

Integrated Adaptive QoS Management in Middleware: An Empirical Case Study

Christopher D. Gill

Jeanna M. Gossett and David Corman*

Washington University, St. Louis, MO
cdgill@cse.wustl.edu

The Boeing Company, St. Louis, MO
{jeanna.m.gossett, david.corman}@boeing.com

Joseph P. Loyall, Richard E. Schantz, and
Michael Atighetchi

Douglas C. Schmidt

BBN Technologies, Cambridge, MA
{jloyall, schantz, matighet}@bbn.com

Vanderbilt University, Nashville, TN
schmidt@dre.vanderbilt.edu

Abstract – Distributed real-time and embedded (DRE) systems in which application requirements and environmental conditions may not be known a priori—or which may vary at run-time—can benefit from an adaptive approach to management of quality-of-service (QoS) to meet key constraints, such as end-to-end timeliness. Moreover, coordinated management of multiple QoS capabilities across multiple layers of applications and their supporting middleware can help to achieve necessary assurances of meeting these constraints.

This paper offers two contributions to the study of adaptive DRE computing systems: (1) a case study of our integration of multiple middleware QoS management technologies to manage quality and timeliness of imagery adaptively within a representative DRE avionics system and (2) empirical results and analysis of the impact of that integration on key trade-offs between timeliness and image quality in that system.

Index terms – Empirical Case Studies, Distributed Real-Time and Embedded (DRE) Systems, Adaptive Middleware

A. INTRODUCTION

Distributed Object Computing (DOC) middleware has become a widely accepted paradigm for developing numerous applications in a wide variety of environments, including distributed real-time and embedded (DRE) systems and applications. As DOC middleware has matured and been applied to a variety of use cases, there has been a natural growth in extensions, features, and services to support these use cases. For example, the Minimum CORBA [1] and Real-time CORBA [2] specifications, as well as the Real-Time Specification for Java (RTSJ) [3], are examples of standards that have emerged from research and experience supporting the quality of service (QoS) needs of DRE applications.

Although previous research has shown the benefits of integrating multiple QoS management techniques in standards-based middleware [4] and applying single-layer adaptive resource management techniques real-world DRE systems [5], only limited practical experience is available, however, with integrating resource management techniques *across multiple layers of standards-based DRE systems*. As a step towards filling this gap, this paper presents a case study of the vertical integration of three layers of middleware QoS management technologies within Boeing’s Bold Stroke [6], which is a standards-based DRE avionics platform. Bold Stroke is representative of a broader class of applications (including,

e.g., mission-critical distributed audio/video processing [7] and real-time robotic systems [8]) that require both static and dynamic support for QoS. In this paper, we describe the integration of our three layered QoS management technologies, show empirical results of their use in the Bold Stroke avionics system, and analyze each technology’s contribution to adaptive QoS management.

This paper is organized as follows: Section B describes the Bold Stroke avionics system’s application context; Section C describes each of the three QoS management technologies and examines the issues and optimizations we discovered while in integrating them within the avionics system; Section D describes the methodology and overall design of our experiments; Section E reports our results, and analyzes trade-offs under different adaptation approaches; Section F summarizes the lessons learned from our experiments; and Section G presents concluding remarks.

B. APPLICATION OVERVIEW

We conducted our experiments using the Weapons Systems Open Architecture (WSOA) open experimentation platform (OEP) shown in Figure 1, which consisted of two airborne server and client nodes (a command and control aircraft and an F-15 fighter aircraft respectively) that collaborated over a very low-bandwidth radio data link to re-plan the client’s mission parameters in real-time. Collaborative re-planning enables responding more rapidly to situational changes in-flight, *e.g.*, the server (C2 node) sends links to downloadable imagery to the client (F-15 node), which it then uses for re-planning. In the example scenario we used to evaluate the WSOA OEP, an off-board sensor detects time-sensitive information that initiates re-planning and provides this information to the server node. The server node has authority to initiate re-planning with the client node and sends an alert to the client node, along with a “virtual folder” that contains thumbnails of relevant images and the associated links to the complete images. Personnel on the client and server nodes collaborate to develop a new plan, which the client then performs.

The research presented in this paper applies multi-layer adaptive middleware techniques to alleviate the following limitations that now impede successful mission re-planning:

1. Limits on radio data link bandwidth that constrain the operational utility of existing systems to collaboratively re-plan the missions of airborne nodes.
2. Static resource management schemes that often rely on over-allocation strategies and reduce and sometimes

* This paper is submitted for review by the RTAS '04 program committee only: it is under review by Boeing and DARPA for approval prior to final publication.

eliminate the amount of processor and network resources available for mission re-planning and rehearsal.

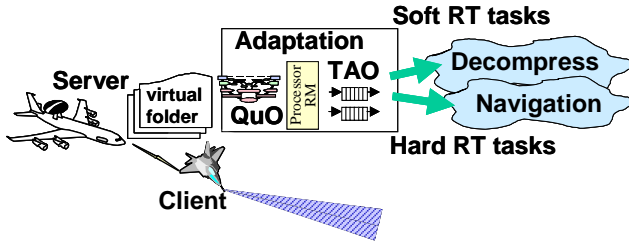


Figure 1: Collaborative Re-planning in WSOA

A key goal of the WSOA OEP evaluation system illustrated in Figure 1 is to use *adaptation* to provide the client the same level of confidence in the re-directed plan as in the original pre-planned version, even in the face of dynamic environmental factors such as variations in network bandwidth and unannounced mission re-planning alerts. Therefore, in addition to providing the client up-to-date information detected by remote sensors (e.g., fresh images of the new destination) and about the environment it will encounter en-route to and from the new destination, the OEP must manage key trade-offs between transmission quality and latency for that information. Our solution is to implement QoS-managed browser-based collaboration capabilities to (1) enable the client and server nodes to view the same displays and information and (2) ensure image quality and transmission latency stay within acceptable bounds, in a manner that is as independent as possible of the available resources (obviously there is a minimum, below which nothing useful can be accomplished).

This common browser view also allows server-side personnel to decorate imagery with annotations that will be visible on the client node rapidly, *i.e.*, within one second. The advantage of this approach is that features can be located on an image via an icon placed at a precise location relative to an easily identified reference point. This capability in turn allows personnel at the client and server nodes to establish a common frame of reference of the plan update and the new destination environment while the client is en-route to that destination, which is far better than the voice-only radio communications previously available in conventional re-planning systems.

Our solution is readily extensible to scenarios encompassing multiple client and server nodes, as well as other applications (such as coordination within teams of autonomous agents in rapidly changing environments or circumventing cascades of failures in distributed critical infrastructure) that require adaptive run-time support for collaborative re-planning.

a. Design and Implementation Overview

In the WSOA OEP application, a server-side operator first uses a user interface to send an alert to the client along with a virtual target folder containing a set of thumbnail images to the client. The collaboration client application (on the fighter aircraft) also contains a virtual folder manager component, which provides it access to and storage of virtual folders and their images. If sufficient memory is available, the virtual

folder manager can hold more than one virtual folder, though only a single virtual folder was downloaded for our OEP evaluation. The client node determines which page of the virtual folder is displayed. Personnel on the client node can navigate the virtual folder both forward and backward using “next” and “previous” buttons on their cockpit display. The virtual folder can also be reset to a *home* page by touching another button. A thumbnail page in the virtual folder allows the operator to select images to download and view without the overhead of downloading each complete image. A bar next to each thumbnail indicates whether its corresponding image has been downloaded: the bar is green if so and if not is red.

