

Washington University in St. Louis

Washington University Open Scholarship

All Computer Science and Engineering
Research

Computer Science and Engineering

Report Number: WUCSE-2004-4

2004-01-26

CAMRIT: Control-based Adaptive Middleware for Real-time Image Transmission

Xiaorui Wang, Huang-Ming Huang, Venkita Subramonian, Chenyang Lu, and Christopher Gill

Real-time image transmission is crucial to an emerging class of distributed embedded systems operating in open network environments. Examples include avionics mission re-planning over Link-16, security systems based on wireless camera networks, and online collaboration using camera phones. Meeting image transmission deadlines is a key challenge in such systems due to unpredictable network conditions. In this paper, we present the design, modeling, and analysis of CAMRIT, a Control-based Adaptive Middleware framework for Real-time Image Transmission in distributed real-time embedded systems. CAMRIT features a distributed feedback control loop that meets image transmission deadlines by dynamically adjusting the quality of image... **Read complete abstract on page 2.**

Follow this and additional works at: https://openscholarship.wustl.edu/cse_research



Part of the [Computer Engineering Commons](#), and the [Computer Sciences Commons](#)

Recommended Citation

Wang, Xiaorui; Huang, Huang-Ming; Subramonian, Venkita; Lu, Chenyang; and Gill, Christopher, "CAMRIT: Control-based Adaptive Middleware for Real-time Image Transmission" Report Number: WUCSE-2004-4 (2004). *All Computer Science and Engineering Research*.
https://openscholarship.wustl.edu/cse_research/1013

Department of Computer Science & Engineering - Washington University in St. Louis
Campus Box 1045 - St. Louis, MO - 63130 - ph: (314) 935-6160.

CAMRIT: Control-based Adaptive Middleware for Real-time Image Transmission

Xiaorui Wang, Huang-Ming Huang, Venkita Subramonian, Chenyang Lu, and Christopher Gill

Complete Abstract:

Real-time image transmission is crucial to an emerging class of distributed embedded systems operating in open network environments. Examples include avionics mission re-planning over Link-16, security systems based on wireless camera networks, and online collaboration using camera phones. Meeting image transmission deadlines is a key challenge in such systems due to unpredictable network conditions. In this paper, we present the design, modeling, and analysis of CAMRIT, a Control-based Adaptive Middleware framework for Real-time Image Transmission in distributed real-time embedded systems. CAMRIT features a distributed feedback control loop that meets image transmission deadlines by dynamically adjusting the quality of image tiles. We derive an analytic model that captures the dynamics of a moderately distributed middleware architecture. A control theoretic methodology is applied to systematically design a control algorithm with analytic assurance of system stability and performance, despite significant uncertainties in network bandwidth. CAMRIT has been successfully implemented as middleware service on top of the TAO real-time CORBA ORB. Experimental results demonstrate that CAMRIT can provide robust real-time guarantees under varying bandwidth for a representative application scenario.

CAMRIT: Control-based Adaptive Middleware for Real-time Image Transmission

Xiaorui Wang, Huang-Ming Huang, Venkita Subramonian, Chenyang Lu, Christopher Gill

{wang,hh1,venkita,lu,cdgill}@cse.wustl.edu

Department of Computer Science and Engineering

Washington University, St.Louis, MO

Abstract

Real-time image transmission is crucial to an emerging class of distributed embedded systems operating in open network environments. Examples include avionics mission re-planning over Link-16, security systems based on wireless camera networks, and online collaboration using camera phones. Meeting image transmission deadlines is a key challenge in such systems due to unpredictable network conditions. In this paper, we present the design, modeling, and analysis of CAMRIT, a Control-based Adaptive Middleware framework for Real-time Image Transmission in distributed real-time embedded systems. CAMRIT features a distributed feedback control loop that meets image transmission deadlines by dynamically adjusting the quality of image tiles. We derive an analytic model that captures the dynamics of a moderately distributed middleware architecture. A control theoretic methodology is applied to systematically design a control algorithm with analytic assurance of system stability and performance, despite significant uncertainties in network bandwidth. CAMRIT has been successfully implemented as middleware service on top of the TAO real-time CORBA ORB. Experimental results demonstrate that CAMRIT can provide robust real-time guarantees under varying bandwidth for a representative application scenario.

Keywords: Quality of service control, distributed real-time and embedded middleware, adaptive real-time image transmission.

I. INTRODUCTION

Recent years have seen rapid growth of a new generation of distributed real-time embedded systems that integrate digital imaging and wireless networking technology. For example, security systems can perform automatic intruder detection through real-time fusion of images from multiple cameras connected through a wireless network [1]. Similarly, to facilitate avionics mission re-planning, personnel on multiple aircraft need to collaborate by exchanging target imagery and display annotations over Link-16 wireless networks [2]. Real-time image transmission is also important in new services on camera-equipped mobile phones (*e.g.*, online collaboration and security monitoring) that rely on “live” image transmission over cellular networks.

These embedded applications are different from traditional imaging applications (*e.g.*, online photo albums) in two ways. First, image transmission in these embedded systems is subject to stringent timing constraints. Second, although higher image

quality usually improves system utility, these next-generation embedded applications often can tolerate some degree of degradation in image quality. For example, late image delivery can be disastrous in a security system because it may result in a delayed security alarm. On the other hand, distributed event detection algorithms usually can maintain a desired probability of event detection even if input images are not perfect. Similarly, meeting deadlines is much more important in avionics mission re-planning than perfect image quality, *i.e.*, as long as key target features are still distinguishable.

These emerging embedded applications are also different from traditional embedded systems, *e.g.*, for process control in factories. While traditional embedded systems usually operate over closed and predictable networks, these new types of embedded systems need to perform image transmission across *open* and *unpredictable* networks. For example, Link-16 is widely used for tactical communication between military aircraft, but has very limited effective bandwidth (*e.g.*, roughly 30 to 340 Kbps divided among all aircraft communicating with a common JTIDS terminal [3]). Furthermore, network bandwidth may vary significantly during a mission due to changes in weather, terrain, and communication distance [2]. Cellular networks and wireless sensor networks share similar characteristics [4]. These bandwidth-constrained and unpredictable networks make real-time image transmission an extremely challenging task.

To address the challenges of real-time image transmission in bandwidth-constrained and unpredictable networks, we have developed *CAMRIT, a Control-based Adaptive Middleware for Real-time Image Transmission*. The CAMRIT project has made three main contributions to the state of the art in performance control for distributed real-time embedded systems.

- 1) *Control Architecture*: We present a novel middleware architecture for *control-based* adaptive management of image transmission. Our architecture features a distributed feedback control loop that supports fine-grained control over the progress of image transmission by dynamically adjusting the quality factor of image tiles.
- 2) *Control Modeling*: We derive an analytic model that captures the dynamics of a moderately complex distributed middleware architecture. Control analysis shows that CAMRIT can assure system stability and transmission latencies under a wide range of available network bandwidth.
- 3) *Middleware Implementation*: CAMRIT has been implemented as a middleware service based on the TAO [5]

real-time CORBA object request broker. Experimental results on a characteristic testbed demonstrate that CAMRIT can provide robust real-time assurance under representative application scenarios.

