Towards Super-Scalable Multicast *

Jörg Liebeherr[†] *Bhupinder S. Sethi*[‡]

[†] Polytechnic University
Department of Electrical Engineering
6 MetroTech Center
Brooklyn, NY 11201
Phone: (718) 260-3493, Fax: (718) 260-3074
email: `jorg@catt.poly.edu`

[‡] Department of Computer Science
University of Virginia
Charlottesville, VA 22903
email: `bss4k@cs.virginia.edu`

Abstract

Multimedia collaborative applications for the Internet require the availability of multicast protocols that enhance the basic connectionless IP Multicast service. A critical requirement of such protocols is their ability to support a large group of simultaneous users. For this it is crucial that there exist mechanisms to efficiently exchange control information between the members of a group. In this paper, we present a new approach for distributing control information within a multicast group. The goal of our approach is to scale to very large group sizes (in excess of 100,000 users). Multicast group members are organized as a logical *n-dimensional hypercube*, and all control information is transmitted along the edges of the hypercube. We analyze the scalability of the hypercube control topology and show that the hypercube balances the load per member for processing control information better than existing topologies. We use actual data traces of the group membership in an MBONE conference to gain insight into the transient changes of the load at each node. We present a set of protocol mechanisms that maintain the hypercube topology in a soft-state fashion without requiring any entity to have global state information.

Key Words: *IP Multicast, Multicast Communications, Implosion Problem, Hypercubes.*

*This work is supported in part by the National Science Foundation under Grant No. NCR-9624106.

1 Introduction

Recently emerging interactive multi-user applications have increased the need for advanced multicast services on the Internet. These services are implemented on top of the basic connectionless IP Multicast service, which does not guarantee reliable or in-sequence delivery [7].

In the basic multicast service, a user joins a multicast group simply by indicating interest in receiving data sent to that group. Any packet that is transmitted to a multicast group is forwarded to all members of the group. Services that include error control, rate control, or in-sequence delivery are not part of the basic service. To implement these advanced services, multicast group members must exchange control information with each other. However, the management of control information in a multicast group is a non-trivial problem, especially, if group sizes are large. Without careful management of control information, a multicast protocol can only support small group sizes.

Consider, for example, the implementation of a reliable multicast service. A unicast protocol with a single sender and a single receiver requires the receiver to send positive or negative acknowledgment packets (ACKs and NACKs) to the sender to indicate reception or loss of data. If the same mechanisms are applied to large groups, the sender would soon be flooded by the number of incoming ACK or NACK packets; this is referred to as *ACK Implosion* [6, 13].

In recent years, many techniques and protocol mechanisms have been proposed to cope with the amount of control information that is exchanged between members of a multicast group [2, 5, 9, 12, 13, 16, 17, 19, 20, 22, 28, 24, 25, 29], mostly in the context of providing a reliable multicast service. Early proposals were targeted at local area networks and exploited the broadcast capabilities in such networks [22, 24]. In packet-switching networks, we find two approaches to deal with the volume of control traffic. In one approach, control information is broadcast to all members of the multicast group; the volume of control traffic is limited based on the volume of data traffic or on the size of the multicast groups [3, 9, 23]. The drawback of this approach is that the each multicast member has to reduce its control traffic as the multicast group grows. In the second approach, the group members are organized in a logical graph, henceforth called *control topology.* Only those group members which are neighbors in the logical graph can exchange control information. By merging control information received from their neighbors, the dissemination of control information can be made efficient. Control topologies that have been considered in the literature are rings [5, 28] and trees [12, 17, 19, 29].

This paper proposes a new approach for disseminating control information between the members of a multicast groups. We present a topology that is derived from an *n-dimensional hypercube.* We claim that the hypercube topology has excellent scalability properties, making it the preferred choice for multicast applications with very large group sizes. The key contributions of this study are as follows:

- We show how to construct a hypercube control topology with simple boolean operations.
- We show that the hypercube balances the load for processing control information at multicast group members. Thus, the creation of bottlenecks in the control topology is

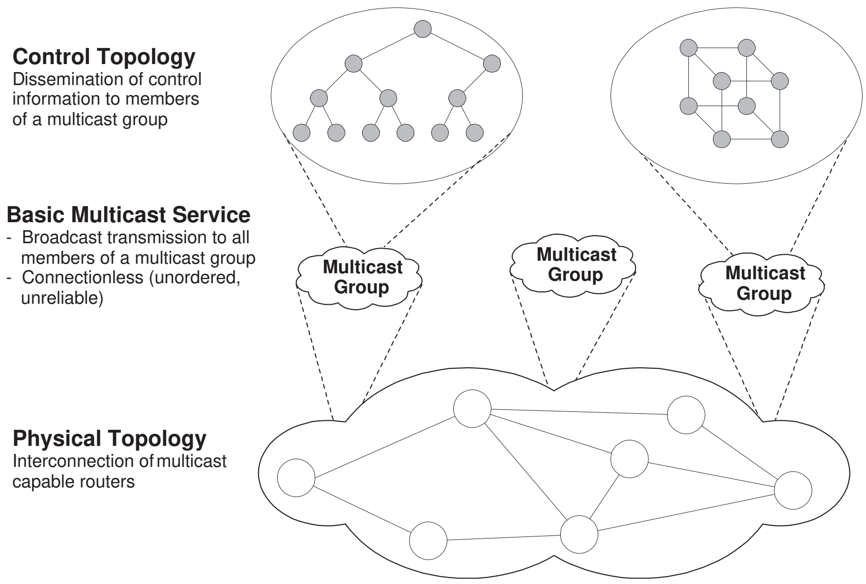

Figure 1: Multicast Framework.

avoided.

- We present a soft-state protocol that maintains the hypercube control topology without requiring any network entity to have global knowledge.

The remainder of the paper is structured as follows. In Section 2 we review the existing proposals for disseminating control information to the members of a multicast group. In Section 3 we present the hypercube as a new solution to disseminate control information in a multicast group. In Section 4 we analyze the scalability properties of a hypercube and compare them with those of other control topologies. In Section 5 we present the basic protocol mechanisms for a soft-state implementation of the proposed architecture. In Section 6 we present an empirical evaluation of various control topologies using actual traces from a large multicast session on the MBONE. In Section 7 we present our conclusions and discuss future work.

2 Control Topologies for Multicast Communications

In this section we review currently used control topologies for disseminating control information in multicast groups, focusing on the ability of these topologies to support large multicast groups. We also introduce some terminology that is used throughout this paper.

An underlying assumption of our work is that communication within a multicast group is *symmetric*, i.e., on the average, each member of the group generates the same amount of traffic. In contrast, an asymmetric groups is characterized by few members of the group generating most or all of the traffic[1].

It is convenient to view the members of a multicast group as a set of nodes V. Nodes are numbered in an arbitrary sequence, that is $V = \{1, 2, \ldots, N\}$. We assume that each node can directly communicate with any other member of the group[2].

2.1 No Control Topology

Several protocols that extend the basic IP Multicast service do not provide a topology for disseminating control information. Instead, the control information from any node is broadcast to all other nodes in the group. Clearly, such a protocol must restrict the volume of control information, since otherwise the scalability is severely impeded. The RTP protocol [23] limits the total amount of control traffic to 5% of the data traffic. In [3], feedback from receivers to the senders is adapted to the size of the multicast group.

Among reliable multicast protocols, the most popular approach to contain control traffic without a control topology is a method known as *NACK suppression* [9, 22] or *damping* [24]. Here, a multicast group member with a control packet to send is forced to queue this packet for a random time interval before it can be transmitted. If a member receives a control packet which matches a queued packet, it cancels the transmission. For large group sizes, however, the random queueing time grows large, resulting in slow feedback times for the control information.

Yet another set of protocols without a control topology employ a central controlling station which coordinates ordering and reliability [4, 8, 10]. Due to the high load at the controlling station, the scalability of such protocols is limited.

2.2 Ring Topology

Ring control topologies have been implemented to provide a reliable multicast service with total ordering of messages [5, 28] . In these protocols, the multicast group is structured as a logical ring, and a token is passed around the ring. Control messages are unicast between the current holder of the token and other nodes. The scalability of ring topologies is only moderate since control messages are always directed to the token holder, thus creating a bottleneck at that node. Also, the time to pass the token around the ring increases for large group sizes, resulting in decreased overall throughput.

2.3 Tree Topology

Tree topologies assume that control information is transmitted along the edges of a *rooted spanning tree*. There is a spanning tree for each multicast member. We use T_k to denote the spanning tree with node k as root. Node l transmits a control message to the root node k

[1] A multicast web server is an extreme example of an asymmetric group; here, the server is the only member of the group that generates traffic.

[2] This can be accomplished by a multicast message which is read only by the destination, or via unicast communication.

by passing the message to its immediate ancestor in T_k, the tree rooted at k. Tree topologies achieve scalability by exploiting the hierarchical structure of a tree. A drawback of tree-based topologies is the overhead in constructing and maintaining the tree. Note that the tree must be dynamically modified both in response to host failures and to members joining and leaving the tree.

Several tree-based control topologies have been proposed for transmission of control information [11, 17, 19, 22, 29], mostly for multicast groups with only a single sender. We discuss the shared K-ary tree topology proposed in [17] which explicitly targets symmetric multicast groups.

The shared tree topology is derived from a single balanced K-ary tree with root r. If some other node $k \neq r$ becomes the root, then the tree is *re-hung* with node k as new root [17]. Re-hanging trees with a new root is illustrated in Figure 2. In Figure 2(a) we show a binary (2-ary) tree with node 1 as root. Figure 2(b) depicts the same tree, 're-hung' for node 10 as root node.

Re-hanging a tree does not increase the number of children of each node. However, after re-hanging, the tree may no longer be balanced. Note that the longest path to the root in the re-hung tree in Figure 2(b) is twice as long as in the original tree in Figure 2(a).

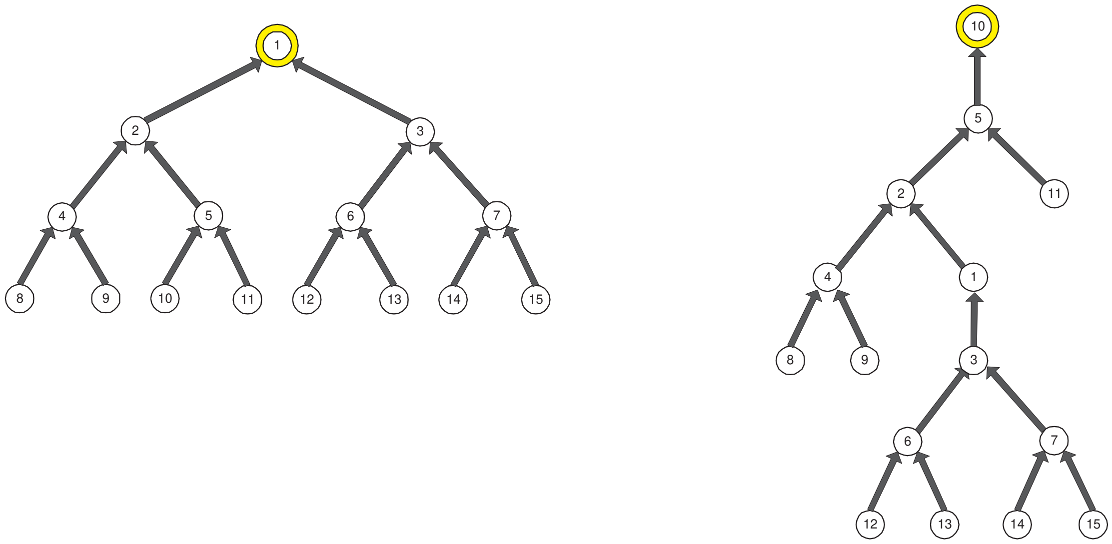

(a) Original Tree Rooted at Node 1.

(b) Re-hung Tree Rooted at Node 10.

Figure 2: Shared Binary Tree.

Among the currently considered topologies, tree-based topologies seem to be most suited to support large multicast groups. However, tree-based topologies are not without problems since re-hanging a shared tree destroys the balance for processing control information among nodes. In the next section we propose a new control topology, derived from a hypercube, that offers better load balancing, and as a result, has better scalability properties than tree-based solutions.

3 The Hypercube Control Topology

In this section we propose the hypercube as a new control topology for multicast communications. We propose organizing the members of a multicast group as the nodes of a logical n-dimensional hypercube. We present a method for embedding spanning trees into the hypercube and use these spanning trees for disseminating control information.

We demonstrate that the hypercube topology is well-suited to support multicast communication in very large groups. In this section we present our method for embedding control trees in a hypercube. In Section 4 we compare the hypercube against the a control topology that is derived from a shared K-ary tree. In Section 5 we present a soft-state protocol for maintaining a hypercube topology in a dynamic multicast group.

3.1 Hypercube and Tree Embeddings

An *n-dimensional hypercube* is a graph with $N = 2^n$ nodes where each node is labeled by a bit string $k_n \dots k_1$ ($k_i \in \{0, 1\}$). Nodes in the hypercube are connected by an edge if their bit strings differ in exactly one position. In Figure 3 we depict hypercubes of dimensions $n = 1$ to $n = 4$.

