

Support for Implementing Scheduling Algorithms Using MESSIAHS

STEVE J. CHAPIN¹ AND EUGENE H. SPAFFORD²

¹Department of Math. and Computer Science, Kent State University, Kent, OH 44242-0001

²Department of Computer Sciences, Purdue University, West Lafayette, IN 47907-1398

ABSTRACT

The MESSIAHS project is investigating mechanisms that support task placement in heterogeneous, distributed, autonomous systems. MESSIAHS provides a substrate on which scheduling algorithms can be implemented. These mechanisms were designed to support diverse task placement and load balancing algorithms. As part of this work, we have constructed an interface layer to the underlying mechanisms. This includes the MESSIAHS interface language (MIL) and a library of function calls for constructing distributed schedulers. This article gives an overview of MESSIAHS, describes two sample interface layers in detail, and gives example implementations of well-known algorithms from the literature built using these layers. © 1994 by John Wiley & Sons, Inc.

1 INTRODUCTION

Recent initiatives in high-speed, heterogeneous computing have spurred renewed interest in large-scale distributed systems, and the desire for better utilization of existing resources has contributed to this movement. A typical departmental computing environment already has a substantial investment in computing equipment, including dozens or hundreds of workstations. Studies have shown that the utilization of this equipment can be as low as 30% of capacity [1, 2].

A solution to this problem is to conglomerate the separate processors into a distributed system, and to recursively join the distributed systems into larger systems to further expand the computational power of the whole. Large-scale distributed

systems can have a combined computing power outperforming that of supercomputers [3].

A central element of effective utilization of such systems is task scheduling, which has two components: macro-scheduling (also defined as global scheduling [4] and task assignment [5]) and micro-scheduling (or local scheduling [4]). Macro-scheduling chooses where to run a process whereas micro-scheduling selects which eligible process to execute next on a particular processor. All further uses of the term *scheduling* in this article refer to macro-scheduling.

The MESSIAHS (Mechanisms Effecting Scheduling Support In Autonomous, Heterogeneous Systems) system [6–9] provides a set of mechanisms that facilitate scheduling in distributed, heterogeneous, autonomous systems. For our purposes, distributed, or loosely coupled, systems communicate via message passing rather than a shared memory bus. Heterogeneous systems may have different instruction set architectures, data formats, and attached devices. All policy decisions in autonomous systems are made locally. Our vision of distributed systems includes all three

Received April 1994

Revised May 1994

© 1994 by John Wiley & Sons, Inc.

Scientific Programming, Vol. 3, pp. 325–340 (1994)

CCC 1058-9244/94/040325-16

attributes, connecting machines of different architectures with individual administrative authorities via a communication network. Section 2 gives precise definitions of autonomy and heterogeneity.

It is vital that a system for distributed computation supports autonomy because of the prevailing decentralization of computing resources. There is usually no longer a single, authoritative controlling entity for the computers in a large organization. A scientist may control a few of his own machines and his department may have administrative control over several such sets of machines. That department may be part of a regional site, which is, in turn, part of a national organization. No single entity, from the scientist up to the national organization, has complete control over all the computers it may wish to use. An example of such usage is when two research organizations pool their resources to solve a common problem.

Heterogeneity is important because it yields the most cost-effective and efficient method for performing some computations. A large computation might have certain pieces best suited for execution on a supercomputer whereas other parts might run best on a hypercube or a graphics workstation. If the distributed system is restricted to only using one architecture, the computation will suffer needless delay. In other cases, tasks such as text processing or high-level language interpretation may be independent of any single architecture.

Our subsequent uses of the terms *distributed system* or *system* refer to a distributed, autonomous, heterogeneous system, and *node* refers to an individual machine within an autonomous system. Our definition of system includes a single machine, as well as two homogeneous workstations communicating via a local area network. This definition also encompasses systems as complex as thousands of machines, including personal computers, workstations, file and computation servers, and supercomputers, spread among several remote sites and connected by a wide area network.

Within this distributed system, each individual system has its own policy for deciding when to accept or remove tasks. The local administrator defines this policy, which is implemented over the MESSIAHS mechanisms via an interface layer. The interface layer provides a virtual machine interface; the underlying mechanism can be presented to algorithm writers in various ways. The language described here provides an interface that is easy to use, yet powerful enough to implement a variety of scheduling algorithms. Primitive opera-

tions are supplied to access system and task state information, manipulate tasks, and control the behavior of the local system.

This approach is distinct from that taken in distributed programming systems such as PVM [10] in which the program distribution is visible to, and even forced upon, the programmer. The MESSIAHS approach more closely reflects that taken in Condor [11], which schedules processes invisibly for the programmer. Program distribution is under the control of the autonomous system, and therefore the administrators, rather than the programmer.

Portions of this article discuss implementation issues of the MESSIAHS prototype. The prototype was written in C for SunOS 4.1, and runs on both Sun 3 and Sun SPARC architectures.

2 THE MESSIAHS ARCHITECTURE

MESSIAHS supports task placement in distributed systems with hierarchical structure based on administrative domains, modeled by directed acyclic graphs. Multiple subordinate systems can be combined into an encapsulating system, yielding the hierarchical structure. The nodes of the graph represent the autonomous systems and edges indicate encapsulation. The graph is directed downward; the edges are directed from encapsulating nodes, or parents, to subordinate nodes, or children. Children of the same parent are siblings. (These correspond to the formal definitions of father, son, and brother described by Aho et al. [12].) The neighbors of a system are its children, parents, and siblings.

Figure 1 shows an example distributed system based on the Purdue University computer sciences department. In the example, the computer sciences department contains two administrative domains, Cypress and General. Cypress in turn encapsulates the research machines belonging to the Cypress project and General contains the gen-

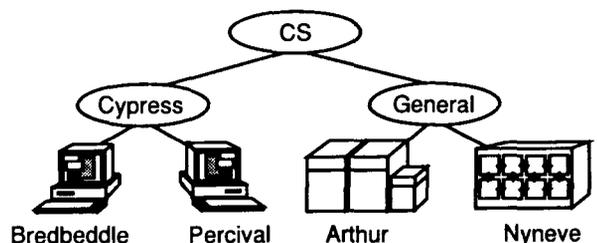


FIGURE 1 A sample distributed system.

eral purpose servers for the department. Bredbeddle and Percival are children of Cypress and therefore are siblings.