II. MIDDLEWARE ARCHITECTURE

The primary goal of CAMRIT is to complete transmitting an image from a server node to a client node within a user specified deadline. At the same time, CAMRIT aims to maximize image quality (under the constraint of the transmission deadline) because an higher quality image usually has higher utility to the application. This requirement excludes trivial solutions such as always sending an image at the lowest quality.

To achieve both goals despite an unpredictable network, CAMRIT employs a feedback control loop that dynamically adjusts image quality based on performance feedback. CAMRIT exploits existing image compression standards that support flexible image quality. For example, the widely adopted JPEG [6] standard provides a user-specified parameter called the *quality factor* which can be any integer from 1 to 100. Since a lower quality factor leads to a smaller image size after compression, the quality factor parameter provides a knob for controlling the time it takes to transmit an image. However, JPEG only supports a *single* quality factor for a whole image. This is insufficient for our feedback control loop, which needs to adjust the quality factor of an image dynamically during its transmission. To support such fine-grained adaptation, CAMRIT splits each image into tiles. Each of which may be compressed with a separate quality factor. The combination of tiling and flexible quality allows us to design an actuator for our feedback control loop.

CAMRIT is designed as a *middleware service* for real-time CORBA. Our approach is to integrate image transmission with DRE software in a common middleware framework. All the tasks in CAMRIT are managed and scheduled according to RMS [7] using the Kokyu [8] dispatcher within the TAO Real Time Event Channel [9]. We note that Kokyu only addresses CPU scheduling, and is not concerned with the image *transmission* delay over a network, which is also a focus of CAMRIT. On a bandwidth-constrained network such as Link-16, the network transmission time often dominates CPU processing time in the end-to-end transmission latency of an image.

A. Service Interface

An application interacts with CAMRIT's ImageTransmissionService interface, specified in CORBA IDL. The following parameters are passed to the service:

- *image_id*: An identifier (e.g., an image file name) for the requested image.
- *deadline*: The relative deadline for delivery of the image. The entire image must be received within *deadline* seconds after the request is sent from the client to the server.

- *num_tile*: The number of tiles into which the image is divided. This parameter allows the application to specify the granularity of control of the image quality, with a trade-off of increased overhead for finer granularity.
- *quality_range*: The defined range of acceptable image quality. This parameter allows configuration of application-specific image quality constraints.

Note that the CAMRIT service implementation serves to hide properties of the underlying network from the the application, particularly the variations in available bandwidth over a network, and delivers the image within the specified deadline. Figure 1 shows the major components of the CAMRIT architecture. We first describe the mechanisms responsible for requesting and transmitting an image, and then discuss the feedback loop for controlling transmission latency.

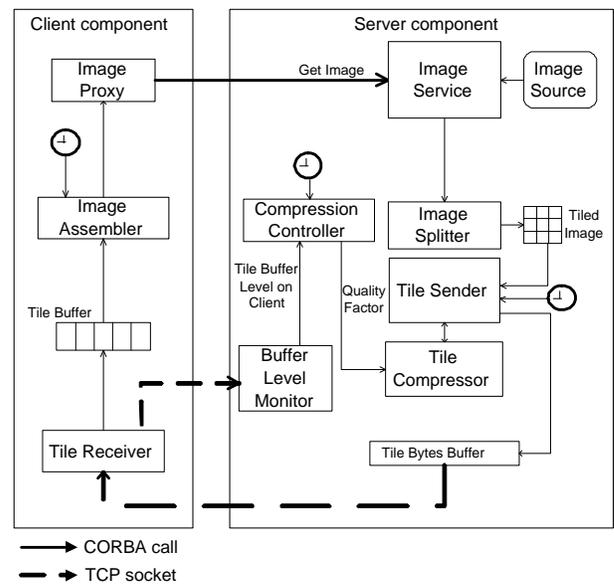


Fig. 1. Overview of the CAMRIT Architecture

B. Image Transmission

The CAMRIT middleware architecture is made up of client and server components, each on a separate endsystem. The *Image Proxy* object in the CAMRIT client component provides the service interface to the application. When it receives a request for an image, this object makes a CORBA call to the *Image Service* object on the server. This CORBA call has the same parameters as the service interface. A *one-way* CORBA call is used to avoid blocking the client thread that executes the call, because transmitting a large image over a bandwidth-constrained network may take a long time (e.g., hundreds of seconds).

The *Image Service* object is implemented as a CORBA servant in the server component, and is advertised to the outside world. When it receives the CORBA call from the client, the *Image Service* object retrieves the requested image (e.g., from an image repository or a camera), and calls the

Image Splitter object to split the retrieved image into a specified number of tiles. Each tile is compressed by the *Tile Compressor* object according to the current quality factor, which is periodically updated by the *Quality Controller* object described in Section II-C. The *Tile Sender* object then sends each compressed tile, as a byte stream through a TCP socket, to the client component.

Since the network is the bottleneck, the available bandwidth must be fully utilized. To accomplish this, the Tile Sender and Tile Compressor are executed by a periodic task. In each invocation, the Tile Sender fills the TCP buffer. The period of this task is chosen such that the TCP buffer never goes empty while an image is being transmitted to the client. Specifically, if B is the TCP buffer size and b_{max} is the maximum bandwidth of the network, the period of the Tile Sender is set to no higher than $\frac{B}{b_{max}}$. This guarantees that the TCP layer in the kernel has enough bytes of data in the TCP buffer to send before the next invocation of the sending task.

Note that we use a non-blocking TCP socket instead of a CORBA call to transmit an image tile. This is because a TCP socket allows finer control granularity over the amount of data pushed into the TCP buffer, and does not incur the transmission overhead of a GIOP header for each tile. For example, the sender may push a fraction of a tile into the TCP buffer. The sending socket is set to NON_BLOCKING mode so that the kernel will inform the application layer through an EWOULDBLOCK error from the *send* system call if the TCP buffer is full.

Pseudo-code for this periodic task is shown below. *Tile.Bytes.Buffer* is a buffer on the server that is used to hold the bytes of a tile (or fraction of a tile) to be sent.

```

Tile.Sender :: handle_timeout ()
{
  while (1)
  {
    return_code = send bytes in Tile.Bytes.Buffer to socket ;
    if ( return_code == EWOULDBLOCK) /* TCP buffer is full */
      exit the current invocation ;
    Compress next tile with current quality factor ;
    Create a header for the tile ;
    Append the new compressed tile to Tile.Bytes.Buffer ;
  }
}

```

The *Tile Receiver* object on the client reads the byte stream from the socket. The boundaries between tiles are indicated in the tile header that precedes each tile. After it receives a whole tile, the Tile Receiver object enqueues the tile into a buffer that holds received but still compressed tiles.

