SAE Aerospace
An SAE International Group

AEROSPACE INFORMATION REPORT

SAE AIR4886

Issued 1996-04
Reaffirmed 2011-11

Statement on Requirements for Real-Time Communication Protocols (RTCP)

FOREWORD

This document contains the requirements for the SAE AS-2A-1 Real-Time Communications Protocols (RTCP) as developed by the RTMT Group. These requirements are an edited compilation of those voiced by numerous military and commercial organizations primarily concerning aircraft applications. However, it is the opinion of the Task Group that these requirements have general applicability to any distributed real-time control systems.

A wide spectrum of industry representation has been tasked to define the requirements for a multinational, multiservice, commercially accepted military standard for real-time communication protocols. This effort was sponsored by SAE Interconnect Networks Committee AS-2. Inputs to this document have been consolidated from participating resources of the Department of Defense's military services, aerospace prime contractors, United Kingdom, France, Germany, electronic and LSI component suppliers, ARINC, commercial aviation equipment manufacturers and consumers, and systems consultants. The requirements defined herein discuss technical aspects based upon the following considerations:

a. Functions
b. Performance
c. Test Requirements
d. Operational Properties
e. Validation Requirements

FOREWORD (Continued)

Contacts:

This document has been prepared by the Real-Time Model Task Group (RTMT) of the SAE AS-2A System, Application and Requirements Subcommittee, which is part of the AS-2 Interconnect Networks Committee of the SAE Avionics Systems Division. Correspondence concerning this document may be addressed to:

Alan Birch, Chairman AS-2A-1 RTMT Group,
GEC-Marconi Avionics Limited,
ISG, Airport Works,
Rochester, Kent, UK. ME1 2XX (44) 1634-844400 Ext 4731
e-mail: alan.birch@gecm.com

Doug Johnson, Chairman AS-2A Subcommittee,
CECOM C2SID,
Department of Army,
AMSEL-RD-C2-TS-I,
Fort Monmouth, NJ 07703-5603 (1) 908 427-3284
email: johnsod1@doim6.monmouth.army.mil

Notice:

Committee approval of the previous version of this document, identified as ARD50007, was obtained by mail ballot completed September 20, 1991. This AIR4886 was approved by the AS-2 Committee by mail ballot, October 26, 1995.

Charter:

The RTMT was chartered by the Systems, Applications, and Requirements Subcommittee (AS-2A) as follows

The Real-Time Model Task Group (RTMT) shall define requirements for communication protocols specifically intended for real-time applications. These requirements are intended to complement SAE AS-2 communication standards and address the logical interface between the User and communication tasks and the functions allocated to communication tasks.

Systems for real-time applications are characterized by the presence of hard deadlines where failure to meet a deadline must be considered a system fault.

In order to fulfill its commitment, RTMT will develop and maintain supporting documentation such as protocol assessment reports, requirements definition and logical interface specifications.

RTMT supports and interacts with developers and users of lower level real-time protocols. RTMT sees as its mission to increase the understanding of the specificity of real-time communication needs. RTMT operates in liaison with ISO, NATO, IEEE, POSIX, JIAWG, NGCR and other standardization groups worldwide.

TABLE OF CONTENTS

TABLE OF CONTENTS (Continued)

1. SCOPE:

The purpose of this document is to establish the requirements for Real-Time Communication Protocols (RTCP). Systems for real-time applications are characterized by the presence of hard deadlines where failure to meet a deadline must be considered a system fault.

These requirements have been driven predominantly, but not exclusively, by aerospace type military platforms and commercial aircraft, but are generally applicable to any distributed, real-time, control systems. These requirements are primarily targeted for the Transport and Network Layers of peer to peer protocols, as referenced in the Open System Interconnect Reference Model (2.2.1 and 2.2.2), developed by the International Standards Organization (ISO). These requirements are intended to complement SAE AS4074 (2.1.1) and AS4075 (2.1.2), and future SAE communications standards. Although information transfer objectives herein concentrate primarily on digital data flow attributes, efforts have been made such as not to preclude sensor/video and voice information transfers. This document provides evaluation criteria for selecting viable alternatives for a real-time communications protocol standard.

2. REFERENCES:

The following publications form a part of this specification to the extent specified herein. The latest issue of SAE publications shall apply. The applicable issue of other publications shall be the issue in effect on the date of the purchase order. In the event of conflict between the text of this specification and references cited herein, the text of this specification takes precedence. Nothing in this specification, however, supersedes applicable laws and regulations unless a specific exemption has been obtained.

2.1 SAE Publications:

Available from SAE, 400 Commonwealth Drive, Warrendale, PA 15096-0001.

2.1.1 AS407 Linear Token-Passing Multiplex Data Bus Standard

2.1.2 AS4075 High Speed Ring Bus Standard

2.1.3 ARD50033 DRAFT Logical Interface Definition for Real-Time Communication Protocols (RTCP)

2.2 U.S. Government Publications:

Available from DODSSP, Subscription Services Desk, Building 4D, 700 Robbins Avenue, Philadelphia, PA 19111-5094.

2.2.1 ISO 7498 Information Processing Systems - Open Systems Interconnection - Basic Reference Model

2.2.2 ISO 7498 Addendum 1: Connectionless-mode Transmission. Information Processing Systems - Open Systems Interconnection - Basic Reference Model.

3. INTRODUCTION:

The RTCP requirements set forth in this document, reflect a consensus of experts' views for satisfying the needs of applications encompassing the entire field of aerospace military platforms and commercial aircraft. Nonaerospace applications are not precluded and will be encouraged to utilize the proposed RTCP standard. Aerospace categories of the RTCP intended applications are outlined in Figure 1.

