

NASA Contractor Report 198498

Full Mesh Audio Conferencing Using the Point-to-Multipoint On-Board Switching Capability of ACTS

Mary L. Rivett and Zubin H. Sethna
NYMA, Inc.
Brook Park, Ohio

June 1996

Prepared for
Lewis Research Center
Under Contract NAS3-27186



National Aeronautics and
Space Administration

FULL MESH AUDIO CONFERENCING USING THE POINT-TO-MULTIPOINT ON-BOARD SWITCHING CAPABILITY OF ACTS

Mary L. Rivett
Senior Software Engineer
NYMA, Incorporated
Brook Park, Ohio

Zubin H. Sethna
Software Engineer
NYMA, Incorporated
Brook Park, Ohio

ABSTRACT

With the growing emphasis on the construction of a Global Information Infrastructure (GII), new satellite systems will be necessary to meet the resulting increased demand for communication services. These systems must be able to provide low-cost on-demand access, integrated digital services, full-mesh single hop operation and high transmission rates. As a result, satellite systems of the future will require revolutionary advances in bandwidth utilization and connectivity.

The Advanced Communications Technology Satellite (ACTS), developed by NASA and launched in September 1993, demonstrates several key technologies which will play a central role in the definition of the GII. These technologies include: 1) the use of the Ka frequency band, which opens up a new portion of the RF spectrum, 2) high-powered narrow hopping beams, that enable the use of smaller, less costly earth stations, and 3) on-board digital switching, which routes traffic to specific end locations. By incorporating these advanced technologies, ACTS not only provides on-demand, high data rate, digital communication services, but also stimulates the development of alternate ways of implementing existing services.

One existing communication service that lends itself to re-implementation using ACTS technologies is audio conferencing. Currently, audio conferencing systems have several shortcomings. First of all, satellite participants cannot be easily supported without undesirable side effects that may decrease the usefulness of the conference. In conferences where some but not all participants require satellite hops, the audio signals between participants have different delays. This causes voice overruns to occur, making conversation difficult. Secondly, in most conferences where all participants require satellite hops, the delays in audio signals between participants, though equal, are quite significant (i.e. 500 to 600 msec.) due to double hop paths. This situation results in lengthy pauses which make conversation unnatural. Finally, audio conferences must be scheduled through a bridge provider, causing the action of starting a conference to become less

spontaneous. In addition, bridges, unlike PABX systems, do not support a camp-on calling feature. Therefore, once in a conference, a user cannot do a "hook flash", get dial tone, and add additional parties at will, which is a completely natural operation.

The purpose of this paper is to describe an implementation of audio conferencing using the ACTS T1-VSAT network. In particular, this implementation evaluates the use of the on-board switching capability of the satellite as a viable alternative for providing the multipoint connectivity normally provided by terrestrial audio bridge equipment. The system that was implemented provides full mesh, full-duplex audio conferencing, with end-to-end voice paths between all participants requiring only a single hop (i.e. 250 msec. delay). Moreover, it addresses the lack of spontaneity in current systems by allowing a user to easily start a conference from any standard telephone handset connected to an ACTS earth station, and quickly add new members to the conference at any time using the "hook flash" capability. No prior scheduling of resources is required and there is no central point of control, thereby providing the user with the spontaneity desired in audio conference control.

The outcome of this effort has shown that satellite on-board switching technology, as demonstrated by ACTS, is indeed suitable for implementing audio conferencing, and may be applicable to similar communication services with multipoint requirements, such as video conferencing. In the future, as the role of satellite systems in the GII is defined, systems with on-board switching technology will prove useful in providing audio conferencing and other innovative services.

INTRODUCTION

With the increasing growth of satellite systems with advanced technologies, such as ACTS, the need arises to demonstrate the usefulness of these technologies for practical real-world applications [2]. An audio conferencing system using the ACTS T1-VSAT network was developed for just this purpose.

Work began on the ACTS audio conferencing system in the fall of 1993 with requirements analysis and preliminary hardware investigations. Implementation and informal testing were completed early in the fall of 1994, with the first three-party satellite conference call taking place on September 23. After a period of formal testing, the system became fully operational in December 1994.

This paper describes the development of the ACTS audio conferencing system, focusing primarily on the system design. An overview of the ACTS T1-VSAT network architecture and operation is provided first, followed by a discussion of audio conferencing state-of-practice. The main body of the paper identifies the goals of the ACTS audio conferencing system, examines key design issues, and provides examples to illustrate its use. It is assumed the reader is familiar with the basic principles of telecommunications and communications satellites.

ACTS T1-VSAT NETWORK OVERVIEW

The ACTS T1-VSAT network is designed to utilize the capabilities of the satellite to provide highly flexible, full mesh communications between T1-VSAT earth stations. This network is comprised of the following components: 1) the satellite, 2) the master control station and 3) the T1-VSAT earth stations. These components are briefly described in this section; however, more detailed descriptions can be found in [3], Section 3 and [1], Sections 3 and 4.

The Satellite

The satellite is equipped with several advanced features.

One of these is a very high-gain, multiple-hopping-beam antenna system, which enables the use of smaller, less costly, 1.2-meter aperture earth stations. The dynamic capabilities of this system, in conjunction with the use of a time-division multiple-access (TDMA) channel allocation scheme, allow capacity to be allocated to users on-demand. In addition, the use of narrow spot beams allows reuse of the allocated frequency spectrum.

Another advanced feature of the satellite is an on-board baseband processor which provides switching at the individual circuit level. On-board switching permits earth stations to be interconnected without transmitting messages to/from a central hub, thereby eliminating a double hop. Both point-to-point circuits (simplex or duplex) and point-to-

multipoint circuits (simplex only) are supported.

Lastly, the satellite operates in the Ka-band (30 GHz uplink / 20 GHz downlink), opening up a new portion of the RF spectrum.

