

# Hierarchical Scheduling of DSP Programs onto Multiprocessors for Maximum Throughput

Phu Hoang                      Jan Rabaey  
Department of Electrical Engineering  
University of California, Berkeley  
Berkeley, CA 94720

## Abstract

*A multiprocessor scheduling algorithm that simultaneously considers pipelining, retiming, parallel execution and hierarchical node decomposition to maximize performance throughput is presented. The algorithm is able to take into account interprocessor communication delays, and memory and processor availability constraints. The results on a set of benchmarks demonstrate the algorithm's ability to achieve near optimal speedups across a wide range of applications of various types of concurrency, with good scalability with respect to processor count.*

## 1. Introduction

In recent years, we have seen a tremendous increase in the computing power of digital signal processors. However, we have concurrently experienced an even greater increase in computational requirements of DSP applications. Eventually, physical constraints will limit the performance of single processors, and the evolution to multiple processors will be inevitable. The major obstacle to the prevalent use of multiprocessor systems, however, has been the absence of adequate tools to automatically schedule the program onto the multiple processors. A careful analysis reveals 4 major issues involved in finding such a schedule: Concurrency, granularity, estimation, and feasibility.

**Concurrency:** The concurrency can be both *spatial concurrency (parallelism)*, where tasks can be executed by several processors simultaneously, and *temporal concurrency (pipelining)*, where chains of tasks can be divided into stages, with every stage handling results obtained from the previous stage. Pipelining is possible here due to the inherent nature of DSP to repeat the same computation to each frame of the input stream. When there is recursion in the program, however, pipelining is no longer applicable as it may alter the computation intended. One efficient way to obtain speedup in this case is to perform a *retiming* transformation[Lei83], which moves logical delays around in the cycle to minimize the execution time.

**Granularity:** The concurrency to be exploited can exist at all levels of granularity, from functions to loops to individual operations, and the appropriate level must be determined for the scheduler to be both efficient, yet effective in exploiting all the necessary concurrency.

**Estimation:** To obtain meaningful results, accurate estimations of the computation times and memory requirements of the tasks are necessary. A clear model of the underlying processor interconnection is also required to calculate costs due to interprocessor communications.

**Feasibility:** The scheduling tool must reject illegal schedules such as those that do not satisfy data dependencies, those that use more processors or more memory than available, or those that attempt to pipeline where forbidden to do so.

The scheduler to be presented in this paper is a compile-time, hierarchical multiprocessor scheduler that strives to maximize the throughput of the resultant implementation while considering all the issues stated above. The algorithm works on a flow graph representation of the program, where nodes represent computations and edges represent data precedences. The flow graph is hierarchical in that the body of a function call or an iteration is represented by a subgraph, which is contracted into a single node at the next hierarchy level. The hierarchical iteration node has parameters indicating the range of the iteration as well as whether the iteration can be executed in parallel or serial. The flow graph representation of the program exposes the available concurrency explicitly, and is crucial for any multiprocessor scheduling technique.

## 2. CAD Environment

The scheduling technique has been implemented as part of a multiprocessor compiler system called McDAS [Hoa92a]. McDAS currently targets the Sequent and SMART (*Switchable Multiprocessor Architecture supporting Real Time applications*) [Koh89] multiprocessor systems. The Sequent system contains 14 Intel 386 processors connected by a single shared bus. The SMART system is composed of a linear array of 10 AT&T's DSP32C DSP processors connected by a shared configurable bus. The input language used is an extended version of Silage [Hil85], a signal-flow language developed especially for DSP specification. The compilation process involves translating the Silage input into a hierarchical flow graph, scheduling the flow graph, and generating C or assembly code for each partitioned subgraph. The basic concept of the scheduling technique as applied to the SMART system was introduced in [Hoa90]. This paper describes the scheduler independent of the underlying architecture. All aspects of the algorithm are described in greater detail, its computational complexity derived, and more extensive results are given.

## 3. Previous Work

Many techniques have been proposed to schedule DSP programs onto multiprocessors. These techniques can be classified into two categories: *Non-overlap execution scheduling*, which schedules one program execution and replicate the schedule for each input frame, and *Overlap execution scheduling*, which allows for overlapped execution of successive iterations. Typical approaches in the first class are based on *List Scheduling*, which exploits concurrency within one program execution to achieve speedup [Cam85][Kas84][Lee87][Pol86][Pri91][Sar89][Sih90]. All exploits only spatial concurrency, except for [Pri91], which can consider systolic implementations as well. The second class, based on *Cyclo-static Scheduling* [Lee85][Par89][Sch85], takes advantage of the stream processing nature of DSP to exploit temporal concurrency between successive iterations. While these can yield optimum solutions, their exponential search strategies preclude a practical implementation of large applications. *Chain partitioning* [Bok88] pipelines serial tasks on a linear array of processors to maximize throughput. Although this technique does not exploit parallelism, it comes closest to our stated objectives and is a basis for the approach we have adopted.

A number of techniques consider granularity [Sar89][Sih90]. However, none allows the scheduler to decompose a node into smaller nodes when it is a scheduling bottleneck. This is a serious drawback, since the granularity of the nodes (and not the performance of the scheduling algorithm) is often the constraining factor towards reaching a high quality solution.

The remainder of the paper is organized as follows: Section 4 discusses our strategy for estimating computation cost, memory usage, and communication delay. Section 5 presents the scheduling algorithm and analyzes its complexity. The scheduling results on a set of benchmarks are summarized in section 6.

