Optimal Use of Mixed Task and Data Parallelism for Pipelined Computations

Jaspal Subhlok
Department of Computer Science
University of Houston
Houston, TX 77098
jaspal@cs.uh.edu

Gary Vondran
Hewlett Packard Laboratories
Hewlett Packard Company
Cambridge, MA 02142
vondran@hpl.hp.com

Abstract

This paper addresses optimal mapping of parallel programs composed of a chain of data parallel tasks onto the processors of a parallel system. The input to the programs is a stream of data sets, each of which is processed in order by the chain of tasks. This computation structure, also referred to as a data parallel pipeline, is common in several application domains, including digital signal processing, image processing, and computer vision. The parameters of the performance for such stream processing are latency (the time to process an individual data set) and throughput (the aggregate rate at which data sets are processed). These two criteria are distinct since multiple data sets can be pipelined or processed in parallel. The central contribution of this research is a new algorithm to determine a processor mapping for a chain of tasks that optimizes latency in the presence of a throughput constraint. We also discuss how this algorithm can be applied to solve the converse problem of optimizing throughput with a latency constraint. The problem formulation uses a general and realistic model of inter-task communication, and addresses the entire problem of mapping, which includes clustering tasks into modules, assignment of processors to modules, and possible replication of modules. The main algorithms are based on dynamic programming and their execution time complexity is polynomial in the number of processors and tasks. The entire framework is implemented as an automatic mapping tool in the Fx parallelizing compiler for a dialect of High Performance Fortran.

Appears in the Journal of Parallel and Distributed Computing. A preliminary version of this work was presented at the Eighth Annual ACM Symposium on Parallel Algorithms and Architectures, under the title: "Optimal Latency–Throughput Tradeoffs for Data Parallel Pipelines".

1 Introduction

A fundamental problem in parallel computing is to find the most efficient mapping of a parallel program onto the processors of a parallel system. In simple data parallel computing, all available processors combine to execute every computation step in a program. Complex applications are often composed of multiple coarse grain tasks, where each task is data parallel and can execute on multiple processors, while different tasks may execute concurrently on different groups of processors. Most modern parallel machines support MIMD (Multiple Instruction Multiple Data) execution, and therefore combined data and task (or function) parallel computing. Compiler and runtime support for task and data parallel computing is an active area of research, and several solutions have been proposed [4, 5, 9, 10, 19, 21]. Recent research has also examined the benefits of mixed task and data parallel programming [3, 7, 18].

This paper specifically addresses the mapping of applications composed of a linear chain of data parallel tasks that act on a stream of input data sets. In this model, each task repeatedly receives input from its predecessor task, performs its computation, and sends the output to its successor task. Tasks are data parallel implying that each task can execute on one or more processors. The first task reads the external input and the last task generates the final output. Figure 1 shows a chain of three tasks.

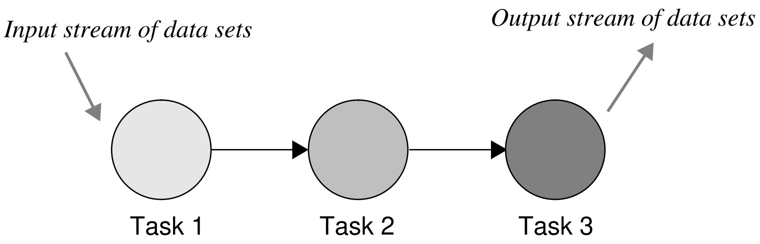

Figure 1: A three task pipeline

A chain of tasks is a relatively simple model of task and data parallel computing, yet it is fairly common in several application areas, including computer vision, image processing, and signal processing [8]. Such applications are typically driven by physical sensors like cameras and antennas. They process streams of relatively small data sets but often have strict throughput and latency requirements. An example application is multibaseline stereo [24], where the first task captures three (or more) images from the cameras, the second task computes a *difference image* for each of 16 disparity levels, the third task computes an *error image* for each difference image, and the final task performs a minimum reduction across error images and computes the final *depth image*.

An application composed of a chain of data parallel tasks can be mapped onto a parallel machine in a variety of ways, and this is illustrated with a three task chain in Figure 2. The figure shows the processor space with different patterns to denote the sets of processors employed to execute each of the three tasks. Figure 2(a) shows a pure data parallel mapping where all tasks execute on all processors. Each data set is processed by each of the three tasks in order. Therefore, at any given time, exactly one task is active and one data set is being processed. Figure 2(b) shows a task parallel mapping where a subset of the processors is dedicated to each task. A data set is processed by the three tasks in order, but task execution employs only a subset of the processors. Hence, up to three data sets may be processed simultaneously by separate tasks executing on disjoint processors. It may be possible to have replicated copies of the data parallel program to process multiple data sets in parallel on different sets of processors, as shown in Figure 2(c). Finally, a mix of task and data parallelism with replication is shown in Figure 2(d).

The goal of this research is to identify the best mapping for a chain of data parallel tasks. However, there are two distinct criterion for judging the quality of a mapping: *latency* and *throughput*. The latency is the time taken to process an individual data set, while the throughput is the aggregate rate at which the data sets are processed. Since multiple data sets may be pipelined or processed in parallel, latency and throughput are not directly related. If the objective is to minimize the latency, a simple data parallel mapping is often the optimal solution, since it makes all processors available to each task, and the tasks can execute at the fastest possible speed. For optimizing the throughput, however, it can be more beneficial to process multiple data sets simultaneously, in parallel or as a pipeline. This causes higher latency for individual data sets, but is likely to improve the throughput, since a single data set can be processed more efficiently by a smaller number of processors. The mapping decisions are particularly important for programs that do not scale well, often because of a relatively small data set size, and are therefore not able to use all processors

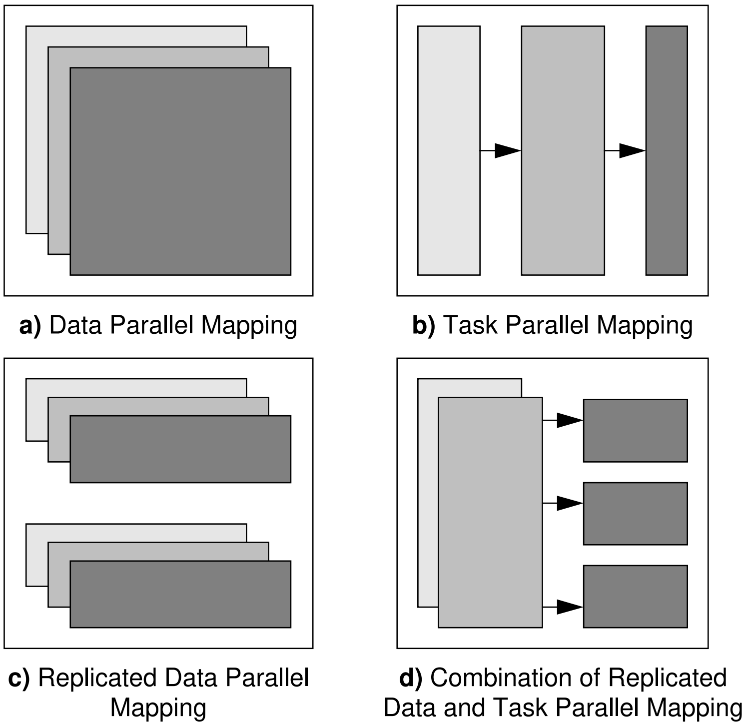

Figure 2: Combinations of data and task parallel mappings

efficiently with pure data parallelism.

This paper presents a solution to the following problem: "Given a chain of data parallel tasks and a minimum throughput constraint, what is the mapping that minimizes the latency ?". The converse problem of finding a mapping corresponding to maximum throughput with a given maximum latency constraint is also addressed.