Server and client node personnel can then draw annotations and move commonly viewed individual cursors during the collaboration. To avoid problems with having both the server and client manipulate the image simultaneously, the client is given control of image download and manipulation during the collaboration, including panning side-to-side, rotation, and zooming. Server and client node personnel can move their respective cursors to indicate a specific location on the image. They are also able to draw circle, line, rectangle, and triangle annotations to designate larger regions on the image. Update messages are sent between the collaboration server and client to update cursor positions and annotations. The server to client update message contains server cursor movements and annotations drawn on the server. The client to server update message contains image manipulation information in addition to client cursor movements and client-drawn annotations. Update messages are only sent as needed and only contain updates since the last such message. Displays on both client and server are updated with the update information to maintain a common synchronized view of the virtual folder.

b. Improvements in the State of the Art

Our DOC middleware approach provides an open systems “bridge” between legacy on-board embedded avionics systems and off-board information sources and systems. The foundation of this bridge is a Real-time CORBA Object Request Broker (ORB) [2] using a pluggable protocol to communicate over a very low bandwidth (approximately 2,400 baud in each direction) Link-16 tactical data network. Link-16 time slots were allocated asymmetrically in the OEP so that the image tiles were downloaded at close to 4,800 baud with a small fraction of the bandwidth allocated to carry tile requests and update messages from the client to the server. We have applied middleware technologies at several architectural layers to manage key resources and ensure the timely exchange and processing of mission critical information. In combination, these techniques support Internet-like connectivity between server and client nodes, with the added assurance of real-time performance in a highly resource-constrained environment.

The WSOA OEP evaluation system leverages existing open systems client and server platforms. On the client side, we used an Operational Flight Program (OFP) system architecture based upon commercial hardware, software, standards, and practices [9] that supports re-use of application components across multiple client platforms. The OFP architecture

includes the Bold Stroke avionics domain-specific middleware layer [10] built upon The ACE ORB (TAO) [11], a widely-used, open-source C++ implementation of Real-time CORBA available at <http://deuce.doc.wustl.edu/Download.html>. This middleware isolates applications from the underlying hardware and operating system (OS), enabling hardware or OS advances from the commercial marketplace to be integrated more easily with the avionics application. This architecture uses the adaptive middleware technologies described in Section C to address the limitations with time-sensitive mission re-planning noted at the beginning of this section.

c. System Resource Management Model

The resource management model for the WSOA OEP evaluation system is illustrated in Figure 2. When client personnel request an image, that request is sent from the *browser application* to a *QuO application delegate* [9], which then sends a series of requests for individual tiles via TAO over a low-bandwidth Link-16 connection to the server. The delegate initially sends a burst of requests to fill the server request queue; after that it sends a new request each time a tile is received. For each request, the delegate sends the tile's desired compression ratio, which is determined by the progress of the overall image download when the request is made.

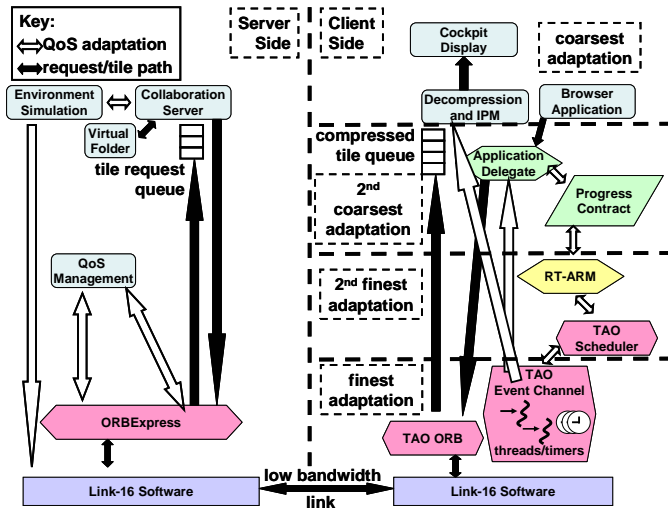


Figure 2: Resource Management Model

On the server, the *ORBExpress* Ada ORB [18] receives each request from the Link-16 connection, and from there each tile goes into a *queue of pending tile requests*. A collaboration server pulls each request from that queue, fetches the tile from the server's *virtual target folder* containing the image, and compresses the tile at the ratio specified in the request. The collaboration server then sends the compressed tile back through *ORBExpress* and across Link-16 to the client. Server-side environmental simulation services emulate additional workloads that would be seen on the command and control (C2) server under realistic operating conditions. Back on the client, each compressed tile is received from Link-16 by TAO and delivered to a servant that places the tile in a queue where it waits to be decompressed. The tile is removed from the

queue, decompressed, and then delivered by client-side operations to Image Presentation Module (IPM) hardware which renders the tile on the cockpit display. The decompression and IPM delivery operations are dispatched by the TAO Event Channel [10] at rates selected in concert by the RT-ARM [13] and the TAO Reconfigurable Scheduler [5][15], as described in Sections C.b and C.c, respectively.

C. OVERVIEW OF ADAPTIVE MIDDLEWARE

To address the challenges described in Section B, we have designed, implemented, and flight-tested an integrated multi-layered QoS enforcement architecture based on the Real-time CORBA standard. A key theme in this architecture is that coarser-grain adaptation is performed by higher layers of the architecture (*i.e.*, closer to the application), with finer grained adaptation at each lower layer (*i.e.*, closer to the OS and hardware). To enhance performance, our architecture tries to handle adaptation at the lowest layer possible, with adaptation moving up to higher layers only if QoS requirements cannot be met via adaptation in the current layer.

Figure 2 illustrates the resource adaptation architecture of the WSOA OEP evaluation platforms and middleware. The *finest* granularity of adaptation in the WSOA system architecture is the lowest priority dynamic scheduling of non-critical operations [5] by the dispatcher of the TAO Real-Time Event Channel, which we developed in previous research [12]. The *second finest* level of adaptation granularity is achieved by a Real Time Adaptive Resource Manager (RT-ARM) [13] and the TAO Reconfigurable Scheduler [5][15], which re-schedule rates of invocation of application components while maintaining deadline-feasible scheduling of critical operations. The *second coarsest* level of adaptation is performed by the Quality Objects (QuO) framework [7], which monitors progress downloading and processing image tiles toward the desired deadline for the entire image. While QuO represents the highest *middleware* layer in the OEP system architecture, the highest layer at which adaptation can be performed is the *application* layer, where the client personnel can specify *coarsest grain* requirements for image quality and timeliness. The remainder of this section describes each middleware layer outlined above in detail, from coarsest to finest grained. The remainder of this section describes the middleware in our multi-layered architecture, ranging from the coarsest to the finest granularity of adaptation.

a. QuO: Second Coarsest Grain Adaptation

QuO is an aspect-oriented middleware framework designed by BBN Technologies to support the development of QoS behavior of a system separate from – but complementary to – the development of its functional behavior. We used the following QuO components in the WSOA OEP testbed:

1. *Contracts* specify desired and available QoS, along with the policies for controlling QoS and adapting to changes.
2. *Delegates* are remote object proxies, with well-defined points to insert adaptive behaviors into end-to-end paths.
3. *System condition* objects provide interfaces to parts of the system that must be measured or controlled by contracts.