Each component system runs a scheduling module that implements the local scheduling policy and manages administrative aspects of the system. These modules exchange data sets describing the status of the systems. On demand, the modules also process scheduling requests, which contain a description of the task for which the sender requests service.

Each module only exchanges status information with modules running on its neighbors. Because of the hierarchical structure of the system, some nodes might be invisible to other nodes. In the example system from Figure 1, Arthur receives information updates only from Nyneve and General and sees no information that can be directly related to Percival or Bredbeddle. The capabilities of Bredbeddle and Percival are subsumed and combined within General's state advertisement.

Individual systems enjoy execution autonomy, communication autonomy, design autonomy, and administrative autonomy [9, 13, 14]. Execution autonomy means that each system decides whether it will honor a request to execute a task; each system also has the right to revoke a task that it had previously accepted. Communication autonomy means that each system decides the content and frequency of state advertisements, and what other messages it sends. A system is not required to advertise all its capabilities nor is it required to respond to messages from other systems. Design autonomy gives the architects of a system freedom to design and construct it without regard to existing systems, yielding heterogeneous systems.

Administrative autonomy means that each system can have its own usage policies and behavioral characteristics, independent of any others. In

particular, a local system can run in a manner counterproductive to a global optimum. In the usual case, scheduling modules will cooperate but administrators must be free to set their local policies or they will not participate in the distributed system. Gantz et al. [1] and Bricker et al. [11] note that users are willing to execute remote jobs on their workstations if the scheduling policy places higher priority on local jobs.

Figure 2 displays the structure of a MESSIAHS scheduling module. The machine-dependent layer handles raw data exchange between scheduling modules, collects the local state information, and interacts with the task manipulation mechanisms specific to the local operating system. The abstract data and task management layer provide an abstract interface for the machine-dependent operations to the data-reporting layer. The shaded layer, data reporting and task manipulation, is the focus of this article. This layer presents the user with the interface to the MESSIAHS mechanisms. The administrator supplies the top-most layer, which embodies the scheduling policy for the system.

MESSIAHS does not determine policy; the three layers provide mechanisms to implement scheduling policies. The interface layer is the administrator's vehicle for expressing and implementing the local policy through the MESSIAHS mechanisms. Sections 4 and 5 describe two interface layers, but we next examine the two lower levels upon which the interface layer is built. This will provide a frame of reference for discussion of the interface layer.

2.1 Machine-Dependent Layer

The machine-dependent layer provides the interface in Table 1 to the management layer of the module. The prototype does not implement those

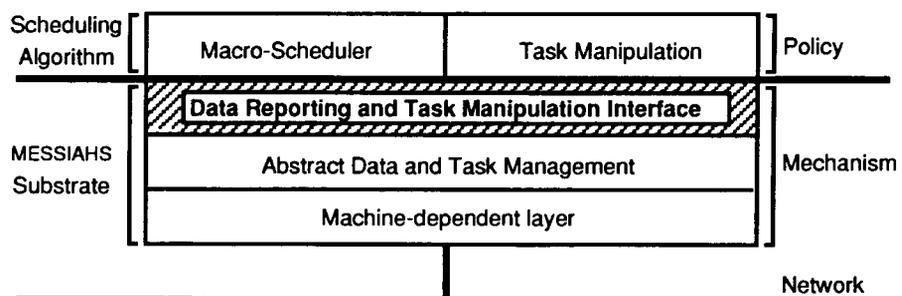


FIGURE 2 Structure of a MESSIAHS scheduling module.

Table 1. Functions in the Machine-Dependent Layer

Purpose	Function Name	Description
Data collection	<code>collect_process_data</code>	Collect data regarding the number of processes and load statistics
	<code>collect_memory_data</code>	Collect data on available and total memory
	<code>collect_disk_data</code>	Collect data on available temporary disk space
	<code>collect_network_data</code>	Collect data on intermodule communication time
Message passing	<code>get_message</code>	Receive a message from the network
	<code>send_message</code>	Send a message over the network
Task management	<code>suspend_task</code>	Pause a running task
	<code>resume_task</code>	Continue executing a suspend task
	<code>kill_task</code>	Halt execution of a task and remove it from the system
	<code>checkpoint_task†</code>	Save the state of a task
	<code>migrate_task†</code>	Checkpoint a task and move it to a target host
	<code>revert_task†</code>	Return a task to its originating system

functions marked with the symbol “†”. As noted earlier, discussion of implementation details pertains to the prototype running on SunOS 4.1.

The functions divide into three main groups: data collection, message passing, and task management. The data collection routines gather information that forms the system description for the local host. The message-passing routines implement abstract message exchange between modules. The task management routines provide access to the underlying operating system process manipulation primitives.

The data collection operations are implemented using the `kvm_open()`, `kvm_read()`, `kvm_nlist()`, and `kvm_close()` routines that access kernel state in SunOS 4.1. The `collect_process_data()` function collects information on the number of processes in the ready queue and the percentage of processor utilization. `collect_memory_data()` determines how much of the physical memory is in use. `collect_disk_data()` finds the amount of public free space on a system, typically in the `/tmp` directory on SunOS. `collect_network_data()` determines the average round-trip time between a host and its neighboring systems within the graph.

An alternative data collection implementation could use the `rstat()` call, which uses the remote procedure call (RPC) mechanisms of SunOS to query a daemon that monitors the kernel state. However, the `rstatd` daemon does not provide information on physical memory statistics or communication time estimates, which are required to

implement the mechanisms. Use of `rstat()` and `rstatd` also involves communication and context switching overhead.

The message-passing routines use the SunOS socket abstraction for communication and the user datagram protocol (UDP) to exchange information between modules. UDP was chosen because it provides an unreliable datagram protocol, which is the minimum level of service required for the update and control channels. The message-passing routines encode the data using the XDR standard for external data representation.

The task manipulation primitives use the SunOS `kill()` system call, which sends a software interrupt, called a signal, to a process. The signals used are `SIGSTOP`, which pauses a process, `SIGCONT`, which resumes a paused process, and `SIGKILL`, which terminates a process. The task migration primitive is not implemented in the prototype, but is a stub procedure for later completion.

