

IQ-Paths: Predictably High Performance Data Streams across Dynamic Network Overlays

Zhongtang Cai, Vibhore Kumar, Karsten Schwan
{ztcai, vibhore, schwan}@cc.gatech.edu
College of Computing, Georgia Institute of Technology

Abstract

Overlay networks are a key vehicle for delivering network and processing resources to high performance applications. For shared networks, however, to consistently deliver such resources at desired levels of performance, overlays must be managed at runtime, based on the continuous assessment and prediction of available distributed resources. Data-intensive applications, for example, must assess, predict, and judiciously use available network paths, and dynamically choose alternate or exploit concurrent paths. Otherwise, they cannot sustain the consistent levels of performance required by tasks like remote data visualization, online program steering, and remote access to high end devices. The multiplicity of data streams occurring in complex scientific workflows or in large-scale distributed collaborations exacerbate this problem, particularly when different streams have different performance requirements. This paper presents IQ-Paths, a set of techniques and their middleware realization that implement self-regulating overlay streams for data-intensive distributed applications. Self-regulation is based on (1) the dynamic and continuous assessment of the quality of each overlay path, (2) the use of online network monitoring and statistical analyses that provide probabilistic guarantees about available path bandwidth, loss rate, and RTT, and (3) self-management, via an efficient packet routing and scheduling algorithm that dynamically schedules data packets to different overlay paths in accordance with their available bandwidths. IQ-Paths offers probabilistic guarantees for application-level specifications of stream utility, based on statistical predictions of available network bandwidth. This affords applications with the ability, for instance, to send control or steering data across overlay paths that offer strong guarantees for future bandwidth vs. across less guaranteed paths. Experimental results presented in this paper use IQ-Paths to better handle the different kinds of data produced by two high performance applications: (1) a data-driven or interactive high performance code with user-defined utility requirements and (2) an adaptive overlay ver-

sion of the popular Grid-FTP application.

1 Introduction

Data-driven high performance applications are important to many constituencies, including corporations in applications like real-time data mining or data integration [3], common end users in telepresence [18], and scientists or engineers in applications like remote data visualization [29] or instrument access [22]. A common characteristic of such applications is their need to meet quality of service (QoS) guarantees and/or offer utility-based services to end users (i.e., meet certain service-level objectives (SLOs)). However, excepting datacenter-based solutions [3] and the few dedicated, high end links existing between select centers of excellence (e.g., via DOE's UltraScience Net [7] or the National Lambda Rail [15]), such guarantees must be provided across shared network infrastructures, where dynamic network behavior and multiple available network paths make it imperative for middleware to assist end user applications in best utilizing available network resources. More specifically, when transporting and manipulating their data, applications should receive specific guarantees from the overlay networks used by middleware, accommodating dynamic variations in network behavior: (1) high performance applications require consistent levels of end-to-end performance, such as limited delays for online collaboration [33] or small jitter for multimedia [8], (2) enterprise applications couple data transport and manipulation with application-level expressions of utility or cost [14], and (3) many application classes can utilize guarantees that differentiate across different traffic types, such as offering stronger guarantees for control vs. data traffic in remote instrument access for high performance codes [24].

Previous work on middleware for data-intensive distributed applications has addressed limitations and runtime variations in network bandwidth with adaptive approaches to matching desired to available network resources. Examples include dynamically adjusting data transfer rates [4], varying compression levels in response to monitored changes in

network bandwidth [32], or changing the nature of the data being sent [4, 13, 33]. Other research has sought to use alternative network connections or new network infrastructures to compensate for problematic connection behaviors [24].

This paper presents the IQ-Paths approach to self-regulating high performance data streaming with defined quality requirements across wide area networks. IQ-Paths offers novel functionality that enhances and complements existing adaptive data streaming techniques. First, IQ-Paths dynamically measures [19] and then, also predicts the available bandwidth profiles on network links. Second, it extends such online monitoring and prediction to the multi-link paths in the overlay networks used by modern applications and middleware. Third, it offers automated methods for moving data traffic across overlay paths. These include splitting a single data stream across multiple paths to improve performance through parallelism and to improve desired end-to-end behavior by dynamically differentiating the amounts and kinds of data traffic imposed onto different paths. Such self-regulating data movement and differentiation utilizes a dynamic packet scheduling algorithm that automatically maps packets to paths to match application-level utility specifications. Finally, an important attribute of IQ-Paths is that unlike other methods for bandwidth prediction based on measurements of average bandwidth, it uses statistical techniques to capture the dynamic or noisy nature of available network bandwidth across overlay paths. This enables it to better map data with different desired utility – service guarantees – to the underlying best effort network infrastructure.

Our research uses IQ-Paths for both scientific and high end media applications. In the scientific domain, real-time remote data visualization for a molecular dynamics (MD) code benefits from IQ-Paths’ ability to better meet its dynamic end user requirements. A specific example is to differentiate the transport of certain elements of the application’s data streams, as with the atoms vs. bond forces visually depicted for each timestep of the MD application. Another example is to use network paths with more stable bandwidths for the critical ‘control’ traffic in the remote visualization software and also for the most time-sensitive data sets in large volume parallel data transfers. Stability is dynamically diagnosed and predicted via the aforementioned statistical techniques. In the multimedia domain, IQ-Paths is shown to deliver improved performance for different encoding levels of MPEG-4 video streams [5].

Results in Section 6 also demonstrate the advantages derived from IQ-Paths’ statistical guarantees. Specifically, there are distinct improvements over earlier work on adaptive methods that provide QoS over wide-area networks by predicting future average network behavior from past history [17]. With such methods, quantities like RTT can be predicted well, but average available bandwidth or packet loss rate are not easily captured (e.g., using predictors

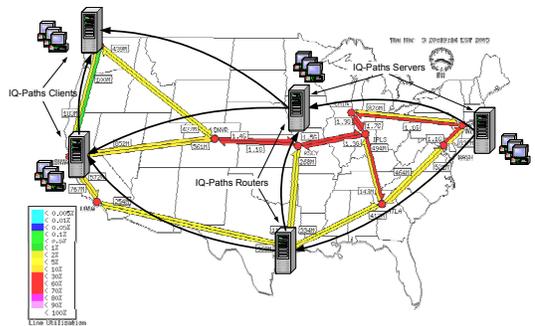


Figure 1: IQ-Paths overlay network: servers, routers, and clients continually assess the qualities of their logical links, admit and map data streams with different desired utility using a self-regulating packet routing and scheduling algorithm.

like MA, AR, or more elaborate methods like ARMA and ARIMA) [34]). This is because noise is a large portion of the signal in the time series of available bandwidth or packet loss rate. As a result, the values for predicted average bandwidths will have large prediction errors. For example, the results reported in [34], based on measurements at over 49 well-connected academic and research institutions, have prediction errors larger than 20% for more than 40% of the predicted values (i.e., $|predictedvalue/actualvalue| > 1.2$), and for 10% of the values, prediction error is larger than 50%. In comparison, IQ-Paths can provide an application with strong guarantees, stating that it will receive its required bandwidth 99% of the time or experience a deadline miss rate (i.e., jitter) of less than 0.1%, for example. Finally, other methods apply low frequency filters [6] to measured values, to reduce prediction error, but unfortunately, this means that they essentially eliminate the noisy nature of (i.e., dynamic variations experienced over) certain network paths. The outcome is that applications cannot adjust to or deal with such variations, by mapping less important or less delay-sensitive data to noisier connections, for example.

