

# Scalable Adaptive Hierarchical Clustering<sup>\*</sup>

Laurent Mathy<sup>1</sup>, Roberto Canonico<sup>2</sup>, Steven Simpson<sup>1</sup>, and David Hutchison<sup>1</sup>

<sup>1</sup> Lancaster University, UK

{laurent, ss, dh}@comp.lancs.ac.uk

<sup>2</sup> University Federico II, Napoli, Italy

roberto.canonico@unina.it

**Abstract.** We propose a new application-level clustering algorithm capable of building an overlay spanning tree among participants of large multicast sessions, without any specific help from the network routers. This algorithm is based on a unique definition of zones around nodes and an innovative adaptive cluster size distribution. The proposed method finds application in many context where large-scale overlay trees can be usefull: application-level multicasting, peer-to-peer networks and content distribution networks (among other things).

## 1 Introduction

More than a decade of research in multicast technologies demonstrates the need for large-scale (application-level) overlay structures.

Tree-based ACKs (TRACKs) have been identified as the most appropriate approach to providing real-time and scalable delivery guarantees to groups of receivers [8]. In this scenario, the overlay provides a control structure. This approach is further re-inforced with the recent emergence of the Source-Specific IP multicast model [4][3] which is an asymmetrical service where only a designated source can send in multicast to the group.

More recently, reasons for the lack of widespread deployment of IP multicast have been identified [2]. These indicate that ubiquitous rollout of IP multicast services may, even if at all possible, take a very long time. In such circumstances, overlays represent an attractive alternative to IP multicast for data dissemination among members of multicast groups. This is the case in Content Delivery Networks (CDN) where application-level multicast overlays are often used, for example, for the distribution of multimedia data from primary to secondary servers.

Peer-to-peer (p2p) applications also rely heavily on overlays. Here, the overlays are used to propagate search strings among the nodes of the p2p network, in order to discover the location of the desired content. Very often, users of p2p networks statically configure a few nodes to peer with, which can result in a non-efficient, almost “chaotic” overlay.

To date, all the above scenarii lack, but would greatly benefit from, effective algorithms to build large-scale, efficient overlays. In this paper, we propose a

---

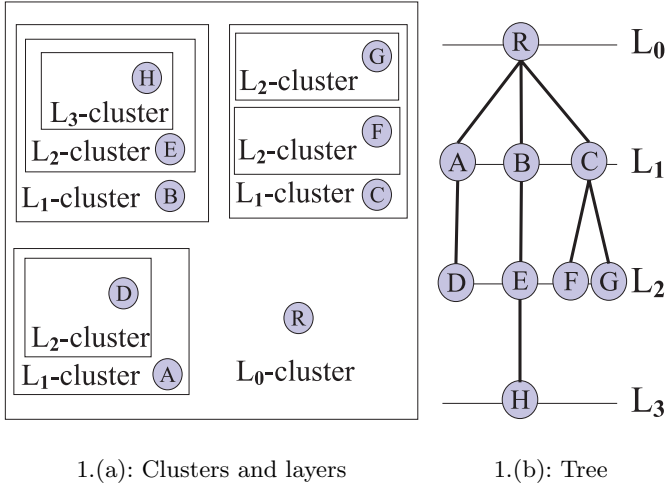
<sup>\*</sup> This work was supported by the BT Alpine Project.

new method designed to build such large-scale overlays, without requiring any special support from the network routers. This method is based on the concept of clustering.

## 2 Adaptive Hierarchical Clustering Algorithm

### 2.1 General Strategy and Goal

The algorithm described in this section is designed to build, recursively, a hierarchy of *clusters*. A cluster is represented by a *cluster head* and is composed of the cluster head and other nodes “close” to the cluster head. The algorithm is “recursive” in the sense that each cluster is divided into sub-clusters, whose (sub-)cluster heads are constituent nodes of the original cluster. The hierarchy of clusters is organised into *layers*, where layer  $L_i$  is composed of the cluster heads of (sub-)clusters that divide  $L_{i-1}$ -clusters (i.e. clusters whose head is in layer  $L_{i-1}$ ). This is illustrated in figure 1.(a). For instance, in this figure, the  $L_1$ -cluster headed by C is composed of C, F and G. This cluster contains two  $L_2$ -clusters, respectively headed by F and G.



**Fig. 1.** Cluster hierarchy.

The principle of the algorithm is that, starting at layer  $L_0$  with a top-level cluster containing all the nodes in the hierarchy and whose cluster head is a well known node called the *root* of the hierarchy, clusters are recursively divided into

sub-clusters, until all clusters obtained are “singleton-clusters” containing only their cluster head<sup>1</sup>.

