

Direct Queries for Discovering Network Resource Properties in a Distributed Environment

Bruce Lowekamp David O'Hallaron Thomas Gross
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

The development and performance of network-aware applications depends on the availability of accurate predictions of network resource properties. Obtaining this information directly from the network is a scalable solution that provides the accurate performance predictions and topology information needed for planning and adapting application behavior across a variety of networks. The performance predictions obtained directly from the network are as accurate as application-level benchmarks, but the network-based technique provides the added advantages of scalability and topology discovery.

We describe how to determine network properties directly from the network using SNMP. We provide an overview of SNMP and describe the features it provides that make it possible to extract both available bandwidth and network topology information from network devices. The available bandwidth predictions based on network queries using SNMP are compared with traditional predictions based on application history to demonstrate that they are equally useful. To demonstrate the feasibility of topology discovery, we present results for a large Ethernet at CMU.

1. Introduction

Significant work has been done promoting the development of network-aware programs as an effective way of adapting to fluctuating network conditions. Providing the

email: {lowekamp, droh, Thomas.Gross}@cs.cmu.edu

Effort sponsored by the Advanced Research Projects Agency and Rome Laboratory, Air Force Materiel Command, USAF, under agreement number F30602-96-1-0287. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon.

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Advanced Research Projects Agency, Rome Laboratory, or the U.S. Government.

information of which applications should be “aware,” however, has proven to be a significant challenge to developing such applications [1]. User-level benchmarks that record the performance of data transfers across the network provide some information, but suffer from poor scaling and the inability to determine topology information.

Scaling is obviously important in distributed system design. A network performance prediction system must have the ability to give predictions for any combination of machines selected from a large set of possible machines at many sites across the network.

Topology information is extremely important for performance-based processor selection. For example, without topology information, it is impossible to determine when multiple communication paths from the same application will be sharing a single network link. Failure to take into account this intra-application sharing can result in an overestimation of the network performance an application can achieve.

An alternative to benchmarking is obtaining performance information through direct queries to the components making up the network. We show that this direct query approach is not only as accurate as benchmarking, but also can lower the load imposed on the network and provide previously unavailable information about network topology. The Simple Network Management Protocol (SNMP)[4] provides an interface to current networking hardware through which queries can be made to obtain information about the hardware's status.

The remainder of this paper will discuss our results with using SNMP to obtain network information. Section 2 describes several techniques for performance prediction, including benchmark- and network-based approaches. Section 3 provides an overview of the structure of SNMP and how it can be used for resource status discovery. Section 4 analyzes the accuracy of using SNMP to obtain predictions of available bandwidth. Section 5 discusses how to extract topology information from the network components. Section 6 discusses changes needed in the network infrastruc-

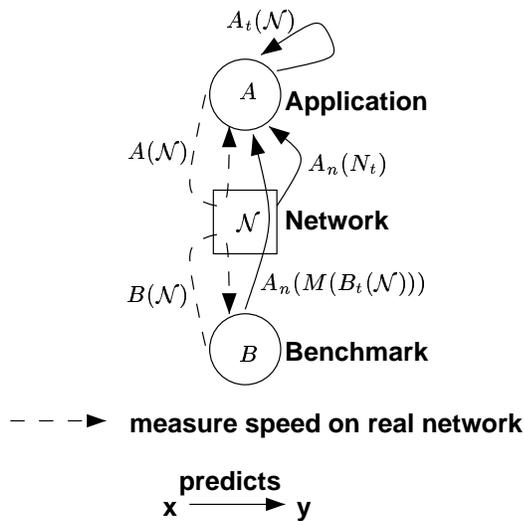


Figure 1. Conceptual diagram of options for prediction. The dashed arrows illustrate the actual application being run on the network as a probe of its status. The solid arrows illustrate the conceptual paths taken to predict an application’s performance using data obtained by running the application itself, by obtaining information directly from the network, and by using a benchmarking program to determine the network’s performance.

ture, including SNMP, so that it is easy for distributed application programmers to use direct network queries to take advantage of the information already stored in the network.

2. Bandwidth prediction techniques

The goal of predicting network performance is to provide accurate predictions of application performance for use with processor selection and application quality parameters. This section discusses three different techniques for performance prediction and introduces a model that describes these techniques. The three techniques represent a continuum of options, from the most straightforward prediction based on the application itself, to the least straightforward benchmark-based prediction. These techniques are illustrated in Figure 1. This section ignores all performance factors except for network bandwidth. In cases where network bandwidth is not the overriding bottleneck, other resource information must be considered.