Our basic approach is to execute the user program with a diverse set of mappings to automatically infer the computation and communication characteristics of the tasks in the program, and use these characteristics to derive an optimal mapping. In our implementation, a program has to be modified and recompiled to change the mapping. However, the methods developed in this research can be the basis of a runtime mapping tool that allows the best mapping to be derived dynamically during execution. This approach assumes that the execution behavior of a program is *not* strongly dependent on the properties of the input data sets. Also, we only address the mapping of the top level modules and not the partitioning of the computations inside these modules.

There is a large volume of literature on mapping and scheduling parallel programs, and some excellent references are [1, 16, 26]. Solutions to the problem of partitioning a chain of tasks among a set of processors have been presented in [2, 12, 13]. Our research extends this body of research to include data parallel tasks that can execute on a variable number of processors. Choudhary et. al. [6] address the problem of optimal processor assignment to data parallel tasks and this paper addresses the same fundamental problem and borrows the basic notation and solution techniques from their work. Our main contribution is that we employ a very general and realistic model for communication cost between parallel tasks, while [6] assumed that the communication cost can be folded into the computation cost. In our experience, a realistic model of communication is very important for a practical automatic mapping system. However, we develop solutions only for chains of tasks, while [6] addresses the more general series-parallel task structures. Ramaswamy et. al. [15] use convex programming based heuristics to address the problem of minimizing the completion time for one data set (latency) for tasks structured as a directed acyclic graph (DAG). An advantage of our approach over techniques like convex programming is that there are no constraints on the forms that execution and communication time functions can take. In particular, our execution time functions may be only available as a table or computed on demand. Finally, none of the earlier approaches address clustering of data parallel tasks into modules, or replication of tasks. Our earlier work [20] addressed clustering and replication but only for unconstrained throughput optimization, while this paper solves the general problem of constrained latency and throughput optimization.

The methods developed in this paper are implemented as an automatic mapping tool for the Fx compiler, and have been used to map several applications, including multibaseline stereo and narrowband tracking radar [8]. Fx is a compiler for a language similar to High Performance Fortran [11] with support for task parallelism [19, 21, 25]. A variation of the Fx task parallelism model is now an approved extension of HPF [11]. The targets for Fx are the Intel Paragon, the Intel iWarp, the IBM SP2, the Cray T3D, and networks of workstations running PVM.

This paper is organized as follows. Section 2 introduces the algorithmic problem to be solved. Section 3 presents an optimal dynamic programming solution to the general latency optimization problem, and Section 4 outlines how it can be used to solve the constrained throughput optimization problem. Section 5 discusses the estimation of execution behavior. Section 6 illustrates the importance of mapping algorithms with examples and discusses some practical issues in automatic mapping. Section 7 contains conclusions.

2 Problem statement

2.1 Execution model

A program in our model consists of a sequence of tasks $t_1, t_2, t_3, ...t_k$, where each task receives a data set from its predecessor, processes it, and sends the output to its successor; the first task reads external input data sets and the last task generates the final output data sets. Each task can execute on one or more processors. The execution time of each task is a function of the number of processors, and will be denoted by $f^{exec}(p)$. The execution time functions for the chain of tasks are represented as follows:

$$f_1^{exec}, f_2^{exec}, f_3^{exec}, ..., f_k^{exec}$$

A pair of adjacent tasks may be assigned to the same set of processors or to different sets of processors. If a pair of tasks is assigned to the same set of processors, the communication between them is a potential internal redistribution of data. This internal communication time is a function of the number of processors assigned to the tasks, and will be denoted by $f^{icom}(p)$. The internal communication time functions for the chain of tasks are represented as follows:

$$f_{1\to 2}^{icom}, f_{2\to 3}^{icom}, f_{3\to 4}^{icom}, ..., f_{(k-1)\to k}^{icom}$$

If a pair of tasks is assigned to different sets of processors, the communication between them involves moving data between different groups of processors. This external communication time is a function of the number of processors assigned to the sending task p_s and the receiving task p_r, and will be denoted by $f^{ecom}(p_s, p_r)$. The external communication time functions for the chain of tasks are represented as follows:

$$f_{1\to 2}^{ecom}, f_{2\to 3}^{ecom}, f_{3\to 4}^{ecom}, ..., f_{(k-1)\to k}^{ecom}$$

It follows from the above discussion that, in general, the communication time between two tasks is a function of the number of processors assigned to the sending task and the receiving task. This general communication time function will be denoted by $f^{com}(p_s, p_r)$. The general communication time functions for the chain of tasks are represented as follows:

$$f_{1\to 2}^{com}, f_{2\to 3}^{com}, f_{3\to 4}^{com}, ..., f_{(k-1)\to k}^{com}$$

The general communication time function essentially represents the applicable internal or external communication time function and can be described as follows:

$$\begin{aligned} f_{i->(i+1)}^{com}(p_i, p_{i+1}) &= f_{i->(i+1)}^{icom}(p_i) && \text{if tasks } t_i \text{ and } t_{i+1} \text{ are mapped on the same set of nodes} \\ &= f_{i->(i+1)}^{ecom}(p_i, p_{i+1}) && \text{if tasks } t_i \text{ and } t_{i+1} \text{ are mapped on different sets of nodes} \end{aligned}$$

where p_i denotes the number of processors assigned to the task t_i.

The response time of a task is the total time that a task spends on one data set that includes the time to receive an input data set, the time to process it, and the time to send the output data set. This response time is a function of the number of processors assigned to that task and the adjacent tasks. Based on the above discussion, the response time functions f_i for the chain of tasks is described as follows:

$$f_1, f_2, f_3, ..., f_k \quad \text{where} \quad f_i(p_{i-1}, p_i, p_{i+1}) = f_{(i-1)\to i}^{com}(p_{i-1}, p_i) + f_i^{exec}(p_i) + f_{i\to(i+1)}^{com}(p_i, p_{i+1})$$

The sender and the receiver tasks are involved in the entire duration of a communication step. The execution of one data set by a part of a chain of tasks is modeled in Figure 3 using $p_1, p_2, p_3, ...p_k$ to denote the number of processors assigned to the corresponding tasks in the chain $t_1, t_2, t_3, ...t_k$.

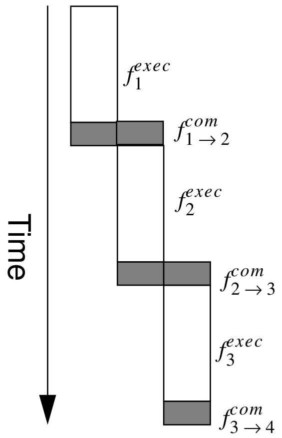

$$f_1(p_1, p_2) = f_1^{exec}(p_1) + f_{1\to 2}^{com}$$

$$f_2(p_1, p_2, p_3) = f_{1\to 2}^{com}(p_1, p_2) + f_2^{exec}(p_2) + f_{2\to 3}^{com}(p_2, p_3)$$

$$f_3(p_2, p_3, p_4) = f_{2\to 3}^{com}(p_2, p_3) + f_3^{exec}(p_3) + f_{3\to 4}^{com}(p_3, p_4)$$

Figure 3: Execution model of a chain of tasks. The total time spent on a data set by (say) task 2 includes the time to receive a data set from task 1, which is a function of p_1 and p_2 denoted by $f_{1\to 2}^{com}(p_1, p_2)$, the time to process the data set, which is a function of p_2 denoted by $f_1^{exec}(p_2)$, and the time to output the data set to task 3, which is a function of p_2 and p_3 denoted by $f_{2\to 3}^{com}(p_2, p_3)$. Thus the total processing time of task 2 is a function of p_1, p_2 and p_3, denoted by $f_2(p_1, p_2, p_3)$.