The Master Control Station (MCS)

The MCS, located at the NASA Lewis Research Center (LeRC), controls the ACTS network by providing synchronization and circuit management information to the satellite and all earth stations via an orderwire channel. The MCS performs demand assignment of network capacity in response to capacity requests from earth stations.

The T1-VSAT Earth Stations

T1-VSAT earth stations are low burst rate terminals (27.5 Mbps uplink and 110 Mbps downlink) which each support 28 DS0 (64 Kbps) channels for the transmission of voice and data.

Each earth station is composed of an outdoor unit and an indoor unit. The outdoor unit consists of the antenna, decoding equipment and RF equipment. The indoor unit is comprised of a Modem-Processor (MP) assembly and the Terrestrial Interface Equipment (TIE). The organization of the indoor unit is shown in Figure 1, and a description of the indoor unit interfaces is given in [3], Appendix D.

The MP contains a modem, uplink and downlink traffic buffers, a TDMA time slot "switch" and a PC-AT compatible control computer. The TIE is a Redcom Laboratories, Inc. Modular Switching Peripheral (MSP), which is a small digital central office switch that provides standard telephone communications interfaces. The TIE




Figure 1 - T1-VSAT Indoor Unit Architecture

switches traffic between the uplink/downlink buffers of the MP and the voice and data equipment of the user, thereby providing the terrestrial hardware interface to the earth station.

The T1-VSAT earth station is controlled by two major software modules ([3], Appendix C). One is the modem management module, whose primary purpose is to control the modem hardware in order to communicate with the satellite. The other is the call management module, which controls the TIE, supplies the interface to the user via a dialing plan and provides end-to-end coordination of call setups and tear downs.

In subsequent sections it will be shown how the ACTS T1-VSAT network can be used to implement a satellite-based audio conferencing system.

AUDIO CONFERENCING STATE-OF-PRACTICE

Most audio conferencing service providers today use advanced, fully digital, full-duplex audio bridges in their conferencing business ([6], p. 33). These bridges are typically used in a star topology, where each conference participant is connected to the bridge. Bridges can be located anywhere in the telecommunications network. Since communication satellites are part of this network, it is possible that a conference participant may be connected to a bridge via one or more satellite hops (Fig. 2).

The audio signal of a conference participant that is connected to a bridge via one or more satellite hops is delayed by 250 to 300 msec. (depending upon the elevation of the satellite above the horizon) for each hop ([8], p. 68). If two participants are connected to the same audio bridge via satellite hops, the total delay in the signal between them is the sum of the delays incurred in each of their signals to the bridge. In contrast, the signal between two participants that are connected to the same bridge via terrestrial links typically has a delay of less than 100 msec. In general, the signals between any one participant and all other participants in the same conference may each have a different delay, depending on how each participant is connected to the bridge.

In conferences where all participants are connected to the bridge via satellite hops, the delays may be equal, but are usually of significant duration. For example, if each participant is connected to the bridge via a single satellite hop, the end-to-end path between all participants is a double hop with a 500 to 600 msec. delay.

Audio conferences generally operate in one of two basic modes: Meet-Me and Dial-Out. The Meet-Me mode requires all conference participants to dial the bridge telephone number at a designated time. There are several

variations of this mode, such as whether or not a passcode is required, and whether or not operator intervention is necessary. In the Dial-Out mode, each conference participant is called from the bridge at a designated time. There are also several variations of this mode, depending upon the degree of operator intervention and auto-dialing desired. Both the Meet-Me and Dial-Out modes require advance scheduling of the audio bridge. (See [6], page 33 for a more detailed description of these modes.)




Figure 2 - Audio Bridge Topology

Given the capabilities of current audio conferencing systems, several shortcomings can be identified:

- 1) Since it is possible for the audio signals between conference participants to have different delays, voice overruns may occur, making conversation difficult. This problem stems from the inability to determine: a) which connections at the bridge involve satellite hops, and b) the amount of delay on those connections. If this could be determined, an equal amount of delay could be induced on other connections as necessary to synchronize the transmission and reception of audio signals among all participants. However, this currently cannot be done automatically.
- 2) Since it is possible for the audio signals between conference participants to have significant delays due to double satellite hops, lengthy pauses result which make conversation unnatural.
- 3) Since audio conferences must be scheduled through the bridge provider, the action of starting a conference becomes less spontaneous. In addition, bridges, unlike PABX systems, do not support a camp-on calling feature. Therefore, once in a conference, a user cannot do a "hook flash", get dial tone, and add additional parties at

will, which is a completely natural operation.

AUDIO CONFERENCING USING ACTS

Requirements

The motivating factors behind the development of an audio conferencing system using ACTS were: 1) the existence of the unique capabilities of the satellite and T1-VSAT network, and 2) the shortcomings of current audio conferencing systems. Therefore, it was with the intention of demonstrating the usefulness of single hop, point-to-multipoint, on-board switched communications and improving audio conferencing state-of-practice, that the requirements for the ACTS audio conferencing system were developed.

These requirements are stated in [5] and summarized below:

- All participants will be inter-connected using single hop paths.
- No prior scheduling of resources will be required to start a conference.
- A new member will be able to be added at any time (i.e. it is not necessary to add all members at the start of the conference).
- A new member will be able to be added by any existing member (i.e. no operator intervention is required).
- A new member will be added to the conference in a minimal amount of time.
- Conference integrity will be maintained during ACTS network changes (e.g. startup of new traffic terminals, shutdown of existing traffic terminals, etc.).
- Simultaneous conference operations will be processed properly network-wide (e.g. new member adds, existing member drops and conference starts).
- All voice paths will use full-duplex channels, as in current audio conferencing systems.

The audio conferencing system which was developed satisfies all of the above requirements. It uses the point-to-multipoint on-board switching capability of the satellite in conjunction with a conference board in the TIE at each earth station to achieve single hop connectivity among all participants. Moreover, it addresses the lack of spontaneity in current systems by allowing a user to easily start a conference from any standard telephone handset connected to an ACTS T1-VSAT earth station, and quickly add new members to the conference at any time using the "hook flash" capability.

