Center Placement Algorithms for Large Multicast Groups

A. S. Weigmann and J. Nonnenmacher and E. W. Biersack
 Corporate Communications Department, Institut EURECOM
 B.P. 193, 06904, Sophia Antipolis Cedex, FRANCE
 {nonnen, erbi}@eurecom.fr

Abstract

An increasing number of distributed applications require a specific form of multicast called dissemination, in which a single source *reliably* transfers data to multiple receivers. Reliability leads for large groups of receivers (100s or 1000s of participants) to the problem of feedback implosion at the source and to a decrease of transmission efficiency. The cluster approach was identified to have excellent scalability with the number of receivers. It partitions the multicast delivery tree into clusters, where a representative in the cluster called center is used for local feedback processing and local transmission. Up to now, clustering/center placement has been done administratively or based on network addresses. Needed are center placement algorithms, allowing the introduction of placement criteria based on the network topology and on delay. In this work, three center placement algorithms designed for static multicast groups are presented and simulation results are shown in order to assess their performance.

Keywords

Multicast, Feedback Implosion, Local Retransmission, Center Location

1 INTRODUCTION

Emerging high speed networks and the widespread availability of multimedia capable workstations have drastically increased the demand for distributed applications. Along going with this evolution comes an increasing need for multicast, allowing to distribute data to many users while efficiently using scarce communication bandwidth.

Many applications require a specific form of multicast called dissemination [1, 19, 2, 20], in which a single source *reliably* transfers data to multiple receivers. Applications include electronic newspaper distribution, WWW in-advance caching, file transfers and shared whiteboard (see also [1, 13]).

While some applications, like the shared whiteboard, support *dynamic* multicast groups,

where receivers dynamically join or leave the group during transmission time, numerous other ones, like newspaper distribution or file transfers, are designed for *static* groups of receivers, which remain unchanged throughout the entire transmission.

IP Multicast [4], which forms the basis of the MBone [6, 16] provides scalable, but not reliable multicast. Reliability requires that either data packets are acknowledged by the receivers (ACKs) or that packet losses are reported using negative acknowledgments (NACKs). For larger groups of receivers (100s or 1000s of participants), this leads to the problem of packet implosion at the source [15]: The source suffers from the significant overhead incurred due to the processing of the ACKs or NACKs from the receivers.

The research community proposed two directions to deal with the scalability problem:

- *timer approaches* [8, 7] that introduce a certain asynchrony among receivers to protect the source from an implosion and
- *cluster approaches* [20, 13, 9] that achieve the same goal by introducing a hierarchy into the multicast tree.

In the cluster based approaches, the multicast delivery tree is portioned into **clusters**. Every cluster has a representative, the **center**. Receivers do not acknowledge received packets directly to the source but only to the center of their cluster.

All centers are organized in a hierarchy. A center is responsible for all receivers in its cluster as well as for other centers located one level below in the cluster hierarchy. A center receives ACK packets from receivers and centers it is responsible for and forwards a single acknowledgment packet to the center in the next higher hierarchy level (ACK accumulation).

Most of the cluster based approaches presented so far either use centers placed administratively in the network [13] or rely on centers placed algorithmically *within* a group (cluster), which is either defined by hand or based on network- (subnet-) addressing [20, 11]. Administratively placed centers or clusters do not scale well for large groups of receivers and clustering based on addressing does not reflect the real network structure (connectivity and delay relations) but only the address structure.

Needed are center placement algorithms, allowing the introduction of placement criteria based on the network topology and on delay relations and not only on addressing schemes. In this work, criteria for a “good” center placement are stated. Three center placement algorithms designed for static multicast groups are presented and simulation results are shown in order to assess their performance.

The remainder of this paper is organized as follows. In section 2, the center placement problem is described in a formal way. In section 3, placement algorithms are presented and section 4 shows simulation results. In section 5, final conclusions are drawn.