2.2 Abstract Data and Communication Management

The middle layer in Figure 2 comprises the abstract data and task manipulation functions. These functions use the basic mechanism provided by the machine-dependent layer to construct higher-level semantic operations. For example, the `send_sr()` routine, which sends a schedule request to a neighbor, is implemented using the `send_message()` function. Table 2

Table 2. Functions in the Abstract Data and Communication Layer

Purpose	Function Name	Description
Data exchange	send_sr	Send a schedule request message
	send_sa	Send a schedule accept message
	send_sd	Send a schedule deny message
	send_trq	Send a task request message
	send_ta	Send a task accept message
	send_td	Send a task deny message
	send_trv	Send a task revoke message
	send_ssq	Send a system status query
	send_ssv	Send a system status vector
	send_tsq	Send a task status query
	send_tsv	Send a task status vector
	send_jr	Send a join request
	send_jd	Send a join deny
	Description	sys_lookup
Vector access	sys_first	Return the first neighbor from the system hash table
	sys_next	Return the next neighbor from the system hash table
	task_lookup	Find the TDV for a task in the task hash table
	task_first	Return the first task from the task hash table
Events	task_next	Return the next task from the task hash table
	register_event	Insert an event into the timeout event queue
	enqueue_event	Enqueue an event
	dequeue_event	Dequeue an event
	new_queue	Allocate an event queue
	qempty	Check if a queue is empty
	set_input_timeout	Enqueue an input timeout
	set_output_timeout	Enqueue an output timeout
	set_recalc_timeout	Enqueue a recalculation timeout
	set_revoke_timeout	Enqueue a revocation timeout
	set_oto_period	Set the output timeout period
set_ito_period	Set the input timeout period	
set_rcto_period	Set the recalculation timeout period	
set_rvto_period	Set the revocation timeout period	

lists the abstract data and task management functions.

The message-passing functions construct a message from the pertinent data and use the `send_message()` function to communicate with a neighboring module. There is one `send` routine for each message type.

MESSIAHS maintains two hash tables containing description vectors: One table containing description vectors of foreign tasks executing on the local host and another table for description vectors of neighboring systems. The hash tables use double hashing as described in Knuth [15, pp. 521–526] for efficiency. The `sys_lookup()` and `task_lookup()` routines search the tables for a particular task or system. The `sys_first()`, `sys_next()`, `task_first()`, and `task_next()` routines iterate over the tables, returning successive description vectors with each call.

The event manipulation routines provide access to the internal event queues used by the module. The `register_event()` function inserts a timed event into the timeout queue and the `enqueue()` and `dequeue()` routines allow direct manipulation of the queues. The `set timeout` routines enqueue timeout events of particular types and the `set period` functions set the timeout periods for the various timers in MESSIAHS. If a timeout period is set to 0, the associated timer is disabled. Input timeouts occur when a neighbor has not sent a status message to the local host within the timeout period. Output timeouts indicate that the local host should advertise its state to its neighbors. Recalculation timeouts cause the local host to recompute its update vectors. When a revocation timeout occurs, the host checks its state to see if tasks should be revoked.

3 SUPPORT FOR SCHEDULING POLICIES

As seen in Figure 2, the scheduling policy is implemented over the interface layer. Through the interface layer, MESSIAHS either directly provides or supports five mechanisms that can be used to construct scheduling policies: system description, decision filters, task revocation, data combination and condensation, and node configuration and behavior customization.

3.1 Intrinsic Mechanisms

MESSIAHS uses a mechanism called description vectors to characterize available resources and requests for resources. A system description vector (SDV) lists the capabilities of an autonomous system and comprises the state advertised between systems. A task description vector (TDV) describes the resources required by a computational job. Description vectors contain a fixed portion that is optimized for task placement support and an extensible portion that administrators can use to implement new scheduling policies or to extend the basic descriptions of requirements or abilities.

To determine the basis for the fixed portion of the description vector, we reviewed 18 algorithms from the existing scheduling literature [2–5, 7, 16–27]. Table 3 depicts the resulting data set. We found that only two characteristics—processor speed and interprocessor communication time estimates—were used by more than four algorithms. Therefore, we included processor speed estimates in the description vector and provided a mechanism to determine intersystem communication time. We also augmented SDVs with other data items that we expect will be useful to writers of future scheduling algorithms. These data support the common case, as represented by the surveyed

algorithms, whereas the extension mechanism allows the inclusion of special purpose data.

The address and module fields uniquely identify a scheduling module: The address specifies a machine and the module indicates the module on that machine. MESSIAHS allows multiple modules to run on a single machine [7]. The nsys field indicates how many systems the vector represents; just as a distributed system encapsulates multiple subordinate systems, the description vector for a system contains information describing its component systems. The ntasks, nactivetasks, and nsuspendedtasks list the number of total tasks, running tasks, and suspended tasks for the system. The willingness gives the rough probability that the system will accept a new task and loadave estimates the computational load on the entire system. The Procclass field is an array of records describing statistical measures of the processor utilization, processor speed, free memory, and disk space.

Execution autonomy mandates the ability to remove a task from execution on the local system. Aborting a running task fulfills the autonomy requirements but does not support load-balancing algorithms based on process migration. Therefore, MESSIAHS includes mechanisms to kill, checkpoint, suspend, resume, and migrate jobs.

In support of administrative and communication autonomy, tunable parameters affect the general behavior of the node. These parameters are independent of any single scheduling policy and effect all policies running on the node. These four parameters are listed in Table 4.

The recalc_timeout and revocation_timeout fields determine how often prescribed events occur. The SPECint92 and SPECfp92 are measures of processor speed using the SPEC benchmark suite [26]. The SPEC benchmark suite consists of application-oriented programs

Table 3. Fixed Portion of a System Description Vector

Field Name	Purpose
address	Address of the system
module	Identifier of module on this system
nsys	Number of systems described by the vector
ntasks	Total number of tasks currently accepted by the system
nactivetasks	Number of active tasks running on the system
nsuspendedtasks	Number of inactive tasks waiting on the system
willingness	Desire of the system to take on new tasks
loadave	An estimate of the load average for the entire system
Procclass	Information on the different classes of processors in the system

Table 4. General State Parameters

Parameter	Purpose
recalc_timeout	Suggested period, in seconds, between recalculations of the local system description
revocation_timeout	Period, in seconds, between checks for possible task revocation
SPECint92	The integer performance rating of the node, per the SPEC integer benchmark
SPECfp92	The floating point performance rating of the node, per the SPEC floating point benchmark

specifically selected to represent real-world workloads.

The machine architecture type (e.g., SPARC or VAX) does not appear as a universal parameter because many jobs are architecture independent. For example, text formatting requests require the presence of a particular text processing package but do not depend on the underlying architecture.

