A Formal Model for Reasoning about Adaptive QoS-Enabled Middleware

Nalini Venkatasubramanian* Carolyn Talcott† Gul Agha‡

Abstract

Systems that provide QoS-enabled services such as multimedia are subject to constant evolution - customizable middleware is required to effectively manage this change. Middleware services for resource management such as scheduling, protocols providing security and reliability, load balancing and stream synchronization, execute concurrently with each other and with application activities and can therefore potentially interfere with each other. To ensure cost-effective QoS in distributed systems, safe composability of resource management services is essential. In this paper we present a meta-architectural framework for customizable QoS-based middleware based on the actor model of concurrent active objects. Using TLAM, a semantic model for specifying and reasoning about components of open distributed systems, we show how a QoS brokerage service can be used to coordinate multimedia resource management services in a safe, flexible and efficient manner. In particular, we show that a system in which the multimedia actor behaviors satisfy the specified requirements, provides the required multimedia service. The behavior specification leaves open the possibility of a variety of algorithms for resource management as well as adding additional resource management activities by providing constraints to ensure their non-interference. *Keywords: meta-object models, distributed systems, theoretical foundations, object-oriented applications, multimedia*

1 Introduction

In the coming years, QoS-enabled distributed servers will be deployed to deliver a variety of interactive services. Applications such as telemedicine, distance learning and electronic commerce exhibit varying requirements such as timeliness, security, reliability and availability. The set of servers, clients, user requirements, network and system conditions, in a wide area infrastructure are changing continuously. Future applications will require dynamic invocation and revocation of services distributed in the network without violating QoS constraints of ongoing applications. To assure safe adaptation to dynamically changing requirements it is important to have a rigorous semantic model of the system: the resources, middleware that provides system management, the application activities, and the sharing and interactions among these. Using such a model, designs can be analyzed to clarify assumptions that must be met for correct operation, and to establish criteria for non-interference. In [34, 31], we presented the TLAM(Two Level Actor Machine) semantic framework for specifying, composing and reasoning about resource management services in open distributed systems. The TLAM is based on the actor model of computation, a natural model for dynamic object-based distributed systems. In the TLAM, a system is composed of two kinds of actors (active objects), base actors and meta actors, distributed over a network of processing nodes. *Base-actors* carry out application level computation, while *meta-actors* are part of the runtime system which manages system resources and controls the runtime behavior of the base level. The two

*Contact Author: Nalini Venkatasubramanian Dept. of Information and Computer Science, Univ. of California Irvine, Irvine, CA 92697-3425, Email: nalini@ics.uci.edu

†Department of Computer Science, Stanford University, Stanford, CA 94305, Email: clt@cs.stanford.edu

‡Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL 61801, Email: agha@cs.uiuc.edu

levels provide a clean separation of concerns and a natural basis for modeling and reasoning about customizable middleware and its integration with application activity. Based on the two-level architecture, a customizable and safe distributed systems middleware infrastructure, called CompOSE|Q (Composable Open Software Environment with QoS) [32] is being developed at the University of California, Irvine, that has the ability to provide cost-effective and safe QoS-based distributed resource management.

In this paper, we use the TLAM framework to develop a formal model for customizable, cost-effective middleware to enforce QoS requirements in multimedia applications. We begin by informally describing the notion of a system providing *QoS-based MM Service*.We then map QoS requirements to resource requirements and focus on modeling and reasoning about the resource management underlying a QoS-based service. For this purpose we define, in a rigorous manner, the notions of a system providing *Resource-based MM Service*, of a system having *Resource-based MM Behavior*, and finally refining the system with an *Adaptive Request Scheduling Policy*. The Resource-based MM Service specification reflects the chosen system resource architecture and allows us to reason about the availability and use of resources. The Resource-based MM Behavior specification models the QoS broker software architecture presented in [31, 33] and places constraints on the actions of the QoS meta actors. Such a behavior specification can serve as a first stage in refining a service specification into an implementation. The *Adaptive Request Scheduling Policy* illustrates such refinement. It specifies one of the resource management policies developed in [31] by giving high-level algorithms for determining meta actor behavior. The main results are:

(1) if a system provides Resource-based MM Service, then (under the assumptions on the mapping from QoS requirements to Resource requirements) it provides QoS-based MM Service;
(2) if a system has Resource-based MM Behavior, and meets certain initialization and non-interference conditions, it provides Resource-based MM Service;
(3) if a system is refined with the Adaptive Request Scheduling Policy, then it implements Resource-based MM Behavior.

A consequence of (2) is that new broker policies can be safely installed as long as they satisfy the behavior constraints. (3) is an example of such a policy.

The rest of this paper is organized as follows. Section 2 reviews the TLAM framework concepts needed to understand the formal model of the QoS Broker. Section 3 recalls the QoS Broker multimedia server meta-architecture of [31]: its physical and software architectures, and a mapping of the software architecture onto the physical architecture. Section 4 describes several resource management policies that we have used in realizing the software architecture. Section 5 shows how we use the TLAM framework to formally model and reason about systems such as QoS Broker. Section 6 discusses related work and outlines areas for future research.

2 The Two Level MetaArchitecture

In this section we briefly summarize the main concepts of the TLAM semantic framework. More detail can be found in [34, 31]. The Actor model is a model of concurrent active objects that has a built-in notion of encapsulation and interaction and is thus well-suited to represent evolution and co-ordination among interacting components in distributed multimedia applications. Traditional passive objects encapsulate state and a set of procedures that manipulate the state; actors extend this by encapsulating a thread of control as well. Each actor potentially executes in parallel with other actors and interacts only by sending and receiving messages. (See [1, 2] for more discussion of the actor model, and for many examples of programming with actors.)

As mentioned earlier, in a TLAM model, a system is composed of two kinds of actors, base actors and meta actors, distributed over a network of processing nodes. *Base-actors* carry out application level computation,

while *meta-actors* are part of the runtime system which manages system resources and controls the runtime behavior of the base level. A TLAM model provides an abstract characterization of actor identity, state, messages, and computation, and of the connection between base- and meta-level computation. Meta-actors communicate with each other via message passing as do base-actors, and meta-actors may also examine and modify the state of the base actors located on the same node. Base-level actors and messages have associated runtime annotations that can be set and read by meta actors, but are invisible to base-level computation. Actions which result in a change of base-level state are called events. Meta-actors may react to events occurring on their node.

A TLAM is a structure of the form

$$TLAM = <Net, TLAS, loc>$$

where Net is the underlying computer network with processor nodes and communication links; and $TLAS$ is a two-level actor system distributed over the network by the map, loc. The semantics of a TLAM model is given by a labeled transition relation on configurations. A TLAM *configuration*, C, has sets of base- and meta-level actors, and a set of undelivered messages. The actors are distributed over the TLAM nodes. Each actor has a unique name (address) and the configuration associates a current state $getS(C, a)$ to each actor name, a. $Cast(C)$ is the set of names of actors in C. $getA(C, a, t)$ is the value, in C, of the annotation of base-actor a with tag t; $setA(C, a, t, v)$ sets the value of the annotation of a with tag t, returning the updated configuration. Thus $getA(setA(C, a, t, v), a, t) = v$. The undelivered messages are distributed over the network – some are traveling along communication links and others are held in node buffers. A transition is either a communication transition, an execution transition, or the idle transition. Communication transitions move undelivered messages around the network and are the same in every TLAM. An execution transition is a computation step taken by a base or meta level actor. The possible execution transitions are determined by *reaction rules* specific to a particular TLAM. These rules determine how individual actors react to messages received, and in the case of meta actors how they react to base-level events. A meta-actor reaction rule has the general form:

$$\langle a : s\rangle[, a \triangleleft M] \xrightarrow[effect]{trigger} \langle a : s'\rangle \ , \ C_n \qquad \text{if } cond$$

where $\langle a : s\rangle$ is an actor with name a in state s and $a \triangleleft M$ is a message to a with content M. ([...] indicates that the message part may be empty.) s' is the new state of actor a and C_n is a configuration describing newly created meta-level actors and messages. $trigger$ if not empty, is a base level event such as message delivery or base-level execution step. $effect$ if not empty, represents a modification of the base level configuration on the node of the meta-actor (for example modifying annotations of existing base-level actors or creating new base-level actors and messages). $cond$ is a predicate that may depend on a, s, M, $trigger$, and/or $effect$.

A *computation path* for a configuration C_0 is an infinite sequence of labeled transitions:

$$\pi = [\ C_i \xrightarrow{l_i} C_{i+1} \ \Big|\ i \in \mathbf{Nat}\ ]$$

where l_i is the transition, indicating the transition rule applied. The semantics of a configuration is the set of fair computation paths starting with that configuration, where fairness means that any enabled communication transition will eventually happen, and enabled reaction rules will either fire or become permanently disabled.

