

A Rigorous Approach to Resource Management in Activity Coordination

Rodion M. Podorozhny, Barbara Staudt Lerner, Leon J. Osterweil

University of Massachusetts, Amherst MA 01003, USA

Abstract. System behaviors can be expected to vary widely depending upon the availability, or shortage, of resources that they require. Thus, the precise specification of resources required by, and resources available to, a system is an important basis for being able to reason about, and optimize, system behavior. Previous resource models for such disciplines as management and workflow have lacked the rigor to support powerful reasoning and optimization. Some resource models for operating systems have been quite rigorous, but have generally been overly narrow in scope. This paper argues that it is important to be able to optimize and reason about the broad and complex resource requirements of modern distributed multi agent systems. This entails precisely modeling a wide range of entity types, including humans, tools, computation platforms, and data. The paper presents a metamodeling approach that can be used to create precise models of a wide range of resource types, and provides examples of the use of this metamodel. The paper also describes a prototype resource allocation and management system that implements these approaches. This prototype is designed to be a separable, orthogonal component of a system for supporting the execution of processes defined as hierarchies of steps, each of which incorporates a specification of resource requirements.

Keywords

Resources, resource specification, process execution, process centered environments

1 Introduction

Much research in such areas as software process, workflow, multi agent scheduling, and computer supported cooperative work focuses on devising mechanisms for adequately specifying coordination of diverse activities to accomplish a complex task. Most of this work incorporates approaches in which the overall task to be accomplished is represented as a synthesis of lower level tasks. In addition, different approaches incorporate, to differing degrees of formality, the specification of the artifacts that these various tasks use and produce. Some approaches also have mechanisms to specify the agents and resources that are to be used to support task execution.

The differences in emphasis on these different components of activity coordination specification are often clearly due to differences in goals. Some specifications are intended to be largely illustrative and advisory, being intended for use in helping humans to come to common understandings. But many specifications are intended to be sufficiently rigorous and detailed that they can be used as prescriptions for computer support and the application of tools. Some of these more rigorous specifications are sufficiently rigorous that they can support powerful reasoning about the activities, thereby allowing them to be used as the basis for strong support of these activities, and for diagnosis and improvement of the activities being represented.

Our own past work has this latter character and goal. We have concentrated on developing and evaluating specification formalisms that are sufficiently rigorous that they can be used to reason about such difficult and complex activities as the processes used to develop software. Our goals have included being able to provide strong tool and automation support to these processes, and also to reason about how to support these processes most efficiently. This past work has demonstrated the need for structuring the tasks that comprise larger overall processes, and the need for being articulate and complete in specifying the artifacts that these tasks use and produce. But our work has also increasingly shown that it is essential to represent the use of resources in descriptions of processes. This need is most acute in the case of processes that are to be carried out through the coordination of multiple agents.

Operating systems research has long ago demonstrated the pivotal importance of reasoning about resources in parallelizing, coordinating, and optimizing the execution of system processes. Clearly, the abundance of resources can enable execution of tasks in parallel, thereby speeding up accomplishment of larger goals. This same phenomenon is also clearly observable in broader classes of activity coordination. If a software design activity requires the use of a particular design tool, then the various members of a design team are able to work in parallel only if there are multiple copies of that tool available. Similarly, if the design activity has been decomposed into separate parallel subtasks, the design process may go faster, but only if more than one qualified designer is available to work on the project.

Correspondingly, the lack of resources causes contention, occasions the need for some tasks to wait for others to complete, and generally slows down accomplishment of larger goals. Often potential delays can be avoided or reduced by using resource analysis to identify ways in which tasks can be parallelized. Thus, for example, if only one design tool is available for use by two designers, delay may be avoided or reduced by scheduling one designer to perform some other task (e.g. requirements revision) while the other designer has use of the one copy of the design tool.

Of course operating systems research has also long ago demonstrated that it is all too easy to create such problems as deadlocks and livelocks if one is not careful in the scheduling of the assignment of resources. These problems are very real for the larger class of activity coordination problems as well. It is not

hard to devise a process in which a requirements analyst is awaiting the results of a prototype activity in order to complete requirements specification, while a prototyper is awaiting the completion of the requirements specification in order to complete the prototype.

In our work, we are interested in creating activity coordination specifications that are sufficiently precise and rigorous that they effectively support reasoning of the sorts just indicated. We would like to be able to identify tasks that are truly parallelizable so that we can schedule them in parallel. We would like to be able to identify where deadlocks, race conditions, and starvation are possible, so that we can assure that they do not occur. We would like to be able to infer when additional resources could potentially speed up execution of the overall activity, and when an activity has an excess of resources, some of which might potentially be reassigned. We would like to be able to manage contention for resources that are in high demand, and those to which access should be managed by disciplines such as transaction mechanisms.

The preceding sorts of reasoning and controls seem to us to be impractical or impossible unless the resources needed to perform tasks are specified. As with most software analyses, these sorts of reasoning can be more powerful and precise when the rigor with which needed resources are specified is more powerful and precise. Thus, in this work we propose a powerful and precise resource specification formalism, and indicate how it can be combined with a suitably powerful and precise process formalism to provide a basis for the kinds of analysis indicated above.