We point out some important aspects of this execution model and their implications. First, the communication cost of sending a data set between a pair of tasks executing on different sets of processors is assumed to be an arbitrary function of the number of processors assigned to the sender as well as the receiver. This coupling makes it harder to develop efficient processor assignment algorithms, but our experience suggests that it is important to model communication costs accurately. Second, the task chain acts on a long input stream of independent input data sets, and therefore, multiple data sets can be processed concurrently. Finally, the execution and communication times are static functions of the relevant numbers of processors. This is an important restriction on the applicability of this approach.

2.2 The mapping problem

We discuss the components of a mapping and show how a mapping is represented.

Processor assignment

If we assume that each task is assigned a disjoint subset of the processors, the mapping problem is that of finding an assignment A from tasks to processors, i.e. $A(i)$ equals the number of processors assigned to the task t_i. For a valid assignment of P processors to k tasks:

$$\sum_{i=1}^{k} A(i) \leq P$$

For the three task example shown in Figure 1, the problem is simply to divide the available processors among the three tasks according to the optimality criterion. Note that the computation is pipelined and the tasks are typically working on different data sets during execution.

Clustering into modules

A pair of adjacent tasks in a chain may execute on the same set of processors. We assume that any two tasks execute on identical or disjoint sets of processors. For the purpose of mapping, subchains of tasks are clustered into modules, and each module is assigned a disjoint subset of the available processors. A *clustering* of a chain of tasks is a list of modules $M_1, M_2, M_3, ...M_l$, where each module contains a subchain $t_i, t_{i+1}, t_{i+2}, ...t_j$ of the task chain and every task belongs to exactly one module. We implicitly assume that only adjacent tasks in a chain can be clustered together. Other clusterings (e.g. tasks 1 and 3 in one cluster and 2 in a different cluster) are rarely profitable and are not considered.

As an example, one way to map the three task pipeline of Figure 1 is to cluster tasks 1 and 2 in one module, task 3 in a different module, and divide the available processors among the two modules, as shown in Figure 4. Each data set will be processed by the first two tasks on the processors executing the first module, and then transferred to the processors executing the second module for processing by the third task.

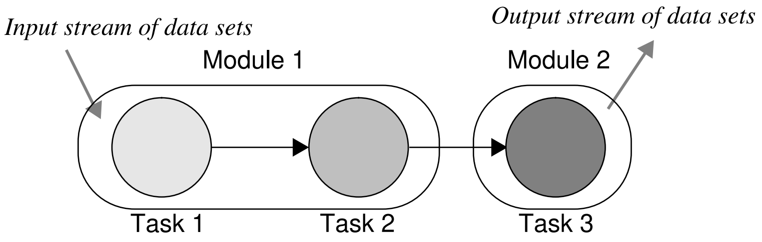

Figure 4: Clustering a three task pipeline into modules

Replication of modules

Suppose a set of p processors were available to a module. It may be profitable to divide the p processors in two or more disjoint groups and process alternate data sets on these disjoint groups of processors, as shown in Figure 5. While the response time for each data set, and therefore, the overall latency, will increase, the throughput may improve since multiple data sets are processed in parallel. This is called replication, and its legality depends on data dependences. For this paper, we assume that the tasks are known to be replicable or non-replicable. Replication is primarily used to improve throughput. The replication problem is to determine the modules that should be replicated, and the number of replicated instances for each such module.

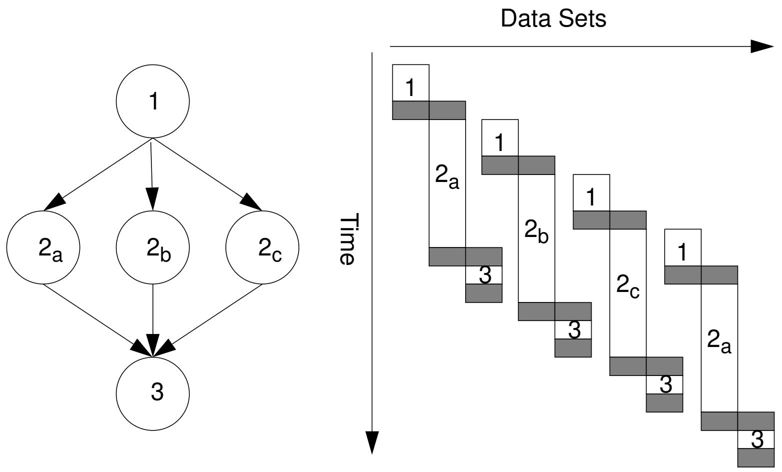

Figure 5: Replication. Module 2 is replicated to 3 instances. At any given time, up to 3 data sets can be processed by module 2, one each by the group of processors executing 2_a, 2_b and 2_c. Module 1 sends each output data set to one of the instances of module 2 in round robin sequence.

Representation of a mapping

A mapping is a result of clustering, replication and processor assignment decisions. The mapping of a chain of tasks $t_1, t_2, t_3, ..., t_k$ can be expressed as a list of modules $M_1, M_2, M_3, ..., M_l$, where M_i is a triplet $(\mathcal{T}, r, p)$, where $\mathcal{T}$ is the subchain of tasks clustered into that module, r is the number of times the module is replicated, and p is the number of processors assigned to each instance of the module.

The execution and communication time functions of a module can be composed from the corresponding functions of the tasks that constitute the module. For a mapping containing l modules, the latency , which is the time for processing an individual data set, can be expressed as:

$$\sum_{i=1}^{l}(f_i^{exec} + f_{i \to i+1}^{ecom})$$

and the throughput, which is the rate of processing data sets, as:

$$1/(\max_{i=1}^{l}(f_i/r_i))$$

where f , f^{exec}, and f^{ecom} correspond to the modules in the mapping and apply to the number of processors in each instance of a module. Also, since the last module does not have a successor, $f_{l \to l+1}^{ecom}$ is undefined and a value of 0 is used.

3 Latency optimization

In this section, we present an algorithm to find the mapping of a chain of tasks that minimizes the latency, while ensuring that a given minimum throughput constraint is satisfied. We first present an algorithm for unconstrained latency minimization without considering clustering or replication. This algorithm is similar to the latency minimization algorithm of Choudhary et.al. [6]. The difference is that we explicitly include communication time between tasks, which changes the details and complexity of the algorithm. We then extend this algorithm, first to include a throughput constraint, and subsequently to include clustering and replication. The final algorithm solves the general constrained latency optimization problem.

3.1 Unconstrained latency minimization

We begin by presenting a dynamic programming algorithm to find an optimal assignment of processors for a chain of tasks, without a throughput constraint. We assume that each task must execute on a disjoint set of processors, and that there is no replication. Thus, we are dividing the available processors among the given set of tasks such that the total execution time of the task chain (for one data set) is minimized. Every task corresponds to a unique module, so we will refer only to tasks and not to modules. For simplicity of presentation, we also assume that a task can use all processors available to it, and that the optimal assignment uses all available processors.

Approximate latency minimization can be done by adding processors one at a time to the task chain, by always selecting the task whose execution time decreases the most by adding a processor. Tasks that take longer to execute and tasks that scale well are favored, since their execution time is likely to decrease more with an additional processor. But this simple method does not always lead to an optimal mapping since it does not consider all possible assignments.

We first state our algorithm intuitively. The optimal assignment of p processors to a subchain of tasks, say from task t_1 to t_j, can be computed without consideration of the remaining parts of the chain, except that the number of processors assigned to the task t_{j+1} must be known, since the communication cost from t_j to t_{j+1} depends on it. This is illustrated in Figure 6. If we have a table of optimal assignments for this subchain from t_1 to t_j for all possible values of the number of processors assigned to this subchain, and all possible values of the number of processors assigned to the next task t_{j+1}, then we can use this information to build a corresponding table for the longer subchain from task t_1 to t_{j+1}. In this manner, we can build optimal assignments for incrementally longer subchains, and finally the optimal assignment for the full chain.

