

Exploiting Application Tunability for Efficient, Predictable Parallel Resource Management

Fangzhe Chang, Vijay Karamcheti, and Zvi Kedem
Courant Institute of Mathematical Sciences
New York University, New York, NY 10012
{fangzhe,vijayk,kedem}@cs.nyu.edu

Abstract

Parallel computing is becoming increasing central and mainstream, driven both by the widespread availability of commodity SMP and high-performance cluster platforms, as well as the growing use of parallelism in general-purpose applications such as image recognition, virtual reality, and media processing. In addition to performance requirements, the latter computations impose soft real-time constraints, necessitating efficient, predictable parallel resource management. In this paper, we propose a novel approach for increasing parallel system utilization while meeting application soft real-time deadlines. Our approach exploits the application tunability found in several general-purpose computations. Tunability refers to an application's ability to trade off resource requirements over time, while maintaining a desired level of output quality. We first describe language extensions to support tunability in the Calypso system, then characterize the performance benefits of tunability, using a synthetic task system to systematically identify its benefits. Our results show that application tunability is convenient to express and can significantly improve parallel system utilization for computations with predictability requirements.

1. Introduction

Parallel computing has become central and mainstream, driven both by a decrease in cost and a growth in its general-purpose applicability. The commodity nature of small-scale symmetric multiprocessors (SMPs) and high-performance cluster platforms [4] has ensured the widespread availability of relatively inexpensive hardware platforms. Moreover, these platforms are increasingly being used to run general-purpose applications such as image recognition, virtual reality, and media processing [6]. However, these latter applications introduce additional demands on the management of

resources in a parallel system. On top of the performance requirements traditionally associated with parallel applications (e.g., floating-point intensive scientific codes), these applications impose additional soft real-time constraints requiring completion of different portions of the application within specific intervals of time. For example, an application that is trying to analyze a live video feed to recognize important artifacts needs to complete its processing by the time the next frame arrives. Thus, such applications necessitate *efficient, predictable parallel resource management* that must simultaneously optimize resource utilization and ensure that applications meet their deadlines.¹

Unfortunately, traditional resource management approaches in both parallel and real-time systems are inadequate for meeting these objectives. The parallel approaches focus primarily on improving application performance and/or system utilization at the cost of providing only *best effort* guarantees to the application. A specific application can experience arbitrary delay which may grow with the number of applications contending for the resources. Clearly, such delays are unacceptable for soft real-time applications that work with continuous media. Real-time systems are better at providing predictable guarantees to applications. However, they do so by being overly conservative, ensuring that enough resources are available for each application to meet its deadline. Admission control is used to ensure an underloaded system, providing application predictability at the cost of system utilization.

In this paper, we propose a novel approach for increasing parallel system utilization while meeting application soft real-time deadlines. Our approach exploits the *application tunability* found in several general-purpose computations such as the ones described above. Tunability refers to an application's ability to trade off resource requirements over time, while maintaining a desired level of output quality. In other words, a large allocation of resources in one stage of

¹In this paper we assume that processors are the primary resource that is being managed: applications request non-preemptive allocation of a specific number of processors for a fixed amount of time.

the computation’s lifetime may compensate, in a parameterizable manner, for a smaller allocation in another stage. For example, the artifact recognition application may first sample different portions of the image to decide on interesting regions, and then run a resource intensive algorithm on these regions; spending more resources on the sampling step reduces the work that will need to be performed in the analysis step. Application tunability provides flexibility to the underlying resource management system, which can now use the choice in resource allocation profiles to increase the number of applications that can be admitted into the system, while still ensuring that an application meets its real-time requirements.

This paper describes support for predictable, tunable applications in the MILAN parallel and distributed computing system [14]. We first describe language extensions to the Calypso programming language for expressing application tunability, and evaluate their effectiveness by describing the construction of a tunable image processing application (junction detection). We then characterize the performance benefits of application tunability, by first proposing competitive heuristics for the underlying scheduling problem, and then using a synthetic task system to systematically identify the benefits and shortcomings of tunability. We find that tunability helps improve both system utilization and job throughput (the number of jobs which meet their deadlines) over a wide range of task parameters, such as arrival rate, deadline laxity, and the shape of the required resource profile. Our results are very encouraging: application tunability is convenient to express, and can significantly improve parallel system utilization for computations with predictability requirements.

The rest of this paper is organized as follows. Section 2 provides relevant background on the MILAN system and the Calypso programming language. In Section 3, we present the overall MILAN resource management architecture for predictable computations. Section 4 describes extensions to the Calypso source language for supporting tunability. The performance impact of tunability is characterized in Section 5. Section 6 places our work in context with related efforts and we conclude in Section 7.

2. Background and Project Context

The MILAN system supports robust, predictable parallel computation over a possibly heterogeneous collection of resources. MILAN takes advantage of two execution techniques with strong theoretical foundations [5]—*two-phase idempotent execution strategy*, and *eager scheduling*—to provide programmers with the view of a fault-free virtual shared memory environment, even when the underlying resources may incur faults and exhibit wide variations in processing speeds. This support is exposed to the programmer

in the form of several programming systems: Calypso [1] described in further detail below, Chime [16] which supports distributed execution of C++ [3] programs, and Charlotte [2] which provides a web-based metacomputing infrastructure. In addition, the MILAN system consists of supporting infrastructure such as ResourceBroker (a system for dynamically managing the association and integration of resources into multiple parallel computations according to user-specified policies).