As will be seen later in this section, the fulfillment of

several of the above requirements posed challenging design problems that demanded innovative solutions.

Constraints

Several constraints governed the development of the ACTS audio conferencing system. These constraints and the reasons for them are described below:

- An unlimited number of conferences may be in progress in the network at the same time; however, each T1-VSAT earth station can only participate in a single conference at one time. This is a result of: 1) the customer decision to support only one Redcom Eight Party Additive Conference board at each earth station, and 2) the decision to use the board in a non-partitioned mode. (For more information regarding the use of this board, see [7], Sections 4.18 and 6.22.)
- The maximum number of connections to a conference is restricted to eight, due to the capacity of the Redcom Eight Party Additive Conference board.
- For any given conference, a single T1-VSAT earth station can only host one local connection to the conference. This is caused by a conference board limitation that restricts the way connections to it can be made. However, through the use of a speakerphone in lieu of a handset, this single local connection can be made to support multiple conference participants. (Note: For clarity in the remainder of this document, it is assumed that the local connection to a conference is via a standard handset; therefore, each local connection indicates one participant.)

These constraints, which are documented in [5], are provided here as supporting information to clarify subsequent discussions of design and implementation details of the ACTS audio conferencing system.

Solutions to Key Design Problems

The design of the ACTS audio conferencing system required the solution of many diverse problems. Three of these proved to be of special interest and are presented here in detail.

1) Achieving connectivity using single hop paths among all participants

The key elements underlying the achievement of efficient connectivity in ACTS audio conferencing are the use of the point-to-multipoint capability of the satellite's on-board switch and the replication of audio mixing hardware at each participant's earth station. The architecture described below

uses these two elements to provide full inter-connection among conference participants in a single hop.

In an ACTS audio conference, connectivity among earth stations is achieved by using point-to-multipoint circuits. Point-to-multipoint circuits are set up by exchanging a pre-defined sequence of orderwires with the MCS. This protocol first allows a simplex circuit to be set up from an origination earth station to a single destination earth station. Then, at any future time, a new destination earth station can be added to the same circuit, with the branching occurring at the satellite ([3], Section 3.9.3.1 and [1], Section 4.2.5). This can be repeated for any number of destination earth stations, thereby achieving 1 to N connectivity (Fig. 3). Assuming a group of N earth stations, full mesh connectivity is achieved by having each earth station in the group set up such a point-to-multipoint circuit to every other earth station in the group (Fig. 4).

The advantage of using point-to-multipoint circuits is that adding each additional destination to such a circuit does not require any additional resources at the origination earth station, nor any additional uplink capacity at the satellite. At each earth station, this method requires resources for 1 outgoing simplex circuit and N-1 incoming simplex circuits.

Since each earth station in an ACTS audio conference supports only one conference participant, the voice of that participant is transmitted on the single outgoing simplex circuit, and the voices of all other participants are received on the incoming simplex circuits. Hence, the use of point-to-multipoint circuits in this manner provides full mesh connectivity among all participants. The necessary circuits required to achieve this connectivity are set up dynamically as each new participant is requested to join the conference.

At each earth station, the incoming simplex circuits are used as inputs to the local audio conference hardware, and the mixed output is presented to the earpiece of the participant's handset. The output of the mouthpiece of the participant's handset is connected to the single outgoing simplex circuit and bypasses the audio conference hardware. This outgoing signal becomes one of the inputs to the audio conference hardware at each of the other conference participants' earth stations. These connections are shown in detail for a single earth station in Figure 5 and for a four-member




Figure 3 - Point-to-Multipoint Circuit

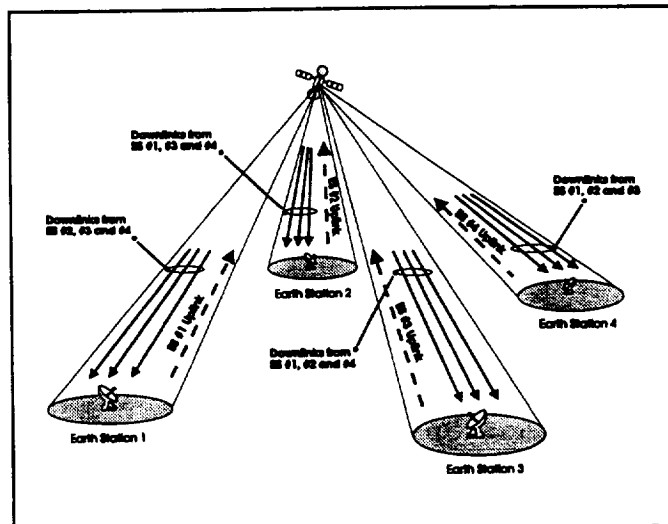


Figure 4 - Full Mesh Connectivity (N=4)

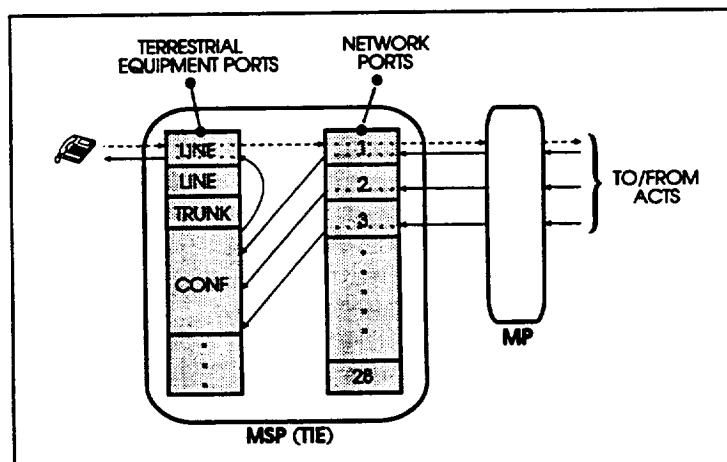


Figure 5 - Detailed Connections at Earth Station in 4-member Conference




Figure 6 - Total Connections for 4-member Conference

audio conference in Figure 6.