2 THE CENTER PLACEMENT PROBLEM

The problem of placing centers in a multicast tree can be looked at from a graph-theoretical point of view. This allows to approach the problem without the need to consider implemen-

tational aspects, which are linked very closely to a specific type of network. Once the center placement problem is solved in theory it can be transferred to reality and implemented.

The graph-theoretical point of view abstracts from the real connection endpoints (receivers) and considers just multicast routers. Regarded is a directed **tree** (routers = vertices, links = edges) rooted at the multicast source. Edge weights represent link delays. **Receivers** can be located on every vertex except on the tree's root. All leaves of the tree are receivers, but vertices that are not leaves can also be receivers.

Centers can be placed on all internal vertices, thus not on leaves. The root is implicitly a center. The **number of centers** in the tree is called **NOC**.

The criteria for center location are based on a given **optimal cluster size Z** and the minimization of the end-to-end delay between the source S and the receivers $\in R$ in a multicast tree T .

The following definitions help to specify the criteria for center placement.

Definition 1 Directly Served

A receiver r_k (a center c_j) is **directly served** by a center c_i when no other center is located on the path in the tree T between c_i and r_k (between c_i and c_j). c_i is not directly served by itself.

Definition 2 Cluster with Center c_i

The **cluster with center c_i** is defined as the set of all the receivers and centers directly served by c_i , c_i itself, and all the edges interconnecting the elements of this set.

Definition 3 Cluster Size $cs(c_i)$

The **cluster size $cs(c_i)$** of the cluster with center c_i is the number of centers and receivers directly served by c_i .

2.1 Cluster size

The cluster size plays a major role in the success of reliable multicast with clustering. The amount of feedback a center c_i is capable to process imposes a scalability limit on the cluster size. Also the number of packets a center has to retransmit for its cluster increases with the number of receivers or further centers directly served by c_i and has a major influence on the performance.

In order to limit the cluster size of any cluster with center $c_i \in C$ to a given optimal cluster size Z , a **cluster size constraint (CSC)** will be introduced:

$$\forall c_i \in C : cs(c_i) \leq Z \quad (1)$$

The **mean squared error of cluster sizes MSE** is defined as

$$MSE = \frac{1}{NOC} \sum_{i=1}^{NOC} (Z - cs(c_i))^2 \quad (2)$$

The minimization of the *MSE* will lead to a optimal center placement with respect to a given optimal cluster size Z . Minimizing the *MSE* is the first criterion for center placement and is referred to as the *Cluster Size Criterion*.

2.2 End-to-End Delay

In multicast connections the end-to-end delay has a big influence on the performance of a reliable transport protocol. At every center c_i feedback from directly served centers and receivers must be processed. The end-to-end delay between a receiver r_j and the source S is composed of the delay of the links the packet transverses and a number of **center delays** cd experienced at every center. The cumulative delay is the **receiver delay** $D(r_j)$, defined as the sum of the edge delays $d(v_k, v_l)$ on the path from the source S to a receiver r_j in the tree T plus the product of the number of centers on the path and a constant delay cd experienced at every center. The end-to-end delay of the whole multicast connection is defined as the **maximal receiver delay** mDR of all r receivers:

$$mDR = \max_{j=1..r} \{D(r_j)\} \quad (3)$$

The objective, referred to as the *Delay Criterion*, is to minimize the delay mDR . Please note, that the objective is not the same as to minimize the maximal number of centers on the path from the source to a receiver. Having an arbitrary center placement, the path to the receiver that yields mDR may be different from the path where the most centers can be found on.

3 CENTER PLACEMENT ALGORITHMS

Static multicast trees occur to be important in the context of reliable multicast as there exist numerous applications like newspaper distribution, WWW in-advance-caching and software distribution. All three examples have in common that the group of receivers is determined preceding the multicast transmission and that no further receivers will join or leave this group after the transmission has started.

Static multicast trees are also the easier case, as the center placement can be done in advance. The placement is done before the actual multicast transmission starts and is not time critical. Also extensive exchange of status information (for example information on the tree structure) between nodes would be possible (although not desirable).

