

Measuring and Modelling the Group Membership in the Internet

Jun-Hong Cui¹, Michalis Faloutsos³, Dario Maggiorini⁴, Mario Gerla², Khaled Bousetta²

Emails: jcui@cse.uconn.edu, michalis@cs.ucr.edu, dario@dico.unimi.it, gerla@cs.ucla.edu, boukha@cs.ucla.edu

¹ Computer Science & Engineering Department, University of Connecticut, Storrs, CT 06029

² Computer Science Department, University of California, Los Angeles, CA 90095

³ Computer Science & Engineering, University of California, Riverside, CA 92521

⁴ Computer Science Department, University of Milan, via Comelico 39, I-20135, Milano, Italy

ABSTRACT

In this paper, we measure and model the distribution of multicast group members. Multicast research has traditionally been plagued by a lack of real data and an absence of a systematic simulation methodology. Although temporal group properties have received some attention, the location of group members has not been measured and modelled. However, the placement of members can have significant impact on the design and evaluation of multicast schemes and protocols as shown in previous studies. In our work, we identify properties of members that reflect their spatial clustering and the correlation among them (such as participation probability, and pairwise correlation). Then, we obtain values for these properties by monitoring the membership of network games and large audio-video broadcasts from IETF and NASA. Finally, we provide a comprehensive model that can generate realistic groups. We evaluate our model against the measured data with excellent results. A realistic group membership model can help us improve the effectiveness of simulations and guide the design of group-communication protocols.

Categories and Subject Descriptors

C.2.2 [Computer-Communication Networks]: Network Protocols—Applications

General Terms

Algorithms, Measurement, Performance, Experimentation

Keywords

Group Membership, Member Clustering, Skewed Distribution, Pairwise Correlation, Maximum Entropy

1. INTRODUCTION

Where should the members be located in a multicast simulation? This is the question that lies in the heart of this work. Multicast

research can greatly benefit from realistic models and a systematic evaluation methodology ([7] [25] [10] [21] [8] [23] [22] [26] [30] [14]). Despite the significant breakthroughs in modelling the traffic and the topology of the Internet, there has been little progress in multicast modelling. As a result, the design and evaluation of multicast protocols is based on commonly accepted but often unproven assumptions. For example, the majority of simulation studies assumes that the users are uniformly distributed in the network. In this paper, we challenge this assumption and study the spatial properties of group members, such as clustering and correlation.

A realistic and systematic membership model can have significant impact on the design and development of multicast protocols. Spatial information can help us address the scalability issues, which has always been a major concern in multicasting. Similarly, reliable multicast protocols need spatial information in order to fine-tune their performance or even evaluate their viability. Furthermore, spatial properties of a group with common interest members transcend the scope of IP multicast. Group communications is an undeniable necessity independently of the specifics of the technology that is used to support it. For example, web caching or application level multicast protocols can and should consider the member locality.

Only recently, properties of group membership have received some attention, but the spatial properties have not been adequately measured and modelled. Several studies show the importance of the spatial distribution of members [30] [14] [21]. However, there does not exist a generative model for such a distribution, which is partly due to unavailability of real data. In more detail, there have been several studies on the temporal group properties [4] [15]. In addition, several studies examine the scaling properties of multicast trees [10] [21] [8] [9] and the aggregatability of multicast state [22] [26] [30] [14]. Philips et al. [21] conclude that the affinity and disaffinity of members can affect the size of the multicast tree significantly. Thaler et al. [26] and Fei et al. [14] observe that the location of members has significant influence on the performance of their state reduction schemes.

In this paper, we study the distribution of group members focusing on their clustering and correlation. A distinguishing point of our work is that we use extensive measurements to understand the real distributions and develop a powerful model to generate realistic distributions. Our contributions can be grouped into two main thrusts.

I. Real data analysis. We measure and analyze the membership of net games and large audio-video broadcasts from IETF and NASA (over the MBONE). We quantify properties of the membership focusing on: a) the clustering, b) the distribution of the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IMC'03, October 27–29, 2003, Miami Beach, Florida, USA.

Copyright 2003 ACM 1-58113-773-7/03/0010 ...\$5.00.

participation, and c) the distribution of the *pairwise* correlation of members or clusters in a group. We observe that the MBONE multicast and gaming groups exhibit differences, which suggests the need for a flexible model to capture both. In our clustering analysis, we use the seminal approach of network-aware clustering [16]. More specifically, we make the following observations.

1. *MBONE multicast members*: The group members are highly clustered and the clusters exhibit strong pairwise correlations in their participation.
2. *Net game members*: The clustering is much less pronounced and there does not seem to be a strong correlation between users. Interestingly, we observe a very strong daily periodicity.

II. GEM: A model for generating realistic groups. We develop **GE**neralized **M**embership model (GEM) that can generate realistic member distributions. These distributions are given as input parameters to the model, enabling users to match the desired distribution. The main innovation of the model is the capability to match pairwise participation probabilities. To achieve this, we use the Maximum Entropy method [31], which, in an under-defined system, chooses the solution with maximum “randomness” or entropy. As a result, GEM can simulate the following membership behavior:

1. Uniform distribution, which is the typical but not always realistic distribution.
2. Skewed participation distribution without pairwise correlations.
3. Skewed participation distribution with pairwise correlations.

We validate our model with very positive results. We are able to generate groups whose statistical behavior matches very well the real distributions.

Modelling location of users with common interests. The analysis and the framework presented here can be of interest even outside the multicast community. Applications with multiple recipients such as web caching and streaming multimedia are also interested in the location of users ([6] [29]). We provide our data and our model to the community with the hope that it can be part of a realistic and systematic evaluation methodology for this kind of research ([1]).

The rest of this paper is organized as follows. Section 2 gives some background on multicast group modelling. Section 3 lists the spatial properties of group members. Section 4 quantifies the spatial group properties using real data from the MBONE and net games. Section 5 describes our powerful group membership model. In Section 6, we validate the capabilities of our model. Finally, we conclude our work in section 7.

2. BACKGROUND

In this section, we give some background on multicast group modelling and related efforts. The properties of multicast group behavior can be classified into two categories: spatial and temporal properties. Spatial properties consider the distribution of multicast group members in the network. Temporal properties concentrate on the distribution of inter-arrival time and life time of group members, in other words, the group member dynamics. In the following, we give an overview of the related work on the modelling of multicast group behavior.

The majority of multicast research assumes simplifying assumptions on the distribution of members in the network. Protocol developers assume almost always that users are uniformly distributed in the network (such as [27], [28], [5], [17], [13], and [10], etc.). This is partly due to the unavailability of real data. On the other hand, it is interesting to observe that skewed distributions have been observed in multiple aspects of communication networks from traffic behavior [18] [20] to preferences for content [11] and peer-to-peer networks [19].

There have been some studies on the temporal group properties, such as [4] and [15]. [4] measured and studied the member arrival interval and membership duration for MBONE. It also showed that, for multicast sessions on MBONE, an exponential function works well for the member inter-arrival time of all type of sessions, while for membership duration time, an exponential function works well for short sessions, but a Zipf [32] distribution works well for longer sessions. [15] conducted a follow-on study for net games. The authors found that player duration time fits an exponential distribution, while inter-arrival time fits a heavy-tailed distribution for net game sessions.