2 Motivation

In order to motivate our work we present an example of how it might be useful in addressing a typical software development problem. To that end, we suggest how the management of resources specified by our modeling formalism can improve the effectiveness of a small software development team. We hypothesize that the team has a variety of human resources, among which are three developers, one of whom is both a qualified coder and designer, one of whom is a qualified coder and tester, and one of whom is qualified to do design work, coding, and testing. Non-human resources include one license for the Rational Rose design support system, and one license for the Visual C++ tool.

Let us further hypothesize that the team is currently in the late stages of development of a product, and is both completing implementation and also carrying out some redesign that is required in order to fix bugs that have been detected in earlier work. Thus, some of the team's activities entail recoding, some requires redesign, and all of it requires retesting. It is not hard to also imagine that the progress of the team is, from time to time, slowed by the lack of availability of resources. In particular there are probably times when it would be desirable to have more qualified designers or coders. More likely, however, there are probably times at which progress could be expedited if the team had available an additional Rose license and an additional Visual C++ license.

We see a number of ways in which the resource management capability that we have developed could be of substantial benefit to this organization. A reasonable scenario might be that the organization has money to purchase only one additional license, and it would be important for it to have some way to determine the best choice, a Rose license or a Visual C++ license. Our resource specification capability is intended to be a facility for supporting informed decisions of precisely this sort. It might be used as the basis for analysis of flow time reduction achievable through one purchase or the other (or both). It might be used to study increases in human resource utilization, or it might be used to suggest ways in which the process itself might be altered to reduce flow time without having to purchase new software licenses at all.

The capability we present here assumes that a process is represented as a set of steps that are connected to each other by control flow and/or dataflow dependencies, and that it supports the careful specification of potential parallel activities. The capability requires that each step has attached to it a specification of the resources that are required in order for the step to be executed. Attaching such resource specifications to the various steps enables a process execution supervisor to assure that steps have what they need before being executed. More important, however, such specifications support determination of how much parallelization can be supported by the available resources, and which infusions of new resources are most likely to be most useful, especially when exceptional cases have arisen. Thus, in our example it may be reasonable to purchase an additional design support tool, either because of the lack of availability of a second designer, or because the set of tasks that are dictated by the process does not require or allow parallel design activities. Precise process and resource specification should help determine such things.

We now introduce a resource modeling and management approach and indicate why it seems to us to be effective in dealing with problems such as these.

3 Overview

Our resource management component is intended to be one component of a larger system that may be used to represent processes, support reasoning about real-time systems, or be integrated into a planning system, for example. Since we believe that the need for powerful and precise resource management is widespread we have designed this component to support the modeling of a wide range of types of resources, from physical entities such as robots, to electronic entities such as programs or data artifacts, and to human entities. The resource management component similarly does not prescribe specific protocols about when resources should be reserved, acquired, and released, but rather leaves the definition of those protocols to the external system with which it is integrated. In this section, we present an overview of our resource management component and introduce some key terminology.

A **resource model** is a model of those entities of an environment that may be required, but for which an unlimited supply cannot be assumed. A

resource model is defined as a collection of **resource classes** and **resource instances**. Resource classes are connected to each other using **IS-A links** to define a resource hierarchy. This allows resources to be described and referenced at various levels of details. Resource instances are themselves connected to the model using IS-A links to identify which resource class(es) they are instances of.

In addition to IS-A links, resources may also be connected by Requires links and Whole–Part links. A **Requires** link indicates that any use of a particular resource also requires use of another particular resource. For example, a piece of licensed software requires use of a computer configured to run that software.

A **Whole–Part** link indicates a relationship in which a resource can be thought of as a single aggregate unit or may be decomposed into smaller pieces that are separately managed. An example of this is a design team that can be assigned a high-level design task, implying responsibility of the team members, while allowing each individual design team member (part) to be assigned other tasks.

There are four primary operations on a resource model:

- **Identification** of resources that can satisfy specific requirements.
- **Reservation** of a resource for a specific start time and duration at some point in the future.
- **Acquisition** of a resource locking the resource for use in a specific activity.
- **Release** of a resource so that it can be used by other activities.

The mechanism described here is quite general and would be used in different ways by different types of systems. A simple process execution environment might rely just on acquisition and release, greedily acquiring resources as necessary. A planning system would reserve resources as the plan is being developed and acquire and release them according to the plan created. Static analysis might examine the resource requests being made by a system to determine whether the necessary resources exist and where delays or even deadlock might arise when resources are contended for. The resource model is also intended to be manipulable by non-programmers. We therefore require both the static definition of the resource model and its dynamic reservation and acquisition status to be visualizable and easily manipulable through a GUI tool.

4 Defining the resource model

A resource model serves two purposes. First, it is a mechanism to describe and organize the resources available within an organization. It is essential that a resource modeling notation be expressive enough to capture the characteristics of resources and relationships among resources that accurately reflect the actual resource entities being modeled. Second, a resource model is used dynamically to control concurrent access to resources, to schedule resources, to analyze the impact of the resource environment on an organization’s ability to carry out an activity. Thus, the dynamic modeling capability of a resource management mechanism is also vitally important. In this section, we present the descriptive

features of the resource modeling mechanism, while in the next we present the dynamic characteristics of our mechanism.