The cluster hierarchy thus built forms a logical tree spanning all the cluster heads (e.g. all the nodes) in the hierarchy (see figure 1.(b)). Consequently, the state to build and maintain this hierarchy can be distributed among all the nodes in the hierarchy such that each node in layer  $L_i$  only needs to record its parent cluster head (i.e. the  $L_{i-1}$ -cluster head whose cluster it belongs to) and the  $L_{i+1}$ -cluster heads that are members of its own cluster. For instance, in figure 1.(a), B records R as its parent cluster and E as its child.

2.2 Workings of the Algorithm

The algorithm is distributed and based solely on unicast communications. In other words, it does not rely on any special network support.

One of the central ideas in the algorithm is that any node (i.e. any cluster head) sees the rest of the world as a set of concentric rings (which we call *zones*), centered on the node itself. Each zone starts where the previous one finishes and the zones are numbered in increasing order, starting at 0 for the smallest ring (see figure 2). The actual size of each ring, as well as the distance measurement used to define it (e.g. delays, throughput, etc.), is unimportant for the general workings of the algorithm. With each zone, a distance called a *radius* is also defined. Again, the size of the radius is unimportant for the workings of the algorithm (but its distance measurement is the same as the measurements used for defining the zones).

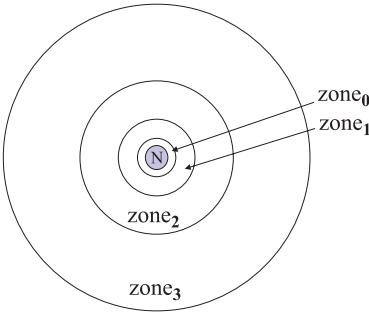


Fig. 2. Zones associated with a node.

The scalable hierachical clustering algorithm works as follows. The cluster hierarchy is rooted on a well known entity called the *root*. A node desiring to join the cluster hierarchy first measures its distance to the root, and then sends

<sup>1</sup> Each node in the hierarchy is therefore the cluster head of a (sub-)cluster.

to the root a JOIN message containing this distance. Based on this distance, the root determines the zone of the joining node. Here, two cases are possible:

1. The joining node is the first node joining in the corresponding zone.
2. Other nodes from the same zone have already joined.

In the former case, the root records the presence of the joining node in the corresponding zone and sends the node a NEW\_CLUSTER\_ACK message, indicating that the joining node has found its place in the hierarchy (this finishes the algorithm for the joining node). The joining node is now the cluster head of one of the sub-clusters dividing the cluster headed by the root (albeit a “singleton-cluster” for the time being).

In the latter case, the root sends to the joining node, in a TRY message, the list of the cluster heads in the same zone as the joining node, along with the radius associated with this zone. The joining node then measures its distance to each of the nodes in the list. Again, we consider two cases:

1. The distance of the joining node to at least one of the cluster heads in the list is smaller than the given radius. The clusters headed by these cluster heads are called *attracting clusters*.
2. The distance of the joining node to all the cluster heads in the list is greater than the given radius.

In the former case, the joining node chooses the closest attracting cluster and joins it: that is, the algorithm starts again with the corresponding cluster head acting as the root. In this case, the joining node is said to “go down one layer” (as the cluster it is heading will potentially be part of the partition of the attracting cluster) and it is important to note that the root does not record the presence of the joining node. In essence, from the root’s point of view, the members of the attracting cluster are “collapsed” into the attracting cluster head, as this cluster head is the only node in the attracting cluster remembered by the root.

In the latter case, the joining node creates a new sub-cluster by sending a NEW\_CLUSTER message to the root (including its distance to the root). The root then keeps a record of the new cluster head (i.e. the joining node) and of its zone and replies with a NEW\_CLUSTER\_ACK message which finishes the algorithm for this joining node.

### 3 Scalability Considerations

From the previous section, it should be clear that the state overhead imposed on each node in the hierarchy is proportional to the number of zones needed for that node to “span” its cluster, times the number of clusters per zone. This number of clusters per zone also influences the scalability of the join procedure, as any joining node must measure its distance to all the cluster heads at the same zone, for all traversed layers. Also, the further away from the central node a zone is, the more nodes – and thus the more clusters – such a zone potentially contains. These observations favour the use of large clusters, within few zones.

On the other hand, large clusters tend to create many layers (as they can contain large sub-clusters which, in turn, will have to be divided), which has a negative impact on the latency of the join procedure.