All three bandwidth prediction techniques rely on the

same basic time series prediction model, which uses a series of measurements to make predictions of future behavior. The difference between the three techniques is what measurements are taken and how they are converted to a prediction of application performance. Ideally, the series of measurements is taken by an independent daemon that collects the data for later use when a user wishes to run an application. A mathematical model is fit to the series. When a user requests a prediction, future performance is extrapolated from the model that has been fit to the past data. Selection and use of time series models has been dealt with by many authors [2, 6, 18]. In our notation, time series models are indicated by a t subscript on the measurement that is used for the series.

2.1. Application-based

The most straightforward measure of an application’s performance is obtained by actually running the application on the network. Similarly, the most straightforward prediction of an application’s future performance on that network is to use the application’s performance history on that network to predict its future performance. The performance of an application A running on a network \mathcal{N} is denoted $A(\mathcal{N})$. The time series model $A_t(\mathcal{N})$ can be used to predict the application’s performance on the network.

Unfortunately, the many combinations of applications, parameters, and resource selections makes gathering enough application history information to provide useful predictions infeasible. For this reason, other prediction techniques must be considered.

2.2. Benchmark-based

Benchmarking solves many of these problems by using a small set of representative applications, called benchmarks or probes, to predict the performance of many applications. The performance of the benchmarking application is denoted $B(\mathcal{N})$. Again, a time series of benchmarks can be used to form a prediction, $B_t(\mathcal{N})$, of how the benchmark B will perform on the network \mathcal{N} in the future.

The challenge with using benchmarks for performance predictions is the mapping from benchmark performance to application performance. One method is to make the assumption that the relative performance of the application and benchmarks are the same, so the best connection for the application is assumed to be the same as the best connection for the benchmark. This approach is often useful for parallel applications where the only concern is moving the data as quickly as possible.

The lack of quantitative information about the application’s performance, however, prevents this technique from being useful in many situations. It does not answer the ques-

tion of which connections are sufficient for the application’s needs, nor does it provide information necessary for setting application quality parameters. Quantitative information is needed for these decisions.

Benchmarking can be used to provide quantitative information, and for some applications, a benchmark will perform similar operations so that the results can be used with a simple rescaling. In other cases, such as using a TCP-based benchmark to predict the performance of a multimedia application that can handle loss, it is necessary to develop a model of the network conditions that caused the benchmark’s performance and then to determine how the application will perform under those same conditions. A mapping function converting the predicted performance of the benchmark to a model of the network can be written as $M(B_t(\mathcal{N}))$. A prediction of application performance based on this network model can be written $A_n(M(B_t(\mathcal{N})))$, where the subscript n is used to denote a performance prediction for the application A based on a network resource model.

Both the application- and benchmark-based techniques are built on the user-level approach of running an application across the network and measuring its performance. Therefore, they both suffer from the same drawbacks. Neither can be used to determine network topology, which is necessary to determine which pairs of machines share the same network component for their connections. Furthermore, predicting performance between all pairs of machines is an $O(P^2)$ operation on the number of processors, P . More complex benchmarking systems use a variety of benchmarks so that they can more accurately predict the performance of different applications, and by grouping machines by location, it is possible to reduce the constant factor of the $O(P^2)$ complexity. But the combination of many applications and $O(P^2)$ pairs of machines makes obtaining and storing this data too costly for systems such as the Globus directory service [7].

However, even these more advanced techniques cannot acquire accurate topology information from the network, which makes predicting the effects of link sharing on the application impossible. Manually providing the topology to the benchmarking tool allows for application-level sharing to be predicted, but it is still difficult to determine which links of the topology are congested.

2.3. Network-based

Rather than obtaining the performance measurement from an application, another approach is to obtain it by querying the network itself. This allows an exact picture of the status of the entire network to be obtained. Network topology can be obtained this way, and the cost is linear in the size of the network. A snapshot of the network \mathcal{N} , con-

sisting of the status of all parts of the network at the same instant, is denoted N . Using a history-based prediction of the network snapshot, N_t , an application’s performance can be predicted as $A_n(N_t)$.

