Washington University in St. Louis

Washington University Open Scholarship

All Computer Science and Engineering
Research

Computer Science and Engineering

Report Number: WUCS-01-05

2001-01-01

Aggregated Hierarchical Multicast for Active Networks

Tilman Wolf and Sumi Y. Choi

Active Networking is the basis for a range of new and innovative applications that make use of computational resources inside network routers. One such application is Aggregated Hierarchical Multicast, which aims at implementing efficient many-to-many communication. In certain scenarios it is possible to transmit less accurate, aggregated data and thus achieve better scalability. Using Active Networks, this aggregation computation can be done transparently by network routers without end system support. We present how aggregated data streams can be structured in a hierarchical fashion to allow easy access of data at the desired aggregation level. We introduce two application examples... [Read complete abstract on page 2.](#)

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

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

Recommended Citation

Wolf, Tilman and Choi, Sumi Y., "Aggregated Hierarchical Multicast for Active Networks" Report Number: WUCS-01-05 (2001). *All Computer Science and Engineering Research*.
https://openscholarship.wustl.edu/cse_research/248

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

Aggregated Hierarchical Multicast for Active Networks

Tilman Wolf and Sumi Y. Choi

Complete Abstract:

Active Networking is the basis for a range of new and innovative applications that make use of computational resources inside network routers. One such application is Aggregated Hierarchical Multicast, which aims at implementing efficient many-to-many communication. In certain scenarios it is possible to transmit less accurate, aggregated data and thus achieve better scalability. Using Active Networks, this aggregation computation can be done transparently by network routers without end system support. We present how aggregated data streams can be structured in a hierarchical fashion to allow easy access of data at the desired aggregation level. We introduce two application examples to illustrate the system design, analyze the performance of the aggregation mechanism, and evaluate it using a prototype implementation.

**Aggregated Hierarchical Multicast for Active
Networks**

Tilman Wolf and Sumi Y. Choi

WUCS-01-05

February 2001

**Department of Computer Science
Washington University
Campus Box 1045
One Brookings Drive
St. Louis MO 63130**

Aggregated Hierarchical Multicast for Active Networks

Tilman Wolf and Sumi Y. Choi
{wolf, syc}@arl.wustl.edu
Department of Computer Science
Washington University in St. Louis, MO, USA

Abstract

Active Networking is the basis for a range of new and innovative applications that make use of computational resources inside network routers. One such application is Aggregated Hierarchical Multicast, which aims at implementing efficient many-to-many communication. In certain scenarios it is possible to transmit less accurate, aggregated data and thus achieve better scalability. Using Active Networks, the aggregation computation can be done transparently by network routers without end system support. We present how aggregated data streams can be structured in a hierarchical fashion to allow easy access of data at the desired aggregation level. We introduce two application examples to illustrate the system design, analyze the performance of the aggregation mechanism, and evaluate it using a prototype implementation.

Keywords: active networks, multicast, data aggregation, active network application

1 Introduction

Many-to-many or multi-source communication is becoming an increasingly important way of communicating. As more and more real-time data is sent by sources distributed over the network, the receiving end-systems become overwhelmed by the amount of data traffic. In a traditional many-to-many communication, each sender is connected to each receiver (be it over unicast or multicast). As a result, a receiver has to deal with as many connections as there are senders. This does not scale well, neither in terms of bandwidth requirements nor in terms of computational demands. Handheld clients with little computational resources or devices connected over low bandwidth wireless links are severely limited in the number of connections they can handle.

A key observation, though, is that for certain applications in many-to-many communications the senders are not of equal importance to the receiver. Thus, it is sufficient to aggregate the information sent by most of the sources, while keeping the full data stream from a few selected sources. To illustrate this concept, we look at two applications: a battlefield information system and an audio conferencing application. We show how aggregation is done for each application and how a hierarchical overlay allows the user to dynamically choose the right level of detail for his needs.