As shown above, an architecture incorporating point-to-multipoint circuits and audio mixing hardware in this manner can be used to provide full conference inter-connection without inducing double hop paths.

2) Adding a new participant in a minimal amount of time

The task of adding a new participant to an audio conference is best explained with an example. Assume a three-party conference is in progress, involving participants at earth stations #1, #2 and #3. Furthermore, assume that the participant at earth station #1 is attempting to add a new member at earth station #4. The following lists show the tasks that must be done at each of the four earth stations, and Figure 7 shows the necessary circuits that must be set up to add the new member.

Earth station #1:

- Notify all earth stations in the network that earth station #4 is to be added to this conference.
- Request earth station #4 to check for the availability of resources to support an incoming simplex circuit.
- If earth station #4 indicates that resources are available, request the MCS to add earth station #4 as an additional destination to the existing outgoing simplex circuit at this earth station.

Earth station #2:

- Request earth station #4 to check for the availability of resources to support an incoming simplex circuit.
- If earth station #4 indicates that resources are available, request the MCS to add earth station #4 as an additional destination to the existing outgoing simplex circuit at this earth station.

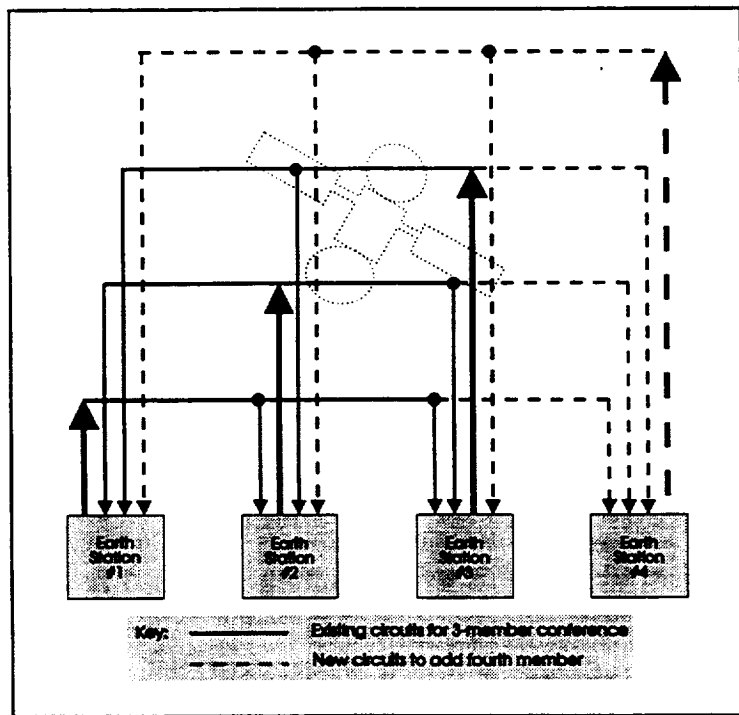


Figure 7 - Circuits to Add Fourth Member to 3-member Conference

Earth station #3:

- Request earth station #4 to check for the availability of resources to support an incoming simplex circuit.
- If earth station #4 indicates that resources are available, request the MCS to add earth station #4 as an additional destination to the existing outgoing simplex circuit at this earth station.

- If the circuit to earth station #2 is set up successfully, request earth station #3 to check for the availability of resources to support an incoming simplex circuit.
- If earth station #3 indicates that resources are available, request the MCS to add earth station #3 as an additional destination to the existing outgoing simplex circuit at this earth station.

Earth station #4:

- Check for the availability of the audio conference hardware resource.
- If available, check for the availability of resources at this earth station to support an outgoing simplex circuit.
- If available, request earth station #1 to check for the availability of resources to support an incoming simplex circuit.
- If earth station #1 indicates that resources are available, request the MCS to setup a new outgoing simplex circuit to earth station #1.
- If the circuit to earth station #1 is set up successfully, request earth station #2 to check for the availability of resources to support an incoming simplex circuit.
- If earth station #2 indicates that resources are available, request the MCS to add earth station #2 as an additional destination to the existing outgoing simplex circuit at this earth station.

Let the length of time necessary to set up a new audio conference circuit in the ACTS network be represented by T . Measurements have shown that the length of time necessary to add a destination to an existing audio conference circuit is approximately half this, or $\frac{1}{2}T$. Continuing the example above, Table 1 shows the time required by each earth station to complete its tasks.

Two approaches for executing these tasks were examined: 1) a serial approach, in which all tasks at one earth station are completed before beginning the tasks at any other earth station, and 2) a parallel approach, in which the tasks at all earth stations are executed at the same time, although the tasks at each individual earth station are executed serially.

Comparing these two approaches, it can be seen that adding a fourth member to a three-party conference using the serial method requires

$$\frac{1}{2}T + \frac{1}{2}T + \frac{1}{2}T + 2T = 3\frac{1}{2}T$$

Earth Station	Task	Time	Total Time
#1	add downlink to #4	$\frac{1}{2}T$	$\frac{1}{2}T$
#2	add downlink to #4	$\frac{1}{2}T$	$\frac{1}{2}T$
#3	add downlink to #4	$\frac{1}{2}T$	$\frac{1}{2}T$
#4	set up new circuit to #1	T	$2T$
	add downlink to #2	$\frac{1}{2}T$	
	add downlink to #3	$\frac{1}{2}T$	

Table 1 - Expected Effort to Add Fourth Member to 3-member Conference

whereas doing the same operation using the parallel method requires

$$\max (\frac{1}{2}T, \frac{1}{2}T, \frac{1}{2}T, 2T) = 2T.$$

In general, attempting to add the N^{th} member to a conference using the serial method requires approximately $(N)T$, but doing the same operation using the parallel method only requires $(N/2)T$. In the ACTS network, under moderate load conditions, T is approximately 9 seconds. Given this value of T and the above example of adding a fourth member to a conference, the difference between the serial method (31.5 secs.) and the parallel method (18 secs.) clearly indicates that a parallel approach is more desirable.