Since QuO is general-purpose framework that can support a variety of adaptation strategies, we developed a *reactive* QoS adaptation policy [17] for the OEP evaluation system that manages the overall trade-offs of timeliness versus image quality. When the client node requests an image from the server node, a QuO delegate breaks the image request up into a sequence of separate *tile* requests—each tile is a smaller-sized piece of the entire image for which a separate compression ratio can be assigned. The number of tiles requested by the delegate is based upon the image size, while the compression level of an individual tile can be adjusted dynamically based upon the deadline for receiving the full image and the expected download time for the tile. The image is tiled from the point of interest first, with the early tiles containing the most important data, so that decreased quality of later tiles will have minimal impact on the overall mission re-planning capabilities.

In the OEP evaluation system, a QuO *delegate* adapts the compression level of the next tile requested. A QuO *contract* monitors progress of the image download through system condition objects and influences the compression level of subsequent tiles based upon whether the image is behind schedule, on schedule, or ahead of schedule. If the *processing* of the image tiles falls behind schedule, the contract prompts the RT-ARM (described in Section C.b) to attempt to adjust invocation rates to allocate more CPU cycles to tile decompression. The delegate first determines the number of tiles into which the image will be broken. Due to constraints on both the server tiling software and the client display software, in the OEP evaluation system the choices were limited to 1, 16, or 64 tiles. Our experiments (described in Section D) revealed that breaking a 512 x 512 pixel image into 64 tiles introduced too much overhead, which increased the download time dramatically. We therefore always requested either 16 tiles or the entire image.

The delegate also determines the initial compression ratio for the image. We used the lowest compression ratio available for the initial tiles, because tiles are requested starting from the region of interest first and subsequent tiles are not as valuable. It therefore is more important for the application to download image tiles at compression ratios greater than or equal to that of the region of interest. After the number and initial compression ratio of tiles have been set, the delegate makes several calls to the server to request the first set of tiles. The number of tiles requested initially is determined by the size of a *tile request queue* that holds outstanding tiles requested from the server, but not yet received by the client. This queue enables the QuO encoded policy to delay requesting tiles until necessary to provide the maximum impact of compression ratio adaptation, while ensuring that there is always a tile request ready for the server to process. Finally, the delegate initiates periodic callbacks to its methods, so that it can perform contract evaluation, adjust compression ratios, and request subsequent tiles as needed to fill the tile request queue. As tiles are received from the server node, QuO system conditions count tiles received, processed, and displayed.

There are four operating regions specified by the QuO contract: *inactive*, *early*, *on time*, and *late*. The *inactive* operating region is entered when the entire image has been

downloaded. The *on time* operating region indicates that the image is on pace to complete before – but close to – its deadline. Similarly, the *early* region indicates that the image is on pace to finish well before its deadline and the *late* operating region indicates that the image will finish after the deadline at the current rate of progress. There is no change in the compression ratio if the current operating region is *on time*. If the current region is *early*, then the compression ratio is lowered to the initial compression ratio, so that the remaining tiles can have the same quality as the initial tiles. If the current operating region is *late*, the compression ratio is increased in increments of 25:1 in the range [50:1, 75:1, 100:1]. After checking progress – and if necessary setting a new compression ratio and notifying the RT-ARM of any changes in the operating region – QuO checks the request queue’s depth and requests additional tiles until the tile request queue is full or the last tile has been requested. QuO can be downloaded in open-source format from quobbn.com.

b. RT-ARM: Second Finest Grain Adaptation

The RT-ARM is a reactive resource adaptation service developed by Honeywell Technologies and used in the WSOA OEP to manage the progress of the thread(s) for decompressing received tiles and delivering them to the application by the client of the OEP. When triggered to react, the RT-ARM manipulates the CPU usage of key operations on the request/tile path, such as tile decompression and delivery of tiles to the IPM processor in the cockpit. The RT-ARM does this by manipulating subsets of task invocation event rates from application-specified available rate sets. If image tile *processing* falls behind schedule, the QuO contract prompts the RT-ARM to adjust ranges of invocation rates to re-allocate more CPU cycles to decompressing remaining tiles. In response to changing environmental conditions, the RT-ARM can trigger such adaptation in two ways: (1) *reactively* when the QuO contract notifies the RT-ARM that the operating region boundary has changed or (2) *proactively* when it periodically checks the status of the system and notices a current or impending violation of the operating region limits. We distinguish the case where the RT-ARM simply evaluates its operating status and takes no action from the case where that evaluation triggers a change in rate ranges and a corresponding re-computation of rates and priorities by the TAO Reconfigurable Scheduler described in Section C.c.

The RT-ARM attempts to keep operations within the *on time* QoS region by shrinking or expanding their respective ranges of selectable rates. This strategy was implemented by computing the average number of dispatches required by an operation at a given time, then discarding the rates that would cause the operation to complete too early or too late. As a result, rates of image processing operations that begin to veer towards the “early” and “late” regions are forced to adapt. If this level of adaptation is insufficient to keep the overall image download on time, QuO steps in and adjusts both the RT-ARM operating region and the compression level of the next tile.

c. TAO Reconfigurable Scheduler: Second Finest Grain Adaptation

The TAO Reconfigurable Scheduler is a CORBA scheduling service implementation designed for flexible support of hybrid static/dynamic scheduling [5], developed by Washington University, St. Louis. It selects a feasible set of rates of operation invocation and assigns priorities to the operations according to the scheduling strategy with which it was configured. When the RT-ARM modifies the ranges of invocation rates, the TAO Reconfigurable Scheduler first provides criticality assurance for the hard real-time operations by ensuring each operation is scheduled at a rate in its available range and that all critical operations can be feasibly scheduled at those rates. The TAO Reconfigurable Scheduler then adds non-critical processing and optimizes processor utilization for the image processing operations by maximizing their rates subject to schedule feasibility. In this application, operations associated with re-planning are non-critical.

In the earlier Adaptive Software Test Demonstration (ASTD) program [14], we tried a simple integration of the TAO Reconfigurable Scheduler with the RT-ARM, in which the RT-ARM would propose a set of rates for operations and TAO's Reconfigurable Scheduler would generate a schedule and then evaluate that schedule's feasibility. Unfortunately, that approach proved computationally inefficient since RT-ARM and TAO's scheduler operated too independently. Those results, however, pointed to the solution pursued in this work: closer integration of adaptation mechanisms. We evolved the TAO Reconfigurable Scheduler so that the rate selection *mechanism* was pushed down into it, while the *policy* for rate selection was supplied by the RT-ARM. Specifically, the RT-ARM provided a specific rate selection strategy to the TAO Reconfigurable Scheduler at system initialization time based upon operation criticality and available rates.

The first revision we made to the TAO Reconfigurable Scheduler for the WSOA OEP case study was to refactor its implementation for greater re-configurability, extending similar efforts started during the ASTD program. Our second revision incorporated rate selection into the schedule generation and feasibility analysis steps to determine an *ordering* of key operation characteristics used by a particular scheduling heuristic, assign both rates and priorities through different forms of sorting, and apply the most efficient sorting algorithm for each case. This strategy allows one scheduler to be used for efficient rate selection and priority assignment, all adaptively at run-time. These revisions are released in TAO's Reconfigurable Scheduler, which can be downloaded at deuce.doc.wustl.edu/Download.html in open-source format along with the rest of the TAO middleware.

