

Automatic Distributed-Memory Parallelization and Code Generation using the Polyhedral Framework

Uday Bondhugula
Department of Computer Science and Automation
Indian Institute of Science, Bangalore 560012

IISc-CSA-TR-2011-3

Sep 2011

Abstract

Compiling for distributed-memory parallel architectures is considered very challenging. In spite of the large amount of work done to address this problem, no practical and efficient solution currently exists.

We present new techniques for compilation of regular sequential programs for distributed-memory parallel architectures – typically, a cluster of multicores interconnected with a high performance interconnect. Compilation for distributed memory requires generation of communication code and its quality is key to scalable performance. Our framework is implemented as a source-level transformer that uses the polyhedral compiler framework, and generates parallel code with communication expressed with the Message Passing Interface (MPI) library. We are able to handle sequences of arbitrarily nested loops with affine dependences, and generated code is parametric in the number of processors and problem sizes. The proposed approach has been implemented into a tool and we report experimental results on a cluster of multicores demonstrating its effectiveness. Compared to all previous approaches, ours is a significant advance either (1) with respect to the generality of input code handled, or (2) efficiency of communication code, or both. To the best of our knowledge, this is the first work reporting end-to-end fully automatic distributed-memory parallelization and code generation for input programs and transformation techniques as general as those we allow.

1. Introduction and Motivation

Shared memory for multiple processing elements is a useful abstraction for parallel programmers. However, due to limitations in scaling shared memory to a large number of processors, the compute power of shared-memory multiprocessor systems is limited. For example, it is currently hard to find shared memory systems with more than 32 or 48 cores. To get greater processing power, multiple processing nodes are connected with a high performance interconnect such as InfiniBand or 10 Gigabit Ethernet to form a cluster. Each node has its own memory space that is not visible to other nodes. The only way to share data between nodes is by sending and receiving messages over the interconnect. The Message Passing Interface (MPI) [24] is the current dominant parallel programming model used to program compute-intensive applications on such distributed-memory clusters.

Distributed memory makes parallel programming even harder from many angles. A programmer has to take care of distribution and movement of data in addition to distribution of computation. Data distribution and computation distribution are tightly coupled – changing the data distribution in a simple way often requires a complete rewrite of compute and communication code. Debugging multiple processes that send and receive messages to and from each other is also significantly more difficult. Parallelizing even simple regular loop nests for distributed memory can be very error-prone and unproductive. In

addition, whenever pipelined parallelism exists, i.e., not all processors are active to start with, or when there is a significant amount of discontinuous data to be transferred to multiple nodes, which is often the case, MPI parallelization can be a nightmare. Hence, a tool that can automatically parallelize for distributed-memory parallel architectures can provide a big leap in productivity.

In this paper, we propose techniques and optimizations for automatic translation of regular sequential programs to parallel ones suitable for execution on distributed-memory machines: typically, a cluster of multicore processors. We use the polyhedral compiler framework to accomplish this in a portable and efficient manner. As a result, we are able to handle sequences of arbitrarily nested loops with regular (affine) accesses, also known as *affine loop nests*. We would like to emphasize that distributed-memory compilation of even this restricted class of codes is very challenging and no automatic solution exists despite decades of research. Codes such as these are common in the scientific and embedded computing domains. Our contributions over previous works are at least one or more of the following: (1) handles imperfect loop nests with affine dependences, (2) significantly lesser communication in the presence of parametric problem sizes and number of processors, and (3) fully automatic end-to-end capability for distributed-memory parallelization. Experiments conducted on a 32-node InfiniBand cluster demonstrates high speedups.

The problem of distributed memory parallelization requires a solution to several sub-problems. New techniques presented in this paper are for efficient communication code generation, i.e., when a transformed parallelized code is given as input. Hence, approaches that determine computation or data partitioning are orthogonal to it. For experimental evaluation, we coupled our proposed code generation technique with a computation partitioning-driven polyhedral parallelizer Pluto [8, 26]. Our framework is thus implemented as a source-level transformer that generates parallel code using the MPI library as its communication backend. Code we generate is parametric in the number of processors and other problem sizes, and provably correct for any number of MPI processes. Besides parallelizing for distributed memory, code we generate is also optimized for execution on all cores of each node, and for locality on each core. Targeting a portable and relatively low-level communication library like MPI allows us to benefit a wide range of architectures.

The rest of this paper is organized as follows. Section 3 discusses challenges. Section 2 provides background and notation. Section 4 and Section 5 describe our solution and optimizations. Section 6 provides experimental results. Discussion of related work is presented in Section 7 and conclusions are presented in Section 8.

2. Background and Notation

The polyhedral compiler framework is an abstraction for analysis and transformation of programs. It captures the execution of a program in a static setting by representing its instances as integer points inside parametric polyhedra. Most publicly available tools and compilers that use this framework extract such a representation from C, C++, and Fortran programs.

Polyhedral representation of programs: Let S_1, S_2, \dots, S_n be the statements of the program. Each dynamic instance of a statement, S , is identified by its iteration vector \vec{i} that contains values for indices of the loops surrounding S , from outermost to innermost. Whenever the loop bounds are affine functions of outer loop indices and program parameters, the set of iteration vectors belonging to a statement form a convex polyhedron called its *domain* or *index set*. Let I_S be the index set of S and let its dimensionality be m_S . Let \vec{p} be the vector of program parameters. Program parameters are not modified anywhere in the portion of code we are trying to model.

A function f on a domain I_S is called an affine function if it can be represented in the following form:

$$f(\vec{i}) = [c_1 \dots c_{m_S}] (\vec{i}) + c_0, \vec{i} \in I_S$$

Regular data accesses in a statement are represented as multi-dimensional affine functions of domain indices. Codes that satisfy these constraints are also known as *affine loop nests*.

Polyhedral dependences: The data dependence graph (DDG) is a directed multi-graph with each vertex representing a statement, and an edge, $e \in E$, from node S_i to S_j representing a polyhedral dependence from an iteration of S_i to an iteration of S_j : it is characterized by a polyhedron, D_e , called the *dependence polyhedron* that captures exact dependence information corresponding to e . The dependence polyhedron is in the sum of the dimensionalities of the source and target iterations spaces, and the number of program parameters. At least one of the source and target accesses has to be a write.

```
for (t=0; t<=T-1; t++)
  for (i=1; i<=N-2; i++)
    for (j=1; j<=N-2; j++)
      a[i][j] = (a[i-1][j-1] + a[i-1][j] + a[i-1][j+1]
                + a[i][j-1] + a[i][j] + a[i][j+1] +
                a[i+1][j-1] + a[i+1][j] + a[i+1][j+1])/9.0;
```

Figure 1. Seidel-style code

For example, for the code in Figure 1, the dependence between the write $a[i][j]$ at $\vec{s} = (t, i, j)$ and the read at $\vec{t} = (t', i', j')$ at $a[i'-1][j'-1]$ is given by the dependence polyhedron, $D_e(\vec{s}, \vec{t}, \vec{p}, 1)$, which is a conjunction of the following equalities and inequalities:

$$\begin{aligned} i' &= i + 1, & j' &= j + 1, & t' &= t, \\ 0 &\leq t \leq T - 1, & 1 &\leq i \leq N - 3, & 1 &\leq j \leq N - 3 \end{aligned}$$

3. Problem and Challenges

When compiling for shared memory, synchronization primitives take care of preserving data dependences when dependent iterations are mapped to different processors. Shared memory support provided by hardware takes care of transparently providing data that had been written to by one processor before a synchronization point, to another one after it. However, in case of distributed-memory systems, this movement of data has to be performed in software via communication over the interconnect.

