

Polyhedral AST generation is more than scanning polyhedra

GROSSER TOBIAS, INRIA and ÉCOLE NORMALE SUPÉRIEURE
VERDOOLAEGE SVEN, INRIA, ÉCOLE NORMALE SUPÉRIEURE and KU Leuven¹
COHEN ALBERT, INRIA and ÉCOLE NORMALE SUPÉRIEURE

Abstract mathematical representations such as integer polyhedra have shown to be useful to precisely analyze computational kernels and to express complex loop transformations. Such transformations rely on Abstract Syntax Tree (AST) generators to convert the mathematical representation back to an imperative program. Such generic AST generators avoid the need to resort to transformation-specific code generators, which may be very costly or technically difficult to develop as transformations become more complex. Existing AST generators have proven their effectiveness, but they hit limitations in the most complex scenarios. Specifically, (1) they do not support or may fail to generate control flow for complex transformations using piecewise schedules or mappings involving modulo arithmetic, (2) they offer limited support to specialize the generated code exposing compact, straightline, vectorizable kernels with high arithmetic intensity necessary to exploit the peak performance of modern hardware, (3) they offer no support for memory layout transformations, (4) they provide insufficient control over the AST generation strategy, preventing their application to complex domain specific optimizations.

We present a new AST generation approach that extends classical polyhedral scanning to the full generality of Presburger arithmetic, including existentially quantified variables and piecewise schedules, and that complements it with new optimizations for the detection of components and shifted strides. Not limiting ourselves to control flow generation, we expose functionality to generate AST expressions from arbitrary piecewise quasi-affine expressions which enables the use of our AST generator for data-layout transformations. We complement this with support for specialization by polyhedral unrolling, user-directed versioning and specialization of AST expressions according to the location they are generated at, and complete this work with fine-grained user control over the AST generation strategies used. Using this generalized idea of AST generation, we present how to implement complex domain specific transformations without the need to write specialized code generators, but instead relying on a generic AST generator parameterized to a specific problem domain.

Categories and Subject Descriptors: D.3.4 [Programming Languages]: Processor – Compilers, Optimizations

General Terms: Algorithms, Performance

Additional Key Words and Phrases: Polyhedral compilation, code generation, unrolling, index set splitting, Presburger relations

ACM Reference Format:

ACM Trans. Program. Lang. Syst. V, N, Article A (January YYYY), 34 pages.

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

1. INTRODUCTION

The development of high-level optimizations for domain specific or general purpose compilers can often be conceptually divided into two parts, the design of a high-level

¹Part of this work was performed as a consultant for LIACS, Leiden.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© YYYY ACM 0164-0925/YYYY/01-ARTA \$15.00

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

optimization strategy and the generation of program code according to this optimization strategy. In many cases, the interesting scientific contribution is the new optimization strategy. However, in practice, significant efforts are put into the generation of efficient program code.

For programs with affine (and even non-affine [Venkat et al. 2014]) control flow, it is common to automatically generate optimized program code from an abstract description [Loechner and Wilde 1997; Pugh and Wonnacott 1994b; Pugh and Wonnacott 1994a] of both the program and the intended transformation. Optimizations [Bondhugula et al. 2008; Kong et al. 2013; Zuo et al. 2013; Bandishti et al. 2012] are described by modifying an abstract schedule that defines the execution order of the individual computations in a program. According to this schedule, imperative program code is (re)generated using a technique called polyhedral scanning [Chen 2012], code generation [Bastoul 2004], or more accurately, Abstract Syntax Tree (AST) generation. Decoupling the optimization and program generation steps not only reduces the time needed to implement a certain optimization strategy, but it also speeds up the evaluation of new optimization strategies [Pouchet et al. 2007; Pouchet et al. 2008]. Furthermore, being able to describe transformations on a highly abstract level enables the development of complex transformations [Grosser et al. 2014] relying on the AST generator to generate efficient imperative code.

Even though AST generators have many benefits, existing approaches focus on control flow generation [Bastoul 2004; Kelly et al. 1996; Chen 2012], provide only rudimentary support for the specialization of the generated expressions, and limited control over code size vs. control overhead. These limitations often prevent their wider usage. Missing support to generate user-provided AST expressions, e.g., to describe memory locations, prevents their application to data-layout transformations [Henretty et al. 2013] or when mapping data to software managed caches [Holewinski et al. 2012]. Also, existing AST generators may produce multiple code versions according to specific parameter values, to reduce control overhead, but existing approaches do not natively support the generation of specialized code. This would be particularly helpful for the separation of full and partial tiles [Ancourt and Irigoien 1991; Goumas et al. 2003; Kim et al. 2007] or to handle boundary conditions separately. Instead, versioning has to be enforced by generating several distinct copies of each statement in the input description [Chen et al. 2008]. Similarly, performing unrolling during AST generation is only possible by duplicating statements in the input description [Chen et al. 2008; Bondhugula et al. 2008; Shirako and Sarkar 2013]. Besides being conceptually unsatisfying, duplicating statements causes serious problems. First, by purposefully hiding the fact that statements are identical, the AST generator is forced to generate duplicate code for them in all cases, missing redundancies in complex expressions and missing opportunities to factor colder parts of the code. Secondly, duplicating statements increases the complexity of polyhedral operations involved in the generation of imperative control flow, supporting optimizations such as full/partial tile separation, and supporting expression specialization or simplification for modulo arithmetic. If we now wish to minimize code size for colder parts of the iteration space (e.g., the partial tiles), we run into the next limitation. Even though AST generators provide basic control over the desirable aggressiveness in separating statements or control flow specialization (conditional hoisting), the level of control is way too coarse-grained in existing methods and tools. Also, no guarantees are given about the maximal number of loop nests and the maximal number of statements generated, which is problematic for scenarios where code size is a major concern, such as AST generation for many-core targets with software-managed caches, embedded processors, and high-level synthesis [Zuo et al. 2013]. Overall, existing approaches and tools are in many ways not yet mature for complex AST generation problems.

This work presents an integrated AST generation approach² that besides classical control flow generation, allows the generation of AST expressions from arbitrary user-provided piecewise affine expressions. We define a fine-grained “option” mechanism that enables the user to request maximal specialization where needed while retaining control over code size. To enable aggressive specialization, we allow the user to instruct the AST generator how to version the code, we provide an integrated polyhedral unrolling facility, and we make sure that AST expressions are specialized according to the context they are generated in. Doing so is essential to correctly model the floor-division and modulo arithmetic arising from abstract transformations of the program, and to cast these expressions to efficient remainder and integer divisions, or to lower-complexity operations, as provided by existing instruction set architectures and programming languages. Finally, we present a miscellaneous set of optimizations that improve the quality of the generated code, in comparison to existing polyhedral scanning tools, but also optimizations made necessary to cover the wider application scenarios of our AST generator.

Our contributions are as follows:

- AST generation with complete support for *Presburger relations* including support for *piecewise schedules*, and their use to express *index set splitting as a schedule-only transformation*.
- An aggressive simplification of *AST expressions* generated from *piecewise quasi-affine expressions* within the context of their position in the AST, including the *detection of modulo and division operations*. The generation of simplified AST expressions is not only used to construct loop bounds and if-conditionals from within the AST generator, but is also exposed to the user, who can use this functionality to generate custom index expressions.
- *Fine-grained options* to control AST generation including an *atomic* option which can be used to control code size and to ensure that no program statements are duplicated.
- *Specialization* through *polyhedral unrolling* and *user directed versioning*, in particular the user can specify a subset of the schedule space (e.g., full tiles) that should be isolated from the rest of the schedule space (i.e., partial tiles).
- Algorithms for *improved stride detection* and *detection of reorderable components*.
- AST generation for structured schedules expressed as *schedule trees*.
- Evaluation in an advanced domain-specific optimizer and comparison to state-of-the-art code generation techniques.

The remaining content of this paper is organized as follows. Section 2 gives a high-level overview of our new AST generation approach and presents new, illustrative use cases. We then present theoretical background in Section 3, the data structures involved in Section 4, our core AST generation approach in Section 5 and its extension to schedule trees in Section 6. We finish with a set of experiments in Section 7, the discussion of related work in Section 8 and the conclusion in Section 9.

2. A NEW APPROACH TO AST GENERATION

To give an idea of the new AST generation concepts proposed in this work, we present them in the context of a complex AST generation scenario. One ideal such scenario is our recent work on hexagonal/parallelogram tiling [Grosser et al. 2014], a domain specific optimization to generate efficient CUDA code for iterative stencil computations. It is implemented on top of PPCG [Verdoolaeghe et al. 2013], a generic C to CUDA trans-

² We implemented this approach as part of `isl`. The majority of what is described in this paper is already publicly available at <http://repo.or.cz/w/isl.git>. Certain recent additions have not yet been made public, but will be made public in the next releases of `isl`.

lator, which uses integer polyhedra to describe both the computation itself and the program transformation to apply. Taking advantage of such an abstract mathematical description, a new tiling scheme is developed that involves complex geometric shapes to address the most important performance issues of compiling iterative stencils for GPUs, including the usage of shared memory, the optimization of data-transfers, the increase of arithmetic intensity, the exploitation of multiple levels of parallelism, and the avoidance of thread divergence. In the following paragraphs, we show how we obtained highly efficient code using our new generic AST generation approach without the need to develop a domain- or optimization-specific generator.

When translating C code to CUDA, we start from code consisting of compute statements and loops. To simplify the exposition, let us first assume in this section that the program consists of a single perfectly nested set of loops, with one outer sequential loop, a set of inner parallel loops and a single compute statement. To generate CUDA code for this computation it is necessary to obtain a set of kernels that can be launched sequentially and that each expose two levels of parallelism: coarse grained parallelism, which will be mapped to so-called CUDA thread blocks, and fine grained parallelism, which will be mapped to so-called CUDA threads. To obtain these two levels of parallelism we divide the set of individual computations (statement instances) enumerated by these loops into subsets (tiles). We do this by computing a polyhedral schedule that enumerates the set of statement instances with two groups of loops. A set of outer loops that enumerate the tiles (tile loops) and a set of inner loops (point loops) that enumerate the statement instances that belong to a certain tile. The first AST generation problem we encounter is that the hybrid-hexagonal schedule defining the tile shapes decomposes the computation into phases and applies to each phase a different schedule. This results in a *piecewise schedule* from which an AST needs to be generated.

As a next step, we map the tile and point loops to a fixed number of thread blocks and threads. We start by looking for a set of parallel point loops and a set of parallel tile loops. We then strip-mine each loop by the number of thread blocks and threads. For instance to map a point loop with n iterations to a set of 1024 kernel threads, we strip-mine the loop by a factor of 1024 such that each 1024th iteration is executed by the same thread. The next step is to produce a piece of CPU code that schedules instances of an accelerated kernel, and the kernel code itself that defines the computation of a specific thread in a specific thread block. No actual loops are generated that enumerate the set of thread blocks and threads, but instead the CUDA run-time and hardware spawns a set of blocks and threads, and provides the block and thread ID as a parameter to each thread executing the kernel code. To model this, we first generate the outer loop, then we use a *nested context* to introduce the block and thread identifiers and, finally, we generate kernel code that can reference values in the outer CPU code, taking into account the AST generation context of the outer C code as well as the constraints on the kernel and thread identifiers. Exploiting this information is very important to generate high-quality code.

When generating kernel code we also need to rewrite all array subscripts in our compute statement. Traditionally this is done textually by replacing all references to old induction variables with expressions that compute the values of the old induction variables from new induction variables. When translating an access $A[i+1]$ where i now is expressed as $c0 + 1$, a classical rewrite would yield $A[(c0 + 1) + 1]$. With our new approach we represent the expression $i + 1$ itself as a piecewise quasi affine expression, perform the translation on the piecewise quasi affine expression, simplify the resulting expression and use our AST generator to *generate an AST expression from this piecewise quasi-affine expression*. As a result we obtain the code $A[c0 + 2]$. In this example the only benefit is increased readability, as any compiler would constant fold

```

for (c2 = 0; c2 <= 1; c2 += 1)
  for (c3 = 1; c3 <= 4; c3 += 1)
    for (c4 = max(((t1 - c3 + 130) % 128) + c3 - 2, ((t1 + c3 + 125) % 128) - c3 + 3);
        c4 <= min(((c2 + c3) % 2) + c3 + 128, -((c2 + c3) % 2) - c3 + 134); c4 += 128)
      if (c3 + c4 >= 7 || (c4 == t1 && c3 + 2 >= t1 && t1 + c3 <= 6
                          && t1 + c3 >= ((t1 + c2 + 2 * c3 + 1) % 2) + 3
                          && t1 + 2 >= ((t1 + c2 + 2 * c3 + 1) % 2) + c3)
          || (c4 == t1 && c3 == 1 && t1 <= 5 && t1 >= 4 && c2 <= 1 && c2 >= 0))
        A[c2][6 * b0 + c3][128 * g7 + c4 - 4] = ...;

```

Fig. 1: Copy code from hybrid hexagonal/parallelogram tiling (a single loop)

the two additions. However, in general, this concept is a lot more powerful. It allows the specialization of expressions according to the context in which they are generated. If, for instance, an access $A[i == 0 ? N - 1 : i - 1]$ is scheduled in a tile where we know i is never 0, we can simplify the access to $A[i - 1]$. This simplification removes the overhead of boundary condition handling from the core computation, a transformation for which a normal compiler misses context information and which traditionally requires specialized statements for boundary and core computations. With our AST generation approach, statements are automatically specialized as soon as boundary computations and core computations are generated as specialized AST subtrees. This is very natural for an AST generator that allows *user-directed versioning*.