The parallel approach is not without its drawbacks however. In situations where adding a member to a conference fails due to unavailability of resources, the serial approach is superior. This is due to the ability of this method to begin failure processing immediately upon detecting a resource problem at any earth station, thereby simplifying cleanup operations and minimizing unnecessary resource allocation at the remaining earth stations. On the other hand, the parallel method is unable to detect failures

this quickly. In this method, when a resource problem is detected at an earth station, failure processing must be delayed until all other earth stations in the conference have completed their specified tasks, due to the complexity of interrupting simultaneous circuit setups in various states. This delay causes cleanup operations to be more complex and results in more frequent allocation of unnecessary resources.

Overall, it was determined that the benefits of using the parallel approach in cases where a member is successfully added to a conference far outweighed the drawbacks incurred in cases where a member fails to be added. As such, to fulfill the requirement of being able to add a member to a conference in a minimal amount of time, the parallel approach was used.

As an example, Figure 8 depicts the chronological sequence of actual orderwires that are sent to accomplish all of the tasks listed above to add a fourth member to a three-party conference, using the parallel method.

Table 2 shows actual timing measurements for adding the second through the seventh member to a conference which was started from ES #1 located at NASA-LeRC.

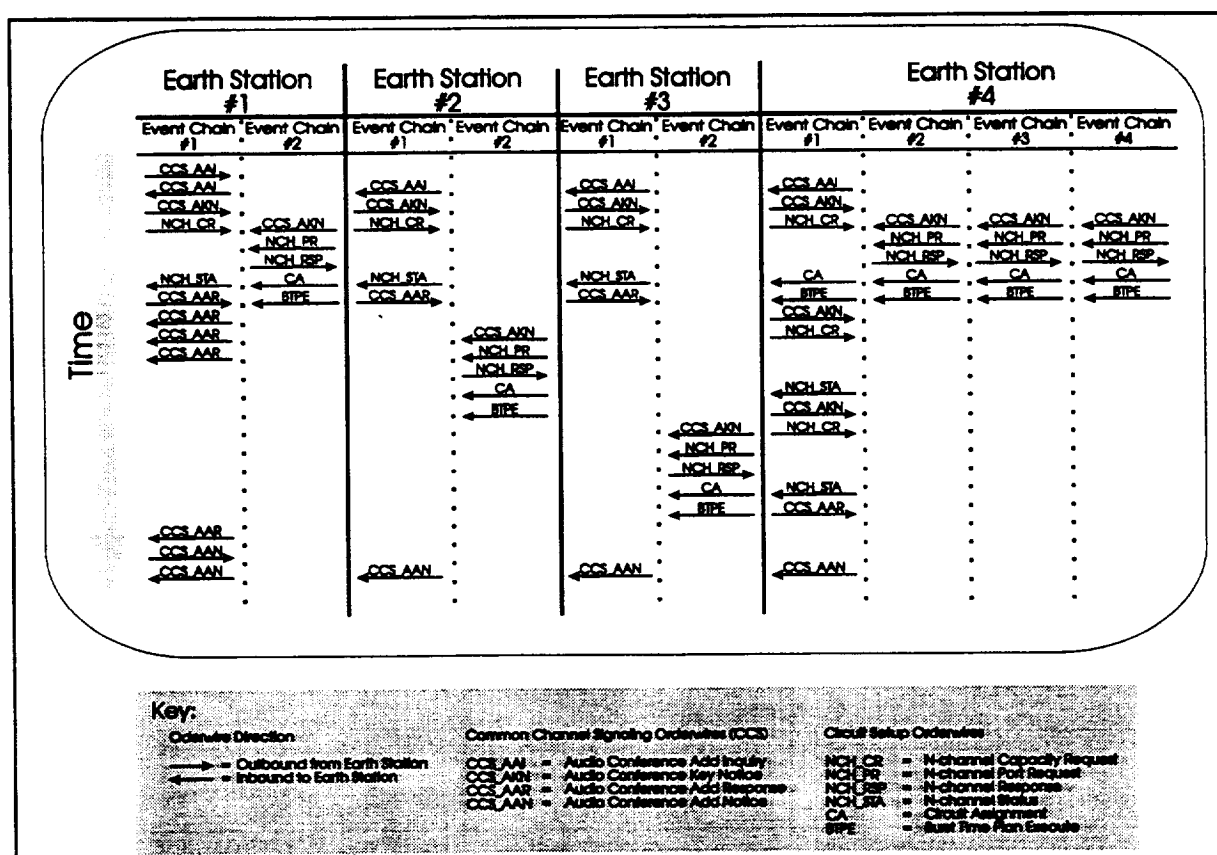


Figure 8 - Orderwire Traffic to Add Fourth Member to 3-member Conference

	2nd member ES #2 NASA-LeRC Clev., OH	3rd member ES #6 NASA-LeRC Clev., OH	4th member ES #21 NASA-LeRC Clev., OH	5th member ES #8 Univ. of Md. Coll. Pk, MD	6th member ES #9 NTIA Boulder, CO	7th member ES #5 Mitre Reston, VA
Trial 1	9.53	13.34	17.28	21.74	24.90	29.12
Trial 2	9.40	14.80	17.16	23.16	25.46	28.17
Trial 3	9.47	13.66	16.85	21.92	24.25	29.40

Table 2 - Actual Timing Measurements for Adding Conference Members

3) Handling simultaneous user actions

The ACTS audio conferencing system supports three basic user actions: *conference start*, *add member*, and *drop member*. When a conference participant initiates one of these actions, a complex chain of lower-level events is triggered at all earth stations in the network, and the associated processing for each of these events is performed. These lower-level events may include processing inbound and outbound ACTS network orderwires (Fig. 8), and sending commands to, and receiving status from, the TIE hardware. Processing of these events results in allocation or deallocation of the necessary resources at those earth stations involved in the conference, setup or teardown of circuits at those earth stations involved in the conference, and network-wide storage or updating of conference information. At any given earth station, these lower-level events are processed asynchronously and are often intermingled with other non-conference-generated events.

