Compiler Support for Parallel Program Performance Prediction *

Rizos Sakellariou† Vikram Adve‡

1 Introduction

Recent years have seen significant strides in individual performance modeling techniques for parallel applications and systems, both in analytical techniques and simulation. Several research projects are now exploring how these techniques can be brought together to support the design of large-scale applications. The POEMS project [4] aims to create a performance modeling environment for the end-to-end modeling of large parallel applications on complex parallel and distributed systems. The wide range of modeling techniques supported by POEMS, and the goal of integrating multiple modeling paradigms make it challenging, if not impossible, for the end-user to generate the required workload information manually for a large-scale application. Thus, since the conception of the project, it has been deemed essential to use compiler support to simplify and partially automate the process of constructing the workload information. To achieve this, we have designed a common task graph based program representation that provides a *uniform* platform for capturing the parallel structure of a program as well as its associated workloads for different modeling techniques.

Task graphs and their equivalents have proved to be a valuable abstraction for representing the execution of parallel programs in a number of different applications [3, 20, 24, 26, 15, 5, 14, 27, 23, 21]. Although the task graphs used in these systems differ in representation and semantics (e.g., whether task graph edges capture purely precedence constraints or also dataflow requirements), there are close similarities. Perhaps most importantly, they all capture the parallel structure of a program separately from the sequential computations, by breaking down the program into computational "tasks", precedence relations between tasks, and (in some cases) explicit communication or synchronization operations between tasks. If task graph representations could be constructed automatically, via compiler support, for common parallel programming standards such as Message-Passing Interface (MPI), High Performance Fortran (HPF), and OpenMP, the techniques and systems described above would become available to a much wider range of programs than they are currently.

Within the context of the POEMS project [4], we have developed a task graph based application representation that is used to support modeling of the end-to-end performance characteristics of a large-scale parallel application on a large parallel system, using a combination of analytical, simulation and hybrid models, and models at multiple levels of resolution for individual components. This paper describes how parallelizing compiler technology can be used to automate the process of constructing this task graph representation for HPF programs compiled to MPI (and, in the near future, for existing MPI programs directly). In particular, this paper makes the following contributions:

- We describe compiler techniques to derive a static, *symbolic* task graph representing the MPI code generated for a given HPF program. A key aspect of this process is the use of symbolic integer sets and mappings to capture a number of dynamic task instances or edge instances as a single node or edge at compile time. These techniques allow the compiler to describe sophisticated computation partitionings and communication and synchronization patterns in symbolic terms.

- We describe how standard analysis techniques can be used to condense the task graph and simplify

*A version of this paper has been presented at the *13th International Workshop on Languages and Compilers for Parallel Computing (LCPC'00)* and is due to appear in the proceedings, to be published as volume 2017 in Springer-Verlag's Lecture Notes in Computer Science (©Springer-Verlag 2001).

†Rizos Sakellariou is with the Department of Computer Science, University of Manchester, Oxford Road, Manchester M13 9PL, U.K. Email: rizos@cs.man.ac.uk

‡Vikram Adve is with the Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL 61801, U.S.A. Email: vadve@cs.uiuc.edu

control flow, whenever less than fully detailed information suffices (as in many performance modeling applications in practice).

- Finally, we describe an approach to instantiate a dynamic task graph representation from the static task graph, based on a novel use of code generation from symbolic integer sets.

In addition to the above, the compiler also uses standard techniques to compute symbolic scaling functions for task computation times and message communication volumes.

The techniques described above have been implemented in the Rice dHPF compiler system, which compiles HPF programs to MPI for message-passing systems using aggressive techniques for computation partitioning and communication optimization [1, 6]. This implementation was recently used in a joint project with the parallel simulation group at UCLA to improve the scalability of simulation of message passing programs [5]. In that work, we showed how compiler information captured in the task graph can be used to reduce the memory and time requirements for simulating message-passing programs in detail. In the context of the present paper, these results illustrate the potential importance of automatically constructing task graphs for widely used programming standards.

The next section briefly describes the key features of our static and dynamic task graph representations. Section 3 is the major technical section, which presents the compiler techniques described above to construct the task graph representations. Section 4 provides some results about the structure of the compiler-generated task graphs for simple programs and illustrates how task graphs have been used to improve the scalability of simulation, as mentioned above. We conclude with a brief overview of related work (Section 5) and a discussion of future plans (Section 6).

2 Background: The Task Graph Representation

This representation uses two flavors of a task graph, the *static task graph* and the *dynamic task graph*. Its design is described in detail in [2] and is briefly summarized here. The specific information we aim to collect for a given program includes: (1) The detailed computation partitioning and communication structure of the program, described in symbolic terms. (2) Source code for individual tasks to support source-code-driven uses such as detailed program-driven simulation of memory hierarchy performance. (3) Scaling functions that describe how computation and communication scale as a function of program inputs and processor configuration. (4) Optionally, the detailed dynamic behavior of the parallel program, for a specified program input and processor configuration.