4.1 Resource Entities

A resource model is described as a collection of resource classes, resource instances, and relationships between resources. A **resource instance** represents a unique entity from the physical environment, such as a specific person, printer, or document. A **resource class** represents a set of resources (other classes and/or instances) that have some common properties. Resource classes are further subdivided into unschedulable resource classes and schedulable resource classes. A **schedulable resource class** is one in which its instances are similar enough that the person using the resource model might not care which instance is chosen. For example, **Printer** would probably be a schedulable resource class. An **unschedulable resource class** is more abstract and is intended more as an organizational convenience when defining the resource model. For example, **Person** might be a sensible resource class, but is too general to be scheduled for a specific task.

Each resource is described using a set of typed attribute-value pairs. There is a small set of predefined attributes required of all resources. In addition, the resource modeler can define user-defined attributes that are unique to the specific type of resource. The attribute values of a resource serve to identify the resource and distinguish it from other similar resources. For example, a printer might have an attribute indicating whether it produces color or black output, but that attribute would not be required of resources that are not printers. These descriptive attributes have static values, that is, the values do not change as a result of use of the resources.

There are three pre-defined descriptive attributes required of all resources: a name, a textual description, and a set of criteria assertions. Criteria assertions are membership tests used to ensure that the resources are placed in the model consistent with semantic rules defined by the modeler. Criteria assertions are described further in Section 4.3.

Schedulable resource classes and resource instances also have a capacity attribute associated with them. Capacity is used to indicate the maximum rate at which a resource can be used. The units by which capacity is measured are specific to the type of resource. For example, a printer's capacity would be measured in pages per minute, while a human's capacity would be measured in hours per week. Capacity is used to support resource sharing and is discussed further in Section 4.2.

4.2 Relationships between Resources

The resource modeling mechanism supports the definition of three types of relationships between resources: IS-A relationships, Requires relationships, and Whole-Part relationships. IS-A relationships provide a hierarchical abstraction mechanism. Requires relationships express functional dependencies between the

use of resources. Whole-part relationships allow resources to be manipulated at either coarse or fine granularities and to allow refinement of resource usage over time.

IS-A hierarchy The relationship between a resource class and its members is expressed with an IS-A link. The resource classes and their IS-A links form a singly-rooted tree. The root of the tree is the predefined resource class named **Resource**. A resource instance may belong to multiple classes, thus the entire resource model forms a singly-rooted DAG. Each child of an IS-A link inherits all attributes of its parents. If the same attribute name appears in multiple parents and is originally defined in separate classes (that is, not in a common ancestor), the instance contains both attributes, qualified by the class name. For example, suppose Amelia is an instance of both the **Programmer** class and the **Translator** class. Suppose each of these classes has a **language** attribute. Amelia might have Java and C++ as the value for the **Programmer.language** attribute and English and Japanese as values for the **Translator.language** class. An example of an IS-A hierarchy is presented in Fig. 1.