The Calypso programming system views computations as consisting of several parallel tasks inserted into a sequential program. These parallel tasks are responsible for performing the computationally intensive work, while the sequential code is responsible for the high-level control-flow and I/O. Within a parallel step, Calypso supports CREW (concurrent read, exclusive write) semantics to shared data structures, with updates visible only at the end of the current step. Additionally, the parallel tasks are idempotent, implying that a code segment can be executed multiple times (with possibly some partial executions), with exactly-once semantics. These multiple executions mask any faults in the underlying resources.

Calypso augments standard C++ with four keywords: `shared`, `parbegin`, `parend`, and `routine`. Globally shared variables are declared using the keyword `shared`. `parbegin` and `parend` help delimit a parallel step consisting of a sequence of `routine` statements:

```
parbegin
  routine [int-exp](int width, int number)
    routine-body1
  routine [int-exp](int width, int number)
    routine-body2
parend;
```

The `routine` statements specify the tasks within the parallel step: `routine-body1` and `routine-body2` are sequential C++ program fragments, `int-exp` specifies an integer expression indicating the number of copies of each routine to be created within the parallel step, and `width` and `number` are arguments provided to each task denoting, respectively, the number of tasks created and the sequence number of the specific task among these tasks. As shown in the code fragment above, each parallel step may consist of multiple `routine` statements. Concurrency exists both inside one routine, as well as among multiple routines within the same parallel step.

3. Resource Management Architecture

The MILAN resource management architecture, shown in Figure 1, consists of two major components: an application-level *QoS agent* and the system-level *QoS arbitrator*.² The QoS agent communicates the application re-

²The component names signify our focus on providing predictable *quality of service* (QoS) for applications. Here, quality refers both to the

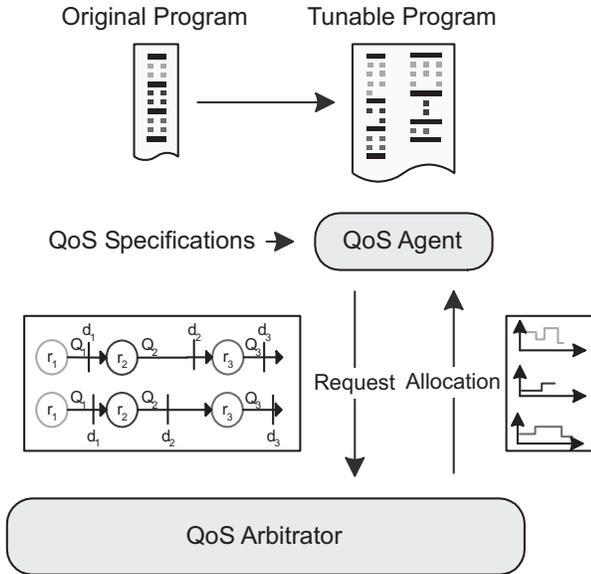


Figure 1. MILAN QoS arbitration architecture

source and predictability requirements to the QoS Arbitrator, which satisfies this request (and those from other applications) by providing an appropriate resource allocation. We discuss these components in detail below.

3.1. Components of the Architecture

QoS agent The QoS agent, automatically generated from the application’s specification by a preprocessing step (see Section 4 for details), describes the application’s real-time constraints, its resource requirements, and more importantly its tunability. Tunability is expressed by indicating choice in the execution path available for the application. As Figure 1 shows, from the perspective of the QoS agent, the application is viewed as an execution path (a chain, or more generally, a dag) comprising several tasks, each with its own resource requirements and deadlines. Resource requirements can be thought of as a vector of values, one for each resource in the system. Each task also has an associated output quality, closely related to the requested resources. The quality value of the execution path is obtained by composing the output qualities of each of the tasks. Tunability is expressed by specifying multiple such execution paths, each with its own resource requirement and deadline profiles, representing alternate ways in which the application can consume resources in order to produce outputs with the desired quality.

The QoS agent acts on behalf of the application to negotiate with the QoS arbitrator an appropriate level of resource reservation/allocation for each task, maximizing the

desired outputs as well as associated timeliness guarantees.

application output quality. In general, this negotiation involves an initial allocation that gets revised as a function of changing application demands and/or changing system conditions. For the results reported in the rest of the paper however, we restrict our attention to a relatively static negotiation model: the QoS agent communicates all the possible application execution paths and their resource requirements up front, and receives in return (from the QoS arbitrator) a resource allocation profile for one of these paths.

QoS arbitrator The QoS arbitrator takes advantage of the flexible program specification provided by QoS agents to enhance system utilization while satisfying the predictability requirements of each application. In MILAN, this flexibility comes from two aspects. First, application tunability provides the freedom to choose a resource allocation over time for each application. And second, the adaptiveness of the underlying fault-masking techniques (two-phase idempotent execution and eager scheduling) provide micro-level flexibility, permitting preemptive allocation, deallocation, and reallocation of resources to each parallel step of the application. In this paper, we restrict our focus to the flexibility obtained from application tunability.