Unknown number of processors: A number of difficulties arise due to the fact that the number of processes we are compiling for is not known at compile time, and has to be treated as a parameter. This is not an issue with shared-memory auto-parallelization. OpenMP support takes care of partitioning a parallel loop with a choice of strategies, and since no software data transfer is performed, the two steps of generation of parallel code (marking a loop as parallel) and

that of distributing the parallel loop across processors are decoupled. No matter how the parallel iterations are scheduled across processors, the hardware would transparently guarantee visibility of correct data after synchronization points.

Generating code parametric in problem sizes and processors is also very important for portability. For proprietary software, a vendor may not be able provide binaries for each configuration of problem and system size, and a user would not be able to recompile for it.

Plugging in the number of processors as a parameter in the polyhedral representation does not help since it introduces non-linearity. For example, a simple loop of N iterations, when divided across $nprocs$ processors would lead to the following SPMD code (even if N to be a multiple of $nprocs$ below):

```
for (i=my_rank*N/nprocs; i<=(my_rank+1)*N/nprocs-1; i++)
```

Thus, there is no turn-key approach to get polyhedral machinery to directly compute data accessed inside such a loop. In addition, the number and identity of communication partners as well as communication data may often depend on the total number of processors as well as other program symbols, actual values of which will only be known at runtime. These are hard problems and we provide a compile-time solution to these.

4. Distributed memory code generation

In this section, we describe all steps involved in obtaining communication code given the original program and a transformation or computation partitioning for it.

4.1 Dependences and communication

When code is partitioned across multiple processing elements, any communication required arises out of data dependences. Recall that there are primarily three types of data dependences: flow (Read-after-Write or RAW), anti (Write-after-Read or WAR), and output (Write-after-Write dependences). It is interesting to contrast the effect of these dependences when compiling for shared-memory versus for distributed-memory systems. Anti and output dependences merely exist because the same memory location is being reused. In case of shared memory auto-parallelization, anti and output dependences are still important – this is because when iterations that are dependent via such a dependence are mapped to different processors, owing to the same shared memory location they access, synchronization is needed. However, in case of distributed memory, each processor has its own address space. This coupled with the fact that there is no flow of data results in anti and output dependences neither leading to any communication nor any synchronization.

Note that our goal is to generate a distributed-memory program that preserves semantics of the original sequential program. Once the parallelized portion of the input code finishes execution, all results are to be available at a single process, say, the master process. Thus, even in the absence of any dependences, communication is needed to make sure that all results will have been gathered at the master process by the time all parallel processes have finished executing. We show that this communication code can be generated efficiently using output (WAW) dependences.

A loop that can be placed at any position and marked parallel has no dependence components along it, and is called an outer parallel or a communication-free parallel loop. Outer parallelism will require no communication except a gather of results at the master process. Wherever pure inner parallelism exists, i.e., communication cannot be avoided via transformation, generating efficient communication code is crucial. Note that inner parallelism, wavefront parallelism, and pipelined parallelism can all be converted into inner parallelism, i.e., one parallel loop followed by a synchronization call when running on shared memory, or communication code in case of distributed memory.

4.2 Computing communication sets

In the rest of this section, by *tile* we refer to the portion of computation under a given iteration of the parallel dimension, i.e., all dimensions surrounding it and including itself serve as parameters for the tile. It may or not have been obtained as a result of loop tiling. It is the smallest piece of computation for which we will define communication sets. It is important to note that constraints that describe a tile's domain are affine at compile time. We classify communication data for a tile into two classes:

1. Writes to locations that are read at another process the next time the same loop is run in parallel, i.e., for another iteration of the surrounding (outer) sequential loop if any.
2. All results need to be available at a root node once parallelized computation has finished executing, i.e., final writes for all data spaces need to be aggregated.

We show that by computing two sets for each data variable, one for each of the above cases, one can determine all that has to be sent out from a process after it has finished executing a tile. We call these the *flow-out* set and the *write-out* set. Each of these sets can be a union of convex polyhedra. The integer points in these polyhedra yield actual data elements to be communicated.

Running example: We use the code in Figure 2 as an example to demonstrate all steps, showing results and code they yield at each step. This is a typical Jacobi-style stencil with time along the vertical axis and space along the horizontal. For simplicity, assume that all dimensions are tiled by a factor of 32. Tiling serves a number of purposes in our context: increasing granularity of parallelism and reducing the frequency of communication, improving locality, and bounding buffer sizes by a factor proportional to tile size where possible.

In the rest of this section, whenever we refer to a set, we mean a union of convex polyhedra with integer points enclosed by them being of interest to us. Whenever a set of linear equalities and inequalities are listed, they represent a conjunction of those. Recall notation introduced in Section 2. \setminus is used as the set difference operator. In addition, some polyhedral operations are notated as below:

- $project_out(D, p, n)$: eliminates n dimensions from set D starting from p^{th} dimension ($p \geq 1$)
- $\mathcal{I}_p(M, \mathcal{D}, l)$: image of \mathcal{D} under a multi-dimensional affine function M while treating l outermost dimensions as parameters
- $pad(D, pos, n)$: extend dimensionality of D by adding n (unconstrained) dimensions at position pos ($pos \geq 1$)

Algorithms for projection are provided by polyhedral libraries. For parametric image, the chosen parameters are not projected out of the image. In particular, if one needs data accessed for a given set of outer loops through an access function M , the outer loops are to be treated as parameters just like other program parameters appearing in loop bounds, and M is used as the function for the image operation. As an example, if \mathcal{D} is

$$\begin{aligned} 1 \leq i \leq N-1, \quad 1 \leq j \leq N-1 \\ 32i_T \leq i \leq 32i_T + 31, \quad 0 \leq j - 32j_T \leq 31 \end{aligned}$$

and $M = (i-1, j-1)$, $l = 1$. Then, $\mathcal{I}_p(M, \mathcal{D}, l)$ is

$$\begin{aligned} 0 \leq d_0 \leq N-2, \quad 32i_T \leq d_0 \leq 32i_T + 31 \\ 0 \leq d_1 \leq N-2 \end{aligned}$$

4.2.1 Flow-out set

The *flow-out* set of a tile is the set of all values that are written to inside the tile, and then next read from outside the tile. Temporary sets used in the algorithm have the following meanings. Let e be a RAW dependence from S_i to S_j .