#### 4. Computation, Memory, and Communication Estimation

For static schedules to produce meaningful results, accurate estimates of the computation costs are crucial. These estimates characterize the underlying architecture of the multiprocessor system for the scheduler. To estimate the computation costs, sample programs were benchmarked, and a cost was assigned to each primitive node such as addition and multiplication. With these values, the cost of every node in the hierarchy of the flow graph can be estimated by traversing the graph bottom up, accumulating computation times of primitive nodes into subgraphs, and so on up to the root graph. In [Hoa92b], we demonstrate the technique to be accurate to within 5% of the actual computation costs on the benchmarked architectures.

An edge between two nodes assigned to different processors represents an interprocessor communication. Each communication is supported by a FIFO in our compilation strategy [Hoa92a]. The buffer memory requirement of a node  $n$  on a processor  $p$ , denoted as  $bm(n, p)$ , is dependent on the size of the data and the difference in the source and destination pipeline stages. The parameter  $bm(n, p)$  is used by the scheduler to prohibit a node from being assigned to a processor if executing this node would overflow the processor's buffer memory. To do this, each processor  $p$  has to also keep track of its remaining buffer memory size during scheduling. This parameter is denoted  $bm_{avail}(p)$ . Section 5 will describe how  $bm(n, p)$  and  $bm_{avail}(p)$  are used by the scheduling algorithm.

The communication costs depend on the amount of data being sent and the distance between the source and destination processors. When a node is scheduled on a processor, the data transfers that are needed to bring the input variables to the processor (if data is non-local) are also scheduled on the appropriate bus or busses. This is done by building time-slots of the proposed data transfers and merging it onto the time-slots already scheduled onto the buses. If any proposed time-slot conflicts with an existing time-slot, it is scheduled on the next available time-slot. This allows the scheduler to take bus congestion into account when calculating the arrival time of the data. The abstraction of the communication delays presented to the scheduler is summed up in a parameter  $\xi(n, p)$ , called the *Earliest Starting time* of node  $n$  on processor  $p$ . This parameter makes sense when all predecessor nodes to  $n$  have been scheduled, and represents the earliest time all output data from these predecessor nodes can all be available at processor  $p$  for node  $n$  to start. More details on the time-slot model and the calculation of  $\xi(n, p)$  can be found in [Hoa92b].

#### 5. Scheduling Strategy

We now present a heuristic scheduling algorithm that simultaneously consider pipelining, parallel execution, and retiming to maximize throughput while meeting processor and memory con-

---

straints. The algorithm starts at the top level of hierarchy, and proceeds top-down. Hierarchical nodes are systematically broken down to smaller nodes when more concurrency can be exploited.

**Definition 1:** The *stagetime*  $T$  is defined as the reciprocal of the throughput of the system.

The stagetime equals the time allocated to each pipeline stage, and thus to each processor in that stage. We minimize this stagetime  $T$  to maximize the throughput. During the minimization process, the following condition must always be satisfied:

**Maximum Granularity Condition:** The computation time of the largest node in the graph (at a given level of granularity)  $W_{max}$  has to be  $\leq T$ .

The reason for this is clear. A node of computation time  $> T$  cannot be scheduled on a processor with only time  $T$  to execute.

The algorithm is based on performing a binary search to find the minimal  $T$  given the number of processor  $P$ . The pseudo-code for the algorithm is shown below.

```

Main() :
1. Assign computation times to nodes
2. LB = Lower bounds on T, UB = Upper bounds on T
3.  $W_{max} = \text{MaxWeight}(G)$ 
4. repeat while (LB < UB) {
    1. proc = Schedule(G, T)
    2. if (proc == P) then  $G_{opt} = G$ 
    3. if (proc > P) then LB = T
    4. if (proc ≤ P) then {
        UB = T
        if (T ==  $W_{max}$ ) then {
            G = Expand nodes with weight  $W_{max}$ 
             $W_{max} = \text{MaxWeight}(G)$ 
        }
    }
    5. T = Max ( $W_{max}$ , (LB+UB)/2)
}


```

The upper bound UB is initialized to  $W_{total}$ , the computation time of the entire graph, and the lower bound LB is set to  $W_{total} / P$ , the ideal stagetime given  $P$  processors. The search repetitively calls a routine *Schedule()*, which determines how many processors are needed to execute the algorithm, given a stagetime  $T$ . (This routine also effectively partitions and schedules the flow graph). The processor returned tells whether the proposed  $T$  is too small or too large. The LB and UB bounds are updated accordingly, allowing convergence to the minimum feasible stagetime.


When decreasing  $T$ , care must be taken to ensure that the *Maximum Granularity Condition* is not violated. When there are nodes in the graph with weights as large as  $T$ , and the search decides that  $T$  can be decreased further, it will break up these nodes to find a better solution. We called these nodes *bottleneck* nodes. In this way, large granularity nodes are only decomposed if they block the minimization process, keeping the number of nodes in the graph at a minimum.

The core routine of the search is *Schedule()*. Given the stagetime  $T$ , *Schedule()* traverses the graph from input to output, partitioning the graph into stages of pipelines. Nodes are scheduled

onto a processor until the total computation costs of the nodes plus the communication cost of output edges exceeds the stagetime  $T$ . When two nodes are assigned to the same processor, communications between them incur no cost. Once a pipeline is filled, the scheduler proceeds to schedule the remaining nodes on the next pipeline stage. At the end, the graph is partitioned into a number of pipeline stages, and the total number of processors needed is returned. For a sequential flow



(a) A naive approach



(b) A better approach

Figure 1 : Scheduling a general graph with  $T = 10$

graph, the node accumulation is straightforward. To schedule general graphs, it is necessary to resolve how to handle the parallelism available in branching paths. One naive algorithm would be to continue accumulating nodes to fill the stagetime, whether they are parallel or not. This would result in a schedule shown in Figure 1a. Values inside the nodes represent estimates of their computation costs, and values on the edges represent the additional delays for communication (for the sake of simplicity, the time-slot model is not used here). A more sophisticated algorithm would exploit the parallelism in the graph to yield a schedule shown in Figure 1b. This schedule uses the same number of processors but has a smaller latency and communication cost.