A *system* is a set of configurations closed under the transition relation. We say that π is a computation path of a system S if it is a computation path of C for some configuration C of S. Properties of a system modeled in the TLAM are specified as properties of computation paths. A property can be a simple invariant that must hold for all configurations of a path, a requirement that a configuration satisfying some condition eventually arise, or a requirement involving the transitions themselves.

Our general approach to modeling systems using a TLAM framework is to specify a service at different levels beginning with the abstract notion of end-to-end service provided by a system in response to a request.

One form of refinement is to express system wide properties in terms of of abstract properties of the underlying network. Another level of specification is the notion of a group of actors providing a service. This moves from specification in terms of global system behavior to specification in terms of constraints on the behaviors of individual actors or collections of actors. The behavioral level can be further refined by specifying protocols and algorithms for the actions of individual actors. At each stage we show that a system satisfying a refinement of a specification also satisfies the refined specification. Properties are checked using the properties of the building blocks for configurations – message contents and actor state descriptions – and of the TLAM reaction rules that determine the behavior of actors in the system. The staging and refinement of specifications provides a form of modularity, scalability, and reusability by reducing the task of implementation to that of implementing individual abstract behaviors.

3 The QoS Broker Meta-Architecture for QoS-based Services

Using the TLAM framework, we develop a meta-architectural model of a multimedia server that provides QoS based services to applications. The physical architecture of the MM server consists of:

- a set of data sources (DS) that provide high bandwidth streaming MM access to multiple clients. Each independent data source includes high capacity storage devices (e.g. hard-disks), a processor, buffer memory, and high-speed network interfaces for real-time multimedia retrieval and transmission;
- a specific node designated as the distribution controller (DC) that coordinates the execution of requests on the data sources; and
- a tertiary storage server that contains the MM objects—replicas of these MM objects are placed on the data source nodes.

All the above components are interconnected via an external distribution network that also transports multimedia information to computers and set-top devices at the client end. A lower speed back-channel conveys subscriber commands back to the data sources via the DC.

The software architecture of a multimedia server consists of two subsystems - the base level and meta-level subsystems corresponding to the application and system level resource management components respectively. The base level component implements the functionality of the MM application and models both MM data objects and their replicas (e.g. video and audio files), and MM requests to access this data via *sessions*. The corresponding base-level entities are *replica actors* and *request actors*. The meta-level component deals with the coordination of multiple requests and sharing of existing resources among multiple requests. To provide coordination at the highest level and perform admission control for new incoming sessions, we introduce a meta-level entity, the *QoS Broker* meta-actor, QB.

The two main functions of the QoS Broker are data management and request management. The multimedia data and request management functions of the QoS broker in turn require a number of services. We describe a representative sample below. The organization of meta-level services provided by a QoS broker is shown in Figure 1. Each of these services in turn can be based on one or more of the core services integrated into the metaarchitecture - remote creation, distributed snapshot and directory services.

We now describe some of the main meta-level services and the supporting services in the above QoS broker system. The *data management* component decides the placement of data in the distributed system, i.e., it decides when and where to create additional replicas of data. It also determines when additional replicas of data actors are no longer needed and can be garbage collected/dereplicated. The *request management* component performs the task of admission control for incoming requests and ensures the satisfaction of QoS constraints for requests that are ongoing in the system. We have developed adaptive admission control mechanisms [33] in the *request*

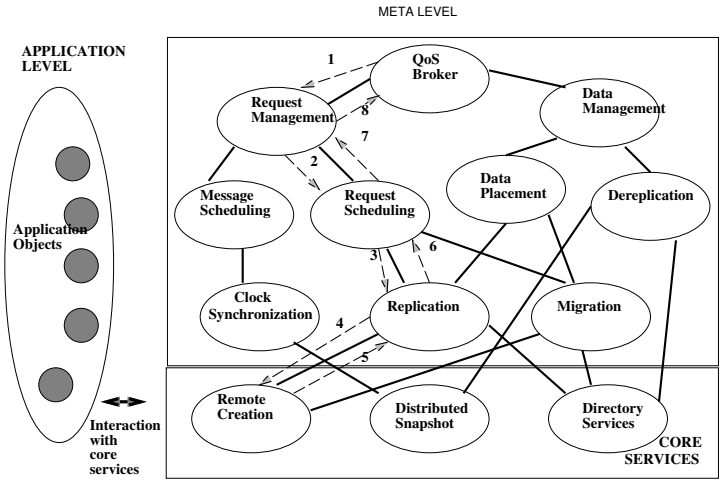

Figure 1: Detailed architecture of the QoS meta-architecture system.

scheduling module that assigns requests to servers and ensures cost-effective utilization of resources. The *message scheduling* module ensures QoS constraint satisfaction of requests that have already been initiated. The data and request management functions in turn require basic services such as clock synchronization, replication, dereplication, and migration. The services we focus on in this paper are:

•*Replication:* to replicate data actors using adaptive and predictive techniques for selecting where, when and how fast replication should proceed. The rate at which replication proceeds also has a direct impact on system performance and application interactivity (QoS).

•*Dereplication:* to dereplicate/garbage-collect data actors and optimize utilization of storage space in the distributed system based on current load in the system as well as expected future demands for the object. Dereplication cannot occur instantly—the service must ensure that a copy that has been chosen for dereplication is removed only after all requests that are currently being serviced by that copy have completed.

In order to map the QoS meta-architecture to the physical system architecture, we distinguish between local and global components and define interactions between local resource managers on nodes and the global resource management component. The global component including the QoS broker and associated meta-actors reside on the distribution controller node. The node local components include, for each DS node:

• a DS meta-actor for load-management on that node. The DS meta-actor contains state information regarding the current state of the node in terms of available resources, replicas, ongoing requests and replication processes etc.

• Request base actors corresponding to the requests assigned to that node.

• Replica base actors that correspond to the replicas of data objects currently available on that node.

4 Resource Management Policies for MM Servers

Apart from the QoS broker, QB, the MM system contains a number of meta-actors whose behaviors are coordinated by QB to provide the resource management services discussed above. In this section, we describe some of the load management policies that have been treated in the formal model. The policies are implemented as meta-level actors and provide a modular and integrated approach to managing the individual resources of a MM server so as to effectively utilize all of the resources such as disks, CPU, memory and network resources. A MM request specifies a client, one or more multi-media objects, and a required QoS. The QoS requirement

in turn is translated into resource allocation requirements. The ability of a data source to support additional requests is dependent not only on the resources that it has available, but also on the MM object requested and the characteristics of the request (e.g., playback rate, resolution). We characterize the degree of loading of a data source DS with respect to request R in terms of its load factor, $LF(R, DS)$, as:

$$LF(R, DS) = max(\frac{DB^R}{DB^{DS}}, \frac{Mem^R}{Mem^{DS}}, \frac{CPU^R}{CPU^{DS}}, \frac{NetBW^R}{NetBW^{DS}})$$

where DB^R, Mem^R, CPU^R, and $NetBW^R$ denote the disk bandwidth, memory buffer space, CPU cycles, and network transfer bandwidth, respectively, that are necessary for supporting request R and similarly Res^{DS} denotes the amount of resource Res available on data source DS. The load factor helps identify the critical resource in a data source, i.e., the resource that limits the capability of the data source to service additional requests. By comparing the load factor values for different servers, load management decisions can be taken by the QoS brokerage service. We briefly describe mechanisms for the scheduling of MM requests and placement of MM data. Optimizations to the basic mechanisms are outside the scope of this paper and are presented elsewhere [33, 31].

Scheduling of Multimedia Requests: The Request Scheduling meta-actor (RS) implements an adaptive scheduling policy that compares the relative utilization of resources at different data sources to generate an assignment of requests to replicas, so as to maximize the number of requests serviced. The data source that contains a copy of the MM object requested and which entails the least overhead (as determined by the computed load factor) is chosen as the candidate source for an incoming request. If no candidate data source can be found for servicing request R, then the meta-actor RS can either reject the incoming request or initiate *replication on demand* - implemented via a replication on demand meta-actor (ROD). The QoS broker also analyzes the rejections over time and triggers appropriate placement policies, implemented via predictive placement and dereplication meta-actors (PP and DR) to reduce the rate of rejection. The replication on demand meta-actor ROD attempts to create a new replica of a requested MM object on the fly. The source DS on which the new replica is to be made is one that has minimum load-factor with respect to the request, R, i.e., with the minimum value of $LF(R, DS)$. By doing so, ROD attempts to maximize the possibility of the QoS broker servicing additional requests from the same replica. In order for this approach to be feasible and attractive, the replication must proceed at a very high rate, thereby consuming vital server resources for replication.

Placement of MM Objects: The predictive placement and dereplication meta-actors (PP and DR) implement a placement policy that determines in advance when, where and how many replicas of each MM object should be placed in a MM server, and when to dereplicate an existing replica. In particular, the goal of the predictive placement procedure is to facilitate the task of the adaptive scheduler meta-actor, by allocating MM objects in such a way as to maximize system-wide revenue, by permitting a maximum number of requests to be admitted and scheduled for service.

