

Network Working Group
Request for Comments: 5401
Obsoletes: 3941
Category: Standards Track

B. Adamson
Naval Research Laboratory
C. Bormann
Universitaet Bremen TZI
M. Handley
University College London
J. Macker
Naval Research Laboratory
November 2008

Multicast Negative-Acknowledgment (NACK) Building Blocks

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (c) 2008 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

This document discusses the creation of reliable multicast protocols that utilize negative-acknowledgment (NACK) feedback. The rationale for protocol design goals and assumptions are presented. Technical challenges for NACK-based (and in some cases general) reliable multicast protocol operation are identified. These goals and challenges are resolved into a set of functional "building blocks" that address different aspects of reliable multicast protocol operation. It is anticipated that these building blocks will be useful in generating different instantiations of reliable multicast protocols. This document obsoletes RFC 3941.

Table of Contents

1. Introduction	2
1.1. Requirements Language	4
2. Rationale	4
2.1. Delivery Service Model	5
2.2. Group Membership Dynamics	6
2.3. Sender/Receiver Relationships	6
2.4. Group Size Scalability	6
2.5. Data Delivery Performance	7
2.6. Network Environments	7
2.7. Intermediate System Assistance	8
3. Functionality	8
3.1. Multicast Sender Transmission	11
3.2. NACK Repair Process	13
3.3. Multicast Receiver Join Policies and Procedures	26
3.4. Node (Member) Identification	26
3.5. Data Content Identification	27
3.6. Forward Error Correction (FEC)	28
3.7. Round-Trip Timing Collection	29
3.8. Group Size Determination/Estimation	33
3.9. Congestion Control Operation	34
3.10. Intermediate System Assistance	34
4. NACK-Based Reliable Multicast Applicability	35
5. Security Considerations	36
6. Changes from RFC 3941	38
7. Acknowledgements	38
8. References	39
8.1. Normative References	39
8.2. Informative References	39

1. Introduction

Reliable multicast transport is a desirable technology for efficient and reliable distribution of data to a group on the Internet. The complexities of group communication paradigms necessitate different protocol types and instantiations to meet the range of performance and scalability requirements of different potential reliable multicast applications and users (see [RFC2357]). This document addresses the creation of reliable multicast protocols that utilize negative-acknowledgment (NACK) feedback. NACK-based protocols generally entail less frequent feedback messaging than reliability protocols based on positive acknowledgment (ACK). The less frequent feedback messaging helps simplify the problem of feedback implosion as group size grows larger. While different protocol instantiations may be required to meet specific application and network architecture demands [ArchConsiderations], there are a number of fundamental components that may be common to these different instantiations.

This document describes the framework and common "building block" components relevant to multicast protocols that are based primarily on NACK operation for reliable transport. While this document discusses a large set of reliable multicast components and issues relevant to NACK-based reliable multicast protocol design, it specifically addresses in detail the following building blocks, which are not addressed in other IETF documents:

1. NACK-based multicast sender transmission strategies,
2. NACK repair process with timer-based feedback suppression, and
3. Round-trip timing for adapting NACK and other timers.

NACK-based reliable multicast implementations SHOULD make use of Forward Error Correction (FEC) erasure coding techniques, as described in the FEC Building Block [RFC5052] document. Packet-level erasure coding allows missing packets from a given FEC block to be recovered using the parity packets instead of classical, individualized retransmission of original source data content. For this reason, this document refers to the protocol mechanisms for reliability as a "repair process." Note that NACK-based protocols can reactively provide the parity packets in response to receiver requests for repair rather than just proactively sending added FEC parity content as part of the original transmission. Hybrid proactive/reactive use of FEC content is also possible with the mechanisms described in this document. Some classes of FEC coding, such as Maximal Separable Distance (MDS) codes, allow senders to dynamically implement deterministic, highly efficient receiver group repair strategies as part of a NACK-based, selective automated repeat-request (ARQ) scheme.

The potential relationships to other reliable multicast transport building blocks (e.g., FEC, congestion control) and general issues with NACK-based reliable multicast protocols are also discussed. This document follows the guidelines provided in [RFC3269].

Statement of Intent

This memo contains descriptions of building blocks that can be applied in the design of reliable multicast protocols utilizing negative-acknowledgement (NACK) feedback. [RFC3941] contains a previous description of this specification. RFC 3941 was published in the "Experimental" category. It was the stated intent of the Reliable Multicast Transport (RMT) working group at that time to resubmit this specification as an IETF Proposed Standard in due course.

This Proposed Standard specification is thus based on [RFC3941] and has been updated according to accumulated experience and growing protocol maturity since the publication of RFC 3941. Said experience applies both to this specification itself and to congestion control strategies related to the use of this specification.

The differences between [RFC3941] and this document are listed in Section 6.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Rationale

Each potential protocol instantiation using the building blocks presented here (and in other applicable building block documents) will have specific criteria that may influence individual protocol design. To support the development of applicable building blocks, it is useful to identify and summarize driving general protocol design goals and assumptions. These are areas that each protocol instantiation will need to address in detail. Each building block description in this document will include a discussion of the impact of these design criteria. The categories of design criteria considered here include:

1. Delivery Service Model,
2. Group Membership Dynamics,
3. Sender/Receiver Relationships,
4. Group Size Scalability,
5. Data Delivery Performance, and
6. Network Environments.

All of these areas are at least briefly discussed. Additionally, other reliable multicast transport building block documents, such as [RFC5052], have been created to address areas outside of the scope of this document. NACK-based reliable multicast protocol instantiations may depend upon these other building blocks as well as the ones presented here. This document focuses on areas that are unique to NACK-based reliable multicast but may be used in concert with the other building block areas. In some cases, a building block may be

able to address a wide range of assumptions, while in other cases there will be trade-offs required to meet different application needs or operating environments. Where necessary, building block features are designed to be parametric to meet different requirements. Of course, an underlying goal will be to minimize design complexity and to at least recommend default values for any such parameters that meet a general purpose "bulk data transfer" requirement in a typical Internet environment. The forms of "bulk data transfer" covered here include reliable transport of bulky, fixed-length, a priori static content and also transmission of non-predetermined, perhaps streamed, content of indefinite length. Section 3.5 discusses these different forms of bulk data content in further detail.

2.1. Delivery Service Model

The implicit goal of a reliable multicast transport protocol is the reliable delivery of data among a group of members communicating using IP multicast datagram service. However, the specific service the application is attempting to provide can impact design decisions. The most basic service model for reliable multicast transport is that of "bulk transfer", which is a primary focus of this and other related RMT working group documents. However, the same principles in protocol design may also be applied to other service models, e.g., more interactive exchanges of small messages such as with whiteboarding or text chat. Within these different models there are issues such as the sender's ability to cache transmitted data (or state referencing it) for retransmission or repair. The needs for ordering and/or causality in the sequence of transmissions and receptions among members in the group may be different depending upon data content. The group communication paradigm differs significantly from the point-to-point model in that, depending upon the data content type, some receivers may complete reception of a portion of data content and be able to act upon it before other members have received the content. This may be acceptable (or even desirable) for some applications but not for others. These varying requirements drive the need for a number of different protocol instantiation designs. A significant challenge in developing generally useful building block mechanisms is accommodating even a limited range of these capabilities without defining specific application-level details.

Another factor impacting the delivery service model is the potential for different receivers in the multicast group to have significantly differing quality of network connectivity. This may involve receivers with very limited goodput due to connection rate or substantial packet loss. NACK-based protocol implementations may wish to provide policies by which extremely poor-performing receivers are excluded from the main group or migrated to a separate delivery

group. Note that some application models may require that the entire group be constrained to the performance of the "weakest member" to satisfy operational requirements. In either case, protocol designs should consider this aspect of the reliable multicast delivery service model.

2.2. Group Membership Dynamics

One area where group communication can differ from point-to-point communications is that even if the composition of the group changes, the "thread" of communication can still exist. This contrasts with the point-to-point communication model where, if either of the two parties leave, the communication process (exchange of data) is terminated (or at least paused). Depending upon application goals, senders and receivers participating in a reliable multicast transport "session" may be able to join late, leave, and/or potentially rejoin while the ongoing group communication "thread" still remains functional and useful. Also note that this can impact protocol message content. If "late joiners" are supported, some amount of additional information may be placed in message headers to accommodate this functionality. Alternatively, the information may be sent in its own message (on demand or intermittently) if the impact to the overhead of typical message transmissions is deemed too great. Group dynamics can also impact other protocol mechanisms such as NACK timing, congestion control operation, etc.

2.3. Sender/Receiver Relationships

The relationship of senders and receivers among group members requires consideration. In some applications, there may be a single sender multicasting to a group of receivers. In other cases, there may be more than one sender or the potential for everyone in the group to be a sender and receiver of data may exist.

2.4. Group Size Scalability

Native IP multicast [RFC1112] may scale to extremely large group sizes. It may be desirable for some applications to scale along with the multicast infrastructure's ability to scale. In its simplest form, there are limits to the group size to which a NACK-based protocol can be applied without the potential for the volume of NACK feedback messages to overwhelm network capacity. This is often referred to as "feedback implosion". Research suggests that NACK-based reliable multicast group sizes on the order of tens of thousands of receivers may operate with acceptable levels of feedback to the sender using probabilistic, timer-based suppression techniques [NormFeedback]. Instead of receivers immediately transmitting feedback messages when loss is detected, these techniques specify use

of purposefully-scaled, random back-off timeouts such that some potential NACKing receivers can self-suppress their feedback upon hearing messages from other receivers that have selected shorter random timeout intervals. However, there may be additional NACK suppression heuristics that can be applied to enable these protocols to scale to even larger group sizes. In large scale cases, it may be prohibitive for members to maintain state on all other members (in particular, other receivers) in the group. The impact of group size needs to be considered in the development of applicable building blocks.