In order to accommodate these conflicting requirements, we propose to define zones based on RoundTrip Times (RTTs) measurements, and whose sizes follow an “exponential distribution” (see figure 2):

$$\text{zone}_0 : 0 < \text{dist} \leq 1 \quad (1)$$

$$\text{zone}_i : (1 + \Delta)^{i-1} < \text{dist} \leq (1 + \Delta)^i, \text{ with } \Delta > 0 \quad (2)$$

This, in turn, allows us to define the size (i.e. radius) of the clusters at zone<sub>i</sub> as:

$$r_i = \frac{(1 + \Delta)^i - (1 + \Delta)^{i-1}}{2} \quad (3)$$

The parameter  $\Delta$  in the formulae could be either fixed or varied according to which layer the cluster, headed by the corresponding node, belongs to. Other size “distributions” for both zones and radii are of course possible, but the ones we propose prevent an explosion of the number of clusters in far zones while keeping the number of zones down and retaining the desirable property that “detail” (i.e. “fine grain positioning”) matters only for nodes close to a cluster head.

## 4 Discussion and Conclusions

In this paper, we have proposed a method to build a hierarchy of nodes, based on the notion of proximity, in a distributed and scalable way. The hierarchy is built through a series of “local” decisions involving only a small subset of the hierarchy’s population for each decision. This, coupled with an innovative adaptive cluster size distribution approach, yields a simple, yet powerful, approach to building overlay, application-level structures without relying on any special support from network routers.

The hierarchy thus built is loopless and spans all the nodes in it. Our scalable adaptive hierarchical clustering algorithm can therefore be seen as a new member in the category of application-level multicast tree building methods (e.g. [1][5][7][6]). The overlay application-level multicasting trees built with our scalable adaptive hierarchical clustering are unconstrained, meaning that nodes in the tree cannot explicitly control their number of children. This may not be a problem for overlay trees built for control purposes [8] but could yield a significant penalty for trees built for data distribution. However, the method presented in this paper can still be very useful in the context of application-level multicast data distribution.

Indeed, a constrained application-level multicast overlay tree can be built by having each cluster head and its sub-clusters (i.e. its members populating the next layer in the hierarchy) run any algorithm that builds a constrained overlay tree [1][7][6]. With this approach, each node in the cluster hierarchy would be a

member of the overlay tree rooted at its parent cluster, as well as the root of the overlay tree spanning its own cluster. This would allow the building of very large constrained overlay trees.

Another application of the scalable adaptive hierarchical clustering presented in this paper is resource discovery. Indeed, a permanent hierarchy of resources could be built, rooted on a well known node, and “searched” by clients with a modified join procedure which does not declare the creation of a new (sub-) cluster when it finishes (see section 2.2). This could even substitute *expanding ring searches* in asymmetric network multicast circumstances or when network multicast is unavailable.

In future work, we will investigate the performance of the proposed algorithm under dynamic conditions (e.g. dynamic group membership, failures, etc.)

## References

1. Y-H. Chu, S. Rao, and H. Zhang. A Case for End System Multicast. In *ACM SIGMETRICS*, pages 1–12, Santa Clare, CA, USA, June 2000.
2. C. Diot, B. Levine, B. Lyles, H. Kassem, and D. Balensiefen. Deployment Issues for the IP Multicast Service and Architecture. *IEEE Network*, 14(1):78–88, Jan/Feb 2000.
3. B. Fenner, M. Handley, H. Holbrook, and I. Kouvelas. Protocol Independent Multicast - Sparse Mode: Protocol Specification (revised). Internet Draft draft-ietf-pim-sm-v2-new-02, IETF, Mar 2001. Work in Progress.
4. H. Holbrook and D. Cheriton. IP Multicast Channels: EXPRESS Support for Large-scale Single-source Applications. *ACM Comp. Comm. Reviews*, 29(4):65–78, Oct 1999.
5. J. Jannotti, D. Gifford, K. Johnson, F. Kaashoek, and J. O’Toole. Overcast: Reliable Multicasting with an Overlay Network. In *USENIX OSDI*, San Diego, CA, USA, Oct 2000.
6. L. Mathy, R. Canonico, and D. Hutchison. An Overlay Tree Building Control Protocol. In *Proc. of Intl. workshop on Networked Group Communication (NGC)*, pages 76–87, Nov 2001.
7. D. Pendarakis, S. Shi, D. Verma, and M. Waldvogel. ALMI: an Application Level Multicast Infrastructure. In *3rd USENIX Symposium on Internet Technologies*, San Francisco, CA, USA, Mar 2001.
8. B. Whetten and G. Taskale. An Overview of Reliable Multicast Transport Protocol II. *IEEE Network*, 14(1):37–47, Jan 2000.