Several studies examining the scaling properties of multicast trees ([10] [21] [8] [9] [23]) and the aggregatability of multicast state ([22] [26] [30] [14]) show that the spatial properties do matter in multicast research. In their seminal work, Chuang and Sirbu [10] discovered that the scaling of the tree cost follows power law with respect to the group size, assuming that group members are uniformly distributed throughout the network. Philips et al. gave an explanation of the Chuang and Sirbu scaling law in [21]. They also considered member affinity¹, and concluded that, for a fixed number of members, affinity can significantly affect the size of the delivery tree. These two works mainly concentrate on multicast efficiency (the gain of multicast vs unicast). Besides defining a metric to measure multicast efficiency, Chalmers and Almeroth ([8] [9]) also examined the shape of the multicast trees through measurements from MBONE, basically focusing on the the distribution and frequency of the degree of in-tree nodes, the depth of receivers, and the node class distribution. In this work, Chalmers and Almeroth also indicate that the multicast efficiency can be affected by the member clustering.

The distribution of the group members affects our ability to aggregate the multicast state significantly. State aggregation has been the goal of several research efforts ([22], [26], and [14]). These papers proposed different state reduction schemes, and showed that group spatial properties, such as clustering of members, correlation between members, affects the performance of their approaches. In [30], Wong et al. did a comprehensive analysis of multicast state scalability considering network topology, group density, clustering/affinity of members and inter-group correlation. They conclude that application-driven membership has significant impact on multicast state distribution and concentration.

3. CHARACTERISTICS OF THE GROUP MEMBERSHIP

In this section, we identify and define several properties of group membership, which we quantify through measurements in the next section. For simplicity, we refer to the hosts or routers in the Internet as “nodes” or “network nodes”.

1. **Member Clustering**: Clustering captures the proximity of the group members. We are interested in the proximity from

¹Member affinity means the members are likely to cluster together, while disaffinity means that they tend to spread out.

a networking point of view, and we use the network-aware clustering method [16] in our measurement.

Earlier studies proposed models to capture the clustering of group members ([26], [30]). However, these studies do not provide measurements of the clustering in the Internet.

Note that the metrics we present below can refer to a node or a cluster. We will use the term “cluster”, since a node is a cluster of size 1. In addition, we focus on clusters in our analysis.

2. **Group Participation Probability:** Different clusters have different probabilities in participating in multicast groups: some clusters are more likely to be part of a group. The uniform distribution of participation is a special case where all clusters have the same probability.

Multiple Group Participation: If we have many groups, we define the participation probability of a cluster as the ratio of groups that the cluster joins.

Time-based Participation: For a single but long-lived group, the participation probability can be defined as the percentage of time that a cluster is part of the group. We find this definition particularly attractive, since our data is often limited in the number of groups. It should be noted that, in our analysis, we use this definition.

Fei et al. [14] proposed a node-weighted model to incorporate the difference among network nodes, where each node is assigned a weight representing the probability for that node to be in a group.

3. **Pairwise Correlation in Group Participation:** This metric captures the joint probability that two clusters are members of a group. The intuition is that common-interest or related users (e.g. friends) will probably share more than one groups. More specifically, we quantify the pairwise correlation between two clusters as follows. Given two clusters C_i and C_j , we denote the participation probabilities of cluster C_i and C_j as p_i and p_j respectively, and the joint participation probability of C_i and C_j is denoted by $p_{i,j}$. The correlation coefficient between C_i and C_j , $coef(i, j)$, is the normalized covariance between C_i and C_j ([24]):

$$coef(i, j) = \frac{(p_{i,j} - p_i \times p_j)}{\sqrt{p_i \times (1 - p_i)} \times \sqrt{p_j \times (1 - p_j)}}. \quad (1)$$

Multiple Group Pairwise Correlation: In the presence of many groups, we can use the multiple group participation probability to compute pairwise correlation.

Time-based Pairwise Correlation: In this work, we measured and analyzed single but long lived sessions (from MBONE and net games). Thus we can use the time-based participation probability that we defined above to compute time-based pairwise correlation.

In the literature, there has been some effort to model the pairwise correlation. In [26], a two-dimensional array of randomly allocated correlation probabilities is used. In [30], the authors simulated the correlation implicitly by encouraging the members of sets of nodes to join the same group, once one of the nodes of the set has joined.

We did not find any previous studies which use real data to verify and quantify the spatial properties. In addition, no previous effort

has provided a comprehensive model for all of the above properties of group membership, as we do here.

4. MEMBERSHIP FEATURES MEASURED FROM MBONE AND NET GAMES

In this section, we measure the properties of multicast group membership in real applications. First, we use data from NASA and IETF broadcasts over the MBONE, which are single-source large-scale application. Second, we measure the membership at net games, which are multiple-source interactive application. The MBONE is an overlay network on the Internet, and it has served as a testbed for multicast researchers since 1992. Net games is one of the most popular multiple-source applications. Though most of net games are implemented using multiple unicasts, we are interested in the membership behavior (or spatial group properties), which is independent of the underlying implementation.

4.1 Measurement Methodology

MBONE. We use data sets provided by Chalmers and Almeroth from University of Santa Barbara ([8] [9]). The data sets are divided into two groups: real data sets and cumulative data sets, which are summarized in Table 1 and Table 2 separately.

The real data sets include IETF43-A, IETF43-V and NASA, and the cumulative data sets include UCSB-2000, UCSB-2001, Gatech-2001 and UOregon-2001. For the details of measurement on MBONE, please see references [8] and [9]. One thing deserving more description is the generation of cumulative data sets: multicast paths are traced using a number of sources (UCSB, Georgia Tech, and Univ. of Oregon) for a series of 22,000 IP addresses that were known to have participated in multicast groups over a two years period, June’97-June’99. In these data sets, although most of the traces were collected recently and reflect the latest multicast infrastructure, the group members represent a relatively random sample taken from the older MBONE. Due to the limited number of real data sets, we use cumulative data sets to get an intuition of how the size of groups affects the property of member clustering.

Net Games. For net games, we use the QStat tool [3] to collect data. QStat is a program designed to poll and display the status of game servers. Some game servers offer a querying mechanism, which can retrieve some specific information, such as the number of players, players’ nicknames, IP addresses, and scores, and the time that each player has been connected, etc. To analyze clustering of members, we need members’ IP addresses (the reason will be clarified in Section 4.2). Not all games, however, provide players’ IP addresses. Quake I is one of the fewer games that allow this. Thus, we choose Quake I though it is a little bit old game. Using QStat, we measured 70 Quake game servers (obtained from a master server) for five days (across a weekend), and the servers are polled every one minute. We select the 10 most popular servers (5 of them providing IP addresses of players) for our analysis, and the selected game servers are illustrated in Table 3.

From each data set of MBONE and net games (note that, one data set corresponds to one multicast session or group), we sample the group membership at regular interval (1 minute). Each sample of group membership is composed of members with IP address or player ID for some net games data sets (in which IP address of players are not provided). To give some intuition of the data sets, we plot some examples chosen from real MBONE data and net games data in Fig. 1, Fig. 2, Fig. 3, and Fig. 4. In all these figures, the X-axis is the time, and the Y-axis represents the number of members (receivers for MBONE or players for net games). We can see clearly how the number of members changes with the time.