The *Image Assembler* is executed as a periodic task. The first instance of this task is released when the first tile of the image is inserted into the tile buffer. In every invocation, it dequeues and decompresses a tile from the tile buffer if it is not empty. When all the tiles of an image have been decompressed, it assembles them back into a whole image and notifies the Image Proxy, which then returns a handle (e.g., the memory address) for the decompressed image to the application.

The CPU utilization is bounded such that each invocation of the Image Assembler is completed before the next invocation. The period of the Image Assembler is selected to meet the end-to-end image deadline, as follows. Suppose the first tile of an image is inserted into the tile buffer t_1 sec after the image request is sent to the server. The period must then satisfy the following condition in order to guarantee the whole image is received and decompressed by the deadline:

$$t_1 + p * num_tile \leq deadline$$

Hence, the upper bound for the Tile Assembler period is:

$$p \leq \frac{(deadline - t_1)}{num_tile} \quad (1)$$

C. Feedback Control Loop

In the transmission scheme described in Section II-B, if the tile buffer on the client never becomes empty during the transmission of an image, the Image Assembler can always decompress all the tiles by the deadline. Therefore, a sufficient condition to meet the deadline is to maintain a buffer level that is no lower than one tile during the transmission of an image. However, this is not easy when network bandwidth is unpredictable. While the Image Assembler dequeues tiles from the tile buffer at a constant rate, the rate at which tiles are inserted into the tile buffer (called the *tile enqueue rate*) depends on the network bandwidth and the size of compressed tiles. To deal with the unpredictable network, we designed a feedback control loop to maintain a specified buffer level (the set point) by periodically adjusting the quality factor of the remaining tiles that are yet to be transmitted. The feedback control loop is composed of a *Buffer Level Monitor*, a *Quality Controller*, and the Tile Compressor described earlier, which serves as an actuator in the control loop.

Each time the Tile Receiver on the client reads a chunk of data from the socket (i.e., completes a `read()` call), it sends the current tile buffer level to the Buffer Level Monitor on the server. Note that the reported buffer level includes the fraction of the tile that is currently being received by the client. For example, if the tile buffer currently contains 3 tiles, and the Tile Receiver has received the first 2KB of another tile of size 5KB, the current buffer level is $3 + 2/5 = 3.4$. The Buffer Level Monitor makes this information available to the Quality Controller. The use of fractional buffer levels as feedback improves control performance because it gives a more precise representation of the buffer level than would integer values.

In our design, the Tile Receiver *pushes* the buffer level to the Buffer Level Monitor when the buffer level changes. An alternative design would be to have the Quality Controller *poll* this value at each sampling instant of the feedback control loop. The advantage of the push approach for the kinds of real-time image transmission applications for which CAMRIT is designed, is reduced control timing jitter – the control computation does not need to wait for transmission of the current buffer level from the client to the server over a slow network. On the other hand, the push approach does

introduce more communication overhead. However, we find this overhead to be negligible in our experiments.

The Quality Controller periodically re-computes the quality factor of the remaining tiles based on the current tile buffer level. Intuitively, if the buffer level is lower than the set point, CAMRIT needs to increase the tile enqueue rate. Since the network bandwidth is fully utilized, the Quality Controller increases the enqueue rate by reducing the quality factor (and hence the size) of the remaining tiles so that more of the compressed tiles can be transmitted over the network within a sampling period. Conversely, if the buffer level is higher than the set point, CAMRIT needs to increase the quality factor. The new quality factor is then used by the Tile Compressor to compress the remaining tiles that are sent in the following sampling period. Clearly, the Quality Controller is critical to the performance of CAMRIT. We discuss the design of the control algorithm next, in Section IV.

III. DYNAMIC MODELS

Modeling the dynamics of the controlled system is crucial for control design. It is also a key challenge in complex distributed middleware systems, whose dynamics are not understood as well as those of many physical control systems. In this section we formulate a dynamic model for a characteristic real-time image transmission system controlled by our feedback control loop.

A. Controlled System Model

As described in the Section II, the controlled variable in our feedback control system is the tile buffer level on the client, and the manipulated variable is the quality factor used by the server to compress tiles. Before formulating the dynamic relationship between the tile buffer level and the quality factor, we first introduce some essential notation:

- T : the sampling period of the feedback control loop.
- $l(k)$: the tile buffer level at the k^{th} sampling point (kT sec after the system starts). As described in Section II, $l(k)$ may include a fraction of a tile.
- l_s : the set point, *i.e.*, the desired tile buffer level.
- r : the constant rate (*i.e.*, the frequency) at which tiles are dequeued from the tile buffer by the Image Assembler. It is equal to the inverse of the period of the Image Assembler task.
- $b(k)$: the network bandwidth in the k^{th} sampling period, $[kT, (k+1)T)$. The value of $b(k)$ is unknown *a priori* in an unpredictable network environment, but its range $[b_{min}, b_{max}]$ is usually known.
- s : the size of an uncompressed tile. This is known and fixed for a given image and number of tiles.
- $s(q)$: the average size of a tile compressed with a quality factor q .
- $q(k)$: the quality factor computed by the controller at the k^{th} sampling point.

In each sampling period, rT tiles are dequeued from the tile buffer. Supposing $n(k)$ tiles are transmitted and inserted

to the tile buffer in the k^{th} sampling period, we then have this equation:

$$l(k+1) = l(k) + n(k) - rT \quad (2)$$

$n(k)$ depends on the size of compressed tiles and the network bandwidth. The size of a compressed tile is a (usually non-linear) function of the quality factor used to compress it. For the purpose of control design, we linearize this function such that

$$s(q) = \frac{sq}{g} \quad (3)$$

where g is a gain that can be estimated through linearization in the steady-state operation region of the system. The details of the linearization are presented in Section III-B.

In our control design, we assume $b(k) = b$ where b is the nominal bandwidth. Although we design the controller based on b , the controller is tuned such that it remains stable as long as the bandwidth stays within the range $[b_{min}, b_{max}]$.

If we ignore control delay, we get a simple first-order model for the controlled system:

$$l(k+1) = l(k) + \frac{bTg}{sq(k)} - rT \quad (4)$$

Unfortunately, this model is inaccurate because control delay plays a major role in the dynamics of our distributed middleware. The control delay $t_d(k)$ in our system is the time interval between the moment when the controller on the server outputs the new quality factor $q(k)$ and the moment when this new quality factor starts to have an effect on the tile buffer on the client.

The control delay is due to residual data in the TCP buffer and the Tile Byte Buffer on the server, as Section II describes. When the controller outputs a new quality factor, these buffers still contain tiles compressed with the *old* quality factor, $q(k-1)$. Hence the system will continue to transmit and enqueue those *old* tiles to the tile buffer on the client until all the data in the TCP buffer and the Tile Byte Buffer have been transmitted to the server.