D. METHODOLOGY FOR EMPIRICAL STUDIES

This section introduces the objectives and approach to a set of adaptive middleware experiments completed during post-flight ground tests of the WSOA OEP in January 2003, which followed the actual flight tests conducted in December 2002.

The four primary goals of our experiments were to (1) quantify the ability of multiple layered QoS management mechanisms within the Bold Stroke middleware framework to maximize image fidelity while meeting download deadlines, (2) offer preliminary assessment of the relative contributions of the different QoS management mechanisms outlined above, (3) profile the temporal performance of those mechanisms, and (4) quantify the relative benefits of this approach compared to the same application running without adaptation. We note that perceivable image quality decreases monotonically as image compression increases over the range from 50:1 to 100:1. Moreover, our assessment of the compression quality achieved for a given image is weighted by whether or not it met its deadline. These experiments also measure trade-offs between timeliness and image quality in a relatively sanitary system environment, to remove all influences outside the scope of the metrics considered here. In doing so, we established a baseline against which realistic parameters (*e.g.*, network latency jitter, traffic loads, or other factors) can be varied in a managed way and *their* contributions to system behavior also quantified.

Section D.a first introduces the metrics we used to evaluate the OEP architecture. Section D.b then describes the design of the experiments themselves, grouped into the following four distinct studies of adaptive QoS management: (1) the OEP system with no adaptation (which serves as an experimental baseline), (2) the QoS management approach described in Section C, with reactive adaptation of both image compression levels and scheduling (rates and priorities) of image tile processing operations, (3) the same approach but with scheduling adaptation turned off, and (4) a simple control-based approach to image compression adaptation that explored the system's response to this kind of control. Finally, Section D.c describes the platform on which the experiments were run. The results of these experiments are presented in Section E.

a. Evaluation Metrics

The key metrics assessed by our experiments were:

1. *Timeliness of image download*, which is measured in terms of whether the entire image was downloaded and displayed before an advertised deadline relative to the time of the image request from the application.
2. *Quality of the downloaded image*, which is measured in terms of the compression ratios of the image tiles, compared to the uncompressed version of each tile, and
3. *Scalability of the resource management approach*, which is measured in terms of the overheads of specific mechanisms in the critical path of the resource management services, *i.e.*, the QuO infrastructure, the RT-ARM service, and the TAO Reconfigurable Scheduler.

The first two metrics assess the ability of the OEP to manage multiple QoS properties simultaneously, as perceived by the collaborative mission re-planning application, while the third metric assesses the underlying middleware infrastructure itself.

In addition to studying our overall resource management approach, we also sought to examine the relative contributions of the individual mechanisms. In particular, we sought to

isolate the impacts of mechanisms for (1) end-to-end reactive image compression management and (2) client-side reactive rescheduling of tile processing operation rates.

b. Experiment Design

Our experiments were conducted using the server and client software systems developed for the WSOA OEP evaluations, including a representative Operational Flight Program (OFP) on the F-15 fighter airplane client and an imagery server on the command and control (C2) server. Resource management was conducted primarily on the client side, which is where we focused the bulk of our analysis. The experiments were run on realistic hardware in the Avionics Integration Center (AIC) laboratory at Boeing, St. Louis. We ran each experiment using the client and server system terminals in that laboratory and ran each set of trials over a range of download deadlines. Each experiment consisted of requesting a virtual folder containing compressed thumbnail renditions of the actual images being downloaded from the server. When the virtual folder arrived at the client, it then immediately requested four images in succession from the server.

Within each experiment, the same trial was then repeated with different deadlines, except for the case of experiments without adaptation where instead we set the compression ratio explicitly, and measured the download time at each of 3 fixed image compression ratios, *i.e.*, 50:1, 75:1, and 100:1. Compression ratios of 50:1 and 100:1 were selected by Boeing system engineers as upper and lower boundaries of image quality for the experiment. There was no noticeable degradation in image quality below 50:1 compression (thus making it a baseline calibration point for adaptation), while degradation was significant at 100:1. Due to time and cost constraints, we did not seek to examine the effects of different characteristics of the images themselves, but instead experimented with an assortment of images so that we could (1) quantify performance of the adaptation techniques over a range of image effects and (2) determine preliminary indications of sensitivity to image makeup for future study.

In the experiments, processing is initiated by transmission of an *Alert* from the server to the client, followed by a virtual folder with two thumbnail images. Each thumbnail serves as an additional icon to distinguish that image from the others in the virtual folder. For evaluating the performance of the WSOA adaptation architecture we confine our attention to the images themselves, though for completeness we also measured thumbnail download latencies and present them in Section E.

To assess the viability of the individual QoS adaptation technologies and the overall WSOA architecture, we ran the four experiment trials described below. In each trial the image was divided into 16 tiles, which were sent from the region of interest outward. For each tile, a message was sent from the client to the server with a request for the tile to be sent *at a given compression ratio*. The server selected the closest achievable compression ratio to that requested, transmitted the tile to the client, and recorded the ratio actually used. When the tile was received by the client, it was queued pending processing by a decompression operation, which then

decompressed the tile and delivered it via an image transfer operation to the IPM for display on the client.

Trial 1: No Adaptation of Compression or Scheduling. We first benchmarked the OEP application performance *without* adaptation to establish a baseline against which we measure improvement for the three other experiment trials. We measured the download time of each of the 4 images at each of three compression ratios (50:1, 75:1, and 100:1).

Trial 2: Reactive Compression + Scheduling Adaptation. We then measured the OEP system with adaptation of both image compression parameters and operation scheduling parameters. We instrumented the system to record the (1) end-to-end performance of the application, (2) performance of particular segments of the data and computation paths affecting end-to-end performance, and (3) overhead for key adaptation mechanisms in the infrastructure.

Trial 3: Reactive Compression Adaptation Only. To assess the relative contributions of compression vs. scheduling adaptation, we ran the same set of experiments used in the second set of trials, but with scheduling adaptation turned off. The need for this set of experiments was reinforced late in the system development phase when Boeing engineers noticed the contribution of scheduling adaptation to end-to-end performance was not evident in the Boeing Windows NT-based Desktop Test Environment (DTE). As the results in Section E reveal, this was solely an artifact of the non-real-time performance of the DTE, *i.e.*, when the VxWorks real-time OS was used in the ground and flight demo environments, the contribution of scheduling adaptation to end-to-end timeliness became clear.

Trial 4: Linear Control Law Experiments. During the experiments in the AIC laboratory, we noticed that the reactive style of compression adaptation used in the system design was resulting in very coarse-grained transitions in the image tile compression ratios, albeit with the resulting performance being suitable to the specific collaboration application. To further explore the potential applicability of our adaptation technologies outside the particular application studied, we conducted a narrowly focused set of experiments to examine the responsiveness of the OEP evaluation system to finer-grained image tile compression management.