The basis for the work lies in the assumption that the interconnection network between the senders and receivers is capable of performing the aggregation of data on-the-fly. Thus, we require what is referred to as an “active network” [8], [2]. An active node in such a network is capable of performing processing of packets as they are being forwarded. Typical implementations range from workstations that act as active routers to high-performance switches that are augmented with per-port network processors. The definition of a unifying node operating system (NodeOS [7]) aims

at making these systems interoperable. While Aggregated Hierarchical Multicast is an application for active networks, we do not go into the details of active networking in this paper.

Section 2 introduces the example applications considered here and shows different information aggregation methods. Section 3 discusses the hierarchical structure of multicast sessions that provides various levels of aggregation detail to the user. Section 4 describes a general-purpose aggregation algorithm that is used in our prototype implementation. Section 5 gives quantitative evaluation of Active Hierarchical Multicast. A summary of the contributions in Section 6 concludes this paper.

2 Information Aggregation and its Applications

As discussed above, the limited scalability of many-to-many communications lies in the demands on the network to deliver numerous data streams to the end system and the demands on the end system to process and display the information. To be more concrete, we look at the following two applications.

2.1 Application I: Battlefield Information System

The battlefield information system is aimed at providing status information of numerous soldiers and equipment to a large number of commanders and observers. We assume that all soldiers are connected to a common interconnection network (i.e., via wireless links) and have the equipment to periodically transmit their status information (i.e., geographic location, vital statistics, and other easily observable data). The observers can receive this data and display the status of all soldiers accordingly. This application is an example of large multicast with many sources.

The challenge of this application lies in the large amount of data that is received by an observer. Considering only a few thousand soldiers, who transmit every few seconds, gives an average of several hundred data packets per second that have to be processed and displayed.

Since it is unlikely that any observer is interested in the exact status of every single soldier at all times, it is possible to present aggregated status information of a group of soldiers to an observer thereby reducing the number of data streams that are sent to a receiver. Geographic information for a group of people, for example, can be represented by its centroid, the weighted geographic average, or the convex hull. This reduces the amount of detail in the representation, but it also reduces the amount of data that needs to be transmitted. Similarly, vital statistics can be aggregated. For example the predicate "healthy" of each soldier can be aggregated using a boolean "and" function.

Finally, there is the question of where in the network the aggregation computations should be done. For this purpose, we use a given hierarchical structure (i.e., the typical chain-of-command). Each node sends its information to its hierarchical parent, where the information is aggregated with that of the node's siblings. This is repeated over all levels of the hierarchy. An illustration of this is shown in Figure 1. More details on how to tap the hierarchy at the right level to get the right detail of information is discussed in Section 3.

2.2 Application II: Audio Conferencing

An audio conferencing application, as illustrated in Figure 2 is another typical example of a many-to-many multicast. Each participant in the conference needs to be able to hear all other participants and therefore needs to receive their audio data stream. In a non-active network, the end system

Figure 1: Battlefield Information Application for Aggregated Hierarchical Multicast. Each node sends its status information to its hierarchical parent, where it is aggregated.

Figure 2: Audio Conferencing Application for Aggregated Hierarchical Multicast. Each end system sends audio that is mixed at nodes where multiple audio streams merge. The end system receives a single, mixed audio stream.

has to receive all these data streams, mix them together, and play the result to the user. In an active network, though, the audio streams can be aggregated as they traverse the network and the end-system is provided with a single mixed audio stream.

2.3 Data Aggregation

Aggregation of many-to-many multicast data streams in the network brings benefits to the network as well as to the end system.

From the view point of the network, the total amount of traffic is reduced, since only one aggregated data stream has to be delivered to a receiver. Even though the aggregated data stream can potentially be larger than any single data stream, it will always be less than the sum of all single data stream. Also, there can be a reduction in transmission frequency for aggregated data streams. For example, in the battlefield information system, it might be necessary to have high frequency updates for the status of an individual soldier who moves around quickly. The centroid of a group, though, moves at a slower rate and therefore needs to be updated less frequently. A qualitative analysis of this result can be found in Section 5.

From the view point of the end system, the received data can be seen as a single unicast connection from the group and displayed directly without the need for complex aggregation processing. In the audio application, for example, the mixed audio stream differs not from a unicast audio stream, except that the audio samples are an aggregate from several sources. This reduces the processing requirements and simplifies the end system application development.