MILITARY PLATFORMS

- Space Craft
  - Manned
  - Unmanned
- Manned Aircraft
  - Fixed Wing
    - Command Post
    - Bomber
    - Transport
    - Fighter
    - Anti-Sub Warfare
    - Ground Attack
    - Airborne Early Warning
  - Rotary Wing
    - Light
    - Heavy
- Unmanned Aircraft
  - Missile
  - Remote Piloted Vehicle

CIVILIAN AIRCRAFT

- Manned Aircraft
  - Fixed Wing
    - Transport
  - Rotary Wing

NOTE: Multiservice
Multinational

FIGURE 1 - Aerospace Categories of RTCP Intended Applications

3.1 Target Applications Descriptions:

The protocol requirements defined in this document are intended for real-time applications. Real-time applications, as defined here, are characterized by the presence of hard deadlines where failure to meet a deadline must be considered a system fault. While these applications are open in that they use published standards to ensure interoperability between equipment and across platforms, they are closed in the sense that valid users and transactions are predetermined during system design for a specific implementation.

3.1 (Continued):

The following descriptors and characteristics of the intended target applications are based on a set of KEY DESCRIPTOR WORDS. These target application key descriptor words were developed by the Real-Time Model Task Group (RTMT) and they represent the consensus of this group. These key words and functional breakdown are included as Appendix A.

3.1.1 Distributed: The RTMT target applications are distributed. This means that in addition to applications and processing being distributed, the following are also distributed: data, communications functions, control, storage, and system management functions.

The characteristics of distributed data are many fold and include the following:

a. Demand driven; that is, the data are sent when requested.
b. Supply driven; that is, the data are sent at the time that the data are originated.
c. For some data categories all data are essential, while for other categories only the latest data (or the latest N sets of data) are essential.
d. For some data categories the integrity of the data is essential, while for other categories integrity is not a driving requirement.
e. For some categories of data, the order of arrival, or knowledge of the order of origination, is essential.
f. Distributed data are of any size and of different data types (objects).

The characteristics of the communication functions being distributed include the data transfer itself, packetization (the segmentation and reassembly of data), flow control and the routing to target destination(s) of the data.

Characteristics of the control being distributed in these RTMT target applications include the following:

a. The control of initiating execution of subsets of Application tasks on distributed resources.
b. Provisions for the distribution of a global reference point, for synchronization purposes, and to allow time stamps on data, should that be desired.
c. Provisions for remote interrupts.

The characteristics of distributed storage imply that the target applications facilitate parallel data searches and stores, and need guaranteed data coherence (that is, the data stored must at any point in time appear consistent to all users).

The implications of system management functions being distributed in these RTMT target applications include the following:

a. Configurations; both static and dynamic.
b. Monitoring of system status, health and errors; the latter being either hard or intermittent.
c. System state support; the global understanding of current system state, for instance, normal or exception states with degraded modes and changing system priorities and objectives.

3.1.2 Uninterrupted Operational Performance: The RTMT intended applications are characterized by requiring uninterrupted operational performance. This uninterrupted operational performance implies the tolerance of the occurrence of both scheduled and unscheduled events. The RTCP may take care of scheduled events either off-line (action taken in anticipation) or on-line (action taken reacting to the event).

The unscheduled events may be either planned, such as interrupts or asynchronous events, or unplanned. The unscheduled unplanned events include both faults and defects.

3.1.3 Real-Time: The RTMT intended applications are real-time systems. In this context real-time means that the systems are characterized by the presence of hard deadlines where failure to meet a deadline must be considered a system fault.

Real-time is a relative concept; timing parameters may vary from system to system. Requirements for each application specify the necessary deadlines, i.e., the acceptable response times.

3.1.4 Highly Automated: The RTMT intended applications are characterized by many functions being automated. To be automated a function must be predefinable, that is, either via analytical methods or parametrically through behavioral models.

3.1.5 Control: The RTMT intended applications may to a large extent be characterized by executing control functions. There are cause/reaction modes of system behavior. This implies closed loop(s) with cyclic data, and a fail safe operation with high integrity. In addition to closed loops the system may contain open loops.

3.1.6 Predefined System Configurations: All RTMT intended applications are characterized by having system configurations that are predefined, that is planned in advance during the system design. Therefore, all system parameters, such as processor or communication network loads, can be calculated and resources allocated. This includes the configurations for both system dependability and system states changing, with anomalies resulting in the exception handling of, for instance, data, applications/processing, system management, and initialization/reinitialization.

3.2 System Definition:

The RTMT intended target applications as described above (see 3.1) can be functionally partitioned into the User Application and supporting communications functions (the RTCP) with a logical interface between the two functional areas as shown in Figure 2. It is intended that this logical interface be independent of underlying interconnect network implementation (e.g., 2.1.1 and 2.1.2).

This document defines the requirements for the communications functions that can support the target applications as described above. It is understood that there may be a User Specific Communications Function integrated with the User Application. Requirements for this integrated User Specific Communications Function are not included in this requirements document.

3.2 (Continued):

These target applications are distributed over a set of hardware equipment executing User Applications and Communication Functions (see Figure 3). The Communication Functions may be partly or completely allocated to the Network Interface Unit (NIU); so may the User Application with the Host.

4. REQUIREMENTS:

4.1 General:

4.1.1 Transfer Topologies: Transfer topologies as defined below apply to all system services, that is, to the data communication service (3.2.1), the synchronization service (3.2.2), and to management services (3.2.3).

RATIONALE: The transfer topologies described below for User to User(s) transfers are also applicable within the RTCP.

From the theoretical point of view, only unidirectional point-to-point transfers are necessary. However for the real-time applications being considered here, other topologies are useful. Please refer to Figure 4 for definitions of topology terms as used here.

4.1.1.1 Point-To-Point: The RTCP shall support unidirectional point-to-point transfers between any given pair of User entities.

It is desired that the RTCP support bidirectional point-to-point transfers between any given pair of User entities.

RATIONALE: The unidirectional point-to-point transfer provides the fundamental service. Sometimes, transfers are simultaneously required in both directions between two User entities. While this could simply be provided by two, independent, unidirectional transfers, it would be advantageous for the two transfers to be invoked by a single request. Since the bidirectional link may have to be implemented as two unidirectional links at a lower level, the advantage is slight and the capability only desired.