Hypercubes have been studied extensively by the parallel computing community; they are deemed attractive as a multiprocessor architecture because of their symmetry, the short distances between nodes, and the number alternative routes. The literature on hypercubes is rich, and we refer to [14, 15, 21] as excellent sources on the topic.

In the following we show that the properties of the hypercube topology can be exploited to address the problem of disseminating control information in large multicast groups. We propose to organize the members of a multicast group as the nodes of a hypercube. Then we embed spanning trees into the hypercube and disseminate control information along the edges of the spanning trees.

Past research on parallel algorithms has produced numerous algorithms for embedding trees in hypercubes (see [15] for an overview). The goal of these these algorithms is to assign a parallel computation, represented as a tree, into a multicomputer with an hypercube interconnection network. These algorithms make a number of assumptions that are not applicable in the context of a multicast group. First, most algorithms assume a static hypercube. Second, with few exceptions [26, 27], these algorithms assume a complete hypercube, i.e., $N = 2^n$. Both assumptions are not realistic for multicast groups with a dynamically changing membership.

Our goal is to exploit the strong symmetry of the hypercube in the context of multicast communications. To achieve this we must devise methods that address the problems of actual multicast applications. First, we have to consider *incomplete* hypercubes with $N < 2^n$ nodes. We add the requirement that the embedded spanning trees in an incomplete hypercube with nodes V only contain nodes in V. We refer to such trees as *completely contained* in the incomplete hypercube. Second, we have to consider that the multicast group membership changes dynamically. Since adding and removing nodes may degenerate the compact structure

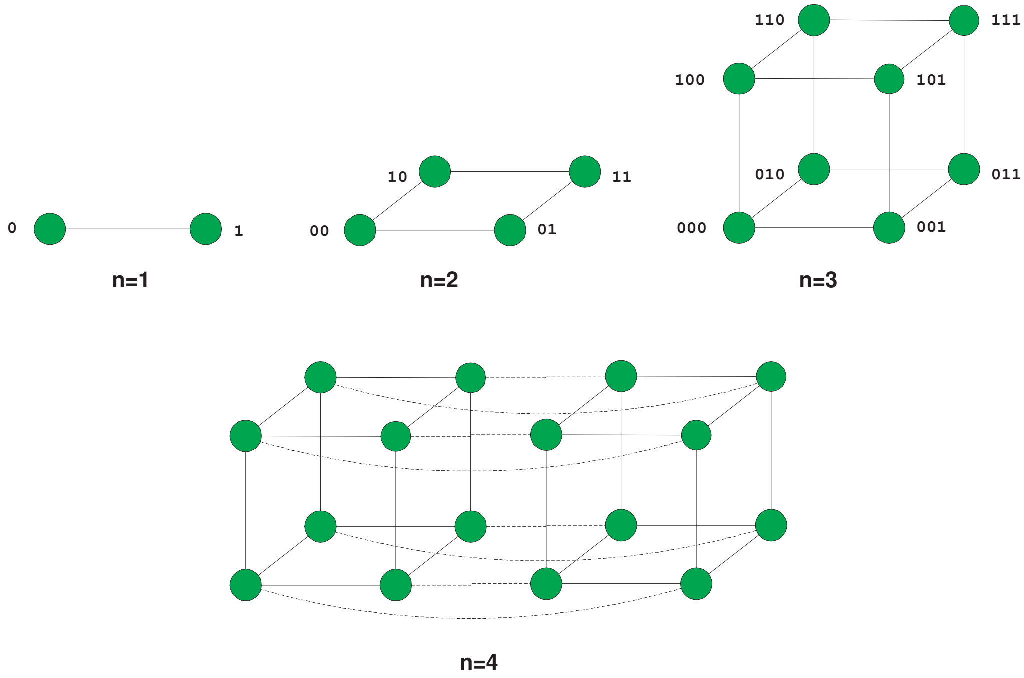

Figure 3: n-dimensional Hypercubes.

of a hypercube, we need to have mechanisms in place that keep the dimension of the hypercube as small as possible; we refer to this property as *compactness*.

3.2 Gray Ordering of Hypercube Nodes

The key to ensure complete containment of all spanning trees and maintain compactness of the incomplete hypercube is the selection of an appropriate ordering of the nodes.

The standard ordering of hypercube nodes interprets the label of a node as a binary number. Specifically, the number $a = \sum_{i=1}^{n} a_i \cdot 2^{i-1}$ is associated with the node labeled $Bin(a) := a_n \ldots a_1$ $(a_i \in \{0, 1\})$. With the ordering imposed by the numbers, compactness can be achieved by ensuring that in a multicast group with N members the positions $Bin(0), Bin(1), \ldots, Bin(N-1)$ are always occupied. However, using this ordering there it is not clear how to construct spanning trees that satisfy the complete containment condition.

As a solution we propose to use a different ordering of the nodes, which is based on an ordering obtained by interpreting node labels using *Gray codes*. A Gray code, denoted by '$G(\cdot)$', is defined via the following properties [21]:

- The values are unique. That is, if $G(i) = G(j) \Rightarrow i = j$.
- $G(i)$ and $G(i+1)$ differ in only one bit, for $0 \leq i < 2^{d-1} - 1$.
- $G(2^{d-1} - 1)$ and $G(0)$ differ in only one bit.

In other words, a Gray code corresponds to a Hamiltonian walk on the hypercube [15].

Let i be a number and $Bin(i)$ its binary representation, it is easy to verify that following generates a Gray code:

$$G(i) := Bin(i) \otimes Bin(i/2)$$

where '$\otimes$' is the XOR operator and '$x/2$' is an integer division by 2. We use $G^{-1}(\cdot)$ to denote the inverse of $G(\cdot)$, that is, $G^{-1}(G(i)) = i$.

By interpreting the node labels in the hypercube as Gray codes, the relationship between i and $G(i)$ defines an ordering of the nodes. Clearly, with this ordering compactness can be enforced by ensuring that the members of a hypercube with K nodes occupy positions $G(0), G(1), \ldots, G(K-1)$.

Example: Consider the ordering of nodes in a 3-dimensional hypercube. Refer to Figure 3 for the labeling of nodes. In the following table, we show the position number i, the binary interpretation $Bin(i)$, and the Gray code $G(i)$:

Position Number i	0	1	2	3	4	5	6	7
$Bin(i)$	000	001	010	011	100	101	110	111
$G(i)$	000	001	011	010	110	111	101	100

Thus, using the standard ordering $Bin(i)$ a multicast group with $K = 5$ members would occupy the following positions in the hypercube: 000, 001, 010, 011, 100. In contrast, using a Gray code $G(i)$ occupies the following positions: 000, 001, 011, 010, 110.

3.3 Tree Embedding in Gray-ordered Hypercubes

We now present an algorithm to embed spanning trees into hypercubes that use Gray codes for ordering the nodes. The embedding of spanning trees in the hypercube is calculated locally: A node with label $G(x)$ directly obtains the address of its parent node in the tree with root $G(r)$. The ability to calculate the embedded tree in a distributed fashion will be exploited in Section 5 where we present protocol mechanisms to maintain a hypercube topology.

Most importantly, for a Gray-ordered hypercube which preserves compactness as shown in the previous subsection, our algorithm, which is presented in Figure 4, always generates a *completely contained* spanning tree. Given two node labels I and R, the algorithm computes the label of the parent node of node I in the spanning tree that is rooted at node R. If each node performs the procedure **Parent** for a root node R, we obtain a spanning tree with root R embedded into the hypercube.

In Figures 5 and 6 we show the embeddings of the spanning trees in a 3-dimensional hypercube for root nodes 1 and 5, respectively. Intentionally, we have depicted an incomplete hypercube with $N = 7 < 2^3$ nodes.

All trees that are constructed by procedure '**Parent**' have the following set of properties. The properties follow directly from the procedure '**Parent**' and are shown without proof.

Property 1: A node and its parent always have a Hamming distance of 1.

Input: Label of the i-th node in the Gray encoding:
$G(i) := I = I_n \ldots I_2 I_1$,
and the label of the r-th node ($\neq i$) in the Gray encoding:
$G(r) := R = R_n \ldots R_2 R_1$.
Output: Label of the parent node of node I in the embedded tree rooted at R.

Procedure Parent (I, R)
  If $(G^{-1}(I) < G^{-1}(R))$
    // Flip the *least significant bit* where I and R differ.
    Parent $:= I_n I_{n-1} \ldots I_{k+1}(1 - I_k) I_{k-1} \ldots I_2 I_1$
      with $k = \min_i (I_i \neq R_i)$.
  Else // $(G^{-1}(I) > G^{-1}(R))$
    // Flip the *most significant bit* where I and R differ.
    Parent $:= I_n I_{n-1} \ldots I_{k+1}(1 - I_k) I_{k-1} \ldots I_2 I_1$
      with $k = \max_i (I_i \neq R_i)$.
  Endif

Figure 4: Tree Embedding Algorithm.

Property 2: The path length between a node and a root is given by their Hamming distance.

Property 3: In a hypercube with N nodes, all trees have a depth of $\lceil \log_2(N) \rceil$. If $N = 2^n$, the embedding results in a binomial tree[3].

Property 4: If **Parent** (I, R) is the p-th node in the Gray encoding, then the following holds: $p \leq \max\{i, r\}$.

Property 4 of the algorithm guarantees complete containment of all embedded trees. Next we analyze the properties of the proposed hypercube control topology and compare it against tree-based solutions.

4 Comparison of Scalability Properties

To gain insight into the scalability property of the hypercube control topology, we conduct a performance comparison with the K-ary shared tree discussed in Subsection 2.3. Our analysis should be seen in the light of our assumption that communication within a multicast group is *symmetric*, that is, on the average each member of the group generates the same amount of traffic.

The scope of our investigation is limited to quantitative aspects of disseminating control information. We do not consider the effects of protocol processing or routing issues. Furthermore, we consider generic transmission of control messages without assumptions on a particular control function (flow control, error control, etc.) as in [20].

[3] A *binomial tree* of height 0 is a single node. For all $i > 0$, a binomial tree of height i is a tree formed by connecting the roots of two binomial trees of height $i - 1$ with an edge and designating one of these roots to be the root of the new tree (cited from [21]).

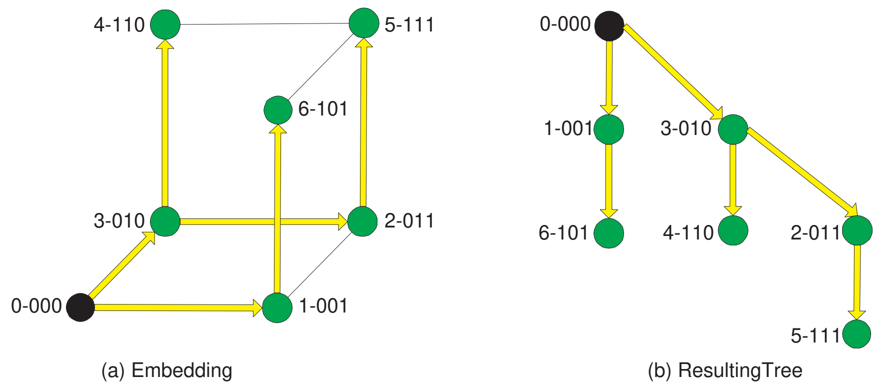

Figure 5: Embedding a Tree with Root $r = 1$.

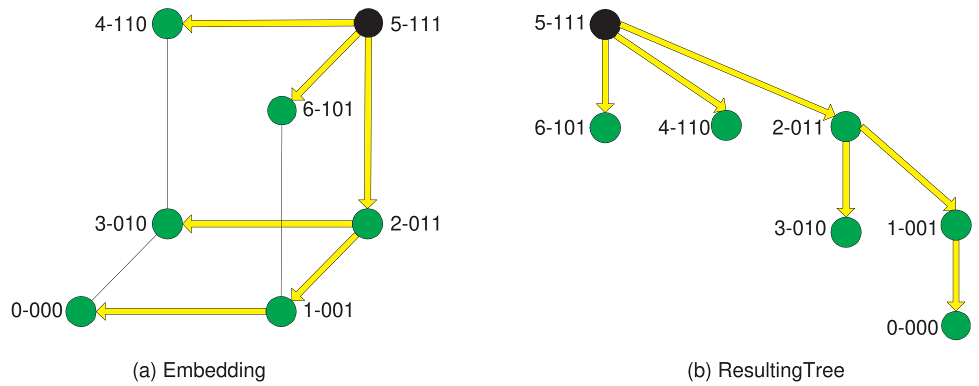

Figure 6: Embedding a Tree with Root $r = 5$.

4.1 Performance Measures

We define a set of performance measures which capture the load for control processing incurred at each member of a multicast group. In all configurations considered, control information is transmitted along the edges of a *rooted spanning tree*. Recall that we use T_l to denote the spanning tree with node $l \in V$ as root.

For most control function, a good indicator for the load at a node in a control tree is the number of direct children. We define:

$$w_k(T_l) := \text{ Number of children of node } k \in V \text{ in tree } T_l.$$

Since control functions may incur a load at a node that is proportional to the number of nodes 'below it' in a tree, we define:

$$\begin{aligned} v_k(T_l) := \;& \text{Number of descendants of node } k \in V \text{ in tree } T_l \text{ (including node } k\text{),} \\ & \text{where the } \textit{descendants} \text{ of node } k \text{ in tree } T_l \text{ are the nodes that} \\ & \text{have node } k \text{ on their path to the root node } l. \end{aligned}$$

Finally, since the path lengths in a control tree indicate delays of passing control information in the tree, we define a measure that expresses this delay:

$p_k(T_l) :=$ Length of the path from node k to root node l in T_l.

