Multimedia
SYSTEMS

Volume 1 · 1993

Tactus: toolkit-level support for synchronized interactive multimedia

Roger B. Dannenberg*, Tom Neuendorffer, Joseph M. Newcomer, Dean Rubine, and David B. Anderson**

Information Technology Center. School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213, USA

Abstract. Tactus addresses problems of synchronizing and controlling various interactive continuous-time media. The Tactus system consists of two main parts. The first is a server that synchronizes the presentation of multiple media, including audio, video, graphics, and MIDI at a workstation. The second is a set of extensions to a graphical user interface toolkit to help compute and/or control temporal streams of information and deliver them to the Tactus Server. Temporal toolkit objects schedule computation events that generate media. Computation is scheduled in advance of real time to overcome system latency, and timestamps are used to allow accurate synchronization by the server in spite of computation and transmission delays. Tactus supports precomputing branches of media streams to minimize latency in interactive applications.

1 Introduction

Recently, many proposals have emerged for extending graphics systems to support multimedia applications with sound, animation, and video (Ripley 1989; Wayner 1991; Digital Equipment Corporation 1992; Yager 1992). Other research has been directed toward real-time transmission of multimedia data over networks (Anderson and Homsy 1991; Little and Ghafoor 1991; Rowe and Smith 1992) and standards for the representation and exchange of multimedia data (Newcomb et al. 1991). New capabilities for real-time interactive multimedia interfaces (Blattner and Dannenberg 1992) create new demands upon application programmers. In particular, programmers must manage concurrent processes that output continuous media. Timing, synchronization, and concurrency are among the new implementation problems.

---
* *Present address:* School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213, USA
** *e-mail addresses:* R.B. Dannenberg: dannenberg@cs.cmu.edu
T. Neuendorffer: tpn+@andrew.cmu.edu
J.M. Newcomer: newcomer@cs.cmu.edu
D. Rubine: dandb+@andrew.cmu.edu
D.B. Anderson: dba+@cs.cmu.edu
Correspondence to: Roger B. Dannenberg, at his present address

Traditionally, object-oriented graphical interface toolkits have presented a high-level programming interface to the application programmer, hiding many details of underlying graphics systems such as X or Display Postscript. However, timing is usually overlooked in these systems. Programmers usually add animation effects by ad-hoc extensions, and synchronization at the level of milliseconds needed for lip-sync, smooth animation, and sound effects is not generally possible.

We have extended an existing toolkit with new objects, abstractions, and programming techniques for interactive multimedia. We also implemented a synchronization server that supports our toolkit extensions. Intuitively, our synchronization server does for time what a graphics server does for (image) space. In our terminology, the application program is the client, which calls upon the server to synchronize and present data. We call the combined toolkit and server the Tactus system.

The Tactus system has a number of novel and interesting features. It works over networks with unpredictable latency, and it can maintain synchronization even when data underflows occur. The techniques are largely toolkit-independent, and the Tactus Server is entirely toolkit-independent. Data are computed ahead of real time to overcome latency problems, but the initial latency of a presentation is due only to computation and bandwidth limitations. Tactus is organized so that pre-existing graphical objects acquire real-time synchronizing behavior without changes to the existing code. The Tactus system also offers a new mechanism called cuts, which allows for user interaction by selecting among precomputed media with very low latency.

1.1 Assumptions

Before describing Tactus, we will present some assumptions and ideas upon which it is based. First, we are interested in distributed systems, and thus we assume that there will be significant transmission delays between servers and clients. Second, we assume that multimedia output will require the merging of multiple data streams; we want more than just "canned" video in a window. By data stream, we mean any set of timed updates to an output device. Data streams include video, audio, animation, text, images, and MIDI.

The assumption that delays will be present imposes limitations on the level of interaction we can expect. Network media servers may take seconds to begin presenting video even though the presentation, once started, is continuous. We intend to support applications where media start-up delays and latency due to computation of 10 to 1000 ms are tolerable. This includes such things as multimedia documents, presentations, video mail, and visualizations. It also includes more interactive systems such as hypermedia, browsers, and instructional systems where user actions determine what to view next. Although we rule out continuous feedback systems such as video games and artificial reality, we want to support rapidly altering the presentation at discrete choice points.

1.2. Principles

To deal with transmission and computation delays, it is necessary to start sending a data stream before it is required at the presentation site. Because of variance in computation, access, and transmission delays, it is also necessary to have a certain amount of buffering in the Tactus Server at the presentation site. When multiple streams are buffered, it is necessary to synchronize their output. With Tactus (see Fig. 1), all data streams are timestamped, either explicitly or implicitly, so that Tactus can determine when each component of a stream should be forwarded to a device for presentation. We assume a distributed time service that can provide client software with an accurate absolute time, with very little skew between machines (Kolstad 1990), although this assumption is not critical for most applications.