D_e^T : dependence polyhedron for edge e in the transformed space
 $I_{S_i}^T, I_{S_j}^T$: domains of S_i and S_j in the transformed space with dimen-

sionalities $m_{S_i}^T$ and $m_{S_j}^T$ respectively

I_e^t : iterations inside the tile reading data written to in the tile as a result of dependence e

O_e^t : source iterations of the tile whose writes are read by iterations outside the tile through dependence edge e

t_k^i : k^{th} dimension of $I_{S_i}^T$

The first step is that of finding the subset of the transformed dependence polyhedron that has both its source and target iterations in the same tile. This can be obtained by intersecting D_e^T with a set of equalities equating the first l dimensions of $I_{S_i}^T$ to those of $I_{S_j}^T$. Let E_l be that set of equalities, i.e.,

$$E_l = \left\{ t_1^i = t_1^j \wedge t_2^i = t_2^j \wedge \dots \wedge t_l^i = t_l^j \right\}$$

We now obtain the set of all iterations of the tile that write to values that are (later) read within the same tile through dependence edge e :

$$\begin{aligned} C_e^t &= D_e^T \cap E_l \\ I_e^t &= project_out(C_e^t, m_{S_i} + 1, m_{S_j}) \end{aligned}$$

Next, subtracting I_e^t from the set of all source dependence iterations in the tile yields those source dependence iterations whose writes are read outside the tile:

$$O_e^t = project_out(D_e^T, m_{S_i} + 1, m_{S_j}) \setminus I_e^t$$

Now, computing the image of the source write access function, M_w , on O_e^t yields the flow-out set for this particular write access and dependence.

$$F_{out}^x = \mathcal{I}_p(M_w^{S_i}, O_e^t, l)$$

Algorithm 1 computes the entire flow-set for a particular variable. Transformed dependence polyhedra and transformed index sets can be generated by taking the original ones and augmenting them with transformation functions that map old iterators to new ones, and then projecting out the old ones. This yields the dependence relation between iterations in the target space D_e^T . Since anti and output dependences

Algorithm 1 Computing flow-out set for variable x

INPUT Depth of parallel loop: l ; set \mathbf{S}_w of (write access, statement) pairs for variable x

- 1: $F_{out}^x = \emptyset$
- 2: **for each** $\langle M_w, S_i \rangle \in \mathbf{S}_w$ **do**
- 3: **for each** dependence $e(S_i \rightarrow S_j) \in E$ **do**
- 4: **if** e is of type RAW and source access of e is M_w **then**
- 5: $E_l = \left\{ t_1^i = t_1^j \wedge t_2^i = t_2^j \wedge \dots \wedge t_l^i = t_l^j \right\}$
- 6: $C_e^t = D_e^T \cap E_l$
- 7: $I_e^t = project_out(C_e^t, m_{S_i} + 1, m_{S_j})$
- 8: $O_e^t = project_out(D_e^T, m_{S_i} + 1, m_{S_j}) \setminus I_e^t$
- 9: $F_{out}^x = F_{out}^x \cup \mathcal{I}_p(M_w^{S_i}, O_e^t, l)$
- 10: **end if**
- 11: **end for**
- 12: **end for**

OUTPUT F_{out}^x

are ignored, multiple copies of the same location may exist with different processors. However, when a later read to the same location happens, the correct written value would end up being transferred due to the presence of a flow edge between the write and the read.

4.2.2 Write-out set

The *write-out* set of a tile is the set of all those data elements to which the last write access across the entire iteration space is performed in the tile. We compute this by looking for any WAW edges leaving the tile. If they do, subtracting the sources of those edges from the set of all points written to in the tile in an iterative manner across all WAW deps leaves us with locations that have been "finalized"

```

for (i=1; i<=T-1; i++){
  for (j=1; j<=N-1; j++){
    u[i][j] = 0.333*(u[i-1][j-1]
      + u[i-1][j] + u[i-1][j+1]);
  }
}

```

Figure 2. Jacobi-style code

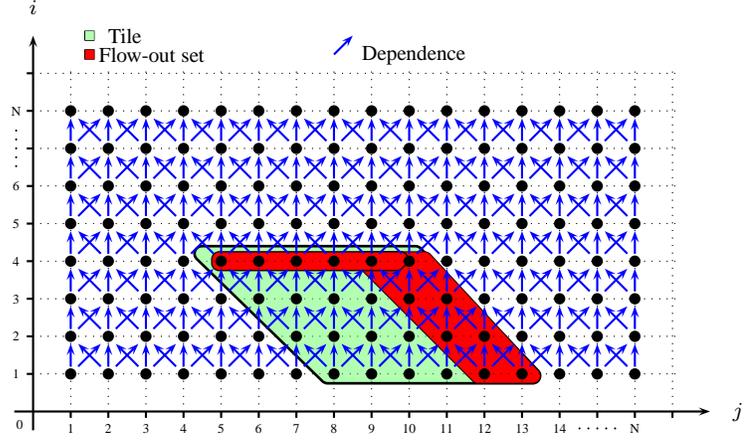


Figure 3. Flow-out set for code in Figure 2

by computation in the tile. A union has to be taken across all write accesses to a given variable in a tile. For edge e associated with variable x , let:

M_s, M_t : source and target write access functions respectively for edge e ,

I_e^t : iterations that write to locations that were written to inside the tile,

Q_e^t : iterations outside the tile that will again write data to locations that were written to inside the tile, and

W : write-out set due to a given write access, and

W_{out}^x : write-out set for x .

Algorithm 2 computes the write-out set for a variable.