Based on the these notions we define more concise measures by taking the average over all trees T_l with $l \in V$, denoted as w_k, v_k, and p_k. These measures are defined as follows:

$w_k := \frac{1}{N}\sum_{l=1}^{N} w_k(T_l)$	Average number of direct children of node $k \in V$ in a spanning tree.
$v_k := \frac{1}{N}\sum_{l=1}^{N} v_k(T_l)$	Average number of descendants of node $k \in V$.
$p_k := \frac{1}{N}\sum_{l=1}^{N} p_k(T_l)$	Average path length from node k to the root.

To further condense the amount of data, we take the averages and maxima of the above values and obtain:

$$w_{avg} := \frac{1}{N}\sum_{k=1}^{N} w_k \qquad v_{avg} := \frac{1}{N}\sum_{k=1}^{N} v_k \qquad p_{avg} := \frac{1}{N}\sum_{k=1}^{N} p_k$$

$$w_{max} := \max_k w_k \qquad v_{max} := \max_k v_k \qquad p_{max} := \max_k p_k$$

We use w_{avg} and v_{avg} as indicators of the average processing load at a node. For w_{avg}, our rationale is that the load of processing control information directly correlates to the number of children of a node. Somewhat differently, v_{avg} correlates the load to the total number of descendants of a node. We interpret p_{avg} as a measure for the delays that occur in disseminating control information; the longer the path length to the root node, the higher the expected delay.

We use the values of w_{max}, v_{max}, and p_{max} to calculate measures for the degree of load balancing of a topology. Specifically, we use the ratios w_{max}/w_{avg}, v_{max}/v_{avg}, and p_{max}/p_{avg} to compare the worst-case node and average node for a control topology. Our expectation is that a control topology with good scalability properties must balance the load incurred at a node, which is reflected in the need for low maximum-to-average ratios.

In the following we present the results of the above measures for the K-ary tree and the hypercube. The complete set of derivations are presented in the appendix of this report.

4.2 Analysis of the Shared K-ary Tree

Recall that the control trees in the shared K-ary tree are obtained from a single K-ary by re-hanging this tree with different nodes as root [17]. In Figure 2 we showed an example of re-hanging a binary tree.

As K is increased, the maximum path length from a node to the (re-hung) root decreases. This, however, increases the load on the node that is the root in the original tree. In the extreme case, we have $N = K$ and obtain a *star topology* where re-hanging the tree always results in $N - 1$ nodes hanging off the original root of a star topology.

Let us assume for simplicity that all leaves of the tree are occupied, i.e., there exists a $d \geq 0$ with $N = \sum_{l=0}^{d} K^l$. Then we obtain:

$$
\begin{aligned}
w_{avg} &= 1 - \frac{1}{N} \\
w_{max} &= K + \frac{1}{N} \\
v_{avg} &= 2d + \frac{K-5}{K-1} + \frac{6d}{N(K-1)} + \frac{4(K-2)}{N(K-1)^2} + \frac{4(d+1)}{N^2(K-1)^2} \\
v_{max} &= V(l_{max}) + 1 \begin{cases} \dfrac{5N}{8} + \dfrac{1}{4} - \dfrac{11}{8N} & \text{if } K = 2 \\[2ex] \dfrac{(K-1)N}{K} + \dfrac{2}{K} - \dfrac{1}{KN} & \text{otherwise} \end{cases} \\
p_{avg} &= 2d - \frac{4}{K-1} + \frac{6d}{N(K-1)} + \frac{4(K-2)}{N(K-1)^2} + \frac{4(d+1)}{N^2(K-1)^2} \\
p_{max} &= 2d - \frac{3}{K-1} + \frac{3(d+1)}{N(K-1)}
\end{aligned}
$$

The derivations for these performance measures are given in Appendix A.

4.3 Analysis of the n-dimensional Hypercube

Calculating the performance measures for the hypercube, where trees are embedded as described in Section 3, requires considerable effort. The complete derivations are shown in Appendix B. For a complete hypercube, that is, $N = 2^n$, we obtain:

$$
\begin{aligned}
w_{avg} &= 1 - \frac{1}{N} \\
w_{max} &= 2 - \frac{\log_2 N + 2}{N} \\
v_{avg} &= \frac{1}{2} \log_2 N + 1 \\
v_{max} &= \frac{1}{8} (\log_2 N)^2 + \frac{3}{8} \log_2 N + 1 \\
p_{avg} &= \frac{1}{2} \log_2 N \\
p_{max} &= \frac{1}{2} \log_2 N
\end{aligned}
$$

4.4 Discussion

We now use the measures from the previous subsections to demonstrate how the control topologies scale when the number of nodes grows large.

Let us first examine the number of children of a node, w_{avg} and w_{max}. Since, for all spanning trees, the average number of children is $w_{avg} \leq 1$, i.e., on the average a node in a control tree has only one child, we only compare the ratios w_{max}/w_{avg}:

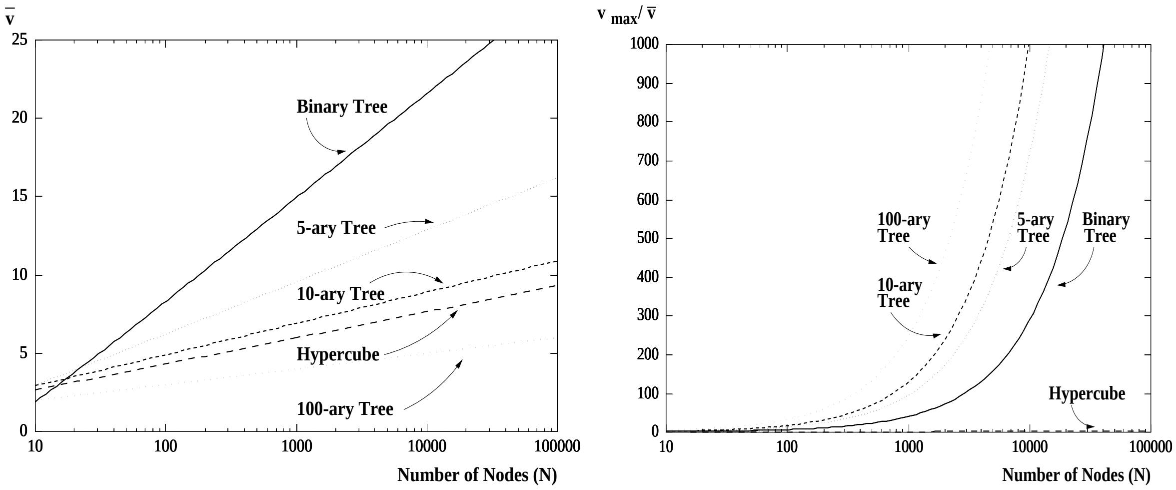

(a) Results for v_{avg}.

(b) Results for v_{max} / v_{avg}.

Figure 7: Comparison of Average Number of Descendants.

K-ary Tree:	$w_{max}/w_{avg} = K$
Hypercube:	$w_{max}/w_{avg} = 2$

Thus, there is a node in a K-ary tree that has K times as many children as the average node resulting in a highly unbalanced load. The hypercube, in contrast, is better load-balanced. Here, the difference between the worst-case and the average case is only a factor of 2. (In the K-ary tree, the maximum is attained for the root in the original K-ary tree. The hypercube attains the maximum at the node with label `00...0`).

In Figure 7(a) we present a graph where we plot the values for v_{avg} by varying the number of nodes N. We present results for the shared K-ary tree (with $K = 2, 5, 10, 100$) and the hypercube. In the figure we see that, in a K-ary shared tree, v_{avg} decreases for increasing values of K. Note that, over the entire range of values, v_{avg} for the hypercube is smaller than the v_{avg} values for the 10-ary tree.

The comparison of the ratios v_{max}/v_{avg}, depicted in Figures 7(b), reveals a problem with load balancing for the shared tree topologies. Since v_{max} increases linearly in N, all shared trees have a bottleneck at the node were the maximum is attained. For large values of K, scalability problems arise even for small group sizes. Even for the binary tree, the maximum-to-average ratio exceeds 100 when the number of nodes has only a few thousand nodes. In a direct comparison with the K-ary tree, the value of v_{max}/v_{avg} for the hypercube appears almost insignificant.

In Figures 8(a) and 8(b) we present the results for the path lengths. It is interesting to note that average path to the root is shorter in the hypercube than in a 10-ary tree. Figure 8(b) shows that load balancing is not an issue when considering the path lengths. As the size of the multicast group is increased, the ratio p_{max}/p_{avg} quickly approaches 1 in all topologies.

In summary, the hypercube appears very suitable to support large multicast groups. A

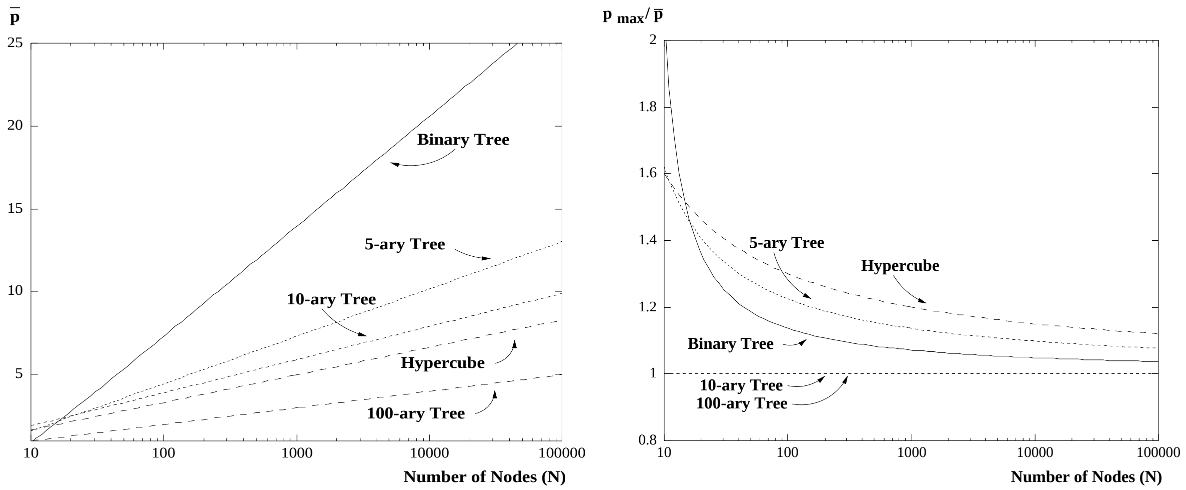

(a) Results for p_{avg}.

(b) Results for p_{max}/p_{avg}.

Figure 8: Comparison of the Average Path Lengths.

comparison with shared K-ary tree showed that the hypercube has all the advantages, and none of the disadvantages of the shared K-ary tree. Particularly, the hypercube provides an excellent balance of the average and worst-case load at the nodes. The load-balancing indicators w_{max}/w_{avg} and v_{max}/v_{avg} clearly demonstrate that the shared tree topology has problems when scaled to very large group sizes. In contrast, the hypercube topology does not have these scalability problems.

5 Protocol Mechanisms for Maintaining a Hypercube Control Topology

We now present a set of protocol mechanisms for maintaining a hypercube control topology. We only discuss the most basic operations for entering a new node ('Join') and eliminating an existing node ('Leave') from the hypercube.

The control topologies considered in this study are built on top of the basic IP multicast service. Recall, that a message sent to an IP multicast address is broadcast to all users listening to this address. No guarantees are made on successful delivery or order of delivery.

The protocol mechanisms to be presented are all *soft-state*; that is, all state information is refreshed periodically without requiring a consistent state at all times.

5.1 Protocol Goals

The goal of the protocol is to organize a group of nodes in a logical hypercube. Every node is aware only of its neighbors in the hypercube. No entity in the system has a notion of the complete state information[4].

We assume that each multicast group has a single multicast channel, referred to as *control channel*. Every member can send and receive on this channel. Obviously, scalability requirements demand that the traffic on this channel be kept minimal[5].

Every node in the hypercube has a *physical address*, given by an IP address/port pair and a *logical address*, given by the node label discussed in Section 3. The logical addresses are totally ordered using the Gray code from Subsection 3.2. The hypercube is said to be in a *stable state* if it satisfies the following three criteria:

1. *Consistent:* No two nodes have the same logical address.
2. *Compact:* In a multicast group with N nodes, the dimension of the hypercube is given by $\lceil \log_2 N \rceil$.
3. *Connected:* Each node knows the physical address of each of its neighbors in the hypercube.

Faults and the mere fact that the protocol is soft-state may lead to a violation of the above conditions, in which case the hypercube becomes unstable. Our protocol returns an unstable hypercube to a stable state. Due to space limitations, we do not discuss recovery from network or node failures.