Table 1: MBONE Real Group Data Sets

<i>Name</i>	<i>Description</i>	<i>Trace Period</i>	<i>Receivers</i>		
			<i>Total</i>	<i>Maximum</i>	<i>Average</i>
IETF43-A	43rd IETF Meeting Audio	Dec. 7-11, 1998	107	93	58.72
IETF43-V	43rd IETF Meeting Video	Dec. 7-11, 1998	129	90	48.59
NASA	NASA Shuttle Launch	Feb. 14-21, 1999	62	62	40.33

Table 2: MBONE cumulative Group Data Sets

<i>Name</i>	<i>Description</i>	<i>Trace Period</i>	<i>Receivers</i>		
			<i>Total</i>	<i>Maximum</i>	<i>Average</i>
SYNTH-1	UC Santa Barbara	Jan. 6-10, 2000	1,871	1653	805.94
SYNTH-2	Georgia Tech	Jul. 12-25, 2001	1,497	1497	958.17
SYNTH-3	University of Oregon	Dec. 18-19, 2001	1,019	1019	492.45
SYNTH-4	UC Santa Barbara	Dec. 19-22, 2001	1,018	1018	474.35

Table 3: Netgames Group DataSets

<i>Name</i>	<i>Game Server</i>	<i>Measurement Period</i>	<i>Players</i>		
			<i>Total</i>	<i>Maximum</i>	<i>Average</i>
QS-1	quake.dircon.co.uk	May 14-18, 2002	352	8	1.71
QS-2	sense-sea.oz.net	May 14-18, 2002	265	11	1.89
QS-3	195.147.246.71	May 14-18, 2002	234	11	1.72
QS-4	ut2003.kos.net	May 14-18, 2002	158	10	2.22
QS-5	zoologi38.zoologi.su.se	May 14-18, 2002	391	11	2.34
QS-6	200.230.198.53: 26004	May 14-18, 2002	1198	10	3.31
QS-7	frodo.trinicom.com	May 14-18, 2002	437	16	4.04
QS-8	bridge.widomaker.com	May 14-18, 2002	417	13	3.67
QS-9	200.230.198.53: 26001	May 14-18, 2002	1298	8	3.50
QS-10	209.48.106.170	May 14-18, 2002	604	15	8.12

MBONE multicast: decelerating increase and “black-out” phases. In Fig. 1 and Fig. 2, we see that the IETF broadcast increases close to monotonically but with decreasing rate of increase. We also notice some short periods (there is also a big period for IETF Video) in which the number of members drop suddenly and then rise again. One possible explanation is the network instability: either the tree was torn down and rebuilt or the measurements got lost. Another possible reason is that these might correspond to breaks of the IETF meeting, such as lunch time. Fig. 3 shows the sampled data sets for NASA broadcast. We see that NASA broadcast has smaller number of drop periods than IETF broadcast. One reason to explain this is that, unlike IETF meeting, NASA shuttle lunch is a more continuous event. The big drop period can be explained by some break of network connection or some unexpected and uninteresting event.

Net games: membership is strongly periodic. Fig. 4 shows very interesting behavior of net games (Quake) players: in each day, there is a big spike in user participation. Moreover, there are more players during the weekend (May 17th and May 18th). This periodicity is natural given the nature of the activity: For a game server, due to the delay constraints of gaming, most of the players come from areas within some range (say, in several hops). Thus the players are more likely active in some relatively fixed period of time in a day. For example, in Fig. 4, we see that late night is a very active period for game players in this server.




Figure 1: The data set sampled from IETF Meeting Audio (IETF43-A).

In the rest of this section, we examine membership properties of the above data.

4.2 Member Clustering

To model member clustering, we employ network-aware clustering. Intuitively, two members should be in the same cluster if they are close in terms of network routing. In the Internet, this kind of grouping can be done based on IP addresses. We adopt the method in [16] to identify member clusters using network prefixes, based on information available from BGP routing snapshots (we use the BGP dump tables obtained from [2]). This way, clustered nodes are likely to be under common administrative control. For details, please see [16].

We briefly outline network-aware clustering for completeness. We first extract the network prefixes/netmasks from BGP dump tables and the IP addresses of members from group membership samples, then we classify all the member IP addresses that have the same longest-match prefix into one cluster, which is identi-

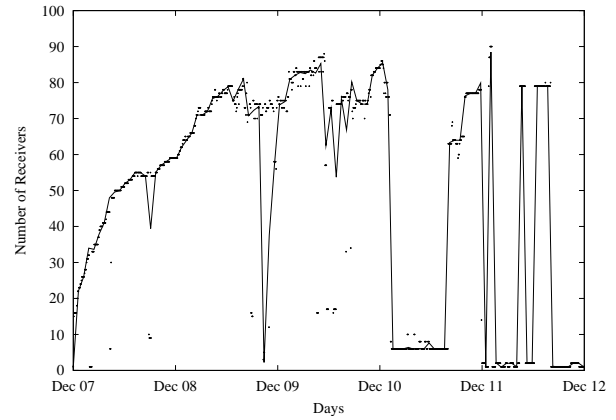


Figure 2: The data set sampled from IETF Meeting Video (IETF43-V).




Figure 3: The data set sampled from NASA Shuttle Lunch (NASA).




Figure 4: The data set sampled from net game server 1 (QS-7).

fied by the shared prefix. For example, suppose we want to cluster the IP addresses 216.123.0.1, 216.123.1.5, 216.123.16.59, and 216.123.51.87. In the routing table, we find the longest-match prefixes are 216.123.0.0/19, 216.123.0.0/19, 216.123.0.0/19, and 216.123.48.0/21 respectively. Then we can classify the first three

IP addresses into a cluster identified by prefix/netmask 216.123.0.0/19 and the last one into another cluster identified by 216.123.48.0/21. It should be noted that other clustering methods, such as network topology based approach, are possible. But network-aware clustering is an easy and effective way for us to do clustering considering the data we have achieved. In later sections, we will show that our analysis and model are not constrained by the clustering method.

For each data set, we want to see the number of group members per cluster. Here, we refer to the number of group members in a cluster as the **size of the cluster** (or cluster size). For all the group membership samples, we examine the Cumulative Distribution Function (CDF) of the cluster size. The results for data sets from MBONE and net games are shown in Fig. 5. Therefore, for a given cluster size in the X-axis, we see how many clusters have at most that size.

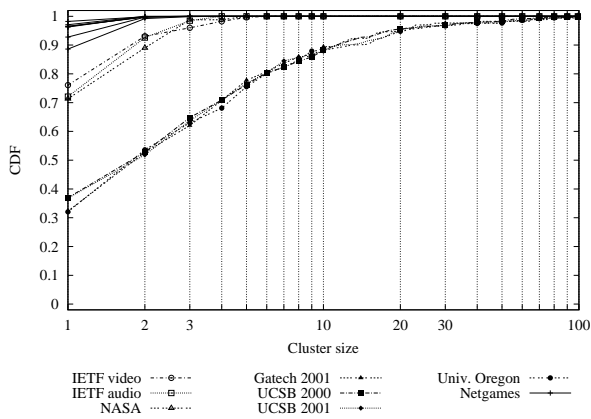


Figure 5: CDF of cluster size for data sets from MBONE and net games. The upper set of curves are for net game data sets (with 5 game servers providing IP addresses of players), the middle set of curves are for MBONE real data sets (IETF43-A, IETF43-A, and NASA), and the lower set of curves are for MBONE cumulative data sets.