Algorithm 2 Computing write-out set for variable x

INPUT Depth of parallel loop: l ; set \mathbf{S}_w of (write access, statement) pairs for variable x

```

1:  $W_{out}^x = \emptyset$ 
2: for each  $\langle M_w, S_i \rangle \in \mathbf{S}_w$  do
3:    $W = \mathcal{I}_p(M_w, I_{S_i}^T, l)$ 
4:   for each dependence  $e(S_i \rightarrow S_j) \in E$  do
5:     if  $e$  is of type WAW and source access of  $e$  is  $M_w$  then
6:       Let  $M_t$  be the target access of  $e$ 
7:        $I_{S_j}^{T'} = \text{pad}(I_{S_j}^T, m_{S_i} + 1, m_{S_j})$ 
8:        $I_e^t = D_e^T \cap I_{S_j}^{T'}$ 
9:        $I_e^{t'} = \text{project\_out}(I_e^t, l + 1, m_{S_i})$ 
10:       $Q_e^t = I_e^{t'} \setminus I_{S_j}^T$ 
11:       $W = W \setminus \mathcal{I}_p(M_t, Q_e^t, l)$ 
12:     end if
13:   end for
14:    $W_{out}^x = W_{out}^x \cup W$ 
15: end for

```

OUTPUT W_{out}^x

4.2.3 Example

For the code in Figure 2, Figure 3 shows the flow-out set. It is obtained as a union of the following two polyhedra:

$$\begin{aligned}
1 \leq i \leq T - 2, \quad 1 \leq j \leq N - 2 \\
32i_T + 30 \leq d_0 + d_1 \leq 32i_T + 31 \\
32i_T \leq d_0 \leq 32i_T + 31
\end{aligned}$$

$$\begin{aligned}
i = 32 * i_T + 31, \quad 1 \leq i \leq T - 2 \\
1 \leq j \leq N - 1, \quad 32i_T \leq d_0 + d_1 \leq 32i_T + 31
\end{aligned}$$

The second one corresponds to the horizontal line, while the first to the two oblique lines. d_0 and d_1 will index the array dimensions in

the copy-out code. Note that for simplicity, the above constraints are expressed in terms of source iterators. They are actually computed in the space of transformed iterators, i.e., in terms of (t_1, t_2) where $t_1 = i, t_2 = i + j$, since tiling has been performed here after a skewing of the space dimension. As for the write-out set, all writes that occur in a tile here are last writes, and they all need to be sent out.

4.3 Packing and unpacking communication sets

With MPI, it is easy to transfer data from, and receive into, contiguous buffers. However, in nearly all cases, we require discontinuous data to be sent and copied back on the receiver side. Hence, after the above communication sets are computed, one has to, (1) pack data to be sent in a contiguous buffer, (2) map to communication library calls, (3) unpack data at receiver side, and (4) determine send and receive buffer sizes for allocation. We construct additional statements to add to the polyhedral representation of the source program for the copy-out, copy-back. The flow-out and write-out sets serve as the domains for the copy statements.

Reasonably tight upper bounds on send and receive buffer sizes can be determined from tile constraints; we do not present details on it here due to space constraints. Write-out sets are gathered at the master process – in our case, this can be chosen to be the MPI process with rank 0.

A naive approach: A naive approach that works is to send the flow-out set to all processors, i.e., each processor sends its flow-out set to all processors. This means that all of the data to be sent is sent, but not just to the processors that need them. Hence, a processor may receive more data than necessary, and a processor that need not receive any data may receive some; the latter is a matter of greater concern. However, this approach provides a very clean way to generate communication code. Two of MPI's collectives, `MPI_Allgatherv` (all-to-all broadcast) and `MPI_Gatherv` perfectly fit. `Allgather` can be used to broadcast flow-out sets to all processors. The `Gather` call with process 0 as root is used to collect write-out sets.

5. Optimizing communication code

Recall that communication sets were defined per ‘tile’, for which we have affine constraints at compile time. This allowed us to use polyhedral machinery to compute them in the first place. Multiple such tiles may get mapped to a single physical processor and communication is done only after all of these tiles have been executed, for every iteration of immediately surrounding sequential loop, if any.

5.1 Precise determination of communication partners

The naive scheme described in the previous section broadcasts flow-out sets to all processors. In cases, where we have inner parallelism, depending on communication latencies and bandwidth, this will likely lead to a bottleneck. Recall again that the problem in determining

communication partners was that the allocation of tiles to processors is not known at compile time. Consider the simple scenario when the number of communication partners itself depends on the total number of processors. Long dependences may traverse any number of processors. However, in many cases such as in the presence of uniform dependences, only near-neighbor communication is needed. Even in these cases, if iteration spaces are shaped peculiarly, one cannot predict near-neighbor communication just based on dependence distances. Hence, even for uniform dependences, the number and identity of communication partners cannot be determined at compile time.

We describe a solution below that achieves the following: the flow-out set is not sent to processors that do not need any value from this flow-out set. More precisely, we guarantee the following:

1. Every element in the flow-out set sent by a processor is needed by at least one other processor
2. Only processors that expect to receive at least one value from another processor receive the flow-out set

We define two functions as part of the output code for each data variable, x , that can be a multidimensional array or a scalar. If t_1, \dots, t_l is the set of sequential dimensions surrounding parallel dimension t_p , the functions are:

1. $\pi(t_1, t_2, \dots, t_l, t_p)$: rank of processor that executes $(t_1, t_2, \dots, t_l, t_p)$
2. $\sigma_x(t_1, t_2, \dots, t_l, t_p)$: set of processors that need the flow-out set for data variable x from the processor calling this function

Generating π and σ : Code for π and σ functions is meant to be generated and added to output code. Constructing π is straightforward. It only requires the lower and upper bound expressions for t_p , and the number of processors. π is also used in computing σ_x . σ_x can be expressed as follows.

$$\sigma_x(t_1, t_2, \dots, t_l, t_p) = \{ \pi(t'_1, t'_2, \dots, t'_l, t'_p) \mid \exists e \in E \text{ on } x, D_e^T(t_1, \dots, t_p, \dots, t'_1, \dots, t'_p, \dots, \vec{p}, 1) \}$$

σ can be constructed as follows for each variable x . For each relevant RAW dependence polyhedron in the transformed space, we eliminate all dimensions that are inner to t_p . We then scan the dependence polyhedron to generate loops for the target iterators while treating source iterators as parameters, i.e., running the generated loop nest at run-time will enumerate all dependent tiles, $(t'_1, t'_2, \dots, t'_l, t'_p)$, given the coordinates of the source tile. However, our goal is not to enumerate dependent tiles, but to determine processors they are mapped to. Hence, σ makes use of π to aggregate a set of distinct values corresponding to processor ranks that the target tiles were mapped to. The overhead of evaluating σ at runtime is minimal since a call to it is made only once per all computation for a given t_p .

Send-synchronous scheme: With σ and π functions, generating more accurate communication code for a parametric number of processors now becomes possible. Processes send out data to the set of processor ranks returned by σ , and receivers are *forced* to receive them. The receivers will use received data in one or more future iterations. Hence, sends and receives are posted in a synchronous manner, relatively speaking. Non-blocking sends and receives are used so that simultaneous progress is made on all sends across all data variables when possible. We wait for their completion (`MPI_Waitall`) and copy-back received data to the right place before the next iteration of the sequential loop outer to the parallel one starts. Communication code can now either be expressed with point-to-point sends and receives to processor ranks obtained from σ .