3.2 Supported Mechanisms

MESSIAHS supports the use of filters to implement scheduling policies. Decision filters take two description vectors as input and return an integer value denoting how well the two vectors match according to the local policy. Larger values indicate closer matches. Scheduling modules employ filters to determine where to attempt scheduling a task (including on the local node) and what tasks are eligible for migration or revocation.

MESSIAHS allows multiple scheduling policies to operate within the system simultaneously, and a single node can support two or more scheduling policies. For example, batch queues for text processing, remote compilation, and remote program execution could all coexist within the same distributed system, each with its own individual scheduling policy. The administrator for each node could determine whether that node would participate as a server for any or all of the services.

Communication autonomy requires that the local policy control the flow of information out of a system. This mandates a mechanism to combine and compact the data set and to allow the advertisement of restricted sets of information. In addition, data condensation is essential to avoid arbitrary limits on scaling the mechanisms. If systems concatenated all the data describing subordinate systems, the resources required to transmit and process a description vector would soon outstrip the capabilities of many networks and processors.

Unfortunately, some information loss is unavoidable if data compression takes place. Recall that in our example system, Arthur has no first-hand information about Bredbeddle or Percival. Therefore, Arthur might misdirect scheduling requests to General based on the union of Percival's and Bredbeddle's abilities. For example, if Percival had 100 megabytes of free disk space and 4 megabytes of memory, while Bredbeddle had 10 megabytes of disk space and 32 megabytes of memory, the scheduling module on Arthur might mistakenly think that resources were available to execute a task requiring 16 megabytes of memory and 50 megabytes of disk space. These misdirected requests cause a small efficiency loss, but no tasks will be misscheduled as a result.

4 THE LANGUAGE

The shaded interface layer shown in Figure 2 provides scheduling algorithms with access to lower-level mechanisms. We have chosen to provide two interface layers: a simple programming language, similar to that used in Univers [27], and a library of high-level language functions. This section describes the MESSIAHS interface language (MIL) and the next section describes the library of function calls.

MIL contains direct support for dynamic scheduling algorithms without precluding support for static algorithms. Static algorithms consider only the system topography, not the state, when calculating the mapping. Dynamic algorithms take the current system state as input, therefore the resultant mapping depends on the state [4]. Figure 3 depicts the structure of an MIL program. The grammars for deriving the various rules, along with explanations of their semantics, appear in the rest of this section.

4.1 Expressions and Types

MIL defines four basic types for data values: integers (INT), booleans (BOOL), floats (FLOAT), and strings (STRING). Integers can be written in decimal or in hexadecimal. Booleans have either the value `true` or `false`. Floats are two-decimal digit sequences separated by a decimal point, e.g., 123.45. Strings are a sequence of characters delimited by quotation marks (“”).

Identifiers are a dollar sign followed by either a single word or two words separated by a period. The latter case specifies fields within description vectors. The legal vectors are the received task description (`task`), the description of a task already executing on the system (`loctask`), the system description of a neighboring system (`sys`), the description of the local node (`me`), and the description being constructed by data combination (`out`). `loctask` is used for the task request filter and the revocation filter. `sys` is used for the data combination rules and the schedule request filter. `out` is used only for the data combination rules and `me` can appear in any of the combination rules, filtering, or task revocation sections.

The following grammar defines the expression types used by the language. This grammar only derives expressions of the base types; in particular, there is no access to the Procclass field of the SDV with MIL.

```

int-binop  → + | - | / | * | mod | & | I |
            max | min
int-expr  → int-expr int-binop int-expr |
            (int-expr) | integer |
            int(float-expr) | id

string-expr → string-expr + string-expr |
              (string-expr) | string | id

float-binop → + | - | / | * | max | min
float-expr  → float-expr float-binop float-expr |
              (float-expr) | float |
              float(int-expr) | id
comp        → < | > | = | >= | <= | <>
bool-binop → and | or | xor
bool-expr  → bool-expr bool-binop bool-expr |
              not bool-expr |
              int-expr comp int-expr |
              float-expr comp float-expr |
              string-expr comp string-expr |
              match(string-expr, string-expr) |
              (bool-expr) | true | false | id

```

```

begin state
  <node state rules>
end
begin combining
  <data combination rules>
end
begin schedfilter
  <sched request filter rules>
end
begin taskfilter
  <task request filter rules>
end
begin revokefilter
  <revocation filter rules>
end
begin revokerules
  <revocation rules>
end

```

FIGURE 3 MIL specification template.

4.2 Access to Intrinsic Mechanisms

MIL includes five task manipulation primitives: `kill`, `suspend`, `wake`, `migrate`, and `revert`. Other operations, such as process checkpointing, are available in the lower-level mechanisms but are not explicitly included in the language. `kill` aborts a task, discards any interim results, and frees system resources used by the task. `suspend` temporarily blocks a running task. `wake` resumes a suspended task. `migrate` checkpoints a task

and attempts to schedule it on neighboring systems. `revert` checkpoints the task and returns it to the originating system for rescheduling. Task revocation rules take the following form, using a boolean guard to determine when to take an action.

```
task-action    → kill |
                suspend |
                wake |
                migrate |
                revert
revocation-rule → bool-expr : task-action ;
```

The node state section is a list of types, identifiers, and constant values. Node state declarations are parameters that affect system state. Unlike the extension variables, they do not directly appear in the SDV. The four node state parameters are `specint92`, `specfp92`, `recalc_timeout`, and `revocation_timeout`. The `specint92` and `specfp92` parameters list the speed of the host in terms of the SPEC benchmarks [26]. The `recalc_timeout` and `revocation_timeout` parameters determine the timeout periods for the associated events.

4.3 Filters and Data Combination

In MIL, a filter is a series of guarded statements, similar to combining rules. In place of an action, filters define integer expressions,

```
filter-stmt → bool-expr : int-expr ;
```

A return value of 0 indicates that there is no match. A negative value indicates an error and a positive value measures the affinity of the two vectors. As noted earlier, higher values indicate a better match. If the guard expression uses an undefined variable, the guard evaluates to `false`. If the integer expression references an undefined variable, the filter returns `-1`, indicating an error. With appropriate extension variables and guards, a single scheduling module can serve multiple scheduling policies as stated in Section 3.2.