5.2 Neighborhoods and Leaders

The *neighbors* of a node are those nodes with logical addresses that differ from the logical address of the node in exactly one bit location. Note that in an m-dimensional hypercube, every node has a maximum number of m neighbors. Every node maintains a table with the physical addresses of all its neighbors; this table defines the node's *neighborhood*. A node is said to have a *complete neighborhood* if it knows the physical address of each neighbor that is part of the hypercube. In a stable state, a node can determine whether it has a complete neighborhood, as long as it knows the largest logical address (according to the ordering of the Gray code) in the group.

If, in an unstable hypercube, two or more nodes have the same logical address, we use the physical address to establish a total ordering. Given any node, its successor in the Gray's ordering is defined to be its *ancestor*. In a stable hypercube, every node except the one with the largest logical address has one ancestor. A node without any neighbor with a logical address greater than its own is defined to be a *leader*. A stable hypercube has exactly one leader.

[4]This is in contrast protocols, such as RMP [28], where every node has information about every member in the group, but also different from protocols, such as SRM [9], where no node has any information about the group.

[5]This requirement can be relaxed so that only a small subset of nodes is listening on the control channel at any given point of time. But, currently, we will assume that every member of the group is listening.

5.3 Basic Message Formats

The maintenance operations of the hypercube are based upon a set of messages, called *basic messages*. The protocol has a total of 5 basic messages. A node transmits a message, either by unicasting to one or all of its neighbors, or by multicasting on the control channel. We do not assume transmission of basic messages to be reliable.

Much of the hypercube maintenance is based on the following two basic messages which are repeatedly sent at a given frequency:

- **Ping Message:** Every node periodically *ping*s its neighbors by sending them a short unicast message. A *ping* message contains the logical and physical address of the sender and the receiver of the message. If a node with an incomplete neighborhood receives a *ping* from a new neighbor, it updates its neighborhood table with the physical address of the new neighbor, and begins *ping*ing this new node. If a neighbor of a node does not send *ping*s for an extended period of time, the node assumes the neighbor to have failed, and removes the entry from its neighborhood.

- **Beacon Message:** A *beacon* is a message that is multicast on the control channel. The *beacon* message contains the current neighborhood of the sender. A node can transmit a *beacon* message only if it (1) considers itself to be a leader, or, (2) determines that it does not have a complete neighborhood. Every node uses the *beacon* messages sent by leader(s) to form an estimate of the largest logical address in the group. This information is sufficient to determine whether it has a complete neighborhood.

  Our protocol uses the *beacon* to accomplish the join procedure of new nodes. (The *beacon* messages are also used for stabilizing the hypercube after network or node failures.)

The remaining three basic messages are used for specific tasks:

- **Leave Message:** A node sends a *leave* message when it wants to release its logical address, either because it wants to leave the group or because it is moving to a new logical address. The node sends this message to all its neighbors, which - upon reception - remove the node from their neighborhood.

  REMARK: Since a leave message is not reliable, a neighbor may not receive a *leave* message. In this case, it will notice the absence of the node through missing responses to its *ping* messages.

- **Kill Message:** A *kill* messages is issued to resolve conflicts when two or more nodes happen to have the same logical address. If a node detects that another node has the same logical address as itself, it compares its physical address with the physical address of the other node. If the other node has a lower physical address, it sends a *kill* message, on receiving which, the first node will leave and rejoin the group. If however, its own physical address is lower, it behaves as if it had received *kill* message.

- **Bid Message:** The *bid* message is used by nodes in response to a *beacon* message and contains both the logical and physical address of the sender. A *bid* indicates that a node is willing to move to the neighborhood of the node which sent the *beacon*.

5.4 Node Behavior

The protocol imposes the following rules on each node.

- Every node listens to the control channel and keeps track of the largest logical address in the hypercube (according to the ordering of the Gray code).
- Every node sends a *ping* to its neighbors at a frequency F_1. If a node determines itself to be a leader, it will start to send multicast *beacons* to the control channel at frequency F_2. A node with an incomplete neighborhood will also send *beacons*, albeit at a lower frequency $F_3(<< F_2)$. The frequencies should be set according to the rate of network faults and node failures.
- Each leader node strives to move to a logical address which is lower than its own in the Gray's ordering. Whenever a node can take a new position with a smaller logical address in the hypercube without creating a hole, it will do so. We will call this behavior the *address minimization principle.* More precisely, if a node receives a *ping* directed at its physical address, but to a logical address lower than its own, it will assume the lower logical address contained in the *ping*.
- If a node receives a message from a node with a logical address identical to its own, it resolves the conflict by removing itself from the hypercube, or by sending a *kill* message to the node, as explained previously.

5.5 The Join Procedure

Figure 9 illustrates the steps that are performed when a new node joins the hypercube. The join procedure is completed in six steps; Figure 9(a) shows the first three steps, and Figure 9(b) shows the last three steps. A node that wants to join starts listening to the multicast channel until it encounters a *beacon* packet from the leader node. The node responds with a *bid* packet to the leader. If the leader receives the *bid* packet, it installs the node into its neighborhood and starts *ping*ing the new node with its new logical address. The logical address of the new node is the successor of the old leader in the ordering of the Gray code; hence, the new node becomes the new leader and begins to *beacon* on the control channel (see second picture in Figure 9). As soon, as the neighbors of the new node see the *beacon*, they will update their neighborhood table and start *ping*ing the new node. The new node, will return the *ping*, thereby completing its neighborhood view.

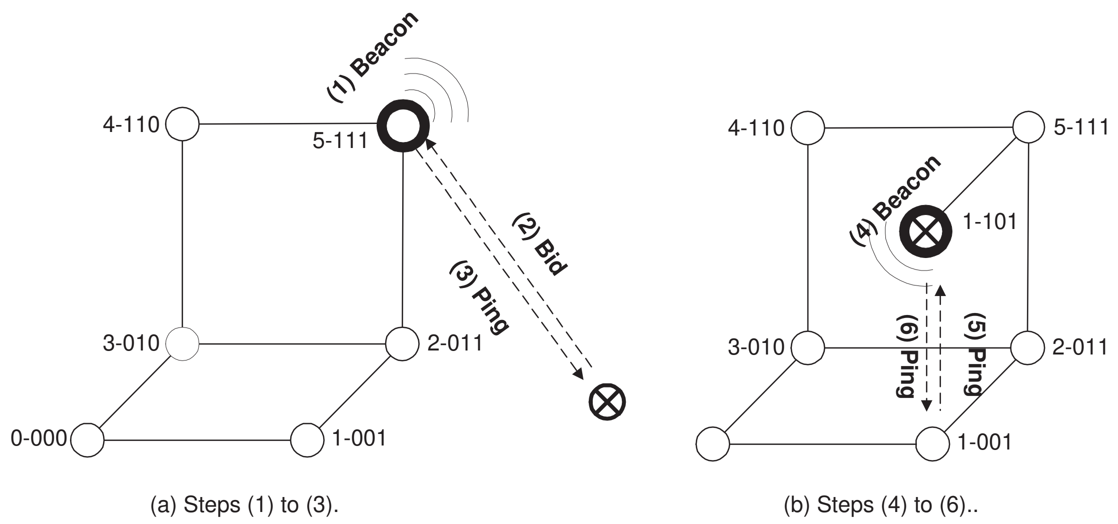

Figure 9: Join Procedure.

5.6 The Leave procedure

The fact that a node has left a position in the hypercube is detected by its neighbors through receipt of a *leave* message, or because no *ping* message was received for an extended period of time. In either situation, the ancestor of the departing node assumes responsibility for restoring the compactness of the hypercube.

To restore compactness, the ancestor waits for a *beacon* from a leader on the control channel. At that point, it installs the leader in its neighborhood, and begins *ping*ing the leader. Because of the address minimization principle, the leader will try to take the empty position. Then, the (old) leader will find its other neighbors in the new neighborhood in a similar fashion as a newly joining node.

The above discussion contains only some of the protocol mechanisms necessary for maintaining the hypercube control topology. A complete description of the protocol will be presented in a forthcoming technical report [18].

6 Empirical Evaluation of Dynamic Behavior

In this section, we present an empirical evaluation of the hypercube control topology. We measure the dynamic behavior of the hypercube control topology, and compare the results with the K-ary shared tree topology.

The basis for our empirical evaluation are data traces of an MBONE session obtained with the *Mlisten* tool [1]. The authors of [1] have made available to us detailed traces of the multicast group membership of MBONE sessions. The traces that are the basis for our evaluation are the video transmission sessions for the *NASA's STS-80C space shuttle mission* measured over a period of 27 days (657 hours) from November 8, 1996 to December 4, 1996.

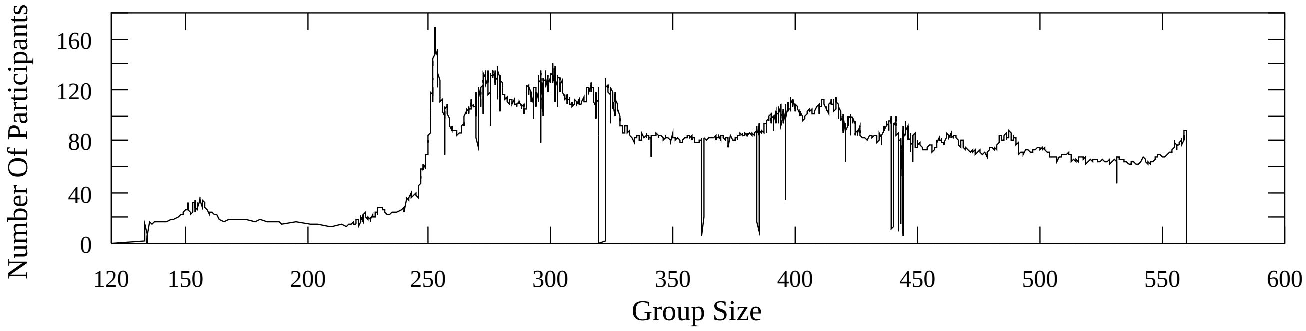

Figure 10: Group Size of the MBONE Session.

The varying size of the group membership over the measurement period from 120 hours to 600 hours is shown in Figure 10. Over the entire length of the experiment, the MBONE sessions had 1,874 different hosts participating with a maximum of 169 simultaneous users (at $t = 252$ hours into the experiment). Note in Figure 10 the occasional drops in the group size, e.g., at $t = 319$ hours. Even though these drops are likely due to artifacts of the measurement experiments, we assume that Figure 10 reflects the actual group size.

We use this trace to track the dynamic behavior of the performance measures from Subsection 4.1 which indicate the average load and the degree of load balancing. [6]

As before, we perform a comparison between the hypercube topology and the shared K-ary tree. In all topologies, we assume that, whenever the group membership has changed, the resulting topology is as compact as possible. In tree topologies, we keep the tree depth minimal, and in the hypercube topology, we attempt to minimize the hypercube dimension. For a hypercube this means we assume that compactness is always maintained. For the tree, our scheme ensures that all nodes at level 0 (root) to d are filled before filling any node at level $d + 1$. This assumption ensures that we always present the most optimistic results for the shared K-ary tree. Actual algorithms for restructuring shared K-ary trees, e.g., [17], may yield worse results.

Next we discuss the outcome of the experiment:

- **Number of children of a node**: In Figure 11 we depict the ratio $w_{max}/\overline{w}$ which reflects the skewness of load balancing. [7] Note that the ratio is lowest for the hypercube. For K-ary trees, the ratio increases as the maximum number of children of a node is increased.
- **Number of descendants of a node:** In Figure 12 we depict the values for the ratio $v_{max}/\overline{v}$. (The values of $\overline{v}$ are ≤ 4 for the hypercube, ≤ 10 for the binary tree, ≤ 6 for the 5-ary tree, and ≤ 3 for the Star network). The hypercube clearly emerges as the topology with the best load balancing properties.
- **Path Length in Control Topology:** Since for all topologies considered, the maximum

[6] We emphasize again that, in this paper, we do not take into consideration any routing aspects.

[7] We mentioned before, that $\overline{w}$ approaches 1 as the number of nodes increases, and is, therefore, not very interesting.

value p_{max} is within a constant from the average $\overline{p}$ (see Section 4), we only show the values for $\overline{p}$. In Figure 13 we see that the average path length in the hypercube topology is less than 4 hops, less than that for the 5-ary tree.

7 Conclusions

We have presented a new approach for disseminating control information between the members of a multicast groups. In our approach, we organize the members of a multicast group in a logical hypercube and assign each multicast group member a number that is derived from a Gray code. The Gray encoding enables each node to locally calculate the next hop for transmitting control information. We analyzed the scalability properties of our approach in symmetric multicast groups, i.e, where each group member is a sender. In a comparison with an K-ary shared tree control topology we showed that the hypercube is superior in balancing the load of control information among all nodes. We outlined a set of protocol mechanism for maintaining the hypercube topology in a packet-switching networks. The protocol was soft-state and did not require global knowledge of the hypercube topology at any network entity.

Our future work on the hypercube control topology will be concerned with both conceptual and practical issues. On a conceptual level, we will account for the geographical location of a host when inserting a node into a hypercube. We are currently working on a Java implementation of the protocol mechanisms. The implementation will give insight into the overhead incurred by the protocol in an actual network. Also, we plan to measure the speed at which the hypercube can recover from network failures.