Exploiting parallelism while pipelining makes the scheduling task much more difficult. Since the number of processors is fixed, not all parallelism can be exploited, and the algorithm must decide which operations deserve extra processors and which do not. The naive approach, on the other hand, does not have to do this as it always puts one processor on a pipeline stage. The exact criteria used for node scheduling is discussed in the next subsection.

### 5.1. Node Scheduling

The goal of the node scheduling algorithm is to start the nodes as soon as possible, taking into account communication delays, memory capacity, and processor availability. At any point in the scheduling process, we keep a list of ready nodes  $\mathcal{R}$  and a list of available processors  $\mathcal{P}$ .  $\mathcal{R}$  contains all nodes whose input nodes are already scheduled. Initially, it contains all input nodes.  $\mathcal{P}$  contains all processors that have been assigned to a pipeline stage, plus an extra processor, called the "new" processor. When a node is scheduled on the new processor, the processor is assigned

the pipeline stage appropriate for the node, and another new processor is added to  $\mathcal{P}$ . This allows the scheduler to use as many processors as it deems appropriate. The scheduling steps are as follows:  $\forall n \in \mathbf{N}, \forall p \in \mathcal{P}$ , we calculate the earliest starting time  $\xi(n, p)$ .  $\xi(n, p)$  is set to  $K_\infty$ , a very large constant, if any of the following is true:

1. There is insufficient buffer memory in processor  $p$  to execute node  $n$ , i.e.  $bm_{avail}(p) < bm(n, p)$ .
2.  $p$  was already assigned a pipeline stage which is different from the stage needed to execute  $n$ .
3. There is insufficient available time left in  $p$  to execute  $n$  within the stagetime limit.

A processor  $p$  for which  $\xi(n, p) < K_\infty$  is called a *feasible* processor for  $n$ . Condition 1 assumes the buffer memory is local to the processor. If it is in a centralized memory, the condition would be if  $bm_{avail}(\mathbf{P}) < bm(n, p)$ , where  $bm_{avail}(\mathbf{P})$  is the remaining buffer memory of the entire system.

**Definition 2:** The *Difference Measure*  $\delta(n)$  is defined as:

$$\hat{\xi}(n) = \text{Min} \{ \xi(n, p) \mid p \in \mathcal{P} \} \quad (\text{EQ 1})$$

$$\hat{p}(n) = \{ p \in \mathcal{P} \mid \xi(n, p) = \hat{\xi}(n) \} \quad (\text{EQ 2})$$

$$\tilde{\xi}(n) = \text{Min} \{ \xi(n, p) \mid p \in \mathcal{P} - \hat{p}(n) \} \quad (\text{EQ 3})$$

$$\delta(n) = \tilde{\xi}(n) - \hat{\xi}(n) \quad (\text{EQ 4})$$

$\delta(n)$  gives a measure of how good the best assignment  $\hat{\xi}$  is, compared to the second best  $\tilde{\xi}$ . A node  $n$  with a large  $\delta(n)$  says that the best assignment is much better than the second best, where as a small  $\delta(n)$  says there exist comparable choices. Thus, it is more urgent to assign nodes with a large  $\delta(n)$ . A candidate node and candidate processor are then chosen as follows: The node  $n$ ,  $n \in \mathbf{N}$ , corresponding to the largest  $\delta(n)$  is chosen as the candidate node, and the processor where it achieves its earliest starting time is chosen as the candidate processor. In case of a tie, the first node with the largest  $\delta(n)$  is chosen. Formally, the candidate node  $n^*$  and processor  $p^*$  pair is given as:

$$n^* = \{ n \in \mathbf{N} \mid \delta(n) \text{ is maximum} \} \quad (\text{EQ 5})$$

$$p^* = \hat{p}(n^*) \quad (\text{EQ 6})$$

Table 1 gives the earliest starting time  $\xi(n_i, p_k)$  of three nodes  $n_0, n_1, n_2$ , on three available processors  $p_0, p_1, p_2$ . The  $\xi(n)$ ,  $\hat{\xi}(n)$ , and  $\delta(n)$  values are also shown. The earliest starting node is  $n_0$  on processor  $p_0$ . However, the candidate node chosen is  $n_2$  on  $p_1$  since its alternative choice,  $p_2$ , is a lot worse.  $\xi(n_2, p_0) = K_\infty$ , signifying that  $p_0$  is not a feasible processor for  $n_2$ .

	$w(n_i)$	$\xi(n_i, p_0)$	$\xi(n_i, p_1)$	$\xi(n_i, p_2)$	$\hat{\xi}(n_i)$	$\tilde{\xi}(n_i)$	$\delta(n_i)$
$n_0$	4	4	8	12	4	8	4
$n_1$	5	7	16	10	7	10	3
$n_2$	3	$K_\infty$	6	12	6	12	6

Table 1 : Earliest Starting Time

Once a candidate pair is chosen, the scheduled node is removed from  $\mathcal{N}$ , and new ready nodes are added. If the new processor in  $\mathcal{P}$  was used, another new processor is added to  $\mathcal{P}$ . Processors assigned to pipeline stages which are no longer considered are removed from  $\mathcal{P}$  to avoid unnecessary computations.  $\xi(n, p)$  and  $\delta(n)$  values that are affected by the assignment are updated, and the next node-processor pair is chosen. The scheduling algorithm ends when all nodes are scheduled. The pseudo-code for `Schedule()` is described below:


