logical\_operator (7.4 ms, 1 483 392 kB)  
seq/cpp/bugs/map\_iterator (7.0 ms, 23 104 kB)  
seq/cpp/bugs/mem\_alloc\_ctor (10.3 ms, 1 003 152 kB)  
seq/cpp/bugs/member\_function\_as\_field (6.0 ms, 17 744 kB)  
seq/cpp/bugs/merge\_bugs/ctor\_ptr (5.9 ms, 594 452 kB)  
seq/cpp/bugs/new\_pod (7.6 ms, 22 180 kB)  
seq/cpp/bugs/non\_default\_global (8.3 ms, 22 888 kB)  
seq/cpp/bugs/rec\_struct (6.9 ms, 14 976 kB)  
seq/cpp/bugs/ref\_to\_pointer (5.5 ms, 8552 kB)  
seq/cpp/bugs/reference\_in\_condition\_expression (6.3 ms, 29 304 kB)  
seq/cpp/bugs/return\_overload (10.0 ms, 113 264 kB)  
seq/cpp/bugs/return\_string (4.2 ms, 5744 kB)  
seq/cpp/bugs/static\_cast\_in\_return (6.3 ms, 39 244 kB)  
seq/cpp/bugs/static\_data\_member\_in\_one\_tu (8.2 ms, 36 124 kB)  
seq/cpp/bugs/static\_member (8.3 ms, 16 484 kB)  
seq/cpp/bugs/strings\_array (7.9 ms, 7520 kB)  
seq/cpp/bugs/temp\_dtor (10.8 ms, 95 608 kB)  
seq/cpp/bugs/template\_integrals (9.0 ms, 16 280 kB)  
seq/cpp/bugs/template\_integrals\_intercept (7.7 ms, 14 440 kB)  
seq/cpp/bugs/templated\_casts (5.6 ms, 300 572 kB)  
seq/cpp/bugs/this (12.5 ms, 150 228 kB)  
seq/cpp/bugs/union (3.7 ms, 9796 kB)  
seq/cpp/bugs/union\_member\_func (9.0 ms, 18 796 kB)  
seq/cpp/bugs/vector\_pointers (5.5 ms, 7284 kB)  
seq/cpp/bugs/zero\_ptr (5.6 ms, 18 024 kB)  
seq/cpp/class (9.0 ms, 86 372 kB)  
seq/cpp/class\_def\_and\_decl\_splitup (9.3 ms, 162 136 kB)  
seq/cpp/const\_pointer (8.5 ms, 15 208 kB)  
seq/cpp/const\_ptr\_param (7.1 ms, 132 720 kB)  
seq/cpp/const\_ref\_arguments (3.8 ms, 19 944 kB)  
seq/cpp/counter (5.6 ms, 243 356 kB)  
seq/cpp/cpp11/auto (5.5 ms, 5860 kB)  
seq/cpp/cpp11/constexpr (8.3 ms, 6944 kB)  
seq/cpp/cpp11/cpp11enums (3.7 ms, 6660 kB)  
seq/cpp/cpp11/decltype (11.5 ms, 234 900 kB)  
seq/cpp/cpp11/declval (3.3 ms, 5704 kB)  
seq/cpp/cpp11/lambda\_to\_funptr (13.1 ms, 24 960 kB)  
seq/cpp/cpp11/range\_for (12.2 ms, 49 608 kB)  
seq/cpp/cpp11/rec\_lambda (13.5 ms, 38 516 kB)  
seq/cpp/cpp11/sizeofpack (3.5 ms, 5768 kB)  
seq/cpp/cpp11/std\_array (2.5 ms, 4908 kB)  
seq/cpp/cpp11/variadic\_templates (7.6 ms, 15 716 kB)  
seq/cpp/ctor\_inherit (14.3 ms, 717 224 kB)  
seq/cpp/ctor\_return (13.3 ms, 3 270 284 kB)  
seq/cpp/dtor (6.1 ms, 215 672 kB)  
seq/cpp/enum (16.8 ms, 23 216 kB)  
seq/cpp/for\_loop (3.2 ms, 5380 kB)  
seq/cpp/friend (6.9 ms, 104 048 kB)  
seq/cpp/func (3.7 ms, 6128 kB)  
seq/cpp/global\_const (4.5 ms, 7512 kB)  
seq/cpp/globals\_in\_classes (9.1 ms, 242 080 kB)  
seq/cpp/header (6.8 ms, 139 888 kB)  
seq/cpp/hello\_stream (3.5 ms, 6764 kB)  
seq/cpp/hello\_world (2.5 ms, 4748 kB)  
seq/cpp/identical\_class\_structure (6.0 ms, 17 620 kB)  
seq/cpp/includes (7.6 ms, 6044 kB)  
seq/cpp/inherit (6.7 ms, 264 620 kB)  
seq/cpp/inherit\_multiple (6.4 ms, 540 308 kB)  
seq/cpp/inherit\_ptr (6.9 ms, 33 396 kB)  
seq/cpp/inherit\_single (8.9 ms, 59 748 kB)  
seq/cpp/inherit\_single\_dynmem (11.5 ms, 188 780 kB)  
seq/cpp/inherit\_virtual (11.5 ms, 337 488 kB)  
seq/cpp/initialization (10.5 ms, 66 924 kB)  
seq/cpp/isomorphic (9.8 ms, 80 288 kB)  
seq/cpp/long\_long (8.4 ms, 42 328 kB)  
seq/cpp/materialize\_with\_bind\_expr (13.7 ms, 283 284 kB)  
seq/cpp/mem\_alloc (8.7 ms, 44 432 kB)  
seq/cpp/mem\_alloc\_ctor (13.2 ms, 943 764 kB)  
seq/cpp/mem\_alloc\_ctor\_inherit (20.8 ms, 1 310 368 kB)  
seq/cpp/member\_func (5.3 ms, 36 372 kB)  
seq/cpp/member\_func\_overload (4.6 ms, 26 056 kB)