We now state the problem formally. Let T represent a chain of k tasks, labeled $t_1, t_2, ..., t_k$, and T_j represent the subchain $t_1, t_2, ..., t_j$ of the chain T. Let A_j represent an assignment of processors for a subchain of tasks T_j, i.e., $A_j(i)$ equals the number of processors assigned to the task t_i. An assignment A_j for a subchain of length j can be appended with a new value to obtain an assignment for a subchain of length $j+1$, i.e., $A_{j+1} = A_j \oplus \{(A(j+1) = p_{j+1})\}$. We will omit the subscript j of A_j for conciseness, whenever the length of the assignment is clear from the context.

Let L be a function that returns the latency of an assignment A for a task subchain T_j. The latency function L for T_j includes the computation time of the tasks in the subchain, the communication time between them, as well as the communication time between the last task t_j and the next task t_{j+1}, if it exists. Hence L is a function of the assignment A, and p_{next}, the number of processors assigned to the next task t_{j+1}. The latency function for an assignment to the subchain T_j can be evaluated as follows:

$$L(A, p_{next}) \quad = \quad \sum_{i=1}^{j} f_i^{exec} A(i) \quad + \quad \sum_{i=1}^{j-1} f_{i \to i+1}^{ecom}(A(i), A(i+1)) \quad + \quad f_{j \to j+1}^{ecom}(A(j), p_{next})$$

We show how an optimal assignment for a subchain of tasks is related to the optimal assignment for the full chain. Let A_j be a unique optimal assignment of p processors to a task subchain T_j, assuming that p_{j+1} processors are assigned to the next task t_{j+1}. We argue that A_j is the assignment for T_j in the final optimal assignment for the full chain T, if p and p_{j+1} are the numbers of processors assigned to T_j and t_{j+1}, respectively, in the optimal assignment for the full chain. This is the case since no other factors influence the final assignment for the subchain T_j. Note that p_{j+1}, the number of processors assigned to the task following the subchain T_j, is needed to compute the optimal assignment for T_j, since the communication time between the last task in T_j and the next task t_{j+1} is part of the latency of the subchain T_j. We have the following result:

Lemma 1 Suppose the function $\mathcal{A}_j(p_{total}, p_{next})$ returns the optimal assignment of p_{total} processors to a subchain T_j of a task chain T, given that the task immediately following the subchain is assigned p_{next} processors. (If no *next* task exists, p_{next} is redundant and assigned ϕ.) If an optimal assignment for the full chain T has p and p_{j+1} processors assigned to the subchain T_j and the task t_{j+1}, respectively, then the assignment of processors to T_j in this optimal mapping is that returned by $\mathcal{A}_j(p, p_{j+1})$.

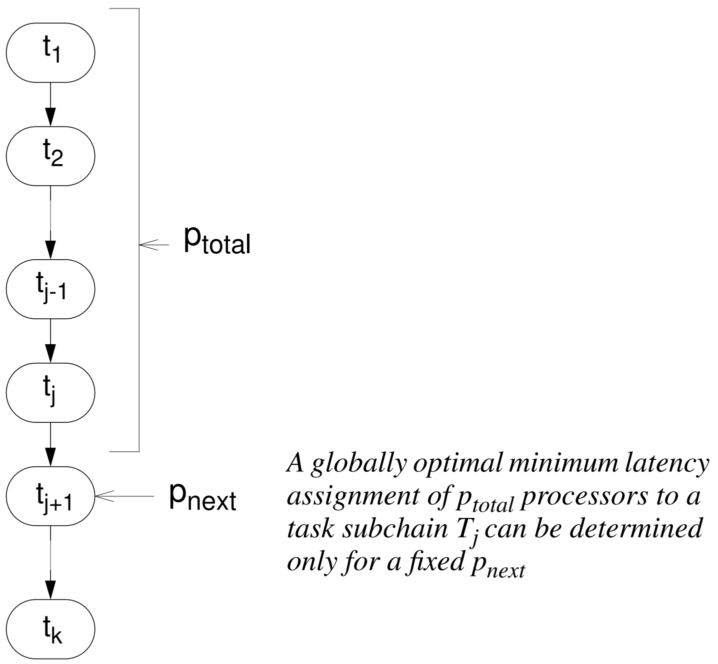

Figure 6: Processor assignment with dynamic programming

The lemma is illustrated in Figure 6. The optimal assignment of P processors to a chain of k tasks, and the corresponding minimum latency, is obtained directly from the function $\mathcal{A}_k$.

$$\begin{aligned} A^{opt} &= \mathcal{A}_k(P, \phi) \\ L^{opt} &= L(A^{opt}, \phi) \end{aligned}$$

We now present a constructive recursive definition of $\mathcal{A}$.

$$
\begin{aligned}
\mathcal{A}_j(p_{total}, p_{next}) &= A_j \equiv \min_{q=1}^{p_{total}} L\,[\,(\mathcal{A}_{j-1}(p_{total}-q, q) \oplus \{(A(j)=q)\}), p_{next}\,] && \text{if } j>1 \\
&= \{(A(1)=p_{total})\} && \text{if } j=1
\end{aligned}
$$

We explain the meaning of the above equation and describe how it is used to compute the processor assignment for minimum latency. Suppose we have a table defining the optimal assignment $\mathcal{A}_{j-1}$, and the corresponding minimum latency, for every possible number of processors assigned to the task subchain T_{j-1} and the next task t_j. To determine an entry in the table for $\mathcal{A}_j$, say $\mathcal{A}_j(p_{total}, p_{next})$, we need to find the optimal latency assignment for the subchain T_j, given that the number of processors assigned to T_j, and the next task t_{j+1}, are p_{total} and p_{next}, respectively. If we assume that q processors are assigned to t_j and the remaining $p_{total}-q$ processors to subchain T_{j-1}, the corresponding optimal assignment and latency for T_{j-1} can be looked up from the table for $\mathcal{A}_{j-1}$. The corresponding latency for the subchain T_j can then be directly computed, given that p_{next} processors are assigned to t_{j+1}. By comparing the latency of the assignments for all possible values of q, we obtain the optimal assignment $\mathcal{A}_j(p_{total}, p_{next})$, and the corresponding minimum latency. The optimal assignment and the minimum latency for the complete chain of k tasks is obtained by repeating the procedure for incrementally longer subchains.

Complexity: The size of a definition table for $\mathcal{A}$ is at most $O(P^2)$. Each entry of a new table $\mathcal{A}_j$ is determined by comparing at most P values, each of which is calculated from $\mathcal{A}_{j-1}$ and the response time functions in $O(1)$ operations. Hence $\mathcal{A}_j$ is obtained from $\mathcal{A}_{j-1}$ in $O(P^3)$ operations. Since there are k such steps, the total complexity of the algorithm is $O(P^3k)$ operations.

3.2 Latency minimization with a throughput constraint

The throughput of a chain of tasks is the rate at which data sets are processed. Intuitively, to increase the throughput of a chain of tasks, processors are added to the slowest or the *bottleneck* task, even though adding them to another task may increase the response time of that task, as well as the total latency, by a larger amount. This is the case since the throughput is determined by the processing rate of the slowest task in a pipeline, and therefore, to improve the throughput, the slowest task has to be speeded up. When a throughput constraint is given, we have to ensure a certain minimum processing speed for every task, and then minimize the total latency of the chain.

The throughput of an assignment A to a subchain T_j, given that the next task t_{j+1} is assigned p_{next} processors, can be computed by the following function:

$$
H(A, p_{next}) = min[\{\min_{i=1}^{j-1} 1/f_i(A(i-1), A(i), A(i+1))\}, \{1/f_j(A(j-1), A(j), p_{next})\}]
$$

We restate that $A(i)$ represents the number of processors assigned to task t_i in assignment A and the response time function f_i represents the sum of the computation and the communication times of the task t_i, and is therefore a function of the number of processors assigned to the task t_i and to the tasks before and after it in the task chain.