There are several issues that have to be addressed in order to achieve efficient data aggregation. For one, it has to be possible to aggregate the data that is transmitted. In many cases, aggregation can be achieved by downscaling information (e.g., scale change for geographic information) or generating an overlay of different data streams (e.g., mixing of audio). Status information can often be aggregated by simple arithmetic and boolean functions. However, certain data, e.g., text messages, cannot be scaled or aggregated effectively without losing crucial components of the data. In such a case, the information can be concatenated and transmitted unchanged, losing the benefits of reduction in bandwidth and processing requirements.

3 Hierarchy of Source-based Multicast Sessions

In a realistic environment, different observers need different levels of aggregation detail and an individual observer might want to change the level of detail dynamically. To accommodate these requirements, we propose a hierarchy of multicast sessions that provide different levels of aggregation to the end system. Each session is augmented by control information that allows the user to change to a higher or lower level of detail if necessary.

3.1 Hierarchy Layout

Each layer in the hierarchy represents a different aggregation level. The lowest layer, layer 0, are the data sources (i.e., soldiers, telephones) that send their unaggregated data. Each node in layer 1 aggregates multiple layer 0 sources to a new stream. This stream is sent upwards to layer 2, where it is aggregated with other layer 1 streams. This continues up to the root node. In general, a layer i node aggregates streams from layer $i - 1$ and sends it to layer $i + 1$. Figure 3 shows this concept.

An observer who wants to get information from a node on a certain layer, can directly connect to that node and receive its aggregated data stream. If more detail is required, the observer can connect to a child of the current node. If less detailed information is required, the observer can move up to the parent of the current node.

Since there are many observers present at any time, it is of course possible that multiple observers are connected to a certain node. To make this scalable, any node offers its data as a multicast session, for which it acts as the source. Any observer can subscribe to this multicast session using well established multicast schemes (e.g., by using multicast routers that exchange IGMP [3] messages and route using MOSPF [6]).

3.2 Session Control

The scheme described above requires that nodes are configured to form the tree and aggregate the right set of lower layer streams. Thus, all nodes other than leaves need to know their children at start time. Once the children are known, the node can subscribe to their multicast sessions and aggregate the information.

Figure 3: Hierarchy of Data Aggregation.

The control information that is distributed with each session needs to contain two components. First, the list of children and their respective session identifiers is included so that an observer can receive data from a lower layer. Second, the session identifier of the parent of the current node is included to allow an observer to step to a higher aggregation layer.

3.3 Naming Issues

For scalability reasons, the control data of a session cannot contain the names/identifiers of all children, grandchildren, etc., since this would lead to an exponential growth in control information. As a result, at a high level in the hierarchy, say layer i , the contributing sources of a data stream are not known. If an observer wants to step towards a source x , he does not know which child of that node contains the aggregated data stream from x .

To solve this problem, we add a naming scheme to the hierarchy. Each node is identified by its name and the concatenation of names of nodes in the path from the root to the node. This results in a hierarchical naming scheme similar to that used in the Domain Name Service [5] in the Internet or that of class names in object oriented programming languages, like Java [4]. Thus, the full name of leaf x contains the list of higher layer nodes that lead to x , which allows the observer to easily navigate through the tree. For convenience purposes and to abbreviate long names, certain intermediary nodes can be aliased with unique names.

4 Aggregation Algorithm

The algorithm we discuss here aggregates several data streams on an active node. Each stream is considered to be a periodic sequence of unreliable datagrams. The basic steps consist of the buffering of packets until they can be merged with other packets, the merging of the data, and the transmission

of the result. Also, there is a mechanism to detect packet loss and avoid indefinite waiting. Our aggregation algorithm is similar to Concast, described in [1]. While Concast is an implementation of a general purpose many-to-one communications paradigm, it does not specifically consider timeout issues due to packet loss and thus does not lend itself well to a real-time environment.