Let $s_t(k)$ and $s_b(k)$ denote the amount of data in the TCP buffer and the Tile Byte Buffer, respectively. The control delay is then

$$t_d(k) = \frac{s_t(k) + s_b(k)}{b} \quad (5)$$

To calculate the control delay, we need to estimate $s_t(k)$ and $s_b(k)$. First, we consider $s_t(k)$. Suppose the TCP buffer size is B , and the period of the Tile Sender task is p_s . The TCP buffer is full (*i.e.*, contains B bits of data) at the end of each invocation of the Tile Sender task. During each period of the Tile Sender, bp_s bits of data are transmitted from the TCP buffer. Therefore, the lower bound for the amount of data that the TCP buffer may hold is $B - bp_s$ bits. Since $s_t(k)$ depends on the specific time when the controller outputs $q(k)$, we approximate $s_t(k)$ with the average of its upper bound and lower bound for our control design:

$$s_t = B - \frac{bp_s}{2} \quad (6)$$

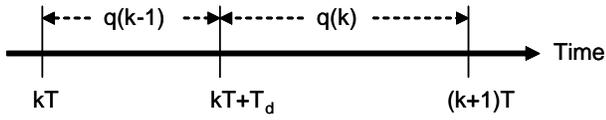


Fig. 2. Quality Factors of Tiles Received in the k^{th} Sampling Period

As Section III-B describes, the Tile Byte Buffer holds the fraction of a compressed tile that cannot fit into the TCP buffer. On average, this buffer contains half of a tile compressed with quality factor $q(k-1)$ at the beginning of the k^{th} sampling period. We approximate $s_b(k)$ with its average value, based on (3):

$$s_b(k) = \frac{sq(k-1)}{2g} \quad (7)$$

As Figure 2 illustrates, if we choose a sampling period $T > t_d(k)$, the tiles placed into the tile buffer in the first $t_d(k)$ seconds of the k^{th} sampling period are compressed with quality factor $q(k-1)$, and the tiles placed there in the remaining part of the sampling period are compressed with quality factor $q(k)$. Therefore, a more accurate model that considers the control delay is

$$l(k+1) = l(k) + \frac{bt_d(k)g}{sq(k-1)} + \frac{b(T-t_d(k))g}{sq(k)} - rT \quad (8)$$

Note that the last term in (8) is non-linear because it includes both $q(k)$ and $t_d(k)$, which is a function of $q(k-1)$ (see (5) and (7)). Since the quality factor does not change significantly in a steady state, we can linearize this model by replacing the $q(k-1)$ in this term with $q(k)$. Finally, let $u(k) = 1/q(k)$ be the control input. We then have an approximate linear model of the controlled system:

$$l(k+1) = l(k) + Au(k) + Cu(k-1) + D \quad (9)$$

where $A = \frac{(bT-s_t)g}{s}$, $C = \frac{s_t g}{s}$ and $D = -rT$.

B. Tile Size and Quality Factor

We now describe how to estimate the gain g . We first compare the size of the compressed sample image $s(q)$ with each quality factor q , and plot the *inverse of the compression ratio* $a(q) = s/s(q)$ as a function of the inverse of the quality factor $u = 1/q$, which is the control input. We measured the relationship between the inverse of the compression ratio and the inverse of the quality factor, for an example image shown in Figure 3. The resulting profile is plotted in Figure 4, which shows that the relationship between those parameters is clearly non-linear. We linearize $a(u)$ in the operational region of the system in steady state, in the following three steps.

- 1) Given the deadline d for transmission of an image, the rate r of the Image Assembler is calculated using (1). In steady state, tiles are transmitted from the server to



Fig. 3. Image 0

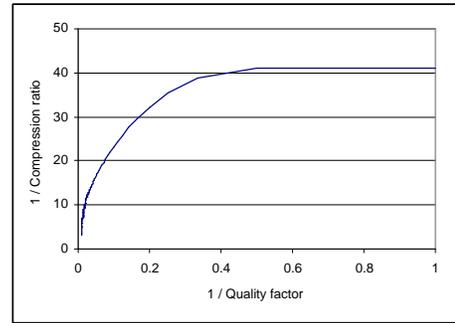


Fig. 4. Quality Factor vs. Compression Ratio

the client at the same rate as r , to maintain a constant tile buffer level.

- 2) We then use (10) to calculate the range of the inverse of the compression ratio, $[a_{min}, a_{max}]$, that can satisfy the tile transmission rate r in steady state based on the range of possible network bandwidth $[b_{min}, b_{max}]$.

$$\frac{ba(u)}{s} = r \quad (10)$$

- 3) Finally, we perform linear regression on the segment of function $a(u)$ where $a_{min} \leq a(u) \leq a_{max}$. The slope of the linear regression is the estimated g .

When an image request is submitted, CAMRIT uses the estimation process above to derive g , based on the specified deadline and the function $a(u)$ from the profiling results for a representative image. While function $a(u)$ may differ for different images, the difference is small for images in a similar application domain (e.g., landscape images taken from airplanes). Furthermore, the feedback control loop can be designed to tolerate a range of variations in g .

As an example, we now show how to estimate g based on hypothetical but plausible system settings, and using the measured profile shown in Figure 4 for the sample image shown in Figure 3. The key parameters for this example are

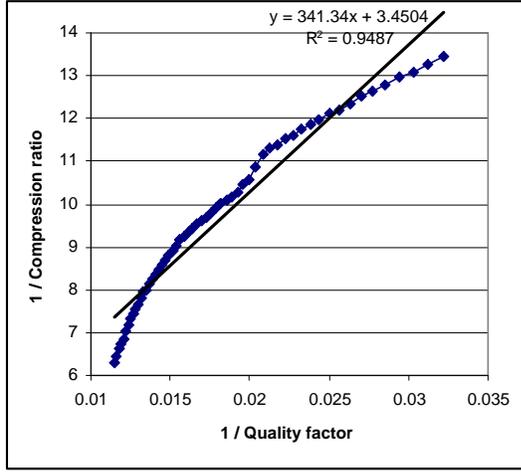


Fig. 5. Linearization of $a(u)$

as follows:

- Image: 640×640 pixel resolution; divided into 64 tiles; each uncompressed tile size $s = 18.75$ KB.
- Deadline: $d = 200$ sec.
- Bandwidth: [4 Kbps, 8 Kbps]. The top of this bandwidth range approximates the maximum data rate of a single link at the lowest Link-16 network capacity of 28.8 Kbps [10], with time slots divided among links to 3 aircraft collaborating with a common JTIDS terminal on the C2; we assume a minimum network bandwidth of half the maximum; we use the midpoint of the resulting range, $b = 6$ Kbps, for our control design.