seq/cpp/meta\_fib (5.4 ms, 5896 kB)  
seq/cpp/multiple\_tu (6.0 ms, 46460 kB)  
seq/cpp/naive (8.8 ms, 70032 kB)  
seq/cpp/nameless\_struct (7.7 ms, 33904 kB)  
seq/cpp/nested\_class (6.8 ms, 243068 kB)  
seq/cpp/new\_string (5.2 ms, 7792 kB)  
seq/cpp/object\_array\_new (8.1 ms, 24208 kB)  
seq/cpp/object\_bug (4.8 ms, 14652 kB)  
seq/cpp/op\_over1 (15.3 ms, 1880756 kB)  
seq/cpp/pairs (5.4 ms, 7816 kB)  
seq/cpp/partial\_interception (4.1 ms, 10744 kB)  
seq/cpp/pod\_struct (9.0 ms, 18280 kB)  
seq/cpp/pod\_struct\_assignment (9.1 ms, 15932 kB)  
seq/cpp/pod\_struct\_bug (11.3 ms, 28032 kB)  
seq/cpp/pod\_struct\_return (4.0 ms, 9688 kB)  
seq/cpp/ptr\_offset (8.4 ms, 29096 kB)  
seq/cpp/rec\_class (8.1 ms, 103804 kB)  
seq/cpp/rec\_memberfunc (5.6 ms, 60036 kB)  
seq/cpp/refs (4.3 ms, 5908 kB)  
seq/cpp/refs\_ctor (7.9 ms, 22896 kB)  
seq/cpp/refs\_obj (7.9 ms, 59760 kB)  
seq/cpp/return (13.5 ms, 193908 kB)