From the description given above, it can be seen that handling a single user action initiated by a conference participant is a very intricate process, involving multiple independent steps and affecting the entire network. This process is further complicated if two or more such actions are unknowingly initiated simultaneously, either by multiple users at the same earth station or by multiple users at different earth stations. In these situations, the corresponding chains of lower-level event processing become intermingled. For certain user actions, this intermingling may result in several problems, such as the following:

- deadlock situations involving resource allocation,
- a loss of integrity in conference data structures, and
- faulty addition or removal of conference members.

It is clear, therefore, that for certain user actions, the corresponding chain of lower-level event processing, or

some portion thereof, cannot be mixed with that of other user actions, if correct behavior of the system is to occur. The solution to this problem lies in the ability to immediately detect occurrences of simultaneous incompatible user actions across the network.

This detection is performed in the ACTS audio conferencing system through the use of a *network lock*, which can best be described as a type of "network-wide distributed semaphore". The implementation of this lock mechanism is based upon the broadcast capability of the network (i.e. the ability to specify that an orderwire be sent to all earth stations in the network) and the fact that orderwires are processed sequentially in the order in which they are received at each point in the network (i.e. at the MCS, the satellite and each earth station). These two items provide underlying support for a network-wide test and set operation. This operation allows a single earth station to manipulate data items at all other earth stations in the network, with the guarantee of no interference from other earth stations that may have simultaneously attempted the same manipulation -- a concept which is fundamental to the operation of the network lock.

The ACTS network lock differs from a traditional semaphore in several ways.

First of all, a semaphore is a localized access control mechanism, represented by a single data structure residing on a single system, whereas the network lock is a distributed access control mechanism, represented by replicated data structures residing on multiple systems (i.e. earth stations). In spite of the dispersed nature of the network lock, however, it is still, in the same manner as a semaphore, considered a single entity. Therefore, whenever a lock operation is requested, it is either granted at all earth stations in the network or at none of them.

Secondly, unlike a semaphore, the network lock does not provide mutually exclusive access to a small critical region

of code or a resource on a single system, but rather to several non-continuously executed sections of code, each occurring on multiple systems. These sections of code represent processing of low-level events generated as a result of a specific audio conference user action. Once such an action is initiated, the network lock is set and each earth station begins processing the events associated with the specific user action. At each earth station, processing of these audio conference events may be intermingled with processing of other non-conference events, even though the network lock is set. However, as long as the network lock is set, no processing of any audio conference events generated by a different user action will occur at any earth station.

Lastly, a semaphore and the network lock differ in the types of operations that can be performed upon them. Since a semaphore guards an object (i.e. a critical region or resource) by providing a single entry point and a single exit

point, only two operations are needed -- one to denote when entry is permitted (i.e. *wait*) and one to denote when exit is complete (i.e. *signal*). There also is no need for a semaphore to store an identifier denoting the holder of the semaphore. In contrast, the network lock guards several non-continuously executed sections of code which do not have a single entry point or single exit point. This requires one operation to denote when initial entry is permitted, one to denote when repeated entry is permitted and one to denote when exit is complete. The initial entry operation (i.e. *program*) also causes an identifier denoting the holder of the lock (i.e. the *access key*) to be stored in the lock, thereby programming the lock so that it can only be passed by that key. The repeated entry operation (i.e. *key_verify*) only allows further event processing if the presented access key passes the lock. The exit operation (i.e. *de-program*) causes the access key to be cleared from the lock, making the lock available for re-programming.