Note that the number of processors assigned to the task immediately following a subchain T_j is also needed to compute the throughput of a subchain. The algorithm for unconstrained latency minimization can be modified in a simple way to compute the minimum latency with a minimum throughput constraint. During the building of the table of optimal assignments, only assignments that satisfy the throughput constraint are considered.

We introduce the constrained latency function $L^{\mathbf{h}}$, which returns the latency of an assignment if the corresponding throughput is greater than $\mathbf{h}$, and *infinity* otherwise.

$$
\begin{aligned}
L^{\mathbf{h}}(A, p_{next}) &= L(A, p_{next}) && \text{if } H(A, p_{next}) \geq \mathbf{h} \\
&= \infty && \text{if } H(A, p_{next}) < \mathbf{h}
\end{aligned}
$$

Using the constrained latency function in place of the unconstrained latency function, the equation system from the previous section can be used for latency minimization with a throughput constraint. We have:

$$
\begin{aligned}
\mathcal{A}_j(p_{total}, p_{next}) &= A_j \equiv \min_{q=1}^{p_{total}} L^{\mathbf{h}}\,[\,(\mathcal{A}_{j-1}(p_{total}-q, q) \oplus \{(A(j)=q)\}), p_{next}\,] && \text{if } j>1 \\
&= A_1 \equiv L^{\mathbf{h}}\{(A(1)=p_{total})\} && \text{if } j=1
\end{aligned}
$$

When the constrained latency function $L^{\mathbf{h}}$ returns infinity, the corresponding assignment is undefined. If all entries for the optimal mapping of a subchain are undefined, then no assignment satisfies the throughput constraint, and therefore, no solution exists.

The algorithm for finding the constrained optimal latency is essentially the same as that for unconstrained latency optimization presented earlier, with L replaced by L^h. The additional check for throughput requires only constant additional operations at every step. Therefore the computation complexity is also the same, that is, $O(P^3k)$ operations.

3.3 Latency minimization with clustering

In the previous subsections, we have assumed that every task is assigned a disjoint set of processors. However, it is often more beneficial to group a set of tasks together in a module and assign a set of processors to be shared by the tasks in the module. The tasks inside a module are sequentialized, that is, all processors are used by one task until it finishes processing a data set, and then by the other tasks in order. Coalescing a pair of tasks in a single module can potentially reduce the communication time between the tasks and improve the latency, but it can also reduce the throughput if the tasks do not scale well, since they may not execute as efficiently on a larger number of processors. We outline the changes to the dynamic programming algorithm to take clustering into account. For brevity, we will not present the modified equation system, and instead refer to [22].

Given a chain of tasks T, the mapping now involves partitioning the tasks into a chain of modules $M_1, M_2, ...M_l$, and allocating available processors to the modules. For a dynamic programming solution, the central subproblem, as before, is to find a mapping of a subchain T_j using p processors which is guaranteed to be the mapping for T_j in a global optimal mapping of T, if T_j is indeed allocated p processors in such a global optimal mapping. We label M_{last} as the last module in a mapping of the subchain T_j. Hence M_{last} contains the task t_j, a zero or more length sequence of tasks preceding t_j, and a zero or more length sequence of tasks following t_j. We label M_{next} as the *next* module, i.e., the one after M_{last}, in the final mapping. The key question is what quantities must be known to decide the optimal mapping of subchain T_j using p processors. The answer is as follows:

1. Number of tasks following T_j that will be assigned to M_{last}. This information is necessary to compute the execution time of M_{last} which is needed to verify if it satisfies the throughput constraint.

2. Number of processors assigned to the following module M_{next}. This information is necessary to compute the communication time of M_{last} which is needed to verify if M_{last} satisfies the throughput constraint. This information is also needed for computing the latency of T_j when the next task t_{j+1} belongs to M_{next} and not to M_{last}.

3. Number of processors assigned to the module M_{last}, i.e., we must also fix the number of processors to be assigned to the last module. The rather non-intuitive reason is as follows. The total processing time of M_{next} depends on the communication time from M_{last} to M_{next}, which in turn, depends on the number of processors assigned to M_{last}. Hence the fact whether M_{next} satisfies a throughput constraint depends on the number of processors assigned to M_{last}. Therefore selecting M_{last} based on optimization for a subchain may not lead to a global optimal as the choice impacts whether later modules will satisfy the throughput constraint. Hence M_{last} must be fixed, implying that the optimization must be done separately for every legal value of M_{last}.

The implication of the above discussion is as follows. The dynamic optimization table must be indexed by the number of processors assigned to a subchain, the number of processors assigned to the last module M_{last} in the subchain, the number of processors assigned to the following module M_{next}, and the number of following tasks clustered with the module M_{last}. As the number of processors is bounded by the total number of processors P, and the number of tasks by the total number of tasks k, the size of the table is $O(P^3k)$. A table entry for subchain T_{j+1} can be computed from the table for T_j with at most $O(P)$ comparisons as before. An important difference due to clustering is that cases where the new task t_{j+1} is clustered with and without the previous task t_j have to be considered. In terms of dynamic programming terminology used in this paper, we have to work with the mapping function $\mathcal{M}_j(p_{total}, p_{last}, p_{next}, modlength)$ instead of $\mathcal{A}_j(p_{total}, p_{next})$, where $modlength$ is the length of the subchain after task t_j that is clustered in the module M_{last}, and p_{last} and p_{next} are numbers of processors assigned to M_{last} and M_{next}, respectively. We state again that the table size is $O(P^3k)$, a new table entry is computed in $O(P)$ steps, and a sequence of k tables have to be computed for a full solution. Hence, the computation complexity of latency optimization with a throughput constraint in the presence of clustering is $O(P^4k^2)$.

3.4 Latency minimization with replication

With replication, processors available to a module are divided into multiple groups that process input data sets in a round robin sequence, as shown in Figure 5 and discussed in Section 2.2. Replication implies that each data set uses fewer processors, and therefore, it generally increases the latency, but it can improve the throughput. Hence, the purpose of replication is to help satisfy a throughput constraint.

We briefly discuss how the mapping algorithm is modified to incorporate replication. We use the algorithm in the previous subsection, with the exception that if a particular assignment of tasks and processors to M_{next} does not satisfy the throughput constraint, we determine the number of copies of M_{next} required to satisfy the throughput constraint, and assume their existence if that does not exceed the available number of processors. The complexity of the mapping algorithm is unchanged. Hence, for the general problem of minimizing latency with a throughput constraint, we have the following:

Lemma 2 *Given a chain of k tasks, P processors, and a minimum throughput constraint, an optimal latency mapping for the task chain can be be determined in $O(P^4k^2)$ time.*

3.5 Execution efficiency of latency optimization

The execution time complexity of the latency minimization algorithm is $O(P^4k^2)$. The practical cost of executing this algorithm is an issue for large P, particularly when the algorithm is used at runtime to map and remap programs.

The reason for the quadratic complexity in P (rather than square or cubic) is that we model communication between two tasks as an arbitrary function of the number of processors assigned to the sender and the receiver tasks. If communication cost is modeled as two independent sending and receiving costs, the complexity of the algorithm is lower, but our experience suggests that this is not a realistic model of communication. We also make no assumptions about the behavior of the execution and communication time functions.

Several optimizations based on the properties of the execution time functions can be used to improve the performance, but a full discussion is beyond the scope of this paper. However, we will briefly discuss the use of small processor blocks, instead of individual processors, as the assignment units for the mapping algorithms. This leads to a simple tradeoff between precision and complexity. Using a block of b processors instead of one would reduce the computation complexity of the general algorithm by a factor of b^4, while ensuring that the number of processors assigned to a task is at most b processors off from the corresponding assignment in an optimal mapping. In general, if we are willing to be off by one processor for every P_0 processors in the system, i.e. by a fixed fraction of the total number of processors, the total complexity is $O(P_0^4k^2)$. This is independent of the actual number of processors, and is a function of the desired precision. We believe that this is a practically useful way to use these algorithms for large numbers of processors.