MIL provides a mechanism to combine description vectors. To support communication autonomy, this mechanism allows the administrator to write rules specifying operations to coalesce the data.

```
int-action    → discard | set int-expr
float-action  → discard | set float-expr
bool-action   → discard | set bool-expr
string-action → discard | set string-expr
```

```
combining-rule → int id bool-expr:
                  int-action ; |
                  float id bool-expr:
                  float-action ; |
                  string id bool-expr:
                  string-action ; |
                  bool id bool-expr:
                  bool-action ;
```

The boolean expression acts as a guard and the action is performed for a particular (type, identifier) pair if the value of the guard is `true`. Administrators may supply multiple rules for the same pair. If multiple rules exist, the module evaluates them in the order written, performing the action corresponding to the first guard that evaluates to `true`.

If no matching rule is found for a pair, the identifier is discarded. Explicit discarding of data items, via the `discard` action, fulfills the constraint of communication autonomy. The `set` value action assigns value to the current pair in the outgoing description vector. An error in evaluating a guard automatically evaluates to `false`. If the evaluation of an action expression causes a run-time error, e.g., a division by 0, the action converts to `discard`.

4.4 Specification Evaluation

The extension and node state rules are interpreted when the specification is first loaded. The data combination rules are applied when a recalculation timeout occurs. When a revocation timeout occurs, the module passes once through the list of revocation rules, repeatedly evaluating each one until its guards return `false`. If the guard evaluates to `true`, the revocation filter is applied to the appropriate list of tasks to provide a target for the revocation action. If no task matches, the module moves on to the next rule in the list.

When a scheduling request arrives, the module iterates over the list of available systems, evaluating the request filter rules in order until a guard that evaluates to `true` is found, or the rules are exhausted. If no matching rule is found, 0 is returned. If a rule is found, its value is returned as the suitability ranking for that system. The module follows a similar procedure for task requests, iterating over the set of available tasks.

4.5 A Small Example

Figure 4 shows a simple MIL specification for a SPARC IPC participating in a distributed L^AT_EX text-processing system. Line 1 in the node state section sets the period for SDV recalculation at 60 seconds. Every minute, each participating system will compute its SDV and forward updates to its neighbors.

The SDV extension variable `hasLaTeX` is true if the system has L^AT_EX available and wishes to act as a formatting server. Clients requesting L^AT_EX processing set the `needsLaTeX` variable to true in their TDV. The combining rule in line 2 sets the outgoing `hasLaTeX` variable if any of the incoming description vectors have it set; the rule on line 3 sets the `hasLaTeX` variable for the local hosts. Hosts providing the L^AT_EX service would use line 3; hosts not providing the service would use line 2 to propagate advertisements by other hosts.

The scheduling filter rule in line 4 compares the available system vectors to the incoming task vector, accepts servers with load averages of less than five, and ranks the systems based on their load average. The guard would fail for a neighbor that had not set the `hasLaTeX` variable, and return false.

5 A LIBRARY OF FUNCTION CALLS

This section describes a library of function calls, called a scheduling toolkit, that provides access to the underlying mechanism. The toolkit provides access to the functions in the low and middle lay-

```

begin state
1.   int    $recalc_period    60;
end
begin combining
2.   bool   $out.hasLaTeX
      $sys.hasLaTeX: set true;
3.   bool   $out.hasLaTeX
      $sys.address == $me.address:
          set true;
end
begin schedfilter
4.   $task.needsLaTeX and $sys.hasLaTeX
      and int($sys.loadave) < 5 :
          6 - int($sys.loadave);
end

```

FIGURE 4 A simple MIL specification.

ers as well as the functions listed in Table 5. The `send_Uvec()`, `send_Dvec()`, and `send_Svec()` functions send update vectors to a system's parents, children, and siblings, respectively.

As shown in Figure 5, statistics vectors (`statvec`) are components of the `procclass` structure that are used to condense the advertised state information for a virtual system. Processors are grouped into process classes on a logarithmic scale, based on their computation speed. The `statvec` fields represent multiple processors using statistical descriptions of their capabilities. Processor speed was chosen as the grouping factor because research of the existing scheduling algorithms indicates that processor speed is the primary consideration for task placement [7, Chapter 2]. The SPEC ratings were chosen as the default speed rating because they are the most widely available benchmark for both integer and floating point performance. Other measures of speed can be included through the extension mechanism.

The `merge_statvec()` function merges two statistics vectors and `merge_procclass()` merges two processor classes into one. The `merge_SDV()` function provides a default mechanism for merging two SDVs into one. The functions in Figure 5 are used to implement MIL, which was described in the previous section.

The programmer uses the toolkit to write a set of event handlers that comprise the scheduling policy. MESSIAHS predefines the set of handlers listed in Table 6, which may be overloaded by the administrator to create a new policy.

6 EXAMPLE ALGORITHMS USING MIL

In addition to the simple L^AT_EX batch processing system described earlier, we present two applications built using MIL. The first demonstrates the task revocation facility as used by a general purpose distributed batch system. The second implements a load-balancing algorithm.

6.1 Distributed Batch

The MITRE distributed batch [1], Condor [11], and Remote Unix [2] systems support general purpose distributed processing for machines running the UNIX operating system. Figure 6 lists a short specification file for a SPARC IPC participating in a distributed batching system. The state rules (lines 1–4) give the speed ratings for an IPC

Table 5. Functions in the MESSIAHS Toolkit

Purpose	Function Name	Description
Data exchange	send_Uvec	Send the <i>U</i> update vector to a parent
	send_Dvec	Send the <i>D</i> update vector to a child
	send_Svec	Send the <i>U</i> update vector to a sibling
Description vector Manipulation	merge_SDV	Merge two SDVs into one
	merge_statvec	Merge two statistics vectors into one
Miscellaneous	merge_procclass	Merge two procclass sets into one
	mk_sid_sb	Return a printable form of the system identification number
	Log	Produce output in the error log
	pLog	Produce output in the error log, including operating system specific error messages

```

struct statvec {
    float min, max, mean, stddev, total;
};

typedef struct statvec Statvec;

struct procclass {
    bit32  nsys;      /* # of machines in this class */
    Statvec qlen;    /* run queue statistics */
    Statvec busy;    /* load on cpu (percentage) */
    Statvec physmem; /* total physical memory */
    Statvec freemem; /* available memory */
    Statvec specint92; /* ratings for SPECint 92 */
    Statvec specfp92; /* ratings for SPECfp 92 */
    Statvec freedisk; /* public disk space (/tmp) stats */
};

typedef struct procclass Procclass;

#define SDV_NPROCCLASS 7
#define SDV_MAXUSERDEF 2048 /* multiple of 2 for cksum */

struct SDV {
    SysId  sid;      /* Autonomous System ID */
    bit32  nsys;     /* number of total systems */
    bit32  ntasks;   /* number of total tasks */
    bit32  nactivetasks; /* number of active tasks */
    bit32  nsuspendedtasks; /* number of suspended tasks */
    float  willingness; /* probability of taking on
                        /* a new task */
    float  global_load; /* global load average */
    Procclass procs[SDV_NPROCCLASS]; /* information on the
                        /* different classes of procs
                        /* in the autonomous system */
    bit32  userdeflen; /* length of user-defined data */
    bit8   userdef[SDV_MAXUSERDEF]; /* user defined data */
};

typedef struct SDV Sdv;

```

FIGURE 5 MESSIAHS data structures.