The three techniques described here for obtaining a prediction of an application’s future performance running on \mathcal{N} , denoted $\hat{A}(\mathcal{N})$, are described by the following equations:

$$\hat{A}(\mathcal{N}) \approx \begin{cases} A_t(\mathcal{N}) & \text{Application-based} \\ A_n(N_t) & \text{Network-based} \\ A_n(M(B_t(\mathcal{N}))) & \text{Benchmark-based} \end{cases}$$

\mathcal{N}	real network
N	network status snapshot
A	application
B	benchmark
$A(\mathcal{N})$	performance of A running on \mathcal{N}
$A_t()$	time series performance prediction of A
$A_n()$	network model performance prediction of A
$M()$	function inferring network status from benchmark performance

The network-based technique offers several improvements over the use of benchmarking or application history:

- Direct measurement of the network status allows the performance of different network operations to be predicted, without the need for many different types of benchmarks.
- Network topology is acquired directly from the network, allowing application-level sharing to be predicted.
- Direct network queries require $O(|\mathcal{N}|)$ operations, imposing significantly lower load on the network than $O(P^2)$ benchmarks to obtain the same information.

The major challenge in using network-based predictions of application performance is the necessity of developing the performance prediction function $A_n()$, which maps network status to the application’s performance. Our current approach to this problem is to establish a prediction function by querying the network, then running the application. Once enough measurements are taken, a predictor can be built that takes the bottleneck link of the network as input and predicts the application’s performance over that link.

3. SNMP overview

SNMP was designed to allow network managers to remotely observe and adjust network components. It defines the structure of and operations on a database that is stored on each network component. The database is organized hierarchically, with portions reserved for various standards

bodies and vendors. Components are free to implement only those portions of the hierarchy that are desired. Each portion of the hierarchy is specified by a document referred to as a Management Information Base (MIB). Although it is more correct to refer to only the database protocol as SNMP, in common usage SNMP is used to describe the collection of MIBs as well as the protocol. We follow common usage unless distinctions are needed for clarity. For more information about SNMP and MIBs, many books have been written for use by network managers [16].

3.1. MIB-II

RFC1213 describes the standard MIB, called MIB-II [9]. It is intended to describe essential information needed for all network components—including hosts, routers, and bridges. It provides information about components’ offered services and networking hardware. It also provides statistics and information about major networking protocols, including IP, TCP, UDP, and SNMP.

Two parts of this MIB are of interest to us. The first is the interface table. The entry for each interface provides the maximum data rate as well as octet counters, which indicate the number of bytes the interface has sent and received. These two items allow the available bandwidth on the link attached to that interface to be determined.

Another useful component of this MIB is the IP routing table, which indicates the interface a device will use when sending IP packets to their destination. This is the first item of importance in determining a network’s topology.

3.2. BRIDGE-MIB

The second most important MIB is the BRIDGE-MIB [5]. This MIB provides information about the status of an Ethernet bridge, which is used to forward packets between different portions of a LAN. The interesting part of this MIB is the forwarding database, which stores the port used to reach each of the Ethernet addresses the bridge has seen. Because bridges operate transparently, making queries from this MIB on each bridge is the only way to obtain the information needed to construct the topology of an Ethernet LAN.

4. Experimental verification

While there are clear advantages to the network-based technique, making predictions about higher-level, or end-to-end, operations using low-level information is inherently difficult [14]. We have verified the network-based technique against an application-based technique. If the two techniques offer similar accuracy, the better scaling and ef-

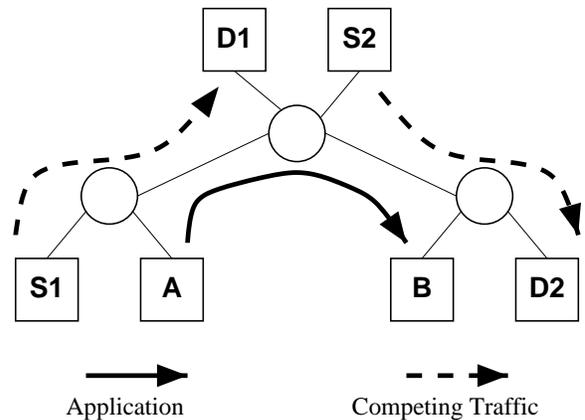


Figure 2. Topology of the testbed used for the prediction experiments. All links are 100Mb. The hosts are 300Mhz DEC Alphas and the routers are Cisco 7206 routers.

iciency advantages of network-based methods make them the better choice for performance prediction.