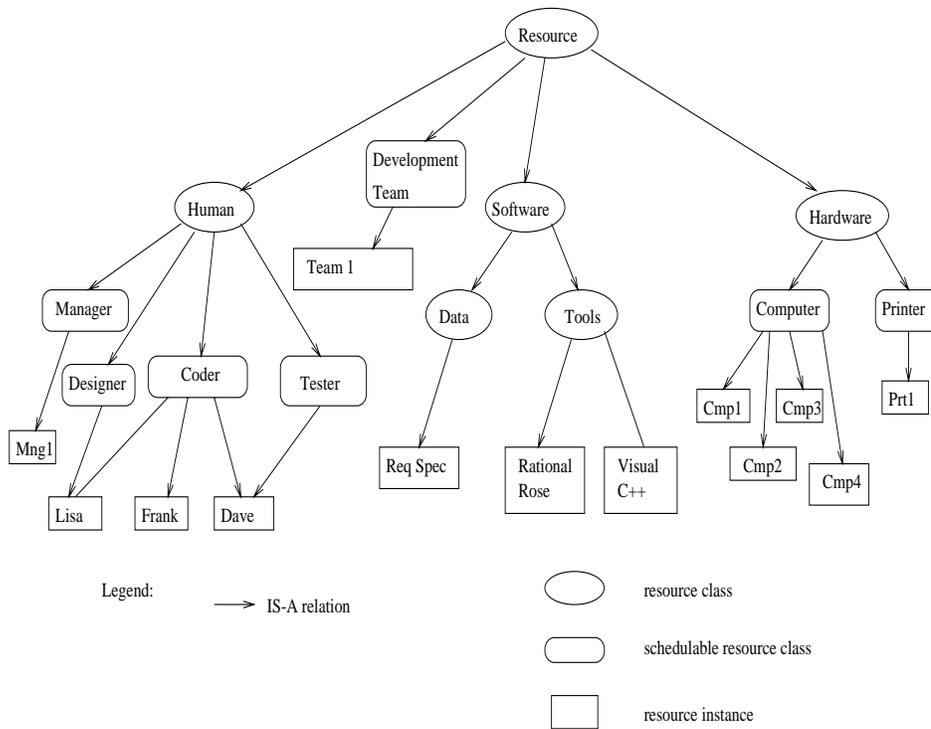


Fig. 1. IS-A hierarchy of an example environment

Requires Functional dependencies between resources are captured with a **Requires** relation. A **Requires** relation denotes that one resource requires another resource in order to perform any work. For example, a piece of software requires a computer with a certain configuration or a delivery person requires a vehicle. The fundamental property captured by the **Requires** relation is that any conceivable use of one resource requires another resource. By capturing these relations in the resource model, the dependency is permanently captured and does not need to be repeated whenever the first resource is used. Arities may be associated with the **Requires** relation to indicate how many instances of a resource are required.

It is also possible to associate a capacity attribute with a **Requires** relation. This is useful in situations in which a resource does not require exclusive use of a second resource. For example, while a software package might require use of a specific computer, it might be possible for multiple people to use the same computer concurrently, if the CPU and memory requirements of that software package are not excessive. When the computer is used for that software package, the capacity available for other activities is reduced by the amount specified in the **Requires** relation.

The **Requires** relation can exist between resource classes, resource instances, or a combination of these. A **Requires** relation between resource instances establishes a tight, static binding between the resources. For example, a particular software package might only be loaded on a specific computer. To use that software package therefore requires use of that specific computer. A **Requires** relation between resource classes establishes the dependency, but defers the binding of which specific required resource is assigned. For example, a delivery person requires a car, but it does not matter which delivery person uses which car in general. A specific task to be performed by a delivery person might place additional requirements on the car. For example, the product to be delivered might need to remain cool, so the delivery car would then require air conditioning. Those additional requirements can be described in the context in which the resources are being used.

An example of the **Requires** relation is presented in Fig. 2. The example shows that a resource instance under schedulable resource class `Coder` requires one resource instance `Visual C++` which, in its in turn requires only one resource instance under schedulable resource class `Computer`.

Whole-Part The semantics of the **Whole-Part** relation is that one resource is logically or physically part of a second resource. The difference between **Requires** and **Whole-Part** is subtle. Both the **Whole** and the **Parts** are modeled as resources to indicate that the resource can be used as a whole, or parts of it can be used individually. Arities and capacities can also be associated with the **Whole-Part** relation. An example of a **Whole-Part** relation is that a team is a **Whole** whose **Parts** are the team members. The team can be given an assignment, implying a joint assignment to the team members. The assignment can be refined later, giving parts of the assignment to the individual members. This capability does

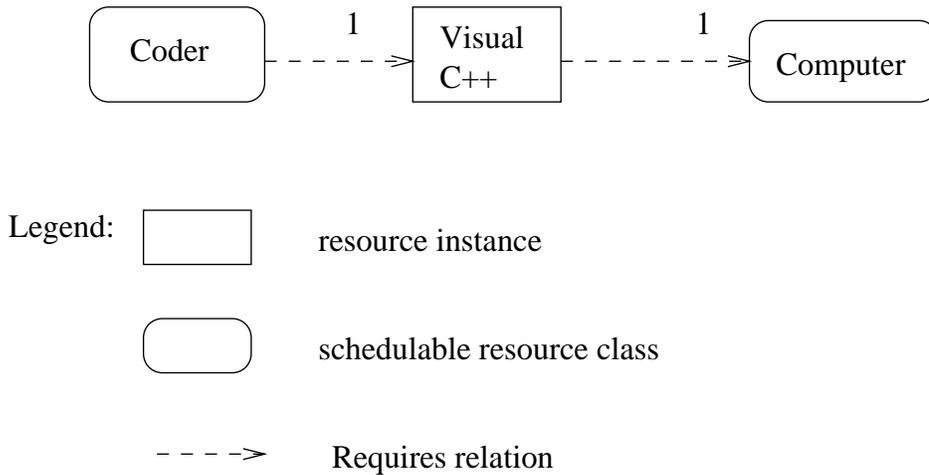


Fig. 2. Requires relation example

not exist for resources connected with the Requires relation.

The Whole-Part relation can be between resource classes or between resource instances. A Whole-Part relation between resource classes specifies a typing relation that must hold for all Whole instances. For example, a Team resource class might have a Whole-Part relation with a Member resource class with an arity of 3–5. Such a relation implies that any Team instance must have a Whole-Part relation with 3–5 Member instances. Thus, when a Whole instance is added to the resource model, it is imperative that the corresponding Whole-Part relation be created to identify the specific team’s members. This differs from the Requires relation in which the binding between resource instances can be deferred until the resources are used, if desired.

An example of Whole-Part relation overlayed on the IS-A hierarchy is presented in Fig. 3. The figure shows that resource instance **Team1** is an aggregate of resource instances **Frank** and **Dave**.