The Static Task Graph: The static task graph (STG) captures the static parallel structure of a program and is defined only by the program *per se*. Thus, it is independent of runtime input values, intermediate program results, and processor configuration. Each node (or task) of the graph may represent one of the following main types: control flow statements such as loops and branches, procedure calls, communication, or pure computation. Edges between nodes may denote control flow within a processor or synchronization between different processors (due to communication tasks). For example, the STG for a simple parallel program is shown in Figure 1, and is explained in more detail in the next section.

A key aspect of the STG is that each node represents a set of instances of the task, one per processor that executes the task at runtime. Similarly, an edge in the STG actually represents a set of edge instances connecting pairs of dynamic node instances. We use symbolic integer sets to describe the set of instances for a given node, e.g., a task executed by P processors would be described by the set: $\{[p] : 0 \leq p \leq P-1\}$, and symbolic integer mappings to describe the edge instances, e.g., an edge from a `SEND` task on processor p to a `RECV` task on processor $q = p + 1$ (i.e., each processor sends data to its right neighbor, if any) would be described by the mapping: $\{[p] \rightarrow [q] : q = p + 1 \wedge 0 \leq p < P - 1\}$. This kind of mapping enables precise symbolic representations of arbitrary regular communication patterns. Irregular patterns (i.e., data-dependent patterns that cannot be determined until runtime) have to be represented as an all-to-all communication, which is the best that can be done statically.

To capture high level communication patterns where possible (e.g., shift, pipeline, broadcast, etc.) we group the communication operations in the program into related groups, each describing a single "logical communication event". A communication event descriptor, kept separate from the STG, captures all information about a single communication event. This includes the communication pattern, the set of communication tasks involved, and a symbolic expression for the communication size. The CPU

components of each communication event are represented explicitly as communication tasks in the STG, allowing us to use task graph edges between these tasks to explicitly capture the synchronization implied by the underlying communication calls. The number of communication nodes and edges depends on the communication pattern and also on the type of message passing calls used. This technique does not work for MPI receive operations that use a wildcard message tag (because the matching send cannot be easily identified). It does work for receive operations that use a wildcard for the sending processor, but the symbolic mapping on the communication edges may be an all-to-all mapping (for the processors that execute the send and receive statements). Making the communication tasks explicit in the STG has proved valuable also because it allows us to describe arbitrary interleavings (i.e., overlaps) of communication and computation tasks on individual processors and across processors.

In addition to the symbolic sets and mappings above, each node and communication event in the STG includes a symbolic scaling function that describes how the task computation time or the message size scales as a function of program variables. Finally, note that the STG of a program containing multiple procedures is represented as a number of unconnected graphs, each corresponding to a single procedure. Each call site is represented by a CALL task that identifies the called procedure by name.

The Dynamic Task Graph: The dynamic task graph (DTG) is a directed acyclic graph that captures the execution behavior of a program on a given input and given processor configuration. This representation is important for detailed performance modeling because it corresponds closely with the actual execution behavior being modeled by a particular program performance model (whether using detailed simulation or abstract analytical models).

The nodes of a dynamic task graph are computational tasks and individual communication tasks. In particular, the DTG does not contain control flow nodes (loops, branches, jumps, and jump targets). It can be thought of as being instantiated from the static task graph by unrolling all the loops, resolving all the branches, and instantiating all the instances of parallel tasks, edges, and communication events.

There are two approaches to making this representation tractable for large-scale programs, and these approaches can be combined: (1) we can condense tasks allocated to a process between synchronization points so that only (relatively) coarse-grain parallel tasks are explicitly represented, and (2) if necessary, we can compute the dynamic task graph "on the fly," rather than precomputing it and storing it offline. We describe techniques to automatically condense the task graph in Section 3.2. The approach to instantiate the task graph on-the-fly is outside the scope of this paper, but is a direct extension of the compile-time instantiation of the DTG described in Section 3.3.

3 Compiler Techniques for Synthesizing the Task Graphs

As noted in the Introduction, there are three major aspects to synthesizing our task graph representation for a parallel program: (1) synthesizing the static, symbolic task graph; (2) condensing the task graph; and (3) optionally instantiating a dynamic task graph representing an execution on a particular program input. Each of these steps relies on information about the message-passing program gathered by the compiler, although for many programs the third step can be performed purely by inspecting the static task graph, as explained in Section 3.3. These steps are described in detail in the following three subsections.