Consider a router that has to merge k data streams, $s_1 \dots s_k$, to a new data stream s_a . Assume we are given an aggregation function F that generates a packet p_{s_a} from packets $p_{s_1}, \dots, p_{s_k}$ ($p_{s_a} = F(p_{s_1}, \dots, p_{s_k})$). If a packet p_{s_i} has not been received during a period and the timeout is triggered after $t_{timeout}$, the merging function F can use either an older data packet of stream s_i or use the neutral element 0_F as a placeholder. Let us also assume for now that all sources send packets of the same size (same amount of information or samples) and with the same period. The resulting procedure for buffering and merging packets is shown in detail in Figure 4. There are four parts to the algorithm:

1. *Part I (lines 7-12): A packet arrives and is stored in the buffer.* This requires that the buffer slot for that stream is not yet used.
2. *Part II (lines 13-26): A packet arrives and its buffer slot is already taken.* This happens, when the algorithm was waiting for packets from other streams that were lost or delayed. With the arrival of a second packet from a stream, we know that it is time to send the aggregated packet. Empty buffer slots are filled with null packets and the data is aggregated and sent.
3. *Part III (lines 27-34): All buffer slots are filled.* In this case, one packet from each flow is available and we can aggregate the data and send out the result.
4. *Part IV (lines 36-46): A timeout occurred.* In this case, missing packets are replaced by null packets and the aggregated result is sent.

The packets are stored in the buffer array, b , that has n slots. The variable $bcount$ keeps track of the number of valid buffer entries. The *timer* is set to $t_{timeout}$ every time the first packet is put into the empty buffer. This way, no packet is ever stored longer than $t_{timeout}$. The *timer* is cleared (set to ∞) when the buffer is cleared.

5 Evaluation

To show the effectiveness of Aggregated Hierarchical Multicast, we first look at the reduction of link bandwidth for a given scenario. Second, we analyze the complexity and correctness of the aggregation algorithm. And third, we show measurements from a prototype implementation of the audio conferencing application.

5.1 Bandwidth and Computation

To analyze the benefits of aggregation in the network, we compare Aggregated Hierarchical Multicast with traditional many-to-one communication, where the receiver aggregates all data. Assume a balanced tree of height h with nodes of degree k . Say the receiver is the root of the tree and it wants to observe all leaves of a subtree with height l . In Aggregated Hierarchical Multicast, each leaf of the subtree sends one message to its parent, which is further aggregated until it reaches the root of the subtree. In traditional many-to-one communication, all leaves have to send a message to the root. Figure 5 shows the fraction of traffic that is necessary to observe a subtree on various

```

1: initialization:
2: clear  $b$ 
3:  $bcount \leftarrow 0$ 
4: set  $timer$  to  $\infty$ 
5:
6: receive packet  $p_{s_i}$  from stream  $s_i$ 
7: if  $b[i]$  is empty then
8:    $b[i] \leftarrow p_{s_i}$  {store packet in buffer}
9:    $bcount \leftarrow bcount + 1$  {increase buffer counter}
10:  if  $bcount = 1$  then
11:    set  $timer$  to  $t_{timeout}$  {set timer if this is the first packet in the buffer}
12:  end if
13: else
14:   {there is already a packet from source  $s_i$ }
15:   for  $j = 1$  to  $k$  do
16:     if  $b[j]$  is empty then
17:        $b[j] \leftarrow 0_F$  {fill empty buffer slots with neutral elements}
18:     end if
19:   end for
20:    $p_{s_a} \leftarrow F(b)$  {merge packets}
21:   send  $p_{s_a}$ 
22:   clear  $b$ 
23:    $b[i] \leftarrow p_{s_i}$  {store packet that was just received in buffer}
24:    $bcount \leftarrow 1$  {adjust buffer counter}
25:   set  $timer$  to  $t_{timeout}$  {set timer (since this is the first packet in the buffer)}
26: end if
27: if  $bcount = n$  then
28:   {there is one packet from each source in the buffer}
29:    $p_{s_a} \leftarrow F(b)$  {merge packets}
30:   send  $p_{s_a}$ 
31:   clear  $b$ 
32:    $bcount \leftarrow 0$  {adjust buffer counter}
33:   set  $timer$  to  $\infty$ 
34: end if
35:
36: on  $timer = 0$  do: {timeout occurred}
37: for  $j = 1$  to  $n$  do
38:   if  $b[j]$  is empty then
39:      $b[j] \leftarrow 0_F$  {fill empty buffer slots with neutral elements}
40:   end if
41: end for
42:  $p_{s_a} \leftarrow F(b)$  {merge packets}
43: send  $p_{s_a}$ 
44: clear  $b$ 
45:  $bcount \leftarrow 0$  {adjust buffer counter}
46: set  $timer$  to  $\infty$ 

```

Figure 4: Packet Merging Algorithm.