In the following sections, three placement algorithms are presented:

- Center Placement by Split and Shift (CPSS)
- Center Placement by Shift and Merge (CPSM)
- Center Placement by Dynamic Programming (CPDP)

Two cases will be considered for every algorithm:

1. **Placement with cluster size constraint (CSC):** Center placement is done with respecting the cluster size constraint (CSC) defined in section 2.1.
Note that if a center is placed on a vertex with an outdegree* greater than Z , the cluster size constraint cannot be met as the smallest feasible cluster size equals the outdegree.* Anyway, this case is only a theoretical one. In practice, and also in the trees used for simulations, the maximal outdegree remains below the optimal cluster size Z .*
2. **Placement without cluster size constraint (CSC):** The cluster size constraint CSC is dropped. This allows more flexibility for center placement at the expense of a non-deterministic maximal cluster size. Note that only the cluster size constraint is removed – the objective to minimize the mean squared error of the cluster size remains valid.

Before moving on to the different placement algorithms, two more expressions are defined. Although the algorithms can also be implemented without them, they allow to increase the efficiency in terms of execution speed of the algorithms both in the simulation environment and in a real implementation.

Figure 1: Furthest Directly Connected Descendant (FDCD). (a) Original tree, all nodes are routers (no receivers), the root is implicitly a center, (b) center on node 1: $cs = 1$, (c) center on node 2: $cs = 1$, (d) center on nodes 3 and 4: $cs = 2$.

Assume the following situation (see also Figure 1): A center c_i is located on a vertex v_i . Centers directly served by c_i are located on descendants of v_i . It can be seen that the size $cs(c_i)$ of the cluster with center c_i changes, if a directly served center is shifted from a vertex v to its children, if (i) vertex v has an outdegree > 1 ; or (ii) it is a receiver. In order to formalize this effect, **directly connected nodes** and **furthest directly connected descendants (FDCDs)** are defined:

*The number of children of a node x in a tree T is called the *outdegree* of x .

*All children of the vertex the center is placed on are either leaves or internal vertices with centers.

*Work done by the Routing community considers average outdegrees between 2 and 8 [18] and maximal outdegrees of 4 and 5 [12] for networks. The outdegree for multicast trees is typically smaller than the outdegree in the network. The optimal value Z is not known yet; it will be in the order of 20–30.

Definition 4 Directly Connected

A vertex $v_j \in V_t$ located in a tree $T = (V_t, E_t)$ is defined as being **directly connected** to an ancestor v_i if there are no other vertices with an outdegree > 1 , centers or receivers on the path between v_i and v_j .

Definition 5 Furthest Directly Connected Descendant (FDCD)

In a tree $T = (V_t, E_t)$ a directly connected descendant $v_j \in V_t$ of a vertex $v_i \in V_t$ is called a **furthest directly connected descendant (FDCD)** of v_i if there exists no other directly connected vertex $v_k, k \neq j$ with a path between v_i and v_k that includes the path from v_i to v_j .

3.1 Split and Shift (CPSS)

The first algorithm called Center Placement by Split and Shift (CPSS) belongs to the group of *greedy* algorithms [3]. *Greedy* algorithms always make the choice that looks best at the moment. That is, they make a locally optimal choice in the hope that this choice will lead to a globally optimal solution.

CPSS starts the center placement by building the first cluster with its center located on the root of the tree. It places directly served centers on the root's furthest directly connected descendants (FDCDs) and subsequently tries to *shift* them towards the leaves. Centers are always shifted to their furthest directly connected descendants (FDCDs). Centers that must be shifted and that are located on a node with an outdegree > 1 are *split* and all FDCDs of this node become centers.

Once the cluster is created, the algorithm is repeated recursively for all placed (directly served) centers.

The directly served centers are shifted due to the following rules:

1. Minimize the maximal receiver delay mDR .
2. If the delay cannot be optimized any further, try to optimize the size of the cluster.

Minimizing the maximal receiver delay mDR does not imply that the total number of centers in the tree has to be as small as possible. It means that the number of centers on the longest path (in terms of delay) has to be as small as possible, as every center on the path to a receiver causes an additional center delay cd . The longest path *before* and the longest path *after* the center placement do not necessarily have to correspond.