Since imagery tiling was done from the point of interest and radiating outward, the net effect of the reactive adaptation policy was to show the largest possible area around the point of interest at highest quality and then degrade the remaining tiles as a step function to a lower resolution. While this approach is suitable for our avionics application, it is reasonable that other applications (such as opportunistic recognition of features from real-time imagery) might show less bias toward a particular single location in an image, and thus could benefit from maximizing the quality of *all* tiles. We therefore experimented with replacing the reactive tile compression adaptation strategy encoded in the QuO contract with a simple controller that sought to minimize image tile compression while still meeting the image download deadline. When each tile was received, the controller calculated a new minimum feasible compression ratio based on the image deadline and the download progress to that point.

All of the experiments conducted in this work were designed to provide a first quantitative picture of the strengths and weaknesses of the WSOA multi-layered adaptive QoS management approach, while respecting the schedule and budget constraints of the OEP program itself. An even broader set of experiments would be appropriate and in fact are motivated by the results shown in Section E, as we discuss in Section G. For these experiments, we found that 38, 42, 46, 50, 54, and 58 seconds represented a covering set of image download deadlines for the trials with both compression and scheduling adaptation, and ran only those deadlines for the two remaining trials with compression adaptation but not scheduling adaptation.

c. Experimental Platform

In the experiments whose results are reported in this paper, the client platform was a 400 MHz Dy-4 PPC 750 processor with 128 MB of memory, running the VxWorks real-time OS, version 5.3.1, with TAO version 1.0.7. The server was hosted on a flight-ready chassis incorporating multiple Alpha processors running the DEC Unix OS, with ORBExpress/RT version . A Boeing-owned operator console with dual Digital Alpha 480 MHz single board computers was used by the server-side operator. The majority of server functionality was inherited from a legacy Boeing project, whose software was tested on Digital Alpha and Sun Solaris variants of the UNIX OS. At the time of system design, only the Alpha platform was available in a ruggedized, flight-worthy package. Alpha UNIX represents a high-performance, soft real-time OS. System components were distributed across both computers, using a simulated Link-16 network over 100Base-T Ethernet cabling.

E. EMPIRICAL RESULTS

This section presents the results of the experiments described in Section D. We first examine baseline end-to-end image latencies for images compressed at the fixed ratios of 50:1, 75:1, and 100:1. We next present latencies when using the adaptation techniques described in Section C. After this, we examine image tile compression adaptation response under different strategies and present image tile queuing latencies measured on the client node. We then explore the overhead of the adaptation techniques. Finally, we present overhead results for adaptive rescheduling of operation rates using the integrated RT-ARM and TAO Reconfigurable Scheduler described in Section C.c.

End-to-End Image Latency at Fixed Compression Ratios. We first examine the total time from initial request to receive and process each image, as well as the arrival latency of the two initial thumbnail images downloaded in the virtual folder. In the following analysis we focus our attention on the download times for the entire images. For completeness, however, we include fixed-compression download times for the thumbnails. From Trial 1, Figure 3 shows quantitative points of reference for the image download times achievable at the endpoints (50:1 and 100:1) and midpoint (75:1) of the image quality range in the OEP evaluation system. We use

these points of reference to compare results achieved with different kinds and degrees of adaptation in the other trials, to assess the effectiveness of adaptation in each case and to establish quantitative bounds on the image quality and download time trade-offs achievable by adaptation in the OEP evaluation system. Over the bandwidth-limited radio data link, images compressed at the highest ratio (lowest image quality) of 100:1 took roughly 40 seconds to download (a lower bound on timeliness), and each factor of 25 reduction in the compression ratio (corresponding to improved image quality) cost another 6 to 7 seconds to download the image, thus establishing a baseline for the trade-off between timeliness and compression. We also note latency variations between the images themselves, which are also seen in the other trials.

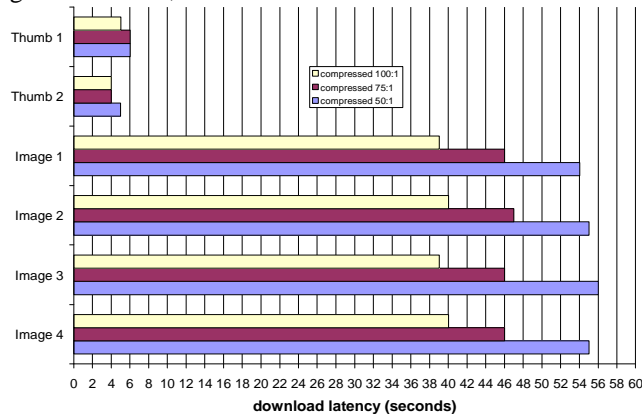


Figure 3: Image Latency Without Adaptation

Image Latency with Adaptation to Specific Deadlines. We next compare end-to-end image download times to respective deadlines in the covering set – the respective deadline is shown by the rightmost column in each group. We note that the 38 second deadline is lower than the measured latency for any image downloaded without adaptation at the highest compression ratio of 100:1, so that meeting it is infeasible. Similarly, a deadline of 58 seconds exceeds the maximum latency of any image at the lowest compression ratio of 50:1, and thus does not require any adaptation.

From Trials 2 and 3 respectively, Figure 4 and Figure 5 show the end-to-end image download latencies for deadlines of 38, 42, 46, 50, 54, and 58 seconds, both with reactive adaptation of compression ratios; in Trial 2, adaptation of operation invocation rates was also performed, while in Trial 3 it was not. These results demonstrate that compression adaptation alone is insufficient to ensure key deadlines are met. Even with adaptation of both image tile compression and operation invocation rates, however, the additional overhead of adaptation can make tight deadlines (e.g., 42 seconds) infeasible even though without adaptation they are (barely) achievable. Interestingly, the benefit of adding adaptation of operation invocation rates outweighs its cost even with tight deadlines, e.g., Figure 4 and Figure 5 show that more images made the 42 second deadline with adaptation of operation invocation rates than without such rate adaptation.

Image Compression Adaptation Response. We now consider the recorded image tile compression levels in each of

the trials. In the cases where the sequence of compression ratios was the same for more than one deadline in a given tile, we plot only the *latest* deadline of each such equivalent set.

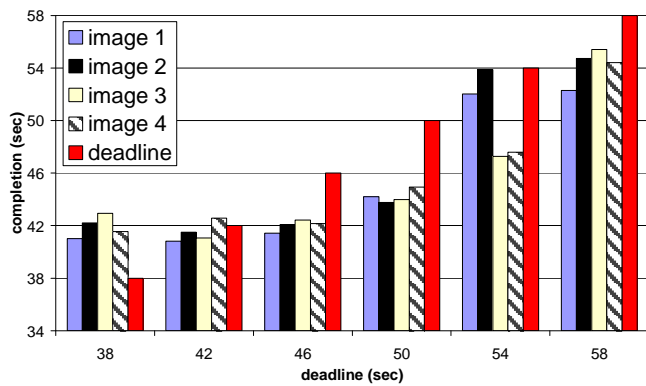


Figure 4: Both Compression and Scheduling

In Trial 3, we confined our attention to control of image tile compression only. It is therefore most appropriate to compare the experiments with simple compression control in Trial 4 to those in Trial 3. Since the scheduling adaptation mechanisms in the RT-ARM were deactivated in both experiments, the effects of scheduling adaptation are suppressed, letting us focus on the effects of compression control in isolation.