Group size scalability may also be aided by intermediate system assistance; see section 2.7 below.

2.5. Data Delivery Performance

There is a trade-off between scalability and data delivery latency when designing NACK-oriented protocols. If probabilistic, timer-based NACK suppression is to be used, there will be some delays built into the NACK process to allow suppression to occur and to allow the sender of data to identify appropriate content for efficient repair transmission. For example, back-off timeouts can be used to ensure efficient NACK suppression and repair transmission, but this comes at the cost of increased delivery latency and increased buffering requirements for both senders and receivers. The building blocks SHOULD allow applications to establish bounds for data delivery performance. Note that application designers must be aware of the scalability trade-off that is made when such bounds are applied.

2.6. Network Environments

The Internet Protocol has historically assumed a role of providing service across heterogeneous network topologies. It is desirable that a reliable multicast protocol be capable of effectively operating across a wide range of the networks to which general purpose IP service applies. The bandwidth available on the links between the members of a single group today may vary between low numbers of kbit/s for wireless links and multiple Gbit/s for high speed LAN connections, with varying degrees of contention from other flows. Recently, a number of asymmetric network services including 56K/ADSL modems, CATV Internet service, satellite, and other wireless communication services have begun to proliferate. Many of these are inherently broadcast media with potentially large "fan-out" to which IP multicast service is highly applicable. Additionally, policy and/or technical issues may result in topologies where multicast connectivity is limited to a source-specific multicast (SSM) model from a specific source [RFC4607]. Receivers in the group may be

restricted to unicast feedback for NACKs and other messages. Consideration must be given, in building block development and protocol design, to the nature of the underlying networks.

2.7. Intermediate System Assistance

Intermediate assistance from devices/systems with direct knowledge of the underlying network topology may be used to increase the performance and scalability of NACK-based reliable multicast protocols. Feedback aggregation and filtering of sender repair data may be possible with NACK-based protocols using FEC-based repair strategies as described in the present and other reliable multicast transport building block documents. However, there will continue to be a number of instances where intermediate system assistance is not available or practical. Any building block components for NACK-oriented reliable multicast SHALL be capable of operating without such assistance. However, it is RECOMMENDED that such protocols also consider utilizing these features when available.

3. Functionality

The previous section has presented the role of protocol building blocks and some of the criteria that may affect NACK-based reliable multicast building block identification/design. This section describes different building block areas applicable to NACK-based reliable multicast protocols. Some of these areas are specific to NACK-based protocols. Detailed descriptions of such areas are provided. In other cases, the areas (e.g., node identifiers, forward error correction (FEC), etc.) may be applicable to other forms of reliable multicast. In those cases, the discussion below describes requirements placed on those general building block areas from the standpoint of NACK-based reliable multicast. Where applicable, other building block documents are referenced for possible contribution to NACK-based reliable multicast protocols.

For each building block, a notional "interface description" is provided to illustrate any dependencies of one building block component upon another or upon other protocol parameters. A building block component may require some form of "input" from another building block component or other source to perform its function. Any "inputs" required by a building block component and/or any resultant "output" provided will be defined and described in each building block component's interface description. Note that the set of building blocks presented here do not fully satisfy each other's "input" and "output" needs. In some cases, "inputs" for the building blocks here must come from other building blocks external to this document (e.g., congestion control or FEC). In other cases NACK-

based reliable multicast building block "inputs" must be satisfied by the specific protocol instantiation or implementation (e.g., application data and control).

The following building block components relevant to NACK-based reliable multicast are identified:

NORM (NACK-Oriented Reliable Multicast)-Specific

1. Multicast Sender Transmission
2. NACK Repair Process
3. Multicast Receiver Join Policies and Procedures

General Purpose

1. Node (Member) Identification
2. Data Content Identification
3. Forward Error Correction (FEC)
4. Round-Trip Timing Collection
5. Group Size Determination/Estimation
6. Congestion Control Operation
7. Intermediate System Assistance
8. Ancillary Protocol Mechanisms

Figure 1 provides a pictorial overview of these building block areas and some of their relationships. For example, the content of the data messages that a sender initially transmits depends upon the "Node Identification", "Data Content Identification", and "FEC" components, while the rate of message transmission will generally depend upon the "Congestion Control" component. Subsequently, the receivers' response to these transmissions (e.g., NACKing for repair) will depend upon the data message content and inputs from other building block components. Finally, the sender's processing of receiver responses will feed back into its transmission strategy.

The components on the left side of this figure are areas that may be applicable beyond NACK-based reliable multicast. The more significant of these components are discussed in other building block documents, such as the FEC Building Block [RFC5052]. Brief

descriptions of these areas and their roles in NACK-based reliable multicast protocols are given below, and "RTT Collection" is discussed in detail in Section 3.7 of this document.

The components on the right are seen as specific to NACK-based reliable multicast protocols, most notably the NACK repair process. These areas are discussed in detail below (most notably, "Multicast Sender Transmission" and "NACK Repair Process" in Sections 3.1 and 3.2). Some other components (e.g., "Security") impact many aspects of the protocol, and others may be more transparent to the core protocol processing. Where applicable, specific technical recommendations are made for mechanisms that will properly satisfy the goals of NACK-based reliable multicast transport for the Internet.

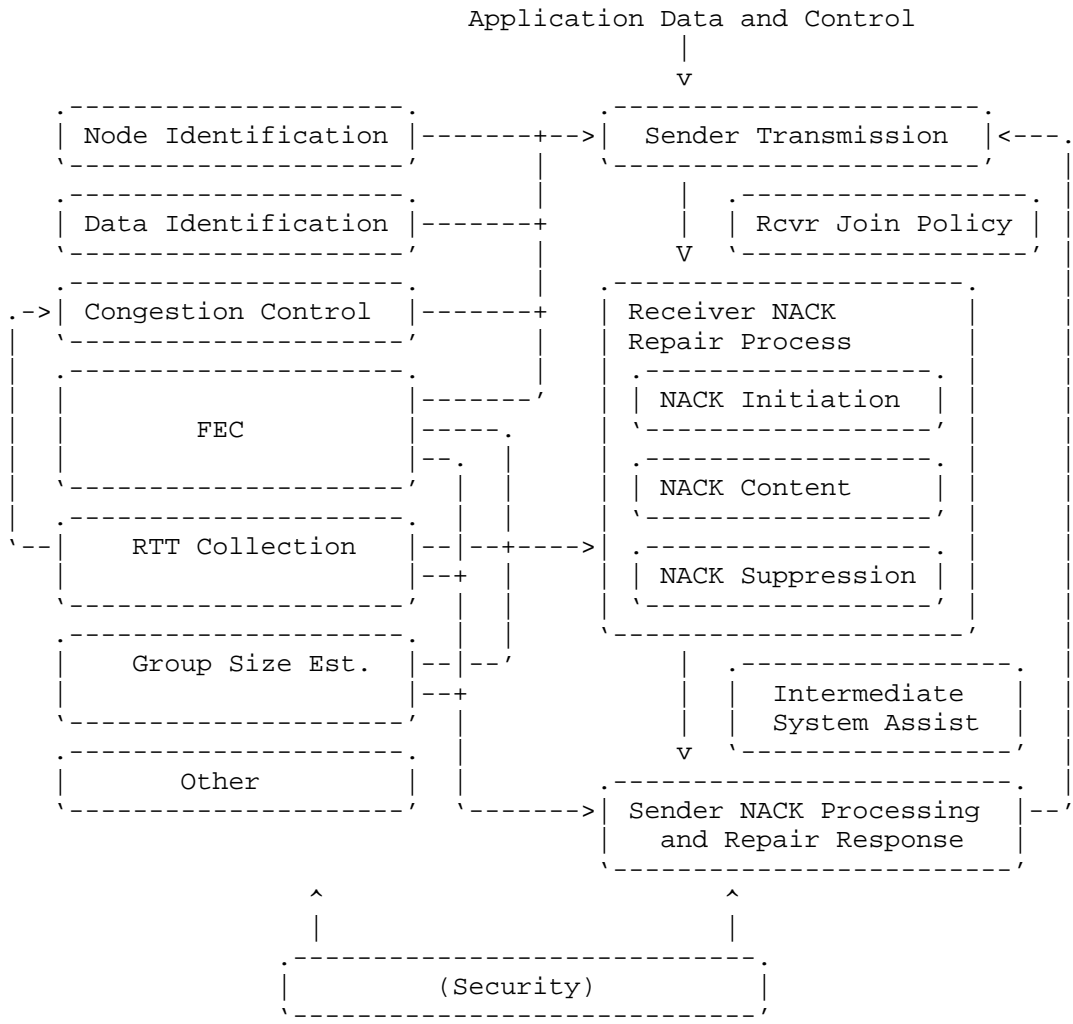


Figure 1: NACK-Based Reliable Multicast Building Block Framework

3.1. Multicast Sender Transmission

Reliable multicast senders will transmit data content to the multicast session. The data content will be application dependent. The sender will transmit data content at a rate, and with message sizes, determined by application and/or network architecture requirements. Any FEC encoding of sender transmissions SHOULD conform with the guidelines of the FEC Building Block [RFC5052]. When congestion control mechanisms are needed (REQUIRED for general Internet operation), the sender transmission rate SHALL be controlled

by the congestion control mechanism. In any case, it is RECOMMENDED that all data transmissions from multicast senders be subject to rate limitations determined by the application or congestion control algorithm. The sender's transmissions SHOULD make good utilization of the available capacity (which may be limited by the application and/or by congestion control). As a result, it is expected there will be overlap and multiplexing of new data content transmission with repair content. Other factors related to application operation may determine sender transmission formats and methods. For example, some consideration needs to be given to the sender's behavior during intermittent idle periods when it has no data to transmit.