In order to keep the number of centers on a certain path as small as possible, the centers on this path must be located as far apart from each other as possible. Following this rule, CPSS always tries to shift the center on the longest path towards the leaves. If a center is shifted onto a leaf, it will be removed since centers located on leaves do not make sense (they cannot serve any other node).

CPSS does not shift the center on the longest path if the cluster size constraint (CSC) is violated. In this case, CPSS tries to shift the center on the second longest path. If this one cannot be shifted either, it tries to shift the one on the third longest path and so on. If it

does not succeed in shifting any center, the cluster is “finished” and the clusters for the just placed directly served centers can be built.

3.2 Shift and Merge (CPSM)

The second algorithm presented, called Center Placement by Shift and Merge (CPSM), also belongs to the group of *greedy* algorithms [3]. Unlike the CPSS algorithm, CPSM starts building clusters at the leaves and then moves subsequently towards the source. It concentrates on optimizing the cluster sizes (and thereby the mean squared error of the cluster size MSE). Less effort is spent in minimizing the maximal receiver delay mDR since CPSM does not explicitly try to reduce the number of centers on the longest path (in terms of delay), like CPSS does.

The first cluster is built by randomly choosing a leaf in the multicast tree and placing a center on this node (Figure 2 (a) and (b), node 4).

Figure 2: Center Placement by Shift and Merge (CPSM), part 1: (a) Tree without centers. Big circles represent receivers, small circles represent routers, (b) placing center on leaf, (c) shifting center towards root and placing new center on leaf in sibling’s subtree.

In the following steps, this center will be *shifted* towards the root until it is located on a vertex that has one or more siblings* (node 2 in the example). Now, in every subtree rooted at one of these siblings, a center has to be placed on a randomly chosen leaf (Figure 2 (c), node 7), which is then again propagated towards the root. If a center encounters a node with one or more siblings (node 5), again, there has to be placed a center on a randomly chosen leaf (node 6) in every subtree rooted at a sibling and so on (Figure 3 (a)).

This leads to a situation, where a vertex (node 3) has two or more children a center is located on. Now a selection process starts, deciding, which of the centers are shifted to the parent node. First, this decision is made for the center on the longest path (in terms of

*Two nodes having the same parent are called **siblings**.

delay), then for the center on the second longest path, then for the one on the third longest path etc. Note, that already placed centers contribute to the length of a path. If more than one center is shifted, the centers will be *merged* to a single center (Figure 3 (b)). If no center

Figure 3: Center Placement by Shift and Merge (CPSM), part 2: (a) Shifting center towards root and placing new center on leaf in sibling's subtree, (b) two centers have been shifted to their parent and have been merged, (c) no center has been shifted, a new center will be placed on the parent.

is shifted, a new center will be placed on the parent (Figure 3 (c), node 1).

The rules for deciding, which centers are moved, are the following:

- A center located on a leaf is always moved.
- If the new cluster size $cs_{new}(\text{parent}(V))$ after shifting a center on node V to its parent is closer to the optimal value Z than the cluster size $cs(V)$ of V before the shifting and not higher than Z , the center will be moved.
- If the cluster size $cs_{new}(\text{parent}(V))$ after shifting a center on node V to its parent is equally close to the optimal value Z as the current cluster size $cs(V)$ (without shifting) and if V lies on the path from parent(V) to the node in the subtree of parent(V) yielding the mDR or no center has been shifted to the parent of V yet, the center will be shifted.
- In all other cases, the center will not be shifted.

3.3 Dynamic Programming (CPDP)

The third algorithm presented is based on the method of dynamic programming, which is related to *divide-and-conquer* [3] and *branch-and-bound* [14] in the sense that it partitions a problem into subproblems and *intelligently* enumerates the feasible points of a combinatorial optimization problem.