Figure 1 illustrates an example of an IQ-Paths overlay, which utilizes automatic network resource profiling, admission control, and self-regulating data routing and scheduling to guarantee different streams’ desired utility requirements. The overlay implemented by IQ-Paths has multiple layers of abstraction. First, its *middleware underlay* – a middleware extension of the network underlay proposed in [21]) – implements the execution layer for overlay services. The underlay is comprised of processes running on the machines available to IQ-paths, connected by logical links and/or via intermediate processes acting as router nodes. Second, underlay nodes continually assess the qualities of their logical links as well as the available resources of the machines on which they reside. The service guarantees provided to applications are based on such dynamic resource measurements, on runtime admission control, resource mapping, and on a self-regulating packet routing and scheduling algorithm. This algorithm, termed PGOS (Predictive Guarantee Overlay Scheduling), provides

probabilistic guarantees for the available bandwidth, packet loss rate, and RTT attainable across the best-effort network links in the underlay.

Key technical advantages of IQ-Paths and its PGOS algorithm include the following:

- *Probabilistic and ‘violation bound’ guarantees:* since the PGOS algorithm uses bandwidth distribution analysis and prediction to capture network dynamics, it can make service guarantees and provide prediction accuracies superior to those provided by prediction methods based on average network behavior: (1) it can ensure that applications receive the bandwidths they require with high levels of assurance (e.g., it can guarantee that an application receives its required bandwidth 99% of the time or that its deadline miss rate is less than 0.1%); (2) in addition, PGOS can also provide deadline violation guarantees that bound the average number of packets that miss their guaranteed QoS (e.g., their deadlines).
- *Reduced jitter:* by reducing jitter in applications like remote data acquisition or display, buffering needs are reduced. This is particularly important for high volume data transfers in time- or delay-sensitive applications.
- *Differentiated streaming services:* different streams can receive different levels of guarantees. As a result, when applications use close to the total available bandwidths of all overlay paths, PGOS can ensure that high utility streams receive stronger service guarantees than others.
- *Full bandwidth utilization:* providing guarantees does not imply sacrificing the bandwidths available to applications (e.g., by purposely under-utilizing some link). Instead, PGOS has sufficiently low runtime overheads to satisfy the needs of even high bandwidth wide area network links.

The remainder of this paper is organized as follows. The next section describes related work, to provide a better perspective on the technical contributions of the IQ-Paths approach. We then outline the software architecture of IQ-Paths, followed by descriptions of its bandwidth prediction methods and of the PGOS algorithm using these methods. Experimental evaluations on an emulated network testbed appear before the paper’s conclusions.

2 Related Work

While the PGOS packet scheduling algorithm is inspired by the DWCS packet scheduling algorithm [31], its use for efficient multimedia data streaming across the Internet leverages substantial prior work on improving the quality of network video streaming [13, 25]. Improvements like those in [25] use TCP-friendly control mechanisms to react to congestion on shorter timescales, with mismatches between the

two timescales absorbed by buffering at the receiver. The control mechanisms used for multimedia data streaming are based on an adaptive layered video streaming algorithm for MPEG-4 with limited buffer size described in in [13], where priorities are used in the VOP (video object plane) to select or discard each VOP element based on average bandwidth prediction, to control the fashion in which fine-grain scalable coding allocates bandwidth to different encoding layers. The contributions of IQ-Paths in this context are its use of statistical bandwidth measurement and prediction to capture network link qualities, and its PGOS self-regulating data routing and scheduling algorithm can utilize both multiple or alternate overlay paths to satisfy different video layers’ utility requirements. The outcome is improved smoothness of video playback, despite the variable-bit-rate nature of layered video.

OverQoS [28] describes the general idea of using overlays and admission control to deliver video across the Internet. In OverQoS, performance gains are achieved by FEC (Forward Error Correction) and conditional packet retransmission in the form of ARQs (Automatic Repeat reQuests). Bandwidth less than the total available bandwidth can be achieved for a subset of the OverQoS flows, with high probability, but potentially at the expense of other flows. In contrast, the PGOS algorithm controls path usage with a more general abstraction that is able to provide statistical guarantees for both single and multiple streams across both single and multiple paths across the overlay.

Both IQ-Paths and OverQoS assume that overlay routing nodes can be placed such that the paths between different pairs of routing nodes do not share common bottlenecks. In practice, such placements require knowledge of the network, by using methods of detecting shared congestion across flows [26], or by using more direct ways of detecting network topologies [27]. Further, the implementation of IQ-Paths could take advantage of underlays [21], which is a general framework that extracts and aggregates topology and other network information from the underlying Internet.

Our general approach of using overlay networks to adapt to network dynamics is shown feasible in [6], which compares the performance of an End System Multicast architecture to that of IP Multicast. The paper also notes that noisy link measurements coupled with aggressive adaptation can cause overlay instability, while conservative adaptations may experience low performance. The proposed solution is to use exponential smoothing to capture the long term performance of a link, thereby distinguishing persistent from temporary changes. Our approach differs in that it exploits knowledge about noise rather than suppressing it, for example, by mapping critical data flows to less noisy links.

While our work complements research on process/job scheduling for Grid services [30]), Data Grid or similarly data-intensive applications in particular must harness both

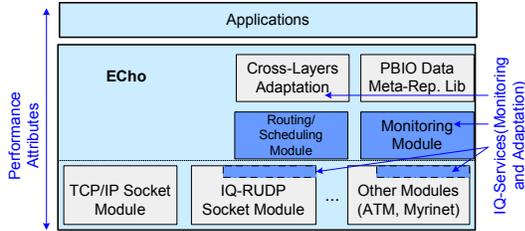


Figure 2: Middleware Architecture.

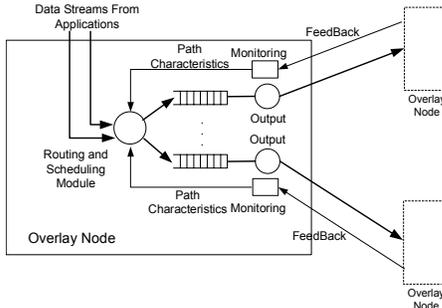


Figure 3: Structure of IQ-Paths Overlay Node.

computational and network resources for their distributed operation [23]. Algorithms like PGOS can provide valuable assistance in these contexts, by integrating it, for instance, into the Dataset Scheduler framework described in [23]. To demonstrate the importance our methods, we have extended the popular GridFTP package [1] with the PGOS routing and scheduling algorithm, to better control how parallel data streams with different service-level objectives are scheduled across multiple network links.

A basic contribution of the PGOS algorithm is its ability to predict future network behavior. [34] points out the difficulty of predicting bandwidth in wide area networks, studying the likelihood of observed bandwidth remaining in a region for which the ratio between the maximum and minimum observed values is less than a factor of ρ . We adopt a similar approach, assuming that it is difficult to predict the exact value of throughput in the next time interval (e.g., in the next second) and instead, providing statistical guarantees for predicting the distribution of throughput in the near future. Interestingly, as shown in [24], it is easier to make guarantees about RTT. Finally, we also leverage the substantial research on measuring available bandwidth described in [12, 19]. Of specific relevance to this paper is recent work presenting more accurate metrics and algorithms to measure the variation of end-to-end available bandwidth [20].