Fig. 1. The Tactus system. Clients send timestamped data (*heavy lines*) to the server ahead of real time. Data are buffered and then delivered to various presentations devices. Some presentation devices (e.g., MIDI as shown here) may accept data early and provide further buffering and more accurate timing than can be provided by the Tactus Server. The clock on the left shows logical time as seen by the client, while the clock on the right shows real time as seen by the Tactus Server

Multiple presentations may be buffered at the Tactus Server. At any time, one is being presented while the others are potential responses to user choices. This avoids the latency of transmitting a presentation over the network after the user makes a choice. Transition points are marked so that smooth cuts are possible (see Sect. 5).

Input is handled in mirror image to output. Input events are timestamped so that they can be related back to the output that was taking place at the time of the input. Since applications compute ahead of real time, it is up to the application to deal with the skew between input and output. For example, output can be "rewound" to an indicated stop point, or input might be applied to future output only.

The task of synchronizing output in a distributed environment is simplified by pre-computing or pre-transmitting data streams and timestamping them. Without additional support, however, this would complicate the work of the application, which then must compute data streams in advance of real time. One way to reduce this problem is to schedule application activity by a clock that is ahead of real time. A good analogy is that if you set your watch ahead by 5 min, you are more likely to show up on time for meetings.

In summary, the three most important principles of Tactus are (1) compute data streams in advance of real (presentation) time, (2) use a server at the presentation site to buffer and synchronize data streams, and (3) buffer all possible responses to user choices to minimize response times. Buffering data at the presentation site can greatly increase the timing accuracy with which data are presented.

1.3 Previous work

Few of these principles are original, but their integration and application are new. Tactus was inspired by Anderson and Kuivila's work on event buffering for computer music systems (1986, 1990). This work is in turn related to discrete-event simulation. Later, Anderson et al. (1990) applied these ideas to distributed multimedia, but not to interface toolkits. Rowe and Smith (1992) developed a system similar to Tactus. Their system implemented sophisticated network protocols for continuous media, but their client toolkit does not offer the features implemented in Tactus. Active objects (see Sect. 3.1) have long been used for animation (Kahn 1979) and music (Cointe and Rodet 1984) systems, but have only recently gained attention in multimedia circles (Gibbs 1991). The Cognitive Coprocessor (Robertson et al. 1989) manages computational latency by adapting animation rate and detail according to principles of human perception. This approach could be used in conjunction with the Tactus architecture.

To our knowledge, we are the first to extend an object-oriented application toolkit with support for managing latency through precomputation and event buffering. CD-ROM-based video systems have used buffering of images at choice points to allow for seek time. Our work focuses more on the implications of all these techniques for application toolkits.

Recently, many commercial multimedia systems have been introduced, including Apple's Quicktime (Wayner 1991), IBM's MMPM/2 (IBM 1992), Microsoft's MPC (Yager 1992), and Dec's XMedia (Digital Equipment Corp. 1992). These systems emphasize storage, playback, and scalability. HyTime (Newcomb et al. 1991) provides a standard representation for hypermedia, but no implementation is specified. These systems could benefit from the synchronization and latency management techniques we propose, and our work sug-

gests how a graphical interface toolkit might be extended to take advantage of commercial multimedia software.

Although this area is relatively new, efforts to establish standards are already underway. The Interactive Multimedia Association has published an RFT (IMA 1992) for multimedia system services that sketches out an architecture for interactive, distributed multimedia and gives some indication of how synchronization and timing services like those in Tactus may fit into future industry standards.

2 The Tactus Server

The Tactus Server provides buffering, resource management, timing, and synchronization to its clients. It can be thought of as analogous to a display server, such as X11, with two important differences. First, the Tactus Server is a real-time server; all data are buffered until the appropriate output time. This supports the smooth uninterrupted flow of continuous-time media to hardware presentation devices. Input data are also timestamped. Second, the Tactus Server manages multiple media, including audio, video, and MIDI in addition to graphics. It is the job of the Tactus Server to present each stream according to a specified time offset.

An important point here is that all the latencies in the system are handled by the Tactus Server in a fashion that is largely transparent to the application. The server handles this adaptively, being able to compensate for transient worst-case delays and to take advantage of any transient improvements. A drawback is that the server must be extended to handle new media and devices, but in practice, only a small amount of code is needed to support a new device. For the most part, Tactus merely forwards data to indicated devices at the correct time.

The server is responsible for the timely delivery of media to devices, so it places certain real-time requirements on the operating system. First, the server should have an accurate source of time. Not only should the server clock be synchronized with the client(s), but corrections for drift should not impact multimedia presentations. For example, if the local clock is set ahead by 100 ms, a noticeable jump may occur in the presentation. Therefore, the local clock should advance smoothly with respect to allowable media jitter. It should be synchronized to the client within a small fraction of the time that media is buffered in the server. In the future, low-cost atomic clocks (Creedy 1993) may eliminate clock drift and synchronization problems.