3.1 Synthesizing the Static Task Graph

Four key steps need be performed in synthesizing the static task graph (STG): (1) generating computation and control-flow nodes; (2) generating communication tasks for each logical communication event; (3) generating symbolic sets describing the processors that execute each task, and symbolic mappings describing the pattern of communication edges; and (4) eliminating excess control flow edges.

Generating computation and control-flow nodes in the STG can be done in a single preorder traversal of the internal representation for each procedure; in our case, the representation is an Abstract Syntax Tree (AST). STG nodes are created as appropriate statements are encountered in the AST. Thus, program statements, such as `DO`, `IF`, `CALL`, `PROGRAM/FUNCTION/SUBROUTINE`, `STOP/RETURN`, trigger the creation of a single node in the graph; encountering one of the first two also leads to the creation of an ENDDO-NODE or an ENDIF-NODE, a THEN-NODE and an ELSE-NODE, respectively. Any contiguous sequence of other computation statements *that are executed by the same set of processors* are grouped into a single

```
CHPF$ DISTRIBUTE A(*,BLOCK)
      DO I=2,N
         DO J=1,M-1
            A(I,J) = A(I-1,J+1)
         ENDDO
      ENDDO
```

(a) HPF source code fragment

```
blk = block size per processor
DO I=2,N
    IF (myid < P-1)
        irecv B(i, myid*blk+blk+1) from myid+1

    ! Execute local iterations of j-loop
    DO J=myid*blk+1, min(myid*blk+blk-1, M-1)
        A(I,J) = A(I-1,J+1)
    ENDDO

    IF (myid > 0) isend B(i, myid*blk+1) to myid-1
    IF (myid < P-1) wait-recv
    IF (myid > 0) wait-send

    ! Execute non-local iterations of j-loop
    J=myid*blk+blk
    IF (J <= M-1)
        A(I,J) = A(I-1,J+1)
ENDDO
```

(b) Unoptimized MPI code generated by dHPF

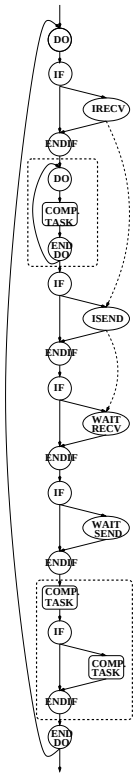

(c) Static task graph

Figure 1: An example of generating the communication tasks.

computational task (contiguous implies that they are not interrupted by any of the above statements or by communication).

Identifying statements that are computed by the same set of processors is a critical aspect of the above step. This information is derived from the computation partitioning phase of the compiler and is translated into a symbolic integer set [1] that is included with each task. By having a general representation of the set of processors associated with each task, our representation can describe sophisticated computation partitioning strategies. The explicit set representation also enables us to check equality by direct set operations, in order to group statements into tasks. These processors sets are also essential for the fourth step listed above, namely eliminating excess control-flow edges between tasks, so as to expose program parallelism. In particular, a control flow edge is retained between two tasks *only if* the intersection of their processor sets is not empty. Otherwise, the sink node is connected to its most immediate ancestor in the STG for which the result of this intersection is a non-empty set.

When the first communication statement for a logical communication event is encountered, the communication event descriptor and all the communication tasks that are pertinent to this single event are built. The processor mappings for the synchronization between the tasks are also built at this time.

For an explicit message-passing program, the computation partitioning information can be derived by analyzing the control-flow expressions that depend on process id variables. The communication pattern information has to be extracted by recognizing the communication calls syntactically, analyzing their arguments, and identifying the matching send and receive calls. In principle, both the control-flow and the communication can be written in a manner that is too complex for the compiler to decipher, and some message passing programs will probably not be analyzable. But in most of the programs we have looked at, the control-flow idioms for partitioning the computation and the types of message passing operations that are used are fairly simple. We believe that the required analysis to construct the STG would be feasible with standard interprocedural symbolic analysis techniques [17].

To illustrate the communication information built by the compiler, consider the simple HPF example, which, along with the MPI parallel code generated by the dHPF compiler, are shown on the left-hand side of Figure 1. The parallelization of the code requires the boundary values of array `A` along the j dimension

to be communicated inside the `I` loop. (In practice, the compiler pipelines the communication in larger blocks by strip-mining the `I` loop [18] but that is omitted to simplify the example.) The corresponding STG is shown on the right-hand side of the figure. Solid lines represent control flow edges and dashed lines represent interprocessor synchronization. In this example, the compiler uses the non-blocking MPI communication primitives. Also, the compiler interleaves the communication tasks and computation so as to overlap waiting time at the `isend` with the computation of local loop iterations, i.e., the iterations that do not read or write any off-processor data. The use of explicit communication tasks within the task graph allows this overlap to be captured precisely in the task graph. The dashed edge between the `isend` and the `wait-recv` tasks is associated with the processor mapping: $\{[p_0] \rightarrow [q_0] : q_0 = p_0 - 1 \wedge 0 \leq q_0 < p - 1\}$, denoting that each processor receives data from its "right" neighbor, except the rightmost boundary processor. The other dashed edge has the inverse mapping, i.e., $q_0 = p_0 + 1$.