These experiments were performed on a dedicated testbed where the conditions could be controlled to represent a wide variety of congestion levels. Because of the breadth of conditions experienced on networks, it is important to test prediction techniques at all levels of congestion [11].

4.1. Experimental setup

The network configuration used in the experiments is shown in Figure 2. The “application” used was a simple 1MB data transfer from A to B using TCP. Every 15 seconds, SNMP was used to measure the available bandwidth on all segments of the path between A and B, followed immediately by the application’s data transfer. To measure available bandwidth over different averaging periods, the SNMP packet counts were obtained 5, 3, 0.5, and 0 seconds prior to the TCP message.

A 1MB data transfer would be a typical benchmark. However, for this experiment, it is considered an application because it is being used to predict its own, rather than other applications’, performance. A real application would also involve computation, which is being ignored for the purposes of this paper.

Synthetic traffic was inserted onto the network between S1 and D1, and S2 and D2, resulting in two congested links competing for bandwidth with the application. The competing traffic was generated using fractional Gaussian noise, a method described by Paxson for representing realistic aggregate traffic encountered on networks [10]. The average

rate of competing traffic on each link was chosen between 0Mbps and 100Mbps (link capacity) and changed an average of every 10 minutes.

4.2. Experimental results

Over 65,000 observations were taken during the experiments. To determine the accuracy of the two prediction techniques, 30 sets of 1500 consecutive observations were chosen at random from the experiment. The first 1000 were used to fit the time series model. The prediction technique was then tested over the next 500 observations. The model was refit for each additional observation, so the time series model was only used to predict one observation interval ahead. Each prediction was compared with the next actual observation.

The implementation of the times series predictors that were used is described by Dinda et al. [6]. Two prediction techniques were used. The first is autoregression (AR). The second is the sliding window average (SW). Wolski examined several prediction models and found these two to be useful for network performance prediction [18].

The application prediction was performed on the series of times recorded for the 1MB data transfer. The relative error between each step-ahead prediction and the next actual observation was recorded.

The SNMP-based predictions were done by applying the same time series model to the series of SNMP available bandwidth measurements. The step-ahead prediction of available bandwidth was mapped to the average observed performance at that available bandwidth level. The relative error was then calculated between this predicted number and the next actual observation.

For these experiments, the mapping from available bandwidth to application performance was calibrated using the 1000 training observations and updated with each of the following 500 test observations. In a real system, the calibration would generally only be done once, with occasional updates if the characteristics of the network hardware or competing traffic change.

A comparison between the relative errors is shown in Figure 3. The important observation is that there is little difference between the application- and network-based techniques. This leads us to conclude that the network-based prediction technique can be used to provide network predictions with accuracy comparable to application-based techniques.

4.3. Saturated networks

Because of the nature of network traffic, a link for which SNMP reports little available bandwidth may actually provide a higher rate than reported to a new application if the

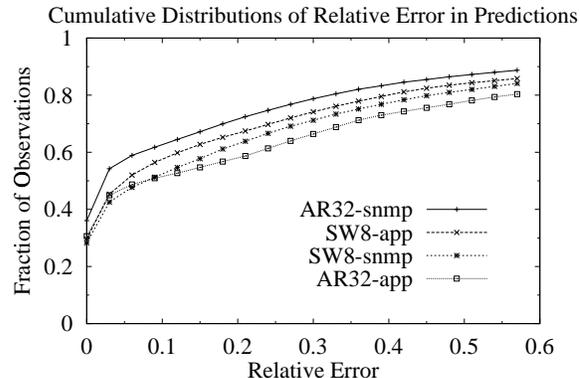


Figure 3. Cumulative relative error distributions for application- and network-based prediction. Both 32nd-order autoregressive (AR32) and 8 observation sliding window mean (SW8) predictions are shown. The SNMP available rate was averaged over 3 seconds for each observation.

competing traffic reduces its rate in response to the new application. This behavior should be represented in the function that maps network status to application performance and is dependent on the type of network and competing traffic involved.