Table 6. Predefined Event Handlers in MESSIAHS

Function Name	Corresponding Event
handle_msg_sr	Schedule request message
handle_msg_sa	Schedule accept message
handle_msg_sd	Schedule deny message
handle_msg_trq	Task request message
handle_msg_ta	Task accept message
handle_msg_td	Task deny message
handle_msg_trv	Task revoke message
handle_msg_ssq	System status query message
handle_msg_ssv	System status vector message
handle_msg_tsq	Task status query message
handle_msg_tsv	Task status vector message
handle_msg_jr	Join request message
handle_msg_jd	Join deny message
handle_input_timeout	Input timeout
handle_output_timeout	Output timeout
handle_recalc_timeout	Recalculation timeout
handle_revoke_timeout	Revocation timeout

```

begin state
1.   float $SPECint92      13.8;
2.   float $SPECfp92      11.1;
3.   int   $recalc_period  30;
4.   int   $revocation_period 30;
end
begin combining
5.   string $out.proctype not match($out.proctype, "SPARC"):
      set $out.proctype + ":SPARC";
6.   string $out.OSname   not match($out.OSname, "SunOS4.1"):
      set $out.OSname + ":SunOS4.1";
7.   string $out.proctype
      not match($out.proctype, $sys.proctype):
      set $out.proctype + $sys.proctype;
8.   string $out.OSname   not match($out.OSname, $sys.OSname):
      set $out.OSname + $sys.OSname;
end
begin schedfilter
9.   $sys.address == $me.address and
      match($sys.proctype, $task.proctype) and
      match($sys.OSname, $task.OSname):
      max(200 - (100 * int($sys.loadave)), 0);
10.  match($sys.proctype, $task.proctype) and
      match($sys.OSname, $task.OSname):
      max(400 - (100 * int($sys.loadave)), 0);
end
begin revokefilter
11.  true: 1;
end
begin revokerules
12.  $me.loadave > 2.0 and $me.nactivetasks > 2: suspend;
13.  $me.loadave < 1.0 and $me.nsuspendingtasks > 0: wake;
end

```

FIGURE 6 MIL remote execution specification.

and the recalculation and revocation timeout periods. The combining rules in lines 5 and 6 ensure that the processor type variable, `proctype`, contains the string `'': SPARC''` and that the operating system variable `OSname` contains the string `'': SunOS4.1''`. Lines 7 and 8 propagate incoming processor and operating system names.

The example schedule request filter (lines 9 and 10) computes a rating function in the range $[0, 200]$ for the local system and $[0, 400]$ for remote systems. The scheduling request rules ensure that the processor type and operating system match, and assign a priority to a match based on the system load average. Because there is no provision for requesting tasks from a busy system, the section for task request rules is empty.

Hosts participating in the batch system preserve autonomy by varying the parameters of the schedule request filter. For example, tasks submitted by a local user can be given higher priority by basing the rating function on the source address of the task.

The task revocation rules (lines 12 and 13) determine, based on the computational load on the node, whether active tasks should be suspended or whether suspended tasks should be returned to execution. The `true` guard in the revocation filter rule (line 10) matches any available task, and the value portion of the rule assigns an equal priority to all tasks under consideration.

6.2 Load Balancing

Several researchers have investigated load balancing and sharing policies for distributed systems, such as those described previously [28–30].

The *greedy load-sharing algorithm* [28] makes decisions based on a local optimum. When a user submits a task for execution, the receiving system attempts to place the task with a less busy neighbor, according to a weighting function. If no suitable neighbor is found, the task is accepted for local execution.

The suggested weighting function to determine if a task should be placed remotely is $f(n) = n \text{ div } 3$, where n is the number of tasks currently executing on the local system. The algorithm searches for neighbors whose advertised load is less than or equal to one third of the local load. Because the greedy algorithm depends on local state, it is dynamic.

The policy specification in Figure 7 implements a variant of the greedy algorithm. The original al-

```

begin state
1.   int $recalc_period 5;
end
begin combining
2.   int $out.minload ($sys.address == $me.address):
      set min($out.minload, $me.ntasks);
3.   int $out.minload true:
      set min($out.minload, $sys.minload);
end
begin schedfilter
4.   $sys.address == $me.address: 1;
5.   $sys.minload <= ($me.ntasks / 3):
      max(100 - $sys.minload, 2);
end

```

FIGURE 7 MIL specification for greedy load sharing.

gorithm used a limited probing strategy to collect the set of candidates for task reception. The version in Figure 7 sets the recalculation and retransmission periods low (line 1) and depends on the SDV dissemination mechanism to determine the candidate systems.

The combination rules (lines 2 and 3) set the `$minload` field to be the minimum of the load advertised by neighbors and the local load. The filter assigns a low priority to local execution (line 4) and rates the neighboring systems on a scale of 2 through 100 (line 5). Any eligible neighbor takes precedence over local execution, but if the resultant candidate set is empty, the local system executes the task.

The greedy algorithm has no provision for task revocation; any tasks accepted run to completion. Thus, systems using the depicted specification yield some execution autonomy in the spirit of cooperation.

7 EXAMPLE ALGORITHMS USING THE TOOLKIT

As stated earlier, MESSIAHS contains a set of event handlers that may be overloaded by the administrator to create a new policy. For example, the MESSIAHS prototype includes a default handler for schedule request message events. The administrator customizes the scheduling policy by writing a filter routine.