Upon job arrival, the QoS arbitrator first performs admission control to check whether or not application resource requirements can be satisfied. Application tunability increases the likelihood that an application can be admitted into the system. The QoS arbitrator scheduling algorithms (discussed in Section 5) first choose the best execution path, and then make an assignment of which processors will execute which application tasks and for what time. These decisions are communicated back to the application QoS agent which configures the application appropriately. In general, the QoS arbitrator also monitors system resources, and triggers renegotiation on detecting a significant change in resource levels (e.g., on a fault, or when new resources become available as in the metacomputing environment). For the purposes of this paper however, we assume that the underlying system is fault-free and has a fixed amount of available resources.

3.2. Junction Detection: A Tunable Application

The junction detection [11] application detects distinguished pixels in an image where the intensity or color changes abruptly. Junction detection is a core component of several image-processing applications, often serving as a precursor to shape construction and classification tasks. Our junction detection algorithm consists of three steps. The first step samples a subset of the pixels in parallel and performs a quick test to determine whether or not the tested pixel is of interest. A pixel is of interest if the difference among intensities/colors of its neighbor pixels is beyond a threshold. The second step draws a region of interest around

a cluster of interesting pixels. The region is essentially a convex hull containing at least a certain number of interesting pixels in close proximity. Finally, the third step runs a compute-intensive algorithm for every pixel in the regions of interest.

Junction detection is a tunable application in that the granularity of sampling in the first step can be parameterized, resulting in different resource requirements. The computation can compensate (with respect to result quality) for a coarser sampling in the first task by possibly drawing additional and/or larger regions of interest. Thus, a smaller resource requirement in the first step (for coarser sampling) is compensated for by requiring a larger resource allocation in the third step. Figure 2 demonstrates this tunability, showing two configurations with different sampling granularities, different thresholds for drawing the regions of interest, and consequently different resource requirements for the third step.

The QoS agent for this application represents tunability in terms of two parameters: the sampling granularity, and a search distance metric which determines how regions of interest are constructed in the second step of the algorithm. The sampling granularity parameter affects the number of pixels sampled in the first step, with coarse sampling being compensated for by a higher value of the search distance (correspondingly, a larger or more regions of interest). The resource requirements, deadlines, and output qualities associated with each alternate execution path are assumed to be available a priori (these can be obtained by profiling on a training set of representative images). The QoS agent communicates this task system to the QoS arbitrator, which chooses the path that will be executed. Note that depending upon system load, different paths may get chosen for junction-detection jobs which arrive at different times. The QoS agent then configures the application to execute along that path. In this case, application configuration just requires setting values for the sampling granularity and search distance parameters.

4. Language Support for Tunability

We extend the Calypso programming language to support expression of tunable, predictable applications. These extensions primarily associate resource requirements and output quality values with individual tasks, specify how tasks may be configured as a function of available resources, and finally specify alternate paths of execution through the program. The Calypso preprocessor uses these extensions to construct a QoS agent for the program which embodies the task graph and tunability aspects of the application. As described in Section 3, during job startup time, this QoS agent negotiates with the QoS arbitrator for an appropriate level of resource allocation.

4.1. Different Models of Tunability

Tunability modifies the behavior of an application in response to available resources. There are two orthogonal dimensions along which such behavior modification can be expressed: granularity of code modification, and the granularity of resource change which can be detected and acted upon by the program. The first dimension refers to how application behavior modification is effected. Application behavior is controlled both at the macro-level by the algorithms it uses (*coarse tunability*), and at the micro-level by certain parameters that affect control-flow decisions (*fine tunability*). The second dimension refers to an application's ability to respond to a change in available resources. Some applications may be able to take advantage of even a very small change in resources (*continuous tunability*), while other applications can only respond to a finite set of specific resource levels (*discrete tunability*).

Applications typically exhibit tunability corresponding to three of the possible four combinations: coarse-discrete, fine-discrete, and fine-continuous. The first two situations are probably most common where an application can use different algorithms (e.g., compression algorithms in a continuous media application) or modify its control parameters in response to a change in resource availability. However, the application can only react to discrete resource levels, not the entire range. An example of the third situation is the sampling step of our junction-detection program: the sampling granularity serves as a knob which can vary application resource requirements over a continuous range.

4.2. Calypso Extensions for Tunability

These extensions to the Calypso programming language identify program control parameters, alternate execution paths, and the corresponding resource and timeliness requirements of program tasks. To keep our preprocessor simple, we focus on supporting only two models of tunability: coarse-discrete (where tunability is achieved by using different algorithms), and fine-discrete (where tunability is effected by modifying parameter values).³ Moreover, we restrict our attention to computations where the sequence of parallel steps is independent of program control flow. Most parallel applications satisfy this assumption.

The language extensions can broadly be classified into three categories: *declaration* constructs (for identifying control parameters), *task* constructs (for associating resource and timeliness requirements with tasks), and *structure* constructs (for specifying alternate execution paths in the code). Control parameters are declared (and option-

³Supporting fine-continuous tunability requires the preprocessor to handle symbolic expressions for resource requirements and deadlines. This is more an implementation rather than a fundamental limitation.