In addition to data content, other sender messages or commands may be employed as part of protocol operation. These messages may occur outside of the scope of application data transfer. In NACK-based reliable multicast protocols, reliability of such protocol messages may be attempted by redundant transmission when positive acknowledgement is prohibitive due to group size scalability concerns. Note that protocol design SHOULD provide mechanisms for dealing with cases where such messages are not received by the group. As an example, a command message might be redundantly transmitted by a sender to indicate that it is temporarily (or permanently) halting transmission. At this time, it may be appropriate for receivers to respond with NACKs for any outstanding repairs they require, following the rules of the NACK procedure. For efficiency, the sender should allow sufficient time between the redundant transmissions to receive any NACK responses from the receivers to this command.

In general, when there is any resultant NACK or other feedback operation, the timing of redundant transmission of control messages issued by a sender and other NACK-based reliable multicast protocol timeouts should be dependent upon the group greatest round-trip timing (GRTT) estimate and any expected resultant NACK or other feedback operation. The sender GRTT is an estimate of the worst-case round-trip timing from a given sender to any receivers in the group. It is assumed that the GRTT interval is a conservative estimate of the maximum span (with respect to delay) of the multicast group across a network topology with respect to a given sender. NACK-based reliable multicast instantiations SHOULD be able to dynamically adapt to a wide range of multicast network topologies.

Inputs:

1. Application data and control.
2. Sender node identifier.

3. Data identifiers.
4. Segmentation and FEC parameters.
5. Transmission rate.
6. Application controls.
7. Receiver feedback messages (e.g., NACKs).

Outputs:

1. Controlled transmission of messages with headers uniquely identifying data or repair content within the context of the reliable multicast session.
2. Commands indicating sender's status or other transport control actions to be taken.

3.2. NACK Repair Process

A critical component of NACK-based reliable multicast protocols is the NACK repair process. This includes both the receiver's role in detecting and requesting repair needs and the sender's response to such requests. There are four primary elements of the NACK repair process:

1. Receiver NACK process initiation,
2. NACK suppression,
3. NACK message content,
4. Sender NACK processing and repair response.

3.2.1. Receiver NACK Process Initiation

The NACK process (cycle) will be initiated by receivers that detect a need for repair transmissions from a specific sender to achieve reliable reception. When FEC is applied, a receiver should initiate the NACK process only when it is known its repair requirements exceed the amount of pending FEC transmission for a given coding block of data content. This can be determined at the end of the current transmission block (if it is indicated) or upon the start of reception of a subsequent coding block or transmission object. This implies the sender data content is marked to identify its FEC block number and that ordinal relationship is preserved in order of transmission.

Alternatively, if the sender's transmission advertises the quantity of repair packets it is already planning to send for a block, the receiver may be able to initiate the NACK process earlier. Allowing receivers to initiate NACK cycles at any time they detect their repair needs have exceeded pending repair transmissions may result in slightly quicker repair cycles. However, it may be useful to limit NACK process initiation to specific events, such as at the end-of-transmission of an FEC coding block or upon detection of subsequent coding blocks. This can allow receivers to aggregate NACK content into a smaller number of NACK messages and provide some implicit loose synchronization among the receiver set to help facilitate effective probabilistic suppression of NACK feedback. The receiver MUST maintain a history of data content received from the sender to determine its current repair needs. When FEC is employed, it is expected that the history will correspond to a record of pending or partially-received coding blocks.

For probabilistic, timer-based suppression of feedback, the NACK cycle should begin with receivers observing backoff timeouts. In conjunction with initiating this backoff timeout, it is important that the receivers record the position in the sender's transmission sequence at which they initiate the NACK cycle. When the suppression backoff timeout expires, the receivers should only consider their repair needs up to this recorded transmission position in making the decision to transmit or suppress a NACK. Without this restriction, suppression is greatly reduced as additional content is received from the sender during the time a NACK message propagates across the network to the sender and other receivers.

Inputs:

1. Sender data content with sequencing identifiers from sender transmissions.
2. History of content received from sender.

Outputs:

1. NACK process initiation decision.
2. Recorded sender transmission sequence position.

3.2.2. NACK Suppression

An effective feedback suppression mechanism is the use of random backoff timeouts prior to NACK transmission by receivers requiring repairs [SrmFramework]. Upon expiration of the backoff timeout, a receiver will request repairs unless its pending repair needs have

been completely superseded by NACK messages heard from other receivers (when receivers are multicasting NACKs) or from some indicator from the sender. When receivers are unicasting NACK messages, the sender may facilitate NACK suppression by forwarding a representation of NACK content it has received to the group at large or by providing some other indicator of the repair information it will be subsequently transmitting.

For effective and scalable suppression performance, the backoff timeout periods used by receivers should be independently, randomly picked by receivers with a truncated exponential distribution [McastFeedback]. This results in the majority of the receiver set holding off transmission of NACK messages under the assumption that the smaller number of "early NACKers" will supersede the repair needs of the remainder of the group. The mean of the distribution should be determined as a function of the current estimate of the sender's GRTT assessment and a group size estimate that is either determined by other mechanisms within the protocol or is preset by the multicast application.

A simple algorithm can be constructed to generate random backoff timeouts with the appropriate distribution. Additionally, the algorithm may be designed to optimize the backoff distribution given the number of receivers ("R") potentially generating feedback. This "optimization" minimizes the number of feedback messages (e.g., NACK) in the worst-case situation where all receivers generate a NACK. The maximum backoff timeout ("T_maxBackoff") can be set to control reliable delivery latency versus volume of feedback traffic. A larger value of "T_maxBackoff" will result in a lower density of feedback traffic for a given repair cycle. A smaller value of "T_maxBackoff" results in shorter latency, which also reduces the buffering requirements of senders and receivers for reliable transport.

In the functions below, the "log()" function specified refers to the "natural logarithm" and the "exp()" function is similarly based upon the mathematical constant 'e' (a.k.a. Euler's number) where "exp(x)" corresponds to 'e' raised to the power of 'x'. Given the receiver group size ("groupSize") and maximum allowed backoff timeout ("T_maxBackoff"), random backoff timeouts ("t") with a truncated exponential distribution can be picked with the following algorithm:

1. Establish an optimal mean ("L") for the exponential backoff based on the "groupSize":

$$L = \log(\text{groupSize}) + 1$$

- Pick a random number ("x") from a uniform distribution over a range of:

$$\frac{L}{T_maxBackoff * (\exp(L) - 1)} \text{ to } \frac{L}{T_maxBackoff * (\exp(L) - 1)} + \frac{L}{T_maxBackoff}$$

- Transform this random variate to generate the desired random backoff time ("t'") with the following equation:

$$t' = T_maxBackoff / L * \log(x * (\exp(L) - 1) * (T_maxBackoff / L))$$

This "C" language function can be used to generate an appropriate random backoff time interval:

```
double RandomBackoff(double T_maxBackoff, double groupSize)
{
    double lambda = log(groupSize) + 1;
    double x = UniformRand(lambda/T_maxBackoff) +
              lambda / (T_maxBackoff*(exp(lambda)-1));
    return ((T_maxBackoff/lambda) *
            log(x*(exp(lambda)-1)*(T_maxBackoff/lambda)));
} // end RandomBackoff()
```

where "UniformRand(double max)" returns random numbers with a uniform distribution from the range of "0..max". For example, based on the POSIX "rand()" function, the following "C" code can be used:

```
double UniformRand(double max)
{
    return (max * ((double)rand()/(double)RAND_MAX));
}
```

The number of expected NACK messages generated ("N") within the first round-trip time for a single feedback event is approximately:

$$N = \exp(1.2 * L / (2 * T_maxBackoff / GRTT))$$

Thus, the maximum backoff time can be adjusted to trade off worst-case NACK feedback volume versus latency. This is derived from the equations given in [McastFeedback] and assumes "T_maxBackoff >= GRTT", and "L" is the mean of the distribution optimized for the given group size as shown in the algorithm above. Note that other mechanisms within the protocol may work to reduce redundant NACK generation further. It is suggested that "T_maxBackoff" be selected as an integer multiple of the sender's current advertised GRTT estimate such that:

$$T_maxBackoff = K * GRTT; \text{ where } K \geq 1$$

For general Internet operation, a default value of "K=4" is RECOMMENDED for operation with multicast (to the group at large) NACK delivery; a value of "K=6" is the RECOMMENDED default for unicast NACK delivery. Alternate values may be used to achieve desired buffer utilization, reliable delivery latency, and group size scalability trade-offs.