**Schedule(G,T):**

1. Input nodes  $\rightarrow \mathcal{N}$ ,  $p_0 \rightarrow \mathcal{P}$
2. Repeat while  $\mathcal{N} \neq \emptyset$ 
  1. For each  $n \in \mathcal{N}$ ,  $p \in \mathcal{P}$  do
    - Calculate  $\xi(n, p)$
    - Calculate  $\delta(n)$
  2. Schedule candidates  $n^*$  and  $p^*$
  3. Update  $\mathcal{N}$ ,  $\mathcal{P}$

The complexity of `Schedule()` is derived as follows: Let the number of nodes be  $N$ , the number of edges be  $E$ , and the number of available processors be  $P$ . Given a fixed processor  $p_k$  for each new node  $n$  put in  $\mathcal{N}$ , the calculation of  $\xi(n, p_k)$  requires calculating the arrival time of each input edge. For the whole graph, all edges are visited, for  $O(E)$  calculations. Adding and removing  $N$  nodes to and from  $\mathcal{N}$  takes  $O(N)$  time. Hence the total computation runs in  $O(N+E)$  time. Since  $P$  processors are considered for each node, the complexity is  $O(P(N+E))$ . Finally, for nodes in  $\mathcal{N}$  which were not chosen in a scheduling step,  $\xi(n_i, p^*)$  is updated in constant time. Since there are at most  $N$  such nodes in  $\mathcal{N}$ , at most  $N$   $\xi(n, p)$  updates are made each scheduling step. For  $N$  scheduling steps,  $N^2$  calculations are required. The total complexity is  $O(P(N+E)+N^2)$ . Finally, since the algorithm can actually continue to add processors beyond  $P$ , the potential number of processors considered can be  $N$ , the number of nodes. The final complexity is  $O(N(N+E))$ .

**5.2. Path Merging**

In its attempt to maximize parallelism, the algorithm may under-utilize a processor. This would be the case if, for example, the node assigned to processor 1 in Figure 1b were to have a cost of 1 instead of 8, as shown in Figure 2. To solve this problem, we can merge these parallel paths so they are forced to execute on 1 processor. This is achieved by adding a dependency edge to serialize the computation, as shown in Figure 2.



**Figure 2: Path Merging to improve processor utilization**

Parallel paths are characterized by *Branch Nodes*, i.e. nodes that have 2 or more output nodes. All parallel paths can be considered by examining only branch nodes.

**Definition 3:** The *Overlap Path time*  $t_{op}$  between 2 parallel paths measures the time both paths are simultaneously active. This is equivalent to the elapsed time between the branch node and the point where either the two paths join again or one of them terminates.

$t_{op}$  is the minimum of the computation costs of the 2 paths, and it measures the amount of parallelism that exists between the paths. The smaller the  $t_{op}$ , the less the parallelism that can be exploited. The pair of parallel paths with the smallest  $t_{op}$  is the best candidate for merging as it under-utilizes its processor the most. To facilitate the calculation of  $t_{op}$ , the transitive closure matrix of the flow graph is used to quickly locate the merging point of paths. In addition, the nodes are leveled from output using their computation costs as weights, allowing the computation cost of a path to be given by the difference between the output levels  $ol(n)$  of the branch and join nodes. An  $O(N^2)$  algorithm for calculating  $t_{op}$  is given as:

**OverlapPathTime(G):**

```

For each branch node  $n_i \in N$  do
  For each pair of nodes  $(n_j, n_k)$  which are successor nodes to node  $n_i$  do
    1. Locate join node  $n_{jk}$  of  $(n_j, n_k)$  as the node with the highest output level which is a
       descendant node to both  $n_j$  and  $n_k$ 
    2.  $t_{op}(n_j, n_k) = \text{MIN}[ol(n_j), ol(n_k)] - ol(n_{jk})$ 

```

After the smallest parallel path is found, a routine *MergePath()* adds a dependency edge from one end of the path to the beginning of the other.  $t_{op}$  is then updated taking the new dependency into account. The *Schedule()* routine is repeated along with *MergePath()* as long as it needs more processors than available and as long as there are still parallel paths to merge. Upon exit of the loop, we either have a feasible schedule using  $P$  or less processors, or an infeasible schedule ( $\text{proc} > P$ ) which uses the minimal number of processors for the given stagemtime  $T$ . For this latter case, a longer stagemtime is needed. The overall routine, called *Partition()*, replaces the *Schedule()* routine in *Main()*:

**Partition(G, T):**

```

1.  $\text{proc} = \text{Schedule}(G, T)$ 
2. Repeat while ( $\text{proc} > P$  &&  $\text{MergePath}(G) == \text{TRUE}$ )
   1. Update Transitive Closure Matrix, Output Level
   2.  $\text{proc} = \text{Schedule}(G, T)$ 

```

The calculation in the inner loop involves updating  $t_{op}$ , finding the minimum  $t_{op}$ , adding a dependency edge, and scheduling the graph. Updating  $t_{op}$  involves updating the output levels and the transitive closure graph. To update the output levels, only the source node of the newly added dependency edge and its ancestors need to have their output levels recalculated. This takes  $O(N)$  time. In the transitive closure graph, an edge has to be added between the source node and its ancestors to each descendants of the destination node. This takes  $O(N^2)$  time. Locating the minimum  $t_{op}$  takes  $O(N^2)$ , and scheduling takes  $O(N(N+E))$  time. This inner loop repeats at most  $N$  times since there are at most  $N-1$  merges possible. Hence, the algorithm runs in  $O(N^2(N+E))$  time. For most examples that we have encountered,  $E \approx N$ , giving an  $O(N^3)$  approximation. Figure 3 shows a sample flow graph and the improvement in processor utilization using path merging. Each node in the graph has a cost of 1.