After having generated basic CUDA code including the rewritten data accesses, we can start to optimize the code. An essential optimization is to switch from the use of slow “global memory” to the use of fast, manually managed “shared memory“. To do so we need to change the code of each tile such that, before the actual computation takes place, the relevant data from global memory is copied into shared memory, and at the end, the modified data is copied back from shared to global memory. To perform the computation in shared memory, we need to adjust all memory accesses such that they point to the new shared memory arrays and the corresponding locations. How exactly the mapping is computed is outside the scope of this paper, but how we generate the relevant code is interesting. We derive from our mapping a set of piecewise quasi-affine expressions that define the new data locations and *generate AST expressions* for them, relying on the AST generator to ensure that efficient code is generated. This approach enables us to use possibly complex mappings, without writing specialized code generation routines. To create the code that moves the data we create new statements that copy data from a given global memory location to a given shared memory location and vice versa. In case there is more data to copy than there are threads we use a modulo mapping to assign data locations to threads. Figure 1 shows the code generated to copy data back to global memory. There are various interesting observations possible. First, we see that our modulo expressions have been mapped to the C remainder operator %, which will be translated to fast bitwise operations. This is only possible because we have context information about the value of $t1$. Otherwise we would need to fall back to expensive `floord` or `intMod` expressions, dealing with arbitrary relative integers, like the state-of-the-art AST generators `CLooG` and `CodeGen+` do. Secondly, we see that we generate a reasonably dense loop nest that enumerates the statements. Because of the presence of existentially quantified variables in the input description, this is by itself non-trivial (see Section 7.3).

Nevertheless, we observe that the generated code is not very efficient. Every loop iteration performs very little computation and evaluates a complex condition. One might hope the condition could be simplified further, but unfortunately the data modified when moving a 5-point stencil forming a cross over a hexagonal tile shape is by itself

```

A[0][6 * b0 + 1][128 * g7 + (t1 + 125) % 128] - 1] = ...;
A[0][6 * b0 + 2][128 * g7 + (t1 + 127) % 128] - 3] = ...;
if (t1 <= 2 && t1 >= 1)
    A[0][6 * b0 + 2][128 * g7 + t1 + 128] = ...;
A[0][6 * b0 + 3][128 * g7 + (t1 + 127) % 128] - 3] = ...;
if (t1 <= 2 && t1 >= 1)
    A[0][6 * b0 + 3][128 * g7 + t1 + 128] = ...;
A[0][6 * b0 + 4][128 * g7 + (t1 + 125) % 128] - 1] = ...;
A[1][6 * b0 + 1][128 * g7 + (t1 + 126) % 128] - 2] = ...;
A[1][6 * b0 + 2][128 * g7 + (t1 + 126) % 128] - 2] = ...;
if (t1 <= 3 && t1 >= 2)
    A[1][6 * b0 + 2][128 * g7 + t1 + 128] = ...;
A[1][6 * b0 + 3][128 * g7 + (t1 + 126) % 128] - 2] = ...;
if (t1 <= 3 && t1 >= 2)
    A[1][6 * b0 + 3][128 * g7 + t1 + 128] = ...;
A[1][6 * b0 + 4][128 * g7 + (t1 + 126) % 128] - 2] = ...;

```

Fig. 2: Copy code from hybrid hexagonal/parallelogram tiling (unrolled)

already non-convex. Applying another level of modulo scheduling makes the necessary compute pattern even more complex, such that obtaining a simpler loop structure is difficult. However, by using *polyhedral unrolling* on the inner three loops and by specializing the statements according to the iteration they are unrolled for we can remove almost all control overhead. The result is shown in Figure 2. The code is very smooth and each array subscript is specialized to the specific location. We can also see that for the conditionally executed statements the subscripts are optimized according to the conditions such that the remainder operations disappear entirely. Unrolling this code is not trivial, as it needs to be performed in the presence of multiple loop boundaries as well as strides and we need to support the generation of guarded instructions when unrolling. The guarded instructions at the innermost level are very cheap on a GPU, as they can be implemented as predicated instructions. In this small example this is not very visible, but for realistic tile sizes a larger number of statements share the same conditions. We perform similar unrolling for the compute code in our kernel to ensure sufficient instruction level parallelism is available.

The code in Figure 2 is now close to optimal. However, so far we only looked at a simplified example, a single tile which does not touch any iteration space boundaries. In case iteration space boundaries are taken into account the generated code is a lot more complex. To ensure we can still use the “close to optimal” code most of the time, we use *user directed versioning* to isolate the core computation (the full tiles) from the set of tiles that need to take into account the boundary conditions (partial tiles). Doing so gives us maximal specialization and best performance. However, we now specialize and unroll not only the core computation, but also the code that was introduced to handle the boundary cases which increases the size of the generated code as well as the time necessary to generate it. When targeting a GPU this may be acceptable, but for FPGAs [Zuo et al. 2013] the cost may be prohibitive. This problem can be easily addressed by using *fine-grained options* to limit the amount of unrolling and specialization in the boundary tiles.

In summary, extending AST generation beyond the creation of control flow makes it possible to use automatic AST generation in complex scenarios. Even though existing AST generators combined with workarounds such as duplicating statements before running the AST generator can be used to solve some of the previously mentioned AST generation issues, such workarounds only exist for some features, they apply only in simple special cases and often inhibit other necessary transformations. By instead

```

for (int i = 0; i < n; ++i) {
S1: s[i] = 0;
    for (int j = 0; j < i; ++j)
S2:     s[i] = s[i] + a[j][i] * b[j];
S3: b[i] = b[i] - s[i];
}

```

Fig. 3: Example Program

carefully integrating several important new extensions into a single AST generation approach, we significantly extend the concept of automatic AST generation such that it is usable in complex AST generation scenarios. We ensure that the different features do not block each other, but when combined provide novel opportunities and solutions to complex AST generation problems. As a result we hope to not only significantly simplify AST generation, but to enable its use in new optimization scenarios.

3. POLYHEDRAL MODEL

The polyhedral model [Feautrier 1992] is a powerful abstraction for analyzing and transforming (parts of) programs that are “sufficiently regular”. The key feature of this model is that it is instance based. That is, each statement instance (i.e., each dynamic execution of a statement inside a loop nest) and each array element is treated individually through the use of a compact representation such as polyhedra [Loechner and Wilde 1997] or Presburger relations [Pugh and Wonnacott 1994b]. A program is typically represented using *iteration domains*, containing the statement instances, *access relations*, mapping statement instances to the accessed array element(s), *dependences*, relating statement instances that depend on each other, and a *schedule*, assigning an execution order to the statement instances.

In terms of AST generation, the most relevant elements are the iteration domain and the schedule, where the iteration domain describes the statement instances that need to be executed and the schedule describes the order in which they should be executed. For the iteration domain, we use the representation proposed by Verdoolaege [2011], where each statement instance is represented by a name (identifying the statement) and a tuple of integers (identifying the instance). For each statement, the instances in the iteration domain are described using a *Presburger formula*. We call such a set a *named Presburger set*. Other representations of iteration domains can easily be converted to such a named Presburger set.

Before defining Presburger formulas, let us first consider *affine* expressions, which are terms composed of variables, integer constants, symbolic constants, addition (+) and subtraction (−). Multiplication by an integer constant is available as syntactic sugar for repeated addition or subtraction. Symbolic constants have a fixed but unknown value and typically represent problem sizes. A *quasi-affine* expression additionally allows integer division by an integer constant ($\lfloor \cdot / d \rfloor$). A Presburger formula is then constructed from quasi-affine expressions, comparison (\leq) and the first order logic operators: conjunction (\wedge), disjunction (\vee), negation (\neg), existential quantification (\exists), and universal quantification (\forall). A *piecewise quasi-affine* expression is a list of pairs of named Presburger sets and quasi-affine expressions. The sets are pairwise disjoint and the value of the piecewise quasi-affine expression at a given point is equal to the value of the quasi-affine expression associated to the set that contains the point.

Binary relations on pairs of named integer tuples can be defined in a similar way and are called *named Presburger relations*. Although we will use a more structured representation for schedules in Section 6, it is instructive to consider the basic case of representing schedules as named Presburger relations proposed by Verdoolaege [2011].

These named Presburger relations associate an integer tuple to each statement instance and the execution order expressed by the schedule is given by the lexicographic order of these integer tuples. Consider for example the program in Figure 3. The iteration domain is

$$\{ S1(i) : 0 \leq i < n; S2(i, j) : 0 \leq j < i < n; S3(i) : 0 \leq i < n \}. \quad (1)$$

One way of expressing the original execution order is the schedule

$$\{ S1(i) \rightarrow (i, 0, 0); S2(i, j) \rightarrow (i, 1, j); S3(i) \rightarrow (i, 2, 0) \}. \quad (2)$$

An alternative execution order may be obtained through the schedule

$$\{ S1(i) \rightarrow (0, i, 0, 0); S2(i, j) \rightarrow (1, i, 0, j); S3(i) \rightarrow (1, i - 1, 1, 0) \}. \quad (3)$$

The purpose of the AST generator is to construct an AST that visits the elements of the iteration domain in the lexicographic order of the integer tuples assigned to the iteration domain elements by the schedule. The construction uses several operations on named Presburger sets and relation available in `isl` [Verdoolaege 2010], including domain and range of a relation, intersection, union, set difference, projection, shared constraints (“simple hull”), simplification (“gist”), coalescing (replacing pairs of disjuncts by single disjuncts without introducing spurious elements) and integer affine hull. It does, however, *not* use the convex hull operation, as this may introduce constraints with large coefficients.

4. DATA STRUCTURES

The core AST generation algorithm translates schedule constraints into lower and upper bounds of the for loops of the generated AST. In order to be able to construct the for loops, the algorithm may need to break down the schedule domain into several pieces, resulting in a tree of for loops generated from outermost to innermost. There may however also be other constraints that cannot be directly encoded in the lower and upper bounds of for loops and that need to be generated as if conditions instead. In general, we want these conditions to be inserted as high up as possible. On the other hand, we do not want to insert too many redundant conditions, while some conditions may only be redundant with respect to the code that is generated underneath those conditions. Moreover, some constraints, especially disjunctive constraints, may only get discovered at later stages and need to be hoisted up. The AST therefore cannot be constructed in a single pure pre-order depth-first traversal of the schedule. Instead, we perform a single depth-first traversal with some pre-order operations, mainly decomposing the schedule, and some post-order operations, actually constructing the AST. The details of the algorithm will be described in Section 5. In this section, we introduce vocabulary and data structures for passing information up and down the depth-first traversal. These are necessary to describe the current part of the decomposed schedule, the information that is passed down, the actual AST that is being constructed and the information that is passed up.

4.1. Executed Relation

In the base case of the AST generation algorithm, the schedule is given by a named Presburger relation mapping statement instances to their relative multi-dimensional execution times. The loops in the generated AST are derived from these multi-dimensional execution times. During the AST generation, it is therefore more natural to consider the inverse of this schedule relation, which we call the *executed relation* and which maps execution time vectors to the statement instances that should be executed at those times. For example, given the schedule relation in (3), the (initial) executed

relation is

$$\{ (0, s_1, 0, 0) \rightarrow S1(s_1); (1, s_1, 0, s_2) \rightarrow S2(s_1, s_2); (1, s_1, 1, 0) \rightarrow S3(s_1 + 1) \}. \quad (4)$$

The levels of the depth-first pass over the schedule correspond to the input dimensions of this executed relation. At each level, the domain of the executed relation is broken up into pieces along that dimension and each piece of the executed relation is considered in turn.

4.2. Stock

The AST node corresponding to a dimension in the domain of the executed relation is constructed upon leaving that level during a depth-first traversal. However, the main information about the AST node is already available when that level is first entered. Some of this information needs to be stored and forwarded through the traversal. We introduce the *stock* to collect this information that can be used to simplify the descendant AST nodes. The stock mainly keeps track of two pieces of information, the conditions on symbolic constants and outer loop iterators that are known to hold at the current position and a mapping from loop iterators to schedule dimensions. At each level of the depth-first traversal, a new stock is created that is initialized from the stock passed down from the higher level.

The conditions come in two groups, the *generated* conditions and the *pending* conditions. The generated conditions are those for which the algorithm has already decided that they *will* be enforced by the outer nodes in the AST. These are typically the loop bounds on the outer loop nodes. The pending conditions are those that *may* end up being enforced by the outer nodes. They may also get dropped if they turn out to be implied by the inner AST nodes. Note that this distinction is only relevant at the point where the actual AST nodes are being constructed. At other points in the AST generation algorithm, we can simply consider the combination of the two groups of constraints.

The mapping from loop iterators to schedule dimensions is needed because unlike other AST generators, ours exploits the fact that a schedule only specifies a *relative* execution order. The loop iterators in the final AST may then not correspond exactly to the schedule dimensions in the input schedule due to, e.g., scaling or strip-mining by the AST generator, while these schedule dimensions may still be referenced from other parts of the schedule through options (see Section 5.6) or advanced schedule tree nodes (see Section 6).