The rate of the Image Assembler (also the steady-state tile transmission rate) is computed using (1). CAMRIT uses 95% of the actual deadline to give some leeway to the transmission, and t_1 is estimated based on the nominal bandwidth and the tile size with the initial quality factor (68 in this example). The resultant $r = 0.34$ tile/sec. According to (10), in order to allow the bandwidth variation from 4 Kbps to 8 Kbps, the range for the inverse of compression ratio needs to be [6.38, 12.75]. Linearization is then performed in this range for $a(q)$ as shown in Figure 5. The slope of the linear regression is $g = 341.34$. The linear regression fits well (with an $R^2 = 94.87\%$) with the original function in this operation region.

IV. CONTROL DESIGN AND ANALYSIS

We now apply linear control theory to design the controller based on the controlled system model described in Section III. The z-transform of the controlled system model (9) is:

$$L(z) = z^{-1}L(z) + Az^{-1}U(z) + Cz^{-2}U(z) + \frac{Dz}{z-1} \quad (11)$$

A block diagram of the closed-loop system is shown in Figure 6. The system has two inputs: the set point of the tile buffer level and a disturbance input $\frac{Dz}{z-1}$ that represents the dequeuing of tiles from the tile buffer by the Image Assembler.

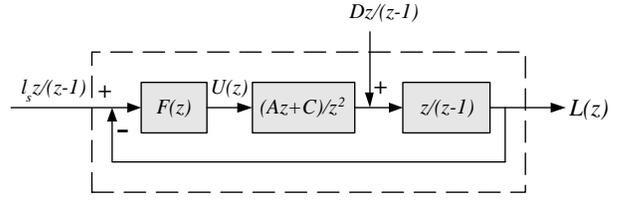


Fig. 6. Block Diagram for the Closed-Loop System

Letting $F(z)$ be the transfer function of the controller, we can derive the closed-loop transfer function in response to the reference input and disturbance, respectively:

$$H_s(z) = \frac{(Az + C)F(z)}{(z - 1)z + (Az + C)F(z)}$$

$$H_d(z) = \frac{z^2}{(z - 1)z + (Az + C)F(z)} \quad (12)$$

Therefore, the close-loop response to both inputs is

$$L(z) = H_s(z) \frac{z}{z-1} l_s + H_d(z) \frac{z}{z-1} D \quad (13)$$

To achieve stability and zero steady state error, we design a *Proportional-Integral (PI)* controller for our system:

$$F(z) = \frac{K_1(z - K_2)}{z - 1} \quad (14)$$

The time-domain form of (14) is:

$$u(k) = u(k-1) + K_1 e(k) - K_1 K_2 e(k-1) \quad (15)$$

where K_1 and K_2 are control parameters that can be analytically tuned to guarantee system stability and zero steady state error using standard control design methods.

We first apply the control design to our example application integrated with the CAMRIT framework. The sampling period is $T=10$ sec. The TCP buffer size is $B = 4$ KB. The period of the Tile Sender task is set to 2.67 sec to fully utilize network bandwidth. The other parameters (including g) are the same as for the example given in Section III-B. From (5), the control delay in the k^{th} sampling period is $T_d = 4 + q(k-1)/27.31$ sec. For example, the control delay is 5.8 sec when $q(k-1)=50$. Compared to a sampling period of 10 sec, the control delay clearly plays a significant role in the system dynamics. From (9), the parameters of the controlled system model are $A=81.922$; $C=54.614$; $D=-3.420$.

Using the Root-Locus method, we select our control parameters as $K_1=0.0068$ and $K_2=0.9$. The corresponding closed-loop poles are $0.278 \pm 0.547i$ and 0.887 . Since all the poles are in the unit circle, the system is stable. From the final value theorem [11], we can prove that the closed-loop system achieves zero steady state error. That is, the tile buffer level will achieve the set point in steady state: $\lim_{k \rightarrow \infty} l(k) = l_s$. If the set point is set to $l_s \geq 1$, the tile buffer will remain non-empty in steady state, and hence the image transmission deadline will be met. Furthermore, by substituting different bandwidths into the system model, we can prove that the

system can maintain stability and zero steady-state error with the same control parameters as long as the network bandwidth remains within the range [4Kbps, 8Kbps]. A detailed analysis is not given here due to space limitations: interested readers are referred to a standard control textbook [11].

In summary, pseudo code for the control algorithm implemented in CAMRIT is as follows:

```

Controller ( $l_s, K_1, K_2$ ) {
   $l$  = current tile buffer level;
   $e = l_s - l$ ;
   $u = u + K_1 * e - K_1 * K_2 * e_{prev}$ ;
   $e_{prev} = e$ ;
   $q = 1/u$ ;
  /* enforce the constraints on acceptable quality factor */
  /* default range is [1,100] */
  if ( $q < q_{min}$ )  $q = q_{min}$ ;
  if ( $q > q_{max}$ )  $q = q_{max}$ ;
  UpdateQF( $q$ );
  /* updated  $q$  will be used by the Tile Compressor */
}

```

V. EXPERIMENTAL EVALUATION

This section describes the testbed and experiments used to evaluate the stability and performance of our feedback controlled adaptation in CAMRIT.

A. WSOA Scenario

The Weapons System Open Architecture (WSOA) [2] program had a primary objective to provide internet-like connectivity, over Link-16, between legacy embedded mission systems in fighter aircraft and off-board Command and Control (C2) systems. This capability was designed to support time-sensitive mission re-planning and redirection of attack nodes, as necessary based on situational events, even if a different mission was already underway.

The following high-level sequence of interactions between the C2 and fighter aircraft constitutes a representative WSOA scenario: 1) The C2 node receives information about a higher priority time critical target and requests a planning session with attack nodes by sending an alert; 2) Upon receiving an alert, a fighter aircraft begins downloading a Virtual Target Folder (VTF). The VTF contains several thumbnail-sized images, each representing a virtual target; 3) Once the fighter receives a folder, the pilot can select a thumbnail image in the folder via a graphical display; 4) A request is then made to the C2 for a larger version of the selected image. The experiments presented in this paper emulate step 4, which is the most time critical part of the application.