Given that ("K*GRTT") is the maximum backoff time used by the receivers to initiate NACK transmission, other timeout periods related to the NACK repair process can be scaled accordingly. One of those timeouts is the amount of time a receiver should wait after generating a NACK message before allowing itself to initiate another NACK backoff/transmission cycle ("T_rcvrHoldoff"). This delay should be sufficient for the sender to respond to the received NACK with repair messages. An appropriate value depends upon the amount of time for the NACK to reach the sender and the sender to provide a repair response. This MUST include any amount of sender NACK aggregation period during which possible multiple NACKs are accumulated to determine an efficient repair response. These timeouts are further discussed in Section 3.2.4.

There are also secondary measures that can be applied to improve the performance of feedback suppression. For example, the sender's data content transmissions can follow an ordinal sequence of transmission. When repairs for data content occur, the receiver can note that the sender has "rewound" its data content transmission position by observing the data object, FEC block number, and FEC symbol identifiers. Receivers SHOULD limit transmission of NACKs to only when the sender's current transmission position exceeds the point to which the receiver has incomplete reception. This reduces premature requests for repair of data the sender may be planning to provide in response to other receiver requests. This mechanism can be very effective for protocol convergence in high loss conditions when transmissions of NACKs from other receivers (or indicators from the sender) are lost. Another mechanism (particularly applicable when FEC is used) is for the sender to embed an indication of impending repair transmissions in current packets sent. For example, the indication may be as simple as an advertisement of the number of FEC packets to be sent for the current applicable coding block.

Finally, some consideration might be given to using the NACKing history of receivers to bias their selection of NACK backoff timeout intervals. For example, if a receiver has historically been experiencing the greatest degree of loss, it may promote itself to statistically NACK sooner than other receivers. Note this requires correlation over successive intervals of time in the loss experienced by a receiver. Such correlation MAY not always be present in multicast networks. This adjustment of backoff timeout selection may

require the creation of an "early NACK" slot for these historical NACKers. This additional slot in the NACK backoff window will result in a longer repair cycle process that may not be desirable for some applications. The resolution of these trade-offs may be dependent upon the protocol's target application set or network.

After the random backoff timeout has expired, the receiver will make a decision on whether to generate a NACK repair request or not (i.e., it has been suppressed). The NACK will be suppressed when any of the following conditions has occurred:

1. The accumulated state of NACKs heard from other receivers (or forwarding of this state by the sender) is equal to or supersedes the repair needs of the local receiver. Note that the local receiver should consider its repair needs only up to the sender transmission position recorded at the NACK cycle initiation (when the backoff timer was activated).
2. The sender's data content transmission position "rewinds" to a point ordinal less than that of the lowest sequence position of the local receiver's repair needs. (This detection of sender "rewind" indicates the sender has already responded to other receiver repair needs of which the local receiver may not have been aware). This "rewind" event can occur any time between 1) when the NACK cycle was initiated with the backoff timeout activation and 2) the current moment when the backoff timeout has expired to suppress the NACK. Another NACK cycle must be initiated by the receiver when the sender's transmission sequence position exceeds the receiver's lowest ordinal repair point. Note it is possible that the local receiver may have had its repair needs satisfied as a result of the sender's response to the repair needs of other receivers and no further NACKing is required.

If these conditions have not occurred and the receiver still has pending repair needs, a NACK message is generated and transmitted. The NACK should consist of an accumulation of repair needs from the receiver's lowest ordinal repair point up to the current sender transmission sequence position. A single NACK message should be generated and the NACK message content should be truncated if it exceeds the payload size of single protocol message. When such NACK payload limits occur, the NACK content SHOULD contain requests for the ordinally lowest repair content needed from the sender.

Inputs:

1. NACK process initiation decision.
2. Recorded sender transmission sequence position.
3. Sender GRTT.
4. Sender group size estimate.
5. Application-defined bound on backoff timeout period.
6. NACKs from other receivers.
7. Pending repair indication from sender (may be forwarded NACKs).
8. Current sender transmission sequence position.

Outputs:

1. Yes/no decision to generate NACK message upon backoff timer expiration.

3.2.3. NACK Message Content

The content of NACK messages generated by reliable multicast receivers will include information detailing their current repair needs. The specific information depends on the use and type of FEC in the NACK repair process. The identification of repair needs is dependent upon the data content identification (see Section 3.5 below). At the highest level, the NACK content will identify the sender to which the NACK is addressed and the data transport object (or stream) within the sender's transmission that needs repair. For the indicated transport entity, the NACK content will then identify the specific FEC coding blocks and/or symbols it requires to reconstruct the complete transmitted data. This content may consist of FEC block erasure counts and/or explicit indication of missing blocks or symbols (segments) of data and FEC content. It should also be noted that NACK-based reliable multicast can be effectively instantiated without a requirement for reliable NACK delivery using the techniques discussed here.

3.2.3.1. NACK and FEC Repair Strategies

Where FEC-based repair is used, the NACK message content will minimally need to identify the coding block(s) for which repair is needed and a count of erasures (missing packets) for the coding

block. An exact count of erasures implies the FEC algorithm is capable of repairing any loss combination within the coding block. This count may need to be adjusted for some FEC algorithms.

Considering that multiple repair rounds may be required to successfully complete repair, an erasure count also implies that the quantity of unique FEC parity packets the server has available to transmit is essentially unlimited (i.e., the server will always be able to provide new, unique, previously unsent parity packets in response to any subsequent repair requests for the same coding block). Alternatively, the sender may "round-robin" transmit through its available set of FEC symbols for a given coding block, and eventually effect repair. For the most efficient repair strategy, the NACK content will need to also explicitly identify which symbols (information and/or parity) the receiver requires to successfully reconstruct the content of the coding block. This will be particularly true of small- to medium-size block FEC codes (e.g., Reed Solomon [FecSchemes]) that are capable of providing a limited number of parity symbols per FEC coding block.

When FEC is not used as part of the repair process, or the protocol instantiation is required to provide reliability even when the sender has transmitted all available parity for a given coding block (or the sender's ability to buffer transmission history is exceeded by the "(delay*bandwidth*loss)" characteristics of the network topology), the NACK content will need to contain explicit coding block and/or segment loss information so that the sender can provide appropriate repair packets and/or data retransmissions. Explicit loss information in NACK content may also potentially serve other purposes. For example, it may be useful for decorrelating loss characteristics among a group of receivers to help differentiate candidate congestion control bottlenecks among the receiver set.

When FEC is used and NACK content is designed to contain explicit repair requests, there is a strategy where the receivers can NACK for specific content that will help facilitate NACK suppression and repair efficiency. The assumptions for this strategy are that the sender may potentially exhaust its supply of new, unique parity packets available for a given coding block and be required to explicitly retransmit some data or parity symbols to complete reliable transfer. Another assumption is that an FEC algorithm where any parity packet can fill any erasure within the coding block (e.g., Reed Solomon) is used. The goal of this strategy is to make maximum use of the available parity and provide the minimal amount of data and repair transmissions during reliable transfer of data content to the group.

When systematic FEC codes are used, the sender transmits the data content of the coding block (and optionally some quantity of parity packets) in its initial transmission. Note that a systematic FEC coding block is considered to be logically made up of the contiguous set of source data vectors plus parity vectors for the given FEC algorithm used. For example, a systematic coding scheme that provides for 64 data symbols and 32 parity symbols per coding block would contain FEC symbol identifiers in the range of 0 to 95.

Receivers then can construct NACK messages requesting sufficient content to satisfy their repair needs. For example, if the receiver has three erasures in a given received coding block, it will request transmission of the three lowest ordinal parity vectors in the coding block. In our example coding scheme from the previous paragraph, the receiver would explicitly request parity symbols 64 to 66 to fill its three erasures for the coding block. Note that if the receiver's loss for the coding block exceeds the available parity quantity (i.e., greater than 32 missing symbols in our example), the receiver will be required to construct a NACK requesting all (32) of the available parity symbols plus some additional portions of its missing data symbols in order to reconstruct the block. If this is done consistently across the receiver group, the resulting NACKs will comprise a minimal set of sender transmissions to satisfy their repair needs.

In summary, the rule is to request the lower ordinal portion of the parity content for the FEC coding block to satisfy the erasure repair needs on the first NACK cycle. If the available number of parity symbols is insufficient, the receiver will also request the subset of ordinally highest missing data symbols to cover what the parity symbols will not fill. Note this strategy assumes FEC codes such as Reed-Solomon for which a single parity symbol can repair any erased symbol. This strategy would need minor modification to take into account the possibly limited repair capability of other FEC types. On subsequent NACK repair cycles where the receiver may receive some portion of its previously requested repair content, the receiver will use the same strategy, but only NACK for the set of parity and/or data symbols it has not yet received. Optionally, the receivers could also provide a count of erasures as a convenience to the sender.

Other types of FEC schemes may require alteration to the NACK and repair strategy described here. For example, some of the large block or expandable FEC codes described in [RFC3453] may be less deterministic with respect to defining optimal repair requests by receivers or repair transmission strategies by senders. For these types of codes, it may be sufficient for receivers to NACK with an estimate of the quantity of additional FEC symbols required to

complete reliable reception and for the sender to respond accordingly. This apparent disadvantage, as compared to codes such as Reed Solomon, may be offset by the reduced computational requirements and/or ability to support large coding blocks for increased repair efficiency that these codes can offer.