4.3. Abstract Syntax Tree

The generated AST contains only syntactical information and has been designed to be easily translatable to both C and compiler IR. Each node of the AST is of one of four types, an *if node*, a *for node*, a *block node* or a *user node*. An if node has an AST expression as condition, a then node and optionally an else node. A for node has initialization, condition and increment expressions and a body node. A block node represents a compound statement and maintains a list of nodes. Finally, the statement expressed by a user node is represented as an AST expression.

An AST expression is itself a tree with operators in the internal nodes and integer constants or identifiers as the leaves. The set of operators contains the standard operators found in C-like programming languages, but also higher level operators such as *min* and *max*. Boolean logical operators and the conditional operator (*cond ? a : b*) are available in two forms, *lazy* and *eager*. We found in our work on low-level compilers [Grosser et al. 2012] that eagerly evaluating operands, instead of using C's lazy evaluation, is often beneficial as it reduces control overhead and simplifies the hoisting of loop invariant subexpressions.

The integer division operator also comes in different forms, one of them corresponding to the mathematical operation $\lfloor a/b \rfloor$. Unfortunately, this operation cannot be translated directly into `a / b` in C because the `/`-operator in C rounds toward zero rather than toward negative infinity. A correct translation to C involves a condition on the sign of a , which can bring significant extra costs on some architectures such as GPU devices. We therefore also have a form of the integer division where the result is known to be an integer (such that rounding becomes irrelevant) and one where the dividend is known to be non-negative. The user can specify a preference for these latter forms in which case the AST expression generator will look for opportunities to use them (see Section 5.10). Similarly, the remainder operator comes in two special forms, one where the dividend is known to be non-negative and one where the result of the operations is only compared against zero. In both cases, the remainder operator can be translated into the `%`-operator in C.

4.4. Annotated AST

The AST nodes are created after having visited all the children in the depth-first traversal of the schedule. The AST generator may however decide to not encode some of the conditions that need to be satisfied by the symbolic constants in the generated AST nodes such that these conditions may be hoisted up to higher levels. An *annotated AST* keeps track of the (purely syntactical) AST itself, as well as such extra pieces of polyhedral information. Besides the conditions described above that still need to be enforced by the AST at higher levels, the annotated AST also keeps track of the conditions that *have* been enforced already. This latter set of conditions can then be used to simplify or even eliminate some of the pending conditions in the stock at higher levels.

5. AST GENERATION

This section describes the core AST generation algorithm. In Section 6, we will see that this algorithm is applied for each band node in the schedule tree. Our core algorithm is derived from the “Quilleré et al.” algorithm [Bastoul 2004; Quilleré et al. 2000], with several significant changes such as isolation, unrolling, if-hoisting, the detection of components, shifted strides and the optimized generation of AST expressions.

5.1. Overview

Algorithm 3 forms the core of the AST generation and creates a (possibly degenerated) for AST node for a given schedule dimension after generating AST nodes for the next schedule dimensions through a call to Algorithm 1. The process of creating such a for node is detailed in Section 5.3. The input is a single-disjunct set corresponding to the current dimension in the schedule domain. The actual schedule domain at this dimension may however consist of several disjuncts. Algorithm 2 takes care of breaking up (or overapproximating) the schedule domain into disjoint single-disjunct pieces and calling Algorithm 3 on each of them. The different ways of breaking up the schedule domain are explained in Section 5.4 and Section 5.5.

Finally, Algorithm 1 is the main driver that is called (recursively) for each dimension in the schedule space (i.e., the domain of the executed relation). It calls Algorithm 2 after detecting some special cases. In particular, as long as the inner level has not been reached, the algorithm first looks for components as described in Section 5.7. In each component, the algorithm checks for shifted strides as explained in Section 5.8, possibly modifying the executed relation if any shifted strides were detected. The domain of the executed relation is then projected onto the outer level dimensions as explained in Section 5.2. Finally, if the user has specified a piece of the schedule space that needs to be isolated (see Section 5.6), then the algorithm splits the schedule domain into four

ALGORITHM 1: Generate Next Schedule Dimension (“next”)

```

if at inner level then
  | return terminate(stock, executed)
end
list := ()
foreach sorted component do
  (stock, executed, f) := detect shifts(stock, executed)
  domain := project(executed, level)
  if has isolation domain then
    | (before, domain, after, other) = split on isolation(domain)
    | list' := base(stock, executed, before, false)
    | list' += base(stock, executed, domain, true)
    | list' += base(stock, executed, after, false)
    | list' += base(stock, executed, other, false)
  else
    | list' := base(stock, executed, domain, false)
  end
  list += transform(list', f)
end
return list

```

ALGORITHM 2: Generate Component (“base”)

```

type := generation type(level, isolated)
if type = unroll then
  | (bound, n) := find lower bound(domain)
  | list := ()
  | for  $0 \leq i < n$  do
    | domain' := slice(domain, bound, i)
    | domain' := shared constraints(domain')
    | list += (create(stock, executed, domain'))
  end
  return list
end
if type = separate then
  | domain list := separate(domain, executed)
else if type = atomic then
  | domain list := (shared constraints(domain))
else
  | domain list := make disjoint(domain)
end
list := ()
foreach domain in sorted domain list do
  | list += (create(stock, executed, domain))
end
return list

```

parts, the part that comes before the isolated part, the isolated part itself, the part that comes after and the part that is incomparable to the isolated part.

When the inner level has been reached, the schedule makes no further distinction between the statement instances in the range of the current executed relation. We therefore generate an AST for each statement separately. Usually, the executed relation will only associate a single statement instance to a given schedule point and we

ALGORITHM 3: Create for node (“create”)

```

domain := bounds intersected with domain of executed
stock' := detect strides(stock, domain)
stock' := check for single iteration(stock', bounds)
list := next(stock', executed)
list := combine(list, stock, stock')
return construct for loop(bounds, stock, list)

```

simply create and return a user node. Otherwise, the AST still needs to iterate over the different instances and it can do so in any order. We therefore extend the domain of the executed relation with a copy of the range and continue processing the new dimensions in the domain of the executed relation until the inner level is reached again, in which case the executed relation is guaranteed to have only a single statement instance associated to a given schedule point.

5.2. Local Schedule Domain Constraints

The lower and upper bounds of a for loop generated at a given level are derived from the constraints in the domain of the executed relation that involve the current schedule dimension. These constraints may also involve schedule dimensions, both those corresponding to outer for loops and those corresponding to inner for loops. Since the lower and upper bounds of a for loop can only refer to iterators of outer loops and (obviously) not to those of inner loops, we first need to project the domain of the executed relation onto its first level dimensions. This operation may introduce additional existentially quantified variables, which cannot be encoded directly in AST expressions. We therefore need to remove them in some way.

One way of removing the existentially quantified variables is to perform quantifier elimination. This process preserves the meaning of the set and therefore ensures that only those values of the current schedule dimension that have any associated statement instances will be executed. On the other hand, quantifier elimination may break up the schedule domain into several pieces. In particular, `isl` performs quantifier elimination by applying parametric integer programming [Feautrier 1988] to compute explicit, quasi-affine representations of the quantified variables. In general, this may split the domain into several parts, each with its own quasi-affine expressions. Other quantifier elimination algorithms lead to similar decompositions of the schedule domain.

Another way of removing the existentially quantified variables is to perform Fourier-Motzkin elimination on them. This may result in an overapproximation of the schedule domain, but it is guaranteed not to break up the schedule domain into several disjuncts. In our AST generator, we take this second option in order to avoid the code size expansion resulting from the first option. That is, we apply Fourier-Motzkin elimination to remove all existentially quantified variables for which we do not have an explicit representation yet. Note that we will only consider the constraints of the possible overapproximation as having been generated at this level, such that any constraint on the actual schedule domain that is not satisfied by this overapproximation will end up getting enforced at a deeper level.

Example 5.1. As an example of a case where these two approaches produce different results, consider the following projection of some schedule domain onto the outermost schedule dimension,

$$\{ (t) : (\exists \alpha : \alpha \geq -1 + t \wedge 2\alpha \geq 1 + t \wedge \alpha \leq t \wedge 4\alpha \leq N + 2t) \}, \quad (5)$$

where N is a symbolic constant. Applying quantifier elimination to the set description in (5) results in

$$\{(t) : (t \geq 3 \wedge 2t \leq 4 + N) \vee (t \leq 2 \wedge t \geq 1 \wedge 2t \leq N)\}. \quad (6)$$

Note that this is the same set, but that it is described in a different way, without existentially quantified variables. Also note that the description now consists of two disjuncts. Applying Fourier-Motzkin elimination on the α variable in (5) on the other hand results in

$$\{(t) : 2t \leq 4 + N \wedge N \geq 2 \wedge t \geq 1\}. \quad (7)$$

This set contains an extra element that is not an element of the set in (5). In particular, it contains the extra element

$$\{(2) : 2 \leq N \leq 3\} \quad (8)$$

Using the set in (6) would produce the code

```
for (int c0 = 1; c0 <= min(2, floord(N, 2)); c0 += 1)
  // body
for (int c0 = 3; c0 <= floord(N, 2) + 2; c0 += 1)
  // body
```

whereas using the set in (7) produces the code

```
for (int c0 = 1; c0 <= floord(N, 2) + 2; c0 += 1)
  // body
```

As explained above, `isl`'s quantifier elimination may replace a quantified variable by a quasi-affine expression. During the simplification of set descriptions, `isl` may perform a similar substitution. We therefore need to take into account that the projection of the schedule domain, even after eliminating the existentially quantified variables using Fourier-Motzkin elimination, may involve such quasi-affine expressions. If any of these quasi-affine expressions depend on the current level, then they are also eliminated using Fourier-Motzkin elimination since a lower bound of a loop cannot depend on the value of the loop iterator. Similarly, a constraint involving such an expression cannot be used to construct an upper bound since it may fail to hold for some values of the iterator while being satisfied by higher values of the iterator. Note that this process also removes any stride information, but this information will be recovered from the executed relation as explained in Section 5.3.

Example 5.2. Consider an iteration domain

$$\{S(i) : 3 \left\lfloor \frac{i+1}{3} \right\rfloor \leq i \wedge i \geq 0 \wedge i \leq 3\} \quad (9)$$

with schedule

$$\{S(i) \rightarrow (i)\}. \quad (10)$$

The schedule domain and its projection onto the outermost (and only) level is

$$\{(i) : 3 \left\lfloor \frac{i+1}{3} \right\rfloor \leq i \wedge i \geq 0 \wedge i \leq 3\}. \quad (11)$$

The constraint $3 \lfloor (i+1)/3 \rfloor \leq i$ involves a quasi-affine expression in terms of the current level and therefore cannot be used in the construction of the for loop bounds. Instead, the expression is eliminated, resulting in the schedule domain

$$\{(i) : i \geq 0 \wedge i \leq 3\}. \quad (12)$$

The eliminated constraint is then taken into account at the innermost level, resulting in the following code.

```
for (int c0 = 0; c0 <= 3; c0 += 1)
  if ((c0 + 1) % 3 >= 1)
    S(c0);
```

Note that when reaching the innermost level, we can no longer perform any approximations and we have to perform quantifier elimination on any remaining existentially quantified variables. This may then result in a disjunctive condition around the generated statement. Also note that the quantifier elimination procedure of the Omega library is different from the one used by `isl`, but it may also result in splitting the domain. A detailed comparison of the two is beyond the scope of the present article. It is not clear at which point and how far `CodeGen+` applies quantifier elimination, but it appears to be tailored to constraints that only involve a single existentially quantified variable.

5.3. For Node Construction

Besides the stock and the executed relation, Algorithm 3 for creating a for node takes as extra input a convex set called “bounds” from which the bounds of the for node will be extracted. The first step is the detection of strides, but this information is not available in bounds since it does not involve any quasi-affine expressions that depend on the current level. Instead, we use the intersection of the bounds set with the domain of the executed relation. The strides are extracted from the integer affine hull [Verdoolaege 2010] of the resulting set. This operation extracts the equality constraints satisfied by the elements of the set and preserves (most of) the stride information. Let

$$h(\mathbf{p}) + u(vi + sf(\boldsymbol{\alpha})) = 0 \quad (13)$$

be one of these equality constraints, with i the current schedule dimension, \mathbf{p} the symbolic constants and outer dimensions and $\boldsymbol{\alpha}$ the existentially quantified variables. Furthermore, v and s have no common factor and f and h are affine expressions such that the coefficients of f have no common factor. If v is nonzero and s is greater than one, then this equality represents a non-trivial stride constraint. Using Bézout’s identity $av + bs = 1$, we can rewrite (13) to

$$ui = -ah(\mathbf{p}) + us(bi - af(\boldsymbol{\alpha})) \quad (14)$$

so that $-ah(\mathbf{p})$ is a multiple of u and i is equal to $o(\mathbf{p}) = -ah(\mathbf{p})/u$ modulo s . The offset $o(\mathbf{p})$ and the stride s are stored in the stock. If there is more than one equality with a non-trivial stride, then the offsets and strides can be combined and the overall stride will be the least common multiple of the strides. In particular, if we have two offset/stride pairs

$$i = o_1(\mathbf{p}) + s_1 f_1(\boldsymbol{\alpha}) \quad (15)$$

$$i = o_2(\mathbf{p}) + s_2 f_2(\boldsymbol{\alpha}) \quad (16)$$

then let g be the greatest common divisor of s_1 and s_2 and let c and d be such that $cs_1 + ds_2 = g$ (Bézout’s identity once more). Multiplying the equation in (15) by $t_1 = ds_2/g$ and the equation in (16) by $t_2 = cs_1/g$, we obtain

$$i = (ds_2 + cs_1)/g i = t_1 o_1(\mathbf{p}) + t_2 o_2(\mathbf{p}) + (s_1 s_2)/g (bf_1(\boldsymbol{\alpha}) + af_2(\boldsymbol{\alpha})). \quad (17)$$

That is, the combined offset is $t_1 o_1(\mathbf{p}) + t_2 o_2(\mathbf{p})$, while the combined stride is $(s_1 s_2)/g$, the least common multiple of s_1 and s_2 .