A second requirement for the operating system is that the server must wake up at times indicated by timestamps and deliver output to devices with low latency. An interesting property of the server is that wakeup times are aperiodic, but known in advance, offering the possibility of online pre-scheduling. The upper bound on latency for some media, such as MIDI, is in the millisecond range, which may not be possible in some systems. An alternative used by the Tactus Server is to forward data to devices ahead of time (just as clients forward data to the server ahead of time) and let the device driver provide accurately timed output. Device drivers can often provide more
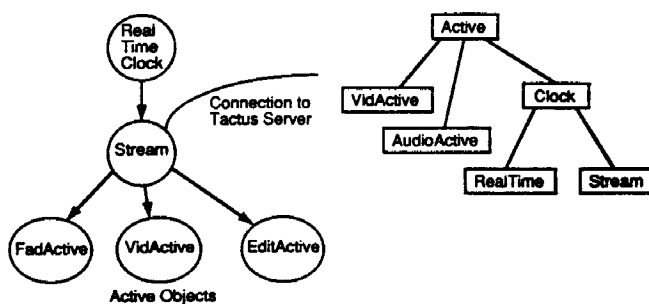

Active Objects

Fig. 2. A Clock Tree. Objects, including clocks, request a wake-up message from their parent in the clock tree. RealTime is at the root of the tree and interfaces with the operating system timing facilities. Stream is a subclass of Clock and manages connections to the Tactus Server. The leaves of the tree are subclasses of Active, which produce and control multimedia data (see Sect. 6). Kick messages flow in the direction of the *arrows*, while RequestKickAt messages are sent in the opposite direction

Fig. 3. Active Object Class Hierarchy. Active objects, which answer Kick (to perform the action) are specialized to create media presenters such as VidActive and AudioActive. Clock objects answer RequestKickAt (to schedule a kick) and send Kick messages to their children at the appropriate times. RealTime and Stream are specializations of Clock

accurate timing by processing at the interrupt level, by using a co-processor, by using DMA, or by running within the operating system kernel. However, this approach adds complexity and reduces the fine-grain control of devices, so it should be used only when necessary. A hybrid approach is sometimes possible where Tactus sends control information only. Devices then access precomputed media data directly from storage.

If the final timing is left to device drivers, there must be a feedback path to notify the Tactus Server when timing problems arise. For example, if a graphics pipeline falls out of real time, the Tactus Server must be informed so that other media output can be delayed as well. More details on device control will follow in Sect. 4.4.

3 The Tactus Toolkit Extensions

The Tactus Toolkit Extensions augment a graphical user interface toolkit such as ATK (Palay et al. 1988) or Interviews (Linton et al. 1989). The Tactus Server could be used without the toolkit extensions, but this would require the clients to compute data in advance of real time, implement various protocols (described in Sect. 4), and interleave computation for various streams. The toolkit simplifies these programming tasks. Our toolkit extensions to ATK include clock objects for scheduling and dispatching messages, active objects that receive wake-up messages and compute media, and stream objects that manage Tactus Server connections and timestamping (see Fig. 2). The class hierarchy is illustrated in Fig. 3.

3.1 Active objects

Active objects form the base class for all objects that handle real-time events and manage continuous time media in the Tactus extensions. Each active object uses a clock object (set via the UseClock method) to tell time and to request wake-up calls. The RequestKick method schedules the active object to be awakened at some future time by sending Request-KickAt to the object's clock. The object's Kick method will then be called by the object's clock at the requested time.

Active objects are intended to take the place of light-weight processes, and often perform tasks over extended periods of time. This is accomplished by having each execution of the Kick method request a future Kick. For example, we have a software video active object whose Kick method is roughly as follows:

```
Kick(self, time) {
    if (read(InputFile, Buffer,
            ImageSize)) {
        XPutImage(Buffer, ...);
        RequestKick(self, time
                + InterFrameTime);
    }
}
```

In this example, XPutImage is the standard X11 call to display a raster image, but it is implemented by a special library that adds a timestamp and sends the message to the Tactus Server. The Tactus Server is then responsible for forwarding the message to the real X server at the proper time. (Since we have not modified the X server to process timestamped events, we have found it helpful to run X at a higher priority in order to help ensure that the display is updated soon after the Tactus Server sends the graphics commands.) Note how RequestKick is used after each video frame to schedule the next one. During a video presentation, the computations of many other active objects may be interleaved with this one.

Rather than develop a special timed graphics library, we support the ordinary xlib calls with a substitute library txlib, which makes a connection to the Tactus Server rather than the X11 server. The txlib sends the timestamp along with each X11 message. The advantage of this approach is that pre-existing graphics code can be used without modification. This point will be further explained in Sects. 3.3 and 6.3. While txlib handles output for X11, similar libraries are used for audio and other media.

3.2 Clock objects

Clocks are a subclass of Active. Each clock object keeps track of all the active objects (children) that have attached themselves via the UseClock method. Since clocks are active objects, they too can be attached to other clocks. Clocks are useful not only for their wake-up service, but also because

they manage mappings from one time system to another. Mappings are linear transformations, meaning that a clock can shift and stretch time as seen by its children. When a change in the mapping of time occurs, children of the clock are notified (whether or not they are waiting for a Kick). We call the time seen by children of clocks *logical time*, as opposed to the *real time*. Logical time allows active objects to compute in "natural" time coordinates. Meanwhile, clocks can be adjusted to achieve "fast forward," "rewind," "pause," and "continue" effects.