4 Throughput optimization

The general throughput optimization problem is to find a mapping that corresponds to the maximum throughput with a maximum latency constraint. In previous work [20], we developed an algorithm for optimizing the throughput, but without a latency constraint. It appears that there is no polynomial solution to the problem of optimizing throughput with a latency constraint.

However, repeated application of the algorithm for latency optimization can be used to find a near optimal solution to the throughput problem. We first present the following result, which essentially states that increasing the minimum throughput constraint when solving the latency optimization problem cannot decrease the optimal latency.

Lemma 3 *Given a task chain and a fixed number of processors, let H be the maximum possible throughput for a maximum allowed latency* $\mathbf{l}$. *Let L_1 be the minimum latency for the chain with a minimum throughput constraint* $\mathbf{h_1}$, *and L_2 the minimum latency with a minimum throughput constraint* $\mathbf{h_2}$. *Then:*

$$\mathbf{h_2} \geq H \geq \mathbf{h_1} \Leftrightarrow L_2 \geq \mathbf{l} \geq L_1$$

This result is fairly obvious and has the following implication. If we find the optimal latency for two different throughput constraints, and the given latency constraint is between the two optimal latency values, then the optimal throughput is between the two throughput constraints. Using this result, we can obtain successively improved approximations to the optimal throughput for a given latency constraint by using binary search and repeatedly solving the optimal latency problem. An initial range for the throughput can be determined by solving the unconstrained throughput problem, which has the same complexity as the general latency problem [20].

Therefore, the general throughput problem can be solved in $O(log(H_0/\delta)P^4k^2)$ time by this procedure, where H_0 is the range of possible throughput values, and δ is the maximum allowed difference between the approximate and the exact optimal throughput.

5 Estimation of execution behavior

The mapping algorithms presented in this paper use the execution time and communication time functions for the tasks. One advantage our algorithms have over those based on mathematical programming (e.g. linear or convex programming) is that they are not tied to any particular execution model. The execution time and communication time functions may be mathematical functions computed at compile time or runtime, or they may be defined pointwise, possibly using interpolation.

In our mapping tool, we model computation and communication times as polynomial functions. The parameters of the model are derived automatically by analyzing the profile information from a set of executions.

The execution time of a task-subroutine on a set of p processors is modeled as:

$$f^{exec}(p) \quad = \quad C_1^{exec} + C_2^{exec}/p + C_3^{exec} * p$$

The constant term C_1^{exec} reflects fixed cost sequential and replicated computation, the second term represents parallel computation whose cost decreases linearly with the number of processors, and the last term represents overheads that increase with the number of processors.

The external communication function denotes the time for transferring data between a pair of tasks mapped on different sets of processors, and is modeled as a function of the number of processors assigned to the sending and the receiving tasks (p_s and p_r) as follows:

$$f^{ecom}(p_s, p_r) \quad = \quad C_1^{ecom} + C_2^{ecom}/p_s + C_3^{ecom}/p_r + C_4^{ecom} * p_s + C_5^{ecom} * p_r$$

The internal communication function denotes the time for a potential redistribution of data when a pair of communicating tasks are mapped to the same set of processors, and is modeled in the same manner. In this case the sending processors are also the receiving processors, and hence there are only 3 terms.

$$f^{icom}(p) \quad = \quad C_1^{icom} + C_2^{icom}/p + C_3^{icom} * p$$

6 Results

The execution model and the mapping algorithms presented in this paper are implemented as an automatic mapping tool in the Fx compiler. The programming model used for integrated task parallelism is discussed in [21], and a variation of this model is also an approved extension of High Performance Fortran [11]. The programming model is not the topic of this paper, and we will simply state that it allows the expression of task parallelism including clustering and replication, and hence is sufficient for programming all mappings discussed in this paper.

We present a set of results using a 64 processor Intel iWarp, and 64 processors of an Intel Paragon. The purpose is to illustrate the mapping procedure, demonstrate that the mapping tool automatically finds optimal or nearly optimal mappings in practice, and to examine the practical reasons for inaccuracies in finding optimal mappings.

6.1 Example Programs

FFT-Hist

The FFT-Hist program takes a sequence of m $n \times n$ complex arrays from a sensor (e.g., a camera). For each of the m input arrays, a two dimensional fast Fourier transform (FFT) is performed, followed by statistical analysis. The main loop of the program and the task graph are shown in Figure 7. For each iteration of the loop, the `colffts` function inputs the array `A` and performs 1D FFTs on the columns. The `rowffts` function performs 1D FFTs on the rows, and the `hist` function analyzes and outputs the results. The `rowffts` and `colffts` functions are parallel with no communication, while `hist` has a significant amount of internal communication. This is an interesting program because it represents the structure of many applications in image and signal processing.

```
do i = 1,m
   call colffts(A)
   call rowffts(A)
   call hist(A)
enddo
```

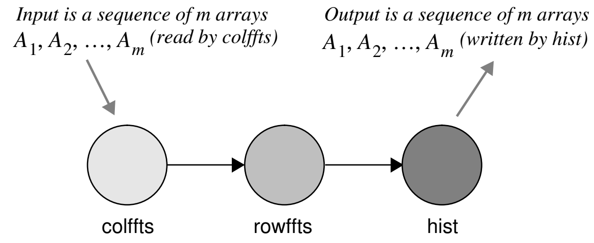

Figure 7: FFT-Hist example program and task graph

Radar

The narrowband tracking radar benchmark was developed by researchers at MIT Lincoln Labs to measure the effectiveness of various multicomputers for their radar applications [17]. The radar program inputs data from a single sensor along $c = 4$ independent *channels*. Every 5 milliseconds, for each channel, the program receives $d = 512$ complex vectors of length $r = 10$, one after the other in the form of an $r \times d$ complex matrix A (assuming the column major ordering of Fortran). The computation phase of the algorithm is composed of column FFTs, scaling and magnitude computation, and thresholding.

The high level Fx version of the radar program operating on a stream of m input data sets has the following form:

```
do i = 1,m
   call dgen(A)
   call compute(A)
enddo
```

Figure 8: Radar example program

Stereo

The multibaseline stereo uses an algorithm developed at Carnegie Mellon that gives greater accuracy in depth through the use of three or more cameras [14, 23].

Input consists of three $m \times n$ images acquired from three horizontally aligned, equally spaced cameras. One image is the *reference image*, the other two are *match images*. For each of 16 disparities, $d = 0, \ldots, 15$, the first match image is shifted by d pixels, and the second image is shifted by $2d$ pixels. A *difference image* is formed by computing the sum of squared differences between the corresponding pixels of the reference image and the shifted match images. Next, an *error image* is formed by replacing each pixel in the difference image with the sum of the pixels in a surrounding 13×13 window. A *disparity image* is then formed by finding, for each pixel, the disparity that minimizes error. Finally, the depth of each pixel is displayed as a simple function of its disparity.

The high level Fx version of the stereo program operating on a stream of m image sets has the following form:

6.2 Modeling and mapping

We have used the mapping tool and the Fx compiler to generate mappings for several task and data parallel programs. In Table 1 we present results from FFT-Hist, narrowband tracking radar and multibaseline stereo applications. The results show the optimal latency obtained with and without a throughput constraint. We observe that adding a throughput constraint generally leads to a mapping with a significantly higher latency. The table also shows the number of processors actually used in the various optimal mappings. We will discuss the mapping of the FFT_Hist program and briefly outline the mapping of the stereo and radar programs.

```
do i = 1,m
   call dgen(INPUT)
   call diff(DIFF,INPUT)
   call error(ERR,DIFF)
   call depth(DEPTH,ERR)
enddo
```

Figure 9: Stereo example program

FFT-Hist

We used our mapping tool on the FFT-Hist program for two data set sizes (256×256 and 512×512) and two parallel machines, the Intel iWarp and the Intel Paragon. In each case, the program was run through a training set of sample mappings to build a computation and communication model for the tasks, as discussed in section 5.