Figure 5: Traffic comparison between Aggregated Hierarchical Multicast and many-to-one communication.

Figure 6: Computation comparison between Aggregated Hierarchical Multicast and many-to-one communication.

levels compared to many-to-one communication. Even for small levels of aggregation ($l \geq 2$), the total traffic is reduced by 60 – 80%.

However, aggregation in the network has its price. Each data stream is aggregated possibly multiple times on the way to its destination compared to a many-to-one scheme, where aggregation happens only once at the destination. Figure 6 shows that for small node degrees ($k = 2$) potentially twice as many aggregation computations are necessary. Higher degree trees have less of a computational overhead. Considering that higher degree trees also require fewer transmissions, they are more favorable for Aggregated Hierarchical Multicast.

5.2 Aggregation Algorithm

Two parts of the aggregation algorithm contribute to its computational complexity. One is the per packet processing and the other is the aggregation of the stored packets. As for the per packet processing, the complexity is constant ($O(1)$). The packet merging, is $O(k)$, since we have to go through all stored packets. If we do an amortized analysis, though, we can add a credit for each packet received, which is used for the $O(k)$ computation. This results in a constant $O(1)$ amortized complexity. Note, that if we have to merge any set of packets with 0_F (due to packet loss or reordering), the amortized analysis does not hold, because no credit was added for null packets. In addition to the complexity of the aggregation, one must consider the complexity of the aggregation function F . It can be expected that for most practical applications, F requires significantly more processing than the buffering of the packets.

To show correctness of the aggregation algorithm, we show for different cases that it operates as desired and discuss the effects of jitter. The consideration of jitter is insofar important that it effects the proper operation of the algorithm. The analysis shows, though, that jitter is not increased by the aggregation step and thus limited to the amount of jitter that is introduced by the data source.

Let p_j^i be the j^{th} packet sent by source s_i . This packet is received by the router at time t_j^i . The interarrival time of packets from one source be Δt . To account for jitter, we define the random variable J that is bound my the maximum jitter j_{max} ($P[-j_{max} \leq J \leq j_{max}] = 1$). We also assume that $j_{max} < \frac{1}{2}$ to avoid reordering of packets and that the average jitter is zero ($E[J] = 0$). Thus, the expression for packet arrival times is

$$t_j^i = i \cdot \Delta t_0^i + J.$$

The time for the first packet, t_0^i , can be set arbitrarily.

As for correctness, the objective of the algorithm is to merge one packet from each source and send out the aggregated result. In the case where there is no packet loss, no jitter, and packets from all streams have the same interarrival time, it is easy to see that within a time of Δt the buffer is filled with one packet from each stream (Part I). When the buffer is full, the packets are aggregated and sent out (Part III). If we allow packet loss, the algorithm needs to detect the loss and replace the missing packet with 0_F before merging. With jitter, the algorithm also has to make sure that a delayed packet is not wrongly assumed to be lost. Therefore we give an upper bound on the allowable jitter for which the algorithm still performs correctly. The following four scenarios discuss variations of packet loss for two streams, derive an upper bound for the jitter, and extend the results to three or more streams.

5.2.1 No packet loss.

As shown in Figure 7, the first packet received from stream A is buffered. As a packet arrives from stream B , both are merged and one aggregated packet is sent out. Due to jitter, the arrival order of packets can change (e.g., from $A-B$ for the first packets to $B-A$ for the third packets in Figure 7). This does not affect the output of the algorithm, since output is only generated when packets from both streams are available. Even if the order of packets changes constantly (e.g. $A-B-B-A-A-B-B-\dots$), the algorithm will still work correctly. This assumes, though, that temporally close packets are matched by the aggregation algorithm and patterns, like $A-B-A-A-B$ do not occur.