The basic duty of a clock is to keep track of the times at which to kick child objects and to issue Kick messages at the appropriate times. A clock keeps a list of children sorted by their local requested kick times. The clock maps the earliest local wakeup time into parent time and issues its own RequestKick so that it will be kicked by its parent. When the clock is kicked, it maps from parent time to local time and kicks the appropriate child.

Clocks form a "clock tree" whose leaves are active objects, whose internal nodes are clocks, and whose root is a special subclass of clock called RealTime. A RealTime object serves as the true source of time for the entire clock tree. It should be noted that the clock tree is entirely independent of the graphical view tree typically found in graphical user interfaces (Cox 1987).

3.3 Stream objects

Stream objects are a subclass of Clock. In addition to scheduling and kicking child objects, stream objects communicate with the Tactus Server and establish timestamps for Tactus messages. Stream objects also schedule their children ahead of real time by the worst-case system delay called *Latency*, a number which is presently determined empirically.

Recall that the Tactus Server expects all messages to have timestamps which serve as the basis for synchronization of multiple media. It might seem logical to use the kick times of active objects, but because kick times are the composition of perhaps several mappings at different levels for the clock tree, the active object kick time may have no simple relationship to real time or to the kick times of other active objects.

Rather than use active object kick times, timestamps are based on the idealized real time of the kick, that is, the requested kick time mapped to real time. Before a kick, the clock tree is inactive. When the kick time arrives, a Kick message is propagated from the RealTime object through the tree to an active object at a leaf of the tree. On the path from root to leaf, exactly one stream object is kicked. The stream sets a globally accessible timestamp and stream identifier before propagating the Kick message. If the active object performs an output action, the output function called accesses the timestamp and stream identifier in order to compose a message for the Tactus Server. In this way, timestamps are implicitly added to client output.

Together, these classes and their specializations serve to insulate the application programmer from the detailed protocols necessary to send streams of data to the Tactus Server. The

For example. objects that formerly displayed text or images can now be called upon to deliver output synchronously with other media.

4 The Tactus system

As described in the introduction. the Tactus system consists of a Tactus Server and a set of extensions to an object-oriented toolkit. In this section. we will describe how the two work together.

4.1 Steady-state media delivery

Steady-state on the client side consists of active objects waiting for wake-up messages. At each wake-up. an object computes data such as a packet of audio or a frame of animation and sends it to the Tactus Server. The wake-up message is scheduled by a stream object that forces the computation to happen ahead of real-time. When no more computation is pending. the stream object computes when the next wake-up will occur and sends a null message to the Tactus Server with that timestamp. This tells the Server not to expect more messages until that time.

As discussed earlier. Tactus may deliver the data to the presentation device slightly ahead of time. relying on the hardware or device driver to delay the presentation until a given timestamp. For example. our MIDI driver maintains buffers of timestamped packets of MIDI data and outputs data at the designated time. In contrast. X11 (our "graphics device driver") has no buffering or timestamping capability yet. so Tactus provides all timing control for graphics. These differences are invisible to clients.

Time is used to regulate the flow of data from clients to Tactus. thus alleviating the need for explicit flow-control messages. The client simply produces "one second of data per second" and sends it to Tactus. When the client is behind. it computes as fast as possible in order to catch up. If the client falls too far behind. the Tactus Server buffers will underflow and a recovery mechanism must be invoked (see below).

4.2 Stream start-up

We anticipate that the worst-case delay from client to device will be quite large (perhaps seconds). This is too large to be acceptable for the normal stream start-up time. Tactus clients typically will start streams with the goal of delivering media to the user as soon as possible. Therefore. the stream object ad-

the starting stream (the low-wat
devices). When the low-water crit
delivering data to devices, and we
the phase transition. a message is
ing how much time elapsed in the
called Glitches.

4.3 Synchronization

Synchronization is achieved by a
Temporal data are thus synchror
to synchronizing directly to othei
Multiple streams may be groupe
tions. It is up to device drivers ei
fied times or to report timing failu
Timing failures in devices are trea
is described next.

4.4 Underflow

An underflow is caused by the s
data. More precisely. underflow
patch a data packet at time T. but
data at a time greater than T. Sir
a timestamp greater than T is des
all data for time T have arrived. I
tus to halt all media presentation
time T can be updated. but we bel
should be supported as well (An

If data have been sent ahead to
logical time in the device (settin
immediately pauses. If this is not
may be somewhat staggered in tin
all device output is stopped and a
This phase resembles the start-u
and the same steps are taken.