4.3 Criteria assertions

A resource class may define a set of criteria assertions. A criteria assertion is a boolean expression over the attribute values of the resource. These assertions serve as membership criteria for all the resource classes and resource instances that are children of the resource class. The additional assertions introduced at a child resource class may not contradict those of a parent resource class. While this is not checkable in general, such contradictions would make it impossible to populate the resource class with instances. Therefore, the resource classes farther from the root along a certain path of IS-A relations have more membership criteria to satisfy and thus represent more specific classes than those closer to the root.

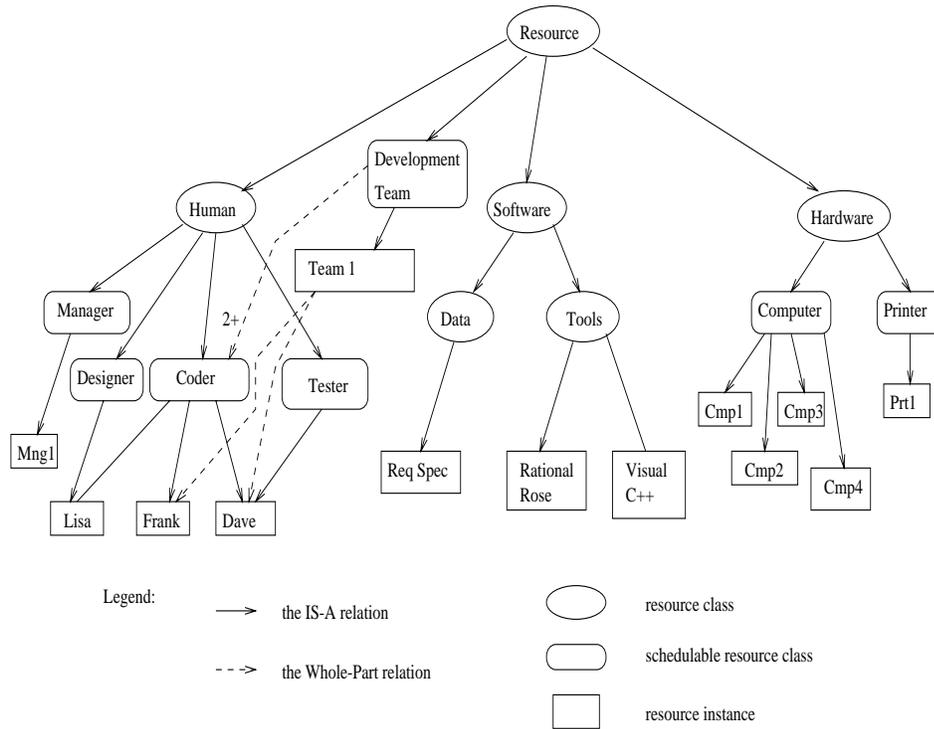


Fig. 3. IS-A and Whole-Part relations

As an example, one might define a resource class `FastPrinter` as a child of the more general `Printer` class. The `FastPrinter` class might have associated with it an assertion that its capacity is greater than 12 pages per minute. Each resource class or instance below it must satisfy this criteria. Failure to do so indicates that the model does not satisfy the semantics that were intended.

The resource manager checks the criteria assertions after any changes to the criteria assertions, to the attribute values of a resource, or the addition of any resources connected (transitively) with IS-A links.

4.4 Modifying the Resource Model

We expect a resource model to be a relatively static entity. Many resources represent physical entities, such as people and equipment. These types of resources change relatively infrequently. As such, we expect these resources to be added to a model by a human resource modeler. We provide a GUI tool to facilitate the creation, reorganization, and visualization of these resource models.

Additionally, we expect some resources to be created, or perhaps destroyed, dynamically during the execution of some activity. For example, a software engineering process produces artifacts, such as design documents, implemented

classes, test cases, etc. Each of these could be modeled as a resource. To facilitate this, we provide an API that allows tools to dynamically create, reorganize, and examine the resource models.

5 Using the resource model

Once a resource model has been defined, it can be used to control resource sharing, to plan future resource usage, or to reason about whether activities could be performed in less time by increasing the number of resources available. In this section, we describe the operations provided by the resource manager that allow such reasoning to be performed. It is important to keep in mind that the activities for which the resources are used are defined and controlled from outside the resource manager. The responsibility of the resource manager is to provide information about the types and availability of resources and to track their usage.

5.1 Identifying Needed Resources

The first step in using a resource model is to determine what resources exist and how well they match the needs of the activity to be performed. To do this, a resource user specifies the name of a resource class and a query over the attribute values contained by resources of that class. The resource manager then returns a collection of resource instances that match that resource class and query. To find matching resources, the resource manager finds all instances connected transitively via IS-A links from the named resource class. The query is applied to each instance to see if it satisfies the necessary conditions and, if so, that instance is added to the set to be returned.

5.2 Acquiring Resources for Use

The resource manager controls the use of resources by keeping track of bindings to activities that are using them. The resource manager has only an abstract notion of the activities that are using them. An activity requests use of a resource by specifying either a specific resource instance by name, or a resource class and query to identify a resource to acquire. If the resource manager can successfully identify a matching resource, the resource is locked for use by the given activity. If more than one resource matches, the resource manager selects one. The acquisition request may also indicate a capacity of the resource that is going to be used if exclusive use of the resource is not required.