Group members form clusters with skewed size distribution. Group members are significantly clustered. In Fig. 5, we can see three different groups of curves: the upper group for net game data sets (with 5 game servers providing IP addresses of players), the middle group for MBONE real data sets (IETF43-A, IETF43-A, and NASA), and the lower group for MBONE cumulative data sets. In each group, the data sets have similar member clustering property: for example, for MBONE real data sets, more than 20% clusters have 2 or more group members; for MBONE cumulative data sets (UCSB-2000, UCSB-2001, Gatech-2001 and UOregon-2001), they also have similar features: more than 60% clusters have 2 or more group members; while for net games, the corresponding group of CDF curves do not show significant member clustering: about 90% “clusters” have size 1.

Cluster size distribution is mainly affected by the group size. We observe that it is primarily the size of the group that affects the range of the distribution. However, the cluster size distribution is similar qualitatively in all three groups of data sets. When comparing the groups of curves, we can conclude that the MBONE cumulative data sets have more significant clustering feature than MBONE real data sets. We attribute this to the larger size of the group. The average number of members for each cumulative data set (from 500 to 1000) is much higher than that for real data sets (around 50). The bigger the group is, the more members tend to

be in one cluster and the more significant of the member clustering feature is. As for net games, the feature of member clustering is even less significant: most “clusters” have only one member, which means that Quake players are more likely scattered over the network. This observation suggests that probably the gain from some multicast or intelligent caching schemes may not provide significant benefits in this case.

The absence of clustering in the net games can be attributed to many factors. One observation is that the maximum number of players (16 in Quake) is controlled by the game servers because of management issues. Thus, the possibility for the members to fall in one cluster becomes smaller. It would be very interesting to examine a net game with a larger user participation, but we were not able to get such data. Another possible explanation may be that the game players are not necessarily from a similar area of the network potentially. This suggests that gaming community is scattered, or alternatively, that net games bring together people from significantly different places.

The practical implications of member clustering. Understanding the clustering properties can help us develop efficient protocols to improve the scalability and the performance of applications. The member clustering captures the proximity of the members in the network especially with the use of network-aware clustering. For example, in a well clustered group, we can potentially develop hierarchical protocols that can exploit the spatial distribution of the members like hierarchical multicasting.

4.3 Group Participation Probability

We find that the **participation probability is non-uniform across clusters or nodes**. This strongly suggests that the uniform distribution used so far for most research is not realistic. In the analysis below, we study the distribution across clusters or nodes that participate at least once in a multicast group. Clearly, there are clusters or nodes in the network that do not appear in any group. In fact, we expect that these clusters (or nodes) are probably large in number, which reinforces the observation that not all clusters (or nodes) are created equal regarding multicast participation.

Note that, given the limited number of groups, we measured the time-based participation probability as defined in the previous section.

We give the CDF of the participation probability of clusters or nodes for MBONE and net games in Fig. 6 and Fig. 7 respectively. Given a probability in the X-axis, we can see how many clusters or nodes have at most that probability to participate in the multicast group.

MBONE: the cluster participation probability is non-uniform. In Fig. 6, we see that the MBONE clusters are not equal in participating in a group. If the clusters had the same probability p of participating in a group, the CDF of the participation probability would appear as a vertical line at the exact p on the X-axis. The current plot of the CDF shows a roughly linear increase with close to 45 degrees slope. This suggests that we have a wide range of participation probabilities: for any value on the X-axis, we can find a cluster with such a participation probability.

Net Games: the node participation probability roughly follows uniform distribution. For net games, since the member clustering feature is not significant at all (about 90% “clusters” have size 1), we simply analyze the node participation probability which is approximated by the frequency of nodes joining the net game session.

Fig. 7 plots the CDF of the participation probability for net games. We observe that the plot is qualitatively different from the MBONE distribution. For all the Quake servers we examined, more than




Figure 6: CDF of the participation probability of clusters for data sets from MBONE (only real data).




Figure 7: CDF of the participation probability of nodes for data sets from net games (including 10 data sets).

95% of nodes only have participation probability less than 0.1%. We can say the uniformly random group membership model is sufficiently realistic for net games (Quake). Interestingly, this suggests that user participation is approximately equally distributed, and we do not find players that would always join a game server. This is consistent with our intuition: a player joins and leaves the game at will and the player's behavior is not heavily affected by other factors as in applications on MBONE, such as IETF43 and NASA, where the delivery content does influence the behavior of members significantly.

4.4 Pairwise Correlation in Group Participation

We study the time-based pairwise correlation of clusters or nodes. For each group, we sample it and compute the participation probability for all clusters or nodes and the pairwise participation probability between any two clusters or nodes. Then, using Equation 1, we calculate the matrix of correlation coefficients. For net games, we analyze the correlation between nodes directly, since most clusters are trivially nodes.

We want to study the correlation of the clusters which is a matrix. A large matrix with probabilities as values is not easy to visualize. To overcome this, we plot the CDF of the correlation coefficient of clusters in Fig. 8 and Fig. 9 for both groups of data sets (MBONE

and net games). In these figures, given a correlation coefficient in the X-axis, we know how many pairs of clusters or nodes have at most that correlation coefficient in multicast participation.

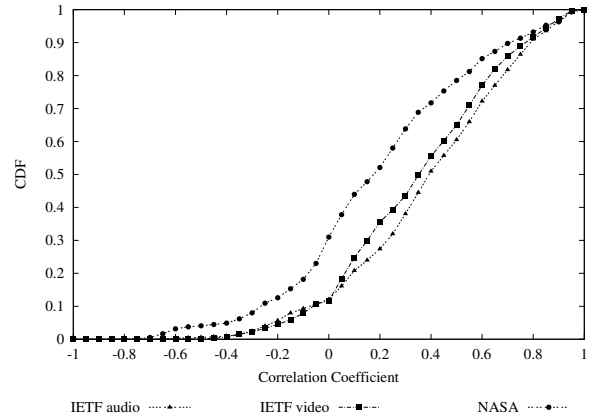


Figure 8: CDF of the correlation coefficient of clusters for data sets from MBONE (real data).




Figure 9: CDF of the correlation coefficient of nodes for data sets from net games (including 10 data sets).

The results for pairwise correlation are consistent with those for participation probability: significant correlation feature for MBONE applications, while very weak correlation phenomena for net games.

MBONE: clusters exhibit strong pairwise correlation. This further argues that the selection of users with fixed equal probability is not realistic for this kind of groups. Fig. 8 shows that most of the correlation coefficients between clusters are not 0 (in fact, for both IETF video and audio data sets, only 1.5% of cluster pairs have 0 correlation coefficient; and for NASA data set, the percentage of 0 correlation coefficient is 8%). This means that most of the clusters are not independent². Moreover, about 90% (for IETF data sets) or 70% (for the NASA data set) of correlation coefficients are greater than 0 (positive correlation). This can be explained by the fact that, in an IETF meeting or a NASA shuttle launch multicasting, many members have very similar interests in specific sessions and thus many clusters tend to be coupled together.

²It is easy to verify that, for any two variables which follow 0-1 distribution, if their coefficient is 0 then they are independent, and vice versa.