An arbitrary allocation: A powerful feature of this scheme is that an arbitrary π function can be used. So far, we have only alluded to a block scheduling of the parallel loop. However, π can be generated to achieve a block-cyclic scheduling, or any other model-driven or even dynamic scheduling. In addition, it is easy to use a multidimensional π whenever we have more than one parallel dimension. Such mappings

to higher dimensional processor spaces achieve better computation to communication ratios, for a given number of processors and problem size.

5.2 Transitivity in dependences

We know that dependences that lead to communication of flow-out sets are RAW dependences. If RAW dependences are transitively covered by other dependences, one would end up communicating from several sources instead of just the last one, i.e., the one that writes last to the location as per the original program. We use a dependence tester that can compute such last writers or the exact data flow so that RAW dependence polyhedra do not contain any redundancy.

Interestingly, one observes the converse of the above effect when dealing with write-out sets. Note that transitivity can be eliminated from write-after-write dependences as well. A WAW dependence can be covered by other WAW dependences as well as through a combination of RAW and WAR dependences. If transitivity is eliminated for WAW dependences and with Algorithm 2 only looking at WAW deps, one would miss writes that happen outside the tile and to the same locations written to in a tile. This leads to a write-out set much larger than the actual one, often, almost the entire set of locations written to. Hence, Algorithm 2 will only be exact if all transitively covered WAW dependences are preserved. In summary, one has to rely on the right dependence testing and analysis techniques. ISL [16] provides functions to compute last writers as desired as part of its dependence testing interface and we use it.

In spite of the above, with our scheme, the amount of data communicated is not optimal when different parts of the flow-out set have different σ s, i.e., different lists of receiving processors. An optimal decomposition of flow-out sets if at all possible at compile time is left for future research.

5.3 Putting all communication code together

We compute domains, schedules for the additional copy-out, copy-back, and communication statements, and they have all components in the polyhedral representation just like the original compute statements. The new program comprising these added statements is given to the code generator, to generate final code in one pass.

Implementation: Our framework is implemented as part of the publicly available source-to-source polyhedral tool chain. Clan [7], ISL [16], Pluto [26], and Cloog-isl [11] are used to perform polyhedral extraction, dependence testing, automatic transformation, and code generation, respectively. The Pluto scheduling algorithm [8, 9] is first used to determine a parallelizing transformation, i.e., a computation partitioning. To implement polyhedral operations for all computations in Section 4 and Section 5, Polylib [27] was used. A powerful feature of our framework is that it will work with any other algorithm for transformation and detection of parallelism. As a result of using the polyhedral representation as both the input and output of our scheme, code can be generated after any sequence of valid transformations have been applied. Nothing prevents our communication code generation scheme to be used in a system that specifies data and computation distributions in a different way as long as these mappings can be expressed as affine functions.

Data distribution free: The distributed memory parallelization tool chain as described above is computation driven. Data moves from processor to another in a manner completely determined by the computation partitioning and data dependences. An initial data distribution can be specified, but it would only affect communication at the start. There exists no owning processor for data. However, last write communication can be modified easily to gather last writes at the right processors based on a user-supplied distribution. At the moment, no pragmas, directives, or distributions are provided to our system, i.e., it is fully automatic. Also, generated code is SPMD. We choose to gather all results at process '0' only to provide exactly the same behavior as the unmodified sequential input program.

5.4 Data allocation

We do not address the issue of data allocation in this paper. We assume that each processor has the entire data space for a variable before it starts executing the portion of the program that was parallelized. It would however only work on portions of the data space it needs to – based on the computation partitioning. Though this is a problem on a very large distributed memory machine, it is not a limitation of the code generation scheme itself, i.e., the problem can easily be addressed in an incremental manner. Keeping a view of the entire space was a natural first choice since our input is an unmodified sequential program. [31] is a recent work that addresses this issue while performing distributed memory parallelization for restricted input, those with uniform dependences. We plan to incorporate such a scheme in the future.

5.5 Improvement over previous schemes

In this section, we describe in detail how our scheme improves over existing ones for communication code generation. We consider three past works that subsume others in the literature. These are that of Amarasinghe and Lam [2], Adve and Mellor-Crummey [1], and Classen and Griehl [10]. The above schemes have the following limitations that we have overcome partly or fully:

1. All three approaches used a virtual processor to physical processor mapping to deal with symbolic problem sizes and number of processors. Communication finally occurs between virtual processors that do not map to the same physical processor. In spite of this, if multiple receiving virtual processors map to the same physical processor and data being sent to two or more of these is not disjoint, the receiving physical processor ends up receiving necessary data multiple times. For example, if a virtual processor V_i is mapped to physical processor P_i and two other virtual processors V_{j1} and V_{j2} are both mapped to physical processor P_j ($P_i \neq P_j$), and $V_i(P_i)$ sends the same data or a large portion of the same data to both $V_{j1}(P_j)$ and $V_{j2}(P_j)$. Avoiding it is not trivial since one has to look for commonality in data being sent out across a set of receiving virtual processors as well as determine the list of receivers – these are only known at runtime if the number of processors and problem sizes are parametric. The sigma function-based solution presented in this section provided an efficient solution to this problem.
2. [2, 1] determine communication sets by directly looking at read and write accesses as opposed to data dependences. Communication is only needed between the last write before a read and the read, and algorithms presented in those works do not appear to consider this issue. Since our approach relies on dependences, this requirement is easily captured in the lastwriter property of flow dependences.

Note that the second limitation also compounds redundancy created due to the first. Not eliminating transitive relations leads to more one-to-many patterns and such one-to-many patterns that in turn leads to greater redundant communication with a simple virtual to physical processor model. The approach of Classen and Griehl [10] does not suffer from the second limitation since it is based on dependences like ours. However, their work was preliminary and conceptual, and reported very limited implementation and experimental evaluation. Communication polytopes are constructed for *each* flow dependence, and so communication code is generated dependence-wise. Since communication sets for multiple dependences may often refer to the same values, a new source of redundant communication is added.

6. Experimental evaluation

Setup: We conducted experiments on a 32-node InfiniBand cluster of dual-SMP Xeon servers. Each node comprises two quad-core Intel Xeon E5430 2.66 GHz processors with a 12 MB L2 cache and 16 GB of main memory. The InfiniBand host adapter is a Mellanox MT25204. All run Linux 2.6.18 64-bit. MVAICH2-1.4 [25] (MPI

Benchmark	Problem size
strmm	10000
tmm	8000
dsyr2k	4096
covcol	N = 8192
seidel	N = 10000, T = 600
jac-2d	N = 10000, T = 1000
fdtd-2d	N = 6000, T = 256
2d-heat	N = 10000, T = 1000

Table 1. Problem sizes used

over InfiniBand) is the MPI implementation used. On this cluster, it provides a point-to-point latency of 3.36 μ s, unidirectional and bidirectional bandwidths of 1.5 GB/s and 2.56 GB/s respectively. All codes were compiled with Intel C/C++ compiler (ICC) version 11.1 with option `-fast` (implies `-O3 -ipo -static` on 64-bit Linux). Portland Group’s compiler `pglhp` 12.1 (with `-O4 -Mmpi`) was used where a comparison with HPF was performed – it was the only publicly available HPF compiler we could find.