8 Acknowledgments

We gratefully acknowledge the help of Prof. Almut Burchard from the Department of Mathematics at Princeton University with devising the tree embedding algorithm. We thank Prof. Mostafa Ammar and Dr. Kevin Almeroth from the College of Computing at Georgia Tech for making available to us the traces of the MBONE measurements.

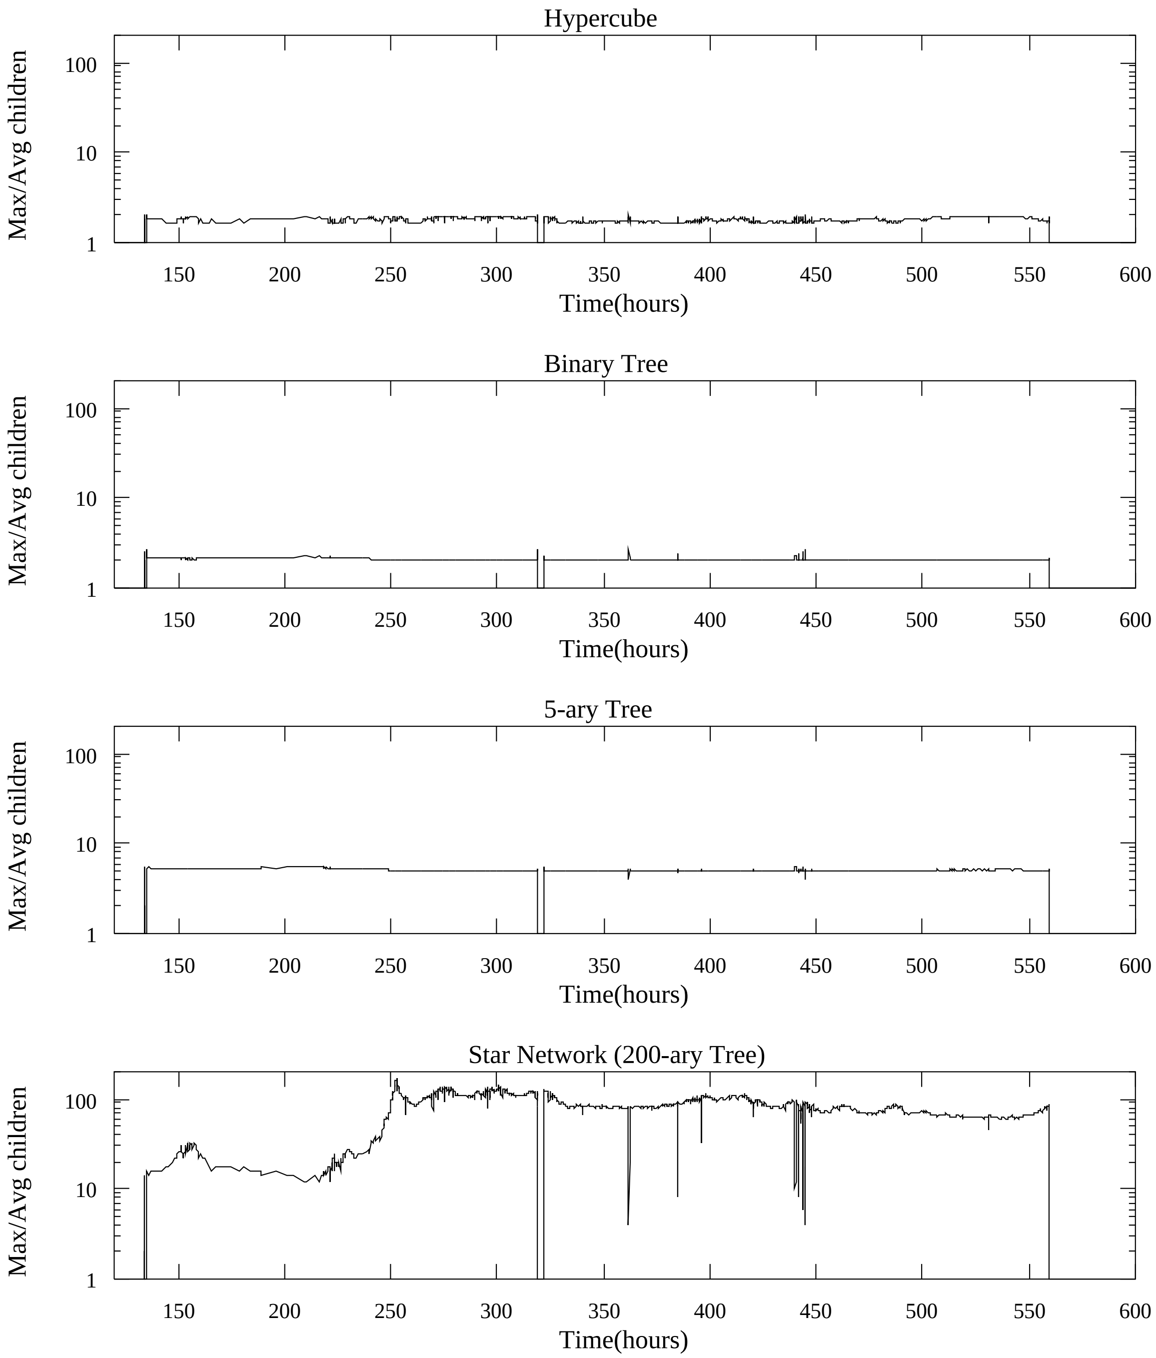

Figure 11: Maximum to Average Children Ratio $w_{max}/\overline{w}$.

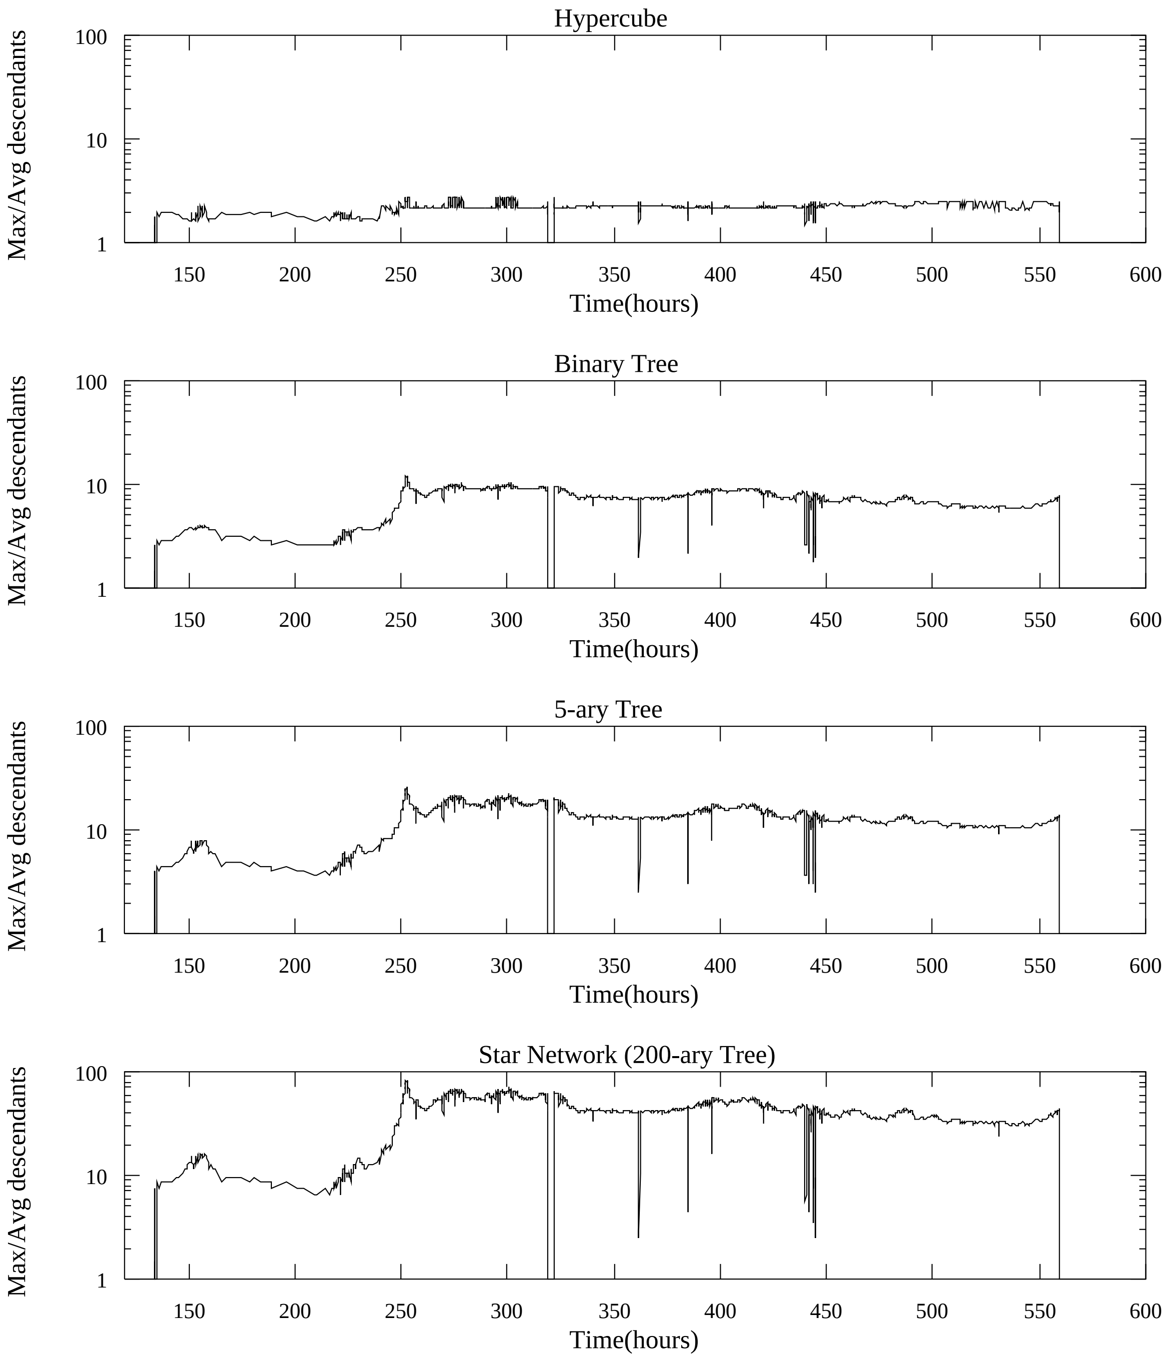

Figure 12: Maximum to Average Descendant Ratio $v_{max}/\overline{v}$.