3 Software Architecture of the IQ-Paths Middleware

The software architecture of the IQ-Paths middleware is depicted in Figure 2. It is derived from our substantial experiences with the IQ-Echo [4] high performance publish/subscribe infrastructure implementing channel-based information subscriptions. IQ-Paths leverages IQ-Echo’s sup-

port for multiple transport protocols (e.g., TCP, RUDP, SCTP) and its monitoring modules for measuring desired network metrics from middleware and in cooperation with certain transport modules (e.g., RUDP). PGOS routing/scheduling module aggregates such runtime measurements in order to schedule application packets across multiple overlay paths. Unlike ECho, however, IQ-Paths is realized at a layer ‘below’ the publish/subscribe model of communication. Namely, IQ-Paths manipulates arbitrary application-level messages flowing from data sources to data sinks. Whether such messages are described as pub/sub events or in other forms is immaterial to the research described here. Similarly, IQ-Paths is not concerned with how source-to-sink links are established. It supports both direct source-to-sink links and more complex linkages that utilize overlay networks to route messages and process them ‘in-flight’ on their paths from sources to sinks. One way for end users to establish such linkages is via IQ-Echo’s ‘derived channel’ abstraction. Another way is to use the deployment features implemented as part of the ‘in-transit’ information flow infrastructure described in [14]. A third way is to directly use IQ-Paths as the transport layer for applications, as with the IQ^{PG}-GridFTP implementation used in the evaluation section of this paper.

The goal of IQ-Paths is to provide a general framework for routing, scheduling, and processing streams of application-level messages. Generality is established by layering IQ-Paths ‘beneath’ the different messaging models used by end users, including the IQ-Echo and in-transit models developed in our own research. A specific example is the IQ^{PG}-GridFTP described in this paper, which (1) replaces its transport level with IQ-Paths and (2) interposes the IQ-Paths message routing and scheduling algorithm between GridFTP’s parallel link layer and lower level message transports. As a result, IQ-GridFTP (1) retains its ability to exploit parallelism in data transport by simultaneously using multiple network links, while more importantly, (2) gaining the ability to adjust the volumes of data being transferred to the current behavior of each single network link between source and sink, and (3) using overlay paths and path bandwidth-sensitive message routing and scheduling to better control how data is streamed across multiple links from source to sink.

Important components of the IQ-Paths middleware described in this paper are its Statistical Monitoring techniques and its Routing/Scheduling algorithms. Figure 3 illustrates the structure of each IQ-Paths overlay node and the dynamic interactions of these software components. Specifically, the Statistical Monitoring component monitors the bandwidth characteristics (i.e., bandwidth distribution) of each overlay path and shares this information with the Routing/Scheduling component. The latter routes applications’ data streams and sub-streams to the appropriate overlay paths and in addition,

for each path, it schedules the data packets mapped to it. The goal, of course, is to route and schedule application-level messages to continuously match the network loads imposed by the middleware to the available network bandwidths present in overlay paths, such that application-level metrics of stream utility are met (e.g., probabilistic guarantees on the timeliness of data delivery).

The remainder of this paper ignores other components of the IQ-Paths middleware, referring the reader to a more complete description of the system in [4]. We next describe the manner in which bandwidth guarantees are attained.

4 Statistical Bandwidth Prediction and Guarantees

The PGOS algorithm presented in Section 5 provides to end users predictive guarantees that with some large probability, application-level messages will experience certain levels of bandwidth across certain overlay paths. Toward this end, for each overlay path, IQ-Paths network monitoring (1) tracks the past *distribution* of path bandwidth in the form of cumulative distribution function(CDF), and (2) uses the percentile points in that distribution as the bandwidth predictor, instead of using average bandwidth. The PGOS algorithm then uses these predictions to judiciously route and schedule streams across overlay paths by finding the path(or set of paths), which, for some large value of P_0 , can ensure that the bandwidth allocated to the stream has the property of $P(bw \geq bw_0) \geq P_0$, where bw_0 is the required bandwidth.

For each specific overlay path, frequent bandwidth variation makes it difficult to predict the exact values of average available bandwidth in the near future, both for very short timescales like milliseconds and for the second timescales at which IQ-Paths operates. Statistical prediction is to leverage rather than suppress such variations, in order to provide to applications higher bandwidth guarantees. In other words, while predicting the exact value of future bandwidth is hard, statistical prediction relaxes the prediction requirement by asking if we can obtain certain amount of bandwidth with high probability. Because of the IID nature of available bandwidth, statistical prediction has much smaller prediction error than average bandwidth prediction. Furthermore, statistical prediction also retains more information including the variation and distribution of the signal, which is directly related with the service-level objectives of many applications.

Figure 4 illustrates the results of predicting average bandwidths vs. the statistical predictions used in our approach. Here, we analyze more than 8GB of IP header trace files from the National Laboratory for Applied Network Research, collected at a number of locations of the Abilene (Internet2) and the Auckland networks. The mean prediction error is the average relative error ($|(predicted\ value - actual\ value)/actual\ value|$) of several widely used average bandwidth predictors (i.e., MA,

EWMA and SMA). From Figure 4, the common average bandwidth predictors have a roughly 20% of prediction error. Similar error ranges are also reported in [34]. In contrast, our statistical prediction method (percentile prediction) achieves less than a 4% prediction failure rate. The percentile prediction failure rate is the number of prediction failures divided by the total number of predictions. For these experiments, we first calculate the distribution of N (e.g., 500 and 1000) samples, where each sample is the bandwidth measured in 0.1 to 1 second. Then, since we are particularly interested in whether a path can guarantee certain throughput for 90% of the time (or for 80%, 70%, etc), we find distribution D 's 10th percentile as X (Mbps), and test whether the next n ($n=5$ to 10) samples are larger than X . If they are, a successful prediction occurs, and if not, a prediction failure occurs.

From these experiments and for representative distributed applications, we determine two facts. First, in practice, an application is typically more interested in whether it can receive its required bandwidth consistently, than in the exact value of the bandwidth the network provides. This is precisely the question answered by our statistical bandwidth prediction methods. Second, such statistical guarantees are easier to make than guarantees about available average bandwidth, because the majority of available bandwidth or maximum throughput on Internet paths is IID [34]. As a result, the exact value of average bandwidth in the near future is hard to predict, but the statistical structure of bandwidth can be predicted well. Simply speaking, if in the last 5 mins., the 10th percentile of bandwidth is 10Mbps, then with a large probability, the bandwidth in the next 1 second will be higher than 10Mbps. The measured low prediction failure rate directly justifies our usage of percentile prediction. For high-performance computing, typically large amounts of data are transferred in some small time interval (e.g. 12.5 MB of data in 0.1 second on a 1Gbps link, or 125MB on a 10Gbps link). Mean prediction time series have a large number of outliers because of the aforementioned IID nature, if the measurement interval is small(e.g., 0.1 second). To avoid these outliers, one has to average the bandwidth over larger interval (e.g., 1 second), which loses the variation (or stability) nature of a particular network and could easily cause inappropriate data movement as illustrated in the next experiment. Statistical prediction arguably avoids this pitfall by profiling the bandwidth distribution and providing accurate prediction without lossy averaging. The outcome is a solid foundation for controlling how large amounts of data are moved around in every small time interval with high confidence levels.

5 The PGOS Overlay Path Guarantee and Scheduling/Routing Algorithm

This section describes the Predictive Guarantee Overlay Scheduling (PGOS) algorithm, first discussing the general al-

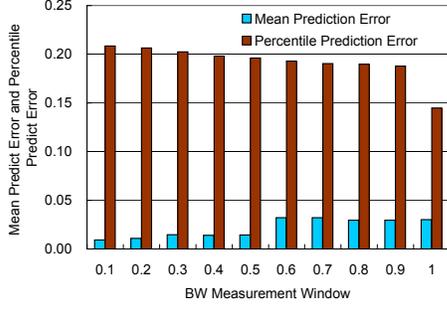


Figure 4: Bandwidth Prediction.