To solve the problem, thus to find the best placement for a tree, dynamic programming

is used in combination with a *local search*. Supposed, the best placements for *all* subtrees, as well as all important parameters of these placements, like the number of clusters in each subtree, the *mDR* in each subtree, etc. are known. Then, the cluster with its center located on the tree's root is built simply by placing directly served centers systematically on the root's descendants and by evaluating the different combinations. In practice, this means: Centers are placed on the root's furthest directly connected descendants (FDCDs) and then propagated towards the leaves. At every node with outdegree > 1 , centers are split. Evaluation of the different center placements is done using a metric composed of the relative increase of the maximal receiver delay *mDR* and the normalized mean squared error of the cluster sizes (*MSE*):

$$Metric = 0.5 \cdot \frac{mDR_{with\ centers} - mDR_{without\ centers}}{mDR_{without\ centers}} + 0.5 \cdot \frac{MSE}{Z^2} \quad (4)$$

Once, the optimal placement for a (sub-)tree has been found, the center configuration as well as the parameters characterizing the placement, like the squared errors of the cluster size, the number of clusters, the *mDR*, etc. will be stored allowing them to be reused as the result of a solved problem.

In order to find the best placement, all feasible combinations must be evaluated and the number of feasible combinations is growing exponentially with the number of nodes in the tree. In order to limit the search to a bearable number of combinations, a **maximal search depth without improvement *MSD*** is introduced. If a certain center has been propagated for *MSD* steps without experiencing any improvement in the quality of the placement in terms of cluster size and delay, this center will be abandoned and not propagated any further. The introduction of this limit restricts the number of evaluated combinations, thereby limiting computational complexity, but also excluding numerous combinations. The optimal solution might not be found, the result has sub-optimal character.

4 RESULTS

In the last sections, three center placement algorithms have been presented. In order to evaluate their performance, they were implemented and executed on a number of networks.

As input data, 10 networks with 200 nodes, an average outdegree of 3.0 and an average link delay of 412.38 were created following Doar [5], who modified the method of Waxman [17] to avoid an increasing outdegree for an increasing number of nodes. The network construction of Waxman/Doar is commonly used by the Multicast Routing community for comparing the performance of different multicast routing algorithms.

For each of the 10 networks, 10 different trees with 10, 20, 40, 60, 80, 100, 120 and 140 receivers each on randomly chosen locations have been constructed using a shortest path tree (SPT) algorithm [5]. This means that for a given number of receivers, 100 trees have been created, which results in a total of 800 trees.

The three center placement algorithms CPSS, CPSM and CPDP were executed on every SPT:

- *with* respecting the cluster size constraint (with CSC)
- *without* respecting the cluster size constraint (without CSC)

For the optimal cluster size Z the values 5, 10, 15, 20, 25 and 30 have been used, for the center delay cd the values 100.0, 200.0 and 300.0, resulting in 18 different parameter sets.

All values presented in the following sections are mean values of 100 evaluated trees. Indicated intervals are 95% confidence intervals [10].

4.1 Evaluation Metrics

In order to compare different algorithms, an exact quantification is needed. As the effects of the cluster size and the maximal receiver delay to a multicast connection are not known exactly, they will be treated separately. For the performance evaluation two metrics are used. The quality of the placement with respect to the cluster size is measured by the normalized MSE , defined as cluster size error CSE :

$$CSE = \frac{\sum_{i=1}^{NOC} (Z - cs(c_i))^2}{Z^2 \cdot NOC} \quad (5)$$

The quality of the center placement with respect to delay is measured by the delay increase cost DIC representing the relative increase of the maximal receiver delay due to the center placement:

$$DIC = \frac{mDR_{with\ centers} - mDR_{without\ centers}}{mDR_{without\ centers}} \quad (6)$$

In the next sections, the proposed algorithms will be evaluated using the above introduced quality metrics.

4.2 Split and Shift (CPSS)

(a) Number of Centers NOC

Figure 4 shows the distribution of the distances in hops between the tree's root and the centers for 40 receivers trees, $Z = 20$ and $cd = 100.0$. It clearly shows that CPSS tends to locate centers further away from the root than do CPSM and CPDP. Note that the distance 0 is caused by the center located on the root. The lower normalized count for distance 0 for CPSS in contrast to CPSM and CPDP results from the higher number of centers for CPSS in general (see also Figure 8).