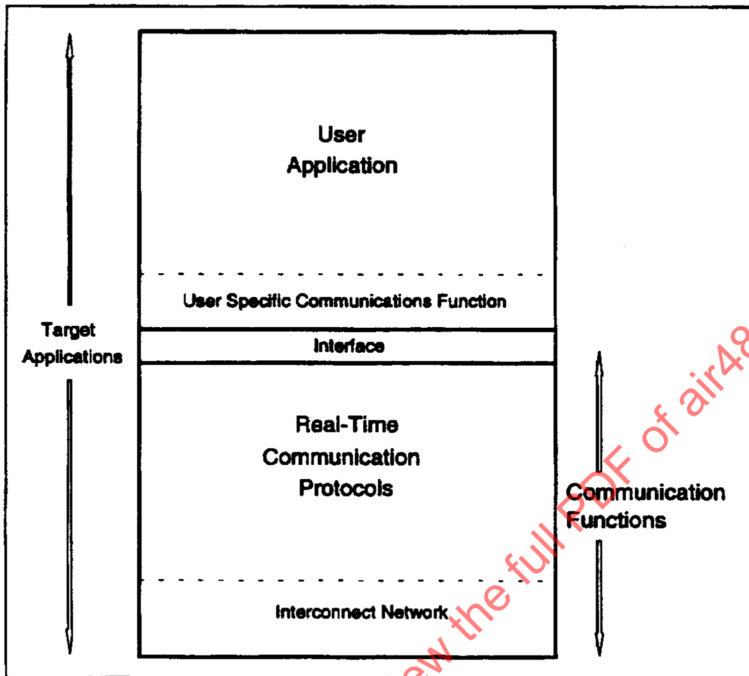

FIGURE 2 - RTCP System Definition

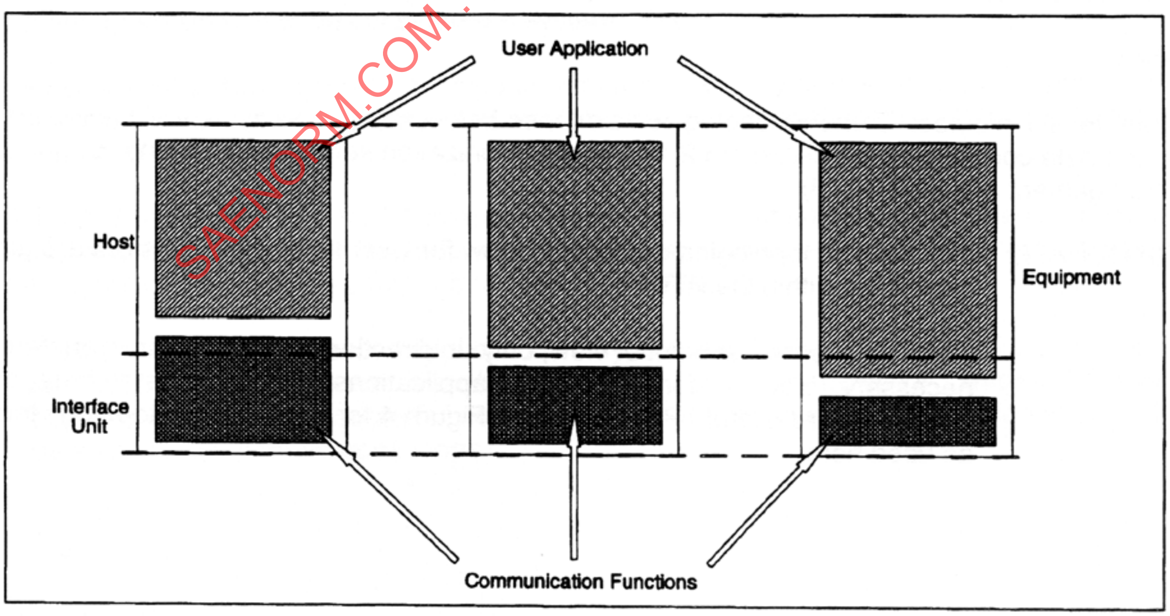

FIGURE 3 - Relationship Between Entities

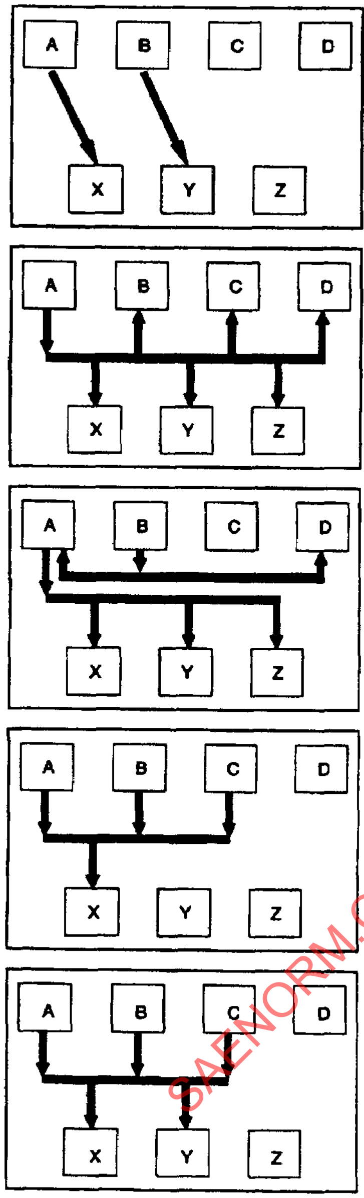

Point-To-Point: A transfer from any one station to any other one station; for instance, from A to X, or from B to Y.

Broadcast: A transfer from any one station to all other stations; for instance, from A to B and C and D and X and Y and Z.

Multicast: A transfer from any one station to a number of, but not all, other stations; for instance, from A to X and Y and Z, or from B to A and D.

Concentration: The transfer to any one station from a number of stations with the merger of the different pieces of information into a larger concatenated piece of information before presentation to the receiving entity; for instance, X receiving data from A, B, and C concatenated into one piece of information 'ABC'.

Multiple Concentration: The transfer to more than one station from a number of stations with the merger of the different pieces of information into a larger concatenated piece of information before presentation to the receiving entity; for instance, X and Y receiving data from A, B, and C concatenated into one piece of information 'ABC'.