Finally, the compiler constructs the symbolic scaling functions for each task and communication event, using direct symbolic analysis of loop bounds and message sizes. For a communication event, the scaling function is simply an expression for the message size. For each DO-NODE the scaling function describes the number of iterations executed by each processor, as a function of processor id variables and other symbolic program variables. In the simple example above, the scaling functions for the two DO nodes are `N-1` and `min(myid*blk+blk-1,M-1)-(myid*blk+1)+1`, respectively. For a computational task, the scaling function is a single parameter representing the workload corresponding to the task. At this stage of the task graph construction, no further handling (such as symbolic iteration counting for more complex loop bounds), takes place.

3.2 Condensing Nodes of the Static Task Graph

The process described above produces a first-cut version of the STG. For many typical modeling studies of parallel program performance, however, a less detailed graph will be sufficient. For instance, a coarse-grain modeling approach could assume that all operations of a single process between two communication points constitute a single task. In order to add this functionality, the compiler traverses the STG and marks contiguous nodes, connected by a control-flow edge, that do not include communication. Such sequences of nodes are then collapsed and replaced in the STG by a single *condensed* task. Note that such a task will have a single point of entry and a single point of exit. For example, the two large dotted rectangles in Figure 1 (c) correspond to sequences of tasks that can be collapsed into a single condensed task. To preserve precision, the computation of the scaling function of the new condensed task is of particular importance. Ignoring conditional control flow for the moment, this scaling function is the symbolic sum of the scaling functions of the individual collapsed tasks, each multiplied by the *symbolic number of iterations of the surrounding loops* [25], where appropriate (only if these loops are also collapsed).

In cases with conditional control flow, tasks can sometimes be condensed with no loss of accuracy, using sophisticated compiler analysis. For instance, no accuracy will be lost in cases where all dynamic instances of the resulting collapsed task have identical computational time; that is, the workload expressions are free of any input parameters whose value changes for different instances of that task. For example, in the ASCI benchmark Sweep3D [4] used in Section 4, the one significant branch is in fact of the first type, which can be pulled out of the task and enclosing loops (the analysis would have to be interprocedural because the enclosing loops are not in the same procedure as the branch). In other cases, condensing would result in some loss of accuracy, and the goals of the modeling study should be used to dictate the degree to which tasks are collapsed together.

3.3 Instantiating the Dynamic Task Graph

As noted in Section 2, the dynamic task graph (DTG) is essentially an instantiation of the STG representing a single execution for a particular input. The DTG is an acyclic graph containing no control-flow nodes. The time for instantiating the DTG grows linearly with the number of task instances in the execution of the program, but much less computation per task is usually required for the instantiation than for the actual execution. This is an optional step that can be performed when required for detailed performance prediction.

The information required to instantiate the DTG varies significantly across programs. For a regular, non-adaptive code, the parallel execution behavior of the program can usually be determined directly from the program input (in which we include the processor configuration parameters). In such cases, the DTG can be instantiated directly from the STG once the program input is specified. In general, and

particularly in adaptive codes, the parallel execution behavior (and therefore the DTG) may depend on *intermediate computational results* of the program. For example, this could happen in a parallel n-body problem if the communication pattern changed as the positions of the bodies evolved during the execution of the program. In the current work, we focus on the techniques needed to instantiate the DTG in the former case, i.e., that of regular non-adaptive codes. These techniques are also valuable in the latter case, but they must be applied at runtime when the intermediate values needed are known. The issues to be faced in that case are briefly discussed later in this section.

There are two main aspects to instantiating the DTG: (1) enumerating the outcomes of all the control flow nodes, effectively by unrolling the DO nodes and resolving the dynamic instances of the branch nodes; and (2) enumerating the dynamic instances of each node and edge in the STG. These are discussed in turn below. Of these, the second step is significantly more challenging in terms of the compile-time techniques required, particularly for sophisticated message passing programs with general computation partitioning strategies and communication patterns.