	Parameters	Time		Parameters	Time
step 1	sampleGranularity=16	8	step 1	sampleGranularity=64	2
step 2	searchDistance=24	1	step 2	searchDistance=40	2
step 3		76	step 3		81

Figure 2. A tunable junction-detection algorithm.

ally initialized) within the `task_control_parameters` block, used as in the following code fragment:

```
task_control_parameters {
    ...
};
```

These parameters are used by the QoS agent, after receiving an allocation of resources from the QoS arbitrator, to appropriately configure the program.

The task construct, `task`, acts as a wrapper around a sequential or parallel Calypso step, and specifies the task deadline, its control parameters, and resource requirements and output quality values for each of the acceptable task configurations. The latter correspond to the assignment of specific values to the control parameters. The following code fragment demonstrates usage of the `task` construct:

```
task name [ deadline ][ parameter-list ]
[ ( { param-values } , { resource-request } , quality ) , ... ]
// Calypso code for the task
taskend
```

The *deadline* denotes the time within which this task should complete. The *parameter-list* is a list of control parameters which will be assigned a value exactly before this task starts at the execution time. The acceptable task configurations are shown as a list in the last argument of the `task` construct. Each configuration consists of the *param-values*, which is a list of value assignments to the control parameters identified in *parameter-list*, the *resource-request*, which is a vector of values of size equal to the number of resource types, and the *quality* which describes the quality value of the task output for the current configuration. For the purposes of this paper, *resource-request* is a processor-time tuple, denoting the number of processors required for the task and the time duration they are required for.

Two structuring constructs are provided to enable selection of an execution path through the program. The `task_select` and `task_loop` constructs are used to represent choice of configurations within a parallel step, and overall iterative structure of the program, respectively. The following code fragment demonstrates their use:

```
task_select
when when-expr
    task, task_select, or task_loop constructs
finally finally-code
...
when when-expr
    task, task_select, or task_loop constructs
finally finally-code
task_selectend ;

task_loop ( loop-expr )
    task, task_select, or task_loop constructs
task_loopend ;
```

`task_select` permits selection among multiple tasks within the same step whose readiness is checked using the *when-expr*. The *finally-code* is executed upon completion of the task and together with the `when` construct permits execution paths to be defined in the program. The `task_loop` construct is used to express an iterative structure surrounding the sequential and parallel steps. Its argument, *loop-expr*, denotes the number of iterations. Both *when-expr* and *loop-expr* can only include constants and control parameters, facilitating their evaluation at scheduling time.

The Calypso language extensions are orthogonal to the parallelism constructs, enabling tunability to be incrementally incorporated into an application program.

4.3. Tunability in the Junction Detection Program

We next demonstrate the use of the above language constructs in the context of the junction detection program described in Section 3. As mentioned earlier, the program is tunable with respect to two parameters, controlling the sampling granularity in the first step, and the search distance in the second step. It is assumed that the resource requirements and deadlines for each step are obtained by separately profiling the program. Here we focus on how application flexibility and these predictability requirements can be expressed.

Figure 3 shows the pseudo code for the tunable version of the junction detection program (the original code

```

task_control_parameters {
    int sampleGranularity;
    int searchDistance;
    int c;
};

task_loop( 1000 )
{
    for ( i=0; i<1000; i++ ) {

        task_sampleImage[10.0][sampleGranularity][({16},{4,8.0},1.0), ({64},{4,2.0},1.0)]
        {
            routine for sampling the image
        }
        taskend;

        task_select
        when (sampleGranularity==16) task_markRegionA[60.0][searchDistance][({24},{2,1.0},1.0)]
        {
            algorithm A for finding region of interest
        }
        taskend; finally {c = 0;}
        when (sampleGranularity==64) task_markRegionB[60.0][searchDistance][({40},{2,2.0},1.0)]
        {
            algorithm B for finding region of interest
        }
        taskend; finally {c = 1;}

        task_selectend;

        task_select
        when (c==0) task_computeJunctions[100.0][]({}, {4,76.0},1.0)]
        {
            routine for computing junctions
        }
        taskend; finally {}
        when (c==1) task_computeJunctions[100.0][]({}, {6,81.0},1.0)]
        {
            routine for computing junctions
        }
        taskend; finally {}

        task_selectend;

    }
}
task_loopeend;

```

Figure 3. Expression of application tunability in the Junction Detection program.

fragment is highlighted using boxes). The sampling granularity and search distance parameters are identified as control parameters whose values will be provided by the QoS agent based upon its negotiation with the QoS arbitrator. The rest of the program uses these parameters to control the application behavior. The tunability in the first step, `sampleImage`, is expressed using the `task` construct: the arguments state that the deadline for the step is 10.0 units, the step behavior is controlled by the `sampleGranularity` parameter, and allowable configurations correspond to `sampleGranularity=16` and `sampleGranularity=64`. The first configuration requires 4 processors for 8.0 units of time, while the second requires 4 processors for 2.0 units of time.