Figure 5 shows for CPSS the number of centers that lie on the paths from the root to the receivers for 40 and 140 receivers trees ($Z = 10$, $cd = 100.0$).

Dist. Source-Centers (Hops), with CSC, 40 Recv, $Z=20$, $CD=100$

Figure 4: Distribution of distances from source to centers in hops. CPSS, CPSM and CPDP with cluster size constraint

Centers on Paths to Recv, CPSS (with CSC), $Z=10$, $CD=100$

Figure 5: Centers on the paths to receivers, CPSS, with cluster size constraint, 40 and 140 receivers.

Cluster Size Distribution, with CSC, 140 Recv, $Z=10$, $CD=100$

Figure 6: Cluster size distribution for CPSS, CPSM and CPDP with cluster size constraint.

(b) Cluster Size Error CSE

CPSS suffers from the fact that numerous clusters near the leaves are sparsely filled. Figure 6 shows the cluster size distributions for the three algorithms CPSS, CPSM and CPDP with cluster size constraint (CSC) for 140 receivers and an optimal cluster size of $Z = 10$ ($cd = 100.0$). Conspicuous are the shapes of the distributions. CPSS often succeeds in reaching the optimal cluster size Z . The high portion of sparsely filled clusters emerges from those clusters located in the leaves areas of the trees. For 140 receivers, $Z = 10$ and

$cd = 100.0$, 63.7% of all clusters created by CPSS have a size of $cs(c_i) \leq 5 = \frac{Z}{2}$ (for comparison: CPSM: 32.2%, CPDP: 21.6%).

(c) Delay Increase Cost *DIC*

In section 2, it has been pointed out that the maximal receiver delay mDR plays a key role in the performance of reliable multicast connections. The relative increase of the mDR (6) is regarded as the basic criterion for evaluating the quality of the placement algorithms in terms of delay.

Figure 7: Receiver delay distribution before and after center placement for CPSS with cluster size constraint for 40 and 140 receivers

Figure 7 shows the impact of centers on the general distribution of receiver delays: * The delays are increased (the shape is moved to the right) and the delays are distributed more equally (flatter shape).

4.3 Shift and Merge (CPSM)

Center Placement by Shift and Merge (CPSM) starts placing centers at the leaves of the tree and concentrates primarily on optimizing the cluster size. Consequently, the performance in terms of the cluster size error CSE is very good and outperforms the Center Placement by Split and Shift (CPSS), while the performance in terms of delay is lower than the one of CPSS.

(a) Number of Centers *NOC*

In contrast to CPSS, CPSM avoids the effect of sparsely sized clusters (see Figure 6). Building clusters from the leaves to the root results in only one really badly sized cluster, the one

*In the figure, receiver delays have been put into bins with a width of 200.0

located at the tree's root. As the average cluster size is closer to the optimal value Z , the total number of clusters in a tree is less than for CPSS (Figure 8). Note, that the number of centers (NOC) in the trees grows approximately linear with the number of receivers $NOC \approx k \cdot r$ with a gradient k inversely-proportional to the optimal cluster size Z : $k \sim \frac{1}{Z}$, $\frac{r}{Z} \gg 1$.

Figure 8: Number of centers NOC . CPSS, CPSM and CPDP with cluster size constraint.

Figure 9: Cluster Size Error CSE . CPSS, CPSM and CPDP with cluster size constraint.

Figure 10: Delay Increase Cost DIC . CPSS, CPSM, CPDP with cluster size constraint (CSC).