Input sequential code without any modification is taken in by our system and compilable MPI code is generated fully automatically in all cases. The entire framework runs fast and the increase in source-to-source transformation time due to distributed memory compilation is less than 1.5s in all cases. We thus did not pay particular attention to optimize compilation time at this point.

Benchmarks: We evaluate performance on selected commonly used routines and applications from dense linear algebra and stencil computations. All of these can be found in the Polybench suite [30]. Heat-2D is available with the Pochoir suite [29]. All computations use double-precision floating point operations. Problem sizes used are given in Table 1. All results are with strong scaling.

Comparison: Regarding experimental comparison with previous approaches, we were unable to find a publicly available system that could perform such code generation. A number of techniques from the literature only a address part of the problem, and rebuilding an end-to-end system with them is infeasible. We believe that the detailed discussion provided in Section 5.5 and related work demonstrates our contributions. Comparison is thus provided with manually parallelized MPI versions of these codes, and with HPF where possible.

Though each node has eight cores and our tool is able to generate MPI+OpenMP code, in order to focus on the distributed-memory part, we run only one OpenMP thread per process, and one MPI process per node. In these figures, `our-commopt` refers to our tool with the optimization described in Section 5. `our-allgather` refers to the basic all-to-all broadcast-based communication scheme described in Section 4.3. `manual-mpi` refers to hand-parallelized MPI version of the code we developed. `seq` refers to original code compiled with `icc` with flags mentioned earlier.

For the first four codes that exhibit outer parallelism, the only communication that occurs is that of write-out sets. We see close to ideal speedup for these. Results with `our-allgather` and `manual-mpi` are not shown since they would yield the same performance. Due to all of these codes involving non-rectangular iteration spaces, manual parallelization still involves significant effort. `pglhp` was unable to correctly compile HPF versions of these – further experimentation revealed that non-rectangularity was the most likely cause.

Figure 4 show GFLOPs performance and scalability on the cluster for codes that do incur flow-out communication as well. All x -axes are on a logarithmic scale. Table 2 shows the actual execution times and reports speedup factors. For `seidel`, the original loop nest has no parallel loops. Our approach includes automatic application of such a transformation and then performing distributed memory code generation. With approaches such as HPF, this code cannot be parallelized unless the programmer manually transforms it first before providing additional directives. Performing manual MPI parallelization for it is extremely cumbersome, even without tiling the time loop.