Two possible optimizations to the basic adaptive and predictive phases are *eager replication* and *lazy dereplication*. With eager replication, replication of video objects occurs in anticipation when there are idle resources, even though the demand for these objects may not be apparent immediately (e.g., replication of video objects during non peak-hours). Video objects slated for eager replication may be chosen as part of the predictive placement procedure, but may actually be replicated only if sufficient video server resources become available. In the lazy dereplication strategy, when a MM object MM_i is dereplicated, MM_i's storage resources are released and marked as being available. However, the disk blocks that were being used for MM_i are rewritten only if there is an immediate need to reuse these blocks for storage of some other object. The implementation of eager repli-

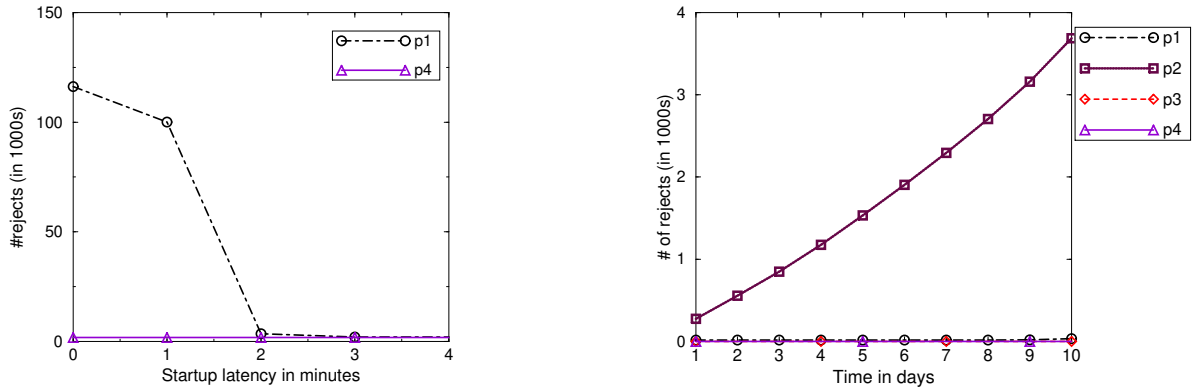

Figure 2: Comparison of the performance of load management policies for request scheduling and video placement in a distributed video server.

cation and lazy dereplication is supported by appropriate classification of video objects in the system (Figure 3 to be discussed later).

Placement mechanisms must be designed to work effectively with request scheduling. In order for predictive placement to execute concurrently with the adaptive scheduling process, a current snapshot of the system is given to the placement processes to provide a consistent view of the system state. The predictive placement and dereplication meta-actors do not consider the exact times at which requests may arrive; the adaptive scheduling meta-actor makes assignment decisions based on the exact arrival times of requests. Another potential for interference arises with the concurrent execution of the ROD and PP policies. The current PP process that initiates replica creation is based on a current snapshot of available systems resources, e.g. disk space. Without proper constraints, this snapshot will not account for replicas being created dynamically by the ROD process. A naive solution is to disable ROD when the PP process is initiated.

Some Performance Results: Performance studies show that application objects can be managed effectively by composing multiple resource management activities managed at the metalevel [33]. Figure 2 illustrates the performance, measured by request rejection rate, of various policies for load management - (a) purely adaptive (on-the-fly) scheduling and placement(P1), (b) purely predictive (decided a priori) scheduling and placement(P2), (c) composite policies that provide adaptive scheduling and predictive placement(P3 and P4, an optimized version of policy P3). The left hand side graph illustrates the request rejection rate under purely adaptive policies for placement and scheduling. *Startup latency* is a QoS factor that indicates how long the user is willing to wait for a replica to be created adaptively. The graph demonstrates that when the startup latency is below a threshold value (2 min), the purely adaptive mechanisms, represented by P1 force a very large fraction of the requests received to be rejected. Assuming that startup latency is sufficiently large, the right hand side depicts the inadequacy of P2, that relies on only predictive policies for scheduling and placement. In comparison, the other 3 policies (P1, P3 and P4), show hardly any rejects (indicated by the overlapping lines in the graph). As can be observed from the performance results, the ability to run multiple policies simultaneously (as in cases P3 and P4) reduced the total number of rejected requests in the overall system. In this paper, we study complex interactions that can arise due to the simultaneous execution of multiple system policies.

5 Reasoning about QoS-Based MM Services

Assuring safe composability of resource management services is essential for efficient management of distributed systems with widely varying and dynamically changing requirements. In this section we show how the

TLAM framework is used to model and reason about the multimedia meta-architecture and resource management policies presented above. Following our basic approach to modeling systems in the TLAM framework we specify QoS services and behaviors at several levels of abstraction and establish theorems relating the different specifications.

In § 5.1 we informally describe the notion of a system providing *QoS-based MM Service*. This is the high-level system wide request based service that the customer sees. In § 5.2 we define the notion of a system providing *Resource-based MM Service*. This reflects the relevant system resources and expresses high-level resource management requirements that must be met in order to provide the QoS-based MM Service. We postulate a function that translates QoS requirements to resource requirements and show that:

- if a system provides Resource-based MM Service, then under the given assumptions on the mapping from QoS requirements to Resource requirements, the system provides QoS-based MM Service.

In § 5.3 we define the notion of a system having *Resource Based MM Behavior*. This specification reflects the QoS broker software architecture and places constraints on the actions of the QoS meta actors. We define initial and non-interference conditions, and show that

- if a system has Resource-based MM Behavior, then if the initial and non-interference conditions hold, the system provides Resource-based MM Service.

In § 5.4 we refine the behavior by requiring the system to act according to given *Resource Based MM Broker Policies*. Here we focus on one specific policy, the *Adaptive Request Scheduling* Policy. We show that

- if a system acts according to the Resource-based MM Broker Policy, (e.g. Adaptive request Scheduling Policy), then it has Resource-based MM Behavior.

5.1 QoS based MM service

We assume that there is a fixed set, $MMObjects$, of MM objects available in the system and let MM range over $MMObjects$. We also assume given a set $MMreqset$ of MM requests—messages used to request MM service—and let $MMreq$ range over $MMreqset$. A MM request message $MMreq$ determines a triple (α_{cl}, MM, qs). This is interpreted as a request to initiate a MM streaming service from the server receiving the request to the client α_{cl}, using the MM object MM, and obeying the QoS requirement qs.

Definition 5.1 (QoS-based MM Service): A system S provides a QoS-based MM Service over the set of MM objects, $MMObjects$, and request messages $MMreqset$ iff for every configuration C of S, if there is an undelivered request message $MMreq$ in C, then along any path π from C exactly one of the following properties hold:

(1) there is a unique transition in π where $MMreq$ is accepted for service, and service is provided with the required QoS until complete, or
(2) there is a unique transition in π where $MMreq$ is rejected, and for this to happen it must be the case that the requested QoS cannot be provided at the time that $MMreq$ arrives.

5.2 Specifying a Resource Based MM Service

We assume given a function $QoSTranslate$ that maps MM requests to resource requirements which, if met, will ensure the requested QoS. Thus real-time requirements typical of MM applications, for example required bit-rate of video, are translated into corresponding resource requirements, for example a bandwidth requirement. (See [21] for examples of such QoS translation functions). For the purposes of this specification, we assume that if resource are allocated for a request, then they are used (as needed) to provide the requested QoS.

Definition 5.2 (Managed Resources): We consider four managed resources: network bandwidth ($NetBW$), CPU cycles (CPU), disk bandwidth (DB), and memory buffer (Mem). We let $Resources$ denote this set of resources and let Res range over $Resources$. We use the notation $Unit_{Res}$ for the units in which we measure the resource Res, and let $QoSTuple = Unit_{DB} \times Unit_{CPU} \times Unit_{NetBW} \times Unit_{Mem}$. For an element qt of $QoSTuple$ we write qt_{Res} to select the component associated to Res.

Definition 5.3 (QoSTranslate requirements): The function $QoSTranslate$ maps MM requests to 4-tuples representing resource allocation requirements for the four managed resources: $QoSTranslate : MMreqset \rightarrow QoSTuple$. For request $MMreq$, we require that $QoSTranslate(MMreq)$ be such that the QoS requirement of $MMreq$ is met if

(a) the resources allocated to $MMreq$ are at least those specified by $QoSTranslate(MMreq)$ together with access to a copy of the MM object of $MMreq$, and
(b) the allocated resources are continuously available during the service phase for $MMreq$.

Availability means that the MM object replica is not deleted (or overwritten) and that the total allocation never exceeds capacity, since over-allocation implies that resources must be taken from some already admitted request thereby possibly violating the QoS constraints for that request.