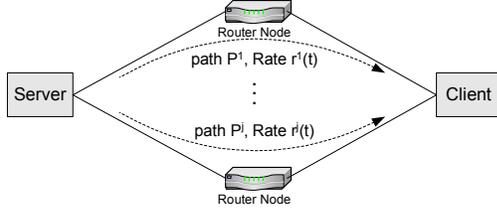


Figure 5: Overlay Routing and Scheduling Algorithm Framework.

gorithm framework, then clarifying the concept of predictive guarantees and describing the algorithm itself. Formal analysis of buffer size under PGOS is given in [5], which shows PGOS also reduces the server/client buffer size requirement and make data transfer less bursty, by using statistical prediction, as compared with using average bandwidth prediction.

5.1 General Framework

An overlay network like the one in Figure 1 may be represented as a graph $G = (V, E)$ with n overlay nodes and m edges. An overlay node may be a server (i.e., data source) running on some host, a client (i.e., data sink), or a daemon for data routing. There may exist multiple distinct paths $P^j, j = 1, 2, \dots, L$, between each server and client, where $P^j = (V^j, E^j)$, $V^j = \{v_0, v_1, \dots, v_k, v_p \neq v_q \text{ if } p \neq q\}$ and $E^j = \{v_0v_1, \dots, v_{k-1}v_k, \text{ where } v_p v_{p+1} \in E, \text{ for all } 0 \leq p \leq k-1\}$. As in [28], we make no assumptions about the placement of overlay nodes in the network. Rather, we assume that the middleware has determined some suitable placement.

For each overlay link, since network bandwidth varies over time, the service time of each application-level message is not known a priori and varies over time. The specific problem addressed by the PGOS algorithm is further illustrated in Figure 5, where multiple streams $S^j, j = 1, 2, \dots, N$ must be transmitted from Server s to Client c with ‘best’ predictive performance guarantees. Figure 6 illustrates a server that deliver multiple streams (in Queue 1, 2, ...) to a client via overlay paths 1, 2, etc. In this model, there is one scheduler and L path services (each service corresponds to one overlay path used to deliver packets, with service rate $r^j(t)$).

Applications specify stream utility in terms of the minimum bandwidths they require, or using Window-Constraints [31] requirement. A Window-Constraint is spec-

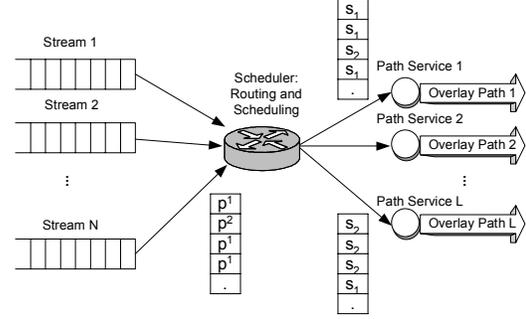


Figure 6: Routing and Scheduling on the Server.

ified by the values x_i , and y_i , where y_i is the number of consecutive packet arrivals from stream S_i for every fixed window, and x_i is the minimum number of packets in the same stream that must be serviced in the window. The dynamics of the underlying network make it difficult to satisfy the minimum bandwidth guarantees required by the utility specifications described above, including for the guarantees associated with each scheduling window t_w . We address this issue by asking applications to specify additional requirements of the following nature: ensure that the minimum bandwidth is met with some large probability P (e.g. 95%, 99%). This also means that 95% of the time, the window constraint will be satisfied. Given these specifications, assuming a packet size of s , and denoting the available bandwidth over a given path by $b^j(t)$, or simply b , (1) the available bandwidth distribution is described as the cumulative distribution function $F^j(b) = P\{avail.bw \in (0, b)\}$, and (2) the service rate of the path service j is described as $r^j = r^j(t)$, where r^j varies over time.

5.2 Predictive Guarantee Overlay Scheduling/Routing Algorithm

The Predictive Guarantee Overlay Scheduling/Routing Algorithm (PGOS) supports two types of guarantees for stream utility specifications: probabilistic and ‘violation bounded’. The former states that with some large probability P , stream S_i will receive the required bandwidth on the selected path. It also means that the stream S_i will receive the required bandwidth for at least $100P\%$ of the time. The latter states that the average number of packets that miss their constraint during each scheduling window can be bounded. Here, we first define a single path selection algorithm for predictive guarantees and then extend it to a scheduling algorithm that operates across multiple overlay paths.

5.2.1 Single path guarantee

The idea of single path selection is to choose the best path among all candidate paths for stream S_i , with some desired guarantee. Single path selection is important because there exist streams that are not easily mapped across multiple paths, an example being a stream with tight deadline/bandwidth requirements which would have to cope

Table 1: Precedence among packets in different streams.

	Packet Ordering
1.	pkts scheduled on current path.
2.	pkts scheduled on other path:
2.1	Earliest deadline first.
2.2	Equal deadlines, highest window constraint first.
3.	pkts not scheduled:
3.1	Earliest deadline first.
3.2	Equal deadlines, highest window constraint first.

with synchronization issues and out of order arrivals when mapped across multiple paths.

Probabilistic guarantee The following is the probabilistic guarantee provided by the PGOS algorithm. Due to limit of space, proofs of Lemmas and Theorem appear in [5]:

Lemma 1 *Suppose during time $(t, t + t_w)$, where t_w is the length of the scheduling window, the available bandwidth distribution of server j is $F^j(b^j)$. Then, with probability $P = 1 - F^j(x_i s / t_w)$, it is guaranteed that x_i packets will be served during the scheduling window t_w .*

Note that this guarantee essentially bounds the probability of insufficient throughput by $F^j(x_i s / t_w)$

‘Violation bound’ guarantees Another useful application-level utility specification is to bound some violation, such as the deadline miss rate. The following is the deadline ‘violation bound’ guarantee provided by PGOS, where Z is the number of packets that miss their deadlines during one scheduling window, given the rate distribution $G^j(r^j)$ in this scheduling window:

Lemma 2 *Given available bandwidth distribution $F^j(b)$, $E[Z]$ is bounded by $x_i \cdot F^j(b_0) - \frac{t_w}{s} \cdot M[b_0]$, where b_0 is the required bandwidth of Stream S_i , $b_0 = x_i s / t_w$, and $M[b_0]$ is the mean of b for all $b \leq b_0$. Both $F^j(b_0)$ and $M[b_0]$ can be easily computed from the available bandwidth distribution.*