The second step, `markRegion`, admits coarse-discrete tunability, allowing use of different algorithms based on the sampling granularity used in the first step. As the code fragment shows, this level of tunability is conveniently expressed using the `task_select` construct. The `finally` code fragments are used to set up another control parameter, `c`, which is used to describe allowable task configurations in the third step. Based on the value of `c` (i.e., the task configuration chosen in the second step), only one of

the `computeJunctions` configurations is allowed. This restriction of configurations based on which configurations were selected in an earlier step make explicit the application's ability to tradeoff resource requirements over its lifetime. A lower allocation of resources in the sampling step require a larger allocation of resources in the junction computation step.

5. Performance Impact of Tunability

To understand the performance impact of tunability, we first propose a greedy heuristic for scheduling tunable parallel applications with predictability requirements, and then using a synthetic task system systematically quantify the benefits and shortcomings of tunability. The QoS arbitrator component of the resource management architecture described in Section 3 will incorporate this heuristic.

5.1. Scheduling Formulation

Without tunability, the underlying scheduling problem we address is one of dynamically scheduling parallel real-time jobs in a system with a fixed amount of homogeneous

processing resources. As described in Sections 3 and 4, we restrict our attention to jobs which can be represented as a chain of tasks. Each task is assumed to be non-preemptible, and either (as in section 5.3) has a fixed resource requirement “shape” (in terms of the number of processors required over a time period) or (as in section 5.4) is malleable (i.e., it can use any number of processors up to its degree of concurrency). The non-malleability assumption is true of most current-day parallel applications written in a message-passing style using systems such as PVM [7] or MPI [9]: these applications can execute on a range of machine sizes, but require that the number of processors assigned to the application remain unchanged over the entire lifetime. Systems such as Calypso support malleable tasks: the programmer specifies only the logical concurrency of the application, which is flexibly mapped to available processors at runtime. We study tunability benefit for non-malleable tasks in section 5.3 and compare these results with that for malleable tasks in section 5.4. In both cases, we assume that information about all tasks of the job is available upon job arrival. The primary objective of this scheduling problem is to maximize the number of on-time jobs. A secondary objective is to maximize system utilization.

With tunability, the only change to the scheduling problem is that a job is now represented by an OR task graph instead of a chain. The multiple paths in the task graph represent the various alternate executions of a tunable program. For uniformity, we assume that all paths through an OR graph have been enumerated, so a tunable application is represented by multiple task chains. For the purposes of this paper, we assume that each chain requires the same total amount of resources and achieves the same output quality value. Note that in practice, task chains of a tunable application are likely to have different overall resource requirements and output qualities: the issue then is of maximizing the achieved job quality.

As with most non-trivial scheduling problems both the above formulations are NP-hard. Consequently, in the next section, we describe a greedy heuristic for the above scheduling problem.

5.2. A Greedy Heuristic

The heuristic greedily allocates resources to jobs using a first fit policy. For a tunable job with multiple schedulable configurations, the heuristic finds among all of them the one that most efficiently uses the system. In more detail, the heuristic keeps track of available *maximal holes* in the processor-time 2D space: each hole is represented by a triple (t_b, t_e, m) (denoting that m processors are available from beginning time t_b until the end time t_e), and is maximal if it is not contained within any other hole. A job is schedulable if all the tasks on its task chain (any

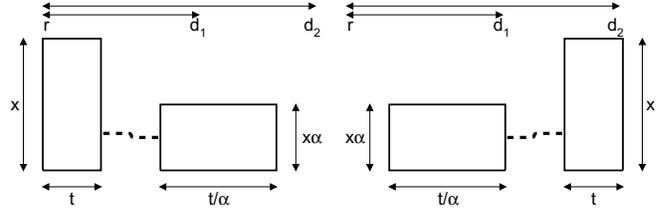


Figure 4. A parameterizable tunable job. The job parameters x , α , and laxity, permit the convenient simulation of a range of job shapes and deadline characteristics.

one of the task chains for a tunable job) can be scheduled into available holes while meeting the task deadlines. Ties between schedulable configurations are broken in favor of chains which maximize system utilization (over a time window defined by the job’s release time and scheduled finish time) and require fewer total resources for some prefix of their tasks. Under the assumptions of our task model, the heuristic finds the job configuration which achieves the earliest finish time.

5.3. Tunability Benefits for Non-malleable Tasks

We use the above heuristic to characterize the performance benefits from tunability. In order to systematically explore the space of application behaviors, we consider a synthetic task system which consists of a parameterizable tunable job shown in Figure 4. The job parameters enable the convenient simulation of a range of job shapes and deadline characteristics.

The parameterizable job consists of two chains, each with two tasks. The two configurations simply transpose the positions of the two tasks. Each task requires the same total amount of resources but with different shapes. One task asks for x processors for time t , whereas the other task requests $x\alpha$ processors for time t/α amount of time. The value of α is chosen in the interval $(0, 1]$ such that both x and $x\alpha$ are integers. Modifying x and α allows the simulation of a variety of task shapes. The task deadlines are set in terms of another parameter, the *laxity* of the job, expressed as the ratio of the slack time in the time period from the release time to the deadline. For a job released at time r , the deadline of the first task is set to $d_1 = r + \max(t, t/\alpha)/(1 - \text{laxity})$; the deadline of the second task is set to $d_2 = r + (t + t/\alpha)/(1 - \text{laxity})$. Since a task can begin execution as soon as its immediate predecessor completes, the task deadline denotes the time by which the task and all its predecessors must finish. The laxity parameter allows the systematic modeling of different amounts of slack time, and hence the constraints that exist for “fitting” a job into the system.