No immediate feedback to i
underflow (presumably. the clier
trying to catch up). When Tactus
a message to the client indicating
sentation was delayed. This infor
the total delay between computat
fault behavior is for the client to
latency; if the Tactus Server stops
the client holds off on computati

Generally, this protocol takes

of frames per second, when the logical time rate (playback speed) is increased.

4.5 Device drivers

Tactus has some interesting implications for presentation sub-systems and device drivers. Devices should start continuous media output at precisely stated times and should synchronize data to the system clock. Devices should also monitor the delivery of output and report exceptions back to Tactus.

A condition similar to underflow arises when a device fails to meet its expected timing requirements. For example, a CD-ROM player might fail to seek within an allocated time. In this case, the device should recognize that timely delivery of media has (or will) fail, and the Tactus Server should be notified. The server treats this condition as a form of underflow and enters the low-water phase. The running phase is reentered when the failing device reports that it is ready to continue, e.g., the seek completes. Another example would be a graphics accelerator that is unable to complete a frame of animation in time.

5 Cuts

Because of various latencies, multimedia systems are often unable to respond to input without obvious "glitches" where, for example, the video image is lost, digital audio pops, and graphics are partially redrawn. This usually happens because there is a time delay between taking down one stream and starting up another. These annoying artifacts could be hidden if the new stream could be started before the the old one is stopped. Tactus supports this model, and a switch from one stream to another is called a *cut*.

In Tactus terminology, a *cut* is made from a primary stream to a secondary stream. To minimize latency, cuts are performed by the Tactus Server on behalf of its client. The client requests a cut, but the request may or may not be honored, depending upon whether the secondary stream is ready to run.

There are two attributes that describe a cut (see Fig. 4). The first determines whether a cut may be taken at any point in time or only at certain time points, and the second describes whether the cut is made to the beginning of a secondary stream or to the current time.

Cuts must be anticipated by the application. Since the application runs ahead of the real presentation time, it will naturally come to choice points before the user has a chance to make a choice. For example, a driving simulator application will generate graphics or video for an intersection before knowing whether the user will say "turn left" or not. While computing ahead, the application will create a cut object to arbitrate between the current (primary) stream and a new "turn left" (secondary) stream.

Within the Tactus Server, the secondary stream will perform a normal stream start-up. If the user requests "turn left," the application (user-cut requests are processed by the application, not by the Tactus Server; this requires a round-trip message to the application, but keeps the input-processing model


Fig. 4. There are four (4) types of cuts. The top two cuts are restricted to discrete time points, whereas the lower two can take place at any time. The first and third cuts here cut to a stream already in progress while the second and fourth types cut to the beginning of a stream

uniform) sends a cut message with a timestamp to the Server. If the message arrives before the time indicated by its timestamp, and if the secondary stream is ready to run, then Tactus switches to the secondary stream at the designated time. A message is returned to the application indicating success or failure. If the cut was a success, then the objects generating the no-longer useful primary stream will be freed.

6 An example application

It is now time to see how clocks, streams, active objects, and the Tactus server work together to produce a synchronized multimedia presentation. We will describe an application we have actually built: a time-line editor for sequencing video and animation.

6.1 The editor

As far as this discussion is concerned, the function of the editor (see Fig. 5) is merely to produce a data structure consisting of a list of animations to run and video segments to show. We will call this data structure the *cue sheet*. An animation sequence represented by the editor consists of a file name, a starting frame, an ending frame, and a duration. An active object of class FadActive takes these parameters and generates a sequence of display updates showing the sequence of frames and some number of interpolated frames, depending upon the duration. Similarly, a video segment is represented by a starting frame, an ending frame, and a duration. An object of class VidActive generates control commands for a laser videodisc player (an all-digital video object has also been implemented) to generate the appropriate sequence of video frames.
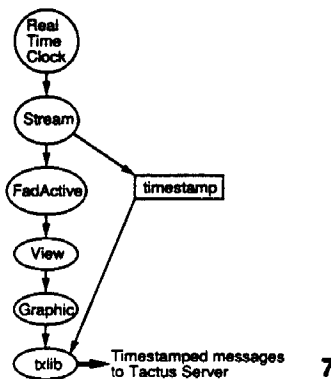
Fig. 5. The Editor Application. The top window is a time-line for animation, video, and music. The presentation below shows a video image and an animated graphical overlay, in this case highlighting the Lunar Excursion Module as it lifts off from the moon

Fig. 6. Timing. The *top line* shows times as specified in a cue sheet. The *middle* shows the timing of kick messages, and the *bottom* shows presentation times. The cue sheet is effectively shifted forward in time from zero to T_{start}. All kicks are advanced in time by the amount *Latency* so that output can be delivered to the server and buffered. The presentation then corresponds exactly to the time-shifted cue sheet. If underflow occurs, Glitches will become non-zero, and both the client computation and the server presentation will move later in time

Fig. 7. Creating timestamped messages. Scheduled computations are initiated by RealTime. The Stream computes a timestamp for later use by txlib. Control then passes to an Active object, which produces output via a View object, which in turn calls a Graphics object. X-window calls from *Graphic* objects are handled by txlib, which adds timestamps and sends messages to the Tactus Server

6.2 Active objects and the clock tree

The structure of the application was shown in Fig. 2. An editor creates three active objects: FadActive for graphical animation control, VidActive for video control, and EditActive for sequencing through the editor data structure. All three of these use a stream object as their clock, and the stream object uses RealTime for its clock.