5.2.2 Guarantees for multiple overlay paths

By combining the properties of multiple paths, PGOS can provide better guarantees to applications than those achievable on single paths. This is particularly relevant to large data transfers, where the parallelism achievable across multiple paths can be used to speed up data transfers as well as desired ‘in flight’ processing. Based on the two types of guarantees developed for each a single path, we now describe an overlay routing and scheduling algorithm that maps multiple streams across multiple paths (Figure 5). The algorithm schedules all packets of streams $S_i, i = 1, 2, \dots, N$ such that the best guarantee is provided for the timely delivery of high utility streams, while other streams are delivered with less stringent guarantees. The PGOS algorithm, therefore, consists of two parts: (1) utility-based resource mapping and (2) path routing and packet scheduling.

```

1  updateCDF(); /*update CDF using bandwidth/lossrate
   measurement in last scheduling window*/
2  if(previous scheduling vectors doesn't satisfy currentCDF){
   /*when new stream joins or CDF changes dramatically*/
3   Find best scheduling share  $Tp_i^j$ ; /* $Tp_i^j$  is the number of
   packets of stream  $i$  scheduled to be sent on path  $j$ */
   /*now rebuild scheduling vectors:*/
4   for( $i = 1; i \leq N; i++$ ){
5     for( $j = 1; j \leq L; j++$ ){
6        $Tp_i^j += Tp_i^j$ ; /*for path lookup vector*/
       /*Insert deadlines corresponding to  $Tp_i(j)$  into  $VD^j$ */
7       UpdatePathDeadlineVector( $VD^j, Tp_i^j$ ).
8     }
9   }
   /*build path lookup vector*/
10   $VP = \text{PathSchedVector}(Tp^j)$ ;
   /*convert deadlines to stream scheduling vector*/
11   $VS[] = \text{StreamSchedVector}(VD[])$ ;
   } /*end of scheduling vectors update(when necessary)*/

12 while(in current scheduling window){
13   /*get next path according to  $V^p$ */
   path=GetNextFreePath( $V^p$ );
   /* get next packet to send based on  $V_s[p]$ */
14   if(getNextScheduledpkt( $V_s[p]$ ))
15     sendpkt(path, pkt);
16   else if(pkt=getNextUnscheduledPkt( $V_s$ )){
   /*other unscheduled pkt. Precedence rule 2 and 3.*/
17     sendpkt(path, pkt);
   }
}

```

Figure 7: Scheduling Algorithm.

Utility-based Resource Mapping The resource mapping part of the PGOS scheduling algorithm (see Figure 7) finds the best proportion of stream S_i to be delivered via path P^j (**resource mapping**). The result is the generation of a *scheduling vector*, which is then used for routing and scheduling stream packets across multiple paths. The resource mapping step is executed when a new stream joins (or an existing stream terminates) or when the CDF of some path changes dramatically. A single resource mapping typically persists across many scheduling windows.

During each scheduling window, PGOS schedules packets based on the current scheduling vector and the stream precedence listed in Table 1. This table maintains the statistically optimal stream division scheme, while also utilizing additional available bandwidth whenever possible. For example, given two overlay paths’ available bandwidth distribution $G^j(j = 1, 2)$, and two streams $S_i(i = 1, 2)$, the table is set up to divide each stream S_i into two sub-streams $S_i^1 + S_i^2$, where S_i^1 will be sent via path 1 and S_i^2 will be sent via path 2, such that their required performance guarantees are met. Note that S_i^j could be a null sub-stream, if necessary. We will send S_1^1 and S_2^1 via path 1, and send S_1^2 and S_2^2 through path 2.

Stream precedence is determined by the probabilities with which different streams' bandwidth requirements must be met. If streams S_i desire to receive their required bandwidths $100P_i\%$ of the time, then PGOS first finds the path that can satisfy the requirement of the most important stream (with highest P_i), then finds the path for the second most important stream, and so on. If there does not exist a single path that can satisfy stream S_i 's requirement, then the stream S_i is divided into multiple parts S_i^j if this can satisfy stream S_i 's requirement. If this still fails due to limited bandwidth, an upcall is made to inform the application that it is not possible to schedule this particular stream. The application can reduce its bandwidth requirement (e.g., from 95% to 90%) or try to adjust its behavior to the limited available bandwidth [4].

When a deadline violation bound guarantee is desired, PGOS works in a fashion similar to the probabilistic guarantees described above. When one or multiple streams join, PGOS begins with new stream with the highest deadline guarantee (i.e., with $Minimum[E[Z_i]]$), and attempts to find a path to meet its guarantee. If such a path does not exist, PGOS divides stream S_i (with x_i packets) into multiple parts S_i^j (with x_i^j packets) such that $\sum_{j=1}^P E[Z_i^j] \frac{x_i^j}{x_i} \leq E[Z_i]$, where $x_i = \sum_{j=1}^P x_i^j$ and $x^j = \sum_{i=1}^N x_i^j$. An alternative approach is to find a feasible division scheme without considering the ordering of $E[Z_i]$ and solve a mixed integer linear programming problem (MILP). However, this is not desirable since it may divide some important stream (e.g., a control stream) into multiple sub-streams, thereby causing synchronization and delays due to potential packet re-ordering across multiple overlay paths. It is also an N-P hard problem. A detailed analysis of alternative approaches and their comparison are beyond the scope of this paper.

Path Routing and Packet Scheduling While it may be computationally complex to find the best possible resource mapping, a more important issue is the affect of complex mappings on PGOS fast path performance. Here, during each scheduling window, PGOS needs to schedule packets according to the resource mappings encoded in scheduling vectors and according to the precedence table (see Table 1). The efficient data structures used by PGOS are depicted in Figure 7: the scheduler has a path routing vector VP , and each path service has one stream scheduling vector VS . The scheduling vector V encodes the currently best resource mapping scheme derived by the resource mapping step. The lookup vector VP is the vector the scheduler uses to switch between the different overlay paths. As derived in the resource mapping step, path j is assigned x^j packets, so path j is assigned x^j virtual deadlines $D_p[k] = t_w/x^j \cdot (k - 1)$. Virtual deadlines are used to maintain the desired resource mapping proportion. That is, VP contains the ordering to be used for visiting each path, based on virtual deadlines.

To illustrate, consider a concrete example with two streams and two overlay paths. Stream S_1 has 5 pack-

ets in one scheduling window that are mapped to path 1. Stream S_2 has 10 packets in one schedule window, where 4 of them are mapped to path 1, while another 6 packets are mapped to path 2. In this example, path 1 has 9 packets to deliver, and path 2 has 6 packets to deliver. Thus, $VP=[1,2,1,2,1,1,2,1,2,1,1,2,1,2,1]$. When the scheduler switches between the overlay paths, the path lookup vector ensures that three fifths of the time, it will visit path 1, and two fifths of the time, it will visit path 2. Stated more generally, when the scheduler visits path j , it uses the stream scheduling lookup vectors VS^j to select the streams to which to send packets. (VS is essentially a lookup table where each row corresponds to one path). The lookup vectors VS^j are based on the deadlines of all of the packets (from multiple streams) to be sent on path j . In the example, path 1 has nine packets, and the deadlines of these 9 packets are for $S_1, S_2, S_1, S_2, S_1, S_2, S_1, S_2,$ and S_1 respectively. Thus, $VS^1 = [1, 2, 1, 2, 1, 2, 1, 2, 1, 2, 1]$.

While using VS for path mapping, PGOS schedules packets based on both VP and VS . That is, once it has selected path j , PGOS sends packets over it according to VS . Specifically, it selects a packet to send based on the stream scheduling lookup vector VS^j and the precedence table (Table 1). First, it sends the packet scheduled on the current path j some other path that has the earliest deadline. Equal deadlines are broken by the window constraint x/y (highest window constraint first) and further ties are broken arbitrarily. When all scheduled packets have been sent out and there are still free paths to utilize, PGOS sends out other unscheduled packets according to their deadlines and window constraints. Whenever a path is blocked, the scheduler switches to the next path immediately, in order to best utilize other available resources. Because of the high cost of blocking, timeouts and exponential backoff are used to avoid sending multiple packets to a blocked path.