Net games: pairwise user participation exhibits weak correlation. In Fig. 9, we can see that some game servers have more significant correlation features than others. However, for all the servers, more than 55% (up to 80% for some servers) of node-pairs have correlation coefficient as 0. In other words, most pairs of nodes are independent. The explanation for the difference between MBONE applications and net games is very similar to the arguments for the participation probability distribution. The number of players in any single game is limited. In addition, repeated users do not seem to want to join the same server as some other particular player joins. Again, we can say that the simple uniform random membership model can describe the membership of net games and the absence of pairwise correlation.

Note that we did not analyze the correlation and participation for the cumulative data sets, since by their nature, they do not provide sufficient details to generate the required distributions.

4.5 Does Member Clustering Affect Skewed Distribution and Pairwise Correlation?

In the above, we measured and analyzed the participation probability and correlation of clusters in stead of nodes for MBONE application, and we observed skewed distribution and strong pairwise correlation. An interesting question would be: does member clustering affect the observed properties? Or in other words, are the new properties caused by our additional processing, i.e., member clustering, or they come by the nature of the applications? In this section, we show the analysis results of MBONE applications without member clustering. That is, we plot the CDF curves of participation probability and pairwise correlation of nodes (without clustering). And the results are showed in Fig. 10 and Fig. 11.




Figure 10: CDF of the participation probability of nodes for data sets from MBONE (only real data).

We can see that Fig. 10 and Fig. 11 have very similar curves to Fig. 6 and Fig. 8, though the values are slightly different due to the absence of member clustering. This confirms that the the properties of skewed distribution and strong pairwise correlation come with the MBONE applications instead of the processing of clustering (which does affect the values though).

5. GEM: A GROUP MEMBERSHIP MODEL

From our measurement and analysis, we conclude that real membership distribution does not follow the simple uniform random distribution. We propose a comprehensive group membership model, called **GEM** (GEneralized Membership model) to generate mem-




Figure 11: CDF of the correlation coefficient of nodes for data sets from MBONE (real data).

bership distributions that conform to realistic distributions. In fact, the distribution is an input parameter for our model.

5.1 An Overview of GEM

GEM considers all the group membership properties we discussed earlier: a) member clustering, b) group participation probability, and c) pairwise correlation in group participation. GEM has the following inputs:

- A network topology
- Clustering method (that determines how to form clusters in the given topology)
- Target group behavior: the distribution of group participation probability, the pairwise correlation in group participation of clusters, and the distribution of member cluster size (i.e., the number of member nodes in a cluster)

GEM generates multicast groups whose members follow the given distributions and constraints. Fig. 12 is a block diagram of GEM, illustrating how GEM works at a high level.

GEM works in the following steps:

1. *Cluster creation:* GEM classify nodes into many disjoint clusters using the specified clustering algorithm. Leaving this as an input parameter gives GEM a lot of flexibility and the option to ignore clustering altogether as we discuss below.
2. *Membership distribution:* GEM creates groups and chooses their cluster members among the clusters in a randomized fashion. The selection of clusters follows the given distributions for clusters participation and pairwise correlation. Note that, in this paper, we also refer the clusters which are chosen for a group as “member clusters” of the group.
3. *Node selection:* In each chosen cluster, GEM random selects nodes based on member cluster size.

First, the model is not tied to a particular membership distribution. This makes the usefulness of the model extend beyond the accuracy of a set of measurements. At the same time, our measurements provide guidelines for the choice of realistic distributions. In the tool we develop based on GEM, the measured distributions will be provided as choices to the users ([1]).

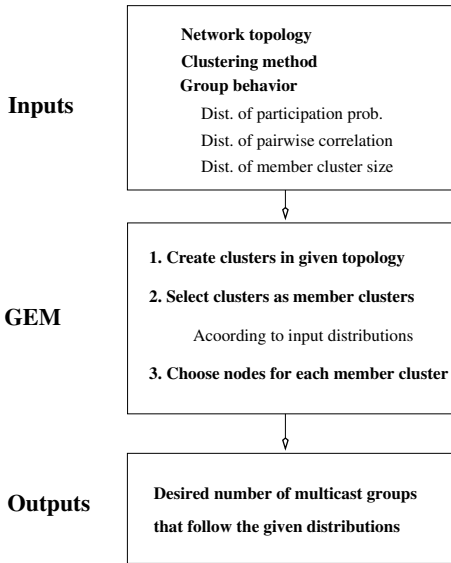


Figure 12: An illustration of GEM.

Cluster and node level membership. Note that as described above, GEM operates first at the cluster level and then node level. First, it generates groups treating a cluster as one entity. The reason for doing this is to simulate the network-level clustering that we have observed. After we identify the member clusters, then we assume that inside a cluster we have a number of active participants according to the measured distributions (that is member cluster size distribution).

5.2 Member Distribution Generation

The core of our model is the selection of the member clusters. The problem can be stated as follows: given a set of clusters, the group participation probability of each cluster, and the pairwise correlation between any two clusters, we want to generate sets of member clusters, which follow the given distributions. In other words, if we generate many multicast groups, the measured distributions should match the targets.

5.2.1 Problem Formulation

Let us start with the following definitions. We assume K clusters: $C_1, C_2, \dots, C_i, \dots, C_K$. Let us denote as p_i the participation probability of cluster C_i , that is, how often the C_i participates in a multicast group. For any two clusters C_i and C_j , there is a correlation coefficient $coef(i, j)$, where $1 \leq i, j \leq K$. Based on p_i, p_j , and $coef(i, j)$, we can easily compute the joint probability $p_{i,j}$ as shown in Equation 2 (derived from Equation 1).

$$p_{i,j} = coef(i, j) \times \sqrt{p_i \times (1 - p_i) \times p_j \times (1 - p_j)} + p_i \times p_j. \quad (2)$$

As a result, we get a symmetric joint probability matrix P_m where $P_m(i, j) = p_{i,j}$ when $i \neq j$ and $P_m(i, i) = p_i$.

A multicast group can be represented by a K -dimensional vector of binary values $x = (x_1, x_2, \dots, x_i, \dots, x_K)$, where $x_i = 1$ if and only if cluster C_i is a member cluster of the group (else $x_i = 0$).

Now, we can formalize the problem as follows. If we assume many groups³, we can model the **participation distribution of the**

³The presentation is easier when we talk about multiple groups, in other words, the multiple group participation. If we have one group we can talk about the time-based participation.

clusters by K random binary variables, $(X_1, X_2, \dots, X_i, \dots, X_K)$, where X_i represents the group participation of cluster C_i .

The generation of multicast groups reduces to generating vectors $x = (x_1, x_2, \dots, x_i, \dots, x_K)$, which follow the given distributions. Namely, we want to select vectors from the distribution $(X_1, X_2, \dots, X_i, \dots, X_K)$, which is defined by:

- $P(X_i = 1) = p_i, \forall i$, which means that X_i follows the given participation probability p_i
- $P(X_i = 1, X_j = 1) = p_{i,j}$, which means that for any two variables X_i and X_j , they have joint distribution $p_{i,j}$.

Note that the problem is in some sense under-defined. Complete knowledge of the distribution of $(X_1, X_2, \dots, X_i, \dots, X_K)$ would require us to know the probability of appearance for every of the $O(2^K)$ binary vectors. In other words, we would need $O(2^K)$ values to be able to generate the desired distribution. We only have partial information: our total input is $O(K + K^2)$. Intuitively, we need to make some assumptions to “fill” the missing information.