6.3 Generating a timestamped message

Let us trace how the first message of a presentation is generated. When the user selects the Start menu item, the editor sets the stream to logical time 0. This happens at real time T_{start}. The editor also sends the EditActive object a message which causes it to request a kick at logical time T_{cue}, the time of the first cue on the editor's cue sheet. For now, let us also assume that the first item is well beyond the *Latency*. Thus, a kick request will be registered with the RealTime object at some time in the future, after which the Start command completes, and the user interface becomes idle.

To summarize so far, the editor has requested that the stream kick it to produce the first action at logical time T_{cue}.

The stream will actually kick the editor's active object ahead of time by *Latency*, a programmer-specified time advance (in the future, time advance may be estimated automatically on the basis of system and network loads). Also, the kick is delayed by *Glitches*, the total time by which presentation has been delayed in the low-water phase. This variable still has its initial value of 0.

Figure 6 illustrates timing in Tactus. The middle line of the figure shows the real time of kicks. At $T_{start} + T_{cue} -$ *Latency* + *Glitches*, the realtime clock will send a Kick message to the stream. Figure 7 shows the sequence of calls that lead from the realtime clock to a timestamped message. The stream object constructs a timestamp of $T_{start} + T_{cue}$ and kicks EditActive, which in turn sends a start message to the FadActive object. EditActive also schedules itself for the next cue in the cue sheet.

The FadActive object initializes its state and prepares to generate a sequence of animated drawings and calls upon a View object to update the screen with the first image of the animation and returns [this is a slightly simplified description; in reality, the FadActive does not know who sent the start message and cannot assume that the stream has properly com—

puted a timestamp. As a precaution, it schedules itself for a kick at the current logical time (according to its clock)]. The kick to EditActive returns, and a new kick to FadActive immediately follows. The View object calls upon its associated Graphic object to perform graphical operations, and the Graphic object makes standard calls to xlib, the X11 client interface library.

When the update is finished, the FadActive object schedules itself for its next update and returns. At this point, the FadActive object is scheduled to wake up on the next animation frame, and EditActive will wake up to send the next cue from the cue sheet. Whenever the stream reaches a state where it will actually suspend until there is another kick from RealTime, the stream computes the stream time at which it will wake up and sends that time to the Tactus Server. This tells the server how far ahead of realtime the computation has progressed. This information (or rather, the lack of it) is in turn used to detect underflow.

On the server side, messages arrive with timestamps and commands for X11 and a videodisc player. The messages are placed in a buffer until the dispatch time arrives. The dispatch time is:

$$T_{dispatch} = Timestamp + Glitches$$

where Glitches measures the total elapsed time in the low-water phase. The bottom line in Fig. 6 shows the server presentation time. Notice from this example, that the active objects can all operate in the relative time of the cue sheet; for example, presentations can start at logical time zero.

Also notice that because timestamps are passed directly from the stream to txlib, no modification to complex graphical objects (the view object in Fig. 7) is needed. Even the full ATK text editor runs under Tactus without modifications. If text were to be scrolled by an active object, the scrolling would be synchronized with other stream elements. Any X application, when linked with the Tactus library, should also run under Tactus, and it is possible to modify other toolkits to use Tactus synchronization.

6.4. Adding an animated cursor

Suppose we want to modify the timeline editor to display a vertical cursor that travels across the timeline during playback. The change is trivial since during playback, any graphics commands are automatically synchronized with other media, and the active object architecture supports interleaving the cursor computation with other output. The changes required are: Subclass EditActive to create a CursorActive object and add it to the clock tree as a child of the stream. The CursorActive object will automatically be notified when the stream starts and stops. When the stream starts, use the RequestKick and Kick messages to wake up every 50 ms or so and redraw the cursor according to the current logical time.

7 Current status

Tactus currently exists as a working prototype on Unix workstations; it supports synchronized digital video, digital audio, animation, and the full range of ATK graphics and text objects. The Tactus Server runs directly on the Mach 3.0 Microkernel (Accetta et al. 1986), and another version runs as an ordinary user process under AIX. Typically, both the client and server run on one machine using a choice of shared memory, Mach IPC, or Unix socket interfaces. We are studying performance issues and porting Tactus to Real-Time Mach (Tokuda et al. 1990) should provide even better predictability. Even though applications remain as Unix processes, the latency management of Tactus allows them to generate accurately timed output.

8 Summary and conclusions

Tactus is a system for synchronizing multiple digital media in a distributed environment where latency is an important factor. In Tactus:
- Data are computed ahead of real time.
- An object-oriented system is used to schedule computation and compute timestamps.
- The data are delivered to a synchronization server and buffered.
- The data are dispatched to presentation devices according to timestamps.
- Latency management is carefully hidden from application programs, which benefit from a simple-to-use interface.