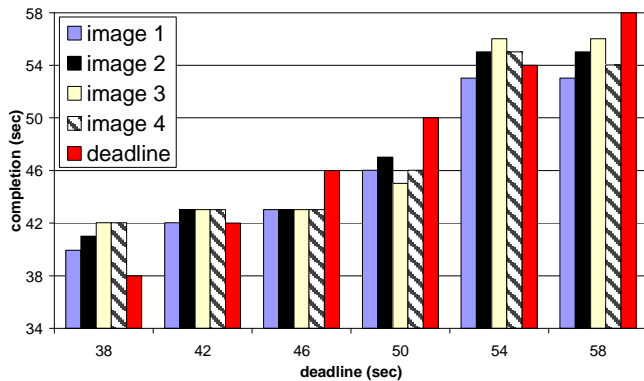


Figure 5: Compression Adaptation Only

From Trial 3, Figure 6 shows the compression level for each of the 16 tiles of image 1 for deadlines of 38, 42, 46, 50, and 54 seconds, with reactive adaptation of both compression ratios and scheduling. Similar results were observed for other images with this strategy and for all images with the strategy with reactive adaptation of both compression ratios and operation invocation rates. From Trial 4, Figure 7 shows the compression for each tile of image 1 for deadlines of 38, 42, 46, 50, 54, and 58 seconds, with simple control of compression, but no scheduling adaptation.

These results show that although it is possible to adapt image download times effectively at coarse-granularity in the compression ratios (100:1, 75:1, and 50:1), the OEP is amenable to much finer-grained compression adaptation management. This is a particularly important result when we consider the excess laxity observed at the 46 and 50 second deadlines in Figure 4. In particular, some of the time by which each image arrived early could potentially be traded back for finer-grained improvements in image quality in practice.

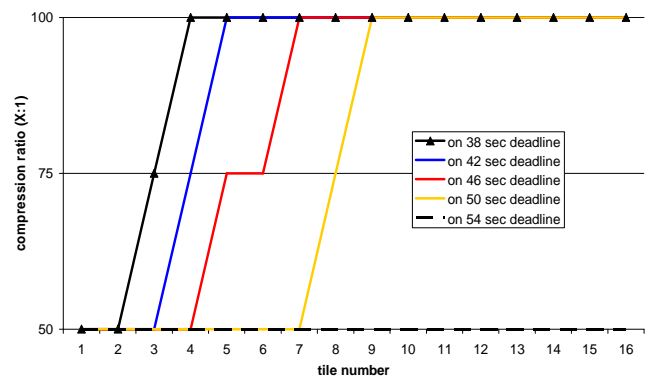


Figure 6: Reactive Compression Adaptation

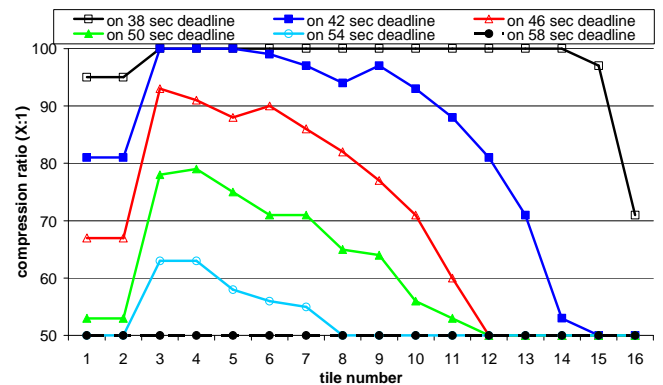


Figure 7: Compression with Simple Control

Client-side Image Tile Queuing Latency. Upon receipt from the network, each tile sent by the server is stored in a queue on the client until it is retrieved from the queue by the tile decompression operation. The rate at which the decompression operation is invoked, and thus at which tiles are retrieved from the queue was fixed at 1 Hz in Trials 1, 3, and 4, and managed adaptively in Trial 2. From Trial 1, Figure 8 shows tile queuing latencies without adaptation of either compression ratios or operation invocation rates. Similar queuing latencies were seen in Trials 3 and 4, where no adaptation of invocation rates was done. From Trial 2, Figure 9 shows the tile queuing latencies measured on the client with reactive adaptation of both compression and operation invocation rates.

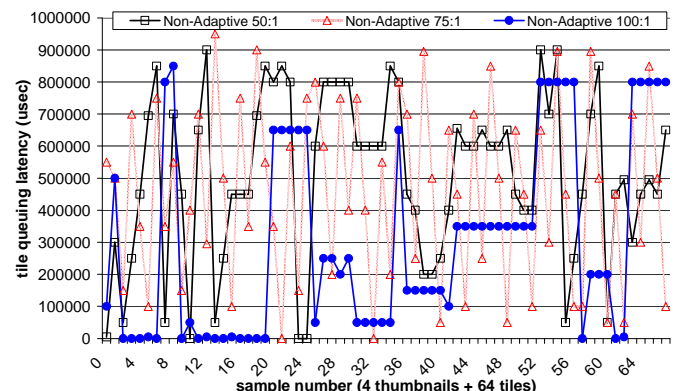


Figure 8: Tile Queuing Latency w/out Adaptation

These results identify the client-side tile receive queue as a crucial stage of the end-to-end QoS performance model for the WSOA OEP. They also highlight the importance of adaptively managing client-side tile processing operations. Adjusting the rates at which those operations are run significantly decreases the time image tiles spend idly in the queue.

Scheduler Re-computation Latency Under RT-ARM Management. Our next area of study was the measurement of schedule re-computation overhead resulting from narrowing of rate ranges by the RT-ARM and priority and rate re-assignment by the TAO Reconfigurable Scheduler. From Trial 2, Figure 10 plots schedule re-computations with adaptation of both compression and scheduling, at deadlines of 48, 42, and 38 seconds. The key insight from these results is that the number and duration of re-scheduling computations is both (1) reduced overall compared to our earlier results in the ASTD program [14] and (2) proportional to the degree of rate adaptation that is useful and necessary for each deadline. All trials showed an initial schedule computation time identical to the initial schedule computation times without rate adaptation.

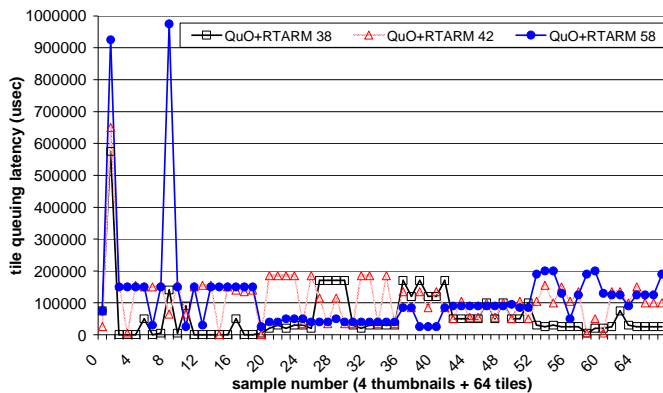


Figure 9. Tile Queuing Latency with Adaptation

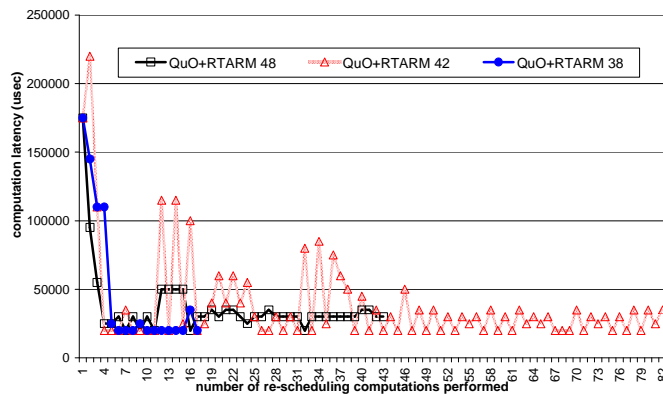


Figure 10. Adaptive Schedule Computation Latency

Overhead of QoS Management Mechanisms. In addition to examining the performance of the application as a whole, we quantify overhead of the individual adaptation services, for preliminary evaluation of scalability and possible optimization, and to guide further expansion of our resource management approach to both systems with constraints at smaller time scales and larger-scale systems of systems. Table 1 summarizes these results. These results suggest scalability of