The following theorem states more precisely the guarantees provided by PGOS:

Theorem 1 *If there is a feasible schedule for PGOS to deliver streams $S_i, i = 1, 2, \dots, N$ over paths $P^j, j = 1, 2, \dots, L$ during scheduling window $(t, t + t_w)$ with bandwidth guarantees, then stream S_i 's window constraint will be met with probability P_i .*

6 Experimental Evaluation

This section evaluates and analyzes IQ-Paths with two types of applications: (1) the SmartPointer system [33] for distributed collaboration and interactive program steering, and (2) the GridFTP [1], a high-performance and reliable data transfer protocol widely used in the Grid community, for reliable parallel data transmission in wide area networks. Interested readers can refer to a technical report [5] for additional studies which demonstrate (1) the importance of choosing the right path based on statistical predictions, and

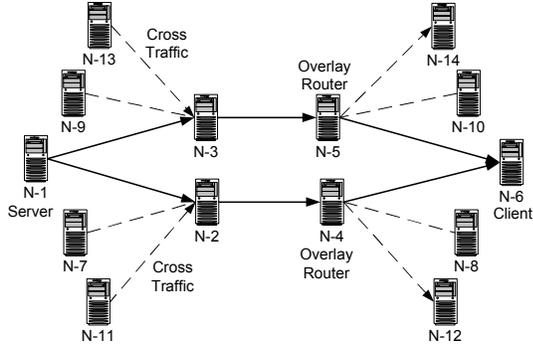


Figure 8: Testbed. The link connecting each pair of nodes is fast ethernet. Cross traffic is injected by Node N-9 to N-14. Overlay routers are placed at Node N-4 and N-5, so that overlay paths and cross traffic paths share the same bottleneck (N-3 to N-5 and N-2 to N-4).

(2) substantially improved service level QoS IQ-Paths offers when applied to MPEG-4 Fine-Grained Scalable video streaming.

Our testbed emulates a realistic wide area setting, using the EmuLab facility [16]. NLNR traces are used to inject representative cross-traffic [9]. If not stated otherwise, the overlay server N-1 has two overlay paths to reach the client N-6 (Figure 8), and the background traffic and data traffic share the common link between N-3 and N-5, and the link between N-2 and N-4. All link capacities are 100Mbps, which is the current up-limit of Emulab.

6.1 PGOS evaluations with SmartPointer

This set of experiments evaluates PGOS’ multi-path message routing and scheduling performance using the SmartPointer distributed collaboration application. The purpose is to see how the algorithm can guarantee some critical stream’s required throughput while also providing high throughput to non-critical streams. Consider the SmartPointer server issuing three streams (Atom, Bond1, and Bond2) to remote clients. Streams Atom and Bond1 are data about all atoms and those bonds that are in the observer’s immediate graphical view volume, whereas stream Bond2 contains the bonds outside the observer’s current view. Therefore, Streams Atom and Bond1 are important and must be delivered in real-time (25 frame/sec) for effective collaboration, but stream Bond2 is less critical (e.g., it may be important when the observer rapidly changes his/her viewing angle.)

Three on-line message transfer algorithms are evaluated and compared to meet this application’s needs: (1) transfer all messages over one single path based on normal Fair Queuing (WFQ), (2) transfer messages over two paths with multi-server Fair Queuing (MSFQ) [2], and (3) transfer messages over two paths using the proposed PGOS routing and scheduling algorithm. The input (utility requirements) to PGOS are 3.249Mbps with 95% predictive guarantee for stream Atom and 22.148Mbps with 95% predictive guaran-

tee for stream Bond1. We also compare these results with a near-optimal off-line algorithm, termed OptSched [5], which assumes that we know available bandwidth a priori. Although this off-line algorithm cannot be used in practice, it can be used to gauge the absolute performance of PGOS.

The results of using these four algorithms appear in Figure 9. Figure 9a depicts the throughput of 3 streams attained by the WFQ algorithm on Path A, which has higher available bandwidth than Path B with larger variance. Multi-Server Fair Queuing (MSFQ) can maintain the proportion of throughput shared by the three streams quite well (see Figure 9b), but because of its inaccurate average bandwidth prediction, it fails to provide the required throughput to the two critical streams Atom and Bond1. Both streams exhibit substantial throughput fluctuation. In comparison, the PGOS algorithm successfully provides very stable throughput to these two critical streams. Furthermore, note that in Figure 9c, the throughput of stream Bond2 is not compromised. This stream is divided by PGOS into two substreams (Bond2-PathA and Bond2-PathB), and the average throughput of stream Bond2 is almost the same as that achieved by MSFQ.

The cumulative distributions of throughput of the three streams under the three algorithms are given in Figure 10. PGOS provides the two critical streams at least 99.5% of their required bandwidth for 95% of the time. MSFQ can only provide about 87% of their required bandwidth for 95% of the time. For example, stream Bond1 requires 22.148Mbps, and the actual 95th percentile of the achieved bandwidth is 22.068Mbps under PGOS, but it is only 19.248Mbps under MSFQ. The standard deviations of bandwidth experienced by the two critical streams appear in Figure 11. Although stream Bond2 has slightly larger standard deviation with PGOS, the two critical streams Atom and Bond1 experience much lower standard deviations.

Both Fair Queuing and Multi-Server Fair Queuing try to allocate bandwidth in a proportional based manner according to predicted bandwidth, but they require exact values of end-to-end bandwidth, which are hard to attain. As a result, although both of these two algorithms can successfully maintain the proportion of the bandwidth allocated to multiple streams, they cannot provide specific bandwidth to a particular stream. In comparison, PGOS relaxes the prediction assumption, only asking if we can obtain certain bandwidth with some high probability. This is not only easier to predict, but also directly provides the functionality needed by applications.

All three algorithms experience certain overheads when routing single streams over multiple paths, because of packet reordering and delays of head-of-line packets. PGOS reduces this overhead by using a single path for one stream whenever possible, especially for streams with higher priorities. Simply speaking, unlike MSFQ which provides the two critical

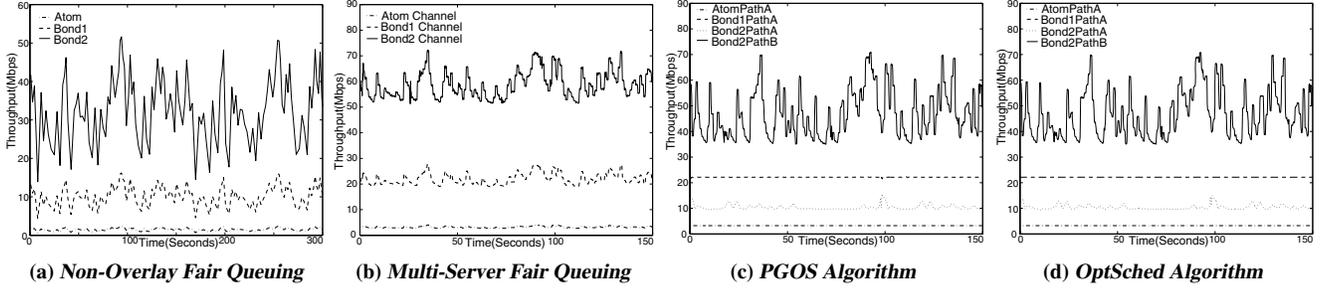


Figure 9: Throughput Time Series Comparison of Three Algorithms.