As a link grows more congested, the available bandwidth reported by SNMP approaches zero. In extreme cases, additional information may be obtained by the count of dropped packets, which is also available through SNMP. However, it is doubtful that such a link would be used for any performance-sensitive distributed applications, therefore we have not examined the usefulness of incorporating this information into the model.

5. Topology discovery

Knowledge of network topology is essential for application mapping because the links that are shared by different components of the same application (internal sharing) have a large effect on the performance of that application. Topology knowledge also simplifies scheduling algorithms because it is possible to schedule in a hierarchical fashion, rather than having to look at all combinations of machines. Finally, it is impossible to use direct network queries without first learning the topology to determine which components are involved in the path about which queries are being made.

Topology discovery is difficult because user-transparency has been a great driving force behind

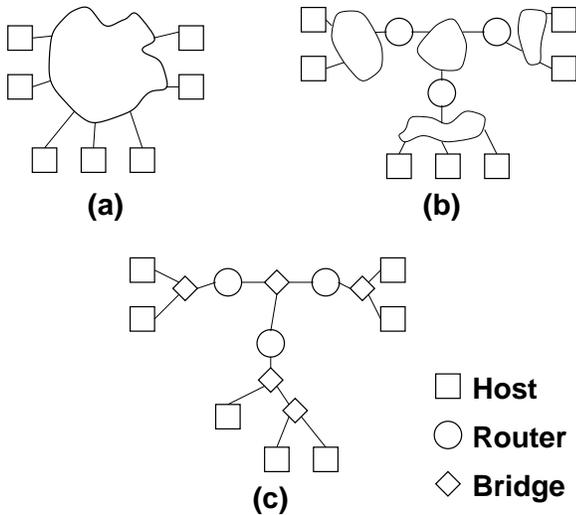


Figure 4. A view of networking at different levels of detail. (a) The view presented to the user. (b) The view at the IP routing layer, where each host and router explicitly forwards packets to the next component in the path. (c) The view including Ethernet bridges, where each bridge learns where the hosts and routers are and transparently forwards the packets towards their destinations.

the success of networking. As a result, there are no good protocols for determining topology. However, the necessary information can be extracted using SNMP with enough perseverance.

5.1. Network structure

Figure 4(a) shows the network view that is presented to the user and that is preserved by most programming libraries. In Figure 4(b), the IP routers connecting these machines are exposed. These are the easiest components to detect. In fact, the `traceroute` program can be used to detect routers between hosts.

The second level of transparency is exposed in Figure 4(c). Here, the bridges that form the Ethernet LANs connecting the machines are exposed. This is the most difficult level of topology to penetrate, although it is the most common LAN infrastructure. The difficulty comes from the beauty of the transparent bridging protocol. The algorithms that the bridges use to determine how to form the LAN and how to forward packets require no global knowledge, nor do the hosts talk directly to the bridges [12]. Thus, the goal of transparency is completely met, at the expense of the ease of determining the topology from the bridges. It

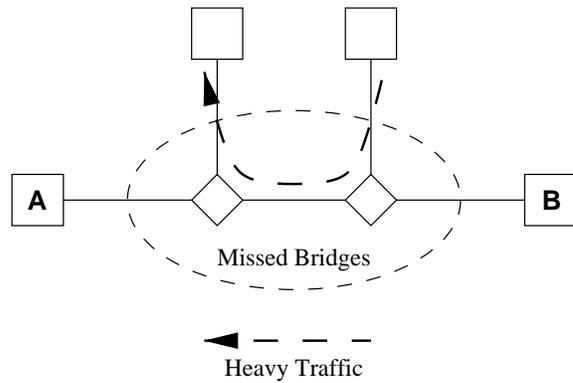


Figure 5. An example of a network configuration where missing bridges can produce misleading results. In this case, missing the two bridges between A and B misses the congested network link between them and may vastly overestimate the available bandwidth.

should be noted that modern networks are typically built with “switches,” which are essentially bridges with many ports.

Despite the difficulties, it is necessary to locate all components of the network topology before using direct queries on that network. Figure 5 shows an example where available bandwidth predictions will be useless because a congested link occurs between two undiscovered bridges.

5.2. IP routing

IP routing topology is easy to determine because the routing table each host and router stores and reports via SNMP explicitly lists the next hop on the route used to reach each destination. As a result, determining IP routing topology is a simple matter of following the routers hop-to-hop from source to destination.