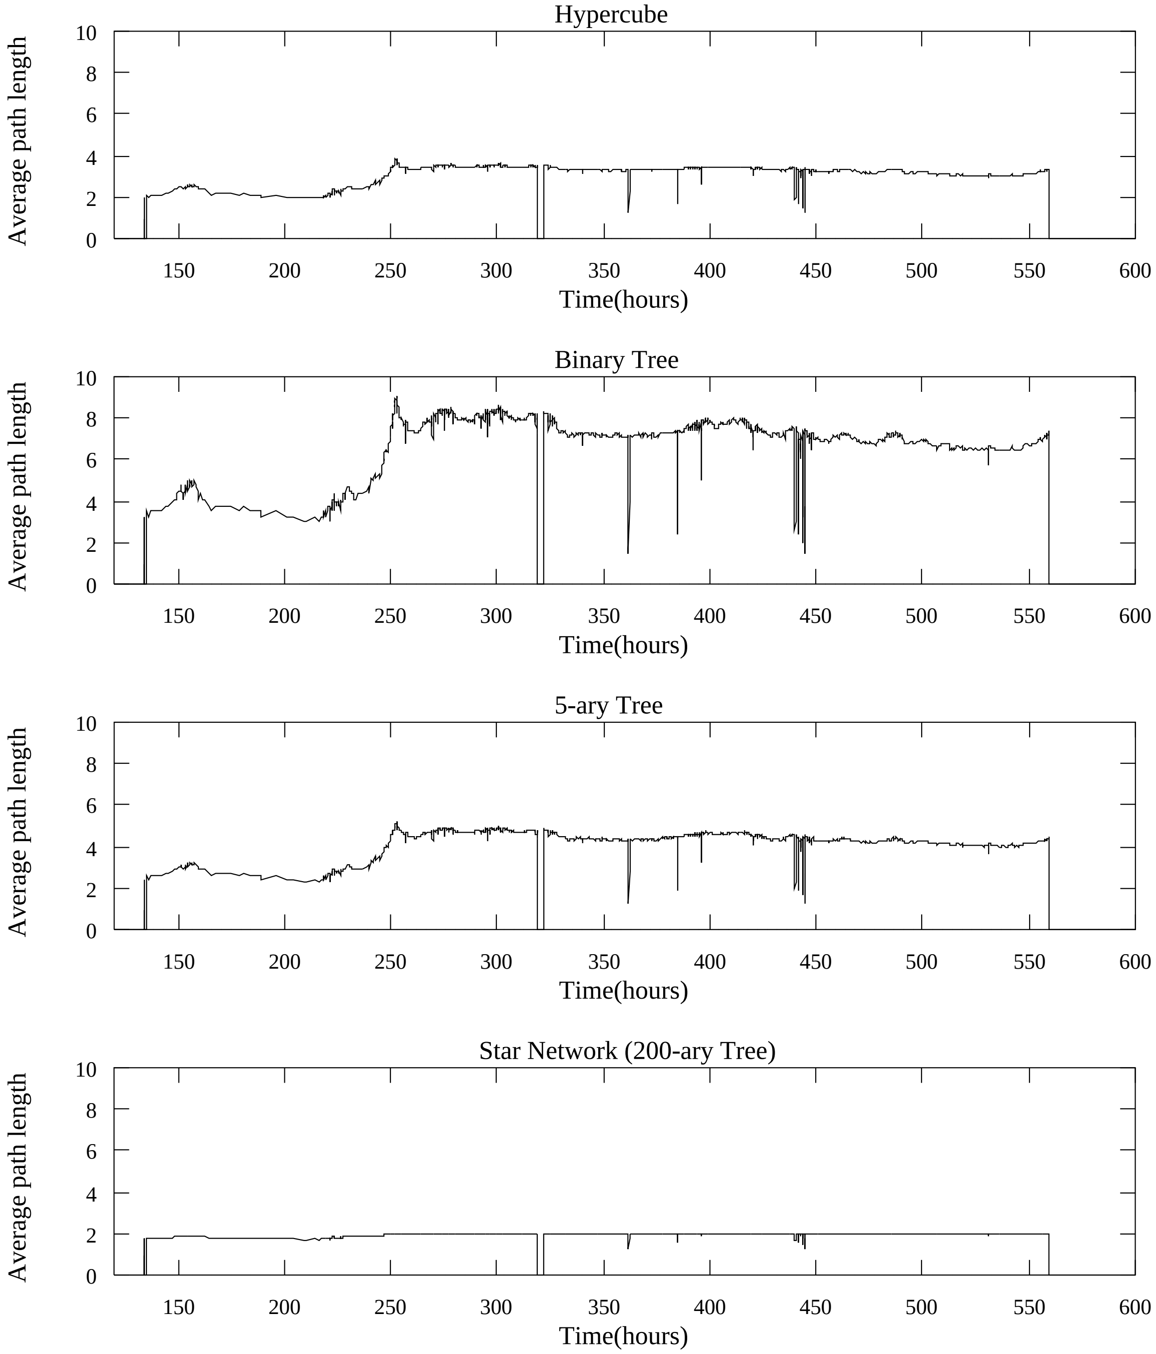

Figure 13: Average Path Length $\overline{p}$.

References

[1] K. Almeroth and M. Ammar. Multicast Group Behavior in the Internet's Multicast Backbone (MBone). *IEEE Communications Magazine*, 35(6), June 1997.

[2] S. Armstrong, A. Freier, and K. Marzullo. Multicast Transport Protocol. Technical Report RFC 1301, Internet Engineering Task Force, February 1992.

[3] J. Bolot. End-to-End Packet Delay and Loss Behavior in the Internet. *Proc. ACM Sigcomm '93*, 23(4):289–298, September 1993.

[4] C. Bormann, J. Ott, H. Gehrcke, T. Kerschat, and N. Seifert. MTP-2: Towards Achieving the S.E.R.O. Properties for Multicast Transport. In *Proc. ICCN '94, San Francisco*, 1994.

[5] J. M. Chang and N. F. Maxemchuck. Reliable Broadcast Protocols. *ACM Transactions on Computing Systems*, 2(3):251–273, August 1984.

[6] J. Crowcroft and K. Paliwoda. A Multicast Transport Protocol. In *Proc. ACM Sigcomm '88*, pages 247–256, August 1988.

[7] S. E. Deering and D. R. Cheriton. Host Groups: A Multicast Extension to the Internet Protocol. Technical Report RFC 966, Internet Engineering Task Force, December 1985.

[8] M. F. Kaashoek et. al. An Efficient Reliable Broadcast Protocol. *Operating Systems Review*, 23(4):5–20, October 1989.

[9] S. Floyd, V. Jacobson, McCanne S, C.-G. Liu, and L. Zhang. A Reliable Multicast Framework for Light-Weight Sessions and Application Level Framing. In *Proc. ACM Sigcomm '95*, pages 342–356, August 1995.

[10] A. Frier and K. Marzullo. MTP: An Atomic Multicast Transport Protocol. Technical report, Cornell University, 1990.

[11] H. Garcia-Molina and A. Spauster. Ordered and Reliable Multicast Communication. *ACM Transactions on Computer Systems*, 9(3):242–271, August 1991.

[12] H. W. Holbrook, S. K. Singhal, and D. E. Cheriton. Log-based Receiver-Reliable Multicast for Distributed Interactive Simulation. In *Proc. ACM Sigcomm '95*, August 1995.

[13] M. G. W. Jones, S. A. Sorensen, and S. Wilbur. Protocol Design for Large Group Multicasting: The Message Distribution Protocol. *Computer Communications*, 14(5):287–297, 1991.

[14] S. Lakshmivarahan and S. K. Dhall. *Analysis and Design of Parallel Algorithms: Arithmetic and Matrix Problems.* McGraw-Hill, New York, 1990.

[15] F. T. Leighton. *Introduction to Parallel Algorithms and Architectures: Arrays, Trees, Hypercubes.* Morgan Kaufman Publishers, San Mateo, 1992.

[16] B. N. Levine and J.J. Garcia-Luna-Aceves. A Comparison of Known Classes of Reliable Multicast Protocols. In *Proc. IEEE International Conference on network Protocols (ICNP '96)*, 1996.

[17] B. N. Levine, D. B. Lavo, and J.J. Garcia-Luna-Aceves. The Case for Reliable Concurrent Multicasting Using Shared Ack Trees. In *Proc. ACM Multimedia '96*, November 1996.

[18] J. Liebeherr and B. S. Sethi. The HyperCast Protocol for Super-Scalable Multicast. Technical report, University of Virginia, 1997. In Preparation.

[19] A. Paul, K. K. Sabnani, J. C.-H. Lin, and S. Bhattacharyya. Reliable Multicast Transport Protocol (RMTP). *IEEE Journal on Selected Areas in Communications. Special Issue for Multipoint Communications*, 15(3):407 – 421, April 1997.

[20] S. Pingali, D. Towsley, and J. F. Kurose. A Comparison of Sender-Initiated and Receiver-Initiated Reliable Multicast Protocols. In *Proc. ACM Sigmetrics '94*, May 1994.

[21] M. J. Quinn. *Parallel Computing: Theory and Practice*. McGraw-Hill, New York, 2nd edition, 1994.

[22] S. Ramakrishnan and B. N. Jain. A Negative Acknowledgment With Periodic Polling Protocol for Multicast over LANs. In *Proc. IEEE Infocom '87*, pages 502–511, March/April 1987.

[23] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson. RTP: A Transport Protocol for Real-Time Applications. Technical Report RFC 1889, Internet Engineering Task Force, January 1996.

[24] W. T. Strayer, B. D. Dempsey, and A. C. Weaver. *XTP: The Xpress Transfer Protocol*. Addison-Wesley Publishing, 1992.

[25] R. Talpade and M. H. Ammar. Single Connection Emulation: An Architecture for Providing a Reliable Multicast Transport Service. In *Proc. 15th IEEE Intl Conf on Distributed Computing Systems (ICDS-15)*, June 1995.

[26] N.-F. Tzeng and H.-L. Chen. Structural and Tree Embedding Aspects of Incomplete Hypercubes. *IEEE Transactions on Computers*, 43(12):1434–1438, December 1994.

[27] N.-F. Tzeng and H. Kumar. Traffic Analysis and Simulation Performance of Incomplete Hypercubes. *IEEE Transactions on Parllel and Distributed Systems*, 7(7):740–754, July 1996.

[28] B. Whetten, S. Kaplan, and T. Montgomery. A High Performance Totally Ordered Multicast Protocol. In *Proc. Infocom '95*, 1995.

[29] R. Yavatkar, J. Griffioen, and M. Sudan. A Reliable Dissemination Protocol for Interactive Collaborative Applications. In *Proceedings of ACM Multimedia 95*, November 1995.

APPENDIX: Proofs of Scalability Properties

This extended appendix contains the formal proofs of the scalability properties of the shared K-ary tree and the hypercube as presented in Section 4. We derive the following performance measures for both the shared K-ary tree and the n-dimensional hypercube.

$$\begin{array}{ccc} w_{avg} & v_{avg} & p_{avg} \\ w_{max} & v_{max} & p_{max} \end{array}$$

Recall the defintions of these performance measures from Section 4. The performance measures are computed as averages over rooter spanning trees; For each node, there is one spanning tree where this node is the root. The spanning trees with some node 'x' as root are obtained as follows:

- **Shared K-ary Tree:** A given K-ary tree, the 'original tree' (see Section 2), is re-hung with node x as root.
- **Hypercube:** A tree with node x as root is embedded in an n-dimensional hypercube using the Algorithm in Figure 4.

In the following, we assume for shared K-are trees that the tree is 'full', that is the total number of nodes is $N = \sum_{l=0}^{d} K^l$. We also assume that the hypercube is complete, that is, a K-dimensional hypercube is assumed to have $N = 2^K$ nodes.

A The Shared K-ary Tree

A.1 Number Of Children

A.1.1 Total Average w_{avg}

The quantity w_{avg} is simply obtained from a general property of trees, which is independent of the specific topology. Clearly, each node in a tree except the root node is a child node. Thus, the number of children nodes in a tree is given by $N-1$, and the the average number of children per node in each tree is $(N-1)/N$. Therefore, w_{avg} is obtained as

$$\begin{aligned} w_{avg} &= \frac{K(N-1)+(K+1)}{N} && (1) \\ &= 1-\frac{1}{N} && (2) \end{aligned}$$

A.1.2 Maximum Average w_{max}

The maximum $w_{max} = \max_k \{w_k\}$ is attained for a node in the original tree with maximum degree. (The degree of a node is the number edges connected to the node.) In a K-ary tree, the maximum degree of any node is $K+1$. If the tree is re-hung, a node with maximum degree has $K+1$ children if it is the root of the re-hung tree, and K children otherwise.

$$w_{max} = K + \frac{1}{N} \tag{3}$$

A.2 Number of Descendants

Before we calculate the values v_{max} and v_{avg}, we derive $A(l)$, the number of nodes in a subtree of a node which is at level l in the original tree. Since, by assumption, the node has K descendants at level $l+1$, K^2 descendants at level $l+2$, etc. , a simple geometric progression yields

$$A(l) = K + K^2 + \ldots + K^{d-l} = \frac{K^{d-l+1} - K}{K-1} \tag{4}$$

Let the value of $A(l)$ averaged over all nodes be denoted as A_{avg}

$$A_{avg} \quad := \quad \frac{1}{N} \sum_{l=0}^{d} K^l A(l) \tag{5}$$

$$\tag{6}$$

Since the number of nodes at level l is K^l and since $N = \sum_{l=0}^{d} K^l$ we can rewrite this expression as

$$A_{avg} \quad = \quad d - \frac{1}{K-1} + \frac{d+1}{N(K-1)} \tag{7}$$

A.2.1 Total Average v_{avg}

We now use $A(l)$ and A_{avg} from Eqns. (4) and (7) to calculate the average number of descendants of a tagged node, which is at level l in the original tree, averaged over all re-hung trees, denoted as $V_{avg}(l)$. $V_{avg}(l)$ only counts the nodes in the subtree of the tagged nodes and does not include the tagged node itself.

Let us consider a K-ary tree ('original tree') and tag a node at level l. Such a tree is shown in Figure 14. Here, the tagged node is labeled as *node x*, its ancestor node is *node A*, one of its children is labeled as *node D*, and the root note is labeled *node R*. When such a tree is re-hung, three types of trees may result:

1. *The tagged* node x *is the root in the re-hung tree (see Figure 15(a)).*
   Here, the tagged node has $N-1$ descendants. Among all N possible re-hung trees, there is only one tree which has the tagged node as root.

2. *The root node of the re-hung tree was a descendant of the tagged node in the original tree. (For example, the root is* node D *or a descendant of* node D*).*
   There is one re-hung tree for each of the $A(l)$ descendants of *node x*. One such tree is depicted in Figure 15(b). In each tree, the number of descendants of node *node x* consists of $N - A(l)/K - 1$ nodes: all nodes in the tree except the tagged node and the nodes in the subtree of *node x* (in the original tree) which contains the root.

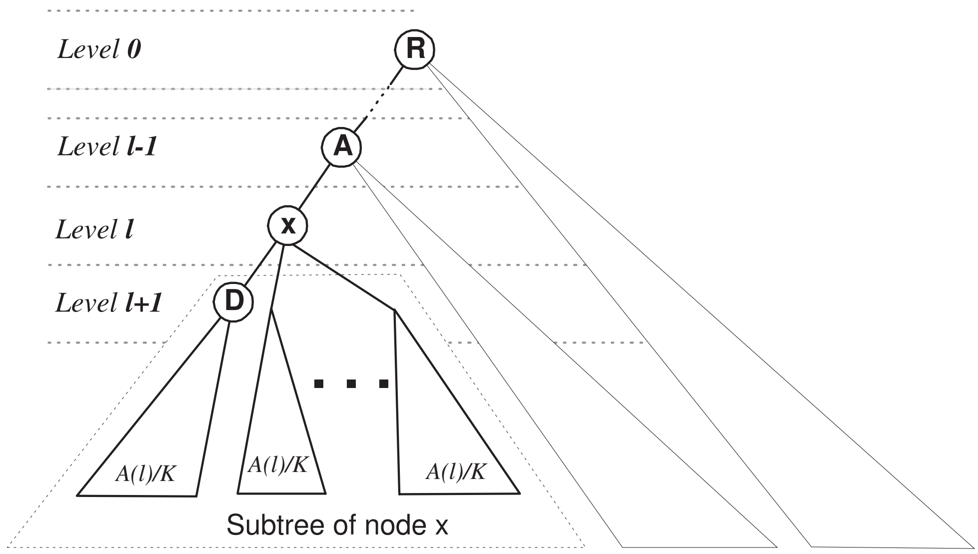

Figure 14: Calculation of $V_{avg}(l)$.