MM System Architecture.

The physical layer of Section 3 is represented as a TLAM network by mapping nodes in the MM server to TLAM nodes. Recall that there are several kinds of nodes: a set of data source nodes that holds replicas and provides the actual MM streaming, a distribution controller node responsible for coordinating the data source nodes, and a set of client nodes from which MM requests arise. (There is also a tertiary storage node that contains the MM objects which we omit here to simplify the presentation.) We let $DSnodes$ be the set of data source nodes and let DS range over $DSnodes$. We assume, given, a function $capacity$ such that $capacity(DS, Res) \in Unit_{Res}$ for any data source node DS and resource Res.

The MM state of the system is modeled by functions characterizing the state of replicas on each data source node and functions characterizing the state of each request that has arrived. The resource based MM service is then specified in terms of constraints on the values of these given functions and the way the values may change, such that if the MM admission and placement processes obey these constraints, then the QoS-based service requirements will be met and the underlying streaming mechanisms will be able to provide the desired QoS-based service. In particular, any admission and placement policies that obey the constraints can be used to implement the MM service.

Definition 5.4 (Functions characterizing replicas): There is at most one replica of a given MM object on any DS node, and thus it can be uniquely identified by the node and object. There are three functions characterizing the state of the replica of a MM object MM on a DS node DS in system configuration C:

- $replState(C, DS, MM) \in ReplStates = \{InQueue, InProgress, replCompleted\}$ is the replication status of the multimedia object MM on node DS. $InQueue$ indicates that replication has been requested but not initiated, $InProgress$ indicates that replication is in progress, and $replCompleted$ indicates that replication is complete.
- $replnBW(C, DS, MM) \in Unit_{NetBW}$ is the minimum bandwidth available to complete a replication, meaningful only if replication is in progress or in queue.
- $replClass(C, DS, MM) \in \{0, 1, 2, 3\}$ is the replication class of the multimedia object MM on node DS. Class 0 indicates that the replica is not present on the node. A replica of class 1 is guaranteed to be available. A replica of class 2 is considered marked as dereplicable but remains available until all requests

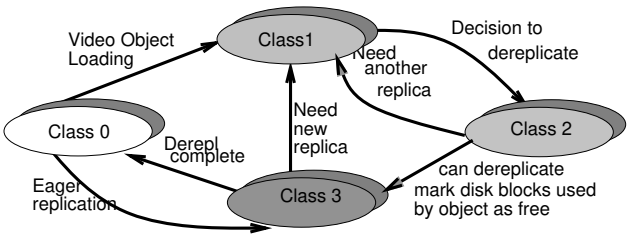

Figure 3: State transition diagram specifying the allowed changes in the class of a MM object replica. States are labelled by values of the class values and the arrows indicate allowed changes in the value as the system evolves. Class 0 indicates that the replica is not present on the node. A replica of class 1 is guaranteed to be available. A replica of class 2 is considered marked as dereplicable but remains available until all requests assigned to it have completed. The Class2 to Class3 transition is allowed only when there are no active(ongoing) requests assigned to the replica. A replica of class 3 cannot be assigned until it is changed to a 1 replica.

assigned to it have completed. A replica of class 3 exists on the node in the sense that it has not been overwritten, but there is no guarantee it will remain that way and can not be considered available until its class is changed.

The constraints, $\phi_{\text{repl}}(S)$, on the replica functions require that as the system evolves, the replication state of a replica moves from $InQueue$ to $InProgress$ to $replCompleted$ and that the $replClass$ function satisfies the constraints specified by the transition diagram given in Figure 3. For example the diagram specifies that if $replClass(C, DS, MM) = 0$, and $C \rightarrow C'$, then $replClass(C', DS, MM) \in \{0, 1, 3\}$. Also, if $replClass(C, DS, MM) = 2$, and $replClass(C', DS, MM) = 3$, then there are no active requests assigned to this replica.

Definition 5.5 (Functions characterizing requests): Each MM request that has been delivered in a system has a uniquely associated base actor that represents the request during admission control and servicing. We let $ReqActors$ be a subset of the base actor identifiers set aside for association with MM requests and let α^{req} range over $ReqActors$. We assume given, the following functions characterizing requests:

- $reqClientId(C, \alpha^{req})$ – identifies the client making the request.
- $reqObjId(C, \alpha^{req})$ – the MM object requested.
- $reqQoS(C, \alpha^{req})$ – the 4-tuple returned by $QoSTranslate$.
- $reqState(C, \alpha^{req}) \in ReqStates = \{Waiting, Granted, Denied, Servicing, reqCompleted\}$
- $reqReplica(C, \alpha^{req}) \in DSnodes + \{\texttt{nil}\}$ – the DS node to which the request has been assigned if any.

The constraints, $\phi_{\text{req}}(S)$, on the request functions require that as a system evolves, the values of $reqClientId$, $reqObjId$, and $reqQoS$ are constant throughout the life of a request actor, and once defined, the replica associated to a request actor remains constant. Furthermore, the state of request actor must move from $Waiting$ to $Granted$ or $Denied$, and from $Granted$ to $Servicing$ to $reqCompleted$, and if the state is $Servicing$, then the replication of the associated replica is competed.

The final definition needed before we state the full specification deals with the use of resources as determined for a given configuration by the replica and request functions.

Definition 5.6 (TotalResource Property ($\phi_{\text{res}}(S)$)): $\phi_{\text{res}}(S)$ states that for every configuration in the system, every data source node, and every managed resource, the sum of the resource allocation over requests on the node do not exceed the nodes total capacity for that resource. The resources currently allocated on a DS node include

resources allocated to streaming accepted MM requests, as given the $reqQoS$ function for requests assigned to that node, as well as replications that are currently ongoing.

Using the characterizing functions and constraints discussed above, we now give the requirements for a Resource-based MM service.

Definition 5.7 (Resource-Based MM Service): A system S provides Resource-based MM service with respect to requests in $MMreqset$, functions $QoSTranslate$, $capacity$, and the functions characterizing replica and request state as specified above iff

- S satisfies the constraints $\phi_{\text{req}}(S)$ (Definition 5.5), $\phi_{\text{repl}}(S)$ (Definition 5.4), and $\phi_{\text{res}}(S)$ (Definition 5.6), and
- for $C \in S$, if there is an undelivered message, $MMreq$, with parameters (α_{cl}, MM, qs), then along any computation path from C there is a (unique) transition which delivers $MMreq$ and creates a new request actor, α^{req}, such that in the resulting configuration C': $reqClientId(C', \alpha^{req}) = \alpha_{cl}$, $reqObjId(C', \alpha^{req}) = MM$, $reqQoS(C', \alpha^{req}) = QoSTranslate(qs)$, $reqState(C', \alpha^{req}) = Waiting$, and $reqReplica(C', \alpha^{req}) =$ `nil`.

Theorem 5.8 (QoS2Resource): If a system S provides Resource-based MM service as defined in 5.7 and the function $QoSTranslate$ satisfies the requirements of definition 5.3, then S provides QoS Based Service according to definition 5.1.

5.3 A Resource Based MM Behavior

At the behavior level we make explicit the QoS meta actors that cooperate to provide the Resource-based MM service. Constraints on their behavior are expressed in terms of abstract meta-actor states and reaction rules specifying allowed actions.

Representing Request and Replica functions (MM resource state)

The actual MM resource state of a configuration, modeled previously by the request and replica functions, is recorded in annotations of base-level actors implementing the MM streaming service. We partition the MM base actors of a configuration into three groups. $ReqActors$ correspond to delivered MM requests and are located on the control node. $DSReqActors$ contains the local DS node representations of the request actors representing granted requests. We define $NodeReq(C, \alpha^{req})$ to be α^{req}_{ds} if α^{req} represents a granted request with DS representative α^{req}_{ds}, and `nil` if the request status is waiting or denied. $DSReplActors$ correspond to the replica actors on the DS nodes. We define $NodeRepl(C, DS, MM)$ to be the replica base actor corresponding to the replica of MM on DS, if a replica is present, and `nil` otherwise.

Definition 5.9 (Replica and request functions): For each replica function $replX$, X one of $State$, BW, $Class$, there is an annotation tag X to represent that function: $replX(C, DS, MM)$ is the value of the X annotation of $NodeRepl(C, DS, MM)$ in C. Similarly, for each request function $reqX$ (for X one of $ClientId$, $ObjId$, QoS, $State$, $Replica$) there is an annotation tag X and $replX(C, \alpha^{req})$ is defined to be the value of the X annotation of $NodeReq(C, \alpha^{req})$ in C for a granted request, and to be the X annotation of α^{req} in C otherwise.