FIGURE 4 - Definition of Topology Terms

4.1.1.2 Broadcast: The RTCP shall support one-to-all transfers (broadcast), always unidirectional, between any given User entity (acting as the unique source of the transfer) and all other User entities (acting as simultaneous data sinks of the transfer).

RATIONALE: For real-time applications, certain information, created at a given source, may be needed by all sinks. If the RTCP offered only point-to-point transfers, this operation would waste a lot of bandwidth since the same information would have to be transmitted (on the same media) many times, repeated for each User. Therefore, it seems reasonable for the RTCP to offer a special type of logical transfer topology which will use a unique transmission (of one message only), but simultaneously destined to all the User entities.

4.1.1.3 Multicast: The RTCP shall support one-to-many transfers (multicast), always unidirectional, between any given User entity (acting as the unique source of the transfer) and a number of (not all) other User entities (acting as simultaneous and defined targets for the transfer).

RATIONALE: For real-time applications, certain information, created at a given source, may be needed by some, but not all, sinks. The rationale is very similar to that for broadcast above. If the RTCP offered only point-to-point transfers, this operation would waste bandwidth since the same information would have to be transmitted (on the same media) several times, repeated for each User that requires the information. Therefore, it seems reasonable for the RTCP to offer logical transfer topologies which will allow transfer of information to groups of Users simultaneously.

4.1.1.4 Concentration: The RTCP shall support many-to-one transfers (concentration), unidirectional only, between a number of defined User entities (acting as simultaneous, or successive, sources of the transfer) and any other given User entity (acting as the single sink of the transfer).

RATIONALE: This is the exact opposite from broadcast/multicast. A key User in the system may need, at more or less regular intervals of time, several pieces of information, coming from different sources. These pieces of information are merged into a single large piece of information and then provided to the User.

That sink (i.e., the RTCP collecting entity) will then be involved in the job of arranging together all the incoming pieces of information in a neat conventional package (ordered and formatted properly) before passing it to its User (the sink host). This frees the User from unnecessary protocol processing.

4.1.1.5 Multiple Concentration: It is desired that the RTCP support many-to-many transfers (multiple concentration), unidirectional only, between any given pair of groups of defined User entities.

RATIONALE: This rationale is as for Concentration. The only difference is that there is now more than one sink for the data, allowing redundancy, for instance.

4.1.2 Semantic Importance: The RTCP shall be required to provide System Services for different categories of message classes distinguished by their degree of semantic importance.

Semantic importance here is defined as a measure of relative value of completion of a process associated with the transfer. Therefore, a measure of relative value, i.e., semantic importance, shall be associated with the completion of each requested data transfer.

The assignment of semantic importance to a message shall be the responsibility of the system designer.

RATIONALE: Semantic importance class assignments for data transfers in the RTCP provide the means for support of varying system states. Semantic importance should not be confused with priority; priority provides the means for scheduling. Semantic importance may have no relevance under fault-free conditions, being mainly used to determine the action to be taken when particular system faults occur. Two such examples are load shedding and missed deadlines.

If the available system resources are reduced below the minimum required to perform the defined functionality, some of the less important functions will not be performed (load shedding). The transfers associated with these functions will also be shed. This may be indicated by a simple semantic importance number which is compared with the system state to determine whether or not the transfer should take place.

If a transfer misses its deadline the system has failed. However, the information may still be of value such that the recovery action could be to complete the transfer as soon as possible. Alternatively, if the completion of a transfer has no value, it may be important to cancel the transfer.

Semantic importance can also be used to dynamically reschedule mission critical transfers e.g., for critical phases of a mission.

The value of completion of a transfer, i.e., semantic importance, may be a complex function varying with time. The term ‘semantic importance’ is used here because there have been previous references in the literature using this term to convey the meaning defined above. The RTMT, while recognizing that the word ‘semantic’ does not imply that the content and meaning of related messages are known to the RTCP, elected to use this term to avoid a proliferation of terms having similar meanings.

The task group has not reached consensus regarding the number of classes of semantic importance supported by the RTCP or the components of the relative value classification. Further definition will be included in ARD50033 (2.1.3).

4.1.3 Performance Requirements: The RTCP shall be fully deterministic; that is, its exact behavior can be predicted for any given set of conditions.

In order to meet deadlines, the RTCP shall guarantee the completion of requested services which have different latency requirements.

All aspects of the RTCP shall be schedulable.

RATIONALE: Since ‘real-time’ is a relative concept, it is not feasible to specify required performance numbers for the RTCP, either in the form of delay or bandwidth from User to User. These will vary from system to system, dependent on the application. The applications envisioned may have latency requirements in the order of hundreds of microseconds. However, it is necessary that the RTCP be predictable in order to guarantee its operation in all circumstances.

It is also necessary to influence the order in which service requests are processed; requests with lower latency requirements should be preferred over those with higher latency requirements. Alternatively, it is possible to define the order in which service requests are to be processed; higher priority requests must be serviced before those of lower priority. If the RTCP protocol processing was implemented as a FIFO with requests serviced in the order they were received, high priority messages may be stuck behind one of a lower priority, or a low latency message behind those with less strict requirements. Thus, prioritized access at the media level can be completely defeated if a time critical message sits in a queue behind a message that cannot gain access to the transmission medium.

4.1.4 Automatic Function Allocation: The RTCP shall support the allocation to RTCP processing of certain automatic functions as instructed by the User. The RTCP shall have the capability to execute these functions without specific interaction from the User for each initiation, or, at the time of the triggering event.

RATIONALE: For some of the transfers supported by the RTCP, it does not make sense to involve the User; the transfers are simple enough that they can be handled by the RTCP automatically once they have been set up. In addition, the RTCP may be able to handle them more efficiently. These automatic functions may include the periodic transmission of synchronization and timing information or messages communicating the current health of the system. While most of these are likely to be in the form of periodic messages, it is anticipated that some may be triggered by specific events; for instance, if so-and-so happens send an alert.

4.1.5 RTCP Layers: For purposes of developing implementations, the RTCP shall be considered as a single functional whole, rather than as an assemblage of interfacing layers.