This section presents the implementation of three scheduling algorithms using the toolkit. Figure 8 lists the code for arrival balanced scheduling

```

nt = 0;
for (i = 0; i < SDV_NPROCCLASS; i++) {
    pp = &(psdv->procs[i]);
    if (pp->nsys > 0) {
        float ps, la, d;

        la = pp->qlen.min;
        ps = pp->specint92.mean;
        d = (la + 1) * ptdv->runtime * ptdv->specint92;
        value = (int) (1000 * ps / d);
        if (value > nt) {
            nt = value;
        }
    }
}
return nt;

```

FIGURE 8 Toolkit implementation of the ABS algorithm.

[16], Figure 9 lists the code for the greedy algorithm, and Figure 10 lists the code for the BOS algorithm [31].

The greedy algorithm was described in Section 6. Arrival balanced scheduling assigns a task to the processor that will complete it first, as estimated by the scheduling host. The estimated runtime of the task, the current load on the target host, and the speed of the target host are used to estimate the finishing time of the task. The BOS algorithm employs a simple centralized scheme using a heuristic approach to schedule a task force on a set of homogeneous processors. The algorithm generates an initial mapping, then uses a bounded probabilistic approach to move towards the optimal solution.

The implementations of three algorithms demonstrate that the underlying mechanisms are easy to use and are flexible enough to support a variety of algorithms. The longest of the three algorithms,

```

if (sidmatch(psdv->sid, pmysdv->sid)) {
    return 1;
} else if (psdv->global_load <= (pmysdv->ntasks / 3)) {
    gl = (int) psdv->global_load;
    nt = (int) psdv->ntasks;
    value = ((100 - gl) * 1000) + (999 - nt);
    return(value);
} else {
    return 0;
}

```

FIGURE 9 Toolkit implementation of the greedy algorithm.

```

i = 100000;
if (is_sibling(pste)) {
    return (0);
}

/* add in 'self' and 'target' */
if (pste == pmyste) {
    i -= ((pmysdv->ntasks + 1) * (pmysdv->ntasks + 1));
} else {
    i -= (pmysdv->ntasks * pmysdv->ntasks);
    gl = (pste->sdv.ntasks + 1) * (pste->sdv.ntasks + 1);
    nt = MAX(pste->sdv.nsys, 1);
    i -= (gl / nt);
}

for (pshte = sys_first(); pshte != (Shte *) NULL;
     pshte = sys_next(pshte)) {

    if (pshte->entry != pste) {
        nt = MAX(pshte->entry->sdv.nsys, 1);
        gl = pshte->entry->sdv.ntasks *
            pshte->entry->sdv.ntasks;
        i -= (gl / nt);
    }
}

return i;

```

FIGURE 10 Toolkit implementation of the BOS algorithm.

BOS, represents less than one-half of 1% of the code for the scheduling support module. Writing a new algorithm involves editing a code skeleton and inserting the algorithm code in a C switch statement. This process takes only a few minutes for a programmer familiar with the MESSIAHS code. In contrast, writing a scheduler from scratch, including data collection, data communication, and task management, would take man-months of effort.

This ratio of schedule code size to support code size is consistent with that seen in other distributed scheduling support systems, such as Condor. However, MESSIAHS has ease-of-use advantages because of its separation of mechanism and policy, and because of its support for customizable scheduling policies.

Performance measurements were taken for each of the three algorithms based on simulated tasks [7, Chapter 6]. These results indicate, but do not prove, that the overhead incurred by use of the prototype is minor, typically less than 10% for dynamic algorithms and less than 40% for static algorithms. The 40% slowdown for a static algo-

rithm may be acceptable in some environments because the MESSIAHS version of the algorithm works in an environment that the original static algorithm could not.

In addition, it appears that the MESSIAHS mechanisms perform better as the ratio of inter-task delay to update frequency increases. This increased ratio means that update information travels farther within the distributed system between task arrivals, and thus the scheduling modules are working with more up-to-date information.

8 CONCLUDING REMARKS

The mechanisms provided by the MESSIAHS system, MIL, and the scheduling toolkit support global task scheduling and load sharing in scalable distributed systems. These mechanisms also protect the autonomy of the individual systems while uniting heterogeneous machines into a coherent distributed system.

The language presented here is simple and expressive. It addresses two neglected areas of distributed scheduling, heterogeneity and autonomy. MIL supports a broad range of existing scheduling algorithms while enabling rapid development, prototyping, and analysis of new policies.

Because of its simplicity, MIL is somewhat limited. It cannot store history and has no control flow or looping constructs. Because of this, scheduling algorithms that accept multiple tasks and a set of system descriptions as input cannot be expressed precisely using this language. MIL also assumes that neighbors can be trusted to tell the truth in their SDV advertisements and depends on a model of timely information exchange.

A more complex approach that addresses these limitations, implemented as a set of library calls for high-level languages, is the scheduling toolkit described in Section 5. The toolkit is a more complex interface to the underlying mechanisms than MIL is, but is also more expressive and efficient than MIL. Algorithms developed using MIL can be implemented and refined using the toolkit. Preliminary performance results obtained from the toolkit demonstrated that an overhead of less than 10% is achievable for dynamic scheduling algorithms.

The prototype continues to evolve. The existing task environment is incompletely defined; in particular, the performance results were obtained using simulated tasks. The primary focus of current research on the prototype is to add support for

task migration and execution while still preserving as much autonomy as possible.

In summary, MESSIAHS embodies mechanisms supporting task placement in distributed, heterogeneous, autonomous systems. This support includes extensible mechanisms for implementing the local scheduling policy. This article briefly described the MESSIAHS scheduling support mechanisms, defined a simple language and a library of function calls for constructing schedulers, and gave sample implementations of representative scheduling policies using these tools.

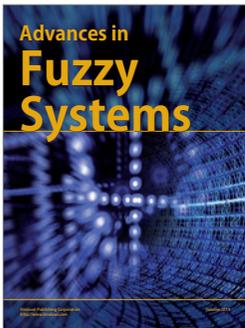
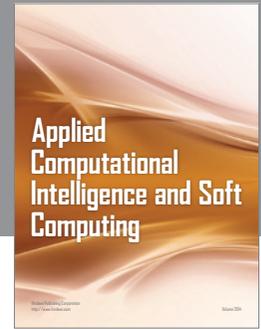
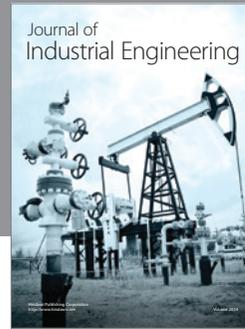
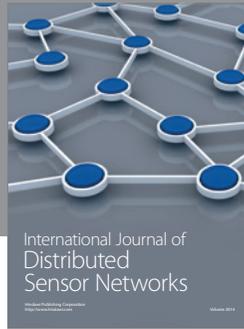
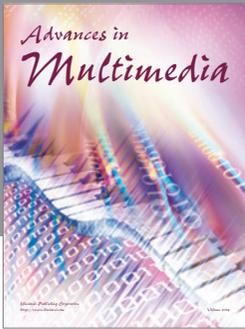
ACKNOWLEDGMENTS

This work was sponsored in part by NASA GSRP grant NGT 50919.