If the acquisition request is for a resource that requires another resource, the acquisition only succeeds if all transitively-required resources can also be acquired. Similarly, if the resource is a Whole, the corresponding Parts must also be available for the acquisition to succeed.

Note that in general, the resource manager is not actually able to prevent unauthorized use of the resources as their use is done externally to the resource

manager. For example, a person can be given an assignment without the resource manager being informed. While this cannot be prevented, it does compromise the resource manager's ability to assist in planning and scheduling.

5.3 Reserving Resources for Future Use

Acquisition results in immediate locking of a resource for use. Reservation supports the ability to plan future use of a resource. As with acquisition requests, reservation requests must specify the resource instance to reserve or a resource class and query as well as the capacity required. In addition, a reservation request must identify the time in the future that the reservation is for and the duration of time that it is for. If there is exactly one matching resource instance available at that time, that resource instance is reserved. If there is a schedulable resource class for which all its resource instances satisfy the query, the reservation is made at the level of the schedulable resource instance. In this way, the selection of a specific resource is deferred until acquisition time. This has the potential to increase the overall utilization of the resources since a later request might require a specific resource and by not binding a specific resource instance at reservation time, it is more likely that the more specific request can be satisfied.

The resource manager creates schedulable resource classes dynamically to allow late binding of reservation requests even when the matching resource instances do not wholly comprise an existing schedulable resource class. For example, suppose our resource model had a class named `Printer` and all printer instances in the building were direct children of the `Printer` class. A printer reservation request might use a query that stipulated a specific value for the `location` attribute. If there was a single printer at that location, that instance would be reserved. If there were multiple printers at that location and at least one printer at a different location, a schedulable resource class would be created dynamically to represent the co-located printers and a reservation would be placed on the class.

Reservations automatically reserve required resources and part resources in a manner similar to acquisitions.

In order to use a reserved resource, an acquisition request must be made identifying the reservation that is to be converted into an acquisition.

5.4 Releasing Resources

If a resource is acquired or reserved, it can be made available again for others to use by releasing the resource. A reserved resource is automatically released if the duration of its reservation expires and it has never been converted to an acquisition.

5.5 Refining Usage of Part Resources

In general, a resource cannot be reserved or acquired if there is insufficient capacity available given its current reservations and acquisitions. An exception to this

is in the use of Part resources. The main purpose of the Whole-Part relation is to allow a Whole resource to be reserved/acquired resulting in all its parts being similarly reserved/acquired, but then to allow the initial reservation/acquisition to be refined to a collection of reservations/acquisitions on the Parts. For example, imagine a project management tool that assigns the design of a software system to a team of designers. This ensures that each member of the team is reserved for that task. As the design progresses, the large design assignment will be refined into more specific tasks, each of those given to a smaller team or an individual.

To support refinement, Parts may be reserved/acquired within the context of an existing reservation/acquisition of the Whole. In these cases the reservation/acquisition are made with respect to the capacity already reserved by the Whole reservation/acquisition and not with respect to the generally available capacity of the Parts.

Note that unlike aggregation found in object-oriented design notations, our Whole-Part relation does not denote that the part is entirely owned by the Whole. In this way, an individual might devote 50% of their time to one team and 50% of their time to another team, for example.

6 Related work

Other resource modeling and specification work has been done in such resource sensitive application areas as software process, operating systems, artificial intelligence planning and management. The approaches in these areas have some similarities to our own work, as they concern themselves with such similar problems as the coordination of activities that can span long time periods.

6.1 Related work in software process

There have been a number of software process modeling and programming languages and systems that have addressed the need to model and manage resources. Among the most ambitious and comprehensive have been APEL [10], and MVP-L [15], both of which have attempted to incorporate general resource models and to use resource managers to facilitate process execution. APEL's approach is similar to ours in that APEL deals with resource management as a separate issue that is orthogonal to other process issues. APEL provides a way to specify an organizational model that includes human resources and their aggregates (teams). It also introduces the notion of a position that is very similar to our notion of a class in that it tags human resources according to their skill sets. APEL's roles define the capacity in which a resource is used by a specific activity. APEL's resource modeling approach, however, seems to be less general and comprehensive than our model of resources. In addition it does not seem to incorporate any provision for support of scheduling.

There are a number of other process languages that provide for the explicit modeling of different sorts of resources. Merlin [6], for example, provides rules for

associating tools and roles (or specific users) with a work context (which may be likened to a Little-JIL step). Some others that offer similar limited capabilities are ALF [5], Statemate [9], and ProcessWeaver [1]. In all of these cases, however, the sorts of resources that are modeled are rather limited in scope.

6.2 Related work in operating systems

The problem of scheduling resources has been extensively studied in the field of operating systems (for example, [4], Chapter 6.4). The most common resources in this problem domain include peripheral devices and parts of code or data that require exclusive access. The differences between the needs of resource management in operating systems and software engineering (or artificial intelligence) arise from the fact that operating system resources:

- are used for much shorter periods of time (hence, more elaborate notions of availability are not usually needed).
- are generally far less varied (e.g. Humans are not considered to be resources in operating systems)
- resource usage is much less predictable. Independent programs compete for the same resources and are executed at unpredictable times.