RATIONALE: The ISO OSI/RM defines a functional model of a communications system. However, this has been used to define an implementation model, with specific protocols for each layer and real interfaces between each with corresponding performance penalties. To prevent the same problem from occurring with the RTCP, the functional model should be as simple as possible, to ensure the implementations of the RTCP are considered as an integrated whole, not a number of disjoint bits and pieces.

4.1.6 RTCP Implementation: The RTCP definition shall not prevent its implementation in either hardware or software.

RATIONALE: For performance reasons, part of the RTCP will most likely be implemented in hardware. However, for implementations that do not have tight performance requirements, the implementation through software may be more cost effective.

4.2 System Services:

4.2.1 Data Communication Services: Data Communication Services shall provide the means necessary to transmit Data Units between User entities.

The RTCP shall transmit the Data Units in a transparent fashion between these User entities. Its role is to free the User entities from having to supply transfer execution details. In particular, the RTCP shall ensure complete management of the transmission of Data Units.

The real-time communications system shall provide the User with two forms of data transfer services:

Category 1: The parameters for a series of transfers are given before any data is passed and thereafter the User provides data only.

Category 2: The data and parameters for the transfer are provided by the User for each transfer.

4.2.1 (Continued):

RATIONALE: From the viewpoint of a real-time communications system User, all data transfers can be considered to be connections with other Users, since resources will have to be allocated and the Users must co-operate whether implicitly or explicitly to transfer the data. The way in which the User chooses to control these data transfers, however, results in two sets of differing requirements for services offered by the real-time communications system and covered by these Real-Time Communication Protocols requirements. These two sets of services may be simply (but not definitively) differentiated by the Category 1 transfers having resources specifically assigned to each such that the User provides, or is provided, only data for a transfer, while the Category 2 transfers share some allocated resources such that the User must provide, or be provided with, data and transfer control information.

4.2.1.1 Category 1 Transfers: The RTCP shall provide Category 1 transfer services. These connection oriented transfers are characterized by having specific resources allocated in the communicating peer entities. Some consequences are:

a. Agreement to partake in the exchange of data is needed; there exists an effective co-operation between communicating entities.
b. No target address is required per message. Knowledge of the communicating entities belong to the ‘Connection’.
c. The connection has to be built with all its parameters before being used.

Two main classes of Category 1 transfers can be considered:

a. Static: These are characterized by parameter values which do not vary for any invocation. These Category 1 transfers may be in the Nonexistent, Dormant or Active state (see Figure 5).

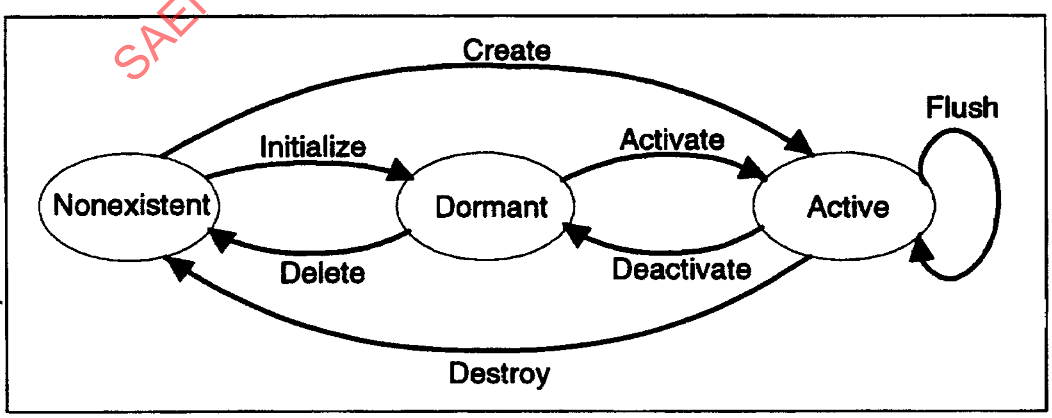

FIGURE 5 - Static Connections State Diagram

4.2.1.1 (Continued):

b. Dynamic: These are characterized by parameter values which may vary between invocations. These Category 1 transfers may be in the Nonexistent or Active state (see Figure 6).

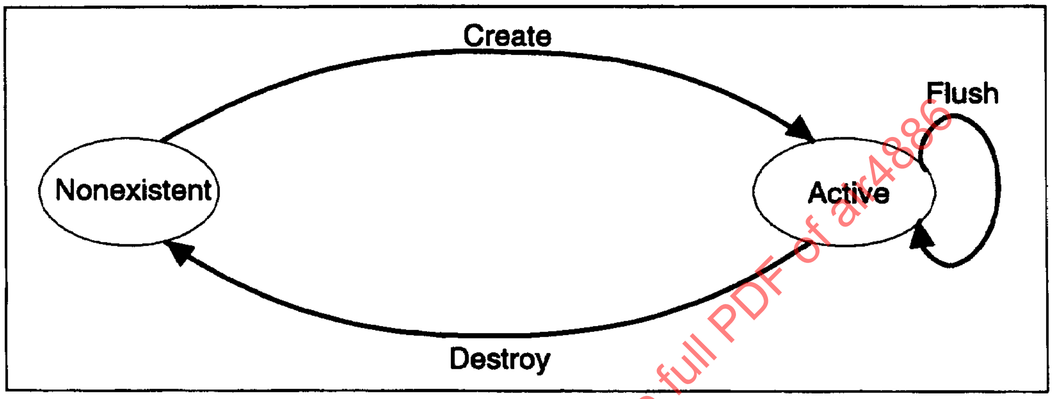

FIGURE 6 - Dynamic Connection State Diagram

A Category 1 transfer is Nonexistent when the parameter values are not explicitly known by the RTCP. In the Dormant state, the parameter values have been loaded and distributed through the RTCP, but all resources have not been allocated. In the Active state, all resources have been allocated and transfers may take place.

Both classes may be invoked for the life of the system or only for short periods during the system life.

RATIONALE: A good way to illustrate a Category 1 transfer is a telephone call. The request for a transfer is set up by dialing the number. Once the recipient answers and agrees to converse, no further addressing information is required; you simply talk into the telephone. Resources are allocated by the telephone company to ensure your connection (you hope).