After receipt and accumulation of NACK messages during the aggregation period, the sender can begin transmission of fresh (previously untransmitted) parity symbols for the coding block based on the highest receiver erasure count if it has a sufficient quantity of parity symbols that were not previously transmitted. Otherwise, the sender MUST resort to transmitting the explicit set of repair vectors requested. With this approach, the sender needs to maintain very little state on requests it has received from the group without need for synchronization of repair requests from the group. Since all receivers use the same consistent algorithm to express their explicit repair needs, NACK suppression among receivers is simplified over the course of multiple repair cycles. The receivers can simply compare NACKs heard from other receivers against their own calculated repair needs to determine whether they should transmit or suppress their pending NACK messages.

3.2.3.2. NACK Content Format

The format of NACK content will depend on the protocol's data service model and the format of data content identification the protocol uses. This NACK format also depends upon the type of FEC encoding (if any) used. Figure 2 illustrates a logical, hierarchical transmission content identification scheme, denoting that the notion of objects (or streams) and/or FEC blocking is optional at the protocol instantiation's discretion. Note that the identification of objects is with respect to a given sender. It is recommended that transport data content identification is done within the context of a sender in a given session. Since the notion of session "streams" and "blocks" is optional, the framework degenerates to that of typical transport data segmentation and reassembly in its simplest form.




Figure 2: Reliable Multicast Data Content Identification Hierarchy

The format of NACK messages should enable the following:

1. Identification of transport data units required to repair the received content, whether this is an entire missing object/stream (or range), entire FEC coding block(s), or sets of symbols,
2. Simple processing for NACK aggregation and suppression,
3. Inclusion of NACKs for multiple objects, FEC coding blocks, and/or symbols in a single message, and
4. A reasonably compact format.

If the reliable multicast transport object/stream is identified with an <objectId> and the FEC symbol being transmitted is identified with an <fecPayloadId>, the concatenation of <objectId::fecPayloadId> comprises a basic transport protocol data unit (TPDU) identifier for symbols from a given source. NACK content can be composed of lists and/or ranges of these TPDU identifiers to build up NACK messages to describe the receiver's repair needs. If no hierarchical object delineation or FEC blocking is used, the TPDU is a simple linear representation of the data symbols transmitted by the sender. When the TPDU represents a hierarchy for purposes of object/stream delineation and/or FEC blocking, the NACK content unit may require flags to indicate which portion of the TPDU is applicable. For example, if an entire "object" (or range of objects) is missing in the received data, the receiver will not necessarily know the appropriate range of <sourceBlockNumbers> or <encodingSymbolIds> for which to request repair and thus requires some mechanism to request repair (or retransmission) of the entire unit represented by an <objectId>. The same is true if entire FEC coding blocks represented by one or a range of <sourceBlockNumbers> have been lost.

Inputs:

1. Sender identification.
2. Sender data identification.
3. Sender FEC object transmission information.
4. Recorded sender transmission sequence position.
5. Current sender transmission sequence position. History of repair needs for this sender.

Outputs:

1. NACK message with repair requests.

3.2.4. Sender NACK Processing and Repair Response

Upon reception of a repair request from a receiver in the group, the sender will initiate a repair response procedure. The sender may wish to delay transmission of repair content until it has had sufficient time to accumulate potentially multiple NACKs from the receiver set. This allows the sender to determine the most efficient repair strategy for a given transport stream/object or FEC coding block. Depending upon the approach used, some protocols may find it beneficial for the sender to provide an indicator of pending repair transmissions as part of its current transmitted message content. This can aid some NACK suppression mechanisms. The amount of time to perform this NACK aggregation should be sufficient to allow for the maximum receiver NACK backoff window (" $T_{\text{maxBackoff}}$ " from Section 3.2.2) and propagation of NACK messages from the receivers to the sender. Note the maximum transmission delay of a message from a receiver to the sender may be approximately " $(1 * \text{GRTT})$ " in the case of very asymmetric network topology with respect to transmission delay. Thus, if the maximum receiver NACK backoff time is " $T_{\text{maxBackoff}} = K * \text{GRTT}$ ", the sender NACK aggregation period should be equal to at least:

$$T_{\text{sndrAggregate}} = T_{\text{maxBackoff}} + 1 * \text{GRTT} = (K+1) * \text{GRTT}$$

Immediately after the sender NACK aggregation period, the sender will begin transmitting repair content determined from the aggregate NACK state and continue with any new transmission. Also, at this time, the sender should observe a "hold-off" period where it constrains itself from initiating a new NACK aggregation period to allow propagation of the new transmission sequence position due to the repair response to the receiver group. To allow for worst case asymmetry, this "hold-off" time should be:

$$T_{\text{sndrHoldoff}} = 1 * \text{GRTT}$$

Recall that the receivers will also employ a "hold-off" timeout after generating a NACK message to allow time for the sender's response. Given a sender " $\langle T_{\text{sndrAggregate}} \rangle$ " plus " $\langle T_{\text{sndrHoldoff}} \rangle$ " time of " $(K+1) * \text{GRTT}$ ", the receivers should use hold-off timeouts of:

$$T_{\text{rcvrHoldoff}} = T_{\text{sndrAggregate}} + T_{\text{sndrHoldoff}} = (K+2) * \text{GRTT}$$

This allows for a worst-case propagation time of the receiver's NACK to the sender, the sender's aggregation time, and propagation of the sender's response back to the receiver. Additionally, in the case of unicast feedback from the receiver set, it may be useful for the sender to forward (via multicast) a representation of its aggregated NACK content to the group to allow for NACK suppression when there is not multicast connectivity among the receiver set.

At the expiration of the "<T_sndrAggregate>" timeout, the sender will begin transmitting repair messages according to the accumulated content of NACKs received. There are some guidelines with regards to FEC-based repair and the ordering of the repair response from the sender that can improve reliable multicast efficiency:

When FEC is used, it is beneficial that the sender transmit previously untransmitted parity content as repair messages whenever possible. This maximizes the receiving nodes' ability to reconstruct the entire transmitted content from their individual subsets of received messages.

The transmitted object and/or stream data and repair content should be indexed with monotonically increasing sequence numbers (within a reasonably large ordinal space). If the sender observes the discipline of transmitting repair for the earliest content (e.g., ordinally lowest FEC blocks) first, the receivers can use a strategy of withholding repair requests for later content until the sender once again returns to that point in the object/stream transmission sequence. This can increase overall message efficiency among the group and help keep repair cycles relatively synchronized without dependence upon strict time synchronization among the sender and receivers. This also helps minimize the buffering requirements of receivers and senders and reduces redundant transmission of data to the group at large.

Inputs:

1. Receiver NACK messages.
2. Group timing information.

Outputs:

1. Repair messages (FEC and/or Data content retransmission).
2. Advertisement of current pending repair transmissions when unicast receiver feedback is detected.

3.3. Multicast Receiver Join Policies and Procedures

Consideration should be given to the policies and procedures by which new receivers join a group (perhaps where reliable transmission is already in progress) and begin requesting repair. If receiver joins are unconstrained, the dynamics of group membership may impede the application's ability to meet its goals for forward progression of data transmission. Policies that limit the opportunities for receivers to begin participating in the NACK process may be used to achieve the desired behavior. For example, it may be beneficial for receivers to attempt reliable reception from a newly-heard sender only upon non-repair transmissions of data in the first FEC block of an object or logical portion of a stream. The sender may also implement policies limiting the receivers from which it will accept NACK requests, but this may be prohibitive for scalability reasons in some situations. Alternatively, it may be desirable to have a looser transport synchronization policy and rely upon session management mechanisms to limit group dynamics that can cause poor performance in some types of bulk transfer applications (or for potential interactive reliable multicast applications).

Inputs:

1. Current object/stream data/repair content and sequencing identifiers from sender transmissions.

Outputs:

1. Receiver yes/no decision to begin receiving and NACKing for reliable reception of data.

3.4. Node (Member) Identification

In a NACK-based reliable multicast protocol (or other multicast protocols) where there is the potential for multiple sources of data, it is necessary to provide some mechanism to uniquely identify the sources (and possibly some or all receivers) within the group. Receivers that send NACK messages to the group will need to identify the sender to which the NACK is intended. Identity based on arriving packet source addresses is insufficient for several reasons. These reasons include routing changes for hosts with multiple interfaces that result in different packet source addresses for a given host over time, network address translation (NAT) or firewall devices, or other transport/network bridging approaches. As a result, some type of unique source identifier <sourceId> field SHOULD be present in packets transmitted by reliable multicast session members.

3.5. Data Content Identification

The data and repair content transmitted by a NACK-based reliable multicast sender requires some form of identification in the protocol header fields. This identification is required to facilitate the reliable NACK-oriented repair process. These identifiers will also be used in NACK messages generated. This building block document assumes two very general types of data that may comprise bulk transfer session content. One type is static, discrete objects of finite size and the other is continuous non-finite streams. A given application may wish to reliably multicast data content using either one or both of these paradigms. While it may be possible for some applications to further generalize this model and provide mechanisms to encapsulate static objects as content embedded within a stream, there are advantages in many applications to provide distinct support for static bulk objects and messages with the context of a reliable multicast session. These applications may include content caching servers, file transfer, or collaborative tools with bulk content. Applications with requirements for these static object types can then take advantage of transport layer mechanisms (i.e., segmentation/reassembly, caching, integrated forward error correction coding, etc.) rather than being required to provide their own mechanisms for these functions at the application layer.