When three packets from one stream are received in a row without any packets from the other stream (e.g., $A-B-B-B-A$), the middle B packet has to be merged with 0_F , since no packet from the other stream is available. But this only occurs for significant jitter, because three packets have

Figure 7: Aggregation algorithm with no packet loss.

Figure 8: Aggregation algorithm with packet loss and $\Delta t < t_{timeout}$.

to be received in a shorter time that two packets from another stream can be spaced apart. Thus, we have to have

$$2 \cdot (\Delta t - j_{max}) > \Delta t + 2 \cdot j_{max},$$

which gives us an upper bound for j_{max} of $j_{max} < \frac{1}{4} \cdot \Delta t$.

5.2.2 Packet loss and large $t_{timeout}$.

With the possibility of packets getting lost, the setting for the timeout value $t_{timeout}$ becomes important. As with any timer, the goal is to set the timeout no larger than it is necessary to detect packet loss. But it should not be so small that a timeout occurs when packets are just delayed due to jitter. The first observation from this is that maximum time between packets that should be merged together is $\Delta t + 2 \cdot j_{max}$. Thus, $t_{timeout} \leq \Delta t + 2 \cdot j_{max}$.

If the timeout is set to $t_{timeout} = \Delta t + 2 \cdot j_{max}$, it almost never times out on packet loss. As shown in Figure 8, if a packet gets lost, say B_1 , most likely the next packet of A (A_2 in this example) arrives before the timer expires and replaces A_1 . It can also happen that packets are merged out of order. In Figure 8, this happens for B_3 and A_4 . For some applications this is an undesired effect and can be avoided by reducing the $t_{timeout}$.

Figure 9: Aggregation algorithm with packet loss and $\Delta t \geq t_{\text{timeout}}$.

5.2.3 Packet loss and small t_{timeout} .

The packet loss scenario for a smaller timeout is shown in Figure 9. Here, a timeout occurs before A_2 arrives, which causes A_1 to be merged with 0_F and sent. This is the same behavior as in Figure 8, except that the packet is delayed less, because the timeout happens before A_2 arrives. Also, it maintains the order of merged packets for this example. The loss of A_3 causes a timeout and B_3 to be merged with 0_F , not with A_4 as it happened above. A_4 and B_4 are then properly merged together.

If the timeout becomes too small (e.g., $t_{\text{timeout}} < \frac{1}{2} \cdot \Delta t$), the algorithm might not function properly, since it times out on almost every packet, merges it with 0_F , and effectively sends twice as many packets. The ideal combination of timeout and jitter is where packets from different streams are close enough together that they can be merged together without packets from the same stream replacing each other. The maximum time between two samples from different streams is $\frac{1}{2} \cdot \Delta t + 2 \cdot j_{\text{max}}$. The minimum time between two samples of the same stream is $\Delta t - 2 \cdot j_{\text{max}}$. Thus, the algorithm works best if

$$\frac{1}{2} \cdot \Delta t + 2 \cdot j_{\text{max}} < \Delta t - 2 \cdot j_{\text{max}}.$$

This leaves an upper bound on j_{max} of $j_{\text{max}} < \frac{1}{8} \cdot \Delta t$. The timeout setting for this case should be $t_{\text{timeout}} = \frac{1}{2} \cdot \Delta t + 2 \cdot j_{\text{max}}$.

5.2.4 Three or more streams.

So far we have looked only at the case of two streams that are merged. For a larger number of streams, the above analysis changes slightly. For the lossless case, the worst case pattern becomes $A-B-B-C-C-A-\dots$ (this pattern also applies for streams with more than three sources). The minimum distance between the first B and the second C is $2 \cdot (\Delta t - 2 \cdot j_{\text{max}})$. The maximum distance between the two A s is $\Delta t + 2 \cdot j_{\text{max}}$. Thus, the jitter is limited by

$$2 \cdot (\Delta t - 2 \cdot j_{\text{max}}) > \Delta t + 2 \cdot j_{\text{max}}.$$

This gives a limit of $j_{\text{max}} < \frac{1}{6} \cdot \Delta t$ for the jitter. The analysis for the timeout in the packet loss case also changes slightly. The maximum time between a set of samples from n different streams is