3. *The root of the re-hung tree is neither the tagged node nor a descendant of the tagged node in the original tree (see Figure 15(c)).*
   There are a total of $N - A(l) - 1$ such re-hung trees. Here, the descendants of *node* x in the re-hung tree are the descendants of *node* x in the original tree.

With these considerations, we obtain $V_{avg}(l)$ by averaging over all N trees

$$V_{avg}(l) = \frac{1}{N}\left[A(l)\,(N - A(l) - 1) + A(l)(\,N - \frac{A(l)}{K} - 1\,) + (N-1)\right] \tag{8}$$

$$= \frac{1}{N}\left[(N-1)\,(2A(l)+1) - A(l)^2\left(1 + \frac{1}{K}\right)\right] \tag{9}$$

Let V_{avg} denote the average of $V_{avg}(l)$ over all nodes. We obtain

$$V_{avg} = \frac{1}{N}\sum_{l=0}^{d} K^l V_{avg}(l) \tag{10}$$

With some algebraic manipulations we obtain

$$V_{avg} = (2A_{avg}+1) - \frac{2A_{avg}+1}{N} - \frac{(A_{avg})^2}{N}\frac{K+1}{K} \tag{11}$$

Recall that v_{avg} is defined to include the descendants and the node itself, i.e., $v_{avg} = V_{avg}+1$. Substituting for A_{avg} we obtain

$$v_{avg} = 2d + \frac{K-5}{K-1} + \frac{6d}{N(K-1)} + \frac{4(K-2)}{N(K-1)^2} + \frac{4(d+1)}{N^2(K-1)^2} \tag{12}$$

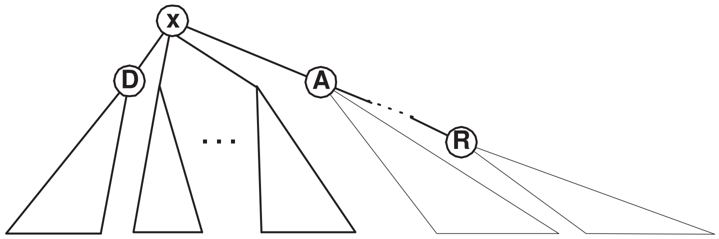

(a) *Node x* is root in the re-hung tree.

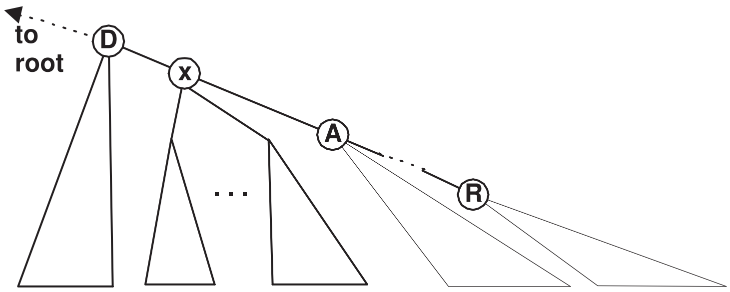

(b) A node in the subtree of *Node x* is root in the re-hung tree.

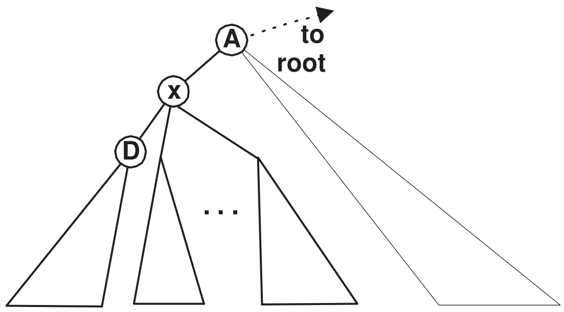

(c) The root in the re-hung tree is neither *Node x* nor a node in the subtree of *Node x*.

Figure 15: Re-hung trees for calculating $V_{avg}(l)$.

A.2.2 Maximum Average v_{max}

The derivation of v_{max} will show that, somewhat counter intuitively, the maximum average number of descendants is not always attained for the root node of the original tree. We first derive an expression for $\max_l \{V_{avg}(l)\}$ with $V_{avg}(l)$ from Eqn. (9). The expression in Eqn. (9) is maximized by maximizing

$$E(l) = 2(N-1)A(l) - A(l)^2(1 + \frac{1}{K}) \tag{13}$$

If $l \neq 0$ and $l \neq d$, then $l-1$ and $l+1$ will exist, and we have the following conditions for a maxima at level l,

$$E(l) > E(l+1) \qquad \text{and} \qquad E(l) > E(l-1) \tag{14}$$

Substituting for $E(l)$, we have

$$2(N-1)A(l) - \frac{K+1}{K}A(l)^2 \;>\; 2(N-1)A(l+1) - \frac{K+1}{K}A(l+1)^2 \tag{15}$$

$$2(N-1)A(l) - \frac{K+1}{K}A(l)^2 \;>\; 2(N-1)A(l-1) - \frac{K+1}{K}A(l-1)^2 \tag{16}$$

Rearranging terms, and cancelling, we are left with the following simplified conditions

$$\frac{K+1}{K}(A(l) + A(l+1)) < 2(N-1) < \frac{K+1}{K}(A(l) + A(l-1)) \tag{17}$$

Hence, our conditions are

$$A(l) + A(l+1) < \frac{2K(N-1)}{K+1} < A(l) + A(l-1) \tag{18}$$

Substituting for $A(l), A(l-1)$ and $A(l+1)$ and simplifying, we have

$$K^{d-l+1} + K^{d-l} < \frac{2K(N-1)(K-1)}{K+1} + 2 < K^{d-l+1} + K^{d-l+2} \tag{19}$$

Since $N = \sum_{l=0}^{d} K^l$, we can replace K^{d+1} by $N(K-1)+1$ and obtain

$$\frac{(N(K-1)+1)(K+1)^2}{2K[K(N-1)(K-1) + (K+1)]} < K^l < \frac{(N(K-1)+1)(K+1)^2}{2[K(N-1)(K-1) + (K+1)]} \tag{20}$$

Let $f(N)$ represent the R.H.S of the second condition. Then we can write

$$\frac{f(N)}{K} < K^l < f(N) \tag{21}$$

Note that $f(N)$ is monotonically decreasing in N. Since, for the above conditions to hold, we must have at least a 2-level tree, i.e., $N > 1 + K + K^2$. Thus, the bounds on $f(N)$ are

$$\frac{(K+1)^2}{2K} < f(N) < \frac{K^3(K+1)^2}{2[K^2(K-1)+1]} \tag{22}$$

And, our condition for l becomes

$$\frac{K^2(K+1)^2}{2[K^2(K-1)+1]} < K^l < \frac{(K+1)^2}{2K} \tag{23}$$

Solving these inequalities gives

$$l_{max} = \begin{cases} 1 & \text{if } K = 2 \\ 0 & \text{otherwise} \end{cases} \tag{24}$$

Substituting l_{max} the formula for $V_{avg}(l)$ and with $v_{max} = V(l_{max}) + 1$ we obtain

$$v_{max} = V(l_{max}) + 1 = \begin{cases} \frac{5N}{8} + \frac{1}{4} - \frac{11}{8N} & \text{if } K = 2 \\ \frac{(K-1)N}{K} + \frac{2}{K} - \frac{1}{KN} & \text{otherwise} \end{cases} \tag{25}$$

A.3 Path Length to the Root

We first derive $p(l)$, the average path length of all nodes to the root of a re-hung tree where the root is at level l in the original tree. For $l = 0$ we have

$$p(0) = \frac{1}{N}\sum_{l=0}^{d} lK^l \tag{26}$$

$$\tag{27}$$

This can be simplified to

$$p(0) = d - \frac{1}{K-1} + \frac{d+1}{N(K-1)} \tag{28}$$

Note that $p(0)$ is identical to A_{avg} from Eqn. (7). Later in this section, we will take advantage off this equality.

For $l > 0$, we can derive a relationship between $p(l)$ and $p(l-1)$. Consider the tree shown in Figure 14. Let us tag *node x* which is at level l. To obatin $p(l)$ we calculate the average path length to *node x* if *node x* is the root of the rehung tree. Let us assume that we know $p(l-1)$, i.e., the average path length to the root if the parent of *node x* (which is *node A* in Figure 14) is the root. Let us know look at the path lengths to the root if *node x* is the root. The paths of all nodes in the subtree of *node x* including em node x's path are one edge shorter, and all the paths from all other nodes are one edge longer. Since we derived $A(l)$, the number of nodes in the subtree of a node at level l, in Eqn (4), we can write

$$p(l) = \frac{1}{N}\left[Np(l-1) - (A(l)+1) + (N - A(l) - 1)\right] \tag{29}$$

$$= p(l-1) - \frac{2}{N}\frac{K^{d-l+1}-1}{K-1} + 1 \tag{30}$$

Resolving the recursive expression yields

$$p(l) = p(0) - \frac{2}{N(K-1)}\left[\frac{K^{d+1}-1}{K-1} - \frac{K^{d-l+1}-1}{K-1} - l\right] + l \tag{31}$$

$$= p(0) - \frac{2}{N(K-1)}[N - A(l) - 1 - l] + l \tag{32}$$

A.3.1 Maximum Average Path Length p_{max}

Since the average path length monotonically increases in l, the maximum occurs at $l = d$, i.e.,

$$p_{max} = p(d) = p(0) - \frac{2}{N(K-1)}[N - A(d) - 1 - d] + d \tag{33}$$

With Eqn. (28) and $A(d) = 0$, we obtain

$$p_{max} = 2d - \frac{3}{K-1} + \frac{3(d+1)}{N(K-1)} \tag{34}$$

$$\tag{35}$$

A.3.2 Total Average Path Length p_{ave}

To find the average path lengths over all roots, we start with Eqn. (32) and take the average over all nodes.

$$p_{avg} = \frac{1}{N}\sum_{l=0}^{d} K^l p(l) \tag{36}$$

$$= p(0) - \frac{2}{N(K-1)}[N - 1 - \sum_{l=0}^{d} K^l A(l) - \sum_{l=0}^{d} lK^l] + \sum_{l=0}^{d} lK^l \tag{37}$$

By inspecting Eqn. (7) and Eqn. (28), we have $p(0) = A_{avg}$.
Hence, we can simplify p_{avg} to

$$p_{avg} = p(0) - \frac{2}{N(K-1)}[N - 1 - A_{avg} - p(0)] + p(0) \tag{38}$$

$$\tag{39}$$

By inserting the expression for $p(0)$ we obtain

$$p_{avg} = 2d - \frac{4}{K-1} + \frac{6d}{N(K-1)} + \frac{4(K-2)}{N(K-1)^2} + \frac{4(d+1)}{N^2(K-1)^2} \tag{40}$$

B The n-dimensional Hypercube

In this section, we prove the properties for trees that are embedded into an n-dimensional hypercube. Recall from Section 3 that all trees embedded into a hypercube with our embedding algorithm are binomial trees. In Figure 16 we illustrate the topology of binomial trees. In the following we assume that the hypercube is full, i.e., $N = 2^n$. Then, unlike the shared n-ary tree, all trees have the same topology.

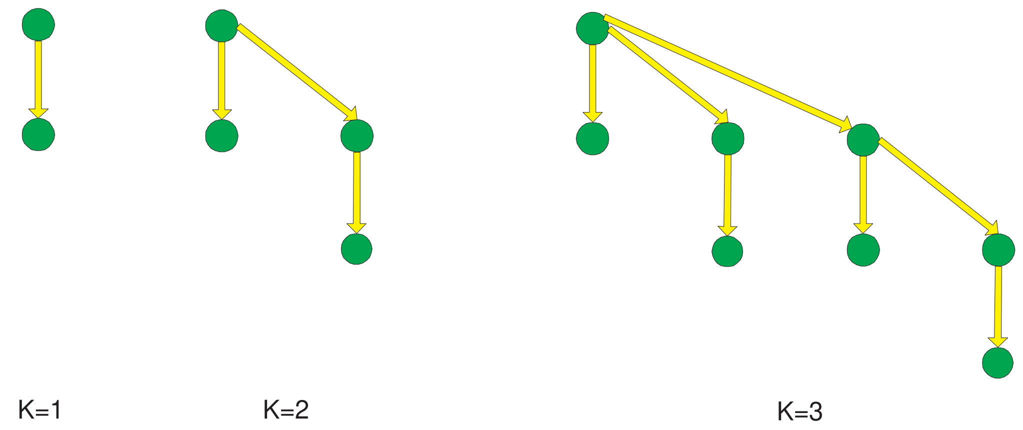

Figure 16: Binomial Tree with eight nodes.