As noted, some applications may alternatively desire to transmit bulk content in the form of one or more streams of non-finite size. Example streams include continuous quasi-real-time message broadcasts (e.g., stock ticker) or some content types that are part of collaborative tools or other applications. And, as indicated above, some applications may wish to encapsulate other bulk content (e.g., files) into one or more streams within a multicast session.

The components described within this building block document are envisioned to be applicable to both of these models with the potential for a mix of both types within a single multicast session. To support this requirement, the normal data content identification should include a field to uniquely identify the object or stream (e.g., <objectId>) within some reasonable temporal or ordinal interval. Note that it is not expected that this data content identification will be globally unique. It is expected that the object/stream identifier will be unique with respect to a given sender within the reliable multicast session and during the time that sender is supporting a specific transport instance of that object or stream.

Since "bulk" object/stream content usually requires segmentation, some form of segment identification must also be provided. This segment identifier will be relative to any object or stream

identifier that has been provided. Thus, in some cases, NACK-based reliable multicast protocol instantiations may be able to receive transmissions and request repair for multiple streams and one or more sets of static objects in parallel. For protocol instantiations employing FEC, the segment identification portion of the data content identifier may consist of a logical concatenation of a coding block identifier <sourceBlockNumber> and an identifier for the specific data or parity symbol <encodingSymbolId> of the code block. The FEC Basic Schemes building block [FEC Schemes] and descriptions of additional FEC schemes that may be documented later provide a standard message format for identifying FEC transmission content. NACK-based reliable multicast protocol instantiations using FEC SHOULD follow such guidelines.

Additionally, flags to determine the usage of the content identifier fields (e.g., stream vs. object) may be applicable. Flags may also serve other purposes in data content identification. It is expected that any flags defined will be dependent upon individual protocol instantiations.

In summary, the following data content identification fields may be required for NACK-based reliable multicast protocol data content messages:

1. Source node identifier (<sourceId>).
2. Object/Stream identifier (<objectId>), if applicable.
3. FEC Block identifier (<sourceBlockNumber>), if applicable.
4. FEC Symbol identifier (<encodingSymbolId>).
5. Flags to differentiate interpretation of identifier fields or identifier structure that implicitly indicates usage.
6. Additional FEC transmission content fields per FEC Building Block.

These fields have been identified because any generated NACK messages will use these identifiers in requesting repair or retransmission of data.

3.6. Forward Error Correction (FEC)

Multiple forward error correction (FEC) approaches using erasure coding techniques have been identified that can provide great performance enhancements to the repair process of NACK-oriented and other reliable multicast protocols [FecBroadcast], [RmFec],

[RFC3453]. NACK-based reliable multicast protocols can reap additional benefits since FEC-based repair does not generally require explicit knowledge of repair content within the bounds of its coding block size (in symbols). In NACK-based reliable multicast, parity repair packets generated will generally be transmitted only in response to NACK repair requests from receiving nodes. However, there are benefits in some network environments for transmitting some predetermined quantity of FEC repair packets multiplexed with the regular data symbol transmissions [FecHybrid]. This can reduce the amount of NACK traffic generated with relatively little overhead cost when group sizes are very large or the network connectivity has a large "delay*bandwidth" product with some nominal level of expected packet loss. While the application of FEC is not unique to NACK-based reliable multicast, these sorts of requirements may dictate the types of algorithms and protocol approaches that are applicable.

A specific issue related to the use of FEC with NACK-based reliable multicast is the mechanism used to identify the portion(s) of transmitted data content to which specific FEC packets are applicable. It is expected that FEC algorithms will be based on generating a set of parity repair packets for a corresponding block of transmitted data packets. Since data content packets are uniquely identified by the concatenation of <sourceId::objectId::sourceBlockNumber::encodingSymbolId> during transport, it is expected that FEC packets will be identified in a similar manner. The FEC Building Block document [RFC5052] provides detailed recommendations concerning application of FEC and standard formats for related reliable multicast protocol messages.

3.7. Round-Trip Timing Collection

The measurement of packet propagation round-trip time (RTT) among members of the group is required to support timer-based NACK suppression algorithms, timing of sender commands or certain repair functions, and congestion control operation. The nature of the round-trip information collected is dependent upon the type of interaction among the members of the group. In the case of "one-to-many" transmission, it may be that only the sender requires RTT knowledge of the GRTT and/or RTT knowledge of only a portion of the group. Here, the GRTT information might be collected in a reasonably scalable manner. For congestion control operation, it is possible that each receiver in the group may need knowledge of its individual RTT. In this case, an alternative RTT collection scheme may be utilized where receivers collect individual RTT measurements with respect to the sender(s) and advertise them to the group or sender(s). Where it is likely that exchange of reliable multicast data will occur among the group on a "many-to-many" basis, there are alternative measurement techniques that might be employed for

increased efficiency [DelayEstimation]. In some cases, there might be absolute time synchronization available among the participating hosts that may simplify RTT measurement. There are trade-offs in multicast congestion control design that require further consideration before a universal recommendation on RTT (or GRTT) measurement can be specified. Regardless of how the RTT information is collected (and more specifically GRTT) with respect to congestion control or other requirements, the sender will need to advertise its current GRTT estimate to the group for various NACK timeouts used by receivers.

3.7.1. One-to-Many Sender GRTT Measurement

The goal of this form of RTT measurement is for the sender to estimate the GRTT among the receivers who are actively participating in NACK-based reliable multicast operation. The set of receivers participating in this process may be the entire group or some subset of the group determined from another mechanism within the protocol instantiation. An approach to collect this GRTT information follows.

The sender periodically polls the group with a message (independent or "piggy-backed" with other transmissions) containing a "<sendTime>" timestamp relative to an internal clock at the sender. Upon reception of this message, the receivers will record this "<sendTime>" timestamp and the time (referenced to their own clocks) at which it was received "<recvTime>". When the receiver provides feedback to the sender (either explicitly or as part of other feedback messages depending upon protocol instantiation specification), it will construct a "response" using the formula:

$$\text{grttResponse} = \text{sendTime} + (\text{currentTime} - \text{recvTime})$$

where the "<sendTime>" is the timestamp from the last probe message received from the source and the (" $\text{currentTime} - \text{recvTime}$ ") is the amount of time differential since that request was received until the receiver generated the response.

The sender processes each receiver response by calculating a current RTT measurement for the receiver from whom the response was received using the following formula:

$$\text{RTT_rcvr} = \text{currentTime} - \text{grttResponse}$$

During each periodic "GRTT" probing interval, the source keeps the peak round-trip timing measurement ("RTT_peak") from the set of responses it has received. A conservative estimate of "GRTT" is kept

to maximize the efficiency of redundant NACK suppression and repair aggregation. The update to the source's ongoing estimate of "GRTT" is done observing the following rules:

1. If a receiver's response round-trip time ("RTT_rcvr") is greater than the current "GRTT" estimate, the "GRTT" is immediately updated to this new peak value:

$$\text{GRTT} = \text{RTT_rcvr}$$

2. At the end of the response collection period (i.e., the GRTT probe interval), if the recorded "peak" response ("RTT_peak") is less than the current GRTT estimate, the GRTT is updated to:

$$\text{GRTT} = \text{MAX}(0.9 * \text{GRTT}, \text{RTT_peak})$$

3. If no feedback is received, the sender "GRTT" estimate remains unchanged.
4. At the end of the response collection period, the peak tracking value ("RTT_peak") is reset to ZERO for subsequent peak detection.

The GRTT collection period (i.e., period of probe transmission) could be fixed at a value on the order of that expected for group membership and/or network topology dynamics. For robustness, more rapid probing could be used at protocol startup before settling to a less frequent, steady-state interval. Optionally, an algorithm may be developed to adjust the GRTT collection period dynamically in response to the current estimate of GRTT (or variations in it) and to an estimation of packet loss. The overhead of probing messages could then be reduced when the GRTT estimate is stable and unchanging, but be adjusted to track more dynamically during periods of variation with correspondingly shorter GRTT collection periods. GRTT collection MAY also be coupled with collection of other information for congestion control purposes.

In summary, although NACK repair cycle timeouts are based on GRTT, it should be noted that convergent operation of the protocol does not depend upon highly accurate GRTT estimation. The current mechanism has proved sufficient in simulations and in the environments where NACK-based reliable multicast protocols have been deployed to date. The estimate provided by the given algorithm tracks the peak envelope of actual GRTT (including operating system effect as well as network delays) even in relatively high loss connectivity. The steady-state probing/update interval may potentially be varied to accommodate different levels of expected network dynamics in different environments.

3.7.2. One-to-Many Receiver RTT Measurement

In this approach, receivers send messages with timestamps to the sender. To control the volume of these receiver-generated messages, a suppression mechanism similar to that described for NACK suppression may be used. The "age" of receivers' RTT measurement should be kept by receivers and used as a metric in competing for feedback opportunities in the suppression scheme. For example, receiver who have not made any RTT measurement or whose RTT measurement has aged most should have precedence over other receivers. In turn, the sender may have limited capacity to provide an "echo" of the receiver timestamps back to the group, and it could use this RTT "age" metric to determine which receivers get precedence. The sender can determine the "GRTT" as described in 3.7.1 if it provides sender timestamps to the group. Alternatively, receivers who note their RTT is greater than the sender GRTT can compete in the feedback opportunity/suppression scheme to provide the sender and group with this information.