Assuming Maximum Entropy. For the missing constraints, we will assume that they have maximum entropy. Entropy is a measure of randomness of a system, and it is the “opposite” of order. In addition, nature tends to increase its entropy. A table with nicely stacked papers in alphabetical order has high order (low entropy). A wind from the window can shuffle the papers, which leads to high entropy. It is unlikely that a subsequent wind will restack and alphabetize the papers.

In our approach, we use entropy to replace the missing information. Given an unconstrained choice, we will choose according to the maximum Entropy (ME) [31]. This is a non-trivial but solved problem in statistical analysis [31]. Let us denote $p^*(x)$ the Maximum Entropy distribution. Intuitively, we can see this as a multidimensional problem with only a few constrained dimensions. The ME distribution $p^*(x)$ satisfies the constraints along the specified dimensions, and it is as unstructured as possible in the unconstrained dimensions. If we see entropy as lack of information, the Maximum Entropy distribution represents all the “known information” and nothing more than that. Our member cluster generation algorithm combines two “conflicting” forces: it maximizes the entropy (randomness), while it tries to match given distributions. Note that it is possible to use a distribution with “any” entropy. However, to compute any-entropy distribution is demanding. Moreover, using maximum entropy is more meaningful: if we do not know, we assume the distribution as random as possible.

5.2.2 Algorithms for member cluster generation

Our problem formulation describes the following three desired distributions in the order of increasing constraints.

1. Uniform distribution without correlation: all clusters have equal probability to join. This is the widely-used multicast membership model.
2. Non-uniform distribution without correlation: participation probability is higher for some clusters.
3. Non-uniform distribution with pairwise correlation: as above plus some pairs of nodes appear more often together.

Note that for the first two cases, it is not necessary to use a maximum entropy distribution, since there are no correlations among clusters. Straightforward algorithms can be used to generate member clusters as shown below.

Uniform distribution without correlation. In this case, $p_i = p$, and $p_{i,j} = p_i \times p_j = p^2$ (or $coef(i, j) = 0$) for any i and j , where

$1 \leq i, j \leq K$. Among the above three cases, this is the case with maximum entropy, because there are almost no constraints: member clusters are chosen uniformly among all clusters, and clusters are independent of each other. The member cluster generation algorithm is straightforward in this case, and it is described in Algorithm 1.

Algorithm 1 Member Cluster Generation (Case 1)

Require: For K variables, $X_1, X_2, \dots, X_i, \dots, X_K$, $P(X_i = 1) = p$, and $P(X_i = 1, X_j = 1) = p^2$ (or $\text{coef}(i, j) = 0$), where $1 \leq i, j \leq K$. Notes: X_i represents the group participation of cluster C_i with values of 0 (not join) or 1 (join)

Ensure: A K -dimension vector, $(x_1, x_2, \dots, x_i, \dots, x_K)$, which follows given distribution.

- 1: **for** $i = 1$ to K **do**
 - 2: generate a random number between 0 and 1, let it be u
 - 3: **if** $u < p$ **then**
 - 4: $x_i = 1$ (cluster C_i joins multicast group)
 - 5: **else**
 - 6: $x_i = 0$ (cluster C_i will not join multicast group)
 - 7: **end if**
 - 8: **end for**
-

Non-uniform distribution without correlation. In this case, $p_{i,j} = p_i \times p_j$ (or $\text{coef}(i, j) = 0$) for any i, j , where $1 \leq i, j \leq K$, while p_i is usually unequal between different clusters. Compared with case 1, this case needs to consider non-uniform distribution, that is, p_i for different clusters. However, all the clusters are still independent to each other. Thus, the member cluster generation algorithm is still straightforward. It is described in Algorithm 2.

Algorithm 2 Member cluster Generation (Case 2)

Require: For K variables, $X_1, X_2, \dots, X_i, \dots, X_K$, $P(X_i = 1) = p_i$, and $P(X_i = 1, X_j = 1) = p_i \times p_j$ (or $\text{coef}(i, j) = 0$), where $1 \leq i, j \leq K$. Notes: X_i represents the group participation of cluster C_i with values of 0 (not join) or 1 (join)

Ensure: A K -dimension vector, $(x_1, x_2, \dots, x_i, \dots, x_K)$, which follows given distribution.

- 1: **for** $i = 1$ to K **do**
 - 2: generate a random number between 0 and 1, let it be u
 - 3: **if** $u < p_i$ **then**
 - 4: $x_i = 1$ (cluster C_i joins multicast group)
 - 5: **else**
 - 6: $x_i = 0$ (cluster C_i will not join multicast group)
 - 7: **end if**
 - 8: **end for**
-

Non-uniform distribution with pairwise correlation. We need to consider pairwise correlation between any two clusters, which means that $p_{i,j} = p_i \times p_j$, for $1 \leq i, j \leq K$ does not hold necessarily (i.e. $\text{coef}(i, j) \neq 0$). In this case, we have to calculate the maximum entropy distribution $p^*(x)$ which is subject to the given constraints. Then, we use Gibbs Sampler ([24]) approach to sample it, i.e., to obtain instances of membership values, $(x_1, x_2, \dots, x_i, \dots, x_K)$.

Calculating the Maximum Entropy distribution. Given the constraints $P(X_i = 1) = p_i$, and $P(X_i = 1, X_j = 1) = p_{i,j}$, where $1 \leq i, j \leq K$, or in other words, given a probability matrix $P_m = [p_{i,j}]$, the maximum entropy $p^*(x)$ is the solution to the following problem:

$$p^*(x) = \arg \max \left\{ - \int p(x) \log p(x) dx \right\}, \quad (3)$$

subject to

$$\int x_i x_j p(x) dx = p_{i,j}, \text{ when } i \neq j, \quad (4)$$

and

$$\int x_i p(x) dx = p_i, \quad (5)$$

and

$$\int p(x) dx = 1. \quad (6)$$

By Lagrange multipliers (an optimization technique), the solution for $p^*(x)$ is:

$$\begin{aligned} p^*(x; \Lambda) &= p^*(x_1, x_2, \dots, x_K; \Lambda) \\ &= \frac{1}{Z(\Lambda)} \exp \left[- \sum_{i=1}^K \lambda_{i,i} x_i - \sum_{i=1, j=1, i \neq j}^{i=K, j=K} \lambda_{i,j} x_i x_j \right], \end{aligned} \quad (7)$$

where $\Lambda = [\lambda_{i,j}]$ is the Lagrange parameter, and

$$Z(\Lambda) = \sum_x \exp \left[- \sum_{i=1}^K \lambda_{i,i} x_i - \sum_{i=1, j=1, i \neq j}^{i=K, j=K} \lambda_{i,j} x_i x_j \right], \quad (8)$$

is the partition function that depends on Λ , and it has the following properties:

$$\frac{1}{Z} \frac{\partial Z}{\partial \lambda_{i,j}} = -p_{i,j}. \quad (9)$$

Unfortunately, a closed form solution for Λ is not available in general. Thus, we seek a numerical solution by solving the following equations iteratively.

$$\frac{d\lambda_{i,j}}{dt} = p_{i,j}^t - p_{i,j}, \quad (10)$$

where $p_{i,j}^t$ is the intermediate probability computed at step t , based on the intermediate distribution $p^t(x; \Lambda)$. For details of the derivation of the equations, refer to [12]. In Algorithm 3, we briefly describe the iterative procedure to construct the maximum entropy distribution $p^*(x)$, and then use Gibbs Sampler ([24]) to obtain instances of member clusters. The key method in the algorithm is basically a known statistical technique [31]. Understanding it and implementing it is a non-trivial task. The user of our GEM tool does not have to thoroughly understand the intricacies of this technique.