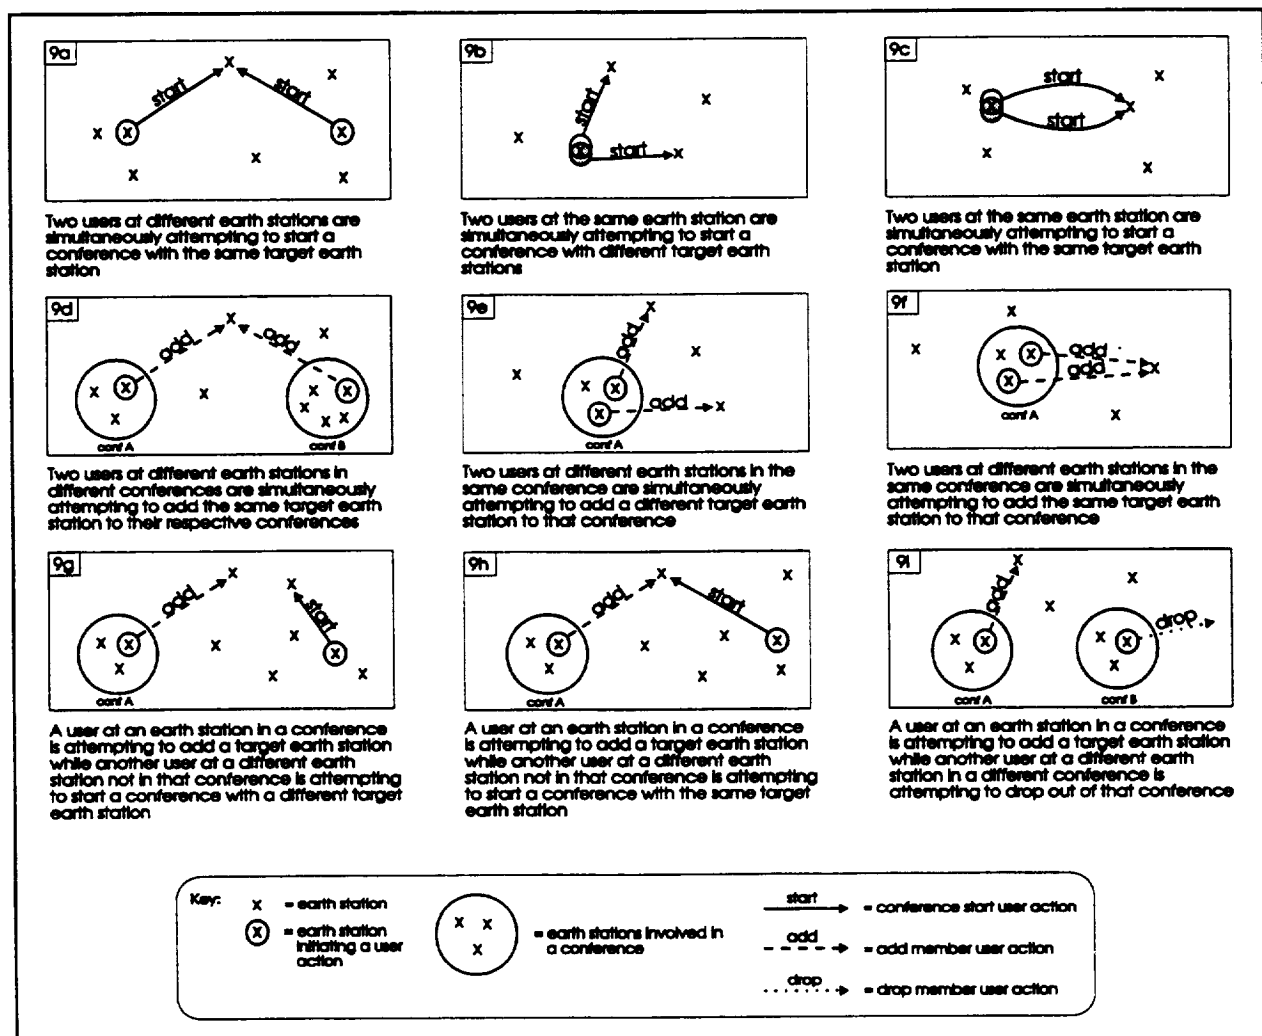


Figure 9 - Examples of Simultaneous User Actions

In summary, the network lock provides the ability to detect occurrences of simultaneous user actions across the network. Figure 9 shows some examples of simultaneous user actions. When situations such as these are detected, processing of one of the user actions proceeds; the other(s) is(are) either disallowed or queued for later processing, as appropriate for the action. More specifically, only one *conference start* user action or one *add member* user action may be in progress in the network at one time. If more than one of these actions are requested simultaneously, the others are rejected. However, multiple *drop member* user actions may occur simultaneously. If a *drop member* user action occurs during a *conference start* or *add member* user action, it is queued until the protected action is complete.

Usage

To start an ACTS audio conference or add a member to an existing ACTS audio conference, a user dials the following:

ff-nn-r-xxxx

where the fields have the following meaning:

ff	feature code (use 50 to start a conference; use 51 to add a member)
nn	number of earth station at which new member is located (must be a 2-digit number, so numbers below 10 must be dialed with a leading '0')
r	code to determine routing of call at earth station at which new member is located (must be a 1-digit number; use 0 to access a phone directly connected to earth station; no others currently supported)
xxxx	extension number (4-digit number of telephone at earth station at which new member is located)

If a user is starting a conference, the dialing sequence is dialed immediately after picking up a telephone and hearing the dial tone. If the conference cannot be started, for any one of a number of reasons, the user hears the intercept tone until the telephone is hung up. Otherwise, the satellite portion of the call is set up, and one of the following occurs:

- If the telephone at the new member's earth station is in use, the user hears the slow busy tone until the telephone is hung up.
- If the telephone at the new member's earth station is not in use, the user hears the ringback tone until either the telephone is hung up, or the new member picks up the ringing telephone, in which case the call is connected.

If a user is adding a member to an existing conference, the dialing sequence is dialed immediately after executing a hook flash. This causes the user to be temporarily disconnected from the conference, dial tone to be heard and dialing to be enabled. At any time after executing the hook flash, the user may return to the conference by executing another hook flash. (Note: If the user executes a hook flash to return to the conference before the outcome of the add member request is determined, no notification of the outcome is provided. For the remainder of this discussion, it is assumed that the user does not return to the conference until notification is received.) If the new member cannot be added, for any one of a number of reasons, the user hears the intercept tone until another hook flash is executed, returning the user to the conference. (The conference is unaffected by the failure of the add member request.) Otherwise, the additional satellite portion of the call is set up, and the user hears the accept tone series until another hook flash is executed, returning the user to the conference. At this point, one of the following occurs:

- If the telephone at the new member's earth station is in use, the user hears the remaining portion, if any, of a 10 second slow busy tone, to which all other conference members have already been listening. Throughout this period, all conference members remain fully connected to each other. After 10 seconds, the tone ceases, and the satellite portion of the call to the new member is torn down.
- If the telephone at the new member's earth station is not in use, the user hears the remaining portion, if any, of a 30 second ringback tone, to which all other conference members have already been listening. Throughout this period, all conference members remain fully connected to each other. If the new member picks up the ringing telephone within the 30 second period, the new member is connected to the conference. Otherwise, after 30 seconds, the tone ceases, and the satellite portion of the call to the new member is torn down.

To exit a conference at any time, a user hangs up the telephone. If only one member remains, that member hears silence for 10 seconds, and then the intercept tone, until that member also hangs up the telephone. If more than one member remains, the conference remains in existence with one less member.

The following examples demonstrate the use of the ACTS audio conferencing system:

- To start an audio conference with the person at earth station 17, extension 9740, a user at any

other earth station dials the following:

50-17-0-9740

- To add the person at earth station 2, extension 1428, to an existing conference, a user who is already in that conference executes a hook flash, waits for the dial tone, and dials the following:

51-02-0-1428

The user can execute another hook flash to return to the conference at any time, but should wait for notification of success or failure of the add member request.