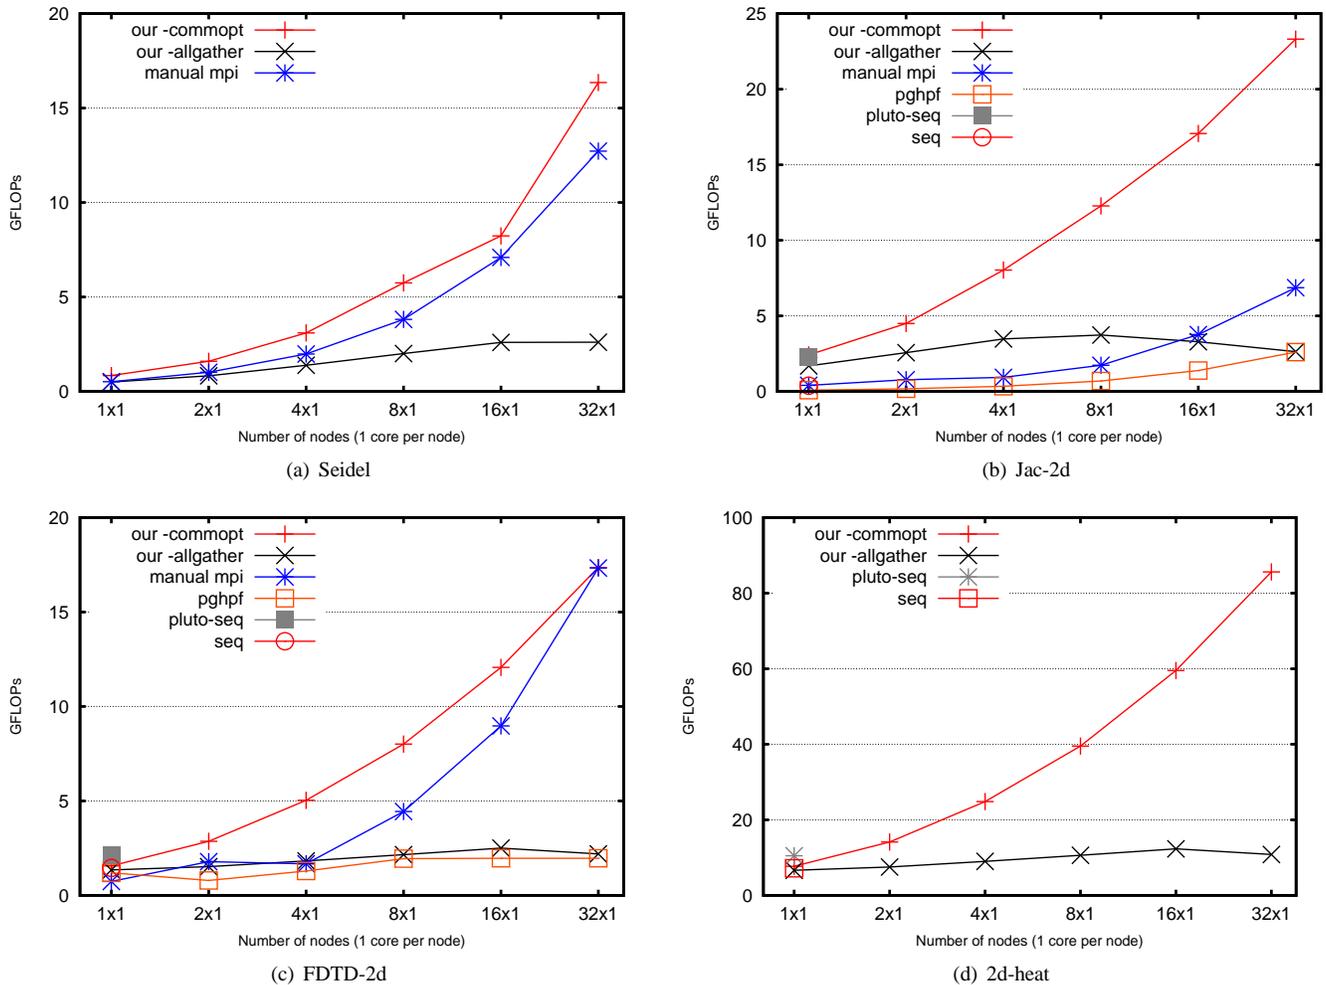


Figure 4. Performance of parallelized code on a 32-node cluster

Benchmark	seq (icc)	pluto-seq	Execution time for our-commopt (number of procs)							Speedup: our-commopt-32 over	
			1	2	4	8	16	32	seq	our-commopt-1	
strmm	30.4m	247s	240s	124.6	63.5s	33.6	17.3	9.4s	194	26.3	
tmm	35.5m	91.8s	96.4s	51.3s	27.4s	15.28s	7.14s	3.74s	570	24.5	
dsyr2k	127s	39s	38.8s	22.4s	13.5s	6.8s	3.8s	1.57s	80.8	24.7	
covcol	462s	30.9s	30.7s	16.7s	8.8s	4.6s	2.48	1.3s	355	23.8	
seidel	17.3m	643.5s	692s	338.7s	174.34s	94s	65.58s	33.03s	31	20.8	
jac-2d	21.9m	206.7s	218s	111.2s	62.27s	40.73s	29.3s	21.45s	61.3	9.6	
fdtd-2d	139s	129.7s	95.2s	70.7s	40.26	25.3s	16.8s	11.68s	11.9	11	
2d-heat	19m	266s	280s	157s	81s	52s	33s	24s	47.5	11.7	

Table 2. Summary of performance

As can be seen, automatically generated code performs much better as a result of it being fully tiled (both space and time dimensions) which in turn leads to better locality and a reduced frequency of communication. It realizes a pipelined parallelization of 3-d tiles. The same is also true for `jac-2d` as well as `fdtd-2d`, improved locality and reduced frequency of communication leads to a better solution. This explanation is also supported by the fact that ‘manual-mpi’ exhibits super-ideal improvement when going from 16 to 32 processors (for `fdtd-2d`), and in general performs relatively better with higher number of processors – a decrease in working set size hides poor locality for ‘manual-mpi’. Manually parallelized code for `jac-2d` performs significantly poorer due to lesser computation per communication call when compared to `fdtd-2d` for example. Our code shows uniformly good scalability throughout. Being able to perform distributed mem-

ory code generation in conjunction with complex transformations is thus a key strength of our tool.

Figure 5 shows the split between compute time and other overhead, i.e., time spent in communication and in packing to and unpacking from communication buffers. Results from our fully optimized codes (‘our-commopt’) were used for this plot.

7. Related work

Several attempts have been made at achieving distributed memory parallelization. Most works [4, 2, 3, 5, 14] addressed the problem in a limited way with the following limitations: (1) applicable to restricted input such as perfectly nested loops with uniform dependences, (2) address only a few steps of the actual parallelization and code gener-

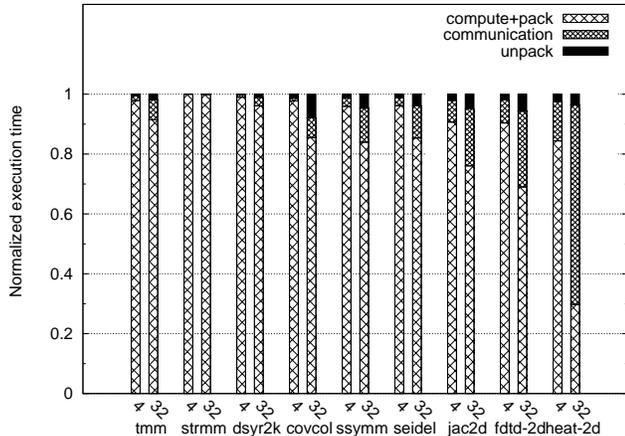


Figure 5. Breakdown of compute, communication, and pack/unpack times while running code generated by our scheme on 4 and 32 processors

ation problem, (3) lead to excessive communication with a symbolic number of processors or problem sizes.

Researchers have looked at the steps of data decomposition and computation decomposition while addressing distributed memory compilation [28, 3, 17]. Computation transformation approaches in the polyhedral framework have themselves undergone advances through [13, 21, 20, 8] that result in better parallelization for shared memory. Note that the affine partitioning works related to SUIF [21, 20] do not address distributed-memory code generation – they are transformation and parallelization algorithms. The Pluto scheduling algorithm has been shown to be a significant improvement on those [8, 9], and we use it to apply parallelizing and locality enhancing transformations before communication code generation techniques described in this paper are applied. As we have shown, generating efficient communication code on top of any auto-transformation algorithm involves a number of non-trivial problems. Without a good scheme, even the best computation partitioning is unlikely to provide good parallel speedup.

Since the primary contributions of this paper are related to communication code generation as opposed to computation or data transformations, the closely related works from literature are those of Amarasinghe and Lam [2], Adve and Mellor-Crummey [1], and Classen and Griebel [10]. [2] handled only perfectly nested loops, while [1] and [10] are based on the polyhedral framework. As explained in Section 5.5, all of these works result in a significantly large amount of redundant communication than ours, in particular with parametric problem sizes and number of processors. However, dHPF [23] implements a number of optimizations (such as multipartitioning [12]) that are useful for any distributed-memory compilation system. Our system does not implement such an allocation scheme yet, but can do so. The discussion at the end of Section 5.1 makes this evident.

Griebel [15] provides a discussion on distributed-memory auto-parallelization using the polyhedral framework. The work proposes a technique for scheduling and allocation keeping distributed memory architectures in mind. However, communication code generation is not discussed.

Works that translate OpenMP to MPI address a subset of problems that we addressed. The latest among them is [19]. Unlike our work, it is restricted to a subset of affine loop nests that transfer the same set of data every invocation of the parallel loop, and communication set construction is primarily done at runtime. In addition, with OpenMP to MPI approaches, one may have to provide an optimized/transformed OpenMP code to get good performance, adding significant complexity to input taken in by such systems. A future comparison with it if available will be interesting.

Baskaran et al. [6] presented a compiler-assisted dynamic scheduling scheme that constructs and schedules the inter-tile dependence graph on a multicore. Our communication code optimization scheme in Section 5.1 can be viewed as a compiler-assisted scheme to determine communication partners at runtime. Kim et al. [18] present automatic pipelined parallelization for distributed memory with speculation. Their scheme is completely orthogonal to ours in the kind of codes it is applicable to and beneficial for, and the way parallelism is extracted. The RSTREAM compiler provides some support for distributed memory execution [22]. However, due to its reliance on PGAS as its target instead of a message passing one, it does not have to deal with communication code generation. One would expect this to result in higher communication overhead – a comparison would have nevertheless been interesting if it was available.