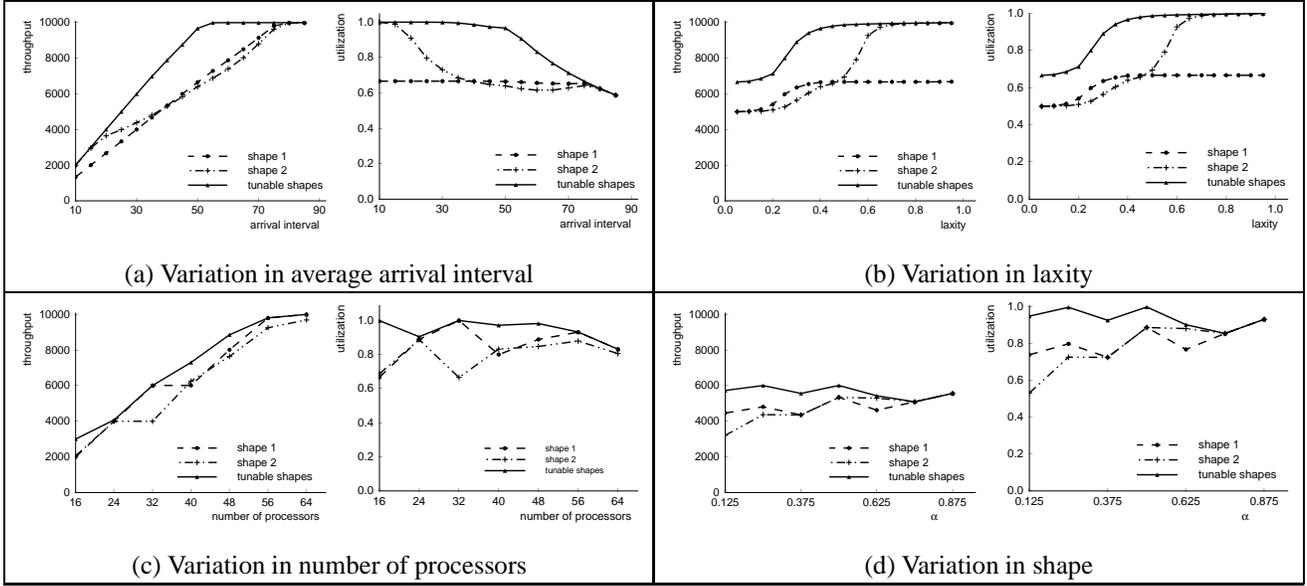


Figure 5. Performance impact of tunability as average arrival interval, laxity, number of processors in the system, and job shapes are varied, in terms of throughput (left) and system utilization (right).

With the above parameterizable job, we construct three task systems: the first task system is tunable, consisting of both job configurations, while the other two systems are non-tunable, containing one configuration apiece. Jobs in each task system are assumed to arrive according to the Poisson distribution. We quantify the benefits of tunability in terms of two metrics—*system utilization* and *job throughput*—measuring the performance of the tunable task system as compared to the non-tunable task systems as a function of four parameters: *mean arrival interval*, *laxity*, the *number of processors* in the system, and α which controls the job shape. Figures 5 show the system utilization (left) and throughput (right) as these parameters are varied one at a time, keeping the others fixed. All experiments reported in this section assume $x = 16$, $t = 25$, and 10,000 job arrivals.

Sensitivity to inter-arrival time In Figure 5 (a), the arrival interval varies from 10 to 85 (note that $t = 25$). When the arrival interval is small, the system is overloaded and only a small portion of the tasks are admitted in all three systems. Tunability has negligible performance impact, since the system is already being 100% utilized. When the arrival interval is very high, the system is underloaded and all three task systems can admit all the jobs. Tunability does not yield much benefit since resources are abundant compared to requests. It is in the middle range of arrival intervals however, that the tunable system achieves the largest improvement in both utilization and throughput: at its peak, it can admit up to 3000 more jobs and achieve up to 30% better system utilization. Even under moderate overload, a non-tunable system is unable to utilize system resources as

efficiently as a tunable system.

Sensitivity to laxity In Figure 5 (b), the laxity varies from 0.05 (a slack time of 5%) to 0.95 (a slack time of 20 times the processing time). when laxity is small, the job deadlines are tight and there is little space left to fulfill resource requirements, the tunable system yields small improvement. This improvement goes up with the laxity. When the laxity is above 60%, shape 2 packs really well and catches up with the tunable system. In contrast, shape 1 requires a larger number of processors for its first task, preventing its packing (due to the greedy nature of the heuristic) even when deadlines are loose. The latter situation demonstrates the performance handicap of an inflexible (i.e., non-tunable) application.