Operating systems research on resource management per se usually focuses on scheduling techniques of a specific resource (such as a CPU or a hard disk as in [7] or [17]). The similarities in purposes of resource modeling between software engineering and operating systems fields appear only in research on admissions control ([3], [2], [16], [13]) which is more in the domain of networks research. For instance, both of these fields require a resource environment abstraction that can model several kinds of resources. They also require a way to satisfy a general resource request (as opposed to the request for a specific instance). This means that a search mechanism is needed in both cases. The resource modeling approaches in the field of networks research are somewhat similar to our approach in that they also often introduce a hierarchy of resources and provide some functionality (i.e., operations for search, reservation and acquisition of resources). It is interesting to note that the authors of [3] and [2] also saw the need to make the resource model independent from the model representing tasks (applications in their terminology). Resource models in this domain, however, seem to lack flexibility and generality.

6.3 Related work in AI planning systems

Probably, the closest resource modeling approach to ours is suggested in the **DITOPS/OZONE** system. **OZONE** is a toolkit for configuring constraint-based scheduling systems [18]. **DITOPS** is an advanced tool for generation, analysis and revision of crisis-action schedules that was developed using the **OZONE** ontology. The closeness is evidenced by the fact that **OZONE** also incorporates

a definition of a resource, contains an extensive predefined set of resource attributes, uses resource hierarchies, offers similar operations on resources, and also resource aggregate querying. We believe that our resource modeling approach places a greater emphasis on human resources in the predefined attributes and allows for an implementation that is easier to adapt to different environments.

The Cypress integrated planning environment is another example of a resource-aware AI planning system. It integrates several separately developed systems (including a proactive planning system (SIPE-2 [20]) and a reactive plan execution system). The ACT formalism [12] used for proactive control specification in the Cypress system has a construct for resource requirements specification. It allows the specification of only a particular resource instance. The resource model does not allow for resource hierarchies and the set of predefined resource attributes is rigid and biased towards the problem domain (transportation tasks).

6.4 Related work in management

An example of a resource modeling approach in a management system is presented in the Toronto Virtual Enterprise (**TOVE**) project [8]. This approach suggests a set of predefined resource properties, a taxonomy based on the properties and a set of predicates that relate the state with the resource required by the activity. The predicates have a rough correspondence to some methods of our resource manager. It is very likely that our resource manager would satisfy the functionality requirements for a resource management system necessitated by the activity ontology suggested in the **TOVE** project.

6.5 Related work in other distributed software systems

The **Jini** distributed software system [19], which is currently being developed by Sun Microsystems, seems to employ a resource modeling approach that seems somewhat similar to ours. The **Jini** system is a distributed system based on the idea of federating groups of users and the resources required by those users. The overall goal of the system is to turn a network into a flexible, easily administered tool on which resources can be found by human and computational clients. One of the end goals of the **Jini** system is to provide users easy access to resources. **Jini** boasts the capability for modeling humans as resources, allows for resource hierarchies, provides ways to query a resource repository using a resource template that is very similar to resource queries in our suggested approach. Because information about **Jini** is limited it is difficult to say what kind of a resource model is used. It is also difficult to see how easily **Jini's** resource model can be adapted to new environments.

7 Evaluation and Future Work

In an earlier paper [14] we describe some of the experience we have had in applying this resource management system. Most of this experience has come from

integrating this system with the Little-JIL [11] process programming system, and using that integrated system to program processes for robot coordination, data mining, and negotiation, as well as software engineering processes such as collaborative design and regression testing. This past experience has confirmed that the features of the system we have developed are of substantial value, and that the approaches we are taking seem appropriate. Our experiences have resulted in the creation and modification of our initial notions and decisions. For example, the **Whole-Part** relation was incorporated into the resource modeling capability after the need became apparent in trying to address the problem of programming the design of software by teams. These experiences have served to reinforce our commitment to further experimentation and evaluation, which we believe will lead to further improvements to this system.

Our experiences have also encouraged us to continue to enhance the features of this system. We are planning to focus more effort on the design and implementation of better languages for the definition of the resource model and the specification of resource requirements. Currently we rely primarily on the use of Java for these specifications, and we seek languages that will be more accessible to non-programmer users.

We would also like to see the creation of a simulation system that would use the process and resource specification systems we have developed to draw inferences about projected resource needs and utilization. A simulation system of this sort would enable users to gauge the gains and losses that would be likely to result from making changes in resource availability, or from making changes in the process itself.

We would also like to see the resource management system enhanced with a capability for modeling the skilling-up of humans as a result of their execution of certain process steps. Our current system assumes that agents have fixed skill levels. But clearly humans will learn as they go through a lengthy process. As a consequence their skill levels should go up. This is clearly a desirable modeling feature and should be incorporated in a future version of our system.

As a more distant goal we see the addition of a best match resource querying mechanism. Ideally, it should be able to find a closest match to a resource query based on the user specified criteria. Currently, the query is satisfied by returning all the matching resources. For such a mechanism to operate we also need to devise metrics that would allow to measure “distances” between resources (to be exact, their attribute values) and resource queries.