8. Conclusions

We presented techniques and optimizations for translation of sequential affine loop nests to code suitable for execution on distributed-memory parallel architectures. Communication code generation and optimizations to minimize associated overhead were the key problems addressed. The scheme we proposed constructs communication sets while completely relying on data dependences. Helper routines generated by the compiler by scanning dependence relations and evaluation of those routines at runtime provided an efficient way to determine communication partners in the presence of symbolic problem sizes or number of processors, and for arbitrary allocations. These techniques were developed within a polyhedral abstraction of the input program allowing sequences of complex transformations to be automatically applied before code is generated. We have implemented them in a source-to-source transformation tool for an end-to-end fully automatic application. Experiments conducted on a 32-node InfiniBand cluster demonstrated good results.

Given how difficult manual parallelization for distributed memory is, we believe our tool will be very useful. Directly targeting a widely implemented message-passing communication backend like MPI makes our system very portable and in a position to transparently benefit from advances in message passing hardware and software. It will be publicly available shortly.

References

- [1] V. S. Adve and J. M. Mellor-Crummey. Using integer sets for data-parallel program analysis and optimization. In *PLDI*, pages 186–198, 1998.
- [2] S. P. Amarasinghe and M. S. Lam. Communication optimization and code generation for distributed memory machines. In *PLDI*, pages 126–138, 1993.
- [3] J. Anderson, S. Amarasinghe, and M. Lam. Data and Computation Transformations for Multiprocessors. In *ACM SIGPLAN PPOPP*, pages 166–178, July 1995.
- [4] J. M. Anderson and M. S. Lam. Global optimizations for parallelism and locality on scalable parallel machines. In *PLDI*, pages 112–125, 1993.
- [5] P. Banerjee, J. A. Chandy, M. Gupta, J. G. Holm, A. Lain, D. J. Palermo, S. Ramaswamy, and E. Su. The paradigm compiler for distributed-memory multicomputers. *IEEE Computer*, 28(10):37–47, 1995.
- [6] M. Baskaran, N. Vydyanathan, U. Bondhugula, J. Ramanujam, A. Rountev, and P. Sadayappan. Compiler-assisted dynamic scheduling for effective parallelization of loop nests on multicore processors. In *ACM SIGPLAN PPOPP*, pages 219–228, 2009.
- [7] C. Bastoul. Clan: The Chunky Loop Analyzer.
- [8] U. Bondhugula, M. Baskaran, S. Krishnamoorthy, J. Ramanujam, A. Rountev, and P. Sadayappan. Automatic transformations for communication-minimized parallelization and locality optimization in the polyhedral model. In *ETAPS CC*, Apr. 2008.
- [9] U. Bondhugula, A. Hartono, J. Ramanujam, and P. Sadayappan. A practical automatic polyhedral program optimization system. In *PLDI*, June 2008.

- [10] M. Classen and M. Griebel. Automatic code generation for distributed memory architectures in the polytope model. In *IEEE IPDPS*, Apr. 2006.
- [11] CLooG: The Chunky Loop Generator. <http://www.cloog.org>.
- [12] A. Darte, J. Mellor-Crummey, R. Fowler, and D. Chavarría-Miranda. Generalized multipartitioning of multi-dimensional arrays for parallelizing line-sweep computations. *JPDC*, 63:887–911, Sep 2003.
- [13] P. Feautrier. Some efficient solutions to the affine scheduling problem: Part I, one-dimensional time. *IJPP*, 21(5):313–348, 1992.
- [14] G. I. Goumas, N. Drosinos, M. Athanasiaki, and N. Koziris. Message-passing code generation for non-rectangular tiling transformations. *Parallel Computing*, 32(10):711–732, 2006.
- [15] M. Griebel. *Automatic Parallelization of Loop Programs for Distributed Memory Architectures*. University of Passau, 2004. Habilitation thesis.
- [16] Integer Set Library. Sven Verdoolaege, An Integer Set Library for Program Analysis.
- [17] K. Kennedy and U. Kremer. Automatic data layout for distributed-memory machines. *ACM TOPLAS*, 20(4):869–916, 1998.
- [18] H. Kim, A. Raman, F. Liu, J. W. Lee, and D. I. August. Scalable speculative parallelization on commodity clusters. In *MICRO*, pages 3–14, 2010.
- [19] O. Kwon, F. Jubair, R. Eigenmann, and S. Midkiff. A hybrid approach of openmp for clusters. In *ACM SIGPLAN PPOPP*, pages 75–84, 2012.
- [20] A. Lim, G. I. Cheong, and M. S. Lam. An affine partitioning algorithm to maximize parallelism and minimize communication. In *ACM ICS*, pages 228–237, 1999.
- [21] A. Lim and M. S. Lam. Maximizing parallelism and minimizing synchronization with affine partitions. *Parallel Computing*, 24(3-4):445–475, 1998.
- [22] B. Meister, A. Leung, N. Vasilache, D. Wohlford, C. Bastoul, and R. Lethin. Productivity via automatic code generation for pgas platforms with the r-stream compiler. In *Workshop on Asynchrony in the PGAS Programming Model*, 2009.
- [23] J. Mellor-Crummey, V. Adve, B. Broom, D. Chavarría-Miranda, R. Fowler, G. Jin, K. Kennedy, and Q. Yi. Advanced optimization strategies in the Rice dHPF compiler. *Concurrency: Practice and Experience*, pages 741–767, 2002.
- [24] MPI: A Message-Passing Interface Standard - version 2.2. <http://www.mpi-forum.org/docs/>.
- [25] MVAPICH: MPI over InfiniBand, 10 GigE/iWARP and RoCE. <http://mvapich.cse.ohio-state.edu>.
- [26] PLUTO: A polyhedral automatic parallelizer and locality optimizer for multicores. <http://pluto-compiler.sourceforge.net>.
- [27] PolyLib - A library of polyhedral functions. <http://icps.u-strasbg.fr/polylib/>.
- [28] J. Ramanujam and P. Sadayappan. Compile-time techniques for data distribution in distributed memory machines. *Parallel and Distributed Systems, IEEE Transactions on*, 2(4):472–482, 1991.
- [29] Y. Tang, R. A. Chowdhury, B. C. Kuzmaul, C.-K. Luk, and C. E. Leiserson. The pochoir stencil compiler. In *SPAA*, pages 117–128, 2011.
- [30] Polybench. <http://polybench.sourceforge.net>.
- [31] T. Yuki and S. Rajopadhye. Canonic multi-projection: memory allocation for distributed memory parallelization. Number CS11-106, Sept. 2011.