When a particular Category 1 transfer is not required, it is beneficial to free the resources allocated to enable their use by other transfers. If the parameter values are fixed (Static), then some of the resources could remain allocated while others are utilized elsewhere (Dormant). Quick re-allocation from the dormant state can be achieved without resending the parameter values. If the parameter values may vary (Dynamic), then all resources must be de-allocated as different values may be used next time, requiring the parameter values to be sent again.

4.2.1.2 Category 2 Transfers: The RTCP shall provide Category 2 data transfer services. These connectionless transfers are characterized by having no specific resources allocated. Some consequences of this are immediate:

a. No agreement to partake in the exchange of data is made between communicating entities; there is a generic co-operation.
b. A target address is required per message.
c. All parameters required are supplied per message.
d. Local resources are shared in the communicating entities allowing diversity.

RATIONALE: A good way to illustrate a Category 2 data transfer is the mailing of a letter. The transfer is initiated by addressing the letter and putting it in the mailbox. This is (hopefully) transferred to the correct person who will (hopefully) read it. Nothing is known of the letter's progress unless specific instructions are included (e.g., registered delivery, RSVP).

It should be noted that, because the target applications of the real-time communications system are closed, the Category 2 transfers are designed into the system. Thus the probability of the transfers reaching the required User(s) may be 1, the same as for the Category 1 data transfers. Of course, to limit the resources required, these transfers may still be designed to reach the required User(s), with a probability of less than 1, if this is acceptable to the system.

4.2.1.3 Message Size: The real-time communications system shall allow the User to define a parameter for a Category 1 transfer defining the size of the message to be passed. It is desired that the bounds of this parameter be unrestricted. A System Design parameter shall also be defined restricting the maximum message size of a Category 2 transfer.

RATIONALE: The message size parameters may be applied to either Category 1 or Category 2 transfers. For Category 1 transfers, the message size will be defined before any transfers take place to allow the resources (memory etc.) to be allocated. In this way, the message size of a Category 1 transfer is unlimited, provided the required resources can be allocated. These resources must be allocated in all Users of the particular Category 1 transfer.

Since the Category 2 transfer does not define its resource requirements beforehand, the maximum size of the message may be limited, firstly by the resources allocated to be shared by the Category 2 transfers and secondly by the amount of these resources already being utilized by other Category 2 transfers. Since there is no co-operation between Users, different Users may not be able to allocate the necessary resources when the data transfer is received. For this reason, there is often a System Design message size parameter applied to the Category 2 transfers generally to limit the resources required by any individual transfer.

4.2.1.4 Packet Size: The real-time communications system shall provide support for a System Design parameter restricting the size of packets transmitted on the physical network. It is desired that the maximum size of a Category 2 message shall be independent of this parameter.

RATIONALE: The maximum packet size is a system design parameter and so applies to both types of transfers. While the User may define the size of a message for either type of transfer to be as big as is required, other constraints may be applied to the way in which this message is transferred through the real-time communication system. One example of this is the data bus packet size, which defines the maximum size of an information unit for transfer across the data bus. This system design parameter may be required to limit the maximum size to less than that of the data bus standard to ensure latency requirements are maintained.

Messages larger than this size of either transfer type will be segmented into a number of data bus packets and reassembled into the message at the receiving node.

The real-time communication system should be capable of segmenting and reassembling Category 2 data transfers; otherwise the reduction of the maximum packet size may require a complete system redesign because of the reduced maximum size of these transfers. Thus, the two parameters need to be independent.

4.2.1.5 Flow Control: The real-time communication system shall allow the User to request flow control for a Category 1 data transfer.

RATIONALE: Flow control of Category 1 data transfers allows the receiving User to prevent the data-providing User from sending further data until such a time when it may be accommodated. Flow control cannot exist for Category 2 data transfers because the receiving User does not know that further data is to be sent. For this transfer category it is acceptable to ignore data received but not for Category 1 transfers.

The flow control between Users is concerned with resources outside the real-time communications system. Thus, the real-time communications system parameters do not apply to the use of flow control and the delay of information could be infinite; it is up to the system designer to define what is allowable. The real-time communications system is just providing a service to the receiving User to exercise flow control without the real-time communications system requiring extra resources to accept data from the data-providing User without passing it on to the receiving User simply because that User is not prepared to accept it.

4.2.1.5 (Continued):

Resources are allocated within the RTCP to transfer the data placed in the interface by User A according to the parameters defined when the Category 1 transfer was activated. See Figure 7. User B must accept data from the RTCP as was also part of the agreement. However, User B may not be able to handle the instantaneous data rate supplied by User A. For example, if the rate defined was 1000 blocks per second, it may be that the 1000 blocks occur in 1 ms followed by 999 ms silence. If User B can only handle 10 blocks at a time then, having received 8 blocks (assume 2 more in transit) in a short period of time, User B uses the flow control to prevent further transfers until some of the buffer space has been freed.

Flow control may also be required by layers within the real-time communication system for similar reasons. However, the layers will also have restrictions placed upon them by other data transfer parameters (e.g., latency) such that the maximum extent to which flow control may be used is limited. Thus sufficient resources will have to be supplied to support the transfer.