Important considerations in the design of Tactus have been:
- The system should support interactivity by helping clients compute multimedia information.
- The server should handle synchronization at the point of presentation.
- Both discrete events and continuous media should be supported.
- Existing graphics libraries should be usable with a minimal amount of change.

Consequently, a large effort has gone into designing the client side of Tactus as well as the server. This has paid off by simplifying the task of writing multimedia applications.

There are problems Tactus does not address. In particular, we have attempted to mask rather than solve network problems. We assume that enough network bandwidth is available on average for multimedia presentations. The amount of precomputation and buffering can be tailored to actual network parameters. Even if communication is not a problem, Tactus can mask computational and storage latency.

Tactus also does not address issues of resource allocation. We assume that adequate bandwidth and processor resources are available on average. If this is not the case, no other approach will be successful. Ideally, Tactus should be used in conjunction with network protocols and operating systems that can make some resource guarantees. It is up to the application programmer to specify requirements or reserve resources. One-

area for future research is to examine how secondary streams for cuts should be scheduled when resources are limited.

Tactus has important implications for the design of multimedia systems. Few device drivers or devices support timestamped data, yet this seems to be a useful technique for solving current synchronization problems. For example. MIDI data should be transmitted within about 1 ms of the specified time. This may be difficult even for the Tactus Server. Thus, we adopt a multi-level approach. The Tactus Server buffers the high-latency (100–1000 ms) connection from the client to produce synchronized low-latency (e.g. 10 ms) data for devices. In the case of MIDI, Tactus dispatches data slightly ahead of real-time to enable the very low-latency (less than 1 ms) device driver to provide the final timing.

One could argue that if all device drivers supported timing, the Tactus Server would not be necessary. However, the Tactus Server performs other important functions. It performs local synchronization failure detection and recovery across multiple media and devices. (This could also be performed by the client, but only with many round-trip messages to a potentially overloaded client.) In the case of a remote client, the Tactus Server provides the logic and buffers for cuts. The Tactus Server also supports the synchronization of multiple distributed clients. Finally, the server is a logical place to provide for resource management.

Designers of network time protocols should keep in mind that abruptly setting clocks ahead or behind can cause problems for multimedia and other real-time systems. It seems that network time systems are designed to minimize clock skew at the expense of jitter caused by local clock adjustment. For multimedia synchronization, this jitter should also be minimized.

Tactus uses global clocks for flow control between client and server. Clients know when data will be output, so they send data just before the data are needed. No explicit flow control messages are necessary.

Media delivery subsystems should be designed with facilities for external control and synchronization, whereas most current designs are closed and self-regulating. For example, a video playback sub-system can be integrated with Tactus if it accepts absolute start times, has the potential to pause when other media underflow, and can report to Tactus when a pause due to underflow is about to occur. In contrast, the practice of synchronizing digital video to a local audio device makes it impossible to synchronize to external systems or other channels of audio.

Finally, Tactus raises some interesting and unaddressed scheduling problems (Dannenberg 1989) not solved by traditional periodic scheduling models. Tactus events are aperiodic, but the precise time and content of a Tactus event dispatch are known significantly in advance.

References

Accetta M, Baron R, Bolosky W, Golub D, Rashid R, Tevanian A, Young M. (1986) Mach: a new kernel foundation for UNIX development. In: Proceedings of Summer Usenix. Usenix

Anderson DP, Kuivila R (1986) Accurately timed generation of discrete musical events. Comput Mus J 10(3):48–56

Anderson DP, Kuivila R (1990) A system for computer music performance. ACM Trans Comput Syst 8(1):56–82

Anderson DP, Govindan R. Homsy G (1990) Abstractions for continuous media in a network window system. Tech Rep UCB/CSD 90/596, Computer Science Division (EECS), U.C. at Berkeley

Anderson DP. Homsy G (1991) A continuous media I/O server and its synchronization mechanism. Computer 24(10):51–57

Blattner MM, Dannenberg RB (ed) (1992) Multimedia interface design. ACM Press, New York

Cointe P. Rodet X (1984) Formes: an object and time oriented system for music composition and synthesis. In: 1984 ACM Symposium on LISP and Functional Programming. ACM, New York, pp 85–95

Cox BJ (1987) Object-oriented programming: an evolutionary approach. Addison-Wesley, Reading, Mass

Creedy S (1993). Time is right for atomic ticker by Westinghouse. Pittsburgh Post-Gazette, April 24, 1993, pp. 1–2

Dannenberg RB (1989). Real-time scheduling and computer accompaniment. In: Mathews, MV, Pierce RJ (eds) Current directions in computer music research. MIT Press, Cambridge, Mass, pp 225–262

Digital Equipment Corporation (1992) XMedia tools, version 1.1A. Software Product Description SPD 36.55.02