seq/cpp/return\_types (9.3 ms, 132664 kB)  
seq/cpp/scalar\_value\_init\_expr (5.6 ms, 10068 kB)  
seq/cpp/scope (10.8 ms, 33948 kB)  
seq/cpp/scoped\_enum (19.0 ms, 21928 kB)  
seq/cpp/simple\_sort (17.0 ms, 30108 kB)  
seq/cpp/sort (20.4 ms, 263820 kB)  
seq/cpp/static\_method (7.3 ms, 6032 kB)  
seq/cpp/std/chrono (9.8 ms, 14752 kB)  
seq/cpp/std/distance (5.2 ms, 7804 kB)  
seq/cpp/stl\_vector\_iterator (4.1 ms, 6996 kB)  
seq/cpp/struct (17.8 ms, 107144 kB)  
seq/cpp/template1 (13.5 ms, 34444 kB)  
seq/cpp/template2 (9.9 ms, 67944 kB)  
seq/cpp/template3 (19.4 ms, 83604 kB)  
seq/cpp/template\_assignment\_op (5.5 ms, 26220 kB)  
seq/cpp/this (10.0 ms, 118136 kB)  
seq/cpp/union\_member (10.8 ms, 75412 kB)  
seq/cpp/user\_defined\_cast (9.5 ms, 38296 kB)  
seq/cpp/vector\_ctor (6.5 ms, 19912 kB)  
seq/cpp/vector\_push\_back (3.0 ms, 5476 kB)  
seq/cpp/xvalue (5.3 ms, 270648 kB)



# Bibliography

- [1] Erlang programming language. URL <http://www.erlang.org/>. [Online; accessed 2 February 2017].
- [2] F# software foundation. URL <http://fsharp.org/>. [Online; accessed 2 February 2017].
- [3] Haskell language. URL <https://www.haskell.org/>. [Online; accessed 2 February 2017].
- [4] Ocaml. URL <http://ocaml.org/>. [Online; accessed 2 February 2017].
- [5] The scala programming language. URL <http://scala-lang.org/>. [Online; accessed 2 February 2017].
- [6] Craig Chambers. Slides: Lattice-theory data flow analysis, 2006. URL <https://courses.cs.washington.edu/courses/cse501/06wi/slides/01-18.slides.pdf>. [Online; accessed 1 December 2016].
- [7] Stephen Chong. Slides: Dataflow analysis, 2011. URL <http://www.seas.harvard.edu/courses/cs252/2011sp/slides/Lec02-Dataflow.pdf>. [Online; accessed 1 December 2016].
- [8] Insieme Team. Insieme compiler project, 2017. URL <http://insieme-compiler.org/>. [Online; accessed 2 February 2017].
- [9] Mark P Jones. Type classes with functional dependencies. In *European Symposium on Programming*, pages 230–244. Springer, 2000.
- [10] Herbert Jordan. *Insieme — A Compiler Infrastructure for Parallel Programs*. PhD thesis, University of Innsbruck, 2014. URL [http://www.dps.uibk.ac.at/~csaf7445/pub/phd\\_thesis\\_jordan.pdf](http://www.dps.uibk.ac.at/~csaf7445/pub/phd_thesis_jordan.pdf). [Online; accessed 1 December 2016].
- [11] Herbert Jordan, Simone Pellegrini, Peter Thoman, Klaus Kofler, and Thomas Fahringer. Inspire: The insieme parallel intermediate representation. In *Proceedings of the 22nd international conference on Parallel architectures and compilation techniques*, pages 7–18. IEEE Press, 2013.
- [12] Nenad Jovanovic. Lattice tutorial, 2005. URL [http://old.iseclab.org/people/enji/infosys/lattice\\_tutorial.pdf](http://old.iseclab.org/people/enji/infosys/lattice_tutorial.pdf). [Online; accessed 1 December 2016].
- [13] Ryan Lemmer. *Haskell Design Patterns*. Packt Publishing Ltd, 2015. ISBN 9781783988723.
- [14] Miran Lipovaca. *Learn You a Haskell for Great Good!: A Beginner’s Guide*. No Starch Press, San Francisco, CA, USA, 1st edition, 2011. ISBN 1593272839, 9781593272838. URL <http://learnyouahaskell.com>.
- [15] Flemming Nielson, Hanne R. Nielson, and Chris Hankin. *Principles of Program Analysis*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1999. ISBN 3540654100.