Example 5.3. Consider the schedule domain

$$\{ (i) : \exists \alpha, \beta : 0 \leq i \leq 100 \wedge n - i + 6\alpha = 0 \wedge m - i + 10\beta = 0 \}, \quad (18)$$

with n and m symbolic constants. For the constraints $n - i + 6\alpha = 0$, we have, using the notation of (13), $h(\mathbf{p}) = n$, $u = 1$, $v = -1$, $s = 6$ and $f(\alpha) = \alpha$. We may take $a = -1$ and $b = 0$ to find $o_1(\mathbf{p}) = n$ with $s_1 = 6$. We similarly find $o_2(\mathbf{p}) = m$ and $s_2 = 10$. We have $g = 2$ and may take $c = 2$ and $d = -1$, resulting in a combined stride of 30 and a combined offset of $-5n + 6m$. To see that this combined offset satisfies both stride constraints, note that the combination of the original two stride constraints implies that $n - m$ is a multiple of two. That is, the bounds set (from which the existentially quantified variables have been eliminated) is of the form

$$\{ (i) : 0 \leq i \leq 100 \wedge 2 \left\lfloor \frac{m - n}{2} \right\rfloor = m - n \}. \quad (19)$$

We therefore have that $m - (-5n + 6m) = 5(n - m)$ is indeed a multiple of 10.

After detecting the strides, we add the constraints enforced by the for loop generated for the current schedule dimension to the stock that will be used in the construction of descendant nodes such that these constraints may be used to simplify those descendant nodes. The actual construction of the for loop corresponding to the current schedule dimension is only performed after these descendant nodes have been created. At that point we will need to make a distinction between the constraints that are enforced by the bounds of the for loop and the constraints that may need to be enforced by an extra if conditional. We therefore take this difference into account while updating the stock, even though this distinction has no influence on the descendant nodes.

In order to determine which constraints are enforced by the for loop corresponding to the current schedule dimension, we first need to know if we are even going to construct a for loop. In particular, based on the constraints in the current stock, the constraints in bounds and the stride constraints (if any), we may be able to determine that the current schedule dimension can attain only a single value. In this case, we will generate a special “degenerate” for loop which we allow the user to translate to an assignment of the initial value to the loop iterator. Note that this single value will in general be specified as a piecewise quasi-affine expression in the symbolic constants and the outer loop iterators. If it turns out that this expression consists of a single quasi-affine expression, then we do not generate any for loop at all, but instead substitute the current schedule dimension for this single quasi-affine expression in the executed relation. We refer to this case as an eliminated for loop. The reason for only performing this substitution when the single value is described by a single quasi-affine expression is that otherwise we would be introducing additional disjuncts in the executed relation. In the eliminated case, we eliminate the current schedule dimension from bounds and add the result to the pending constraints. In the other cases, we add the constraints in bounds that do not involve the current schedule dimension to the pending constraints and we add the remaining constraints in bounds as well as the stride constraint (i.e., the fact that the schedule dimension is equal to the offset plus a multiple of the stride) to the generated constraints.

Example 5.4. Consider the schedule domain

$$\{ (i) : i \geq 1 \wedge n - 1 \leq i \leq n \wedge 4 \left\lfloor \frac{i - 2}{4} \right\rfloor = i - 2 \}, \quad (20)$$

where n is a symbolic constant. The stride constraint $4 \lfloor (i - 2)/4 \rfloor = i - 2$ does not appear in bounds, but it is added back for the purpose of looking for a single value.

From these constraints, we can see that i attains a single value of $4 \lfloor (n+2)/4 \rfloor - 2$ and that this value is represented as a single quasi-affine expression. Substituting this value in the schedule domain, we obtain

$$\{ () : n \geq 2 \wedge \lfloor \frac{n}{4} \rfloor \leq n - 2 \}. \quad (21)$$

These constraints are then added as guards to the annotated AST at the innermost level.

Note that degenerate loops that are not also eliminated are fairly rare in practice. We only report our handling of this situation for completeness, but it may very well not be close to optimal. In most of the remaining rare cases, the loop could in fact be eliminated, but we simply fail to derive an appropriate single quasi-affine expression. In fact, these cases used to occur more frequently, but most of these have been resolved through an improved detection of single quasi-affine expressions.

After generating an annotated AST for the body of the for node, we know which constraints on the symbolic constants and outer loop iterators are enforced by this subtree and we can (optionally) use them to simplify the pending constraints. Additionally, the pending constraints are simplified with respect to the generated constraints. The simplified pending constraints are then combined with the constraints hoisted from the annotated AST for the body (see Section 5.9) and an additional set of *implied* constraints. The implied constraints are those constraints that are implied by the generated constraints, but that may not be implied by their AST expression counterparts. In particular, these are the constraints implied by the stride constraints. In case of a degenerate loop, this also includes the constraints implied by the generated constraints. Since we are allowing the user to consider a degenerate for loop as an assignment, the fact that the upper bound is greater than or equal to the lower bound (i.e., that there even is a single iteration of the loop) is not enforced by this assignment and therefore needs to be considered separately. The combination of the simplified pending constraints, the hoisted constraints and the implied constraints will then definitely be generated either at the current level or hoisted up to a higher level. From this point on, in particular for the construction of the AST expressions for the if conditions that have not been hoisted out of the body and those for the bounds of the for loop, they may therefore be considered as generated constraints.

Example 5.5. As an example of the effect of exploiting enforced constraints, consider the iteration domain

$$\{ S(i, j) : 0 \leq i < m \wedge 0 \leq j < n \}, \quad (22)$$

where m and n are symbolic constants, with schedule

$$\{ S(i, j) \rightarrow (i, j) \}. \quad (23)$$

Projection of the schedule domain onto the outer dimension yields

$$\{ (i) : 0 \leq i < m \wedge n \geq 1 \}. \quad (24)$$

The single pending constraint at this level is therefore $n \geq 1$. At the inner level the schedule domain, simplified with respect to the stock constraints, is

$$\{ (i, j) : 0 \leq j < n \}. \quad (25)$$

The for loop generated at this inner level enforces the constraint $n \geq 1$ so it can optionally be used to simplify the pending constraint at the outer level. If we exploit this enforced constraint, we generated the code

```

for (int c0 = 0; c0 < m; c0 += 1)
  for (int c1 = 0; c1 < n; c1 += 1)
    S(c0, c1);

```

Otherwise, the pending constraint is turned into an if condition and we generate the code

```

if (n >= 1)
  for (int c0 = 0; c0 < m; c0 += 1)
    for (int c1 = 0; c1 < n; c1 += 1)
      S(c0, c1);

```

Example 5.6. As an example of constraints implied by a stride constraint, consider the schedule

$$\{ S(t) \rightarrow (t) : \exists \alpha : 2t - n = 4\alpha \wedge 0 \leq t \leq 100 \}, \quad (26)$$

with n a symbolic constant. Stride detection finds a stride of 2 and an offset of $n/2$. The stride constraint $(n/2 - t) \bmod 2 = 0$ is encoded as

$$n - 2t - 4 \left\lfloor \frac{n + 2t}{4} \right\rfloor = 0. \quad (27)$$

This constraint implies that n is a multiple of 2. Since this stride constraint is added to the generated constraints, this fact may be simplified away at the deeper levels. However, it is not implied by the actually generated for loop. We therefore eliminate t from (27) and add the resulting constraints to the constraints that need to be generated at the outer level. The final code is

```

if (n % 2 == 0)
  for (int c0 = (n / 2) + 2 * floord(-n - 1, 4) + 2; c0 <= 100; c0 += 2)
    S(c0);

```

Note that we have exploited the fact that n is a multiple of 2 while generating the loop initialization.

Let us now consider the construction of the initialization and the condition of the generated for loop from the lower and upper bounds on the current schedule dimension. Since the bounds set may be an overapproximation of the schedule domain, it may in rare cases *not* involve lower and/or upper bounds. If they are missing, then we derive a single piecewise quasi-affine bound from the domain set using parametric integer programming [Feautrier 1988]. If this set does not have a lower bound, then an error is reported. If it has no upper bound, then an infinite for node is generated. In the standard case where there is one or more lower bound constraint $h(\mathbf{p}) + vi \geq 0$ with $v > 0$, each of the constraints is converted to a lower bound $\ell(\mathbf{p}) = \lceil -h(\mathbf{p})/v \rceil$ and the lower bound on the for node is set to the maximum (as an AST expression) of these lower bounds. If the loop is strided, however, then we need to make sure that this lower bound has the right value modulo the stride. We therefore first replace each of the lower bounds $\ell(\mathbf{p})$ by $o(\mathbf{p}) + s \lceil (\ell(\mathbf{p}) - o(\mathbf{p}))/s \rceil$.

Example 5.7. Consider the iteration domain

$$\{ S1(i) : 0 \leq i \leq M; S2() \}, \quad (28)$$

where M is a symbolic constant, with schedule

$$\{ S1(i) \rightarrow (i, 0); S2() \rightarrow (0, 1) \}. \quad (29)$$

Assume that at the outer level, we want to generate a single loop for both statements, as explained in Section 5.4, with bounds set $\{ (i) : i \geq 0 \}$. This bounds set does not

have any upper bound on the current schedule dimension so we consider the schedule domain

$$\{ (0); (i) : 0 \leq i \leq M \} \quad (30)$$

instead. From this set, we can derive the upper bound

$$\begin{cases} 0 & \text{if } M \leq 0 \\ M - 1 & \text{otherwise.} \end{cases} \quad (31)$$

The generated code is as follows.

```
for (int c0 = 0; c0 <= (M <= 0 ? 0 : M); c0 += 1) {
  if (M >= c0)
    S1(c0);
  if (c0 == 0)
    S2(0);
}
```

Example 5.8. Continued from Example 5.3. The bounds set in (19) has only a single lower bound on the current schedule dimension, $i \geq 0$. The default for loop initialization would therefore be $\max\{\lceil -0/1 \rceil\} = 0$. This value may however not satisfy the stride constraint, depending on the values of m and n . It is therefore replaced by

$$-5n + 6m + 30 \left\lceil \frac{5n - 6m}{30} \right\rceil = -5n + 6m + 30 \left\lceil \frac{5n - 6m + 29}{30} \right\rceil. \quad (32)$$

Depending on a user setting, the for node upper bound condition is constructed either as a single comparison of the loop iterator to a minimum of upper bounds, derived analogously to the lower bounds, or as a conjunction of comparisons, each derived directly from an upper bound constraint. The single upper bound is expected by the OpenMP support of some compilers while the conjunction is more efficient on FPGAs [Zuo et al. 2013]. Finally, the independent constraints are added to the annotated AST holding the for node.

Example 5.9. Consider upper bounds of the form

$$M \geq c_3 + 1 \wedge c_1 \geq 3c_3 + 8 \quad (33)$$

If the user has selected the generation of a single upper bound, then an upper bound condition of the form $c_3 < \min((c_1 + 1) / 3 - 2, M)$ is generated, while in the other case, an upper bound condition of the form $M \geq c_3 + 1 \ \&\& \ c_1 \geq 3 * c_3 + 8$ is generated.

5.4. Separation

The schedule domain(s) computed in Algorithm 1 may be arbitrary Presburger sets, which in `isl` are represented in disjunctive normal form. The creation of a for node in Algorithm 3, however, takes a single-disjunct set as input. The responsibility of the intermediate Algorithm 2 is then to replace the schedule domain by an ordered sequence of disjoint single-disjunct domains. There are essentially two ways to obtain such single-disjunct domains, either the entire domain is approximated by a single-disjunct domain or the domain is broken up into single-disjunct parts. The first option may require the introduction of additional guards in the descendants of the currently constructed for node, possibly leading to run-time overhead. The second option can lead to code duplication since different instances of the same domain may end up in different single-disjunct schedule domain parts.

For each schedule dimension, the user can specify which of these options to take. Alternatively, the user may also specify that the schedule dimension should be unrolled (Section 5.5) or she may leave the option unspecified, in which case the schedule domain is broken up into a list of single-disjunct domains in a pragmatic way.

Let us consider the two main options in some more detail. If the “separate” option is specified, then the schedule domains are computed for each statement separately. Each of these schedule domains is broken up into disjoint single-disjunct sets and a common refinement is computed. This is the standard separation of the “Quilleré et al.” algorithm [Quilleré et al. 2000; Bastoul 2004]. In the “atomic” case, the *shared constraints* are used. That is, the constraints of the disjuncts are considered in turn and only those are kept that are satisfied by the entire schedule domain. This process may in some cases result in the absence of a lower and/or upper bound, a case we discussed in Section 5.3.