3.7.3. Many-to-Many RTT Measurement

For reliable multicast sessions that involve multiple senders, it may be useful to have RTT measurements occur on a true "many-to-many" basis rather than have each sender independently tracking RTT. Some protocol efficiency can be gained when receivers can infer an approximation of their RTT with respect to a sender based on RTT information they have on another sender and that other sender's RTT with respect to the new sender of interest. For example, for receiver "a" and senders "b" and "c", it is likely that:

$$\text{RTT}(a \leftrightarrow b) \leq \text{RTT}(a \leftrightarrow c) + \text{RTT}(b \leftrightarrow c)$$

Further refinement of this estimate can be conducted if RTT information is available to a node concerning its own RTT with respect to a small subset of other group members and if information concerning RTT among those other group members is learned by the node during protocol operation.

3.7.4. Sender GRTT Advertisement

To facilitate deterministic protocol operation, the sender should robustly advertise its current estimation of "GRTT" to the receiver set. Common, robust knowledge of the sender's current operating GRTT estimate among the group will allow the protocol to progress in its most efficient manner. The sender's GRTT estimate can be robustly advertised to the group by simply embedding the estimate into all pertinent messages transmitted by the sender. The overhead of this can be made quite small by quantizing (compressing) the GRTT estimate

to a single byte of information. The following C-language functions allow this to be done over a wide range ("RTT_MIN" through "RTT_MAX") of GRTT values while maintaining a greater range of precision for small values and less precision for large values. Values of 1.0e-06 seconds and 1000 seconds are RECOMMENDED for "RTT_MIN" and "RTT_MAX" respectively. NACK-based reliable multicast applications may wish to place an additional, smaller upper limit on the GRTT advertised by senders to meet application data delivery latency constraints at the expense of greater feedback volume in some network environments.

```

unsigned char QuantizeGrtt(double grtt)
{
    if (grtt > RTT_MAX)
        grtt = RTT_MAX;
    else if (grtt < RTT_MIN)
        grtt = RTT_MIN;
    if (grtt < (33*RTT_MIN))
        return ((unsigned char)(grtt / RTT_MIN) - 1);
    else
        return ((unsigned char)(ceil(255.0 -
            (13.0 * log(RTT_MAX/grtt)))));
}

double UnquantizeRtt(unsigned char qrtt)
{
    return ((qrtt <= 31) ?
        (((double)(qrtt+1))*(double)RTT_MIN) :
        (RTT_MAX/exp(((double)(255-qrtt))/(double)13.0)));
}

```

Note that this function is useful for quantizing GRTT times in the range of 1 microsecond to 1000 seconds. Of course, NACK-based reliable multicast protocol implementations may wish to further constrain advertised GRTT estimates (e.g., limit the maximum value) for practical reasons.

3.8. Group Size Determination/Estimation

When NACK-based reliable multicast protocol operation includes mechanisms that excite feedback from the group at large (e.g., congestion control), it may be possible to roughly estimate the group size based on the number of feedback messages received with respect to the distribution of the probabilistic suppression mechanism used. Note the timer-based suppression mechanism described in this document does not require a very accurate estimate of group size to perform adequately. Thus, a rough estimate, particularly if conservatively managed, may suffice. Group size may also be determined administratively. In absence of any group size determination

mechanism, a default group size value of 10,000 is RECOMMENDED for reasonable management of feedback given the scalability of expected NACK-based reliable multicast usage. This conservative estimate (over-estimate) of group size in the algorithms described above will result in some added latency to the NACK repair process if the actual group size is smaller but with a guarantee of feedback implosion protection. The study of the timer-based feedback suppression mechanism described in [McastFeedback] and [NormFeedback] showed that the group size estimate need only be with an order-of-magnitude to provide effective suppression performance.

3.9. Congestion Control Operation

Congestion control that fairly shares available network capacity with other reliable multicast and TCP instantiations is REQUIRED for general Internet operation. The TCP-Friendly Multicast Congestion Control (TFMCC) [TfmccPaper] or Pragmatic General Multicast Congestion Control (PGMCC) [PgmccPaper] techniques can be applied to NACK-based reliable multicast operation to meet this requirement. The former technique has been further documented in [RFC4654] and has been successfully applied in the NACK-Oriented Reliable Multicast Protocol (NORM) [RFC3940].

3.10. Intermediate System Assistance

NACK-based multicast protocols may benefit from general purpose intermediate system assistance. In particular, additional NACK suppression where intermediate systems can aggregate NACK content (or filter duplicate NACK content) from receivers as it is relayed toward the sender could enhance NORM group size scalability. For NACK-based reliable multicast protocols using FEC, it is possible that intermediate systems may be able to filter FEC repair messages to provide an intelligent "subcast" of repair content to different legs of the multicast topology depending on the repair needs learned from previous receiver NACKs. Similarly, intermediate systems could monitor receiver NACKs and provide repair transmissions on-demand in response if sufficient state on the content being transmitted was being maintained. This can reduce the latency and volume of repair transmissions when the intermediate system is associated with a network link that is particularly problematic with respect to packet loss. These types of assist functions would require intermediate system interpretation of transport data unit content identifiers and flags. NACK-based protocol designs should consider the potential for intermediate system assistance in the specification of protocol messages and operations. It is likely that intermediate systems assistance will be more pragmatic if message parsing requirements are modest and if the amount of state an intermediate system is required to maintain is relatively small.

4. NACK-Based Reliable Multicast Applicability

The Multicast NACK building block applies to protocols wishing to employ negative acknowledgement to achieve reliable data transfer. Properly designed NACK-based reliable multicast protocols offer scalability advantages for applications and/or network topologies where, for various reasons, it is prohibitive to construct a higher order delivery infrastructure above the basic Layer 3 IP multicast service (e.g., unicast or hybrid unicast/multicast data distribution trees). Additionally, the multicast scalability property of NACK-based protocols [RmComparison], [RmClasses] is applicable where broad "fan-out" is expected for a single network hop (e.g., cable-TV data delivery, satellite, or other broadcast communication services). Furthermore, the simplicity of a protocol based on "flat" group-wide multicast distribution may offer advantages for a broad range of distributed services or dynamic networks and applications. NACK-based reliable multicast protocols can make use of reciprocal (among senders and receivers) multicast communication under the any-source multicast (ASM) model defined in RFC 1112 [RFC1112], and are capable of scalable operation in asymmetric topologies, such as source-specific multicast (SSM) [RFC4607], where there may only be unicast routing service from the receivers to the sender(s).

NACK-based reliable multicast protocol operation is compatible with transport layer forward error correction coding techniques as described in [RFC3453] and congestion control mechanisms such as those described in [TfmccPaper] and [PgmccPaper]. A principal limitation of NACK-based reliable multicast operation involves group size scalability when network capacity for receiver feedback is very limited. It is possible that, with proper protocol design, the intermediate system assistance techniques mentioned in Section 2.4 and described further in Section 3.10 can allow NACK-based approaches to scale to larger group sizes. NACK-based reliable multicast operation is also governed by implementation buffering constraints. Buffering greater than that required for typical point-to-point reliable transport (e.g., TCP) is recommended to allow for disparity in the receiver group connectivity and to allow for the feedback delays required to attain group size scalability.

Prior experimental work included various protocol instantiations that implemented some of the concepts described in this building block document. This includes the Pragmatic General Multicast (PGM) protocol described in [RFC3208] as well as others that were documented or deployed outside of IETF activities. While the PGM protocol specification and some other approaches encompassed many of the goals of bulk data delivery as described here, this NACK-based building block provides a more generalized framework so that different application needs can be met by different protocol

instantiation variants. The NACK-based building block approach described here includes compatibility with the other protocol mechanisms including FEC and congestion control that are described in other IETF reliable multicast building block documents. The NACK repair process described in this document can provide performance advantages compared to PGM when both are deployed on a pure end-to-end basis without intermediate system assistance. The round-trip timing estimation described here and its use in the NACK repair process allow protocol operation to more automatically adapt to different network environments or operate within environments where connectivity is dynamic. Use of the FEC payload identification techniques described in the FEC building block [RFC5052] and specific FEC instantiations allow protocol instantiations more flexibility as FEC techniques evolve than the specific sequence number data identification scheme described in the PGM specification. Similar flexibility is expected if protocol instantiations are designed to modularly invoke (at design time, if not run-time) the appropriate congestion control building block for different application or deployment purposes.

5. Security Considerations

NACK-based reliable multicast protocols are expected to be subject to the same security vulnerabilities as other IP and IP multicast protocols. However, unlike point-to-point (unicast) transport protocols, it is possible that one badly behaving participant can impact the transport service experience of others in the group. For example, a malicious receiver node could intentionally transmit NACK messages to cause the sender(s) to unnecessarily transmit repairs instead of making forward progress with reliable transfer. Also, group-wise messaging to support congestion control or other aspects of protocol operation may be subject to similar vulnerabilities. Thus, it is highly RECOMMENDED that security techniques such as authentication and data integrity checks be applied for NACK-based reliable multicast deployments. Protocol instantiations using this building block MUST identify approaches to security that can be used to address these and other security considerations.