5.2.3 Discussions

ME Values in 3 Cases. In the above three cases, the ME value decreases from case 1 to case 3. For case 1, we do not know anything except the portion of clusters to join the multicast group. For case 2, we know a little bit more: different cluster has different group participation probability. While for case 3, besides the knowledge in case 2, we even know the pairwise correlation, which adds more constraints to the ME distribution, as leads to a smaller ME value compared to case 2.

Input Distribution Validation. It can be shown that the technique we use in case 3 will lead provably to the correct distribution, if such a solution exists. Note that in some cases, the given distribution may be self-contradicting. For example, cluster A and cluster B have high positive correlation, and so do cluster B and C. However, the correlation of C and A is negative. It may not be possible to satisfy such a constraint. Thus, GEM needs to check whether the input distributions are valid or not. In fact, given p_i , p_j , and $p_{i,j}$,

Algorithm 3 Member Cluster Generation (Case 3)

Require: For K variables, $X_1, X_2, \dots, X_i, \dots, X_K$, $P(X_i = 1) = p_i$, and $P(X_i = 1, X_j = 1) = p_{i,j}$, where $1 \leq i, j \leq K$. Notes: X_i represents the group participation of cluster C_i with values of 0 (not join) or 1 (join)

Ensure: A set of K -dimension vectors, $(x_1, x_2, \dots, x_i, \dots, x_K)$, which follows given distribution.

- 1: initialize Λ as Λ^0
 - 2: $E^0 = 0$ //initialize judgement term as 0
 - 3: $t = 0$ //initialize step counter
 - 4: **repeat**
 - 5: $t = t + 1$
 - 6: compute Z^t based on equation 8, then we know $p^t(x; \Lambda)$.
 - 7: use Gibbs Sampler ([24]) to sample $p^t(x; \Lambda)$ and compute $p_{i,j}^t$.
 - 8: $\lambda_{i,j}^t = \lambda_{i,j}^{(t-1)} + (p_{i,j}^t - p_{i,j})$ //update Λ
 - 9: compute $E^t = \exp(-(\sum_{i,j} (p_{i,j}^t - p_{i,j})^2))$
 - 10: **until** $|E^t - E^{t-1}| < \epsilon$ //the iteration reaches steady state
 - 11: compute Z^t based on equation 8, then we know $p^t(x; \Lambda)$.
 - 12: use Gibbs Sampler ([24]) to sample $p^t(x; \Lambda)$, and output the series of generated vectors x .
-

the only constraint we need to consider is that $p_i + p_j \leq 1 + p_{i,j}$ should hold (this can be easily deduced from the basic definitions of probability and joint probability).

Bounded Size Group Generation. Sometimes, we want to control the size of the groups in our simulations. We can impose constraints in our approaches to stop the creation of a group once a size is reached. However, this may affect the correlations and probabilities in case 3. A suggested method is to recompute the probabilities and correlations under the given condition, i.e., bounded group size. But this will introduce more computation. Thus, a future work direction is to investigate the possibility of an efficient heuristic algorithm design for bounded size group generation.

6. EXPERIMENTAL VALIDATION

In this section, through experiments, we validate that GEM can generate realistic groups with great success and the generated groups match very well the real data.

6.1 Experiment Design

Based on our measurement results and analysis, we see that GEM can model both MBONE applications and net games: to model net games, Algorithm 1 can be used, since uniform distribution without correlation approximately characterizes net games's spatial properties; while to model MBONE applications, such as IETF and NASA, GEM needs to employ Algorithm 3, because in these applications, the features of non-uniform distribution and correlation are very significant. Since Algorithm 1 is pretty straightforward, to validate GEM, in this section, we only consider MBONE applications (IETF and NASA), which will be modelled by GEM with Algorithm 3.

Based on our measurement results of MBONE applications, we give inputs, such as the distributions of group participation probability and pairwise correlation of clusters, and the distribution of member cluster size, to GEM, and GEM generates multicast groups using Algorithm 3. Note that, in our validation experiments, we only consider possible member clusters, which has participation probability greater than 0, as is the same as the methodology in our measurement. Then, we analyze the generated multicast groups similarly, and compute the distributions of group par-

ticipation probability and pairwise correlation of clusters, and the distribution of member cluster size. And finally, we compare the results from real measurement data and simulation data (generated by GEM).

6.2 Results and Analysis

We give the comparison results for data set IETF43-Video as an example. The curves for CDF of cluster size, participation probability, and correlation coefficient are shown in Fig. 13, Fig. 14, and Fig. 15 respectively. To show the limitation of the uniform random model, we also plot the curves for its corresponding simulation data in Fig. 14 and Fig. 15.




Figure 13: The comparison of cluster size distribution for IETF43-Video between measurement and modelling data.

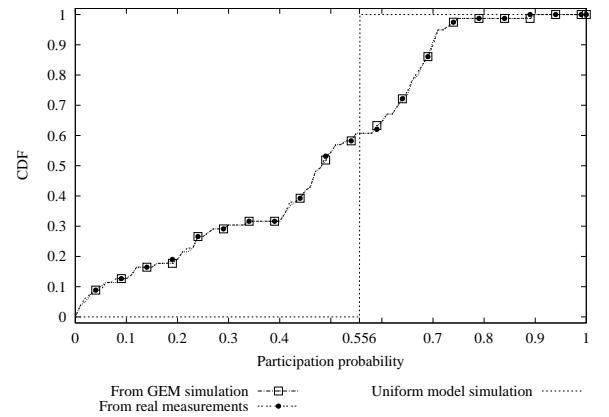


Figure 14: The comparison of group participation probability distribution of clusters for IETF43-Video between measurement and modelling data (GEM and Uniform model).

From Fig. 13 and Fig. 14, we can see that, for cluster size and participation probability, the two CDF curves for measurement and GEM simulation data are “perfectly” matched. For the correlation coefficient (shown in Fig. 15), the CDF curves for measurement and GEM simulation data are “nearly” matched. The difference is because CDF is an accumulative function, and the CDF of correlation coefficient involves much more variables (the number of variables is K^2 in total) than the CDF of participation probability (the number of variables is K). And simulation errors are propagated accumulatively. In addition, to facilitate the computation




Figure 15: The comparison of correlation coefficient distribution of clusters for IETF43-Video between measurement and modelling data (GEM and Uniform model).

in ME method, we approximate some very small (in terms of the absolute value) correlation coefficients as zero (which means the corresponding constraints could be ignored). But this will add experiment errors to the simulation. Another observation in Fig. 14 and Fig. 15 is that the uniform random model could not generate realistic multicast groups, as shows the limitation of it capabilities.

For other sessions, such as IETF Audio and NASA, we get similar results. Thus, based on our experimental results, we can conclude that our model can produce group members which follows given distribution successfully.