Gibbs S (1991). Composite multimedia and active objects. In: Paepcke A (ed) OOPSLA '91 Conference Proceedings, ACM/SIGPLAN. ACM Press, New York, pp 97–112

IBM (1992) The OS/2 multimedia advantage. IBM Corp

IMA (1992) RFT: multimedia system services version 2.0. IMA

Kahn K (1979) Director guide. Technical report, MIT AI Laboratory, Memo 482B. MIT Press, Cambridge, Mass

Kolstad R (1990) The network time protocol. UNIX Rev 8(12):58–61

Linton MA, Vlissides JM, Calder PR (1989) Composing user interfaces with interviews. Computer 22(2):8–22

Little TDC, Ghafoor AS (1991) Spatio-temporal composition of distributed multimedia objects for value-added networks. Computer 24(10):42–50

Newcomb SR, Kipp NA Newcomb VT (1991) The HYTIME multimedia/time-based document structuring language. Commun ACM 34(11):67–83

Palay AJ, Hansen M, Kazar M, Sherman M, Wadlow M, Neuendorffer T, Stern Z, Bader M, Peters T (1988) The Andrew Toolkit – an overview. In: Proceedings of the USENIX Technical Conference, Winter 1988. USENIX, pp 9–21

Ripley GD (1989) DVI – a digital multimedia technology. CACM 32(7):811–822

Robertson GG, Card SK, Mackinlay JD (1989) The cognitive coprocessor architecture. In: Proceedings of the ACM Symposium on User Interface Software and Technology. ACM Press, New York, pp 10–18

Rowe LA, Smith BC (1992) A continuous media player. In: Third International Workshop on Network and Operating System Support For Digital Audio And Video. IEEE Comput Commun Soc pp 334–344

Tokuda H, Nakajima T, Rao P (1990) Real-time mach: toward a predictable real time. In: Proceedings of the USENIX Mach Workshop. USENIX

Wayner P (1991) Inside Quicktime. Byte 16(12):189

Yager T (1992) The multimedia PC: high-powered sight and sound on your desk. Byte 17(2):217

DAVID B. ANDERSON has been a member of the research staff of the School of Computer Science at Carnegie Mellon University since 1988, where he is also a Ph.D. candidate. He received B.S. degrees in Mathematics and Computer Science from Brigham Young University in 1981. He has previously worked as an architect and manager of the Andrew Toolkit project and enjoys doing user-oriented systems work. His research interests include user interface toolkits, operating systems, and compound document architectures.

JOSEPH M. NEWCOMER received his PhD from CMU in 1975 in computer technology. Since then he has participated in compiler research and product development, personal computer software, MIDI software products, and realtime systems. He has numerous articles on programming methodology published in major microcomputer technical magazines, including one on techniques for debugging real-time embedded systems. He is currently a consultant and software developer in Pittsburgh, Pennsylvania. He participated in the design of Tactus in 1991.

DEAN RUBINE is currently a research faculty member at the School of Computer Science at Carnegie Mellon. His interests include continuous-time multimedia systems and applications, human-computer interaction techniques, user interfaces for non-programmers, audio digital signal processing, pattern recognition, real-time control systems, and computer music. Rubine holds B.S. and M.S. degrees from MIT and a Ph.D. from Carnegie Mellon, all in computer science. His Ph.D. thesis addresses the problems of training computers to understand human gestures. In addition to his work on the Tactus system, Rubine co-invented the VideoHarp, a new music instrument controller, researched algorithms for the analysis and synthesis of musical tones, implemented programming languages for the real-time control of synthesizers, and created a system for developing multimedia applications by direct manipulation. He is also an amateur musician, sound engineer, and video producer.

TOM NEUENDORFFER has been working on multimedia user interfaces since 1985. One of the primary developers of the Andrew Toolkit, he worked on many facets of the design and implementation of its embedded object architecture. He was also responsible for the Andrew Development Environment Workbench (ADEW), the first publicly distributed X-based graphical interface builder. Other accomplishments include the FAD animation editor, the GLO interface builder, and what was probably the first working electronic keyboard ever sent via electronic mail. Prior to 1985, he worked for the Information Science Department of the University of Pittsburgh, where he received an M.S. in 1980.

ROGER B. DANNENBERG is a Research Computer Scientist on the faculty of the Carnegie Mellon University School Computer Science Department. He received a Ph.D. from Carnegie Mellon in 1982 after receiving a B.A. from Rice University (1977) and an M.S. from Case-Western Reserve University (1979). His research interests include programming language design and implementation, and the application of computer science techniques to the generation, control, and composition of computer music. Dannenberg's current work in addition to Tactus includes research on music understanding, the automated accompaniment of live musicians, and the design and implementation of Nyquist, a very high-level, functional language for signal processing and control. He was co-director of the Piano Tutor Project which applied music understanding and expert system technology to music education. Dr. Dannenberg is a member of Phi Beta Kappa, Sigma Xi, Tau Beta Pi, Phi Mu Alpha, ACM, IEEE, and the International Computer Music Association.