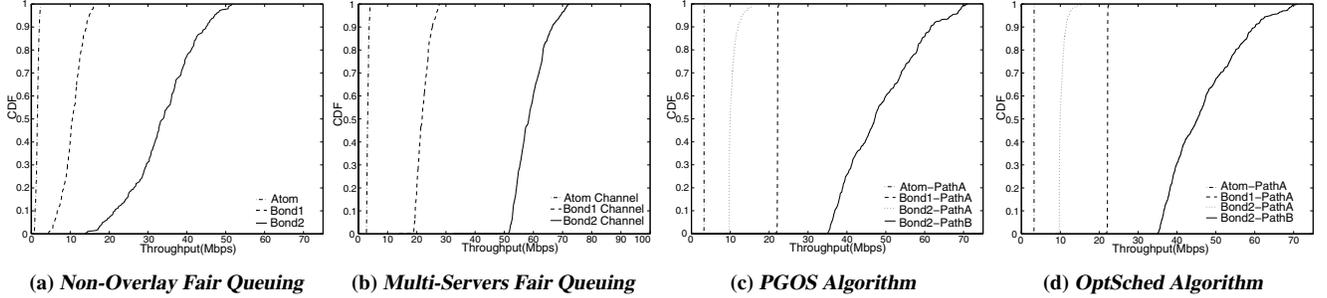


Figure 10: Throughput CDF Comparison of Three Algorithms.

streams less than required bandwidths when the network is congested and more than required bandwidths when the network is free of congestion, PGOS routes and schedules packets such that the two important streams obtain stable required bandwidths no matter whether or not one path is congested. As a result, the application frame jitter is also reduced from 2.0ms (with MSFQ) to 1.4ms (with PGOS).

In summary, these experiments show that with PGOS routing/scheduling, critical streams' required throughput can be guaranteed most of the time. This is done without compromising the average throughput experienced by non-critical streams. A case in point in our experiments is non-critical stream Bond2, which still receives almost the same average throughput under PGOS as under MSFQ.

6.2 GridFTP Experiments

GridFTP [1] is widely accepted as one of the common data transfer services available for high performance applications, with extension to the FTP protocol including parallel data-transfer, SPAS(Striped Passive), and SPOR(Striped Data Port). In this subsection, we present our experiences with IQ^{PG} -GridFTP, which strengthens our previous work [4] by including support for PGOS-enabled parallel file transfers. IQ^{PG} -GridFTP generalizes the publicly available wu-ftp [11] server to support the GridFTP protocol extensions for parallel transfers and implements the Partitioned and Blocked data layout options to distribute file contents across the connections in addition to the PGOS layout. A partitioned data layout is one where contiguous chunks of file are distributed evenly across all the connections for transfer, while a blocked data layout is one where data blocks (each of size block-size) are distributed in a round-robin fashion.

We use a climate database in our experiment as simulation of the Earth System Grid II [10]. Each record in this database has three data components: (1) the numeric data (approximately 172.8KB, denoted by 'DT1'), and (2) and (3) are low resolution images (128KB, denoted by 'DT2') and high resolution images (384KB, denoted by 'DT3'), respectively. GridFTP and IQ^{PG} -GridFTP are configured to concurrently transfer file records over two overlay paths. For such transfers, we want to ensure that the numeric data and low resolution images receive their required bandwidths of at least 25 records/second for real-time data streaming. In addition, we also want to fully utilize bandwidth to transfer high-resolution data.

Experimental results are depicted in Figures 12 and 13. From these measurements, it is apparent that IQ^{PG} -GridFTP can ensure that the streams DT1 and DT2 receive their required bandwidths consistently, while stream DT3 is transferred as fast as possible. In comparison, standard GridFTP splits the dataset into blocks allocated to the multiple connections for transfer, but when the available bandwidth of any path is low, all types of data have to compete with each other. This causes the important data streams to not receive their required bandwidths during these periods. Quantitatively, stream DT1 achieves 33.94Mbps average throughput using GridFTP with a large standard deviation (1.4297), while using IQ^{PG} -GridFTP, it achieves 34.55Mbps average throughput with a small standard deviation (0.4040). Similar results are observed for stream DT2. Note that here the network can provide almost the total throughput required by the application. If the network cannot provide such throughput, then the two streams DT1 and DT2 obtain even less bandwidth using GridFTP, as they have to compete with stream DT3 for the

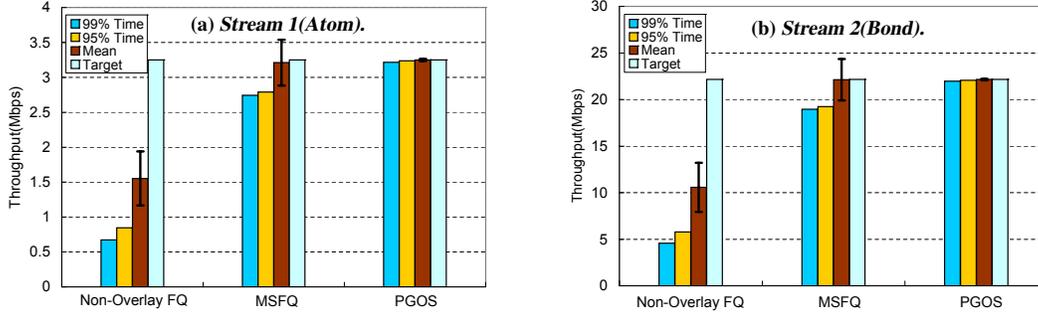
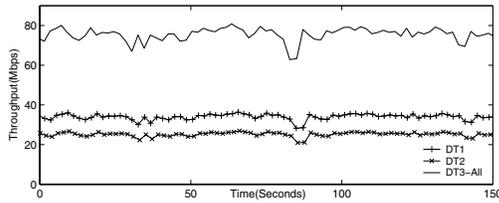
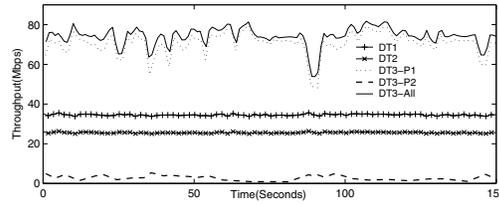


Figure 11: Throughput Achieved by Three Algorithms: Target, Mean, 95% of the time, 99% of the time, and Standard Deviation (represented by the vertical bars).



(a) GridFTP Throughput.



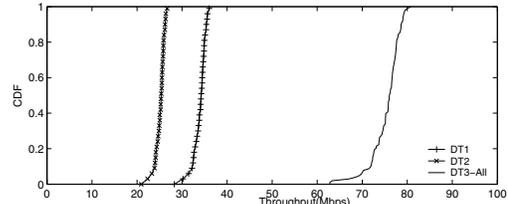
(b) IQ^{PG} -GridFTP Throughput. Line DT3-All is the throughput achieved by stream DT3 (sum of throughput on two paths: DT3-P1 and DT3-P2).

Figure 12: Throughput Achieved by GridFTP and IQ^{PG} -GridFTP same limited bandwidth. In summary, with PGOS, IQ^{PG} -GridFTP can protect more important streams from competing with other less important streams, while also scheduling less important streams to be delivered when there exists sufficient bandwidth. Applications have full control over deciding how much bandwidth will be allocated for a particular stream and what kind of guarantee is for each stream.

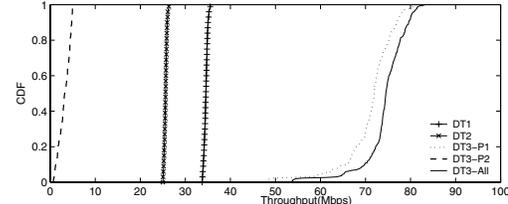
7 Conclusions and Future Work

IQ -Paths is a data streaming middleware that uses overlay networks to better serve the needs of distributed applications running on wide area networks. IQ -Paths employs statistical methods to provide to applications predictive guarantees for the bandwidths available to them from the underlying network, for all paths in the overlay connecting data sources to sinks. In addition, its PGOS scheduling algorithm both suitably routes packets across overlay paths and schedules them across single and multiple (concurrent) paths, coupling parallelism in data transfer with statistical bandwidth guarantees.

The statistical prediction technique used in IQ -Paths not only measures average available bandwidth, but also captures the dynamic or noisy nature of the bandwidth available on overlay paths. As a result, IQ -Paths can provide to appli-



(a) GridFTP Throughput CDF.