5.3. Bridged Ethernet

The Ethernet bridging algorithm is much more complex than the IP routing algorithm. A bridge learns how to forward packets by listening to all traffic on the links to which it is attached. Whenever it sees a packet, it stores its source address and the link it was received on. This information forms the forwarding database used when forwarding packets to their destination. There are three options for forwarding a packet that a bridge sees on one of its ports:

- If the destination machine is on the same port as the packet was received on, do nothing.

- If the destination machine is on a different port than the packet was received on, the packet is resent, unchanged, on the port for the destination machine.
- If no port is known for the destination machine, then the packet is “flooded,” meaning it is sent to all ports of the bridge. Hopefully the destination machine will reply to the packet’s sender, at which point the bridge will learn its location.

This algorithm is known as transparent bridging, which is currently used almost exclusively on Ethernet LANs. More information can be found in Perlman’s book [12].

Because this algorithm is completely transparent to the hosts, it is difficult to find bridges automatically. One solution is to obtain a list of bridges from an external source, such as the local network manager. Secondly, because bridges only learn a host’s location when they receive a packet from that host, care must be taken to ensure that the forwarding entries are present in the bridge’s database.

This situation motivates a rather complex algorithm for determining the bridging topology. Before beginning, the routing topology must be determined, as in Figure 4(b). Once that has been accomplished, the bridging topology, as found within each cloud in that diagram, can be determined.

The algorithm begins with a set of endpoints, E , consisting of all of the hosts on the Ethernet for which the topology is desired, as well as the routers used to connect this Ethernet to other networks. Also known is the set of bridges, B , used in this network. Bridging topology is defined to be a tree, with the members of E forming the leaves and the members of B forming the internal vertices. The basic approach to determining this topology is to go through the members of B , querying for the ports to which they forward packets routed to members of $E \cup B$. This information tells us the edge from each vertex that is used to reach every other vertex. This knowledge is sufficient to construct the complete topology of the tree.

The difficulty of this algorithm is not in deriving the topology from the bridges’ forwarding databases. Rather, it is in ensuring that the needed entries exist in the forwarding database. A two phase approach is used to obtain this information with reasonable efficiency.

Because there are $O(|B|(|E| + |B|))$ queries to make, and implementation difficulties make each query time consuming, it is important to reduce the number of queries wherever possible. The goal of the first phase is to determine the set of nodes, $B_u \subseteq B$, that are used in the network topology connecting E . This is done by querying each member of B for its forwarding port for each member of E . If the bridge uses the same port to reach all members of E , then the bridge cannot be part of the topology. If a bridge forwards packets to E using different ports, however, then it is needed to form the interconnection topology.

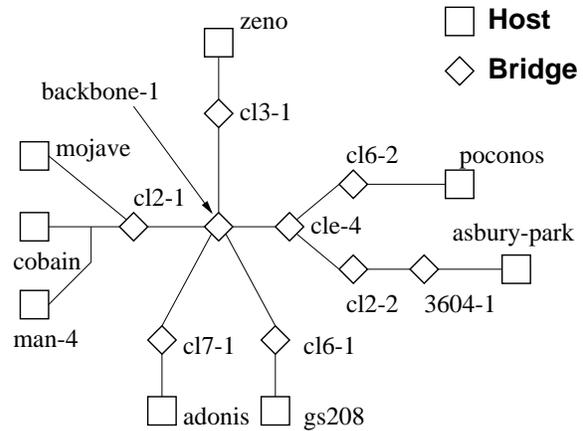


Figure 6. Topology determined from the CMU Computer Science Department’s network. The process was begun with the hosts seen here and the 44 bridges used in our department. Note that cobain and man-4 were found to share a single Ethernet through a hub. The actual department network is much more complex—the algorithm prunes the graph to include only those bridges used to connect the set of hosts being used.

Because bridges learn passively, each bridge in B must have seen a packet from each member of E to have an entry in its forwarding database for that host. To ensure that this table is complete, all members of E periodically ping all other members of E before and during the data collection. This guarantees that if a bridge is on the topology between any two members of E , it will have seen packets from both members and will have their entries in its forwarding database. Note that it is generally not possible for users to have routers send pings, but routers do respond to pings, so if all hosts are sending pings to a router, the router’s replies to the pings will ensure that its entry is present in the forwarding databases of the bridges.