our approach will be reasonably good overall. It is important to note that the timing capabilities of the VxWorks OS where these experiments ran was only accurate to within 5 ms, which is relevant to the overhead measurements in Table 1, many of which are in the range of 0 to 10’s of ms.

Table 1. QoS Management Latency

Mechanism	Trial 2	Trials 1, 3, 4
QuO Contract	0 – 30 msec	0 – 10 msec
Region Transition	0 – 10 msec	< 5 msec
QuO Delegate	0 – 20 msec	0 – 5 msec
RT-ARM Triggering	0 – 10 msec	N/A
Initial Schedule	185 msec	N/A

F. LESSONS LEARNED FROM EMPIRICAL RESULTS

This section summarizes the implications of the empirical results presented in Section E. We describe lessons learned from our experiments with the adaptation techniques that were described in Section C, in terms of (1) end-to-end performance with and without adaptation, (2) image compression response under different adaptation strategies, (3) tile queuing latencies on the client, and (4) overheads shown in Table 1 in Section E.

Adaptation of both tile compression and operation rates improves timeliness, but at some overhead cost. As shown in Figure 4, image 4 missed the 42 second deadline by a small margin with adaptation of *both* compression ratios and operation scheduling. However, that same image missed that deadline with all of the adaptive strategies, even though this deadline is achievable with a fixed compression ratio of 100:1 as shown in Figure 3. Imprecision of the adaptation strategies contributed to missing the deadline – reactive adaptation always started with the first two tile requests being at the lowest compression ratio of 50:1 and control adaptation started at a lower compression ratio (and finished at a lower compression ratio after the deadline was missed). We also surmise that the overhead of adaptation – though small – contributed to the difficulty in attaining this deadline. It is possible that a variation on the adaptation strategy would exhibit better results in similar situations. For example, while our adaptation policy could degrade all but the initial tiles containing the area of interest, it did not consider *dropping* any of the later tiles. The tightest feasible deadlines, *i.e.*, 42 seconds, could only be met by compressing the whole image at 100:1 as Figure 3 shows. However, with looser deadlines it might be preferable to get the first tiles at high quality and drop the last few tiles rather than degrade the whole image.

Choice of adaptation strategy is important. Overall, the strategy without scheduling adaptation sent fewer tiles at the lowest compression ratio of 50:1 before changing to the highest compression ratio of 100:1. This effect reflects an attempt by the strategy to compensate for fixed rates of tile processing operations, which was somewhat (but not entirely) successful per the latency-to-deadline comparison in Figure 5.

The principal feature of interest with the simple control strategy is the more continuous arc of the compression levels

shown in Figure 7, in contrast to the coarser-grained transitions shown in Figure 6. In future work the linear interpolation used by the controller itself should be replaced with the more rigorous control-theoretic approach we describe in Section G. However, the experimental application and supporting middleware infrastructure appear to be amenable to fine-grained adaptation as shown by the fairly continuous response of the image tile management infrastructure.

Operation rate adaptation reduces image tile queueing latencies. The main feature of interest in the image tile queueing measurements on the client is the much larger magnitude and jitter of queueing latencies without adaptation seen in Figure 8, compared to Figure 9 which shows tile queueing measurements for the strategy with adaptation of both compression ratios and tile processing operation scheduling parameters. The other two strategies without scheduling adaptation (*i.e.*, with reactive adaptation or simple control of image tile compression only) showed similar results to those without any adaptation at all, which singles out operation scheduling adaptation as a key contributor to end-to-end QoS. It is especially interesting that improvements were seen in both the precision and tightness of the latency bound – operation rate adaptation can therefore give increased confidence in *how close to that bound we can come* in improving image quality *without risking missed deadlines*.

Overhead for adaptive QoS management is acceptable. The first feature of interest for the overhead results reported in Table 1 is the relatively low latency of QuO contract evaluation, region transitions, and delegate processing. With scheduling adaptation, contract evaluations had the highest latencies but were bounded by 30 msec, and most of these evaluations took much less time than that. Without scheduling adaptation, the latencies are bounded by 10 msec and the common case is that the latencies are negligible. The version of QuO used for these experiments was designed for predictable low-latency response in distributed real-time embedded systems [7], and these results confirm the efficacy of that design.

The second feature of interest in these results is the difference in contract evaluation latency between these two strategies. Due to the low latencies seen with adaptation of compression only, we suspect that much of the increased latency seen when scheduling adaptation is added arises from preemption by OFP operations. We also observed an increased number of contract evaluations with rate adaptation enabled, however, so further studies are motivated to assess relative scalability in terms of both load and responsiveness.

We also note the relatively low latency of RT-ARM triggering operations, bounded by 10 msec, so that in concert the QuO and RT-ARM adaptation mechanisms imposed suitably low overheads. When computing the initial assignment of priorities and rates to operations, the TAO Reconfigurable Scheduler showed highly predictable timing of 185 msec. With the same initial set of scheduling parameters when no scheduling adaptation was involved, there was one invocation of the scheduler at system initialization. We note that in comparison to the latency of other adaptation mechanisms, initial schedule computation latency is an order

of magnitude greater. However, the optimizations described in Section C.c significantly reduce the post-initialization cost of rescheduling, which we consider next.

Scheduler re-computation latency under RT-ARM management is reasonable. The main feature of interest in Figure 10 is the downward settling of schedule computation times toward a similar level of overhead as for QuO contract evaluation, as the RT-ARM narrowed ranges of available rates and thus reduced the input set over which the scheduler performs its computation. We also observed an interesting *phase transition* in the number of re-computations between the infeasible and barely feasible deadlines. If we arrange trials in descending order according to the number of re-computations in each, we get 42, 46, 48, 50, 52, 54, and then 58 seconds, and then finally 38 second and 1 second deadlines showed the same minimal number of computations. The duration of the experiment for the 42 second deadline was comparable to that for other deadlines. The results therefore show that the RT-ARM was in fact triggering scheduling adaptation more frequently. As Section E showed, 42 seconds was the most difficult feasible deadline, with image 4 missing that deadline in all of the adaptive strategies. The RT-ARM performed very little adaptation for infeasible deadlines since this is futile. The RT-ARM performed more adaptation for feasible deadlines, however, but at a level proportional to the difficulty of achieving the deadline.