Representing MM Meta-Actor Behavior

Following the QoS Broker software architecture discussed in section 3 there are five broker meta actors residing on the DC node. In addition there is a DSNode manager meta actor $DSma(DS)$ on each DS node DS. The QoS meta actors, their services and their possible states are summarized in figure 4. We represent the QoS

ActorName	Service	States
QB	QoS Broker	$QBB(MMState, Status)$
RS	Request Scheduler	$IdleB_{\rm rs}$, $WaitB_{\rm rs}(MMState, ReqActors, MMState)$
ROD	Replication on Demand	$IdleB_{\rm rod}$, $WaitB_{\rm rod}(MMState)$
DR	DeReplication	$IdleB_{\rm dr}$, $WaitB_{\rm dr}(MMState, \mathbf{P}_\omega(DSnodes))$
PP	Predictive Placement	$IdleB_{\rm pp}$, $WaitB_{\rm pp}(MMState, \mathbf{P}_\omega(DSnodes))$
$DSma(DS)$	DSnode QoS mgr	$DSB(DS)$

Figure 4: QoS meta actors, services and states.

meta actors knowledge of the MM state as a function from requests (represented by actors in $ReqActors$) and replicas (represented by pairs (DS, MM)) to a function from annotation tags to corresponding values or `nil` if undefined. We let $MMState$ denote this set of functions, and we let mms, mms', $mmsU$ range over $MMState$.

The QoS broker QB coordinates the QoS resource management services: scheduling, replication, predictive placement, and dereplication. Since these activities use and modify the actual MM state, care must be taken to avoid interference among these activities. The QoS broker uses a function, $status$, to keep track of which processes are ongoing. $status(RS) = (\alpha^{req}, rod)$ for rod a boolean, indicates that RS has been requested to schedule α^{req} and scheduling is in progress, with RS allowed to invoke ROD only if rod is true, $status(RS) =$ `nil` indicates that there is no outstanding request from QB to RS and consequently no undelivered messages to or from RS. For $X \in \{PP, DR\}$, $status(X) =$ `true` indicates the process X is ongoing and $status(X) =$ `false` indicates the process X is not active and there are no outstanding requests from QB to X. We let $Status$ denote the set of status functions and let $status$ range over $Status$. The rules for QB behavior assure non-interference amongst the QoS broker services, by not allowing two replication services to run concurrently or a dereplication service to run concurrently with either scheduling or replication. Although simple and effective, this constraint reduces the potential concurrency of the QoS activities. In the full version of the paper, we discuss some ways to safely relax these constraints.

Meta-level Messages

The QoS meta-actor rules specify the reaction of a meta-actor upon receiving a QoS message. Figure 5 summarizes the internal QoS messages. In addition to the internal QoS messages there are the messages used to communicate with clients, and transition rules for these messages. A client α_{cl} may send MM requests to QB of the form $QB \triangleleft \texttt{mmReq}(\alpha_{cl}, MM, qs)$ and the reply has one of the following forms:

$$\alpha_{cl} \triangleleft \texttt{granted}(MM, qs) \mathbin{@} QB \quad \text{or} \quad \alpha_{cl} \triangleleft \texttt{denied}(MM, qs) \mathbin{@} QB.$$

The Transition Rules

In the following we briefly summarize the behavior specified by QoS Broker transition rules. The complete set of rules appears in the full paper. A sampling of rules for RS are given in the appendix.

QB rules. The transition rules for the QoS broker, QB, provide the overall organization of the QoS service activities. If neither dereplication nor request scheduling are in progress then an MM request, $QB \triangleleft$

request	reply
$RS \triangleleft \texttt{schedule}(\alpha^{req}, mms, b) \mathbin{@} QB$	$QB \triangleleft \texttt{scheduleReply}(mmsU) \mathbin{@} RS$
$DR \triangleleft \texttt{derepl}(mms) \mathbin{@} QB$	$QB \triangleleft \texttt{dereplReply}(replU) \mathbin{@} DR$
$PP \triangleleft \texttt{place}(mms) \mathbin{@} QB$	$QB \triangleleft \texttt{placeReply}(replU) \mathbin{@} PP$
$DSma(DS) \triangleleft \texttt{assign}(reqU) \mathbin{@} RS$	$RS \triangleleft \texttt{assignAck}() \mathbin{@} DSma(DS)$
$ROD \triangleleft \texttt{repl}(mms, MM, qt) \mathbin{@} RS$	$RS \triangleleft \texttt{replAck}(replU) \mathbin{@} ROD$
$DSma(DS) \triangleleft \texttt{repl}(replU) \mathbin{@} X$	$X \triangleleft \texttt{replAck}() \mathbin{@} DSma(DS)$
for $X \in \{ROD, DR, PP\} \wedge DS \in DSnodes$	
notifications	
$QB \triangleleft \texttt{notify}([(DS, MM) = [State = replCompleted]]) \mathbin{@} DSma(DS)$	
$QB \triangleleft \texttt{notify}([\alpha^{req} = [State = reqCompleted]]) \mathbin{@} DSma(DS)$	

Figure 5: Internal QoS Messages. These are classified either as requests with corresponding reply messages, or as notifications, which need no reply. The general form of a message is $X \triangleleft \texttt{mid}(\ldots) \mathbin{@} Y$ where X is the intended receiver of the message, Y is the sender, `mid` is the message type, and $(\ldots)$ contains parameters.

$\texttt{mmReq}(\alpha_{cl}, MM, qs)$, can be processed. A message $RS \triangleleft \texttt{schedule}(mms', \alpha^{req}, rod)$ is sent, where mms' is QBs current model of the MM state augmented with request information associated to the new request actor α^{req}, and rod indicates whether replication-on-demand is enabled for the scheduler.

QB may invoke predictive placement if dereplication is not in progress and replication-on-demand is not enabled, by sending a message of the form $PP \triangleleft \texttt{place}(mms) \mathbin{@} QB$. The firing of this rule depends only on the state of QB and, unlike most of the other rules, does not consume a message. Initiation of dereplication is analogous.

When a reply to an outstanding scheduling, placement or dereplication request arrives, QB updates its state using the update MM state contained in the reply message. In the case of a scheduling request, QB also sends a reply to the requesting client indicating whether the request has been granted or denied. Similarly, when a DS notification arrives the QoS broker uses the contained MM state to update its MM state. The `notify` messages are just to inform QB that some resources have been released. Notification is needed because the resources can not be considered available for reuse until such a notification is received by QB.

DS manager rules. When a DS node manager $DSma(DS)$ receives an assignment request with MM state $mmsU$ it creates a new request actor, sets the annotations of this actor and of the replica actor for the requested MM object using $mmsU$ (which contains the MM request information, and name of the associated request actor), and sends an `assignAck` reply to RS. When a DS node manager $DSma(DS)$ receives a replication request it uses the MM state replica information to update the annotations of its replica actors and then sends a `replAck` reply to the requester (which could be ROD, PP, or DR).

When servicing of a request with request actor $\alpha^{req}_{\mathrm{ds}}$ completes on a DS node, an event $\texttt{reqCompletes}(\alpha^{req}_{\mathrm{ds}})$ is signaled. The DS node manager then updates the annotations of $\alpha^{req}_{\mathrm{ds}}$ to record the completion, and sends a notification to QB with the state update for the request actor associated to this request.

Similarly, when replication of an MM object on a DS nodes completes the replica actor annotations are updated, and a notification is sent to QB. Also, if any requests are waiting for this completion they are moved

from $Granted$ state to $Servicing$ state.

Rules for Policies – RS, PP, ROD, DR .

RS. Suppose RS receives a scheduling request, $RS \triangleleft \texttt{schedule}(mms, \alpha^{req}, rod)$. If there is some DS node to which the request represented by α^{req} in mms can be assigned without violating resource constraints, then RS picks one and notifies the DS manager of that node. If there is no such DS node and ROD is disabled (rod is false), then RS sends the broker a denied update. If there is no DS suitable node to which the request can be assigned but rod is true, then RS sends a replication-on-demand request to ROD containing mms, along with the requested MM object and QoS requirement. If RS receives an acknowledgment to an outstanding assignment request to a DS node, then it sends a granted reply to the broker with MM state that contains the update information for α^{req} representing the request along with any replication update that has been done. If RS receives a replication update from an outstanding request to ROD that allows the request it is attempting to schedule to be granted, then RS picks a suitable DS node and sends a corresponding assign request to the nodes DS manager. If the replication update from ROD does not allow the request to be granted, then RS sends a denied reply to QB.

ROD. When ROD receives a request, $ROD \triangleleft \texttt{repl}(mms, MM, qt) \ @ \ RS$ for replication of MM object, MM, supporting QoS resource requirements qt, it looks, using mms, for a DS node that doesn't have the needed MM object and that has the required resources available. If one is found, a replication request is sent to that DS node and ROD waits for an acknowledgment. When the acknowledgment message is received, a reply is sent to RS with MM state containing the replica update information. If no such DS node is found, then a failure reply is sent to RS.