### 5.3. Retiming

The scheduling algorithm presented so far works well for flow graphs which do not contain cycles. For those that do, some enhancements to the algorithm are needed. In this section, a modi-

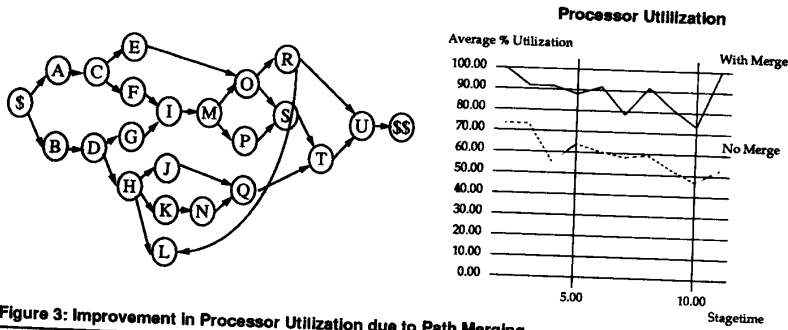


Figure 3: Improvement in Processor Utilization due to Path Merging

modification to the algorithm is presented to allow the retiming of flow graph cycles to improve scheduling.

Consider Figure 4a. In order to execute node A, we need the data d from 2 samples back in time, which we will denote as  $d@2$ . If we apply our scheduling strategy to the flow graph without considering the feedback path, we may pipeline the computation as shown in Figure 4b. Since node D is now 3 samples behind node A,  $d@2$  will not be available when node A needs it. Thus it is necessary to limit the number of pipeline stages on the forward path to ensure that proper data is available when it is needed. Referring to Figure 4a, assume node A is applied to sample 0 at time 0. Since the sample period is the stagetime T, node A will be applied to sample 2 at time 2T. To execute this node, the output data of node D on sample 0 must be available. This implies that the execution of the nodes A, B, C, D on sample 0 (indeed, on all samples) must complete within time 2T. In general, if  $\Delta_c$  is the number of delays in the cycle, then the cycle must execute within time  $\Delta_c T$ . The following definitions formalize this bound of the execution of cycles.




Figure 4: Retiming flow graphs

**Definition 4:** Let G be the flow graph at some granularity level, and let C be any cycle in G. Let  $\Delta_c$  be the number of delays in cycle C,  $W_c$  be the total computation time of cycle C, and  $E_c$  be the total time a valid schedule would need to execute cycle C.

**Theorem 1:** For any cycle C,

$$W_c \leq E_c \leq \Delta_c T \tag{EQ 7}$$

The proof for  $E_c \leq \Delta_c T$  is an obvious generalization of the argument above.  $W_c \leq E_c$  because a schedule may not necessarily execute cycle C contiguously.

With this result, we get the following condition, called the *cycle scheduling bound condition*:

**Lemma 1.1:** Let  $n_0, n_1, \dots, n_N$  be the set of nodes in cycle  $C$ , and let  $n_0$  be the first scheduled node among them. Let  $t_{start}(n_0)$  be the starting time of node  $n_0$ , and  $t_{completion}(n_k)$  be the completion time of any node  $n_k$  in the cycle. A valid schedule must satisfy:

$$t_{completion}(n_k) - t_{start}(n_0) \leq \Delta_c T \quad (\text{EQ 8})$$


This follows directly from the fact that  $t_{completion}(n_k) - t_{start}(n_0) \leq E_c$ , for every node  $n_k$  in the cycle. Lemma 1.1 will be used by the scheduling algorithm to ensure the cycle scheduling bound is satisfied when nodes of a cycle are scheduled.

**Lemma 1.2:** From theorem 1, we also have for every cycle  $c$ ,  $T \geq W_c / \Delta_c$ . In particular,  $T \geq T_{CB}$ , where  $T_{CB}$  is the *Stagetime Cycle Bound* given by:

$$T_{CB} = \text{Max}_c W_c / \Delta_c \quad (\text{EQ 9})$$

$W_c$  is a function of the granularity level of the graph. As the granularity gets finer, the computation times of cycles decrease as hierarchical nodes in the cycles are replaced by only those nodes of the subgraph which belong to the cycle. In this case,  $T_{CB}$  also decreases. Some cycles are entirely imbedded in hierarchical nodes, and are only uncovered when these nodes are expanded. For that case,  $T_{CB}$  may actually increase. Once all cycles are uncovered, however,  $T_{CB}$  decreases monotonically to a bound known as the *Iteration Period Bound*  $T_{IPB}$  [Sch85], the stagetime cycle bound for the finest granularity graph. This bound is the minimum achievable latency between same iterations, and is a theoretical lower bound for our solution.

When the scheduler makes a partition on the forward path, it is in effect putting a logical delay there. To maintain correct functionality of the flow graph, it is necessary to remove a delay in the feedback path. This can be interpreted as moving delays in the cycle around to maximize throughput, a concept known as retiming [Lei83]. By choosing  $T \geq T_{CB}$ , one might conjecture that sufficient time is allocated to each pipeline stage to guarantee that at most  $\Delta_c$  partitions are made in scheduling cycle  $c$ , thereby automatically adhering to the *cycle scheduling bound*. This, unfortunately, is only guaranteed if the nodes in the cycle are scheduled contiguously, the stagetime is exploited completely, and communication delays are not considered. In practice however, data transfers are often present, and using the stagetime completely is difficult due to the granularity of the nodes. An example of a cycle scheduling violation due to large granularity of the nodes is shown in Figure 5. From Figure 5a, the stagetime cycle bound  $T_{CB}$  is found to be  $30/3 = 10$ .