Once B_u has been determined, the second phase of this algorithm begins by expanding the list of machines being pinged by the hosts to $E \cup B_u$. This forces the bridges to learn about each other’s location, information not needed for transparent bridging. Finally, all members of B_u are queried for where they are forwarding packets to all other members of B_u .

Following the completion of this algorithm, the ports used by each bridge to forward packets to $E \cup B_u$ are known. Because the bridging topology is defined to use a tree, it is easy to extend this information to complete the topology.

Figure 6 shows the bridging topology between several machines at CMU. This structure was determined using the above algorithm, beginning with a list of the 44 bridges used in our departmental Ethernet. This topology was verified by our network manager after the algorithm was run.

The $O(|B|(|B| + |E|))$ complexity of this algorithm and the high cost of each access to a bridge's forwarding database make this algorithm impractical for routine use. For example, the topology in Figure 6 took approximately 30 minutes to discover. Fortunately, typical Ethernet LANs do not change topology very often, and it is easier to verify that the topology is still accurate than to discover it. For these reasons, the bridging topology should be stored in a database where applications and prediction systems can make use of it, with it being rediscovered only when necessary.

6. Practical considerations

Our research has demonstrated that SNMP, already supported by almost all of the current networking infrastructure, is sufficient for obtaining the information needed to determine and predict performance directly from the network. Although it is not an ideal interface for this purpose, it allows the network-based approach to performance prediction to be explored and utilized on existing networks. Demonstrating the value of this approach by using it in real systems and applications should result in the development of more appropriate interfaces for network components. However, both administrative and technical considerations must be addressed to provide a better interface for performance prediction purposes.

The administrative complication is mostly accessibility. Typically, SNMP access is only allowed from machines on the local network, and it is usually impossible to make SNMP queries to network components on an ISP's network. Security and privacy are the two primary reasons for this. Security is actually a technical concern; because the designers of SNMP were unable to agree on a workable security protocol, there is little security in current implementations, therefore a minimal security level is achieved by restricting access to local hosts. ISP's are generally concerned about privacy, not wishing to divulge information about the congestion levels of their services. Furthermore, because SNMP queries can be expensive, no one wants to open their network up to excessive load or even denial-of-service attacks with SNMP. We are currently pursuing combining network-based data with benchmark-based data to provide predictions in environments where direct network queries are only available for portions of the network.

Although there are standards documents that describe the behavior of SNMP implementations, the standards and their implementations have not resulted in consistent interfaces

between different manufacturers. For instance, the forwarding databases in Ethernet bridges are particularly troublesome. Some allow queries to be made for the forwarding port of a specific address. Other implementations are designed only for traversal, requiring the same query to be reformulated as a query for the subsequent entry from the numerically preceding address. Furthermore, some bridges remove the forwarding database if queries are made to it too rapidly, apparently as a security measure.

Finally, although SNMP provides much of the information needed for distributed computing, it is difficult to get it in the form required. For example, there are traffic counters for each port, but determining a traffic rate requires multiple, carefully timed queries. It would be much more appropriate to have the router calculate its own time-averaged rate. A limited form of active networking would be ideal for this type of information.

7. Related work

A variety of systems exist for providing network status information. NWS [19] and Prophet [17] provide applications with benchmark-based predictions. SPAND [15] records similar data by storing applications' actual performance during execution and making this data available to help future applications.

All prediction systems described here utilize time series prediction techniques [2]. Wolski has studied several different time series models for their usefulness in predicting network performance [18].

SNMP has much broader uses than those that are described here. It is used to control and monitor a wide range of network resource properties [16]. Busby has explored using SNMP to gather information about both network and CPU resource as an addition to NWS [3]. The techniques evaluated in this paper for SNMP-based bandwidth measurement were developed for the Remos system [8].

Recent network research has focused on modeling the self-similarity in network traffic, and these models may lead to more realistic traffic than Poisson processes. We used fractional Gaussian noise to generate self-similar traffic for our experiments [10]. More realistic wavelet models have been investigated more recently [13], but we have not yet rerun our experiments with the new models.

8. Conclusions

Tools for discovering network performance and topology are an important part of the support infrastructure for distributed computing. This paper demonstrates that the network-based approach has significant advantages over the benchmark-based approach for scaling and topology discovery, without sacrificing accuracy. Importantly, these

queries can be implemented today, using the SNMP structure already included in networking hardware for the purpose of network management. The benefits of the network-based approach outweigh the inherent limitations of using low-level information to predict end-to-end performance.