PP, DR. Based on the information in the MM state of a `place` request, PP may decide to reclassify some replicas from 0, 2, or 3 to 1, and in the case of moving from 0 to 1 initiate replication. It then notifies each DS node of any changes on that node, waits for acknowledgments from these DS nodes, and then sends a reply to QB containing the updated replica state. Similarly, upon receiving a dereplication request, DR may decide to reclassify some replicas from 1 to 2. It then notifies each DS node of any changes on that node, waits for acknowledgments, and then sends a reply QB containing the replica state update. (Note that in practice, PP and DR also use other information for predicting requests which we do not model here as it is not for this level of abstraction.)

Definition 5.10 (Resource-Based MM Behavior): A system S has *Resource-Based MM Behavior* with respect to the underlying system architecture (DS Nodes and capacity function), the QoS meta actors QB, RS, ROD, DR, PP, and the DS node managers $DSma(DSnodes)$, if

- for each configuration C in S, the state of X in C is appropriate for X according to Figure 4, for X one of QB, RS, ROD, DR, PP, or $DSma(DS))$ for $DS \in DSnodes$;
- every computation π of S obeys the transition rules for QoS meta-actors discussed above and guarantees termination of servicing and replication-in-progress states;
- for any transition $\tau = C \longrightarrow C'$ of S $\texttt{replCompletes}(\alpha)$ is an event of τ only if $getA(C, \alpha, State) = InProgress$, for any replica actor α, and $\texttt{reqCompletes}(\alpha^{req}_{\mathrm{ds}})$ is an event of τ only if $getA(C, \alpha^{req}_{\mathrm{ds}}, State) = Servicing$ for any DS request actor $\alpha^{req}_{\mathrm{ds}}$.

To state the "Resource-based MM Behavior provides QoS-based MM Service" theorem it is necessary to specify the conditions under which this in fact holds. For this purpose we define the requirements for QoS Initial configurations and QoS NonInterference. QoS Initial configurations are those in which no QoS meta activity is going on. QoS NonInterference expresses constraints on the environments in which the QoS system can operate correctly.

Definition 5.11 (QoS Initial): $QoSInitial(C)$ holds for a configuration C just if: QB's status function says there has no active processes; RS, ROD, DR, PP are Idle; and there are no undelivered internal QoS messages in C.

Definition 5.12 (NonInterference Requirement): S satisfies the *QoSBroker NonInterference Requirement* iff transitions that do not involve a QoS meta actor as the principle actor obey the following constraints

- Neither QoS annotations nor the state of actors in $ReqActors$, $DSReqActors$, or $DSReplActors$ are modified.
- No resource dedicated to QoS is used.
- No internal QoS messages (Definition 5) are sent.

Theorem 5.13 (Resource-based MM Behavior implies Resource-based MM Service): If a system S satisfies

- S has Resource-based MM behavior (Definition 5.10)
- S satisfies the QoS NonInterference Requirement (Definition 5.12)
- Every configuration C in S is reachable from a configuration satisfying the QoSInitial conditions (Definition 5.11)
- $QoSTranslate$ satisfies the QoSTranslate requirements (Definition 5.3)

then S provides Resource-based MM Service (Definition 5.7) with respect to the given functions $QoSTranslate$ and $capacity$, with $MMreqset$ being messages of the form $QB \lhd \texttt{mmReq}(\alpha_{cl}, MM, qs)$, and replica and request functions defined in terms of annotations according to Definition 5.9.

The proof of this theorem (which appears in the full paper) is organized as follows. First the possible system configurations are characterized in terms of combinations of meta actor states and undelivered messages is established. Then a notion of pending update for the broker model of the MM state is defined and it is shown that the broker MM state model modified by the pending updates is an accurate model of the actual MM state, and that the pending updates preserve the QoS replica, request, and resource constraints (with MM states in place of configurations). Finally, fairness of the actor model and the definition of QoS reaction rules are used to establish that the QoS broker meta actor, QB is always eventually able to receive a MM request.

Theorem 5.14 (Resource-based MM Behavior implies QoS-based MM Service): If a system S satisfies the premisses of Theorem 5.13, then S provides QoS-based MM Service (Definition 5.1) with respect to the given functions $QoSTranslate$ and $capacity$, with $MMreqset$ being messages of the form $QB \lhd \texttt{mmReq}(\alpha_{cl}, MM, qs)$.

Proof: By Theorem 5.13 and Theorem 5.8. □

5.4 Inserting an Adaptive Resource Scheduling Policy into the QoS Broker

In this section, we illustrate how we refine the Resource-Based MM Behavior Specification, by introducing the load-factor based adaptive scheduling algorithm to refine the behavior of the request scheduler meta-actor. To show this is correct we need only show that the resource-based behavior requirements are met. This follows from the fact that the algorithm meets the constraints implicit in the request scheduler transition rules.

The adaptive scheduling algorithm performs the task of data server selection. It is executed as a single transition step when the request scheduling meta-actor RS receives a scheduling request from the QoS Broker QB. The algorithm takes as input a MM request and dynamically chooses the best data source (server) on which to schedule that request. The choice of data server is adaptive in that the current load in the system when the request arrives is used as the basis for server selection.

Recall that the request scheduler receives scheduling requests of the form $RS \lhd \texttt{schedule}(mms, \alpha^{req}, rod)$ from the QoS broker, where mms is the QoS brokers perception of the current MM state. In particular mms

contains information about existing replicas and request assignments that reflects the availability of resources at each data source. We adapt the load-factor function (Section 4) to take an MM state, a request actor and a DS node as arguments, using correspondingly adapted functions for calculating the resource allocation in a given configuration. This is then used to define the $Candidates$ function that determines the best candidate DS node(s) for assignment of the request according to the load factor criteria.

Definition 5.15 (Candidate function): We calculate the available resources on each node for the load factor calculation by subtracting the resources currently allocated from the total capacity of a data source.

$$Available(mms, DS, Res) = capacity(DS, Res) - ResAlloc(mms, DS, Res)$$

where $ResAlloc(mms, DS, Res)$ is amount of resource Res allocated to requests on DS, according to the information in mms. The adaptive load factor calculation is then defined as follows.

$$LF(mms, \alpha^{req}, DS) = max(\frac{qt_{DB}}{DB^{DS}}, \frac{qt_{Mem}}{Mem^{DS}}, \frac{qt_{CPU}}{CPU^{DS}}, \frac{qt_{NetBW}}{NetBW^{DS}})$$

where

$$qt = mms(\alpha^{req}, QoS)$$
$$R^{DS} = Available(mms, DS, R) \quad \text{for} \quad R \in Resources$$

The candidate data sources, $Candidates(mms, \alpha^{req})$ based on the adaptive load factor calculation are those such that the load factor is minimal and not infinite, and a replica of the requested MM object exists on the data source.

We define Adaptive Request Scheduling MM Behavior by replacing the RS rules for scheduling requests discussed in § 5.3 by rules based on the adaptive load factor calculation. We summarize the modified rules below (see Appendix 7 for the two versions of the rules for the request scheduler). The rules for ROD are similarly modified to use the load factor calculation to find a candidate node for replication.

RS adaptive. Suppose RS receives a scheduling request, `schedule`(mms, α^{req}, rod) and D = $Candidates(mms, \alpha^{req})$. If D is none empty then RS picks a a node from D and notifies the DS manager of that node. If D is empty and rod is false, then RS sends the broker a denied update. If D is empty and rod is true, then RS sends a request to ROD containing mms, along with the requested MM object and QoS requirement.

Adaptive Request Scheduling MM Behavior is defined by a simple modification of the definition of Resource-based MM Behavior.

Definition 5.16 (Adaptive RS MM Behavior): A system S has *Adaptive Request Scheduling MM Behavior* with respect to the underlying system architecture, the QoS meta actors QB, RS, ROD, DR, PP, and the DS node managers $DSma(DSnodes)$, if it satisfies the conditions for Resource-based MM behavior, modified by replacing the request scheduler and replication-on-demand rules as discussed above.

The correctness theorem for Adaptive RS MM Behavior is the following.

Theorem 5.17 (Adaptive RS MM Behavior implies QoS-based MM Service): If

- S has Adaptive RS MM behavior (Definition 5.16)
- S satisfies the QoS NonInterference Requirement (Definition 5.12)
- every $C \in S$ is reachable from a configuration satisfying the QoSInitial conditions (Definition 5.11)
- $QoSTranslate$ satisfies the QoSTranslate requirements (Definition 5.3)

then S provides QoS-based MM Service (Definition 5.1).