**Figure 5: Partitioning Cycles**

Using a stagetime  $T = T_{CB} = 10$ , we find that we need 4 partitions, which violates the Cycle scheduling bound condition (Figure 5b). The weight of the nodes are too big to fit in the remaining

available time, leaving holes in the stagetime, which we called *slacks*. If the nodes are broken down into smaller nodes, as in Figure 5c, a feasible partition with no slacks can be obtained (Figure 5d). Communication costs are not considered in this example.

Using the analysis above, the scheduling algorithm presented so far can easily be modified to handle cycles. From Lemma 1.2, the stagetime lower bound LB is modified to  $\text{Max}(W_{total} / P, T_{IPB})$ , to include the theoretical lower bound due to cycles. Each time a node  $n_i$  in a cycle is scheduled in `Schedule()`, Lemma 1.1 is used to ensure that the cycle scheduling bound for the node is satisfied. If it is violated, the large nodes in the cycle are decomposed in an attempt to minimize slacks in the stagetime. This is repeated until there is no more violation or no more nodes in the cycle can be decomposed. In the latter case, LB is immediately updated to the current T, as no feasible solution can be found at any granularity level for this stagetime T. This extension to `Partition()` is called `CyclePartition()`, and is given as:

`CyclePartition(G,T):`

1. (proc, ViolateFlag) = `Partition(G, T)`
2. **Repeat while** (ViolateFlag == TRUE && `IsGraphFlat(G) == FALSE`)
  - G = Expand nodes in critical cycles
  - Find all cycles in G
  - Update  $W_{max}$  and  $T_{CB}$
  - (proc, ViolateFlag) = `Partition(G,T)`

In this routine, `Partition()` is called repetitively on finer and finer granularity flow graphs to satisfy the cycle scheduling bound. When the node decomposed in the cycle is an iteration, each decomposition increases the number of processors allocated to the iteration by 1, and terminates when the number of processors reaches P. Thus, `Partition()` repeats at most P steps when decomposing iterations. When decomposing functions, it may repeat more, and there is no maximum bound. Benchmarks show that the typical number of function node decompositions is less than five. The main computations in each pass is the search for all cycles in the flow graph, and the `Partition()` routine. For cycle detection, we use Johnson's algorithm for finding all the elementary cycles of a directed graph [Joh75], which is the fastest algorithm known. It has a time of complexity of  $O((N+E)(C+1))$ , where C is the number of cycles in the graph. From the last section, `Partition()` was analyzed to run in  $O(N^2(N+E))$  time. Hence, the running time of `CyclePartition()` is  $O(P(N+E)[N^2 + (C+1)])$ .

The updated main algorithm is shown below. Note that  $T_{CB}$  is updated each time a node is decomposed, and the proposed stagetime T is modified to include the  $T_{CB}$  lower bound.

`Main():`

1. Assign computation times to nodes
2. LB = Lower bounds on T, UB = Upper bounds on T
3.  $W_{max} = \text{MaxWeight}(G)$
4.  $T_{CB} = \text{CycleBound}(G)$
5. **repeat while** (LB < UB)
  1. `CyclePartition(G,T)`
  2. **if** (proc == P) **then**  $G_{opt} = G$
  3. **if** (proc > P) **then** LB = T
  4. **if** (proc ≤ P) **then** {
    - UB = T

```

    if (T ==  $W_{max}$ ) then {
      G = Expand nodes with weights  $W_{max}$ 
      Update  $W_{max}$  and  $T_{CB}$ 
    }
    if (T ==  $T_{CB}$ ) then {
      G = Expand nodes in critical cycles
      Update  $W_{max}$  and  $T_{CB}$ 
    }
  }
  5. T = Max ( $W_{max}$ ,  $T_{CB}$ , (LB+UB)/2)
}

```

#### 5.4. Node Decomposition

Node decomposition is the means by which the scheduling algorithm traverse from a coarse grain flow graph to a finer grain flow graph. In previous subsections, we saw that nodes are broken up on two occasions: When they are larger than the available stagetime  $T$ , and when they are part of a cycle that violates the cycle scheduling bound. The first case is denoted as *bottleneck node decomposition* and the second case as *critical cycle decomposition*. There are two types of hierarchical nodes that are candidates for decomposition: Function nodes and Iteration nodes. Function nodes are decomposed by replacing the nodes with their subgraphs. Iteration nodes are replaced by parallel or serial sub-nodes (depending on the data dependencies between iterations), each computing a subset of the iteration range. The method of breaking up the iterations differs between the decomposition of bottleneck nodes and the decomposition of critical cycles. In both cases, the size of the sub-nodes is determined *dynamically* during the scheduling procedure.

When decomposing bottleneck iteration nodes, the sub-nodes adapt to the available time remaining in the stage. The number of iterations assigned to the sub-nodes are not fixed until the first sub-node is scheduled. At that time, the first sub-node is assigned as many iterations as can fit in the remaining stage, and subsequent sub-nodes are assigned as many iterations as can fit in a new, empty stage with stagetime  $T$ . This partitioning strategy allows the iterations to float across processors from one pass of the scheduler to the next to fit the changing stagetime  $T$ . When decomposing iteration nodes in critical cycles, the goal is to obtain a partition which satisfies the *cycle scheduling bound* condition. For nodes with serial dependency, decomposition does not decrease the cycle computation time, but may improve the stagetime slacks. Therefore, the same technique as discussed above is used. For nodes with parallel dependency, decomposition actually decreases the computation time of the cycle, making it easier to meet the *cycle scheduling bound*.




Figure 6: Decomposing parallel iterations in cycles.