Major challenges remain for the distributed computing community to ensure the availability of the information needed for network-based predictions. Network administrators and companies must be encouraged to provide easier access to the performance data needed to properly schedule applications. Furthermore, working with the networking community may lead towards the incorporation of more beneficial and convenient features for performance measurement into SNMP and network devices.

The techniques described here are being implemented in the Remos system [8]. Remos is available from our website, <http://www.cs.cmu.edu/Groups/CMCL/remulac>.

References

- [1] J. Bolliger and T. Gross. A framework-based approach to the development of network-aware applications. *IEEE Transactions on Software Engineering*, 24(5):376–90, May 1998.
- [2] G. E. P. Box, G. M. Jenkins, and G. Reinsel. *Time Series Analysis: Forecasting and Control*. Prentice Hall, 3rd edition, 1994.
- [3] R. E. Busby, Jr. Predictive notification of QoS shortfall with the network weather service. <http://www.cis.ksu.edu/~rbusby/NWS.html>.
- [4] J. Case, K. McCloghrie, M. Rose, and S. Waldbusser. Structure of management information for version 2 of the simple network management protocol (SNMPv2). RFC1902, January 1996.
- [5] E. Decker, P. Langille, A. Rijsinghani, and K. McCloghrie. Definitions of managed objects for bridges. RFC1493, July 1993.
- [6] P. A. Dinda and D. R. O'Hallaron. An evaluation of linear models for host load prediction. In *Proceedings of the 8th IEEE International Symposium on High Performance Distributed Computing (HPDC)*, August 1999.
- [7] S. Fitzgerald, I. Foster, C. Kesselman, G. von Laszewski, W. Smith, and S. Tuecke. A directory service for configuring high-performance distributed computations. In *Proceedings of the 6th IEEE International Symposium on High Performance Distributed Computing*, pages 365–75. IEEE, August 1997.
- [8] B. Lowekamp, N. Miller, D. Sutherland, T. Gross, P. Steenkiste, and J. Subhlok. A resource monitoring system for network-aware applications. In *Proceedings of the 7th IEEE International Symposium on High Performance Distributed Computing (HPDC)*, pages 189–196. IEEE, July 1998.
- [9] K. McCloghrie and M. Rose. Management information base for network management of TCP/IP-based internets: MIB-II. RFC1213, March 1991.
- [10] V. Paxson. Fast, approximate synthesis of fractional gaussian noise for generating self-similar network traffic. Technical Report LBL-36750, Lawrence Berkeley National Laboratory, April 1995.
- [11] V. Paxson and S. Floyd. Why we don't know how to simulate the internet. In *Proceedings of the 1997 Winter Simulation Conference*, pages 1037–44, 1997.
- [12] R. Perlman. *Interconnections: Bridges and Routers*. Addison-Wesley, 1992.
- [13] V. J. Ribeiro, R. H. Riedi, M. S. Crouse, and R. G. Baraniuk. Simulation of nongaussian long-range-dependent traffic using wavelets. In *Proceedings of ACM SIGMETRICS '99*, pages 1–12. ACM, May 1999.
- [14] J. Saltzer, D. Reed, and D. Clark. End-to-end arguments in system design. *ACM Transactions in Computer Systems*, 2(4):277–288, November 1984.
- [15] S. Seshan, M. Stemm, and R. H. Katz. SPAND: Shared passing network performance discovery. In *Proceedings of the USENIX Symposium on Internet Technologies and Systems*, pages 135–46, December 1997.
- [16] W. Stallings. *SNMP, SNMPv2, and RMON*. Addison-Wesley, 2nd edition, 1996.
- [17] J. B. Weissman and X. Zhao. Scheduling parallel applications in distributed networks. *Cluster Computing*, 1(1):95–108, May 1998.
- [18] R. Wolski. Dynamically forecasting network performance using the network weather service. Technical Report CS-96-494, UCSD, 1996.
- [19] R. Wolski. Forecasting network performance to support dynamic scheduling using the network weather service. In *Proceedings of the 6th High Performance Distributed Computing Conference (HPDC)*, pages 316–25, August 1997.