Proof: By Theorem 5.14 we only need to show that under the QoS-Initial and QoS-NonInterference assumptions a system that has Adaptive RS MM Behavior also has Resource-based MM Behavior. For this, it is sufficient to check that the each transition arising from an Adaptive RS or ROD rule is a transition allowed by the corresponding generic Resource-based rule. This holds because (1) every DS node in $Candidates(mms, \alpha^{req}) = \emptyset$ can be assigned the request associated to α^{req} in mms by the generic rules, and (2) if $Candidates(mms, \alpha^{req}) = \emptyset$ then there is no DS node that can be assigned the request associated to α^{req} in mms by the generic rules. □

6 Related Work and Future Research Directions

In general, the dynamic nature of applications such as those of multimedia under varying network conditions, request traffic, etc. imply that resource management policies must be dynamic and customizable. Current mechanisms, which allow arbitrary objects to be plugged together, are not sufficient to capture the richness of interactions between resource managers and application components. In this paper we have shown, using the QoS broker MM architecture, how the TLAM framework can be used to specify and reason about distributed middleware services and their composition, and have also indicated how specifications in the framework can lead to implementations.

Commercially available object-based middleware infrastructures such as CORBA and DCOM represent a step toward compositional software architectures but do not deal with interactions of multiple object services executing at the same time, or the implication of composing object services. Architectures that provide real-time extensions to CORBA [27, 37] necessary to support timing-based QoS requirements [40] have been proposed and used to study performance optimizations [12], and patterns for extensible middleware [26]. In the Java Development Environment, the ability to deal with real-time thread management is dependent on the underlying threads implementation, making QoS support complicated to achieve. Various systems such as the Infospheres Infrastructure [7] and the Globe System [30] explore the construction of large scale distributed systems using the paradigm of distributed objects. Globus, a metacomputing framework, defines a QoS component called Qualis [17] where low level QoS mechanisms can be integrated and tested. Related work in the area of multimedia QoS management includes projects such as Omega [23], *QualMan* [22] and systems that implement a variety of algorithms for MM server management [38, 35, 20, 14, 10, 36, 9].

Reflection allows application objects to customize the system behavior [28] as in Apertos [13] and 2K [16]. The Aspect Oriented Programming paradigm [15] makes it possible to express programs where design decisions can be appropriately isolated permitting composition and re-use. Some of the more recent research on actors has focused on coordination structures and meta-architectures [2] and runtime systems such as Broadway [29] and the Actor Foundry.In other reflective models for distributed object computation [24, 8, 4], an object is represented by multiple models allowing behavior to be described at different levels of abstraction and from different points of view. Many of the middleware systems described focus heavily on implementation issues while the focus of the work presented in this paper is on developing formal semantics and reasoning for a QoS-based middleware environment. The TLAM approach models runtime services and the application using a single framework and uses the framework to reason about interactions between application actors, between metalevel services and between base and meta-level actors.

Much of the work on formal models for QoS has been in the context of QoS specification mechanisms and constructs. In some implementation driven methods of QoS specification, the specification of QoS requirements is intermixed with the service specification [18, 19]. Other approaches address the representation of QoS via multiparadigm specification techniques that specify functional behavior and performance ' constraints distinctly using multiple languages [3, 39, 5]. A translation of formal QoS specifications into monitors and controllers is discussed in [6]. Synchronizers and RtSynchronizers [11, 25] allow us to express QoS constraints via coordi-

nation constraints in the actor model.

We are actively working on extending the existing meta-architecture to support more services. Further work is required to provide a generalized model that captures the architectural resources required in the server and network to support the session connection. Specifying and reasoning about enforcement of timing-based QoS requirements of multiple sessions involves a more thorough treatment of time and synchronization. For end-to-end QoS, it is necessary to determine how *real-time scheduling* strategies for time constrained task management interact with strategies for other tasks such as CPU intensive calculations, or network communication with clients. We are currently working on formalizing other components of the MM meta-architecture such as message scheduling and synchronization. We are also actively working on invariants that can be established to minimize the overhead due to interference between meta-level services such as request scheduling and dereplication. Supporting requirements such as fault-tolerance, availability and hard real-time QoS will require further extensions of the existing MM metaarchitecture. We are currently studying the composability of multiple protocols and mechanisms to address these requirements in the MM metaarchitecture.

References

[1] G. Agha. *Actors: A Model of Concurrent Computation in Distributed Systems*. MIT Press, Cambridge, Mass., 1986.

[2] G. Agha, S. Frølund, W. Kim, R. Panwar, A. Patterson, and D. Sturman. Abstraction and Modularity Mechanisms for Concurrent Computing. *IEEE Parallel and Distributed Technology: Systems and Applications*, 1(2):3–14, May 1993.

[3] G. Blair, L. Blair, H. Bowman, and A. Chetwynd. *Formal Specifications of Distributed Multimedia Systems*. UCL Press, 1988.

[4] G. Blair, G. Coulson, P.Robin, and M. Papathomas. An architecture for next generation middleware. In *Middleware '98*, 1998.

[5] L. Blair and G. Blair. Composition in multiparadigm specification techniques. In *ECOOP Workshop on Aspect Oriented Programming*, July 1999.

[6] L. Blair and G. Blair. Composition in multiparadigm specification techniques. In *IFIP Workshop on Formal Methods for Open Object-based Distributed Systems, FMOODS'99*, Feb. 1999.

[7] K. Chandy, A. Rifkin, P. A. Sivilotti, J. Mandelson, M. Richardson, W. Tanaka, and L. Weisman. A world-wide distributed system using java and the internet. In *Proceedings of IEEE International Symposium on High Performance Distributed Computing (HPDC-5), Syracuse, New York*, Aug. 1996.

[8] F. Costa, G. Blair, and G. Coulson. Experiments with reflective middleware. In *European Workshop on Reflective Object-Oriented Programming and Systems, ECOOP'98*. Springer-Verlag, 1998.

[9] A. Dan and D. Sitaram. An online video placement policy based on bandwidth to space ratio (bsr). In *SIGMOD '95*, pages 376–385, 1995.

[10] A. Dan, D. Sitaram, and P. Shahabuddin. Dynamic batching policies for an on-demand video server. *ACM Multimedia Systems*, 4:112–121, 1996.

[11] S. Frølund. *Coordinating Distributed Objects: An Actor-Based Approach to Synchronization*. MIT Press, 1996.

[12] A. Gokhale and D. C. Schmidt. Evaluating the Performance of Demultiplexing Strategies for Real-time CORBA. In *Proceedings of GLOBECOM '97*, Phoenix, AZ, 1997.

[13] J. ichiro Itoh, R. Lea, and Y. Yokote. Using meta-objects to support optimization in the Apertos operating system. In *USENIX COOTS (Conference on Object-Oriented Technologies*, June 1995.

[14] K. Keeton and R. Katz. The evaluation of video layout strategies on a high-bandwidth file server. In *Proceedings of the Fourth International Workshop on Network and Operating System Support for Digital Audio and Video, Lancaster, UK*, pages 237–250, Nov. 1993.

[15] G. Kiczales, J. Lamping, A. Mendhekar, C. Maeda, C. V. Lopes, J.-M. Loingtier, and J. Irwin. Aspect-Oriented Programming. In *Proceedings of ECOOP'97 European Conference on Object-Oriented Programming*, June 1997.

[16] F. Kon, A. Singhai, R. H. Campbell, D. Carvalho, R. Moore, and F. J. Ballesteros. 2K: A Reflective, Component-Based Operating System for Rapidly Changing Environments . In *Proceedings of ECOOP'98 Workshop on Reflective Object-Oriented Programming and Systems*, Brussels, Belgium, July 1998.

[17] C. Lee, C. Kesselman, J. Stepanek, R. Lindell, S. Hwang, B. S. Michel, J. Bannister, I. Foster, and A. Roy. The quality of service component for the globus metacomputing system. *Proceedings of IWQoS '98*, pages 140–142, 1998.

[18] P. Leydekkers and V. Gay. Odp view on qos for open distributed mm environments. In J. d. Meer and A. Vogel, editors, *4th International IFIP Workshop on Quality of Service, IwQos96 Paris, France*, pages 45–55, Mar. 1996.

[19] F. H. d. S. Lima and E. R. M. Madeira. Odp based qos specification for the multiware platform. In J. d. Meer and A. Vogel, editors, *4th International IFIP Workshop on Quality of Service, IwQos96 Paris, France*, pages 45–55, Mar. 1996.

[20] P. Lougher and D. Shepherd. The design of a storage server for continuous media. *The Computer Journal - Special Issue on Distributed Multimedia Systems*, 36(1):32–42, February 1993.

[21] K. Nahrstedt. *Network Service Customization: End-Point Perspective*. PhD thesis, University of Pennsylvannia, 1995.

[22] K. Nahrstedt, H.-H. Chu, and S. Narayan. Qos-aware resource management for distributed multimedia applications.

[23] K. Nahrstedt and R. Steinmetz. Resource management in networked multimedia systems. *IEEE Computer*, 28(5):52–65, May 1995.