B. Experimental Platform

Figure 7 shows our experimental configuration, which consists of two machines each running RedHat Linux 9.0 with the 2.4.20 kernel. The C2 aircraft was simulated using a more powerful machine - with a 2.53GHz Pentium 4 CPU - while the fighter aircraft was simulated using a 400MHz Pentium 2 machine. The following software was used to perform the experiments:

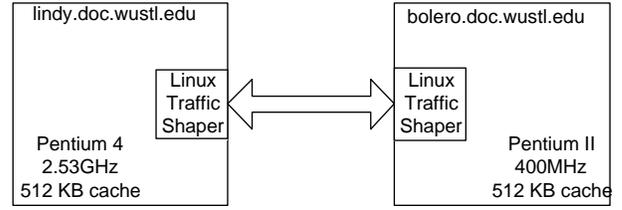


Fig. 7. Experimental Configuration

- ACE 5.3.5 + TAO 1.3.5 : TAO is a widely used open-source real-time CORBA standard object request broker [5]. TAO also provides a Real Time Event Channel [12] that is integrated with the Kokyu dispatching and scheduling framework [8]. This integrated middleware framework allow us to (re)schedule rates of invocation of application components, while maintaining deadline-feasible scheduling of critical operations.
- ImageMagick++ 5.5.7 : We used this library to compress and decompress images.
- Shaper 1.3 for Linux : Shaper is a linux script for traffic shaping. It allows us to specify the maximum bandwidth for network connection between two hosts.

We used Shaper to control the bandwidth between the two machines, *i.e.*, to simulate the performance of a Link-16 or other bandwidth-constrained unpredictable network over an underlying Ethernet connection. We set the range of bandwidth allowed by the traffic shapers to approximate the effective bandwidth of a plausible Link-16 configuration, *e.g.*, with a maximum network capacity of 28.8 Kbps [10], divided between the client and server. Taking into account the slotted nature of Link-16 communication channels and other Link-16 parameters, and the characteristics of the traffic shaper we used, we chose a maximum bandwidth of 8 Kbps for our experiments.

C. Number of Tiles

The CAMRIT API allows the user to specify the number of tiles into which each image was divided. The choice of this parameter involves a tradeoff between the control granularity and the network overhead. With more tiles is an image, the granularity of control is finer because there are more opportunities for the controller to adjust to network variations.

However, every tile contains a JPEG header which includes a non-trivial amount of information, *e.g.*, quantization and huffman tables. As the number of tiles increases, the overhead to send the JPEG and network headers also increases. To understand the relationship between header overhead and the number of tiles, we used a sample image of 640×640 pixels shown in Figure 3, and calculated the achieved compression ratio after tiling under different quality factors. The result of this study is shown in Figure 8.

From that study, we determined that the increase in compression ratio when going from 16 to 32 tiles was about 0.4%; from 32 to 64 tiles, it increased about 0.8%; from 64 to 128

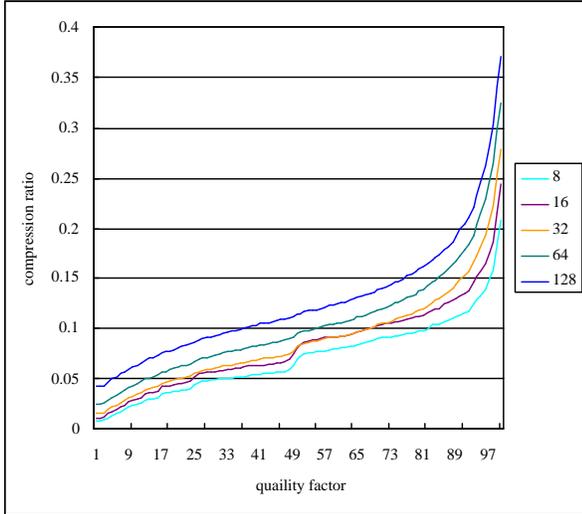


Fig. 8. Number of Tiles, Quality Factor and Compression Ratio

tiles, it increased about 1.6%. In other words, if we chose 64 tiles over 16 tiles, the overhead would cost about a little more than 1 second of transmission time; in the case of 128 tiles versus 16 tiles, the overhead would cost about 3 seconds of transmission time. In general, this study indicates that the overhead of tiling is acceptable when an image is divided into tens of tiles. We settled for 64 tiles for our experiments, to achieve a reasonable balance between control granularity and overhead.

D. Experimental Parameters

Our experiments used the same parameters as the examples in sections III-B and IV. To test CAMRIT's ability to handle different images, our experiments used two other aerial images than the one shown in Figure 3, whose profile was used to tune the control parameters. These two images are called Image 1 and Image 2 respectively.

The set point for the tile buffer level was $l_s = 5$ in our experiments. Note that there is a tradeoff in the choice of the set point. If the set point is too high, the quality factor for tiles transmitted in the first several sampling periods will be unnecessarily low because system has to fill an initially empty buffer with more tiles (with lower quality factors) before it reaches a steady state. On the other hand, if the set point is too low, a fluctuation in the network bandwidth may cause the buffer level drop to zero.

E. Experimental Results

CAMRIT uses (10) to calculate $q(0)$ based on its deadline, the nominal bandwidth (6 Kbps), and the profiled image quality function shown in Figure 4. The resulting initial quality factor is $q(0) = 68$ in all of the following experiments. Note that while $q(0)$ provides a reasonable initial value for the control input, that initial value is usually not correct for meeting the deadline because the actual bandwidth may differ from the nominal one.

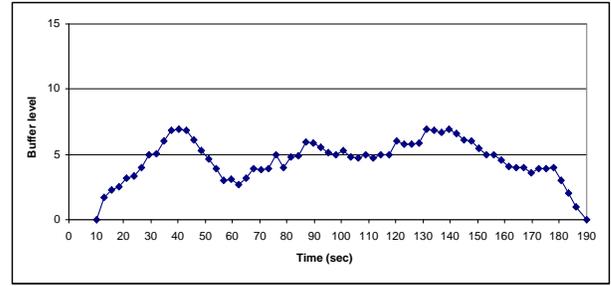


Fig. 9. Tile Buffer Levels During Typical Transmission of Image 1

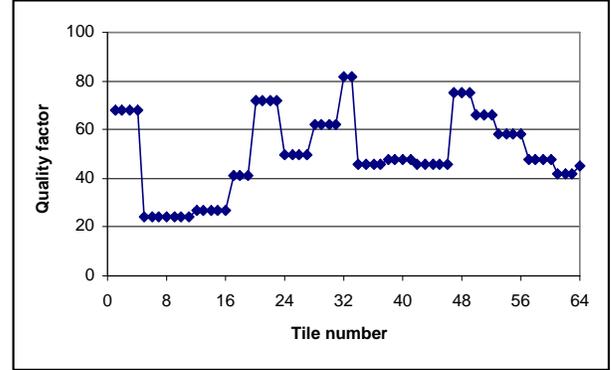


Fig. 10. Quality Factors During Typical Transmission of Image 1

The tile buffer level and quality factors during a typical transmission of Image 1 over a 6 Kbps network are shown in Figures 9 and 10, respectively. The buffer level is recorded by the Image Assembler before everytime it attempts to dequeue a tile. Time 0 in Figure 9 represents the time instant when the image request is sent to the server. The tile buffer is initially empty until the first tile is inserted at around 11 sec. This 11 sec delay includes the time it takes CAMRIT to send the image request to the server, divide the image into tiles on the server, and transmitting the first tile. Since the buffer level is low initially, CAMRIT reduces the quality factor from 68 to about 20 so that the buffer level rises to 5 tiles (the set point) in about 20 sec. The buffer level remains close to 5 tiles until the last image is transmitted to the client near the end of the run. The transmission of the whole image is completed at time 190 sec. This is consistent with our expectation because 190 sec (95% of the deadline) is used to compute the rate of the Image Assembler. Both tile buffer level and quality factor have some oscillation due to system noise. For example, the sizes of different tiles may be different (corresponding to different g values in our model) even if they are compressed using a same quality factor. However, despite the noise the tile buffer is always above 2.5 throughout the transmission. This is important because CAMRIT can guarantee an image transmission deadline is met as long as the tile buffer always contains at least one tile.