7. CONCLUSIONS AND FUTURE WORK

In this paper, we measured and modelled the spatial properties of multicast members. Our contributions can be summarized as follows:

1. We identify and quantify the spatial properties of members: member clustering, group participation probability, and pairwise correlation, in net games and large audio-video broadcasts from IETF and NASA.
2. We observe that MBONE multicast members are highly clustered and the clusters exhibit skewed distribution and strong pairwise correlations in their participation.
3. We find that net game members are not as clustered and there does not seem to be a strong correlation between users. The uniform random membership can roughly model net games (Quake) However, this could be partially due to the small scale of the games we were able to measure. Interestingly, we observe a very strong daily periodicity.
4. We develop GEM, a powerful group membership model, that can generate realistic member distribution. GEM combines of two conflicting forces: it maximizes the entropy of the distribution, while it tries to match given constraints.
5. We validate our model with great success: the generated groups match very well the real data.

We provide our data and our model as a contribution to the community with the hope that it can be part of a realistic and systematic multicast evaluation methodology ([1]).

Future Work. We would like to continue this work in two main directions. First, we want to study more applications with multiple recipients, and explore their spatial properties. We believe that different applications have different spatial properties, but we hope that we can well characterize existing applications, which will assist research and development of multicast and other group communication technologies. Second, we would like to integrate in our model temporal properties of group membership.

8. ACKNOWLEDGEMENT

We give our special thanks to Robert Chalmers and Kevin Almeroth from University of California, Santa Barbara. Without their data, we would not be able to analyze multicast groups from MBONE.

9. REFERENCES

- [1] Group Membership Measurement and Modelling Project Page at UCLA.
<http://www.cs.ucla.edu/NRL/hpi/MCModel/index.html>.
- [2] NLANR Global ISP interconnectivity by AS number.
<http://moat.nlanr.net/AS/>.
- [3] QStat. <http://www.qstat.org>.
- [4] K. Almeroth and M. Ammar. Collection and modeling of the join/leave behavior of multicast group members in the MBone. *Proceedings of High Performance Distributed Computing Focus Workshop (HPDC '96)*, Aug. 1996.
- [5] T. Billhartz and et al. Performance and resource cost comparisons for the CBT and PIM multicast routing protocols. *IEEE Journal on Selected Areas in Communications*, 1997.
- [6] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker. Web caching and zipf-like distributions: Evidence and implications. In *INFOCOM (1)*, pages 126–134, 1999.
- [7] R. Caceres, N. Duffield, J. Horowitz, D. Towsley, and T. Bu. Multicast-based inference of network-internal characteristics: Accuracy of packet loss estimation. *IEEE INFOCOM*, 1999.
- [8] R. Chalmers and K. Almeroth. Modeling the branching characteristics and efficiency gains of global multicast trees. *Proceedings of the IEEE INFOCOM' 2001*, Apr. 2001.
- [9] R. Chalmers and K. Almeroth. On the topology of multicast trees. *UCSB Technical Report*, Mar. 2002.
- [10] J. Chuang and M. Sibrú. Pricing multicast communications: A cost-based approach. *Proceedings of the Internet Society INET'98 Conference*, July 1998.
- [11] M. Crovella and A. Bestavros. Self-similarity in World Wide Web traffic, evidence and possible causes. *SIGMETRICS*, pages 160–169, 1996.
- [12] J.-H. Cui, D. Maggiorini, M. Faloutsos, M. Gerla, and K. Boussetta. Modelling the spatial properties of multicast group members in the internet. Technical report, UCLA CSD Technical Report, Feb. 2003.
- [13] M. Faloutsos, A. Banerjee, and R. Pankaj. QoS-MIC: Quality of Service sensitive Multicast Internet protocol. *ACM SIGCOMM'98, Vancouver, British Columbia*, Sept. 1998.
- [14] A. Fei, J.-H. Cui, M. Gerla, and M. Faloutsos. Aggregated Multicast with inter-group tree sharing. *Proceedings of NGC2001*, Nov. 2001.
- [15] T. Henderson and S. Bhatti. Modelling user behaviour in networked games. *Proceedings of ACM Multimedia'01*, Sept. 2001.

- [16] B. Krishnamurthy and J. Wang. On network-aware clustering of web clients. *Proceedings of SIGCOMM'00*, 2000.
- [17] S. Kumar, P. Radoslavov, D. Thaler, C. Alaettinoğlu, D. Estrin, and M. Handley. The MASC/BGMP architecture for inter-domain multicast routing. In *Proceedings of ACM SIGCOMM'98*, pages 93–104, Sept. 1998.
- [18] W. Leland, M. Taqqu, W. Willinger, and D. Wilson. On the self-similar nature of ethernet traffic. *IEEE Transactions on Networking*, 2(1):1–15, Feb. 1994. (earlier version in SIGCOMM '93, pp 183-193).
- [19] M. Jovanovic. Modeling large-scale peer-to-peer networks and a case study of gnutella. *Master thesis, University of Cincinnati*, 2001.
- [20] V. Paxson and S. Floyd. Wide-area traffic: The failure of poisson modeling. *IEEE/ACM Transactions on Networking*, 3(3):226–244, June 1995. (earlier version in SIGCOMM'94, pp. 257-268).
- [21] G. Philips and S. Shenker. Scaling of multicast trees: comments on the Chuang-Sirbu scaling law. In *Proceedings of ACM SIGCOMM'99*, pages 41–51, Sept. 1999.
- [22] P. I. Radoslavov, D. Estrin, and R. Govindan. Exploiting the bandwidth-memory tradeoff in multicast state aggregation. Technical report, USC Dept. of CS Technical Report 99-697 (Second Revision), July 1999.
- [23] P. Rajvaidya and K. Almeroth. Analysis of routing characteristics in the multicast infrastructure. *IEEE INFOCOM*, 2003.
- [24] S. M. Ross. *Introduction to Probability Models (Seventh Edition)*. Academic Press, 2000.
- [25] J. K. Shapiro, J. Kurose, D. Towsley, and S. Zabel. Topology discovery service for router-assisted multicast transport. *Proceedings of OpenArch 2002*, 2002.
- [26] D. Thaler and M. Handley. On the aggregatability of multicast forwarding state. *Proceedings of IEEE INFOCOM*, Mar. 2000.
- [27] B. M. Waxman. Routing of multipoint connections. *IEEE Journal on Selected Areas in Communications*, 6(9):1617–1622, Dec. 1988.
- [28] L. Wei and D. Estrin. The trade-offs of multicast trees and algorithms. *Proceedings of the International Conference on Computer Communications and Networks (ICCCN)*, 1994.
- [29] A. Wolman, G. M. Voelker, N. Sharma, N. Cardwell, A. R. Karlin, and H. M. Levy. On the scale and performance of cooperative web proxy caching. In *Symposium on Operating Systems Principles*, pages 16–31, 1999.
- [30] T. Wong and R. Katz. An analysis of multicast forwarding state scalability. In *Proceedings of the 8th International Conference on Network Protocols (ICNP)*, Japan, Nov. 2000.
- [31] S. C. Zhu, Y. N. Wu, and D. Mumford. FRAME: Filters, Random field And Maximum Entropy: — Towards a Unified Theory for Texture Modeling. — *International Journal of Computer Vision*, 27(2) pp.1-20, March/April, 1998.
- [32] G. K. Zipf. Human behavior and the principle of least effort. *Reading, MA: Addison-Wesley*, 1949.