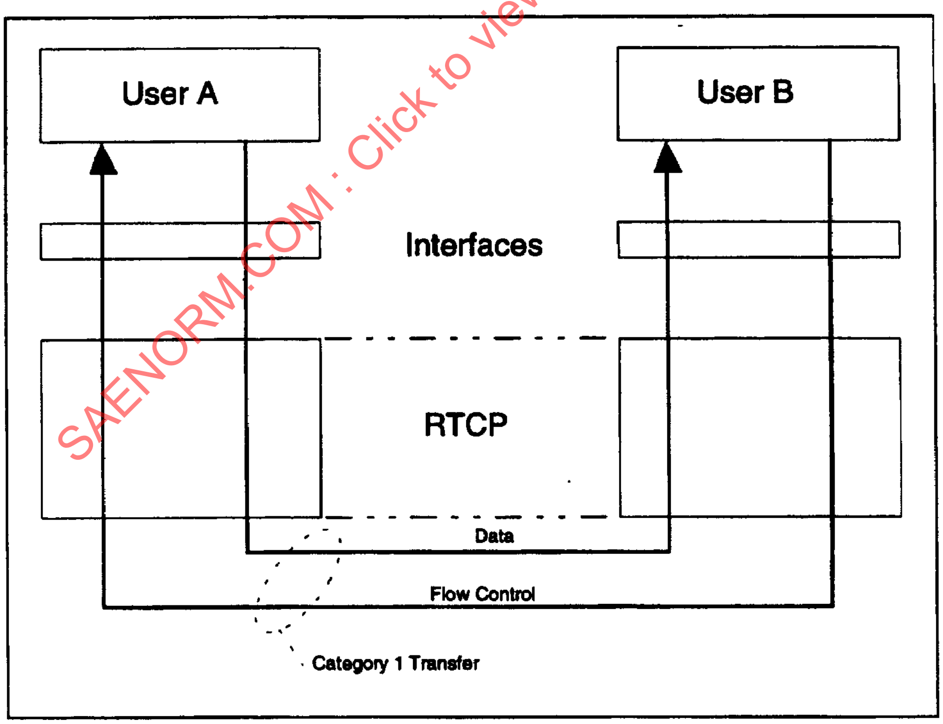

FIGURE 7 - Flow Control

4.2.1.6 Flow Restriction: It is desirable that the real-time communications system shall allow the User to define a parameter for a Category 1 transfer defining the bandwidth of the series of messages to be passed. In which case, the User shall be restricted to this message flow bandwidth in subsequent transfers; the real-time communications system shall be capable of enforcing this flow restriction and shall reserve sufficient network bandwidth to guarantee the transfer of this message flow.

It is also desirable that a system design parameter be defined restricting the maximum message bandwidth of all Category 2 transfers.

If these two desires cannot be implemented, it is required that the real-time communications system report to the User when excessive bandwidth requests result in the inability to meet all the data transfer parameters.

RATIONALE: Flow control helps a receiving User manage its own resources, while flow restriction prevents a data-providing User from hogging more than its fair share of a global resource (i.e., the RTCP).

While the Category 1 data transfer may be considered as a pipe from one User to another, the use of a multiplex network means that some resources (in this case network bandwidth) must be shared and will not be allocated solely to one data transfer.

The use of flow restriction is mainly applicable to Category 1 data transfers and is used to define the communication bandwidth resource which is allocated to the particular transfer. While it is normally only considered as limiting the rate at which data may be transferred, it should also guarantee that the requested bandwidth is always available. Thus, each Category 1 transfer must compete with other transfers for the available bandwidth resource. Flow restriction is especially relevant for the lower layers of the real-time communications system in the case of a transfer where the message size is greater than the network packet size for the system. This will result in multiple packets being queued for transmission which could completely swamp the network. The flow restriction would allow the packets to only be presented for transfer at a particular rate, thus allowing other transfers to take place.

The use of flow restriction may also be applied to Category 2 transfers, but only globally to all transfers of this category, not to individual transfers. The amount of bandwidth allocated to the Category 2 transfers would be allocated and the individual transfers would vie to use it.

The main problem of defining and implementing flow restriction is in the measurement of data bandwidth used/available; this is especially true for the Category 2 data transfers, the network traffic of which must be monitored.

From the above it is clear that flow restriction may not be feasible but is desirable.

4.2.1.7 Acknowledged and Nonacknowledged Services: It is required that the RTCP provide support for both acknowledged and nonacknowledged Category 1 and Category 2 communications. This acknowledgement denotes that the information has been copied into the receiving User's buffer (Ack2), see Figure 8.

RATIONALE: Information transfer acknowledgements are only guaranteed if between Users. This means that acknowledgement is initiated when the information is presented to the destination User, not just when it enters the destination User node. User-to-User acknowledgement does not normally impact the RTCP, although it may be implemented using RTCP functions as follows.

Within the RTCP, acknowledgement services will be required (e.g., to confirm the transfer of information for a Category 1 transfer.) These acknowledgements will be local to the RTCP and not visible to the data-providing User (Ack1), see Figure 8.

Use may be made of this acknowledgement (Ack1) to provide a notification to the data-providing User that the information has been transferred to the required User/RTCP interface, i.e., the point at which the RTCP has completed the information transfer (Ack2).

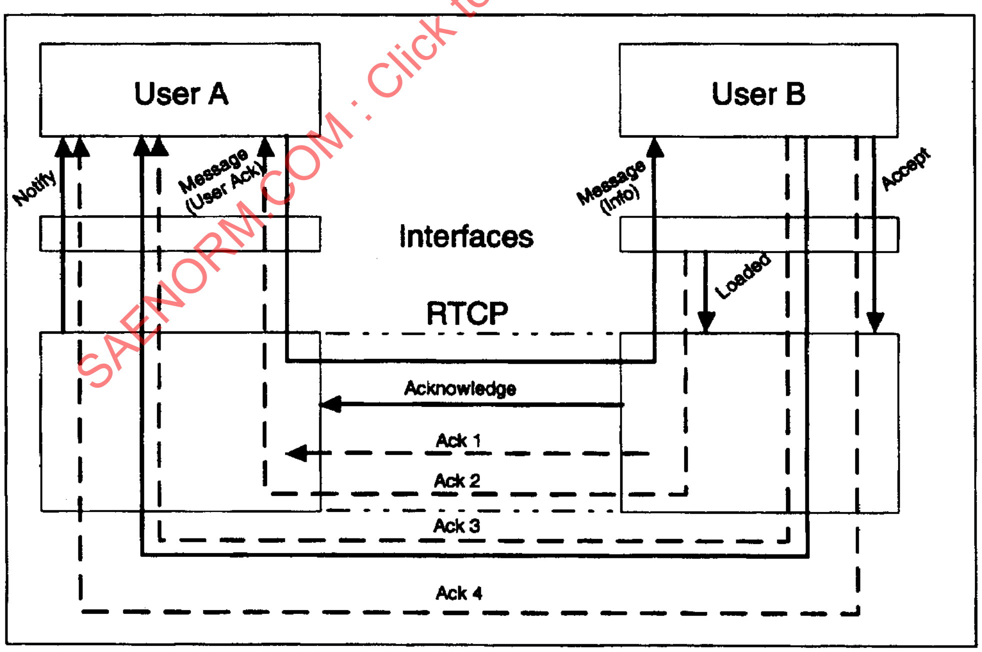

FIGURE 8 - Types of Acknowledgements

4.2.1.7 (Continued):

The notification (Ack2) does not indicate that the receiving User has accepted or used the information (Ack3). This type of acknowledgement is normally considered to be provided by the destination User to indicate that the information is being used. This message will not be recognized as an acknowledgement by the RTCP.