NACK-based reliable multicast is compatible with IP security (IPsec) authentication mechanisms [RFC4301] that are RECOMMENDED for protection against session intrusion and denial of service attacks. A particular threat for NACK-based protocols is that of NACK replay attacks, which could prevent a multicast sender from making forward progress in transmission. Any standard IPsec mechanisms that can provide protection against such replay attacks are RECOMMENDED for use. The IETF Multicast Security (MSEC) Working Group has developed a set of recommendations in its "Multicast Extensions to the Security Architecture for the Internet Protocol" [IpsecExtensions] that can be

applied to appropriately extend IPsec mechanisms to multicast operation. An appendix of this document specifically addresses the NACK-Oriented Reliable Multicast protocol service model. As complete support for IPsec multicast operation may potentially follow reliable multicast deployment, NACK-based reliable multicast protocol instantiations SHOULD consider providing support for their own NACK replay attack protection when network layer mechanisms are not available. This MAY be necessary when IPsec implementations are used that do not provide multicast replay attack protection when multiple sources are present.

For NACK-based multicast deployments with large receiver groups using IPsec, approaches might be developed that use shared, common keys for receiver-originated protocol messages to maintain a practical number of IPsec Security Associations (SAs). However, such group-based authentication may not be sufficient unless the receiver population can be completely trusted. Additionally, this can make identification of badly behaving (although authenticated) receiver nodes problematic as such nodes could potentially masquerade as other receivers in the group. In deployments such as this, one SHOULD consider use of source-specific multicast (SSM) instead of any-source multicast (ASM) models of multicast operation. SSM operation can simplify security challenges in a couple of ways:

1. A NACK-based protocol supporting SSM operation can eliminate direct receiver-to-receiver signaling. This dramatically reduces the number of security associations that need to be established.
2. The SSM sender(s) can provide a centralized management point for secure group operation for its respective data flow as the sender alone is required to conduct individual host authentication for each receiver when group-based authentication does not suffice or is not pragmatic to deploy.

When individual host authentication is required, then it is possible receivers could use a digital signature on the IPsec Encapsulating Security Protocol (ESP) payload as described in [RFC4359]. Either an identity-based signature system or a group-specific public key infrastructure could avoid per-receiver state at the sender(s). Additionally, implementations MUST also support policies to limit the impact of extremely or exceptionally poor-performing (due to bad behavior or otherwise) receivers upon overall group operation if this is acceptable for the relevant application.

As described in Section 3.4, deployment of NACK-based reliable multicast in some network environments may require identification of group members beyond that of IP addressing. If protocol-specific security mechanisms are developed, then it is RECOMMENDED that

protocol group member identifiers are used as selectors (as defined in [RFC4301]) for the applicable security associations. When IPsec is used, it is RECOMMENDED that the protocol implementation verify that the source IP addresses of received packets are valid for the given protocol source identifier in addition to usual IPsec authentication. This would prevent a badly behaving (although authorized) member from spoofing messages from other legitimate members, provided that individual host authentication is supported.

The MSEC Working Group has also developed automated group keying solutions that are applicable to NACK-based reliable multicast security. For example, to support IPsec or other security mechanisms, the Group Secure Association Key Management Protocol [RFC4535] MAY be used for automated group key management. The technique it identifies for "Group Establishment for Receive-Only Members" may be application NACK-based reliable multicast SSM operation.

6. Changes from RFC 3941

This section lists the changes between the Experimental version of this specification, [RFC3941], and this version:

1. Change of title to avoid confusion with NORM Protocol specification,
2. Updated references to related, updated RMT Building Block documents, and
3. More detailed security considerations.

7. Acknowledgements

(and these are not Negative)

The authors would like to thank George Gross, Rick Jones, and Joerg Widmer for their valuable comments on this document. The authors would also like to thank the RMT working group chairs, Roger Kermod and Lorenzo Vicisano, for their support in development of this specification, and Sally Floyd for her early inputs into this document.

8. References

8.1. Normative References

- [RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, August 1989.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4607] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", RFC 4607, August 2006.

8.2. Informative References

- [ArchConsiderations] Clark, D. and D. Tennenhouse, "Architectural Considerations for a New Generation of Protocols", Proc. ACM SIGCOMM, pp. 201-208, September 1990.
- [DelayEstimation] Ozdemir, V., Muthukrishnan, S., and I. Rhee, "Scalable, Low-Overhead Network Delay Estimation", NCSU/AT&T White Paper, February 1999.
- [FEC Schemes] Watson, M., "Basic Forward Error Correction (FEC) Schemes", Work in Progress, July 2008.
- [FecBroadcast] Metzner, J., "An Improved Broadcast Retransmission Protocol", IEEE Transactions on Communications Vol. Com-32, No. 6, June 1984.
- [FecHybrid] Gossink, D. and J. Macker, "Reliable Multicast and Integrated Parity Retransmission with Channel Estimation", IEEE Globecom 1998, 1998.
- [FecSchemes] Lacan, J., Roca, V., Peltotalo, J., and S. Peltotalo, "Reed-Solomon Forward Error Correction (FEC) Schemes", Work in Progress, November 2007.
- [IpssecExtensions] Weis, B., Gross, G., and D. Ignjatic, "Multicast Extensions to the Security Architecture for the Internet Protocol", Work in Progress, June 2008.

- [McastFeedback] Nonnenmacher, J. and E. Biersack, "Optimal Multicast Feedback", IEEE Infocom p. 964, March/April 1998.
- [NormFeedback] Adamson, B. and J. Macker, "Quantitative Prediction of NACK-Oriented Reliable Multicast (NORM) Feedback", IEEE MILCOM 2002, October 2002.
- [PgmccPaper] Rizzo, L., "pgmcc: A TCP-Friendly Single-Rate Multicast Congestion Control Scheme", ACM SIGCOMM 2000, August 2000.
- [RFC2357] Mankin, A., Romanov, A., Bradner, S., and V. Paxson, "IETF Criteria for Evaluating Reliable Multicast Transport and Application Protocols", RFC 2357, June 1998.
- [RFC3208] Speakman, T., Crowcroft, J., Gemmell, J., Farinacci, D., Lin, S., Leshchiner, D., Luby, M., Montgomery, T., Rizzo, L., Tweedly, A., Bhaskar, N., Edmonstone, R., Sumanasekera, R., and L. Vicisano, "PGM Reliable Transport Protocol Specification", RFC 3208, December 2001.
- [RFC3269] Kermodé, R. and L. Vicisano, "Author Guidelines for Reliable Multicast Transport (RMT) Building Blocks and Protocol Instantiation documents", RFC 3269, April 2002.
- [RFC3453] Luby, M., Vicisano, L., Gemmell, J., Rizzo, L., Handley, M., and J. Crowcroft, "The Use of Forward Error Correction (FEC) in Reliable Multicast", RFC 3453, December 2002.
- [RFC3940] Adamson, B., Bormann, C., Handley, M., and J. Macker, "Negative-acknowledgment (NACK)-Oriented Reliable Multicast (NORM) Protocol", RFC 3940, November 2004.
- [RFC3941] Adamson, B., Bormann, C., Handley, M., and J. Macker, "Negative-Acknowledgment (NACK)-Oriented Reliable Multicast (NORM) Building Blocks", RFC 3941, November 2004.

- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC4359] Weis, B., "The Use of RSA/SHA-1 Signatures within Encapsulating Security Payload (ESP) and Authentication Header (AH)", RFC 4359, January 2006.
- [RFC4535] Harney, H., Meth, U., Colegrove, A., and G. Gross, "GSAKMP: Group Secure Association Key Management Protocol", RFC 4535, June 2006.
- [RFC4654] Widmer, J. and M. Handley, "TCP-Friendly Multicast Congestion Control (TFMCC): Protocol Specification", RFC 4654, August 2006.
- [RFC5052] Watson, M., Luby, M., and L. Vicisano, "Forward Error Correction (FEC) Building Block", RFC 5052, August 2007.
- [RmClasses] Levine, B. and J. Garcia-Luna-Aceves, "A Comparison of Known Classes of Reliable Multicast Protocols", Proc. International Conference on Network Protocols (ICNP-96) Columbus, OH, October 1996.
- [RmComparison] Pingali, S., Towsley, D., and J. Kurose, "A Comparison of Sender-Initiated and Receiver-Initiated Reliable Multicast Protocols", Proc. INFOCOMM San Francisco, CA, October 1993.
- [RmFec] Macker, J., "Reliable Multicast Transport and Integrated Erasure-based Forward Error Correction", IEEE MILCOM 1997, October 1997.
- [SrmFramework] Floyd, S., Jacobson, V., McCanne, S., Liu, C., and L. Zhang, "A Reliable Multicast Framework for Light-weight Sessions and Application Level Framing", Proc. ACM SIGCOMM, August 1995.
- [TfmccPaper] Widmer, J. and M. Handley, "Extending Equation-Based Congestion Control to Multicast Applications", ACM SIGCOMM 2001, August 2001.

Authors' Addresses

Brian Adamson
Naval Research Laboratory
Washington, DC 20375

EEmail: adamson@itd.nrl.navy.mil

Carsten Bormann
Universitaet Bremen TZI
Postfach 330440
D-28334 Bremen, Germany

EEmail: cabo@tzi.org

Mark Handley
University College London
Gower Street
London, WC1E 6BT
UK

EEmail: M.Handley@cs.ucl.ac.uk

Joe Macker
Naval Research Laboratory
Washington, DC 20375

EEmail: macker@itd.nrl.navy.mil