The mapping tool used the algorithms presented in Section 3 to determine the mapping corresponding to the optimal latency, with and without a minimum throughput constraint. The optimal latency mappings obtained for the 512×512 FFT-Hist on an Intel Paragon are shown in Figure 10, and the corresponding latency and throughput are shown in Table 1. We briefly discuss the intuitive reasoning that leads the tool to infer these optimal mappings.

The mapping obtained for minimizing latency with no throughput constraint is the simple data parallel mapping. This is expected since clustering into multiple modules and replication typically improve the throughput, but also increase the latency, since the number of processors that execute a task is reduced.

The second mapping corresponds to a minimum throughput constraint of 2 data sets per second. In this case, all tasks are again clustered together, but the module is replicated. Replication improves the throughput since the module instances can use 32 processors more efficiently than 64. The optimal latency is significantly higher than the optimal latency without a throughput constraint.

The third mapping was obtained for a more restrictive throughput constraint of 4 data sets per second. This mapping has two modules, *module 1* contains `colffts`, and *module 2* contains `rowffts` and `hist`. *Module 1* is replicated to 2 instances and each instance is assigned 8 processors, while *Module 2* has 3 instances and 14 processors per instance. The primary reason for having two modules is that the communication step between *colffts* and *rowffts* is somewhat more efficient if they are on different groups of processors. The modules are replicated to satisfy the throughput constraint. We were somewhat surprised that a few processors were not used, but experimentation with nearby mappings that used all processors confirmed that the selected mapping was indeed the most efficient.

We have used this example to illustrate some of the factors that influence the mapping, but the optimal mapping is really determined by a quantitative comparison of the impact of several factors and can be hard to predict. For the FFT-Hist example, the nature of the mappings obtained for different throughput constraints was very different for the 256×256 and 512×512 problems, and was also different for the same problem executing on the iWarp and the Paragon.

Radar

The radar program contains a chain of two tasks as shown in Figure 8. The data parallel implementation of the radar program shows linear speedup but scales only to 10 processors. The reason for this is that the program is parallelized along the dimension with 10 elements. The task and data parallel version that optimizes latency is simply the pure data parallel version. Throughput can be improved by executing different data sets on different groups of processors in a round robin fashion. Since the data parallel version uses only 10 processors, additional throughput is obtained by using more processors and hence the latency is not reduced. The mapping decisions are reflected in the performance results shown in Table 1.

Stereo

The stereo program contains a chain of tasks as shown in Figure 9. The data parallel implementation of stereo uses all 64 processors, although scalability is limited due to increase in the communication overhead for larger numbers of

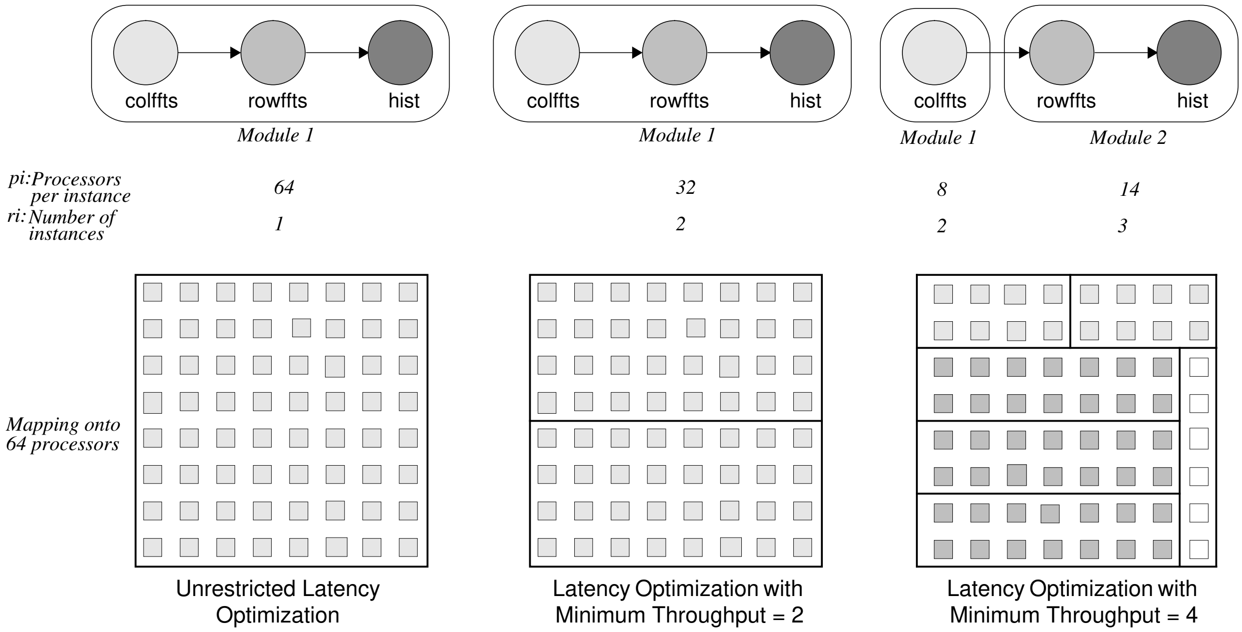

Figure 10: Mappings of a 512x512 FFT-Hist program on an Intel Paragon

processors. The task and data parallel mapping that optimizes latency is again the pure data parallel mapping. The mapping that minimizes latency with a throughput constraint is obtained by clustering tasks in a single module and replicating the module to satisfy the throughput constraints. Since each data set uses fewer processors in this scenario, the latency is adversely affected. The mapping decisions are reflected in the performance results shown in Table 1.

Problem Description		Name	Throughput	Optimal Latency		Percent	Measured	Processors
Name of Program	Array Size	of Machine	Constraint datasets/sec	Predicted sec	Measured sec	Difference	Throughput datasets/sec	Used
FFT-Hist	256x256	Paragon	None	.232	.256	10.3	3.90	58
			8	.357	.293	-17.9	13.3	64
FFT-Hist	256x256	iWarp	None	.463	.444	-4.1	2.25	32
			8	.484	.491	1.5	8.14	64
FFT-Hist	512x512	Paragon	None	.578	.502	-13.1	1.99	64
			2	.853	.807	-5.4	2.48	64
FFT-Hist	512x512	iWarp	None	.739	.739	0.0	1.35	64
			2	.845	.845	0.0	2.37	64
Radar	512x10x4	Paragon	None	.041	.043	4.9	23.4	10
			50	.041	.043	4.9	70.2	30
Radar	512x10x4	iWarp	None	.049	.049	0.0	20.3	10
			50	.049	.049	0.0	60.9	30
Stereo	256x240	Paragon	None	.275	.275	0.0	3.64	64
			12	.486	.514	5.7	11.67	64
Stereo	256x240	iWarp	None	.115	.115	0.0	8.7	64
			12	.152	.149	-2.3	13.4	64

Table 1: Performance results

6.3 Accuracy of performance estimates

Table 1 also shows that the predicted and measured latency generally differ. The maximum estimation error for these examples was 17.9 percent. The mapping algorithms used are optimal, hence the inaccuracy in prediction is entirely because of the inaccuracy in the modeling of the execution behavior. Modeling of execution is not a contribution of this paper, and we have selected a relatively simple model. A small to moderate difference between predicted and measured values is expected, and in fact the reason that some cases have no error is that the optimal mapping happened to coincide with one of the data points used for execution modeling. Some of the factors that cause inaccuracy in performance prediction are as follows:

1. The execution performance sometimes changes abruptly as the number of processors is changed. For example, some tasks execute significantly more efficiently with power of 2 number of processors.