Example 5.10. Continued from Example 5.7. Consider the schedule domain at the outer dimension in (30), repeated here

$$\{ (0); (i) : 0 \leq i \leq M \}. \quad (34)$$

In case of separation, we consider the schedule domains for each statement separately, i.e., $\{ (0) \}$ and $\{ (i) : 0 \leq i \leq M \}$, and apply the standard separation algorithm, breaking up these sets into disjoint single-disjunct sets, resulting in

$$\{ (0) : M \leq -1 \}, \{ (0) : M \geq 0 \} \quad \text{and} \quad \{ (i) : 1 \leq i \leq M \}. \quad (35)$$

The generated code is as follows.

```
if (M <= -1) {
  S2(0);
} else {
  S1(0);
  S2(0);
  for (int c0 = 1; c0 <= M; c0 += 1)
    S1(c0);
}
```

In the atomic case, we consider the constraints shared by the disjuncts in (30). In this example, there is only one such constraint, i.e., $i \geq 0$. Note that the equality constraint $i = 0$ in the first disjunct is equivalent to $i \geq 0 \wedge i \leq 0$. The generated code for this case is shown in Example 5.7.

5.5. Unrolling

Unrolling in the AST generation works by taking slices of the schedule domain for successive values of the current schedule dimension and by calling “create” for each of these slices. By construction, the schedule dimension has a fixed quasi-affine value in each of the slices and no actual for node will be created. Two factors play an important role in unrolling: stride detection and the selection of the most appropriate lower bound. Stride detection is performed as explained in Section 5.3. If any stride is found, then it is substituted ($i = o(\mathbf{p}) + si'$) in the schedule domain for the purpose of selecting a lower bound.

Example 5.11. If the schedule domain is of the form

$$\{ i : 0 \leq i < 1024 \wedge i \bmod 256 = 0 \}, \quad (36)$$

then it is replaced by $\{ i' : 0 \leq i' < 4 \}$.

The lower bound identification requires a single disjunct so we consider once more the shared constraints of the schedule domain, although in this case we also allow constant shifts of the constraints. For each lower bound constraint $h(\mathbf{p}) + vi \geq 0$ with $v > 0$, we compute the maximum value of $i + 1 - \lceil -h(\mathbf{p})/v \rceil$ over the schedule domain. If the maximum exists and has value n then we know we can cover the schedule domain with at most n slices of the form $i = \lceil -h(\mathbf{p})/v \rceil + t$ with $0 \leq t < n$. We take the lower bound with the smallest such n . If no suitable lower bound can be found, then we report an error.

Example 5.12. For the schedule domain

$$\{i : 0 \leq i < 1000 \wedge N \leq i < N + 4\}, \quad (37)$$

we would prefer the lower bound N (with 4 slices) over the lower bound 0 (with 1000 slices).

5.6. Isolation

Isolation in Algorithm 1 is controlled by an option specified by the user. If set, the option describes a part of the schedule domain that should be isolated from the other parts of the schedule domain. The typical use cases are the isolation of full tiles from partial tiles or the isolation of a vectorizable loop from its prologue and epilogue. The *atomic/separate/unroll* option can be specified separately for the isolated part and the rest of the schedule domain. For any given level, the isolated domain is first projected onto the first level dimensions as in Section 5.2. In particular, the inner dimensions are projected out and all existentially quantified variables and all quasi-affine expressions involving the current dimension are eliminated. We subsequently replace the set by its shared constraints to ensure that the center part is contiguous. An intersection with the current schedule domain yields the center part of the isolation. The “before” part is obtained by first constructing a set of iterations that are executed before some iteration in the center part and then subtracting that center part. Similarly, the “after” part is obtained by first constructing a set of iterations that are executed after some iteration in the center part and then subtracting both the center part and the “before” part. The “other” part is obtained by subtracting the before, center and after parts from the current schedule domain and consists of those iterations that are incomparable to the center part. If any *atomic/separate/unroll* option is specified for the rest of the schedule domain, then it is applied to the before, after and other part separately.

Example 5.13. Assume we have an iteration domain

$$\{S(i) : m \leq i < n\}, \quad (38)$$

where m and n are symbolic constants, with initial schedule

$$\{S(i) \rightarrow (i)\} \quad (39)$$

and that we want to strip-mine the loop so that we can allow the backend compiler to vectorize the inner loop. We first modify the schedule to

$$\{S(i) \rightarrow (4 \lfloor \frac{i}{4} \rfloor, i)\} \quad (40)$$

and then we want to single out those iterations that result in an inner loop of exactly four iterations. In particular, we need to make sure that the first schedule dimension $t = 4 \lfloor i/4 \rfloor$ belongs to the schedule domain of the second dimension and that so does $t + 3$. We therefore isolate the values of the first schedule dimension that satisfy

$$\{(t) : m \leq t \wedge t + 3 < n\}. \quad (41)$$

Projecting out inner dimensions and replacing the set by its shared constraints does not modify the isolated set. The before part is

$$\{(t) : n \geq 4 + m \wedge t \leq m - 1\}, \quad (42)$$

the after part is

$$\{(t) : n \geq 4 + m \wedge t \geq n - 3\} \quad (43)$$

and the other part is

$$\{(t) : m - 3 \leq t \leq n - 1 \wedge n \leq m + 3 \wedge n \geq m + 1\}. \quad (44)$$

The generated code (shown below) consists of a single loop performing the prologue computation, two—now easily vectorizable—loops that enumerate the center part, and one loop that forms the epilogue computation. There is also an additional loop nest for the other part, which is executed in case the values of n and m yield an empty center part. In certain cases, it is possible to avoid the generation of dedicated code for the other part by instead enumerating the relevant iterations as part of the before and after parts. We currently do not perform such an optimization.

```
{
  if (n >= m + 4)
    for (int c1 = m; c1 <= 4 * floord(m - 1, 4) + 3; c1 += 1)
      S(c1);
  for (int c0 = 4 * floord(m - 1, 4) + 4; c0 < n - 3; c0 += 4)
    for (int c1 = c0; c1 <= c0 + 3; c1 += 1)
      S(c1);
  if (n >= m + 4 && 4 * floord(n - 1, 4) + 3 >= n) {
    for (int c1 = 4 * floord(n - 1, 4); c1 < n; c1 += 1)
      S(c1);
  } else if (m + 3 >= n)
    for (int c0 = 4 * floord(m, 4); c0 < n; c0 += 4)
      for (int c1 = max(m, c0); c1 <= min(n - 1, c0 + 3); c1 += 1)
        S(c1);
}
```

5.7. Components

It is important to understand that the schedule only defines a *relative* execution order of statement instances with respect to other statement instances and that the generated loops do not need to correspond exactly to the schedule dimensions. In some cases, we can generate much simpler code by taking a closer look at this relative execution order.

Example 5.14. As a simple example, take an iteration domain

$$\{S_0(); S_1(i) : 0 \leq i < 10\} \quad (45)$$

with schedule

$$\{S_0() \rightarrow (0); S_1(i) \rightarrow (i)\}. \quad (46)$$

The schedule domains for the two statements separately are $\{(0)\}$ and $\{(i) : 0 \leq i < 10\}$. The latter is also equal to the combined schedule domain. Applying separation would therefore split off iteration 0 from S_1 , while using an atomic domain would generate a single loop with a condition on S_0 for the iterator to be equal to 0. A closer look at the schedule reveals, however, that $S_0()$ is not scheduled *after* any iteration of S_1 , so that we can simply generate code for S_0 first, resulting in

```

S0();
for (int i = 0; i < 10; ++i)
    S1(i);

```

In the general case, we construct a graph with as nodes the statements and an edge between node a and node b if any instance of b is scheduled after any instance of a and if these instances are mapped to the same value of the outer schedule dimensions. The strongly connected components in this graph may be scheduled in topological order. The above property is evaluated by considering each of the schedule dimensions in turn. If the value for b is always smaller, then the property does not hold. If it may be greater, then the property holds. Otherwise, if the value may be the same, then we continue to the next schedule dimension. If we reach the inner level, then we assume that the instances may be scheduled in any order.

5.8. Shifted Stride Detection

The stride detection of Section 5.3 only works if there is a consistent stride across the entire schedule domain. In some cases, the schedule domains for individual statements may be strided, while the combined schedule domain is not.

Example 5.15. Take an iteration domain

$$\{ A(i) : 0 \leq i < 10; B(i) : 0 \leq i < 10 \} \quad (47)$$

with schedule

$$\{ A(i) \rightarrow (2i); B(i) \rightarrow (2i + 1) \}. \quad (48)$$

and shifted linear schedules [Darte et al. 2001] may be of this form. While the schedule domains per statement have stride 2 (with offsets 0 and 1), the combined schedule domain is simply $\{ (t) : 0 \leq t < 20 \}$. If we replace the above partial schedule by the equivalent $\{ A(i) \rightarrow (2i, 0); B(i) \rightarrow (2i, 1) \}$, then the schedule domain in the outer dimension becomes $\{ (t) : 0 \leq t < 20 \wedge t \bmod 2 = 0 \}$ and the stride can be exploited.

To be able to apply the transformation illustrated in the example, we look for a statement where the current schedule dimension does not have an obviously fixed value. Note that it does not make sense to look for any strides if all of them have a fixed value. We then consider pairs of statements where the first is the selected statement and the second is any statement. If for all of these pairs we find that the *differences* between the values of the schedule dimension are of the form $m_i x + r_i$ with m_i and r_i constant values and $m_i > 1$, then we take the greatest common divisor m of the m_i and reduce the r_i modulo m . If m is greater than one and not all of the r_i are equal, then we apply the transformation $(\dots, t, \dots) \rightarrow (\dots, t - r_i, r_i, \dots)$ to the schedule domain of statement i .

Example 5.16. Continued from Example 5.15. The differences between the values of the schedule dimension for B and A in (48) are $\{ (t) : -17 \leq t \leq 19 \wedge t \bmod 2 = 0 \}$, such that $m_1 = 2$ and $r_1 = 1$. For pairs of A instances, we similarly find $m_0 = 2$ and $r_0 = 0$.

Note that this transformation is mainly meant to make our AST generator generate nice code even when presented with schedules from outside users that contain such hidden strides. The scheduler that comes with `isl` also implements a Feautrier style scheduler, but it will detect and adjust such shifted strides during the scheduling itself. On the other hand, the Pluto style scheduler [Bondhugula et al. 2008] of `isl` may occasionally construct a schedule with several *distinct* shifted strides. Detecting them during scheduling is therefore complicated, but they may still be detectable in the AST

generator since the detection is applied on different components separately. Furthermore, the transformation can also be applied if the schedule itself does not contain an explicit scaling factor but if instead the scaling can be derived from the iteration domain.

5.9. Combining Annotated ASTs

In several parts of the AST generator, we need to combine a list of annotated ASTs into a single annotated AST. This single annotated AST may be used in the same or in an outer stock. In particular, we may want to move to a stock of an outer loop, in which case the extra loop iterator needs to be projected out from the guard and the enforced condition attached to the ASTs. The shared constraints in these projections are hoisted to the single annotated AST and the original guards are simplified with respect to the hoisted guard. Note that if there is more than one element in the list of annotated ASTs, then some of the constraints may not be explicitly available in the guards of each element. We therefore first intersect the guards with both the stock domain and the enforced constraints of each individual annotated AST. The shared constraints are then still selected from the original guards, but they only need to satisfy the intersections to be considered for hoisting. In general, the shared constraints form a single disjunct set. However, as a special case, if all annotated ASTs in the list have the same guard, then we allow this guard to be hoisted even if it is described in terms of multiple disjuncts.

Example 5.17. Assume that at the outer level, we have decided to construct a for loop with bounds $\{(t) : 1 \leq t \leq 8\}$. Assume furthermore that at the inner level, we have constructed three annotated ASTs with the following guards (g_i) and enforced constraints (f_i).

$$\begin{aligned} g_1 &= \{(t) : t \leq 6\} & g_2 &= \{(t) : t \geq 2\} & g_3 &= \{(t) : t \geq 2\} \\ f_1 &= \{(t) : t \geq 2\} & f_2 &= \{(t) : t \leq 5\} & f_3 &= \{(t) : t \leq 8\}. \end{aligned} \quad (49)$$

If we only consider the guards g_i themselves, then we cannot hoist any guards since $t \leq 6$ is not satisfied by g_2 and g_3 while $t \geq 2$ is not satisfied by g_1 . After intersection with the stock domain and the corresponding enforced constraints, we find

$$g'_1 = \{(t) : 2 \leq t \leq 6\} \quad g'_2 = \{(t) : 2 \leq t \leq 5\} \quad g'_3 = \{(t) : 2 \leq t \leq 8\}. \quad (50)$$

Considering the original two constraints $t \leq 6$ and $t \geq 2$ once more, we now find that $t \geq 2$ is satisfied by all three g'_i and that we may therefore hoist this constraint, which can then be used to simplify the original guards to

$$g''_1 = \{(t) : t \leq 6\} \quad g''_2 = \{(t)\} \quad g''_3 = \{(t)\}. \quad (51)$$

The simplified guards (if non-trivial) then need to be expressed as if-nodes around the corresponding AST node. We do not, however, treat every annotated AST in the list separately. Instead, we incrementally build up a tree of if-nodes, keeping track of a list of if-nodes that can still be extended without changing the execution order. For each of these if-nodes, we keep track of the condition enforced by the node and its parents as well as the condition enforced by the corresponding else-branch. For each annotated AST in the original list, we look for the deepest if-node such that one of the attached conditions is implied by the current guard. The guard is then simplified in terms of this condition and the AST node (with an enclosing if-node if needed) is inserted in the corresponding branch. Finally, the list of if-nodes is updated to reflect the insertion.