As can be seen in figure 9 performs CPSM much better than CPSS in terms of the cluster size error CSE . CPSM prefers to merge clusters instead of just expanding clusters along

the longest path (in terms of delay), like CPSS does.* This is in harmony with the objective of the algorithms: While CPSM optimizes the mean squared error of the cluster size (*CSE* criterion), CPSS tries to optimize the delay criterion (*DIC*), this can also be noticed in the better delay performance of CPSS (Figure 10). In contrast to the *CSE* criterion, which offers approximately constant values for different tree sizes, the *DIC* criterion produces lower values for smaller trees than for bigger ones. This effect cannot be attributed to the algorithms, but is the result of the testing environment: Random networks with $|V| = 200$ nodes have been used for the simulations. The larger the trees get, the higher the ratio receivers/nodes becomes. In total, more centers are placed in the trees, also on the *mDR*-paths, and the performance is decreasing.

4.4 Dynamic Programming (CPDP)

CPDP uses the method of dynamic programming. Instead of performing a *full* search on all nodes of the multicast tree for finding the optimal center locations, the amount of time needed to find a good solution is reduced by performing a *local* search, not considering combinations that are unlikely to lead to a good result.

The next paragraphs will show that there exists a close correlation between the extent of the local search and the quality of the resulting placement.

Figure 11: Number of centers *NOC*. Center Placement by Dynamic Programming (CPDP), different weights for the two criteria

(a) Number of Centers *NOC*

Comparing CPSM and CPDP (Figure 8), a slightly lower number of centers can be observed for CPDP. This little difference needs no further explanation, since the exact number of

*Note that this strategy can also be altered easily (for example, in order to put more emphasis on the delay criterion) by modifying the move-rules described in section 3.2.

Figure 12: Cluster Size Error CSE and Delay Increase Cost DIC . CPDP (a) and CPDP (b), without cluster size constraint

centers placed by CPDP depends on the placement metric and thus might differ in both directions when the weights of the different criteria are altered.

CPDP's optimization function can be tailored by weighting the two criteria, cluster size and the delay. While CPDP (a) in Figure 11 weights both criteria equally (0.5 and 0.5), CPDP (b) uses a weight of 0.9 for the delay criterion and only 0.1 for the cluster size criterion. Giving more emphasis to the delay criterion results in a slightly higher number of clusters, a higher cluster size error CSE and a much lower delay increase cost DIC (Figure 12).

(b) Cluster Size Error CSE

Figure 13 shows the quality of the placement for CPDP without cluster size constraint (CSC) in terms of the cluster size error CSE for $Z = 5, 10, 20, 30$. We see a high correlation between the optimal cluster size Z and the required maximal search depth without improvement, MSD . While for $Z = 5$ $MSD = 2$ suffices* (Figure 13 (a)), there exists already a big penalty in the quality between a MSD of 2 and 4 for $Z = 10$ (Figure 13 (b)). Conspicuous is the fact, that for $Z = 10$ (Figure 13 (b)) the mean squared error of the cluster size in the trees with 10 receivers does not equal 0, which would be indicated by a CSE value of 0. This shows that the best solution, the one with all 10 receivers in one cluster of size 10, is not always found if the maximal search depth MSD is chosen too low.

*There is about no difference between the two lines for $MSD = 2$ and $MSD = 4$.

Figure 13: Cluster Size Error CSE for CPSM and CPDP without cluster size constraint

5 CONCLUSIONS

In the last sections, three center placement algorithms have been presented and evaluated. Two of them (CPSS and CPSM) start placing centers on either end of the tree and then proceed consecutively to the other end. The third one (CPDP) uses dynamic programming combined with a local search. CPSS performs very well in terms of delay, but does not try to optimize the cluster sizes resulting in a high cluster size error CSE . CPDM performs very well in terms of the cluster size error CSE , its delay increase cost DIC is worse than the one of CPSS. As long as the maximal search depth MSD is chosen high enough, CPDP has a good performance, both in terms of delay and cluster size, as its optimization function can be altered very easily to meet the requirements. However, CPDP has a high computational

cost which grows exponentially with the optimal cluster size Z . The results obtained lead to the following conclusions:

- For small optimal cluster sizes ($Z \in [5 \dots 10]$), CPDP is recommended, since its optimization function can be altered easily to meet predefined requirements.
- For bigger values of Z , CPDP's complexity is too high and preference should be given to CPSS or CPSM.
- If delay aspects have highest priority and cluster sizing only plays a minor role, CPSS should be used.
- If cluster size is important, CPSM should be used as it obtains small values for CSE and has a low computational cost.