G. CONCLUDING REMARKS

This paper described and quantified the integration of several adaptive middleware technologies, including QuO, RT-ARM, and several layers of The ACE ORB (TAO) (*e.g.*, its Scheduling and Event Services). The paper’s R&D contributions involved (1) presenting an architecture for multi-layer adaptive middleware that is applicable to QoS-managed DRE systems and (2) conducting and analyzing empirical results that quantify the benefits and costs of this architecture for a representative DRE application, *i.e.*, the WSOA OEP mission re-planning and real-time avionics mission computing environment.

The profile of overheads for key QoS management mechanisms offers insight into the types of evolution and further experiments needed to apply our approach in other contexts, such as QoS-managed real-time audio/video streaming for disaster response, coordination of multiple autonomous vehicles, and optimization of sensing-decision-actuation paths in industrial control systems. For example, the latencies seen in the face of preemption by other operations establish bounds on responsiveness of adaptation. However, these bounds do not imply limitations on the ability of the computational and adaptation load to scale with the resources made available. In particular, for system-of-systems approaches (of which the WSOA OEP described here is a canonical example), it is reasonable to expect the resources supplied by additional platforms integrated with the system would be commensurate with the new demands they place on the system.

Careful analysis and integration of mechanisms across service boundaries are crucial to achieving necessary performance of adaptation middleware for DRE systems, *e.g.*, the re-engineered interaction between the RT-ARM and TAO Reconfigurable Scheduler described in Section C.c shows improved performance over that observed previously [19] in a simple integration of those technologies. The settling effect in the adaptation performed by the RT-ARM in conjunction with the TAO Reconfigurable Scheduler—and the questions it raises of stability of particular schedules with respect to further adaptation—is intriguing and worthy of further study. These results demonstrate that we have achieved a more efficient *mechanism* for interactive scheduling adaptation, but now need to focus new work on the *policies* for and *environmental influences* on that interaction. We have recently been investigating implementing various control strategies within the QuO adaptive framework [20][21].

The main conclusion we draw from the results in this paper is that our integrated QoS-management middleware infrastructure showed successful adaptation of multiple QoS parameters, with a quantitative improvement in management of the trade-off between image quality and download times in comparison to the same approach without adaptation. Factors in the actual DRE system environment are important, and can have a significant impact on the behavior of the system. For example, our trials revealed interesting latency effects due to interactions (notably preemption) between avionics mission computing operations and the adaptation infrastructure mechanisms. It is therefore an important achievement to have flown and measured the WSOA OEP evaluation system in a representative avionics mission-computing context.

Our future work will expand these studies to examine the effects of influences such as image contrast and size (*e.g.*, to determine why image 3 took longer to download at a compression ratio of 50:1 than any of the other images, and yet took less time to download at a compression ratio of 100:1 than either image 2 or 4), network latency, and traffic loads on WSOA OEP performance. We also plan to study the effects of applying formal feedback control to run-time adaptation of operation scheduling [16] and compression parameters in the WSOA OEP architecture.

ACKNOWLEDGEMENTS

This work was supported under the AFRL WSOA program and the DARPA Quorum and PCES programs. We wish to express our gratitude to all the program managers whose involvement was crucial to this work, especially Kenneth Littlejohn (AFRL), Gary Koob (DARPA), Lt. Col. Glen Logan (OS-JTF), and Lt. Col. Glenn Palmer (CRSIP) for their support and guidance throughout.

REFERENCES

- [1] Object Management Group. "Minimum CORBA - Joint Revised Submission," OMG Document orbos/98-08-04.
- [2] Object Management Group. "Realtime CORBA Joint Revised Submission," OMG Document orbos/99-02-12.
- [3] Bollella, *et al.*, *The Real-Time Specification for Java*, Addison Wesley Longman, 2000.
- [4] DARPA, "The Quorum Program", 1999.
- [5] Gill, Schmidt, and Cytron, "Multi-Paradigm Scheduling for Distributed Real-Time Embedded Computing", IEEE Proceedings 91(1), Jan 2003.
- [6] D. Corman, J. Gossett, D. Noll, "Experiences in a Distributed, Real-Time Avionics Domain - Weapons System Open Architecture, ISORC, Washington DC, USA, April 2002.
- [7] D. A. Karr, C. Rodrigues, Y. Krishnamurthy, I. Pyarali, and D. C. Schmidt, "Application of the QuO Quality-of-Service Framework to a Distributed Video Application," *3rd International Symposium on Distributed Objects and Applications*, Rome, Italy, September 2001.
- [8] D. B. Stewart and P. K. Khosla, "Real-Time Scheduling of Sensor-Based Control Systems," in *Real-Time Programming* (W. Halang and K. Ramamritham, eds.), Tarrytown, NY: Pergamon Press, 1992.
- [9] Loyall, Gossett, Gill, Schantz, Zinky, Pal, Shapiro, Rodrigues, Atighetchi and Karr, "Comparing and Contrasting Adaptive Middleware Support in Wide-Area and Embedded Distributed Object Applications", 21st ICDCS, April, 2001.
- [10] Sharp, David C., "Reducing Avionics Software Cost Through Component Based Product Line Development", Software Technology Conference, April 1998.
- [11] Schmidt, Levine, and Mungee. "The Design and Performance of the TAO Real-Time Object Request Broker", *Computer Communications* 21(4), April 1998.
- [12] Harrison., Levine, and Schmidt, "The Design and Performance of a Real-time CORBA Event Service," *OOPSLA '97*, October 1997.
- [13] Huang, Jha, Heimerdinger, Muhammad, Lauzac, Kannikeswaran, Schwan, Zhao, and Bettati, "RT-ARM: A Real-Time Adaptive Resource Management System for Distributed Mission-Critical Applications", Workshop on Middleware for Distributed Real-Time Systems, IEEE RTSS-97, San Francisco, California, 1997.
- [14] Doerr, Venturella, Jha, Gill, and Schmidt, "Adaptive Scheduling for Real-time, Embedded Information Systems," *18th IEEE/AIAA DASC*, St. Louis, Oct. 1999.
- [15] Gill C., Levine D., Schmidt, D. "The Design and Performance of a Real-Time CORBA Scheduling Service," *The International Journal of Time-Critical Computing Systems* 20(2), Kluwer, March 2001.
- [16] Wang, Lu, and Gill, "Feedback Control Real-Time Scheduling in ORB Middleware", 9th IEEE RTAS, Washington, D.C., May 2003.
- [17] Cross and Lardieri, "Proactive and Reactive Resource Allocation," Pattern Languages of Programs Conference (PLoP '02), Allerton Park, IL, September 2002
- [18] Objective Interface, "*ORBExpress*", www.ois.com
- [19] Christopher Gill, Douglas Schmidt, and Ron Cytron, "Middleware Scheduling Optimization Techniques for Distributed Real-Time and Embedded Systems", 7th IEEE Workshop on Object-oriented Real-time Dependable Systems (WORDS 2002), January 7-9, 2002 San Diego, CA.
- [20] Abdelwahed, Neema, Loyall, and Shapiro. "Multilevel Online Hybrid Control Design for QoS Management," Real-time Systems Symposium (RTSS), Cancun, Mexico, December 2003.
- [21] Loyall, Schantz, Zinky, and Bakken, "Specifying and Measuring Quality of Service in Distributed Object Systems", Proceedings of the 1st IEEE International Symposium on Object-oriented Real-time distributed Computing (ISORC), April 1998.