Example 5.18. Assume an outer generated domain of

$$D = \{(i_0, i_1) : 1 \leq i_0 \leq u \wedge i_0 \leq i_1 \leq n \wedge 1 \leq l \leq u \leq n\}, \quad (52)$$

with l , u and n symbolic constants, and a list of four annotated ASTs with the following simplified guards.

$$\begin{aligned}
 g_1 &= \{ (i_0, i_1) : l \leq i_0 \wedge i_0 + 1 \leq i_1 \} & (\text{s0}) \\
 g_2 &= \{ (i_0, i_1) : i_0 < l \} & (\text{s3}) \\
 g_3 &= \{ (i_0, i_1) : l \leq i_0 \wedge i_0 + 1 \leq i_1 \wedge u < n \} & (\text{s2}) \\
 g_4 &= \{ (i_0, i_1) \} & (\text{s1})
 \end{aligned} \tag{53}$$

When we consider g_1 , the list of if-nodes is still empty, so we create a first if-node with condition $c_0 = g_1$ and complement $\bar{c}_0 = D \setminus g_1$ and add the graft to the final list of grafts. Guard $g_2 \cap D$ is not a subset of c_0 , but it is a subset of \bar{c}_0 ($i_0 < l$ is satisfied by the complement of $l \leq i_0$). We therefore simplify g_2 in the context of \bar{c}_0 , which has no effect in this case, and extend the else branch of this if-node with the second graft. We also extend the list of if-nodes with condition $c_1 = g_2$ and complement $\bar{c}_1 = \bar{c}_0 \setminus g_2$. The following guard $g_3 \cap D$ is a subset of neither c_1 nor \bar{c}_1 , but it is as subset of c_0 . We therefore simplify g_3 in the context of c_0 , which yields $g'_3 = \{ (i_0, i_1) : u < n \}$ and extend the then branch of the if-node with the third graft. Since we are extending the first element in the list of if-nodes, we drop all subsequent elements in the list of if-nodes so that they can no longer be extended. This ensures that this combination step never changes the order of the statements, even though it may be slightly too conservative. We also extend the list of if-nodes with a new second if-node with condition $c'_1 = g'_3$ and $\bar{c}'_1 = c_0 \setminus g'_3$. The final guard $g_3 \cap D$ is a subset of none of the if-node guards or complements. We therefore clear the list of if-nodes and add the graft directly to the final list of grafts. The generated code then has the following form.

```

if (c0 >= lb && c1 >= c0 + 1) {
  s0(c0, c1);
  if (n >= ub + 1)
    s2(c0, c1);
} else if (lb >= c0 + 1)
  s3(c0, c1, lb, c0, c1);
for (int c3 = max(lb, c0); c3 <= ub; c3 += 1)
  s1(c0, c1, c3);

```

5.10. Generating AST Expressions

Most of the operations performed by the AST generator are performed on `is1` data types that represent Presburger sets and relations or piecewise quasi-affine functions. It is only during the final construction of the if and for-nodes that these objects are converted to syntactic AST expressions. Internally in `is1`, quasi-affine functions are expressed in terms of greatest integer parts ($\lfloor \cdot \rfloor$). In principle, these expressions can be translated directly into their AST expression counterparts, but as explained in Section 4.3, for some use cases it is important to know if the first argument of an integer division is non-negative or if the division is exact. Moreover, we typically want an expression of the form $m \lfloor (a(i)/m) \rfloor$ to be translated to $a(i) - (a(i) \bmod m)$, provided again that $a(i)$ is non-negative.

Whenever generating an if or for-condition or a for initialization or upper bound expression from an expression involving greatest integer parts, we first check for opportunities to extract modulo expressions and then check the signs of the remaining greatest integer parts. Note that when generating a conjunction of constraints, we first generate expressions for the constraints not involving greatest integer parts and then add those constraints to the stock so that they can be used to simplify the remaining constraints. It does not matter in which order these conditions get evaluated in the generated code. We do need to consider a particular order to ensure that there are no

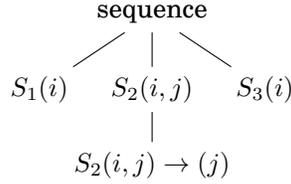


Fig. 4: Schedule tree representation of the original schedule

cycles in constraints being used to simplify other constraints. The functions for generating AST expressions from isl objects are also available to the user and can for example be used to generate an AST expression from an access function.

Example 5.19. Consider the code in Figure 5b. The constraints for the loop epilogue are of the form $n \geq 2 \wedge n - 4 \lfloor n/4 \rfloor \geq 2$. We first generate a condition for the constraint $n \geq 0$ so that it can be used to simplify the condition generated from $n - 4 \lfloor n/4 \rfloor \geq 2$.

If we are generating an equality constraint, we first check if the equality encodes a stride. If so, the stride can be expressed in the AST using an expression of the form $x \% m == 0$. In this case, the sign of x is of no importance. For other constraints or expressions in general, if we find a subexpression of the form $f m \lfloor a(i)/m \rfloor$ and we can prove that $a(i)$ is non-negative based on the constraints in the stock, then the expression is replaced by $f a(i) - f \cdot (a(i) \bmod m)$. If $a(i)$ may be negative, but $-a(i) + m - 1$ can be proved to be non-negative, then it is replaced by $f \cdot (m + 1 - a(i)) - f \cdot ((m + 1 - a(i)) \bmod m)$ instead, exploiting the fact that $\lfloor a/b \rfloor = -\lceil -a/b \rceil = -\lfloor (-a + b - 1)/b \rfloor$. Moreover, the $a(i)$ inside the argument of \bmod can be replaced by any $a'(i) = a(i) + m e(i)$. We therefore look for constraints $h(i) \geq 0$ among the shifted shared constraints of the context with coefficients that are either equal or opposite to those of $a(i)$ modulo m . Since $h(i)$ is known to be non-negative in the context, it can be used directly as $a'(i)$. If no such constraint can be found, we check if $a(i)$ or $-a(i) + m - 1$ themselves can be proved to be non-negative by solving an ILP problem. The latter test is also used to check if the first arguments of the remaining integer divisions are non-negative. These simple heuristics appear to work out fairly well in practice.

Example 5.20. Continued from Example 5.19. The constraints $n - 4 \lfloor n/4 \rfloor \geq 2$ has a subexpression of the form $f m \lfloor a(i)/m \rfloor$, with $f = 1$, $m = 4$ and $a(i) = n$. Since we have added the constraint $n \geq 2$ to the generated constraints of the stock, we can take $h(i) = n - 2$. We may therefore replace $4 \lfloor n/4 \rfloor$ by $n - (n \bmod 4)$ and obtain the constraint $n \bmod 4 \geq 2$.

6. SCHEDULE TREES

The description of the AST generator in Section 5 takes an unstructured schedule as input. Although such unstructured schedules are sufficient for basic applications, a structured schedule representation such as the schedule trees of Verdoolaege et al. [2014] can be more convenient to use. In more advanced applications such as PPCG, the constructed schedule cannot even be expressed as a (single) unstructured schedule. Our AST generator therefore also accepts such schedule trees as input.

The nodes of a schedule tree may be of different types. The main node type is the *band* node. Each band node contains an unstructured schedule of the type accepted by the AST generator in Section 5. Each band node has its own isolation domain (if any) as well as its own values for the atomic/separate/unroll choice. An ordering of groups of statements may be expressed through a *sequence* node, the children of which are *filter*

nodes that partition the iteration domain. For example, a schedule tree representation of the schedule in (2) is shown in Figure 4. Other nodes types include *set* nodes, which are similar to sequence nodes but express that the children may be executed in any order, *context* nodes, which introduce additional symbolic constants and/or additional constraints on these symbolic constants, and *extension* nodes, which extend the iteration domain in terms of the schedule domain. For a detailed discussion we refer to Verdoolaege et al. [2014]. The ability to specify AST generation options for each band separately allows for great flexibility for more complicated programs with statements that, after scheduling, are spread over different, possibly nested, loop nests.

AST generation for a schedule tree performs a depth-first traversal of the schedule tree, where Algorithm 1 is applied on each band node. Termination of this algorithm is replaced by a visit to the child of the band node in the schedule tree, while the original termination is performed at the leaves of the schedule tree. The handling of the other node types is fairly straightforward. Sequence and set nodes simply handle their children in order and combine the resulting annotated ASTs. Filter nodes and context nodes impose the given constraints on either the range of the executed relation or its symbolic constants. When annotated ASTs are combined on leaving a context node, any symbolic constants introduced by the context node are projected out from the guard and the enforced condition.

The core algorithm needs to be slightly adjusted to the presence of descendant nodes in the schedule tree. In particular, the shifted stride detection of Section 5.8 is skipped if any of the descendants (or in fact the current band node itself) refers to the schedule dimensions of outer nodes, i.e., extension nodes and band nodes with an isolation domain. For the band nodes that are nested inside other band nodes, the isolation domain is expressed in terms of the schedule dimensions of both the current and the outer band nodes. The reason for skipping shifted stride detection is that the schedule space is modified on a per statement basis, such that there is no way to express the mapping between the original and the transformed schedule space. Other changes to the schedule space are taken into account through the mapping from loop iterators to schedule dimensions in the stock.

Similarly the detection of components of Section 5.7 needs to take into account the descendants of the current node, while checking if an instance of statement b is scheduled after any instance of statement a . In particular, if the child of the band node is a sequence node and a and b are mapped to different children, then their order determines whether the property holds. If the child is a set node and a and b are mapped to different children, then the property does not hold since the instances of a and b may be scheduled in any order with respect to each other. As long as the property has not been evaluated (and therefore some instances of a and b may be scheduled together), we continue our traversal of the schedule tree. If we reach the leaves, then we assume that the instances may be scheduled in any order. Note that it is important that the entire schedule tree is available for us to be able to evaluate this property. Otherwise, we would have to conservatively assume that the property holds for instances that are co-scheduled by the known part of the schedule.

7. EXPERIMENTAL RESULTS

7.1. Verified Correctness

One of the major goals of our AST generator is that it should accept any well-formed input and that it should produce correct code. In order to test the correctness of our AST generator, we collected inputs from the CLooG and CodeGen+ distributions and verified that we produce correct code for each of them. Since the types of inputs of these tools form a strict subset of the inputs accepted by our AST generator, the inputs

```

if (n >= 2)
  for (c1 = 2; c1 <= n; c1 += 2) {
    if (c1%4 == 0)
      S1(c1);
    if ((c1+2)%4 == 0)
      S2(c1);
  }

```

(a) CLooG 0.18.1 generated code

```

for (int c0 = 2; c0 < n - 1; c0 += 4) {
  S1(c0);
  S0(c0 + 2);
}
if (n >= 2 && n % 4 >= 2)
  S1(-(n % 4) + n + 2);

```

(b) isl generated code

```

#define intMod(a,b) (a) >= 0 ? (a) % (b) ? (b) - abs((a) % (b)) % (b)
for(t1 = 2; t1 <= n; t1 += 2)
  if (intMod(t1,4) == 0)
    S0(t1);
  else
    S1(t1);

```

(c) codegen+ generated code

Fig. 5: Code for example from [Chen 2012, Figure 8(b)] (style edited)

can be easily converted to schedule trees with three nested nodes, a context, a domain and a band node. To verify the correctness of the output, we parse the output using pet [Verdoolaege and Grosser 2012] and verify that the statement instances executed by the output are the same as those specified by the input and that they are executed in an order that matches the input schedule. Moreover, if the input schedule is single-valued (which is usually the case), then we check that each statement instance is executed exactly once. We perform this test for various settings of the options that control the shape of the output. Although we are aware of work that checks that different versions of CLooG produce equivalent output [Verdoolaege et al. 2009], we are not aware of prior work that systematically verifies that the generated AST matches the input schedule.

Performing the same check on CodeGen+ generated output from CLooG inputs, we find that for 13 test cases (out of 94) CodeGen+ produces output containing N/A. This means that CodeGen+ was unable to express the statement instance in terms of the loop iterators in the scheduled code, possibly due to CodeGen+ applying the schedule as a preprocessing step and doing the AST generation on the schedule domain, which prevents the recovery of the statement instance, in case of non-injective schedules. We also found one test case (*darte*) where the generated code contains a condition on a loop iterator outside of the loop over which it iterates and one test case (*walters*) where some statement instances in the input do not appear in the generated code.³ Note that the N/A issue also appears in four of the (disabled) test cases that come with the CodeGen+ distribution. One of the original codegen test cases (*p.delft2*) also produces the error “guard condition too complex to handle”. Performing a similar test on CLooG with CodeGen+ inputs is not feasible since older versions of CLooG (prior to our enhancements) would not allow existentially quantified variables in the input.