Table 1: Worst-case, average, and best-case delay for aggregation algorithm

	worst case	average	best case
no packet loss	$\min(\Delta t + 2 \cdot j_{max}, t_{timeout})$	$\min((1 - \frac{1}{n}) \cdot \Delta t, t_{timeout})$	0
packet loss	$\min(\Delta t + 2 \cdot j_{max}, t_{timeout})$	$\min(\Delta t, t_{timeout})$	$\min(\Delta t - 2 \cdot j_{max}, t_{timeout})$

$(1 - \frac{1}{k}) \cdot \Delta t + 2 \cdot j_{max}$. The minimum between two samples from the same stream remains $\Delta t - 2 \cdot j_{max}$. Thus the bound for j_{max} is given by

$$(1 - \frac{1}{k}) \cdot \Delta t + 2 \cdot j_{max} < \Delta t - 2 \cdot j_{max}.$$

This solves to $j_{max} < \frac{1}{4k} \cdot \Delta t$ with the same timeout of $t_{timeout} = (1 - \frac{1}{k}) \cdot \Delta t + 2 \cdot j_{max}$. This result indicates that the tolerance for jitter decreases as the number of streams increases.

5.2.5 Delay and jitter propagation.

The worst-case, average, and best-case delay introduced by the aggregation algorithm is shown in Table 1. The delay from processing of F is not considered here, since it can be assumed to be constant and just added to the expressions in Table 1. With a timeout setting of $t_{timeout} = (1 - \frac{1}{k}) \cdot \Delta t + 2 \cdot j_{max}$, as recommended above, the worst case is always bound by the timeout. Also, the average and best-case are bound by the timeout in case of packet loss. This shows that the delay per aggregation stage is roughly limited to Δt for small jitter.

Since it is important to keep a bound on jitter for the correctness of the algorithm, we now look at the jitter that occurs on the output. (Note, that we cannot just subtract the best case delay from the worst case delay to obtain jitter, since these cases assume different offsets between streams.) If no packet loss occurs, the jitter of the output remains $\pm j_{max}$. This can be easily shown by the following argument. Assume there was no jitter and stream X sends the packet that fills the buffer and triggers the aggregation. Taking jitter into consideration, the aggregation cannot be delayed by more than j_{max} , since no packet can arrive later than j_{max} after the expected time for the packet from X . Also, the buffer cannot fill earlier than j_{max} before the expected time for the packet from X . Thus, the total jitter is $\pm j_{max}$. If we allow packet losses, a similar argument can be made. The aggregation cannot happen any earlier than j_{max} before the expected time for the packet from X . If a loss occurs, the aggregation might be delayed by $t_{timeout}$, which is typically $(1 - \frac{1}{k}) \cdot \Delta t + 2 \cdot j_{max}$. This introduces significant jitter, but packet loss should only occur infrequently.

5.2.6 Other issues.

Finally, we look at a few issues that have been neglected so far. In the above analysis, we assumed each packet contains the same amount of information. This might not be a realistic assumption. If packets contain different amounts of information, we need to buffer them in a way that we can receive several, possibly small data packets from one stream before we start aggregating. For this purpose per-stream ring buffers can be used. If data is received from a stream, it is put in the ring buffer (Part I of the algorithm). If the buffer is filled, we treat it like Part II of the algorithm. For Part III, we change to condition when to start aggregating. We require that the average amount of packet data is available from each stream. Requiring less than that, can lead to a fragmentation effect and the generation of n packets per Δt .

Figure 10: Processing Delay for 400 byte Packets and Node Degrees of $k = 1 \dots 5$.

Figure 11: Median Processing Delay for Varying Packet Sizes and Node Degrees.

Another issue is maintaining the order of packets when aggregating. For some applications it is undesirable to merge packets with different “timestamps” (as shown for packets A_4 and B_3 in Figure 8). For this purpose, we can also use the ring buffers described above. If a packet loss or reordering is detected, we just fill the buffer with 0_F up to where data is available again. When merging, we merge with these 0_F s instead of the newer data.