REFERENCES

- [1] Azer Bestavros. Speculative data dissemination and service to reduce server load, network traffic and service time for distributed information systems. In *Proceedings of ICDE'96: The 1996 International Conference on Data Engineering*, New Orleans, Louisiana, USA, March 1996.
- [2] D. R. Cheriton. Dissemination-oriented communication systems. DSG Working Paper, 1992.
- [3] T. H. Cormen, C. E. Leiserson, and R. L. Rivest. *Introduction to Algorithms*. The MIT Press, 1990.
- [4] S. Deering. Host extensions for ip multicasting. *Internet Request for Comments*, RFC 1112, August 1989.
- [5] M. Doar and I. Leslie. How bad is naïve multicast routing. In *Proceedings of INFOCOM'93*, volume 1, pages 82–89. IEEE, 1993.
- [6] Hans Eriksson. MBONE: The multicast backbone. *Communications of the ACM*, 37(8):54–60, August 1994.
- [7] Sally Floyd, Van Jacobsen, Steven McCanne, Ching-Gung Liu, and Lixia Zhang. A reliable multicast framework for light-weight sessions and application level framing. In *Proceedings of SIGCOMM'95*, volume 25, pages 342–356. ACM, October 1995.
- [8] Matthias Grossglauser. Optimal deterministic timeouts for reliable scalable multicast. In *Proceedings of INFOCOM'96*, volume 3, pages 1425–1432, San Francisco, CA, USA, March 1996. IEEE.
- [9] Markus Hofmann. A generic concept for large scale multicast. In *International Zürich Seminar*, Zürich, Switzerland, February 1996.
- [10] Robert V. Hogg and Allen T. Craig. *Introduction to Mathematical Statistics*. Macmillan Publishing Co., Inc., New York, 1989.
- [11] H. W. Holbrook, S. K. Singhal, and D. R. Cheriton. Log-based receiver-reliable multicast for distributed interactive simulation. In *Proceedings of SIGCOMM'95*, volume 25. ACM, October 1995.
- [12] Vachaspathi P. Kompella, Joseph C. Pasquale, and George C. Polyzos. Multicasting for multimedia applications. In *Proceedings of INFOCOM'92*, pages 2078–2085. IEEE, 1992.

- [13] John C. Lin and Sanjoy Paul. Rmtp: A reliable multicast transport protocol. In *Proceedings of INFOCOM'96*, volume 3, pages 1414–1424, San Francisco, CA, USA, March 1996. IEEE.
- [14] Ch. H. Papadimitriou and K. Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1982.
- [15] Sridar Pingali, Don Towsley, and James F. Kurose. A comparison of sender-initiated and receiver-initiated reliable multicast protocols. In *Proceedings of Sigmetrics, Conference on measurement and modeling of Computer Systems*, pages 221–230, Santa Clara, CA, USA, May 1994. ACM.
- [16] Macedonia. M. R. and D. P. Brutzmann. Mbone provides audio and video across the internet. *IEEE Computer*, April 1994.
- [17] B. M. Waxman. Routing of multipoint connections. *IEEE JSAC*, 6(9):1617–1622, December 1988.
- [18] Liming Wei and Deborah Estrin. The trade-offs of multicast trees and algorithms. In *Proceedings of ICCCN'94*, San Francisco, CA, USA, Sept 1994.
- [19] T. Yan and H. Garcia-Molina. SIFT - A tool for wide-area information dissemination. In *Proceedings of the 1995 USENIX Technical Conference*, pages 177–186, 1995.
- [20] Rajendra Yavatkar, James Griffioen, and Madhu Sudan. A reliable dissemination protocol for interactive collaborative applications. In *Proceedings of ACM Multimedia*, pages 333–344, San Francisco, CA USA, 1995. ACM.