7.2. Generated Code Quality

We illustrate the improvements in code quality of our AST generator through examples from related work. Figure 5 compares the outputs on an input with iteration domain $\{s_0(i) : \exists \alpha : 1 \leq i \leq n \wedge i = 4\alpha; s_1(i) : \exists \alpha : 1 \leq i \leq n \wedge i = 4\alpha + 2\}$ and schedule $\{s_0(i) \rightarrow (i); s_1(i) \rightarrow (i)\}$, an example reconstructed from [Chen 2012, Figure 8(b)] with the iteration space extended to negative numbers. Notice that thanks to the shifted

³These issues have been reported to the authors six months before submitting this paper for review.

<pre> for(i=1; i<=n-2; i++) { S1(i,i); S2(i,i); for(j=i+1; j<=n-1; j++) S2(i,j); S2(i,n); S3(i,n); } S1(n-1,n-1); S2(n-1,n-1); S2(n-1,n); S3(n-1,n); S1(n,n); S2(n,n); S3(n,n); for (i=n+1; i <= m; i++) S3(i,j); </pre> <p>(a) CLooG 0.14.1</p>	<pre> for(i=1; i<=m; i++) { if(i>=n +1) { S2(i,n); } else { S0(i,i); S1(i,i); if (i>=n) S2 (i,i); } for(j=i+1; j<=n-1; j++) S0(i,j); if(n >= i+1) { S0(i,n); S2(i,n); } } </pre> <p>(b) codegen+</p>	<pre> for(i=1; i<=n; i+=1) { S1(i,i); for(j=i; j<=n; j+=1) S2(i,j); S3(i,n); } for(i=n+1; i<=m; i+=1) S3(i,n); </pre> <p>(c) isl codegen</p>
---	---	---

Fig. 6: Code for youcefn taken from [Bastoul 2004, Figure 6] (style edited)

<pre> // Simple S(n % 128); // Shifted S(((t1 + 121) % 128) + 7); // Conditional if ((t1 + 121) % 128 <= 123) S(((t1 + 125) % 128) + 3); </pre> <p>(a) isl</p>	<pre> // Simple for(i = intMod(n,128); i <= 127; i += 128) S(i); // Shifted for(i = 7+intMod(t1-7,128); i <= 134; i += 128) S(i); // Conditional for(i = 7+intMod(t1-7,128); i <= 130; i += 128) S(i); </pre> <p>(b) codegen+</p>
---	--

Fig. 7: Modulo conditions (examples not supported by CLooG)

stride detection of Section 5.8, the modulo operation is removed from the inner loop in the `isl` output. During the construction of the $n \% 4 \geq 2$ condition, we exploit the fact that the body is only executed if the condition $n \geq 0$ also holds, as explained in Section 5.10. For the input in [Chen 2012, Figure 8(a)], `isl` produces the same AST as `CodeGen+`.

Figure 6c illustrates the detection of components of Section 5.7. The example is the one used in [Bastoul 2004, Figure 6]. Without the detection of components, some separation is needlessly applied as illustrated in Figure 6b for `CodeGen+` and in Figure 6a for `CLooG 0.14.1`. The code in [Bastoul 2004, Figure 7] is similar to the code that `isl` produces, but was obtained either manually or using a preliminary implementation of an “unisolate” technique that was never made public. As a further illustration, for the `darte` input, each statement only occurs twice in the `isl` output even with full separation. In the `CLooG` output, each statement occurs 5 times, and is surrounded by several modulo conditions, some of which are redundant. As explained in Section 7.1, `CodeGen+` produces incorrect output in this example.

7.3. Modulo mappings and existentially quantified variables

<pre> // Two e.q. variables for (int c0 = 0; c0 <= 7; c0 += 1) if (2 * (2 * c0 / 3) >= c0) S(c0); // Multiple bounds for (i = 0; i <= 1; i += 1) for (j = max(-(-t1 + t2 - 3)%128) + t2 - 387, t1-384); j <= min(-t2 + 127, t2 - 383); j += 128) if (j + 256 == t1 (j + 384 == t1 && t1 >= 126)) S(i, j); </pre>	<pre> // Two e.q. variables S(0); S(2); S(3); S(4); S(5); S(6); S(7); // Multiple bounds if (t1 >= 126) S(0, t1 - 384); S(0, t1 - 256); if (t1 >= 126) S(1, t1 - 384); S(1, t1 - 256); </pre>
(a) isl	(b) isl unrolled

Fig. 8: Existentially quantified variables (examples not supported by CLooG/codegen+)

Generating a valid AST for any valid Presburger relation and ensuring that we use efficient remainder operations whenever possible is one of the design goals of our AST generation algorithm. One area not sufficiently covered in Section 7.1 are existentially quantified variables, as they can result from modulo mappings from global to shared memory or from a full iteration space to a set of thread ids. We start with a simple modulo operation $\{[i] \rightarrow [i] : i = n \bmod 128\}$ to verify that modulo operations can be detected at all. Since older versions of CLooG (prior to our enhancements) do not allow existentially quantified variables, we do not compare against it in this section. For `isl` and `CodeGen+`, Figure 7 shows that `isl` uses a single statement with a remainder operation, whereas `CodeGen+` generates a loop. Using a loop is very inefficient, not only due to the call to `intMod` and the general loop overhead, but especially because the expression $n\%128$ is invariant of any possibly surrounding loop and has almost zero cost as the loop invariant code motion pass of a compiler normally moves it out of the loop body. Two slightly more complicated examples are mappings from a set of iterations to a set of threads $t1$ with $\{0 \leq t1 < 128\}$. The first mapping is the one-to-one mapping $\{[i] \rightarrow [i] : 7 \leq ii \leq 134 \wedge i \bmod 128 = t1\}$ which `isl` again translates into a single instruction, the second is the mapping $\{[i] \rightarrow [i] : 7 \leq i \leq 130 \wedge i \bmod 128 = t1\}$ which maps 124 iterations to 128 threads. `isl` lowers this mapping to a single conditional statement. `CodeGen+` generates for both cases a full loop nest. It is interesting to note that all previously shown loops are degenerate loops with just a single iteration. `CodeGen+` is not able to detect those loops, whereas `isl` is designed to always recognize degenerate loops (see Section 5.3).

For the previous test cases only a single existentially quantified variable was introduced due to the one modulo operation in the schedule. For more complex use cases, e.g., the modulo mapping of access functions that already contain modulo expressions or nested modulo mappings, it is often possible that multiple existentially quantified variables are introduced. The first test case $\{[i] \rightarrow [i] : \exists(\alpha, \beta : i = 2\alpha + 3\beta \wedge 0 \leq \alpha < 3 \wedge 0 \leq \beta \wedge 0 \leq i < 8)\}$ involves two existentially quantified variables in a single equality. `CodeGen+` aborts here with the message *guard condition too complex to handle*. In Figure 8 we see that `isl` is able to generate valid code (see Section 5.2 for details), which can be unrolled both for better efficiency and to better understand the computation that is performed. The next test case is $\{[i, j] \rightarrow [i, j] : \exists(\alpha, \beta : 0 \leq i \leq 1 \wedge t1 = j + 128\alpha \wedge 0 \leq j + 2\beta < 128 \wedge 510 \leq t2 + 2\beta \leq 514 \wedge 0 \leq 2\beta - t2 \leq 5)\}$, which was reduced from the example in Figure 1. `CodeGen+` aborts with *Can't generate multiple wildcard GEQ guards right now*. `isl` either generates a loop with multiple loop bounds and remainder conditions or, if unrolled, a set of conditional statements. As Chen [2012] does not discuss how existentially quantified variables are handled, the scope of support

AST generation options	heat 2D		heat 3D	
	GFLOPS	speedup	GFLOPS	speedup
<i>a</i> : no options enabled	1.9	1.0x	4.9	1.0x
<i>b</i> : all optimizations enabled	26.4	13.9x	19.6	4.0x
<i>c</i> ₁ : all, except full/partial separation	19.4	10.2x	18.2	3.7x
<i>c</i> ₂ : all, except IO unrolling	4.5	2.3x	9.6	2.0x
<i>c</i> ₃ : all, except compute unrolling	14.1	7.4x	10.1	2.1x
<i>c</i> ₄ : all, except modulo detection	27.5	14.1x	16.9	3.4x

Table I: AST generation strategy based performance (GFLOPS)

in CodeGen+ is unclear. When inspecting the source code of CodeGen+ we found several code paths that require a single existentially quantified variable per constraint. isl has no limitations on the number of existentially quantified variables per constraint (see Section 5.2).

7.4. Performance implications of AST generation strategies

To understand the performance implications of our new AST generation strategies, we analyze their impact on the run-time of generated code. We ensure a realistic scenario by analyzing a full end-to-end domain specific compiler. As compiler we choose the stencil compiler introduced in Section 2. We remind the reader that this compiler is based on the general purpose compiler PPCG. To create code that is optimized for the domain of stencil computations the computation of a generic execution schedule is replaced with the computation of a hybrid hexagonal/parallelogram execution schedule specifically optimized for the domain of stencil computations. Besides the domain specific schedule, the only other domain specific piece is the parametrization of our AST generator to isolate (Section 5.6) full tiles from partial as well as to unroll (Section 5.5) compute and IO code. AST expression generation (Section 5.10) is used to specialize the access functions of statements, e.g., after unrolling or separation.

We perform the evaluation on a NVIDIA NVS 5200M GPU using a heat 2D and a heat 3D stencil as benchmark. As performance results have shown large differences between two and three dimensional stencils, we choose two benchmarks to cover the most common dimensionalities. We limit ourselves to a single type of stencil, as the general tendency between different types of stencils does not vary enough to give additional insights for this analysis. For further performance results on different hardware and different stencil types, we refer to [Grosser et al. 2014]. Table I shows the results of our analysis. We see in *a* that normal AST generation with no further specialization enabled yields very low performance with just 1.9 GFLOPS in the 2D case and 4.9 GFLOPS in the 3D case compared to *b* where we enable all optimizations and obtain 26.4 GFLOPS in the 2D case and 19.6 GFLOPS in the 3D case, a 13.9x speedup in the 2D case and a 4.0x speedup in the 3D case. To understand better where this speedup comes from we individually disable certain optimizations. In *c*₁ we disable full/partial tile separation, which reduces the performance by 27% for heat-2D and 7% for heat-3D. The larger change on 2D is due to the higher percentage of full tiles. In 3D, already a large amount of time is spent in partial tiles, so optimizations that speed up the execution of full tiles are less visible. In *c*₂ and *c*₃ we see that for heat 3D disabling either unrolling of IO or unrolling of compute reduces the performance by 50%. For the 2D case, disabling unrolling of the compute code also reduces the performance by 50% and, even more importantly, without unrolling of the IO code over 80% of the performance is lost. This large performance difference is both due to the increased ILP after unrolling and because of the simplifications enabled by unrolling

(see Figure 2). In c_4 we see that without modulo detection the performance for heat-3D is reduced by 14% and, surprisingly, slightly increased by 4% in 2D. The increase for heat-2D is due to register spilling caused by loop invariant code motion which again was made possible due to the simpler code after modulo detection. Allowing `nvcc`, the NVIDIA compiler, to use more registers prevents register spilling and modulo detection becomes again beneficial with a new peak performance of 28.4 GFLOPS for heat 2D, a 8% performance improvement over the previous peak value. Overall, we see that just generating control flow using polyhedral scanning is by far not enough to generate high-performant GPU code. Instead, both polyhedral unrolling and specialization for full and partial tiles are highly important to obtain code of competitive performance.

7.5. Generation Time

Although we have mostly focused on ease of use and quality of the generated AST, for completeness we also report some AST generation times. For this experiment, we take 64 of the test cases distributed with `CLooG` (those that can be handled by both `CLooG 0.14.1` and `CodeGen+`) and sum the total AST generation time. For `CLooG 0.14.1` (before our enhancements, using `PolyLib` as a backend), we obtain 0.3s using fixed size integer computations and 1.0s for arbitrary precision integers. For `CLooG 0.18.1` (including some of our enhancements and using `isl` for set operations with arbitrary precision integers), we obtain 0.9s. For `CodeGen+`, we find 3.1s and for `isl`, 1.5s. We attribute the time difference with respect to `CLooG` to the fact that we have not yet implemented some of the heuristics of Vasilache et al. [2006] and that we are much more aggressive in our optimizations, resulting in better output code.

8. RELATED WORK

There are two major approaches to generic AST generation, one that is based on a library for Presburger relations and that focuses on lifting control overhead up [Kelly et al. 1995; Chen 2012] and one that is based on rational polyhedra and that mainly tries to eliminate overhead top-down [Bastoul 2004; Quilleré et al. 2000]. Our approach can be seen as a combination, using the same separation algorithm of [Bastoul 2004; Quilleré et al. 2000], but built on top of a library for Presburger relations. Historically, we started by porting `CLooG` to `isl` and improving `CLooG`. Later, we built a new AST generator on top of `isl`. Both the original approaches only allow single disjunct contexts and schedules, with `CLooG` also not supporting existentially quantified variables. `CodeGen+` can handle such variables in certain cases, but as Chen [2012] does not discuss how such variables are handled in general, the extend of support is unclear. In contrast, our AST generator supports the full generality of Presburger arithmetic, including existentially quantified variables and piecewise schedules.