B.1 Number Of Children

B.1.1 Total Average w_{avg}

As in the Subsection A.1, this quantity follows from a general property of trees. For any tree, w_{avg} can be shown to be

$$w_{avg} = 1 - \frac{1}{N} \tag{41}$$

B.1.2 Maximum Average Number of Children w_{avg}

We stipulate that the maximum average occurs for the node that is labeled '00..0', and call this node $\underline{0}$. To obtain w_{max} we now calculate the number of children of this node $\underline{0}$ in the embedded tree with root node R. We distinguish three different label types for root node R and perform the calculation independently for each type. The labels can be as follows are as follows:

must be proven

Case 1: The label of R contains two 1s.

Case 2: The label of R contains a single 1.

Case 3: The label of R is all 0s, i.e., $\underline{0}$ is the root.

We consider each of the cases separately.

<u>Case 1:</u> In this case, the root has j_1 leading zeroes, and j_2 leading zeroes ($j_1, j_2 \geq 0$). Let and 'x' denote either a '0' or a '1', the root R has the following label

$$R = \underbrace{00...00}_{j_1}\ 1\ xx \ldots xx\ 1\ \underbrace{00...00}_{j_2} \tag{42}$$

If S is a child of $\underline{0}$ when R is the root, then the label of S must have one 1 in one of the j_1 leading positions or j_2 trailing positions. At this time, we know that the following two possibilities exist.

$$\begin{array}{rcl} R &=& \underbrace{00...00}_{j_1}\ 1\ xx\ldots xx\ 1\ \underbrace{00...00}_{j_2} \\ \underline{0} &=& \underbrace{00...00}_{j_1}\ 0\ xx\ldots xx\ 0\ \underbrace{00...00}_{j_2} \\ S &=& \underbrace{00.1.00}_{j_1}\ 0\ xx\ldots xx\ 0\ \underbrace{00...00}_{j_2} \end{array} \quad \text{or} \quad \begin{array}{rcl} R &=& \underbrace{00...00}_{j_1}\ 1\ xx\ldots xx\ 1\ \underbrace{00...00}_{j_2} \\ \underline{0} &=& \underbrace{00...00}_{j_1}\ 0\ xx\ldots xx\ 0\ \underbrace{00...00}_{j_2} \\ S &=& \underbrace{00...00}_{j_1}\ 0\ xx\ldots xx\ 0\ \underbrace{00.1.00}_{j_2} \end{array}$$

Thus, node $\underline{0}$ has $j = j_1 + j_2$ children. For each value of j, there are $(j+1)2^{n-j-2}$ different values of R. By multiplying each value with the number of children and summing up the results, we obtain

$$C_1 \ := \ \sum_{j=0}^{n-2} j(j+1)2^{n-j-2} \tag{43}$$

After some algebraic manipulations and with $N = 2^n$ we obtain

$$\begin{aligned} C_1 &= -n(n-1) - 8\left[\frac{n}{4} - \frac{N}{8}\right] - 16\left[\frac{1}{8} - \frac{N}{16}\right] && (44) \\ &= 2N - 2 - n^2 - n && (45) \end{aligned}$$

Case 2: In this case the label of the root R only contains a single 1. Assuming that R has j_1 leading 0s, we obtain

$$R = \underbrace{00...00}_{j_1}\ 1\ \underbrace{00...00}_{n-j_1-1} \tag{46}$$

Similar to the previous case, we obtain the following relationship between R, $\underline{0}$, and a child S of $\underline{0}$ in the embedded tree with node R

$$\begin{array}{rcl} R &=& \underbrace{00...00}_{j_1}\ 1\ \underbrace{00...00}_{n-j_1-1} \\ \underline{0} &=& \underbrace{00...00}_{j_1}\ 0\ \underbrace{00...00}_{n-j_1-1} \\ S &=& \underbrace{00.1.00}_{j_1}\ 0\ \underbrace{00...00}_{n-j_1-1} \end{array} \quad \text{or} \quad \begin{array}{rcl} R &=& \underbrace{00...00}_{j_1}\ 1\ \underbrace{00...00}_{n-j_1-1} \\ \underline{0} &=& \underbrace{00...00}_{j_1}\ 0\ \underbrace{00...00}_{n-j_1-1} \\ S &=& \underbrace{00...00}_{j_1}\ 0\ \underbrace{00.1.00}_{n-j_1-1} \end{array}$$

There are n possible values of R, and for each, there are $n-1$ possible values of S. Hence, the sum of children for this case is

$$C_2 := n^2 - n \tag{47}$$

Case 3: $\underline{0}$ itself is the root, the sum of children is directly obtained as

$$C_3 := n. \tag{48}$$

Taking the average over all three we can write down the solution for w_{max}.

$$\begin{aligned} w_{max} &= \frac{C_1 + C_2 + C_3}{N} && (49) \\ &= \frac{1}{N}\left[(2N - n - 2 - n^2) + (n^2 - n) + n\right] && (50) \\ &= 2 - \frac{n+2}{N} && (51) \end{aligned}$$

B.2 Number of Descendants

B.2.1 Total Average v_{avg}

Consider the binomial trees for $n = 1, 2, 3$ shown in Figure 16, obtained when executing our embedding algorithm for a hypercube with n dimensions. Let us introduce $W^{(k)}_{avg}$ to denote the sum of descendants of all nodes in a binomial tree with $n = k$. By inspection we see that

$$W^{(1)} = 3 \qquad (52)$$

Note from Figure 16 that the binomial tree for $n = k$ is obtained by taking two binomial trees with $n = k - 1$ and interconnecting them at the root. Therefore, we obtain $W^{(k)}$ from $W^{(k-1)}$ as follows:

$$W^{(k)} = 2W^{(k-1)} + 2^{k-1} \qquad (53)$$

We can resolve the recursion in Eqns. (52) and (53) to

$$W^{(k)} = 2^k + k2^{k-1} \qquad (54)$$

Since v_{avg} is simply $W^{(k)}/2^k$ and with $N = 2^n$ we obtain

$$\begin{aligned} v_{avg} = W^{(n)}/2^n &= \frac{1}{2^n}\left[2^n + n2^{n-1}\right] && (55) \\ &= \frac{n}{2} + 1 && (56) \end{aligned}$$

B.2.2 Maximum Average Number of Descendants v_{max}

We stipulate that the maximum average occurs for the node with label address 00..0, and call this node $\underline{0}$. Proceeding in a similar fashion as for the proof of maximum average number of children in Subsection B.1.2, we address the question of calculating, for any embedded tree with root R, the number of descendants of node 00...0, or $\underline{0}$. Again we consider three different scenarios:

Case 1: The label of R contains two 1s.

Case 2: The label of R contains a single 1.

Case 3: The label of R is all 0s, i.e., $\underline{0}$ is the root.

Case 1: The root has the following structure, with j_1 leading zeroes and j_2 trailing zeroes.

$$R = \underbrace{00...00}_{j_1}\ 1\ xx\dots xx\ 1\ \underbrace{00...00}_{j_2} \tag{57}$$

Let S be a descendant of $\underline{0}$ in the embedded tree with R as root. Then S must have all 0s in the middle portion, and have 1s in only the leading and trailing portions as shown below (again, an 'x' indicates either a '1' or a '0').

$$\begin{aligned}
R &= \underbrace{00...00}_{j_1}\ 1\ xx...xx\ 1\ \underbrace{00...00}_{j_2} && (58)\\
\underline{0} &= \underbrace{00...00}_{j_1}\ 0\ 00\dots00\ 0\ \underbrace{00...00}_{j_2} && (59)\\
S &= \underbrace{xx...xx}_{j_1}\ 0\ 00...00\ 0\ \underbrace{xx...xx}_{j_2} && (60)\\
& && (61)
\end{aligned}$$

Thus, node $\underline{0}$ has 2^j descendants where $j = j_1 + j_2$. This includes $\underline{0}$ itself. For each value of j, there are $(j+1)2^{n-j-2}$ different values of R. Thus, for each value of j, the total number of descendants that $\underline{0}$ has, is $(j+1)2^{n-2}$. Summing over all values of j, we have.

$$\begin{aligned}
D_1 &:= \sum_{j=0}^{n-2} 2^j (j+1) 2^{n-j-2} && (62)\\
& && (63)
\end{aligned}$$

After some simplifications we obtain

$$D_1 = 2^{n-3} n(n-1) \tag{64}$$

Case 2: Now consider that root R has exactly one 1 in its label, with $j-1$ leading 0s.

$$R = \underbrace{00...00}_{j_1}\ 1\ \underbrace{00...00}_{n-j_1-1} \tag{65}$$

Here, the relationship between R, $\underline{0}$, and S is the following.

$$\begin{aligned}
R &= \underbrace{00...00}_{j_1}\ 1\ \underbrace{00...00}_{n-j_1-1} && (66)\\
\underline{0} &= \underbrace{0...0}_{j_1}\ 0\ \underbrace{00...00}_{n-j_1-1} && (67)\\
S &= \underbrace{xx...xx}_{j_1}\ 0\ \underbrace{xx...xx}_{n-j_1-1} && (68)\\
& && (69)
\end{aligned}$$

Note that there are n possible labels for R, each with a different position for the 1. For each of these, there are 2^{n-1} possible values of S. Hence, the total number of descendants in this case is

$$D_2 := n2^{n-1} \tag{70}$$

<u>**Case 3:**</u> Finally, we consider that $\underline{0}$ itself is the root. The number of descendants are all nodes in the embedded tree, that is

$$D_3 := 2^n \tag{71}$$

The average for $\underline{0}$ therefore is,

$$\begin{aligned} v_{max} &= \frac{D_1 + D_2 + D_3}{N} && (72) \\ &= \frac{2^{n-3}n(n-1) + n2^{n-1} + 2^n}{N} && (73) \\ & && (74) \end{aligned}$$

With $N = 2^n$, we can simplify this expression to

$$v_{max} = \frac{n^2}{8} + \frac{3n}{8} + 1 \tag{75}$$

B.3 Path Length to the Root

By construction, the path lengths of trees that are embedded in the hypercube only depends on the hamming distance between the root and the node in question. For any root node, the average path length is the average hamming distance from it to other nodes.

B.3.1 Total Average p_{avg}

Note that all embedded tree in a complete hypercube with $N = 2^n$ have the same average path length. The path length of any node to the root node of an embedded tree is the hamming distance of the nodes' labels.

Since for any label of length n, the number of bitstrings that have a hamming distance of k is given by $\binom{n}{k}$, we obtain for the average path length:

$$\begin{aligned} p_{avg} &= \frac{1}{N}\left[\binom{n}{1} \cdot 1 + \binom{n}{2} \cdot 2 + \ldots + \binom{n}{n} \cdot n\right] && (76) \\ &= \frac{1}{N}\sum_{i=1}^{n}\binom{n}{i} \cdot i && (77) \\ & && (78) \end{aligned}$$

Simplifying this expression yields

$$p_{avg} = \frac{n}{2} \tag{79}$$

B.3.2 Maximum Average p_{max}

In a complete hypercube, i.e., $N = 2^n$, the average path length to the root is the same for all embedded trees. Hence, the average path length p_{avg} is the same as the maximum p_{max}.

$$p_{max} = p_{avg} = \frac{n}{2} \tag{80}$$

C Glossary of Notation

Parameter	Definition
V	Set of nodes.
N	Number of nodes ($\|V\| = N$).
n	Dimension of hypercube.
K	"Arity" of a tree, that is, the maximum number of children of an internal node in a tree.
l	Level of a node in question in a tree; we assume $l = 0$ for the root node.
d	Depth of the tree. The depth is defined as the maximum level l of any node in the tree.
T_l	Spanning tree with root node l.
$w_k(T_l)$	Number of children of node $k \in V$ in tree T_l.
$v_k(T_l)$	Number of descendants of node $k \in V$ in tree T_l (including node k).
$p_k(T_l)$	Length of the path from node k to root node l in T_l.
w_k	$= \frac{1}{N}\sum_{l=1}^{N} w_k(T_l)$. Average number of direct children of node $k \in V$ in a spanning tree.
v_k	$= \frac{1}{N}\sum_{l=1}^{N} v_k(T_l)$; Average number of descendants of node $k \in V$.
p_k	$\frac{1}{N}\sum_{l=1}^{N} p_k(T_l)$; Average path length from node k to the root.
w_{avg}	$= \frac{1}{N}\sum_{k=1}^{N} w_k$. Referred to as *total average number of children.*
v_{avg}	$= \frac{1}{N}\sum_{k=1}^{N} v_k$. Referred to as *total average of number of descendants.*
p_{avg}	$= \frac{1}{N}\sum_{k=1}^{N} p_k$. Referred to as *total average of path length to the root.*
w_{max}	$= \max_k w_k$. Referred to as *maximum average number of children.*
v_{max}	$= \max_k v_k$. Referred to as *maximum average of number of descendants.*
p_{max}	$= \max_k p_k$. Referred to as *maximum average of path length to the root.*
$A(l)$	The number of nodes in a subtree of a node at level l (not including the node itself).
A_{avg}	Average value of $A(l)$ averaged over all nodes.