Sensitivity to the number of processors Figure 5 (c) shows the benefits of tunability when the number of processors in the system are increased from 16 to 64 (recall that $x = 16$). In general, when task concurrency is fixed, more processors would seem to imply increased scheduling flexibility for non-tunable systems and consequently reduced benefit from tunability. The graphs show that the tunable system can use the resources in a better way than the non-tunable systems. The non-tunable shapes are not always able to take advantage of more available resources.

Sensitivity to the job shape Figure 5 (d) shows the benefits of tunability as a function of the job shape, determined by α . Tunability improves performance α is not too large (up to 0.625) and the shape of the two tasks is not too close. Tunability has negligible effect when the resource profiles of alternative executions are very similar.

Summary Our simulation results show that tunable task systems yield substantial performance advantages, with respect to both system utilization and throughput. These improvements are most pronounced in mid-portion of each parameter value range: when the system is moderately loaded, when the deadlines are not extremely loose, when the number of processors available is not significantly larger than the concurrency degree of a parallel task, and when the job shapes are not too similar. The above describes the operating point of most practical parallel systems, attesting to the significant performance potential of application tunability.

5.4. Tunability Benefits for Malleable Tasks

Part of the performance benefits for non-malleable tasks result from the restriction that a non-tunable task requires a fixed allocation of processors whereas tunable tasks can take advantage of different numbers of processors. To assess whether tunability also benefits malleable task systems, we slightly modify the previous heuristic. When allocating resources to a malleable task, our heuristic tries various configurations of the task, starting from the highest number of processors the task can use. Figures 6 (a) and Figures 6 (b) compare the benefits from tunability for non-malleable and malleable tasks respectively. The experiments assume the same task system and parameters as earlier.

The two graphs in each of Figures 6 (a) and 6 (b) correspond to the throughput benefits of tunability over the non-tunable jobs—shape 1 and shape 2—as job arrival interval and laxity are varied. We can see, in general, that tunability achieves less benefit for malleable tasks as compared to non-malleable tasks. However, for system configurations that are moderately overloaded and for jobs that have moderate laxity, the tunable task system still yields significant performance improvement over the non-tunable task systems. An explanation for this behavior is that malleability is constrained to a single task whereas tunability goes beyond the task boundary (recall that a job is a task graph) resulting in additional candidate configurations that can take advantage of available resources.

6. Discussion and Related Work

In this paper, we have examined how to efficiently support applications with predictability requirements in a parallel system environment. The ability to efficiently support such applications ensures that general-purpose media-processing applications, which are likely to be one of the largest application domains in the coming years, can benefit from parallel execution. Our solution, centered around the notion of application tunability, captures the flexibility such applications exhibit with respect to resource requirements. Tunability yields benefits both for the system and

the application. While we have demonstrated it in the specific context of predictable parallel computations, the idea of increased flexibility in application specification is relevant in most resource management scenarios; the flexibility can be viewed as increasing an application’s likelihood of obtaining good performance with a fixed set of resources.

Predictability for applications has been well studied in the real-time systems literature. Several well-known scheduling schemes, such as rate-monotonic [12, 17] and earliest-deadline-first (EDF), exist for scheduling real-time tasks in an uniprocessor environment. Gillies [8] has studied the scheduling of tasks in AND/OR graphs. However, the validity of most of these results is restricted to either uniprocessor systems or a sequential task model. Our interest is in managing parallel real-time tasks in a parallel system: in this scenario, well-known results such as EDF no longer provide the optimality guarantees they do in sequential systems.

As described earlier, the primary focus of almost all parallel resource management systems is on optimizing system utilization, often in the context of relatively static resource requirements. The DRMS (Distributed Resource Management System) [15] is an exception, providing system support and resource management for dynamic reconfiguration of parallel programs (with respect to the number of tasks and mapping of tasks to processors). This reconfiguration is closely related to our notion of application tunability; however, unlike the focus of this paper, DRMS exploits this reconfigurability primarily to optimize system utilization and providing only best-effort guarantees to the application.

The Illinois EPIQ project [13, 10] is closely related to our work in that it looks into quality aspects of an application, trading off output quality for resource requirements. The EPIQ view is that the quality of an execution degrades when there are insufficient resources. Our work stresses the tunable aspects of applications, especially the ability to tradeoff resources over time with the same QoS goal. The EPIQ work is also different in focusing on a general framework for specifying and negotiating QoS values. In contrast, we adopt a simpler model of resource versus quality tradeoff to focus on algorithmic issues surrounding the use of these tradeoffs to improve system performance.