REFERENCES

- [1] C. A. Gantz, R. D. Silverman, and S. J. Stuart, "A distributed batching system for parallel processing," *Software—Practice and Experience*, vol. 19, 1989.
- [2] M. J. Litzkow, *USENIX Summer Conference*. Berkeley, CA: USENIX Association, 1987, pp. 381–384.
- [3] A. H. Karp, K. Miura, and H. Simon, "1992 Gordon Bell prize winners," *IEEE Comp.*, vol. 26, pp. 77–82, 1993.
- [4] T. L. Casavant and J. G. Kuhl, "A taxonomy of scheduling in general-purpose distributed computing systems," *IEEE Trans. Software Eng.*, vol. 14, pp. 141–154, 1988.
- [5] V. M. Lo, "Heuristic algorithms for task assignment in distributed systems," *IEEE Trans. Comput.*, vol. 37, pp. 1384–1397, 1988.
- [6] S. J. Chapin and E. H. Spafford, *27th Hawaii International Conference on Systems Sciences*. 1994, pp. 425–434. IEEE Computer Society Press, Los Alamitos, CA.
- [7] S. J. Chapin, "Scheduling support mechanisms for autonomous, heterogeneous, distributed systems," Ph.D. dissertation, Purdue University, 1993.
- [8] S. J. Chapin and E. H. Spafford, "Scheduling support for an internetwork of heterogeneous, autonomous processors," Technical Report TR-92-006, Department of Computer Sciences, Purdue University, West Lafayette, IN, January 1992.
- [9] S. J. Chapin and E. H. Spafford, "An overview of the MESSIAHS distributed scheduling support system," Technical Report TR-93-011 (supercedes TR-93-004), Department of Computer Sciences, Purdue University, West Lafayette, IN, January 1993.

- [10] A. Geist, A. Beguelin, J. Dongarra, W. Jiang, R. Manchek, and V. Sunderam, *PVM 3.0 User's Guide and Reference Manual*. Oak Ridge, TN: Oak Ridge National Laboratory, February 1993.
- [11] A. Bricker, M. Litzkow, and M. Livny, "Condor technical summary," Technical Report 1069, Department of Computer Science, University of Wisconsin-Madison, January 1992.
- [12] A. V. Aho, J. E. Hopcroft, and J. D. Ullman, *The Design and Analysis of Computer Algorithms*. Reading, MA: Addison-Wesley, 1974.
- [13] F. Eliassen and J. Veijalainen, *TENCON '87*. 1987, pp. 277-281.
- [14] W. Du, A. K. Elmagarmid, Y. Leu, and S. D. Ostermann, *Second International Conference on Data and Knowledge Systems for Manufacturing and Engineering*. New York, NY: IEEE, 1989, pp. 113-120.
- [15] D. E. Knuth, *The Art of Computer Programming, Volume III: Searching and Sorting*. Reading, MA: Addison-Wesley, 1973.
- [16] B. A. Blake, "Assignment of independent tasks to minimize completion time," *Software Practice Exp.*, vol. 22, pp. 723-734, 1992.
- [17] F. Bonomi, "On job assignment for a parallel system of processor sharing queues," *IEEE Trans. Comput.*, vol. 39, pp. 858-869, 1990.
- [18] F. Bonomi and A. Kumar, "Adaptive optimal load balancing in a nonhomogeneous multiserver system with a central job scheduler," *IEEE Trans. Comput.*, vol. 39, pp. 1232-1250, 1990.
- [19] A. Drexler, "Job-processor-scheduling für heterogene computernetzwerke (job-processor scheduling for heterogeneous computer networks)," *Wirtschaftsinformatik*, vol. 31, pp. 345-351, 1990.
- [20] D. Hochbaum and D. Shmoys, "A polynomial approximation scheme for scheduling on uniform processors: using the dual approximation approach," *SIAM J. Comput.*, vol. 17, pp. 539-551, 1988.
- [21] C. C. Hsu, S. D. Wang, and T. S. Kuo, *Proceedings of the 13th Annual International Computer Software and Applications Conference*. 1989.
- [22] C. C. Price and M. A. Salama, "Scheduling of precedence-constrained tasks on multiprocessors," *Comput. J.*, vol. 33, pp. 219-229, 1990.
- [23] V. Sarkar and J. Hennessy, *ACM Conference on Lisp and Functional Programming*. 1986, pp. 202-211.
- [24] V. Sarkar and J. Hennessy, "Compile-time partitioning and scheduling of parallel programs," *SIGPLAN Notices*, vol. 21, pp. 17-26, 1986.
- [25] H. S. Stone, "Multiprocessor scheduling with the aid of network flow algorithms," *IEEE Trans. Software Eng.*, vol. SE-3, pp. 85-93, 1977.
- [26] Standard Performance Evaluation Corporation. *The SPEC Newsletter*.
- [27] M. Bowman, L. L. Peterson, and A. Yeatts. "Univers: an attribute-based name server," *Software Practice Exp.*, vol. 20, pp. 403-424, 1990.
- [28] S. Chowdhury, "The greedy load sharing algorithm," *J. Parallel Distrib. Comput.*, vol. 9, pp. 93-99, 1990.
- [29] D. L. Eager, E. D. Lazowska, and J. Zahorjan, "A comparison of receiver-initiated and sender-initiated dynamic load sharing." Technical Report 85-04-01, University of Washington, Department of Computer Science, April 1985.
- [30] M. F. Pucci, "Design considerations for process migration and automatic load balancing." Technical Report, Bell Communications Research, 1988.
- [31] G. J. Bergmann and J. M. Jagadeesh, *Proceedings of the ISMM International Workshop on Parallel Computing*. 1991, pp. 301-304.



Hindawi

Submit your manuscripts at
<http://www.hindawi.com>