Consider a cycle containing a parallel iteration node as shown in Figure 6. Assuming a tight stagetime  $T = T_{CB} = 35$  and no communication delays, scheduling Figure 6a with  $t_{start}(n_A) = 0$  would yield  $t_{completion}(n_E) = 2 * 35 + 10 = 80$ . Since  $t_{start}(n_A) + 2T = 70$ , the *cycle scheduling bound* is violated. In Figure 6b, after decomposition, the stagetime cycle bound is reduced from 35 to 25.  $T$  is no longer tight, and scheduling yields  $t_{completion}(n_E) = 1 * 35 + 20 = 55$ ,  $t_{start}(n_A) + 2T = 70$ , meeting the *cycle scheduling bound* condition. Of course, the price we pay is the 2 additional processors needed to execute the iterations in parallel. Note that the naive scheduling algorithm described in Figure 1a would not be able to improve the stagetime at all, as the parallelism of the iteration is not exploited.

The goal in the parallel iteration node case is to parallelize the node just enough to satisfy the cycle scheduling bound condition, as excessive parallelizing would only use more processors than necessary. To accomplish this, the parallel iteration node is incrementally divided into equal weight sub-nodes until a partition with no violation or the maximum decomposition is reached. Since the cycle scheduling bound is tightest on the sub-node with the maximum computation, a division into equal weight sub-nodes maximizes the chance of meeting the bound.

In our approach, loops can only be divided at the boundary of each iteration. No attempt will be made, for instance, to partition 3.5 iterations of a loop on one processor and the remainder on another.

### 6. Results

The first example is a histogram algorithm to be mapped onto a Sequent multiprocessor. The input is an array of 128 samples, to be partitioned into 32 subclasses. The computation proceeds as follows: First, the minimum and maximum values of the array are determined. Next, the width of each subclass is calculated from the total range. The third task fills in all the subclasses. This is a parallel task as each subclass can scan the array and count how many numbers fall into its subclass. The fourth task determines the baseline class, the subclass whose range covers 0. Finally, the last task computes the number of samples in the 5 subclasses around the baseline class. Figure 7 shows the same flow graph partitioned onto 6 processors of the Sequent. The computation was




Figure 7: Histogram Example scheduled on 6 processors

pipelined into 2 stages, with processor 0 in stage 0 and processors 1-5 in stage 1. Analyzing the completion times of the processors as estimated by the scheduler and as actually measured on the Sequent (Figure 8), we see that the load was evenly distributed, and the estimated time agrees very well with the actual running time. The same example was then scheduled and executed on differ-

ent numbers of processors to assess how the speedup varies with the processor count. Figure 9 shows the speedup as calculated by the scheduler, and the actual speedup on the Sequent, both as functions of the processors. For this example, we observe that the scheduler is able to consistently achieve a faster throughput with each additional processor. For comparison, the ideal speedup is shown. This can only result if there is perfect load balancing and there is no cost for interprocessor communication.




Figure 8: Load Balancing on 6 processors




Figure 9: Speedup vs. Processors

The second example is a CORDIC algorithm mapped onto the SMART multiprocessor. It converts cartesian to polar coordinates iteratively in 20 steps. It takes as input an (X,Y) coordinate, as well as an array of correction angles. After an initialization, there is a sequential loop of 19 iterations. The scheduler is able to achieve good speedup by pipelining the loop and assigning successive loop iterations to successive processors. In contrast to the Histogram example, the CORDIC program is communication intensive. A linear array architecture thus can execute this program more efficiently than a single shared bus architecture. Figure 10 plots the speedup for the CORDIC example. The "stair case" effect of the speedup curves results




Figure 10: Speedup vs. Processors




Figure 11: % Communication vs. Processors

from the fact that the scheduler does not partition loops at the middle of an iteration. As a result, when there are 12 to 18 processors available, the scheduler still has to assign at least 1 processor 2 iterations to execute. The throughput remains constant until enough processors are available to allow each iteration to be executed by its own processor. This occurs at 20 processors, after which a great leap in throughput is attained. At this same time, the communication load on the bus also

increases dramatically as all processors are communicating (Figure 11). The linear array architecture can support this neighbor-to-neighbor communication well, yielding a better speedup than a shared bus architecture.

Table 2 shows the scheduling results of a number of other DSP examples with different types of concurrency and communication patterns. All assume that the number of available processors is 8, and a single shared bus interconnection. F(G) and H(G) give the number of nodes in the flattened and hierarchical graph, respectively. R(G) gives the number of nodes considered by the scheduler during its top-down search. The CPU measurements is on a Sun Sparc II, and include the Silage to Flowgraph compilation.

Example	F(G)	H(G)	R(G)	Concurrency	Speedup	CPU (sec)
DTW	1.7e8	98	15	Pipe/Par	7.08	30.2
Mat. Mult	2.0e6	24	11	Pipe/Par	6.64	5.8
256pt DFT	7.6e5	35	9	Pipe/Par	6.94	8.2
Pitch Extractor	1.2e5	270	22	Pipe/Par	7.53	21.7
2-Norm	1926	23	12	Pipe/Par	7.61	11.2
8pt DCT	87	62	75	Pipe/Par	3.26	89.3

Table 2: Scheduling Results

The Dynamic Time Warp (DTW) algorithm [Sak78] for speech recognition processes 1000 templates to compute the score of each template with respect to an unknown signal. The matrix multiplication example multiplies two 64x64 matrices. The DFT algorithm computes a 256-pt Discrete Fourier Transform in 2 nested loops. The pitch extractor [Slu80] derives a candidate pitch from an input signal and compare it against 20 templates. The 2-norm squares and accumulates each element of a vector (size 128). All of these examples are computationally intensive and require little inter-processor communications. Hence, good speedup was achieved. The DCT example has a 51% communication overhead, and yields a poor speedup on a shared bus. Table 3 shows 3 more examples: The Adaptive Differential Pulse Code Modulator (ADPCM), the Decision Feedback Equalizer (DFE), and the Echo Canceller [Hon84]. All contained single delay recursions in their computation, which prohibited pipelining. Nevertheless, they all contain parallel iterations in the cycle. In the ADPCM and DFE examples, the stagetime obtained was close to the cycle bound, although the speedup was nowhere near 8. Faster speedup for these examples can only be obtained by transforming and reorganizing the algorithm themselves. In the echo canceller case however, 8 processors were not enough to exploit all the parallelism in the cycle. As a result, each processor allocated to the parallel iteration still has to perform a number of iterations sequentially. In this case, additional processors will allow for a further reduction of the stagetime.