7. Conclusion

We have presented a novel approach for improving parallel system utilization while simultaneously meeting the predictability requirements of general-purpose parallel applications such as image recognition, media processing, and virtual reality. Our approach takes advantage of *application tunability*, which refers to a flexibility in the application specification particularly with respect to its ability to trade off resource requirements over time, while maintain-

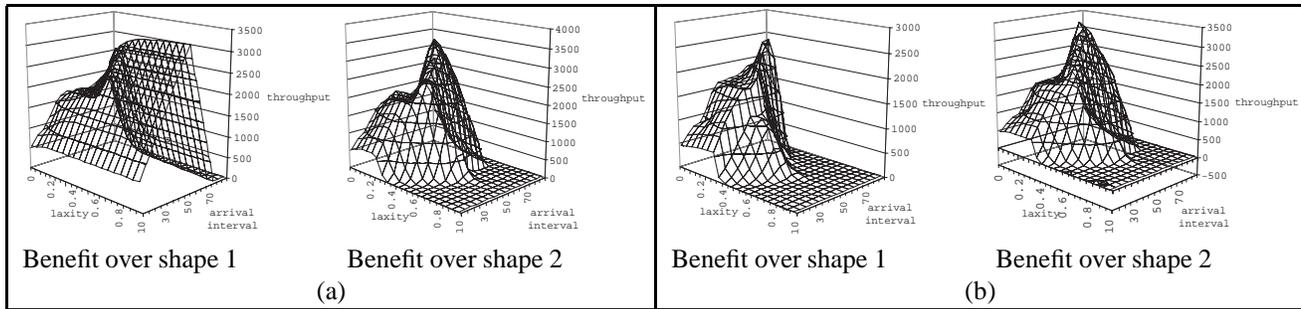


Figure 6. Performance impact of tunability in non-malleable model (a) and in malleable model (b), as job arrival interval and laxity are varied.

ing a desired level of output quality. A resource management system can exploit this flexibility to decide how best to allocate resources to predictable computations so as to maximize both job throughput and overall system utilization. We have described language extensions to the Calypso parallel programming language for expressing tunability, and evaluated its performance impact using a synthetic task system to systematically identify its benefits and shortcomings. Our results show that application tunability is convenient to express, and can yield significant performance benefits.

8. Acknowledgment

This research was sponsored by the Defense Advanced Research Projects Agency and Rome Laboratory, Air Force Materiel Command, USAF, under agreement number F30602-96-1-0320; and by the National Science Foundation under grant number CCR-94-11590. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the author and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Defense Advanced Research Projects Agency, Rome Laboratory, or the U.S. Government.

References

- [1] A. Baratloo, P. Dasgupta, and Z. Kedem. Calypso: A novel software system for fault-tolerant parallel processing on distributed platforms. In *Proc. 4th IEEE International Symposium on High Performance Distributed Computing*, 1995.
- [2] A. Baratloo, M. Karaul, Z. Kedem, and P. Wyckof. Charlotte: Metacomputing on the web. In *Proc. of 9th International Conference on Parallel and Distributed Computing Systems*, Sep. 1996.
- [3] K. Chandy and C. Kesselman. A description of CC++. Technical Report Tech. Rep., CS-92-01, California Institute of Technology, 1992.
- [4] D. Culler et al. Parallel computing on the Berkeley NOW. In *9th Joint Symposium on Parallel Processing (JSP'97)*, 1997.
- [5] P. Dasgupta, Z. Kedem, and M. Rabin. Parallel processing on networks of workstations: Fault-tolerant high performance approach. In *Proc. 15th IEEE International Conference on Distributed Computing Systems*, 1995.
- [6] K. Diefendorff and P. Dubey. How multimedia workloads will change processor design. *IEEE Computer*, 30(9), Sep. 1997.
- [7] G. Geist and V. Sunderam. Network-based concurrent computing on the pvm system. *Concurrency: Practice and experience*, 4(4), 1992.
- [8] D. Gillies and J. Liu. Scheduling tasks with and/or precedence constraints. In *Proc. of 2nd IEEE Conference on Parallel and Distributed Processing*, Dec. 1990.
- [9] W. Gropp, E. Lusk, and A. Skjellum. *Using MPI Portable Parallel Programming with the Message Passing Interface*. MIT Press, 1994. ISBN 0-262-57104-8.
- [10] D. Hull, A. Shankar, K. Nahrstedt, and J. Liu. An end-to-end QoS model and management architecture. In *Proc. of IEEE Workshop on Middleware for Distributed Real-time Systems and Services*, California, Dec. 1997.
- [11] H. Ishikawa and D. Geiger. Segmentation by grouping junctions. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR'98)*, 1998e.
- [12] C. Liu and J. Layland. Scheduling algorithms for multiprogramming in a hard-real-time environment. *Journal of ACM*, 20(1), Jan. 1973.
- [13] J. Liu, K. Nahrstedt, D. Hull, S. Chenand, and B. Li. EPIQ QoS characterization. Draft, July 1997. <http://pertsrserver.cs.uiuc.edu/epiq>.
- [14] The MILAN project: Metacomputing in large asynchronous networks. <http://www.cs.nyu.edu/milan/milan/index.html>.
- [15] J. Moreira and V. Naik. Dynamic resource management on distributed systems using reconfigurable applications. *IBM J. of Research & Development*, 41(3), 1997.
- [16] S. Sardesai and P. Dasgupta. Chime : A versatile distributed parallel processing environment. <http://cs.nyu.edu/milan/milan/index.html>.
- [17] L. Sha and J. Lehoczky. Performance of real-time bus scheduling algorithms. *ACM Perform. Eval. Rev.*, 14(1), 1986.