5.3 Measurements

To evaluate the performance of the described algorithm in a real application, we have implemented a prototype of the audio conferencing application described above. The prototype aggregates PCM μ -law encoded data streams and is limited only insofar that it uses unicast connections between nodes and does not implement all control features. All measurements were performed on a heterogeneous set of machines with Pentium processors running the NetBSD and Linux operating systems.

5.3.1 Processing Delay

The processing delay is the time between the arrival of the packet that triggers aggregation and the transmission of the aggregated packet. It varies depending on the size of the packet, the aggregation complexity, and the number of packets in the buffer. Figure 10 shows the processing times for 1000 aggregations for degrees of $k = 1 \dots 5$ and a packet size of 400 bytes. The processing delay is relatively uniform over the number of packets. Higher degree aggregations require slightly more processing.

To show the effects of varying node degree and packet size, Figure 11 shows the median delay for packets of sizes 80, 240, 400, and 800 byte. The degree of the node varies again from 1 to 5. The increase in processing time is proportional to the packet size, as should be expected. Over the range of data, there is also an increase in processing time due to higher node degrees, but it is not quite as significant.

5.3.2 Jitter

The jitter that is introduced by the aggregation affects the correctness of the aggregation and needs to be bound, as discussed in Section 5.2. The jitter for an audio data source is plotted in Figure 12. It can be seen that the average jitter is limited to 10 μ s. Looking at the jitter after traversing five aggregation steps, Figure 13 shows that it has increased to about 50 μ s.

Next, we look at jitter over a range of transmission frequencies, aggregation layers, and node degrees. Figure 14 shows the jitter as a fraction of the transmission time Δt on different layers of a hierarchy. While jitter increases on higher layers, it is still very small (less than 1%). Additionally, Figure 15 shows that the jitter is not drastically increasing for larger node degrees. Altogether, jitter does increase on higher levels of the hierarchy and for larger degree nodes, but it remains in the order of 1% of the interpacket time and therefore does not pose any problems.

These measurements indicate that the aggregation algorithm is robust and performs well over a wide range of aggregation levels, node degrees, and packet sizes.

6 Summary

We have presented a many-to-many communication scheme that uses data aggregation to scale the view of a receiver to the required accuracy level. By reducing the overall data traffic, while increasing the computational requirements only slightly, Aggregated Hierarchical Multicast improves the scalability of many-to-many communication. Measurements on the prototype implementation have shown that the aggregation algorithm is robust and performs efficiently.

References

- [1] K. L. Calvert, J. Griffioen, B. Mullins, A. Sehgal, and S. Wen. Concast: Design and implementation of an active network service. *IEEE Journal on Selected Areas of Communications*, Q2 2001.
- [2] A. T. Campbell, H. G. De Meer, M. E. Kounavis, K. Miki, J. B. Vincente, and D. Villela. A survey of programmable networks. *Computer Communication Review*, 29(2):7–23, Apr. 1999.
- [3] S. Deering. Host extensions for IP multicasting. RFC 1112, Stanford University, Aug. 1989.

Figure 12: Jitter for 400 byte Packets at Source.

Figure 13: Jitter for 400 byte Packets after 4 Aggregation Steps.

- [4] B. Joy, G. Steele, J. Gosling, and G. Bracha. *The Java Language Specification*. The Java Series. Addison-Wesley, second edition, June 2000.
- [5] P. Mockapetris. Domain names - implementation and specification. RFC 1035, Network Working Group, Nov. 1987.
- [6] J. Moy. MOSPF: Analysis and experience. RFC 1585, Network Working Group, Mar. 1994.
- [7] L. Peterson, ed. NodeOS interface specification. Technical report, AN Node OS Working Group, Feb. 1999.
- [8] D. L. Tennenhouse, J. M. Smith, W. D. Sincoskie, D. J. Wetherall, and G. J. Minden. A survey of active network research. *IEEE Communications Magazine*, 35(1):80–86, Jan. 1997.

Figure 14: Jitter for Various Transmission Frequencies and Aggregation Layers expressed as Fraction of Δt

Figure 15: Jitter for Various Transmission Frequencies and Node Degrees expressed as Fraction of Δt