Interpreting control-flow in the static task graph Enumerating the outcomes of all the control flow nodes in an execution can be accomplished by a symbolic interpretation of the control flow of the program for each process. First, we must enumerate loop index values and resolve the dynamic instances of branch conditions that have not been collapsed in the STG. This requires evaluating the values of these symbolic expressions. We can perform this evaluation directly *at compile-time* when these quantities are determined solely by input values, surrounding loop index variables, and processor id variables. Under these conditions, we know all the required variable values in the expressions, as follows. The input variable values are specified externally. The loop index values are explicitly enumerated for all DO nodes that are retained in the static task graph. The processor id variables are explicitly enumerated for each parallel task using the symbolic processor sets, as discussed below. Therefore, we can evaluate the relevant symbolic expressions for enumerating the control-flow outcomes.

We assumed above that key symbolic quantities were determined solely by input values, surrounding loop index variables, and processor id variables. These requirements only apply to those loop bounds and branch conditions that are retained in the collapsed static task graph (i.e., which affect the parallel task graph structure of the code), and *not* to loops and branches that have been collapsed because they only affect the internal computational results of a task. With the exception of a few common algorithmic constructs, we find these requirements to be satisfied by a fairly large class of regular scientific applications. For example, in a collection of codes including the three NAS application benchmarks (SP, BT and LU), an ASCI benchmark Sweep3D [4], and other standard data-parallel codes such as Erlebacher [1] and the SPEC benchmark Tomcatv, the sole exceptions were terminating conditions testing convergence in the outermost timestep loops. In such cases, we would rely on the user to specify a fixed number of time steps for which the program performance would be modeled.

More generally, and particularly for adaptive codes, we expect the parallel structure to depend on intermediate computational results. This would require generating the DTG on the fly, e.g., when performing program-driven simulation during which the actual computational results would be computed. In this case, the most efficient approach to synthesizing the DTG would be to use program slicing to isolate the computations that do affect the parallel control flow. (This is very similar to the use of slicing for optimizing parallel simulation as described in Section 4.) These extensions are outside the scope of this paper.

Enumerating the symbolic sets and mappings The second challenge is that we must enumerate all instances of each parallel task and each communication edge between tasks. These instances are described by symbolic sets and mappings respectively. In the context of complex computation partitioning strategies and arbitrary communication patterns, this presents a much more difficult problem at compile-time (i.e., without executing the actual program).

We have developed a novel and general solution to this problem that is based on an unusual use of code generation from integer sets. In ordinary compilation, dHPF and other advanced parallelizing compilers use code generation from integer sets to synthesize loop nests that are executed *at runtime*, e.g., for a parallel loop nest or for packing and unpacking data for communication [1, 7, 8, 11]. If we could invoke the same capability but execute the generated loop nests *at compile-time*, we could use the synthesized loop nests to enumerate the required tasks and edges. Implementing this approach, however, proved to be a non-trivial task. Most importantly, each of the sets is parameterized by several variables, including

	tomcatv		jacobi		expl	
Lines of source HPF program	227		64		94	
Lines of output parallel MPI program	1850		1156		3722	
	1st pass	condensed	1st pass	condensed	1st pass	condensed
Total number of tasks	247	193	122	83	225	174
# COMM-NODE	54	54	36	36	114	114
# DO-NODE	18	3	13	1	17	3
# COMP-TASK	39	20	16	5	23	6

Table 1: Size of STG for various example codes before and after condensing.

input variables and loop index variables. This means that the set must be enumerated separately for each combination of these variable values that occurs during the execution of the original program. We solve this problem as follows. We first generate a subroutine for each integer set that we want to enumerate, and make the parameters arguments of the subroutine. Then (still at compile-time), we compile, link, and invoke this code in a separate process. The desired combinations of variable values for each node and edge are automatically available when interpreting the control-flow of the task graph as described earlier. Therefore, during this interpretation, we simply invoke the desired subroutine in the other process to enumerate the ids for a node or the id pairs for an edge.

Overall, the use of code generated from symbolic sets enables us to support a broad class of computation partitionings and communication patterns in a uniform manner. This approach fits nicely with the core capabilities of advanced parallelizing compilers.

4 Results

We have successfully implemented the techniques described above in the dHPF compiler. We have extended the dHPF compiler to synthesize a static task graph for the MPI code generated by dHPF, including the symbolic processor sets and mappings for communication tasks and edges, and the scaling functions for loop nodes. In computing the condensed static task graph, we collapse all DO-NODEs or sequences of computational tasks that do not contain any communication or any IF-NODEs (We would rely on user intervention to collapse IF-NODEs). We also compute the combined scaling function for the collapsed tasks.

We have also partially implemented the support to instantiate dynamic task graphs at compile-time. In particular, we are able to enumerate the task instances and control-flow edges. We also synthesize the code from symbolic integer-sets required to enumerate the edge mappings at compile-time. We do not yet link in this code to enumerate the edges at compile-time.