The primary goal of CAMRIT is to meet image transmission deadlines. Figure 11 shows the transmission delay of

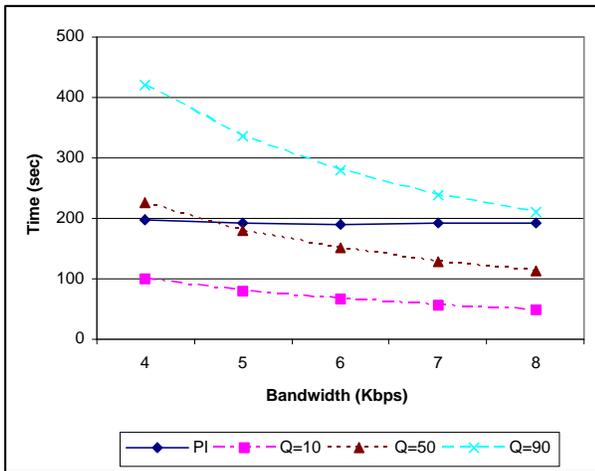


Fig. 11. Transmission Delay under Different Network Bandwidth.

Image 1 under different bandwidths. The transmission delay of CAMRIT (with the feedback loop) is measured through experiments. Each data point of CAMRIT in Figure 11 is the mean of 10 repeated runs. The standard deviation of each data point is within 2.62 sec. The transmission delay results of Image2 are not shown because they are almost identical to those of Image 1. For comparison purposes, we also plot the estimated transmission delays for Image 1 when a fixed quality factor (10, 50, or 90) is used in each run. The transmission delay for an image with a fixed quality factor is estimated by dividing its total (compressed) tile size by the actual network bandwidth¹.

We can see that the transmission delays for images with fixed quality factor vary significantly as the network bandwidth changes. This result confirms the difficulty in selecting a proper quality factor *a priori* when the network bandwidth is unpredictable. A chosen quality factor may be unnecessarily low when transmission completes much earlier than the deadline, or too high causing a deadline miss.

In contrast, the transmission delay under CAMRIT remains close to 190 sec (95% of the original deadline) as the network bandwidth varies from 4 Kbps to 8 Kbps (a 100% variation), and every run meets the deadline of 200 sec. The robust real-time performance is attributed to the feedback control loop that effectively maintains the desired buffer level despite the variation in network bandwidth.

The secondary goal of CAMRIT is to improve the image quality. CAMRIT accomplishes this goal by 1) fully utilizing the network bandwidth and 2) completing the transmission of an image close to the deadline (as shown in Figure 11). The combination of both properties means that CAMRIT sends close-to-maximum amounts of data for a requested image, which generally corresponds to a higher image quality.

Figure 12 shows the average quality factors of both images when they are transmitted by CAMRIT under different net-

¹This estimation is slightly lower than the actual delay because it ignores the overhead of protocol headers.

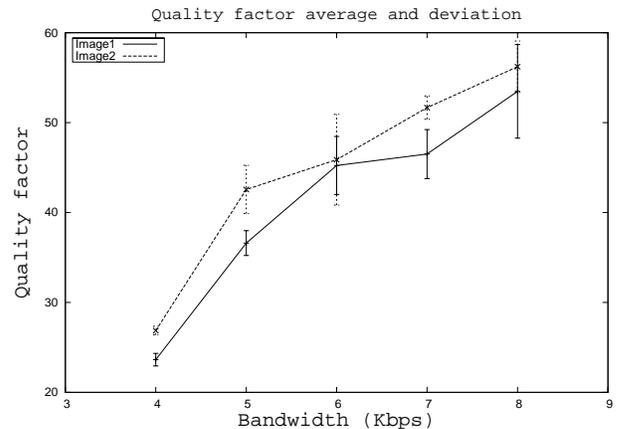


Fig. 12. Average Quality Factor under Different Bandwidth

work bandwidths. Each data point is the mean of 10 repeated runs. The standard deviations are also shown. With CAMRIT the average quality factor improves as more network bandwidth becomes available. This result determines that CAMRIT can automatically adapt to network bandwidth variations by adjusting the quality factor.

VI. RELATED WORK

CAMRIT was originally motivated by the WSOA program [2]. The goal of the WSOA program was to provide a CORBA-based real-time image download capability between aircraft over Link-16, to support collaborative mission re-planning. The WSOA program defined an approach for applying adaptive resource management technologies, to ensure timely exchange and processing of mission critical information (*e.g.*, images). The WSOA resource management approach used heuristics for adaptation that do not provide *a priori* performance analysis, but are evaluated empirically [13]. In sharp contrast, CAMRIT was modeled and designed from the start based on a rigorous control theoretic approach. Therefore, CAMRIT can provide the kind of robust and *analytic* performance guarantees that are crucial in mission-critical real-time systems.

Control theoretic approaches have been applied to a number of computing and networking systems. A survey of feedback performance control for software is presented in [14]. A number of feedback-based real-time processor scheduling algorithms (*e.g.*, [15] [16] [17]) have also been presented in the literature. These algorithms only controlled the allocation of the computing resource on a single node, and do not address transmission delays in distributed systems. Although feedback control real-time scheduling has been extended to handle distributed systems [18] [19], communication delays are not the focus of existing algorithms. Control-theoretic performance management in Internet servers (*e.g.*, web servers [20], e-mail servers [21], and web caches [22]) have also received significant attention recently. These techniques were also only concerned with the performance of server endsystems instead of transmission delays. Control theory has also been applied to

design and analyze network routers (e.g., [23] [24]). CAMRIT is different from these network solutions in that it aims to support real-time image transmission over existing networks through adaptation at the endsystems.

Li and Nahrstedt developed Agilos, a distributed visual tracking system based on control-theoretic adaptation [25]. Agilos embodied a distributed feedback loop that achieved desired image transmission rates through several adaptation mechanisms including image compression. However, Agilos did not control the transmission delay of an image. Moreover, the effect of control delay on the dynamics of distributed systems was not modeled in that project.