In respect to the quality of the generated AST, some extensions proposed in the literature have not been implemented in `isl`. The components of Section 5.7 serve the same purpose as the “unisolat” procedure of Bastoul [2004, Section 4.2]. However, where the unisolat procedure tries to undo some separation, the components allow us to not even apply the separation. Moreover, no implementation of the unisolat procedure was ever made publicly available. Instead, recent versions of `CLooG` implement our components detection. Vasilache et al. [2006] propose several optimizations implemented in `CLooG` to reduce the AST generation time. Some of these optimizations are tailored to their encoding of schedule trees and are not needed when the schedule is represented as an explicit schedule tree. The same authors also propose an algorithm for removing modulo conditions, which on the one hand can be seen as a generalization of the shifted stride detection of Section 5.8, but on the other hand is based on a more restrictive polyhedral formulation. Zuo et al. [2013] describe several fine-tunings of `CLooG` tailored for high-level synthesis, only some of which are available in `isl`.

Kelly et al. [1995] and Chen [2012] as well as Quilleré et al. [2000] and Bastoul [2004] generate AST expressions as necessary to generate control flow for scanning the iteration space, but they do not expose any functionality to generate AST expressions for arbitrary user-provided piecewise quasi-affine expressions. We also are not aware of any work that uses the AST generation context to specialize AST expressions. Especially, no work that uses context information to optimize modulo operations and divisions as they appear in quasi-affine expressions. CodeGen+ always generates expensive `intMod` calls and `CLooG` only introduces a `%` operator in cases where the result of the operator is compared to zero.

Polyhedral unrolling in an AST generator has been proposed (without software being made available) by Vasilache et al. [2006] for the special case of a unimodular schedule where a dimension that has a single lower and single upper bound offset by a constant non-parametric distance can be fully unrolled. In our work we presented polyhedral unrolling for schedules defined by arbitrary Presburger maps, with support for unrolling in presence of multiple lower bounds, unrolling in the presence of strides and unrolling for loops with bounded, non-constant number of iterations using conditional statements. User-directed isolation of arbitrary subsets of the iteration space as such has not been implemented in polyhedral AST generators. The automatic separation used by Bastoul [2004] regularly introduces specialized code versions, but the user can only control the amount of separation and not the subsets that are separated from each other. Full/partial tiling has been discussed as an independent transformation by Ancourt and Irigoin [1991] as well as Goumas et al. [2003] and, combined with unrolling, by Jiménez et al. [2002]. In the context of parametric tiling [Kim et al. 2007; Renganarayanan et al. 2007; Hartono et al. 2010] full/partial tile separation has been researched in AST generators specialized for this use case. We are not aware of any work that uses a generic isolation feature provided by a polyhedral AST generator to perform full/partial tile separation. As parametric tiling techniques commonly rely on polyhedral AST generators, the same isolation techniques may be useful in the context of parametric tiling.

We are not aware of any work that provides configurability on such a fine grained level. Bastoul [2004] originally allowed per-dimension level control over separation and recently gained per-statement control. Chen [2012] allows per loop level control over the amount of control flow. Different AST generation strategies for different subtrees of the generated AST are to our knowledge unique to our work. Also, giving the user the ability to enforce an “atomic” AST generation strategy to minimize code size or to enforce unrolling is new.

9. CONCLUSION

This work significantly widens the scope of polyhedral AST generation. It does so by extending traditional control flow generation to the full generality of Presburger arithmetic. In particular, we provide support for piecewise affine schedules as well as schedules with complex uses of existentially quantified variables, opening AST generation to new application areas and more sophisticated program optimizations, and enhancing its reliability—the ability to predictably generate highly efficient control flow. Our work also improves the quality of the generated imperative code by presenting optimizations for shifted strides and components. We also acknowledge that optimization problems are not limited to control flow restructuring, but also require changes to data access functions: to support such optimizations, we propose facilities to generate efficient AST expressions from piecewise quasi-affine forms. Finally, we improve on the state of the art techniques to recover divisions and modulo expressions in the generated code, and apply these to the optimization of index expression that commonly

appear in the context of explicitly managed caches. Overall, we widened the scope of generic AST generators.

However, to implement domain or target specific optimizations that reach peak performance, it is often necessary to heavily specialize the generated code. For this we allowed the AST generator to be parameterized to perform loop unrolling and partial evaluation of loop iterators in a very general, polyhedral setting. Furthermore, we presented how to separate certain parts of the code and show how to use this separation to generate specialized code for full and partial tiles. By allowing the specialization of user-provided AST expressions according to the context they are generated in, the same feature can also be used to generate specialized code for boundary conditions. As maximal specialization may not always be best, we make AST generation choices such as separation, unrolling and also atomic execution configurable on a fine-grain level. Each individual contribution is by itself useful, but only the integration in a single AST generator ensures their seamless interaction. As demonstrated on hybrid hexagonal/classical tiling, the result is an AST generator that can be used to implement complex domain specific optimizations, outperforming domain specific compilers with specialized code generators.

Acknowledgements

This work was partly funded by a gift received by LIACS from Intel Corporation, a Google European Fellowship in Efficient Computing, by the European FP7 project CARP id. 287767 and the COPCAMS ARTEMIS project.

REFERENCES

- Corinne Ancourt and François Irigoin. 1991. Scanning Polyhedra with DO Loop. In *Symp. on Principles and Practice of Parallel Programming (PPoPP)*. 39–50.
- Vinayaka Bandishti, Irshad Pananilath, and Uday Bondhugula. 2012. Tiling stencil computations to maximize parallelism. In *Supercomputing*. IEEE Computer Society Press, 40.
- Cédric Bastoul. 2004. Code Generation in the Polyhedral Model Is Easier Than You Think. In *IEEE Int. Conf. on Parallel Architecture and Compilation Techniques*. Juan-les-Pins, France, 7–16.
- Uday Bondhugula, Albert Hartono, J. Ramanujam, and P. Sadayappan. 2008. A Practical Automatic Polyhedral Parallelization and Locality Optimization System. In *ACM SIGPLAN Conf. on Programming Languages Design and Implementation (PLDI)*. Tucson, AZ, USA.
- Chun Chen. 2012. Polyhedra scanning revisited. In *Conf. on Programming Language Design and Implementation*. ACM, New York, NY, 499–508.
- Chun Chen, Jacqueline Chame, and Mary Hall. 2008. *A framework for composing high-level loop transformations*. Technical Report. USC.
- Alain Darte, Yves Robert, and Frédéric Vivien. 2001. Loop parallelization algorithms. In *Compiler Optimizations for Scalable Parallel Systems: Languages, Compilation Techniques and Run Time Systems (LNCS)*, Vol. 1808. Springer Verlag, 141–171.
- Paul Feautrier. 1988. Parametric Integer Programming. *RAIRO Recherche Opérationnelle* 22, 3 (1988), 243–268.
- Paul Feautrier. 1992. Some efficient solutions to the affine scheduling problem. Part II. Multidimensional time. *Int. Journal of Parallel Programming* 21 (1992), 389–420. Issue 6.
- Georgios Goumas, Maria Athanasaki, and Nectarios Koziris. 2003. An efficient code generation technique for tiled iteration spaces. *Parallel and Distributed Systems, IEEE Transactions on* 10 (2003), 1021–1034.
- Tobias Grosser, Albert Cohen, Justin Holewinski, P. Sadayappan, and Sven Verdoolaege. 2014. Hybrid Hexagonal/Classical Tiling for GPUs. In *Int. Symp. on Code Generation and Optimization (CGO)*. Orlando, FL.
- Tobias Grosser, Armin Größlinger, and Christian Lengauer. 2012. Polly – Performing polyhedral optimizations on a low-level intermediate representation. *Parallel Processing Letters* 22, 04 (2012).
- Albert Hartono, Muthu Manikandan Baskaran, J Ramanujam, and Ponnuswamy Sadayappan. 2010. DynTile: Parametric tiled loop generation for parallel execution on multicore processors. In *IEEE Int. Symp. on Parallel & Distributed Processing (IPDPS)*. IEEE, 1–12.

- Tom Henretty, Richard Veras, Franz Franchetti, Louis-Noël Pouchet, J. Ramanujam, and P. Sadayappan. 2013. A stencil compiler for short-vector SIMD architectures. In *Int. Conf. on Supercomputing (ICS)*.
- Justin Holewinski, Louis-Noël Pouchet, and P. Sadayappan. 2012. High-performance code generation for stencil computations on GPU architectures. In *Int. Conf. on Supercomputing (ICS)*.
- Marta Jiménez, José M Llabería, and Agustín Fernández. 2002. Register tiling in nonrectangular iteration spaces. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 24, 4 (2002), 409–453.
- Wayne Kelly, Vadim Maslov, William Pugh, Evan Rosser, Tatiana Shpeisman, and Dave Wonnacott. 1996. The Omega Calculator and Library, Version 1.1.0. (1996).
- W. Kelly, W. Pugh, and E. Rosser. 1995. Code Generation for Multiple Mappings. In *Frontiers '95: The 5th Symposium on the Frontiers of Massively Parallel Computation*. McLean, VA.
- DaeGon Kim, Lakshminarayanan Renganarayanan, Dave Rostron, Sanjay Rajopadhye, and Michelle Mills Strout. 2007. Multi-level Tiling: M for the Price of One. In *ACM/IEEE Conf. on Supercomputing (SC '07)*. Article 51, 12 pages.
- Martin Kong, Richard Veras, Kevin Stock, Franz Franchetti, Louis-Noël Pouchet, and P. Sadayappan. 2013. When polyhedral transformations meet SIMD code generation. In *34th ACM SIGPLAN Conf. on Programming language design and implementation*. ACM, 127–138.
- Vincent Loechner and Doran K Wilde. 1997. Parameterized polyhedra and their vertices. *Int. Journal of Parallel Programming* 25, 6 (1997), 525–549.
- Louis-Noël Pouchet, Cédric Bastoul, Albert Cohen, and John Cavazos. 2008. Iterative optimization in the polyhedral model: Part II, multidimensional time. In *ACM SIGPLAN Conf. on Programming Language Design and Implementation (PLDI'08)*. ACM Press, Tucson, Arizona, 90–100.
- Louis-Noël Pouchet, Cédric Bastoul, Albert Cohen, and Nicolas Vasilache. 2007. Iterative optimization in the polyhedral model: Part I, one-dimensional time. In *IEEE/ACM Fifth Int. Symposium on Code Generation and Optimization (CGO'07)*. IEEE Computer Society press, San Jose, California, 144–156.
- William Pugh and David Wonnacott. 1994a. *An exact method for analysis of value-based array data dependences*. Springer.
- William Pugh and David Wonnacott. 1994b. Static analysis of upper and lower bounds on dependences and parallelism. *Transactions on Programming Languages and Systems (TOPLAS)* 16, 4 (1994), 1248–1278.
- Fabien Quilleré, Sanjay Rajopadhye, and Doran Wilde. 2000. Generation of Efficient Nested Loops from Polyhedra. *Int. Journal of Parallel Programming* 28, 5 (Oct. 2000), 469–498.
- Lakshminarayanan Renganarayanan, DaeGon Kim, Sanjay Rajopadhye, and Michelle Mills Strout. 2007. Parameterized tiled loops for free. In *ACM SIGPLAN Notices*, Vol. 42. ACM, 405–414.
- Jun Shirako and Vivek Sarkar. 2013. Oil and Water can mix! Experiences with integrating Polyhedral and AST-based Transformations. (2013).
- Nicolas Vasilache, Cédric Bastoul, and Albert Cohen. 2006. Polyhedral Code Generation in the Real World. In *Int. Conf. on Compiler Construction (CC)*, Vol. 3923. Springer, Vienna, 185–201.
- Anand Venkat, Manu Shantharam, Mary Hall, and Michelle Strout. 2014. Non-affine Extensions to Polyhedral Code Generation. In *Int. Symp. on Code Generation and Optimization*. Orlando, FL.
- Sven Verdoolaege. 2010. isl: An Integer Set Library for the Polyhedral Model. In *Mathematical Software (ICMS'10)*, Vol. 6327. Springer.
- Sven Verdoolaege. 2011. Counting Affine Calculator and Applications. In *First Int. Workshop on Polyhedral Compilation Techniques (IMPACT'11)*. Chamonix, France.
- Sven Verdoolaege and Tobias Grosser. 2012. Polyhedral Extraction Tool. In *Second Int. Workshop on Polyhedral Compilation Techniques (IMPACT'12)*. Paris, France.
- Sven Verdoolaege, Serge Guelton, Tobias Grosser, and Albert Cohen. 2014. Schedule Trees. In *4th Int. Workshop on Polyhedral Compilation Techniques*. Vienna, Austria.
- Sven Verdoolaege, Gerda Janssens, and Maurice Bruynooghe. 2009. Equivalence checking of static affine programs using widening to handle recurrences. In *Computer Aided Verification 21*. Springer, 599–613.
- Sven Verdoolaege, Juan Carlos Juega, Albert Cohen, José Ignacio Gómez, Christian Tenllado, and Francky Catthoor. 2013. Polyhedral parallel code generation for CUDA. *ACM Transactions on Architecture and Code Optimization* 9, 4 (Jan. 2013), 54:1–54:23.
- Wei Zuo, Peng Li, Deming Chen, Louis-Noël Pouchet, Shunan Zhong, and Jason Cong. 2013. Improving Polyhedral Code Generation for High-Level Synthesis. In *Int. Conf. on Hardware/Software Codesign and System Synthesis*.