Because of the aggressive computation partitioning and communication strategies used by dHPF, capturing the resulting MPI code requires the full generality of our task graph representation. This gives us confidence that we can synthesize task graphs for a wide range of explicit message-passing programs as well (including all the ones we have examined so far).

In order to illustrate the size of the static task graph generated and the effectiveness of condensing the task graph, Table 1 lists some particulars for the STG produced by the dHPF compiler for three HPF benchmarks: Tomcatv (from SPEC92), jacobi (a simple 2D Jacobi iterator PDE solver), and expl (Livermore Loop # 18). The effect of condensing the task graph on reducing the number of loops (DO-NODE) and computational tasks (COMP-TASK) can be observed. After condensing, most of the remaining tasks are either IF-NODEs and dummy nodes (e.g., ENDIF-NODE, etc.) or communication tasks (which are never condensed), since we opted for a detailed representation of communication behavior, rather than compromise on the accuracy of the representation. The compiler generated task graphs for the above can be found at `http://www.cs.man.ac.uk/~rizos/taskgraph/`

The most important application of our compiler-synthesized task graphs to date has been for improving the state of the art of parallel simulation of message-passing programs [5]. Those results are briefly summarized here because they provide the best illustration of the correctness and benefits of the compiler-synthesized task graphs. This work was performed in collaboration with the parallel simulation group at UCLA, using MPI-Sim, a direct-execution parallel simulator for MPI programs [10].

Benchmark	% Error in prediction vs. measurement				
	#Procs = 4	8	16	32	64
Tomcatv	−5.44	15.75	11.79	8.50	9.27
Sweep3D	−7.01	−4.97	9.02	9.80	5.13
NAS SP, class A	−2.59		−1.24	7.11	6.10
NAS SP, class C			0.09	−14.01	−1.58

Table 2: Validation of the compiler-generated task graphs using MPI-Sim.

The basic strategy in using the STG for program simulation is to generate an abstracted MPI program from the STG where all computation corresponding to a computational task is eliminated, *except those computations whose results are required* to determine the control-flow, communication behavior, and task scaling functions. We refer to the eliminated computations as *redundant computations* (from the point of view of performance prediction), and we use the task scaling functions to generate simple symbolic estimates for their execution time. The simulator can avoid simulating the redundant computations, and simply needs to advance the simulation clock by an amount equal to the estimated execution time of the computation. The simulator can even avoid allocating memory for program data that is referenced only in redundant computations.

The key to implementing this strategy lies in identifying non-redundant computations. To do so, we must first identify the values in the program that determine program performance. These are exactly those variable values that appear in the control-flow expressions, communication descriptors, and scaling functions (both for task times and for communication volume) of the STG. Thus, using the STG makes these values very easy to identify. We can then use a standard program slicing algorithm [19] to isolate the computations that affect these values. We then generate the simplified MPI code by including all the control-flow that appears in the static task graph, all the communication calls, and the non-redundant computations identified by program slicing. All the remaining (i.e., redundant) computations in each computational task are replaced with a single call to a special simulator delay function which simply advances the simulator clock by a specified amount. The argument to this function is a symbolic expression for the estimated execution time of the redundant computation. Note that the simulator continues to simulate the communication behavior in detail.

We have applied this methodology both to HPF programs (compiled to MPI by the dHPF compiler), and also to existing MPI programs (in the latter case, generating the abstracted MPI program by hand). The benchmarks include an HPF version of Tomcatv, and MPI versions of Sweep3D (a key ASCI benchmark) and NAS SP. Table 2 shows the percentage error in the execution times predicted by MPI-Sim using the simplified MPI code, compared with direct program measurement. As can be seen, the error was less than 16% in all cases tested, and less than 10% in most of these cases. This is important because the simplified MPI program can be thought of as simply an executable representation of the static task graph itself. These results show that the task graph abstraction very accurately captures the properties of the program that determine performance. We believe that the errors observed could be further reduced by applying more sophisticated techniques for estimating the execution time of redundant computations, particularly with simple estimates of cache behavior.

The benefits of using the task graph based simulation strategy were extremely impressive. For these benchmarks, the optimized simulator requires factors of 5 to 2000 less memory and up to a factor of 10 less time to execute than the original simulator. These dramatic savings allow us to simulate systems and problem sizes 10 to 100 times larger than is possible with the original simulator. Also, they have allowed us to simulate MPI programs for parallel architectures with hundreds of processors *faster than real-time*, and have made it feasible to simulate execution of programs on machines with 10,000+ processors. These results are described in more detail in [5].

5 Related Work

There is a large body of work on the use of task graphs for various aspects of parallel systems but very little work on synthesizing task graphs for general-purpose parallel programs. The vast majority of performance models that use task graphs as inputs generally do not specify how the task graph should be

constructed but assume that this has been done [3, 15, 20, 26]. The various compiler-based systems that use task graphs, namely PYRROS [27], CODE [23], HENCE [23], and Jade [21] construct task graphs by assuming some kind of information from the programmer.