VII. CONCLUSIONS

In this paper, we have presented the design, modeling, and analysis of CAMRIT based on a control theoretic approach. A key contribution of this work is an analytic model that captures the dynamics of a moderately complex distributed middleware architecture. CAMRIT has been successfully implemented as a CORBA-based middleware service atop the TAO real-time ORB. Our experiments on a representative testbed demonstrate that CAMRIT can provide robust feedback control of image transmission delays across a range of available network bandwidth, by automatically adjusting image tile quality factors.

A potential extension to this work is to apply a wider range of adaptive control techniques [26] to further improve the robustness of the system under different degrees and kinds of uncertainty. Another important direction of future work is to integrate CAMRIT with feedback control real-time task scheduling [27] in an end-to-end performance control middleware framework for distributed embedded systems.

REFERENCES

- [1] G. J. Pottie and W. J. Kaiser, "Wireless integrated network sensors," *Commun. ACM*, vol. 43, no. 5, pp. 51–58, 2000.
- [2] D. Corman, "WSOA-Weapon Systems Open Architecture Demonstration-Using Emerging Open System Architecture Standards to Enable Innovative Techniques for Time Critical Target (TCT) Prosecution," in *Proceedings of the 20th IEEE/AIAA Digital Avionics Systems Conference (DASC)*, Oct. 2001.
- [3] Rockwell Collins, "JTIDS: Joint Tactical Information Distribution System." www.rockwellcollins.com/ecat/gs/JTIDS.html.
- [4] J. Zhao and R. Govindan, "Understanding packet delivery performance in dense wireless sensor networks," in *Proceedings of the first international conference on Embedded networked sensor systems*, pp. 1–13, ACM Press, 2003.
- [5] Center for Distributed Object Computing, "The ACE ORB (TAO)." www.cs.wustl.edu/~schmidt/TAO.html, Washington University.
- [6] G. K. Wallace, "The jpeg still image compression standard," *Communications of the ACM*, vol. 34, pp. 30–44, Apr. 1991.
- [7] C. Liu and J. Layland, "Scheduling Algorithms for Multiprogramming in a Hard-Real-Time Environment," *JACM*, vol. 20, pp. 46–61, Jan. 1973.
- [8] C. Gill, D. C. Schmidt, and R. Cytron, "Multi-Paradigm Scheduling for Distributed Real-Time Embedded Computing," *IEEE Proceedings, Special Issue on Modeling and Design of Embedded Software*, vol. 91, Jan. 2003.
- [9] T. H. Harrison, D. L. Levine, and D. C. Schmidt, "The Design and Performance of a Real-time CORBA Event Service," in *Proceedings of OOPSLA '97*, (Atlanta, GA), pp. 184–199, ACM, Oct. 1997.
- [10] W. J. Wilson, "Applying layering principles to legacy systems: Link 16 as a case study," in *IEEE International Military Communications Conference (MILCOM)*, 2001.
- [11] G. F. Franklin, J. D. Powell, and M. Workman, *Digital Control of Dynamic Systems*, 3rd edition. Addison-Wesley, 1997.
- [12] C. O’Ryan, D. C. Schmidt, and J. R. Noseworthy, "Patterns and Performance of a CORBA Event Service for Large-scale Distributed Interactive Simulations," *International Journal of Computer Systems Science and Engineering*, vol. 17, Mar. 2002.
- [13] C. D. Gill, J. M. Gossett, D. Corman, J. P. Loyall, R. E. Schantz, M. Atighetchi, and D. C. Schmidt, "Integrated adaptive qos management in middleware: An empirical case study," in *Submitted to the 10th Real-time Technology and Application Symposium (RTAS '04), Embedded Applications Track*, (Toronto, CA), IEEE, May 2004.
- [14] T. Abdelzaher, J. Stankovic, C. Lu, R. Zhang, and Y. Lu, "Feedback performance control in software services," *IEEE Control Systems*, vol. 23, June 2003.
- [15] L. Abeni, L. Palopoli, G. Lipari, and J. Walpole, "Analysis of a reservation-based feedback scheduler," in *IEEE Real-Time Systems Symposium*, Dec. 2002.
- [16] C. Lu, J. A. Stankovic, G. Tao, and S. H. Son, "Feedback Control Real-Time Scheduling: Framework, Modeling, and Algorithms," *Real-Time Systems Journal*, vol. 23, pp. 85–126, July 2002.
- [17] D. C. Steere, A. Goel, J. Gruenberg, D. McNamee, C. Pu, and J. Walpole, "A feedback-driven proportion allocator for real-rate scheduling," in *Operating Systems Design and Implementation*, pp. 145–158, 1999.
- [18] C. Lu, X. Wang, and X. Koutsoukos, "End-to-end utilization control in distributed real-time systems," in *International Conference on Distributed Computing Systems ICDCS 2004*, (Tokyo, Japan), Mar. 2004. To appear.
- [19] J. A. Stankovic, T. He, T. F. Abdelzaher, M. Marley, G. Tao, S. H. Son, and C. Lu, "Feedback Control Scheduling in Distributed Systems," in *The 22nd IEEE Real-Time Systems Symposium (RTSS '01)*, (London UK), Dec. 2001.
- [20] T. F. Abdelzaher, K. G. Shin, and N. Bhatti, "Performance guarantees for Web server end-systems: A control-theoretical approach," *IEEE Transactions on Parallel and Distributed Systems*, vol. 13, no. 1, pp. 80–96, 2002.
- [21] S. Parekh, N. Gandhi, J. Hellerstein, D. Tilbury, T. Jayram, and J. Bigus, "Using Control Theory to Achieve Service Level Objectives in Performance Management," in *Proceedings of the International Symposium on Integrated Network Management*, IEEE/IFIP, 2001.
- [22] Y. Lu, T. Abdelzaher, C. Lu, and G. Tao, "An adaptive control framework for qos guarantees and its application to differentiated caching services," in *Proc. Int. Conf. Quality of Service*, (Miami Beach, FL), pp. 23–32, May 2002.
- [23] C. Hollot, V. Misra, D. Towsley, and W. Gong, "A control theoretic analysis of red," in *Proceedings of IEEE INFOCOM 2001*, Apr. 2001.
- [24] N. Christin, J. Liebeherr, and T. Abdelzaher, "A quantitative assured forwarding service," in *Proceedings of IEEE INFOCOM 2002*, (New York, NY), June 2002. To appear.
- [25] B. Li and K. Nahrstedt, "A Control-based Middleware Framework for QoS Adaptations," *IEEE Journal on Selected Areas in Communications*, vol. 17, pp. 1632–1650, Sept. 1999.
- [26] T. A. Henzinger and S. Sastry, eds., *Hybrid Systems: Computation and Control - Lecture Notes in Computer Science*. New York, NY: Springer Verlag, 1998.
- [27] C. Lu, X. Wang, and C. Gill, "Feedback Control Real-Time Scheduling in ORB Middleware," in *Proceedings of the 9th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, (Washington, DC), IEEE, May 2003.