## Bibliography

- [16] Jesper Nordenberg. My take on haskell vs scala, 2012. URL <http://jnordenberg.blogspot.co.at/2012/05/my-take-on-haskell-vs-scala.html>. [Online; accessed 1 December 2016].
- [17] Bryan O’Sullivan, John Goerzen, and Don Stewart. *Real World Haskell*. O’Reilly Media, Inc., 1st edition, 2008. ISBN 0596514980, 9780596514983.
- [18] Michelle Strout. Slides: Lattice theoretic framework for dfa, 2006. URL <http://www.cs.colostate.edu/~mstrout/CS553Fall06/slides/lecture10-lattice>. [Online; accessed 1 December 2016].
- [19] Samuli Thomasson. *Haskell High Performance Programming*. Packt Publishing Ltd, 2016. ISBN 9781786464217.
- [20] Wikipedia. Lattice (order), 2016. URL [https://en.wikipedia.org/w/index.php?title=Lattice\\_\(order\)&oldid=754335711](https://en.wikipedia.org/w/index.php?title=Lattice_(order)&oldid=754335711). [Online; accessed 12 December 2016].
- [21] Wikipedia. Reaching definition, 2016. URL [https://en.wikipedia.org/w/index.php?title=Reaching\\_definition&oldid=723830514](https://en.wikipedia.org/w/index.php?title=Reaching_definition&oldid=723830514). [Online; accessed 5 June 2016].
- [22] Wikipedia. Control flow graph, 2017. URL [https://en.wikipedia.org/w/index.php?title=Control\\_flow\\_graph&oldid=758176946](https://en.wikipedia.org/w/index.php?title=Control_flow_graph&oldid=758176946). [Online; accessed 3 January 2017].
- [23] Wikipedia. Hasse diagram, 2017. URL [https://en.wikipedia.org/w/index.php?title=Hasse\\_diagram&oldid=757730257](https://en.wikipedia.org/w/index.php?title=Hasse_diagram&oldid=757730257). [Online; accessed 31 January 2017].

# Glossary

**ADT** Algebraic Data Type. 2, 23, 40, 94

**AST** Abstract Syntax Tree. 8, 11, 13, 31

**blob** Binary Long Object. 113

**CBA** Constraint-Based Analysis. 2, 3, 5, 12, 13, 109

**CFG** Control-Flow Graph. 5–13

**DAG** Direct Acyclic Graph. 18

**DFA** Data-Flow Analysis. 1–3, 5, 6, 8, 9, 12, 13, 34, 60–62, 66–68, 70, 74, 81, 85, 88, 94, 98, 109

**ELF** Executable and Linkable Format. 39

**FFI** Foreign Function Interface. 39–41, 49, 53, 72, 75, 96, 101

**FFT** Fast Fourier Transform. 113

**GHC** Glasgow Haskell Compiler. 37, 38, 41

**GMP** GNU Multiple Precision. 38

**GUI** Graphical User Interface. 111, 113

**Haskell RTS** Haskell RunTime System. 40, 41

**HAT** Haskell-based Analysis Toolkit. v, 1, 3, 12, 13, 18, 19, 33, 38, 39, 87, 101

**IMP** Insieme Mangling Prefix. 28

**INSPIRE** INsieme Parallel Intermediate REpresentation. 2, 3, 13, 17–20, 22–24, 26, 28, 30, 34, 37, 40–45, 47, 49, 52, 58, 62–64, 66, 67, 69, 74, 79, 81, 82, 87, 89, 91, 92, 97, 105, 109, 111–115

**IR** Intermediate Representation. 1, 2, 18, 28, 41, 43, 48, 49, 68, 89, 111, 112

**JSON** JavaScript Object Notation. 18, 112–114

*Glossary*

**RegEx** Regular Expression. 113

**TH** Template Haskell. 42