Although the dialing protocol for the ACTS audio conferencing system is ACTS-specific, it can easily be modified to implement the North American Numbering Plan. Additional information pertaining to the usage of the ACTS audio conferencing system can be found in [4], which describes the complete dialing plan available at ACTS T1-VSAT earth stations.

CONCLUSIONS

The ACTS audio conferencing system successfully demonstrates the usefulness of ACTS technology for practical real-world applications, and in addition, it addresses some shortcomings in audio conferencing state-of-practice. The completed system fulfills all of the pre-defined requirements. It allows a conference to be easily started from any standard telephone handset connected to an ACTS T1-VSAT earth station, and new members to be quickly added to the conference at any time using the "hook flash" capability. The system uses the point-to-multipoint on-board switching capability of the satellite plus a conference board at each earth station to achieve single hop voice paths among all participants.

The outcome of this effort has shown that satellite on-board switching technology, as demonstrated by ACTS, is indeed suitable for implementing audio conferencing, and may be applicable to similar communication services with multipoint requirements, such as video conferencing. In the future, as the role of satellite systems in the GII is defined, systems with on-board switching technology will prove useful in providing audio conferencing and other innovative services.

ACKNOWLEDGEMENTS

The authors wish to acknowledge the ACTS Project Office of the NASA Lewis Research Center for providing funding for this project, and Peter A. Lowry in particular, for his

insights and early work on the system architecture. In addition, recognition is given to Kevin M. McPherson for his advocacy and facilitation of the project and Mark J. Vanderaar for his management support. Lastly, special thanks are given to the operators and engineers of the ACTS Master Ground Station for their troubleshooting and testing assistance, and to all those who reviewed the final draft of this paper.

REFERENCES

- [1] COMSAT Laboratories, *Advanced Communications Technology Satellite (ACTS) Ground Segment: LBR TDMA Network Control Performance Specification*, ACTS. Rev. D. Technical Report, NASA-LeRC ACTS Project Office ACTS-DOC-92-011, 1992.
- [2] Gedney, Richard T., and Frank Gargione. "ACTS: New Services for Communications." *Satellite Communications* Sept. 1994: 48-54.
- [3] Harris Corporation, *Statement of Work: LBR Earth Stations for the NASA Advanced Communications Technology Satellite (ACTS) Program*. Rev. B. Technical Report, NASA-LeRC ACTS Project Office, 1991.
- [4] NYMA, Incorporated. *Advanced Communications Technology Satellite (ACTS) Call Manager (CM) Dialing and Call Control Specification*. Rev. 4.0. Technical Report, NASA-LeRC ACTS Project Office CM-USR-002, 1995.
- [5] NYMA, Incorporated. *Advanced Communications Technology Satellite (ACTS) Call Manager (CM) Software Requirement Specification (SRS) for Audio Conferencing (AC)*. Rev. 1.0. Technical Report, NASA-LeRC ACTS Project Office CM-SRS-003, 1994.
- [6] Portway, Patrick S., and Carla Lane, Ed.D, eds. *Guide to Teleconferencing and Distance Learning*. 2nd ed. Livermore, CA: Applied Business teleCommunications., 1994.
- [7] Redcom Laboratories, Incorporated. *MSP Modular Switching Peripheral General Manual*. Victor, NY: Redcom Labs, 1990.
- [8] Tanenbaum, Andrew S. *Computer Networks*. 2nd ed. Englewood Cliffs: Prentice, 1988.

REPORT DOCUMENTATION PAGEForm Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE June 1996	3. REPORT TYPE AND DATES COVERED Final Contractor Report	
4. TITLE AND SUBTITLE Full Mesh Audio Conferencing Using the Point-to-Multipoint On-Board Switching Capability of ACTS			5. FUNDING NUMBERS WU-942-35-04 C-NAS3-27186	
6. AUTHOR(S) Mary L. Rivett and Zubin H. Sethna				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) NYMA, Inc. 2001 Aerospace Parkway Brook Park, Ohio 44142			8. PERFORMING ORGANIZATION REPORT NUMBER E-10311	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) National Aeronautics and Space Administration Lewis Research Center Cleveland, Ohio 44135-3191			10. SPONSORING/MONITORING AGENCY REPORT NUMBER NASA CR-198498	
11. SUPPLEMENTARY NOTES Prepared for the 16th International Communications Satellite Systems Conference cosponsored by AIAA, CASI, AAAF, DGLR and IEICE, Washington, D.C., February 25-29, 1996. Project Manager, Paul Choma, Engineering Support Division, Lewis Research Center, organization code 4410, (216) 433-5750.				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Unclassified - Unlimited Subject Category 17 This publication is available from the NASA Center for AeroSpace Information, (301) 621-0390.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) The purpose of this paper is to describe an implementation of audio conferencing using the ACTS T1-VSAT network. In particular, this implementation evaluates the use of the on-board switching capability of the satellite as a viable alternative for providing the multipoint connectivity normally provided by terrestrial audio bridge equipment. The system that was implemented provides full mesh, full-duplex audio conferencing, with end-to-end voice paths between all participants requiring only a single hop (i.e. 250 msec. delay). Moreover, it addresses the lack of spontaneity in current systems by allowing a user to easily start a conference from any standard telephone handset connected to an ACTS earth station, and quickly add new members to the conference at any time using the "hook flash" capability. No prior scheduling of resources is required and there is no central point of control, thereby providing the user with the spontaneity desired in audio conference control				
14. SUBJECT TERMS Audio conferencing; ACT; Satellite; Full mesh			15. NUMBER OF PAGES 14	
			16. PRICE CODE A03	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT	

**National Aeronautics and
Space Administration**

Lewis Research Center
21000 Brookpark Rd.
Cleveland, OH 44135-3191

Official Business

Penalty for Private Use \$300

POSTMASTER: If Undeliverable — Do Not Return