2. A small change in the number of processors can result in a dramatic change in the cache behavior. This is one of the reasons why the predictions on the iWarp, which has a flat memory, were generally more accurate than the predictions on the Paragon.

3. The communication cost between two modules is often significantly lower when the number of processors in the sender and the receiver are the same, or one is a multiple of the other.

More sophisticated modeling of execution can certainly improve the accuracy of the predictions, but further discussion is beyond the scope of this paper. The algorithms presented in this paper are independent of how execution behavior is modeled, and therefore can be used with other modeling approaches.

7 Conclusions

This paper offers a solution to the problem of finding processor mappings to satisfy the latency and throughput requirements of data parallel pipelines. We have presented an algorithm for minimizing the latency of a chain of data parallel tasks in the presence of a minimum throughput constraint, and discussed how it can be used to solve the converse problem of maximizing the throughput with a maximum latency constraint. The algorithm addresses the general mapping problem, which includes processor assignment, clustering, and replication, and is independent of the details of the computation and communication modeling.

This research is directed towards applications that process a stream of data sets from a physical sensor like a camera, where each data set is relatively small but parallel computing is needed to achieve the desired latency and throughput. Data parallel computing is generally not an adequate solution since small data sets may not execute efficiently on a large number of processors. We have implemented the ideas in this paper as a mapping tool that finds the optimal task and data parallel mapping for such applications. The mapping tool is integrated with the Fx parallelizing compiler for a subset of High Performance Fortran. We have demonstrated with examples that this tool chain can be used effectively to parallelize these applications and to tune them to their performance goals.

8 Acknowledgements

We would like to thank the Fx compiler team at Carnegie Mellon, in particular Peter Dinda, James Stichnoth, Tom Stricker and Bwolen Yang, for building and supporting the compiler infrastructure used in our implementation. Thomas Gross and David O'Hallaron participated in several fruitful discussions about this research.

References

[1] BOKHARI, S. *Assignment Problems in Parallel and Distributed Computing*. Kluwer Academic Publishers, 1987.

[2] BOKHARI, S. Partitioning problems in parallel, pipelined, and distributed computing. *IEEE Transactions on Computers 37* (Jan 1988), 48–57.

[3] CHAKRABARTI, S., DEMMEL, J., AND YELICK, K. Modeling the benefits of mixed data and task parallelism. In *Seventh Annual ACM Symposium on Parallel Algorithms and Architectures* (Santa Barbara, CA, July 1995).

[4] CHANDY, M., FOSTER, I., KENNEDY, K., KOELBEL, C., AND TSENG, C. Integrated support for task and data parallelism. *International Journal of Supercomputer Applications 8*, 2 (1994), 80–98.

[5] CHAPMAN, B., MEHROTRA, P., VAN ROSENDALE, J., AND ZIMA, H. A software architecture for multidisciplinary applications: Integrating task and data parallelism. Tech. Rep. 94-18, ICASE, NASA Langley Research Center, Hampton, VA, Mar. 1994.

[6] CHOUDHARY, A., NARAHARI, B., NICOL, D., AND SIMHA, R. Optimal processor assignment for a class of pipelined computations. *IEEE Transactions on Parallel and Distributed Systems 5*, 4 (April 94), 439–445.

[7] CROWL, L., CROVELLA, M., LEBLANC, T., AND SCOTT, M. The advantages of multiple parallelizations in combinatorial search. *Journal of Parallel and Distributed Computing 21* (1994), 110–123.

[8] DINDA, P., GROSS, T., O'HALLARON, D., SEGALL, E., STICHNOTH, J., SUBHLOK, J., WEBB, J., AND YANG, B. The CMU task parallel program suite. Tech. Rep. CMU-CS-94-131, School of Computer Science, Carnegie Mellon University, Mar. 1994.

[9] FOSTER, I., AVALANI, B., CHOUDHARY, A., AND XU, M. A compilation system that integrates High Performance Fortran and Fortran M. In *Proceeding of 1994 Scalable High Performance Computing Conference* (Knoxville, TN, October 1994), pp. 293–300.

[10] GROSS, T., O'HALLARON, D., AND SUBHLOK, J. Task parallelism in a High Performance Fortran framework. *IEEE Parallel & Distributed Technology 2*, 3 (Fall 1994), 16–26.

[11] HIGH PERFORMANCE FORTRAN FORUM. *High Performance Fortran Language Specification, Version 2.0*, Dec. 1996.

[12] IQBAL, M., AND BOKHARI, S. Efficient algorithms for a class of partitioning problems. *IEEE Transactions on Parallel and Distributed Systems 6*, 2 (February 1995), 170–175.

[13] NICOL, D., AND O'HALLARON, D. Improved algorithms for mapping parallel and pipelined computations. *IEEE Transactions on Computers 40* (Mar 1991), 295–306.

[14] OKUTOMI, M., AND KANADE, T. A multiple-baseline stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence 15*, 4 (1993), 353–363.

[15] RAMASWAMY, S., SAPATNEKAR, S., AND BANERJEE, P. A convex programming approach for exploiting data and functional parallelism. In *Proceedings of the 1994 International Conference on Parallel Processing* (St Charles, IL, August 1994), vol. 2, pp. 116–125.

[16] SARKAR, V. *Partitioning and Scheduling Parallel Programs for Multiprocessors*. The MIT Press, Cambridge, MA, 1989.

[17] SHAW, G., GABEL, R., MARTINEZ, D., ROCCO, A., POHLIG, S., GERBER, A., NOONAN, J., AND TEITELBAUM, K. Multiprocessors for radar signal processing. Tech. Rep. 961, MIT Lincoln Laboratory, Nov. 1992.

[18] SUBHLOK, J., O'HALLARON, D., GROSS, T., DINDA, P., AND WEBB, J. Communication and memory requirements as the basis for mapping task and data parallel programs. In *Proceedings of Supercomputing '94* (Washington, DC, November 1994), pp. 330–339.

[19] SUBHLOK, J., STICHNOTH, J., O'HALLARON, D., AND GROSS, T. Exploiting task and data parallelism on a multicomputer. In *Proceedings of the Fourth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (San Diego, CA, May 1993), pp. 13–22.

[20] SUBHLOK, J., AND VONDRAN, G. Optimal mapping of sequences of data parallel tasks. In *Proceedings of the Fifth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (Santa Barbara, CA, July 1995), pp. 134–143.

[21] SUBHLOK, J., AND YANG, B. A new model for integrated nested task and data parallel programming. In *Proceedings of the Sixth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (Las Vegas, NV, June 1997), pp. 1–12.

[22] VONDRAN, G. Optimization of latency, throughput and processors for pipelines of data parallel tasks. Master's thesis, Dept. of Electrical and Computer Engineering, Carnegie Mellon University, May 1995.

[23] WEBB, J. Implementation and performance of fast parallel multi-baseline stereo vision. In *Computer Architectures for Machine Perception* (New Orleans, LA, Dec. 1993), pp. 232–240.

[24] WEBB, J. Latency and bandwidth consideration in parallel robotics image processing. In *Proceedings of Supercomputing '93* (Portland, OR, Nov. 1993), pp. 230–239.

[25] YANG, B., WEBB, J., STICHNOTH, J., O'HALLARON, D., AND GROSS, T. Do&merge: Integrating parallel loops and reductions. In *Sixth Annual Workshop on Languages and Compilers for Parallel Computing* (Portland, Oregon, Aug 1993).

[26] YANG, T. *Scheduling and Code Generation for Parallel Architectures*. PhD thesis, Rutgers University, May 1993.