Example	Concurrency	Stagetime	Cycle Bound	Speedup
ADPCM	Retime/Par	1381.5	1362.8	3.72
DFE	Retime/Par	1500.8	1238.5	4.22
Echo Canc	Retime/Par	12157	6649	2.34

Table 3: Scheduling Results (continue)

## 7. Conclusion

A flow graph scheduling algorithm that simultaneously considers pipelining, retiming, parallelism and hierarchical node decomposition is presented. The ability to simultaneously consider the many types of concurrency allows the scheduler to find efficient multiprocessor solutions for a wide range of DSP applications. The results on a set of benchmarks demonstrate that the algorithm still achieves near optimal speedups across programs with different types of concurrency.

### References:

- [Bok88] S. Bokhari, "Assignment Problems in Parallel and Distributed Computing," *Parallel Processing and Fifth Generation Computing*, Kluwer Academic Publishers, 1988.
- [Cam85] M.L. Campbell, "Static Allocation For a Data Flow Multiprocessor," in *Proc. 1985 Int. Conf. Parallel Processing*, 1985, pp. 511-517.
- [Hil85] P. Hilfinger, "SILAGE, A High Level Language and Silicon Compiler for Digital Signal Processing," *Proceedings IEEE CICC Conference*, Portland, May 1985.
- [Hoa90] P. Hoang, and J. Rabaey, "Program Partitioning for a Reconfigurable Multiprocessor System," *IEEE Workshop on VLSI Signal Processing IV*, November 1990.
- [Hoa92a] P. Hoang, and J. Rabaey, "A Compiler for Multiprocessor DSP Implementation," *ICASSP*, March 1992.
- [Hoa92b] P. Hoang, "Compiling Real-Time Digital Signal Processing Applications onto Multiprocessor Systems," Ph.D. thesis, University of California, Berkeley, 1992.
- [Hon84] M.L. Honig, and D.G. Messerschmitt, *Adaptive Filters: Structures, Algorithms, and Applications*, Kluwer Academic Publishers, 1984.
- [Joh75] D. Johnson, "Finding All The Elementary Circuits of a Directed Graph," *Siam Journal of Computing*, vol. 4, No. 1, March 1975.
- [Kas84] H. Kasahara, and S. Narita, "Practical Multiprocessor Scheduling Algorithms for Efficient Parallel Processing," *IEEE Transactions on Computers*, vol. c-33, no. 11, November 1984.
- [Koh89] W. Koh, and A. Yeung, P. Hoang, J. Rabaey, "A Configurable Multiprocessor System for DSP Behavioral Simulation," *ISCAS Symposium*, May 1989.
- [Lee87] E.A. Lee, and D.G. Messerschmitt, "Static Scheduling of Synchronous Data Flow Programs for Digital Signal Processing," *IEEE Trans. Comput.*, vol. C-36, January 1987.
- [Lee85] S.H. Lee, and C.J.M. Hodges, T.P. Barnwell III, "An SSIMD Compiler for the Implementation of Linear Shift-invariant Flow Graphs," *ICASSP*, March 1985.
- [Lei83] C.E. Leiserson, and F.M. Rose, J.B. Saxe, "Optimizing synchronous circuitry by retiming," *Proceedings of Third Caltech Conference on VLSI*, Computer Science Press, 1983, pp. 23-36.
- [Par89] K.K. Parhi, "Rate-Optimal Fully-Static Multiprocessor Scheduling of Data-Flow Signal Processing Programs," *ISCAS*, May 1989.
- [Pol86] C. Polychronopoulos, and U. Banerjee, "Speedup Bounds and Processor Allocation for Parallel Programs on Multiprocessors," *Proceedings of the 1986 International Conference on Parallel Processing*, 1986.
- [Pri91] H. Printz, "Automatic Mapping of Large Signal Processing Systems to a Parallel Machine," *Carnegie Mellon University Ph.D. Thesis*, May 1991.
- [Sak78] H. Sakoe, and S. Chiba, "Dynamic Programming Algorithm Optimization for Spoken Word Recognition," *IEEE Trans. Acoustics, Speech, and Signal Processing*, vol. ASSP-26, pp. 43-49, 1978.
- [Sar89] V. Sarkar, "Partitioning and Scheduling Parallel Programs for Multiprocessors," *Research Monographs in Parallel and Distributed Computing*, MIT Press, 1989.
- [Sch85] D.A. Schwartz, "Synchronous Multiprocessor Realizations of Shift-invariant Flow Graphs," *Georgia Tech Ph.D. Thesis*, June 1985.
- [Sih90] Sih, G. and E.A. Lee, "A Multiprocessor Scheduling Strategy," *ERL Internal Report*, UC Berkeley, 1990.
- [Slu80] R.J. Sluyter, and N.J. Kotmans, A.V. Leeuwaarden, "A Novel Method for Pitch Extraction from Speech and a Hardware Model Applicable to Vocoder Systems," *ICASSP*, 1980, pp. 45-48.