Finally, we need to gain more experience in the use of the resource manager in support of other systems. To this end, we intend to continue the development of processes that use resources, to evaluate how well the resource management system addresses the needs of the process execution environment and specific processes, and to determine how well the resource management system can facilitate planning and analysis activities.

8 Acknowledgments

This research was partially supported by the Air Force Research Laboratory/IFTD and the Defense Advanced Research Projects Agency under Contract F30602-97-2-0032. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements either expressed or implied, of the Defense Advanced Research Projects Agency, the Air Force Research Laboratory/IFTD, or the U.S. Government.

References

1. C.Fernström. PROCESS WEAVER: Adding process support to UNIX. In *The Second International Conference on the Software Process*, pages 12–26, 1993.
2. S. Chatterjee and J. Strosnider. A Generalized Admissions Control Strategy for Heterogeneous, Distributed Multi-media Systems. In *Proceedings of ACM Multimedia 95*, San Francisco, CA, November 1995.
3. S. Chatterjee, J. Sydir, B. Sabata, and T. Lawrence. Modeling Applications for Adaptive QoS-based Resource Management. In *Proceedings of the 2nd IEEE High-Assurance System Engineering Workshop*, Bethesda, Maryland, August 1997.
4. H. M. Deitel. *An Introduction to Operating Systems*. Addison-Wesley, Reading, Massachusetts, 1984.
5. G.Canals, N.Boudjlida, J.-C.Derniame, C.Godart, and J.Lonchamp. ALF: A framework for building process-centered software engineering environments. In A. Finkelstein, J. Kramer, and B. Nuseibeh, editors, *Software Process Modelling and Technology*, pages 153–185. Research Studies Press, Ltd., Taunton, Somerset, England, 1994.
6. G.Junkermann, B.Peuschel, W.Schäfer, and S.Wolf. MERLIN: Supporting co-operation in software development through a knowledge-based environment. In A. Finkelstein, J. Kramer, and B. Nuseibeh, editors, *Software Process Modelling and Technology*, pages 103–129. Research Studies Press, Ltd., Taunton, Somerset, England, 1994.
7. P. Goyal, X. Guo, and H. Vin. A Hierarchical CPU Scheduler for Multimedia Operating Systems. In *Proceedings of the Second Symposium on Operating Systems Design and Implementations (OSDI'96)*, pages 107–122, Seattle, Washington, October 1996.
8. M. Gruninger and M. S. Fox. An Activity Ontology for Enterprise Modelling. Submitted to AAAI-94, Dept. of Industrial Engineering, University of Toronto, 1994.
9. D. Harel, H. Lachover, A. Naamad, A. Pnueli, M.Politi, R. Sherman, A. Shtull-Trauring, and M. Trakhtenbrot. STATEMATE: A working environment for the development of complex reactive systems. *IEEE Trans. on Software Engineering*, 16(4):403–414, April 1990.
10. J.Estublier, S.Dami, and A.Amiour. APEL: A graphical yet executable formalism for process modelling. In *Automated Software Engineering*, March 1997.
11. B. S. Lerner, L. J. Osterweil, J. Stanley M. Sutton, and A. Wise. Programming process coordination in Little-JIL. In V. Gruhn, editor, *Proceedings of the 6th European Workshop on Software Process Technology (EWSPT '98)*, number 1487

- in Lecture Notes in Computer Science, pages 127–131, Weybridge, UK, September 1998. Springer-Verlag.
12. K. L. Myers and D. E. Wilkins. The Act Formalism. Working document: Version 2.2, SRI International, Artificial Intelligence Center, September 25 1997. <http://www.ai.sri.com/act/act-spec.ps>.
 13. K. Nahrstedt and R. Steinmetz. Resource Management in Networked Multimedia Systems. *IEEE Computer*, pages 52–64, May 1995.
 14. R. M. Podorozhny, B. S. Lerner, and L. J. Osterweil. Modeling Resources for Activity Coordination and Scheduling. In *Proceedings of the Third International Conference on Coordination Models and Languages*. Springer-Verlag, 1999.
 15. H. Rombach and M. Verlage. How to assess a software process modeling formalism from a project member's point of view. In *The Second International Conference on the Software Process*, pages 147–159, 1993.
 16. S. Floyd and V. Jacobsen. Link Sharing and Resource Management Models for Packet Networks. *IEEE Transactions on Networking*, 3(4):365–386, 1995.
 17. P. Shenoy and H. Vin. Cello: A Disk Scheduling Framework for Next-generation Operating Systems. In *Proceedings of ACM SIGMETRICS'98*, June 1998.
 18. S. F. Smith and M. A. Becker. An Ontology for Constructing Scheduling Systems. In *Working Notes from 1997 AAAI Spring Symposium on Ontological Engineering*, Stanford, CA, March 1997.
 19. J. Waldo. *Jini Architecture Overview*. Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, CA 94303, 1998. <http://www.javasoft.com/products/jini/index.html>.
 20. D. E. Wilkins. *Using the SIPE-2 Planning System: A Manual for Version 4-17*. SRI International Artificial Intelligence Center, Menlo Park, CA, October 1997.