The PlusPYR project [12] has developed a task graph representation that has some similarities with ours along with compiler techniques to synthesize these task graphs. The key difference from our work is that they start with a limited class of *sequential* programs (annotated to identify the tasks) and use dependence analysis to compute dependences between tasks, and then derive communication and synchronization rules from these dependences. Therefore, their approach is essentially a form of simple automatic parallelization. In contrast, our goal is to generate task graphs for existing parallel programs with no special program annotations and with explicit communication. A second major difference is that they assume a simple parallel execution model in which a task receives all inputs from other tasks in parallel and sends all outputs to other tasks in parallel. In contrast, we capture much more general communication behavior in order to support realistic HPF and MPI programs.

Parashar et al. [24] construct task graphs for HPF programs compiled by the Syracuse Fortran 90D compiler, but they are limited to a very simple, loosely synchronous computational model that would not support many message-passing and HPF programs in practice. In addition, their interpretive framework for performance prediction uses functional interpretation for instantiating a dynamic task graph, which is similar to our approach for instantiating control-flow. Like the task graph model, their interpretation and performance estimation are significantly simplified (compared with ours) because of the loosely synchronous computational model. For example, they do not need to capture sophisticated communication patterns and computation partitionings, as we do using code generation from integer sets. Dikaiakos et al. [14] developed a tool called FAST that constructs task graphs from user-annotated parallel programs, performs advanced task scheduling and then uses abstract simulation of message passing to predict performance. Finally, many researchers (e.g., [16, 9, 22]) have developed symbolic compile-time techniques for estimating execution time, communication volume and other metrics. The communication and computation scaling functions available in our static task graph are very similar to the symbolic information used by these techniques, and could be directly extended to support their analytical models.

6 Conclusion and Future Plans

In this paper, we described a methodology for automating the process of synthesizing task graphs for parallel programs, using sophisticated parallelizing compiler techniques. The techniques in this paper can be used *without* user intervention to construct task graphs message-passing programs compiled from HPF source programs, and we believe they extend directly to existing message-passing (e.g., MPI) programs as well. Such techniques can make a large body of existing research based on task graphs and equivalent representations applicable for these widely used programming standards. Our immediate goals for the future are: (1) to demonstrate that the techniques described in this paper can be applied to message-passing programs (using MPI), by extracting the requisite computation partitioning and communication information; and (2) to couple the compiler-generated task graphs with the wide range of modeling approaches being used within the POEMS project, including analytical, simulation and hybrid models.

Acknowledgements: The authors would like to acknowledge the valuable input that several members of the POEMS project have provided to the development of the application representation. We are particularly grateful to the UCLA team and especially Ewa Deelman for her efforts in validating the simulations of the abstracted MPI codes generated by the compiler. This work was carried out while the authors were with the Computer Science Department at Rice University.

References

[1] V. Adve and J. Mellor-Crummey. Using Integer Sets for Data-Parallel Program Analysis and Optimization. In *Proc. of the SIGPLAN Conf. on Programming Lang. Design and Implementation*, Montreal, June 1998.

[2] V. Adve and R. Sakellariou. Application Representations for Multi-Paradigm Performance Modeling of Large-Scale Parallel Scientific Codes. *Int'l Journ. of High Performance Computing Applications*, 14(4), 2000.

[3] V. Adve and M. K. Vernon. A Deterministic Model for Parallel Program Performance Evaluation. Technical Report CS-TR98-333, Computer Science Dept., Rice University, December 1998. Also available at `http://www-sal.cs.uiuc.edu/~vadve/Papers/detmodel.ps.gz`.

[4] V. S. Adve, R. Bagrodia, J. C. Browne, E. Deelman, A. Dube, E. Houstis, J. R. Rice, R. Sakellariou, D. Sundaram-Stukel, P. J. Teller, and M. K. Vernon. POEMS: End-to-End Performance Design of Large Parallel Adaptive Computational Systems. *IEEE Trans. on Software Engineering*, 26(11), November 2000.

[5] V. S. Adve, R. Bagrodia, E. Deelman, T. Phan, and R. Sakellariou. Compiler- Supported Simulation of Highly Scalable Parallel Applications. In *Proceedings of Supercomputing'99*, Portland, Oregon, November 1999.

[6] V. Adve, G. Jin, J. Mellor-Crummey, and Q. Yi. High Performance Fortran Compilation Techniques for Parallelizing Scientific Codes. In *Proceedings of Supercomputing'98*, Orlando, FL, November 1998.