[24] H. Okamura, Y. Ishikawa, and M. Tokoro. Al-1/d: A distributed programming system with multi-model reflection framework. In A. Yonezawa and B. C. Smith, editors, *Reflection and Meta-Level Architetures*, pages 36–47. ACM SIGPLAN, 1992.

[25] S. Ren, G. Agha, and M. Saito. A modular approach for programming distributed real-time systems. *Journal of Parallel and Distributed Computing*, 36(1), July 1996.

[26] D. C. Schmidt and C. Cleeland. Applying patterns to develop extensible and maintainable orb middleware. *Communications of the ACM*, 1998. (to appear).

[27] D. C. Schmidt, D. Levine, and S. Mungee. The design of the tao real-time object request broker. *Computer Communications Special Issue on Building Quality of Service into Distributed System*, 1997.

[28] B. C. Smith. *Reflection and Semantics in a Procedural Langauge*. PhD thesis, Massahusetts Institute of Technology, Jan. 1982.

[29] D. Sturman. *Modular Specification of Interaction Policies in Distributed Computing*. PhD thesis, University of Illinois at Urbana-Champaign, May 1996. TR UIUCDCS-R-96-1950.

[30] M. van Steen, A. Tanenbaum, I. Kuz, and H. Sip. A scalable middleware solution for advanced wide-area web services. In *Proc.Middleware '98, The Lake District, UK*, 1998.

[31] N. Venkatasubramanian. *An Adaptive Resource Management Architecture for Global Distributed Computing*. PhD thesis, University of Illinois, Urbana-Champaign, 1998.

[32] N. Venkatasubramanian. Compose—q - a qos-enabled customizable middleware framework for distributed computing. In *Proceedings of the Middleware Workshop, International Conference on Distributed Computing Systems (ICDCS99)*, June 1999.

[33] N. Venkatasubramanian and S. Ramanathan. Effective load management for scalable video servers. In *Proceedings of the International Conference on Distributed Computing Systems (ICDCS97)*, May 1997.

[34] N. Venkatasubramanian and C. L. Talcott. Reasoning about Meta Level Activities in Open Distributed Systems. In *14th ACM Symposium on Principles of Distributed Computing*, pages 144–152, 1995.

[35] H. M. Vin and P. V. Rangan. Designing a multi-user hdtv storage server. *IEEE Journal on Selected Areas in Communications*, 11(1):153–164, Jan. 1993.

[36] J. L. Wolf, P. S. Yu, and H. Shachnai. Dasd dancing: A disk load balancing optimization scheme for video-on-demand computer systems. In *Proceedings of ACM SIGMETRICS '95, Performance Evaluation Review*, pages 157–166, May 1995.

[37] V. F. Wolfe, J. K. Black, B. Thuraisingham, and P. Krupp. Real-time method invocations in distributed environments. In *Proceedings of the HiPC'95 Intl. Conference on High Performance COmputing*, 1995.

[38] P. Yu, M. Chen, and D. Kandlur. Design and analysis of a grouped sweeping scheme for multimedia storage management. *Proceedings of Third International Workshop on Network and Operating System Support for Digital Audio and Video, San Diego*, pages 38–49, November 1992.

[39] P. Zave and M. Jackson. Requirements for telecommunications services: An attack on complexity. In *IEEE International Symposium on Requirements Engineering*, 1997.

[40] J. Zinky, D. Bakken, and R. Schantz. Architectural support of quality of service. *Theory and Practice of Object Systems*, 3(1), 1997.

7 Appendix: Sample Transition Rules for RS

7.1 Rules expressing behavior constraints RS

(RSscheduleGrant) If there is some DS node to which the request represented by α^{req} can be assigned without violating resource constraints then RS picks one and notifies the DS manager of that node.

$$\langle RS : IdleB_{\mathrm{rs}} \rangle \ , \ RS \triangleleft \texttt{schedule}(mms, \alpha^{req}, rod)$$

$$\longrightarrow$$

$$\langle RS : WaitB_{\mathrm{rs}}(mms, \alpha^{req}, mmsU) \rangle \ , \ DSma(DS) \triangleleft \texttt{assign}(mmsU) \ @ \ RS$$

$$\text{if } DS \in assignable(mms, \alpha^{req}) \ \wedge \ mmsU = assignTo(mms, \alpha^{req}, DS)$$

(RSscheduleDeny) If there is no DS node on which the request represented by α^{req} can be scheduled without violating resource constraints and ROD is disabled, then RS send the broker a denied update.

$$\langle RS : IdleB_{\mathrm{rs}} \rangle \ , \ RS \triangleleft \texttt{schedule}(mms, \alpha^{req}, \texttt{false})$$

$$\longrightarrow$$

$$\langle RS : IdleB_{\mathrm{rs}} \rangle \ , \ QB \triangleleft \texttt{scheduleReply}([\alpha^{req} = [State = Denied]]) \ @ \ RS$$

$$\text{if } assignable(mms, \alpha^{req}) = \emptyset$$

(RSscheduleRod) If there is no DS node on which the request represented by α^{req} can be scheduled without violating resource constraints and ROD is enabled, then RS requests replication meeting the QoS requirements of the request.

$$\langle RS : IdleB_{\mathrm{rs}} \rangle \ , \ RS \triangleleft \texttt{schedule}(mms, \alpha^{req}, \texttt{true})$$

$$\longrightarrow$$

$$\langle RS : WaitB_{\mathrm{rs}}(mms, \alpha^{req}, \texttt{nil}) \rangle \ , \ ROD \triangleleft \texttt{repl}(mms, MM, qt) \ @ \ RS$$

$$\text{if } MM = mms(\alpha^{req}, ObjId) \ \wedge \ qt = mms(\alpha^{req}, QoS) \ \wedge \ assignable(mms, \alpha^{req}) = \emptyset$$

When RS receives an acknowledgement to an outstanding assignment request to a DS node, then it sends a granted reply to the broker with MM state that contains the update information for α^{req} representing the request along with any replication update that has been done.

If RS receives a replication update from an outstanding request to ROD that allows the request it is attempting to schedule to be granted, then RS picks a suitable DS node and sends a corresponding assign request to the nodes DS manager. If the replication update from ROD does not allow the request to be granted, then RS replies to QB with an MM state indicating request denied.

7.2 Adaptive Resource scheduling rules

The Adaptive RS rules are obtained by replacing the Resource-based behavior rules for scheduling:

- **(RSscheduleGrant)** with **(AdaptiveRSscheduleGrant)**
- **(RSscheduleDeny)** with **(AdaptiveRSscheduleDeny)**
- **(RSscheduleROD)** with **(AdaptiveRSscheduleROD)**

where the new rules are defined below.

(AdaptiveRSscheduleGrant) The RS chooses the DS node on which the request represented by α^{req} can be scheduled from the set of candidates if it is non-empty.

$$\langle RS : IdleB_{\mathrm{rs}} \rangle \ , \ RS \triangleleft \texttt{schedule}(mms, \alpha^{req}, rod)$$

$$\longrightarrow$$

$$\langle RS : WaitB_{\mathrm{rs}}(mms, \alpha^{req}, mmsU) \rangle \ , \ DSma(DS_{al}) \triangleleft \texttt{assign}(mmsU) \ @ \ RS$$

if $DS_{al} \in Candidates(mms, \alpha^{req}) \neq \emptyset$

where $\quad mmsU = assignTo(mms, \alpha^{req}, DS_{al})$

(AdaptiveRSscheduleDeny) If there is no DS candidate node for the request represented by α^{req} in the MM state mms and ROD is disabled, then RS sends the broker a denied update.

$$\langle RS : IdleB_{\mathrm{rs}} \rangle \ , \ RS \triangleleft \texttt{schedule}(mms, \alpha^{req}, \texttt{false})$$

$$\longrightarrow$$

$$\langle RS : IdleB_{\mathrm{rs}} \rangle \ , \ QB \triangleleft \texttt{scheduleReply}([\alpha^{req} = [State = Denied]]) \ @ \ RS$$

if $\ Candidates(mms, \alpha^{req}) = \emptyset$

(AdaptiveRSscheduleRod) If there is no DS candidate node for the request represented by α^{req} in the MM state mms and ROD is enabled, then RS requests replication meeting the QoS requirements of the request.

$$\langle RS : IdleB_{\mathrm{rs}} \rangle \ , \ RS \triangleleft \texttt{schedule}(mms, \alpha^{req}, \texttt{true})$$

$$\longrightarrow$$

$$\langle RS : WaitB_{\mathrm{rs}}(mms, \alpha^{req}, \texttt{nil}) \rangle \ , \ ROD \triangleleft \texttt{repl}(mms, MM, qt) \ @ \ RS$$

where $\quad MM = mms(\alpha^{req}, ObjId), qt = mms(\alpha^{req}, QoS)$

if $\ Candidates(mms, \alpha^{req}) = \emptyset$