However, this acknowledgement (Ack3) may also be provided using the RTCP acknowledgement (Ack1), shown as Ack4. To do this, the receiving User must indicate its acceptance of the information to its local RTCP entity, this can be acknowledged (Ack1) to the data-providing User's local RTCP entity, which will finally notify the data-providing User.

4.2.2 Synchronization Service: It is required that the RTCP provide a mechanism that allows the synchronization of tasks processed by the application layer distributed at different remote User entity physical locations.

This system synchronization service can be accomplished by means of remote interrupts (4.2.2.1) or by means of a common knowledge of time (4.2.2.2).

4.2.2.1 Remote Interrupts: It is required that the RTCP provides a remote interrupt service that ensures the transfer of an interrupt vector to a destination entity or to a group of destination entities. It is required that the invocation of this service can be allocated to the RTCP (automatic remote interrupt.)

This does not imply a specific mechanism to implement the service.

RATIONALE: The following provides examples of how interrupts might be used and shows, for instance, that interrupts at the RTCP level might be used by the User:

Remote interrupts to a destination User entity can be implemented via an interrupt request primitive from a User that is the source of the interrupt (Int1), see Figure 9.

In addition, remote interrupts can originate in the RTCP to a remote RTCP entity (Int5).

A User can also request an interrupt to a remote RTCP entity (Int3), or instruct the service-providing RTCP to provide this via automatic invocation (Int4).

In addition, a User can instruct the RTCP to provide this interrupt to a remote User entity (Int2). A User can instruct the RTCP to provide this remote interrupt service as a result of certain events in the network or periodically without involvement of the requesting User at the time, i.e., an automatic invocation.

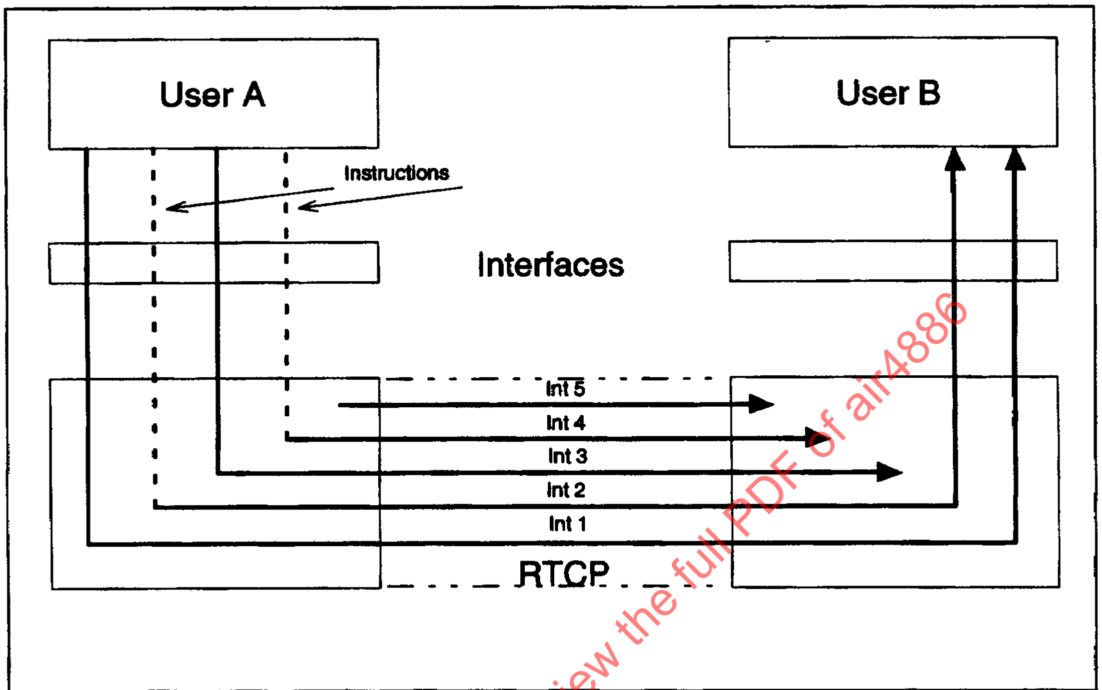

FIGURE 9 - Types of Remote Interrupts

4.2.2.2 Time Distribution: It is required that the RTCP provide a mechanism to initialize and maintain global time to be made available at the distributed RTCP/User interface.

RATIONALE: The distribution of time within a system may be useful for process synchronization, event ordering, etc. However, this does not imply that network time is required for these functions. These functions could be implemented in a number of ways, including special network codes, interrupts, and periodic RTCP messages. The actual method used will depend on the accuracy and stability of time reference required (if at all).

4.2.2.3 Related Data Streams: The RTCP shall provide support to synchronize related data streams.

RATIONALE: Certain related information may be transferred on separate data paths, such as voice and video. If so, synchrony of the sound with the image may not be given. Hence the RTCP must provide methods for reconstituting the timely delivery of related subsets of information.

4.2.3 Management Services: It is required that the RTCP shall support a distributed control mechanism which invokes normal and preplanned exception configurations. The RTCP shall also support a centralized control mechanism. The RTCP shall monitor status and keep fault logs for itself. It shall also support the handling of abnormal system states and exceptions. The RTCP shall support uninterrupted operational performance.