[7] S. Amarasinghe and M. Lam. Communication optimization and code generation for distributed memory machines. In *Proc. of the SIGPLAN '93 Conference on Programming Language Design and Implementation*, Albuquerque, NM, June 1993.

[8] C. Ancourt and F. Irigoin. Scanning polyhedra with do loops. In *Proceedings of the 3rd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, Williamsburg, VA, April 1991.

[9] B. Armstrong and R. Eigenmann. Performance Forecasting: Towards a Methodology for Characterizing Large Computational Applications. In *Proc. of the Int'l Conf. on Parallel Processing*, pages 518–525, Aug. 1998.

[10] R. Bagrodia, E. Deelman, S. Docy, and T. Phan. Performance prediction of large parallel applications using parallel simulation. In *Proc. 7th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, Atlanta, GA, 1999.

[11] P. Banerjee, J. Chandy, M. Gupta, E. Hodges, J. Holm, A. Lain, D. Palermo, S. Ramaswamy, and E. Su. The Paradigm compiler for distributed-memory multicomputers. *IEEE Computer*, 28(10):37–47, October 1995.

[12] M. Cosnard and M. Loi. Automatic Task Graph Generation Techniques. *Parallel Processing Letters*, 5(4):527–538, December 1995.

[13] E. Deelman, R. Bagrodia, R. Sakellariou, and V. Adve. Improving Lookahead in Parallel Discrete Event Simulation of Large Scale Applications using Compiler Analysis. In *Proceedings of the 15th Workshop on Parallel and Distributed Simulation (PADS)*, May 2001.

[14] M. Dikaiakos, A. Rogers, and K. Steiglitz. FAST: A Functional Algorithm Simulation Testbed. In *International Workshop on Modelling, Analysis and Simulation of Computer and Telecommunication Systems – Mascots '94*, 1994.

[15] D. L. Eager, J. Zahorjan, and E. D. Lazowska. Speedup versus Efficiency in Parallel Systems. *IEEE Trans. on Computers*, C-38(3):408–423, March 1989.

[16] T. Fahringer and H. Zima. A static parameter based performance prediction tool for parallel programs. In *Proceedings of the 1993 ACM International Conference on Supercomputing*, Tokyo, Japan, July 1993.

[17] P. Havlak. *Interprocedural Symbolic Analysis.* PhD thesis, Dept. of Computer Science, Rice University, May 1994. Also available as CRPC-TR94451 from the Center for Research on Parallel Computation and CS-TR94-228 from the Rice Department of Computer Science.

[18] S. Hiranandani, K. Kennedy, and C.-W. Tseng. Evaluation of compiler optimizations for Fortran D on MIMD distributed-memory machines. In *Proc. of the 1992 ACM International Conference on Supercomputing*, Washington, DC, July 1992.

[19] S. Horwitz, T. Reps, and D. Binkley. Interprocedural slicing using dependence graphs. *ACM Trans. on Programming Languages and Systems*, 12:26–60, 1990.

[20] A. Kapelnikov, R. R. Muntz, and M. D. Ercegovac. A Modeling Methodology for the Analysis of Concurrent Systems and Computations. *Journal of Parallel and Distributed Computing*, 6:568–597, 1989.

[21] M. S. Lam and M. Rinard. Coarse-Grain Parallel Programming in Jade. In *Proc. 3rd ACM SIGPLAN Symp. on Principles and Practices of Parallel Programming*, pages 94–105, Williamsburg, VA, April 1991.

[22] C. Mendes and D. Reed. Integrated Compilation and Scalability Analysis for Parallel Systems. In *Proc. of the Int'l Conference on Parallel Architectures and Compilation Techniques*, Paris, October 1998.

[23] P. Newton and J. C. Browne. The CODE 2.0 Graphical Parallel Programming Language. In *Proceedings of the 1992 ACM International Conference on Supercomputing*, Washington, DC, July 1992.

[24] M. Parashar, S. Hariri, T. Haupt, and G. Fox. Interpreting the Performance of HPF/Fortran 90D. In *Proceedings of Supercomputing '94*, Washington, D.C., November 1994.

[25] R. Sakellariou. *On the Quest for Perfect Load Balance in Loop-Based Parallel Computations.* PhD thesis, Dept. of Computer Science, Manchester University, Oct. 1996.

[26] A. Thomasian and P. F. Bay. Analytic Queueing Network Models for Parallel Processing of Task Systems. *IEEE Trans. on Computers*, C-35(12):1045–1054, December 1986.

[27] T. Yang and A. Gerasoulis. PYRROS: Static task scheduling and code generation for message passing multiprocessors. In *Proc. of the ACM Int'l Conference on Supercomputing*, Washington, DC, July 1992.