(b) IQ^{PG} -GridFTP Throughput CDF.

Figure 13: GridFTP and IQ^{PG} -GridFTP Throughput CDF Comparison

cations both probabilistic and ‘bounded violation’ delivery guarantees. The former state that with some large probability P , stream S_i will receive the required bandwidth on the selected path. The latter state that the average number of messages that violate some constraint (e.g., miss their deadlines) during each scheduling window is bounded.

This paper uses IQ -Paths to meet the needs of three representative distributed applications: (1) the SmartPointer real-time collaboration system, (2) multimedia data streaming, and (3) GridFTP. The integration of IQ -Paths into these applications is facilitated by its design as a ‘model-neutral’ data streaming layer underlying the application-specific communication models offered by higher middleware layers, including the publish/subscribe model used by SmartPointer, the simple data transfer model used by GridFTP, and the data streaming model used by the multimedia application.

Several extensions of the proposed IQ -Paths framework are of future interest. The path characteristics collected by IQ -Paths can benefit a wide range of high performance, multimedia, and enterprise applications. We will experiment with additional applications in our future work. In addition, it would be interesting to extend this work to content delivery systems that use overlay multicast techniques. For enter-

prise applications, our current research is further enhancing this work by developing runtime methods for fault tolerance. Here, an interesting use of IQ-Paths is to differentiate data traffic required for replication from other traffic, perhaps dynamically varying reliability/performance tradeoffs with selective replication techniques. Another interesting topic is to use IQ-Paths to isolate the effects of fault tolerance or recovery traffic from regular data traffic, perhaps to avoid the additional disturbances arising during recovery. Finally, we will consider additional service objectives, such as message loss rate service guarantees.

8 Acknowledgements

We gratefully acknowledge input from Constantinos Dovrolis on the methods for network bandwidth measurement and packet scheduling used and developed for this paper.

References

- [1] B. Allcock, J. Bresnahan, K. Kettimuthu, M. Link, C. Dumitrescu, I. Raicu, and I. Foster. Zebra: The Globus Stripped GridFTP Framework and Server. In *proc. of Super Computing*, 2005.
- [2] J. M. Blanquer and B. Ozden. Fair queuing for aggregated multiple links. In *Proc. of ACM SIGCOMM*, 2001.
- [3] S. Brin and L. Page. The anatomy of a large-scale hypertextual Web search engine. *Computer Networks and ISDN Systems*, 30(1-7):107-117, 1998.
- [4] Z. Cai, G. Eisenhauer, Q. He, V. Kumar, K. Schwan, and M. Wolf. Iq-services: Network-aware middleware for interactive large-data applications. *Concurrency & Computation. Practice and Experience Journal*, 2005.
- [5] Z. Cai, V. Kumar, and K. Schwan. IQ-Paths: Self-regulating Data Streams across Network Overlays. Technical report, Georgia Tech, 2006.
- [6] Y. Chu, S. Rao, S. Seshan, and H. Zhang. Enabling conferencing applications on the internet using an overlay multicast architecture. *Proc. of SIGCOMM*, 2001.
- [7] DOE. UltraScience Net. <http://www.csm.ornl.gov/ultranet/>.
- [8] W.-C. Feng and J. Rexford. Performance evaluation of smoothing algorithms for transmitting prerecorded variable-bit-rate video. *IEEE Trans. on Multimedia*, Sept 1999.
- [9] N. L. for Applied Network Research". Internet measurement and Internet analysis. <http://moat.nlanr.net/>.
- [10] I. Foster and et al. The Earth System Grid II: Turning Climate Datasets Into Community Resources. In *Annual Meeting of the American Meteorological Society*, 2002.
- [11] W. D. Group. WU-FTP. <http://www.wu-ftpd.org>.
- [12] N. Hu and P. Steenkiste. Evaluation and characterization of available bandwidth probing techniques. *IEEE Journal on Selected Areas in Communications*, Aug. 2003.
- [13] T. Kim and M. Ammar. Optimal Quality Adaptation for MPEG-4 Fine-Grained Scalable Video. In *Proc. of IEEE INFOCOM*, Apr. 2003.
- [14] V. Kumar, B. F. Cooper, Z. Cai, G. Eisenhauer, and K. Schwan. Resource-Aware Distributed Stream Management using Dynamic Overlays. In *Proc. of ICDCS*, 2005.
- [15] N. LambdaRail. <http://www.nlr.net/>.
- [16] J. Lepreau and et. al. The Utah Network Testbed. <http://www.emulab.net/>. University of Utah.
- [17] B. Li and K. Nahrstedt. A control-based middleware framework for quality of service adaptations. *IEEE Journal on Selected Areas in Communications*, Sept. 1999.
- [18] G. M. Mair. Telepresence - The Technology and Its Economic and Social Implications. In *Proc. of IEEE International Symposium on Technology and Society*, 1997.
- [19] M.Jain and C. Dovrolis. End-to-End Available Bandwidth: Measurement methodology, Dynamics, and Relation with TCP Throughput. In *Proc. of SIGCOMM*, Aug. 2002.
- [20] M.Jain and C. Dovrolis. End-to-end Estimation of the Available Bandwidth Variation Range. In *Proc. of SIGMETRICS*, June 2005.
- [21] A. Nakao, L. Peterson, and A. Bavierr. A Routing Underlay for Overlay Networks. In *Proc. of ACM SIGCOMM*, 2003.
- [22] NASA. Using XML and Java for Telescope and Instrumentation Control. In *Proc. of SPIE Advanced Telescope and Instrumentation Control Software*, 2000.
- [23] K. Ranganathan and I. Foster. Decoupling Computation and Data Scheduling in Distributed Data-Intensive Applications. In *proc. of Super Computing*, 2002.
- [24] N. S. V. Rao. Overlay Networks of In-Situ Instruments for Probabilistic Guarantees on Message Delays in Wide-Area Networks. *IEEE Journal on Selected Areas of Communications*, 22(1), 2004.
- [25] R. Rejaie, M. Handley, and D. Estrin. Quality Adaptation for Congestion Controlled Playback Video over the Internet. In *Proc. of ACM SIGCOMM*, 1999.
- [26] D. Rubenstein, J. Kurose, and D. Towsley. Detecting shared congestion of flows via end-to-end measurement. *IEEE/ACM Transactions on Networking*, 10(3), June 2002.
- [27] N. Spring, R. Mahajan, and D. Wetherall. Measuring ISP Topologies with Rocketfuel. In *Proc. of ACM SIGCOMM*, 2002.
- [28] L. Subramanian, I. Stoica, H. Balakrishnan, and R. Katz. OverQoS: An Overlay Based Architecture for Enhancing Internet QoS. In *Proc. of 1st Symposium on Networked Systems Design and Implementation(NSDI)*, Mar. 2004.
- [29] N. Taesombut, X. Wu, and A. A. Chien. Collaborative Data Visualization for Earth Sciences with the OptIPuter, 2005.
- [30] T. Tannenbaum, D. Wright, K. Miller, and M. Livny. Condor - a distributed job scheduler. In *Beowulf Cluster Computing with Linux*. MIT Press, 2001.
- [31] R. West and C. Poellabauer. Analysis of a Window-Constrained Scheduler for Real-Time and Best-Effort Packet Streams. In *Proc. of IEEE Real-Time Systems Symposium*, 2000.
- [32] Y. Wiseman and K. Schwan. Efficient End to End Data Exchange Using Configurable Compression. In *Proc. of ICDCS*, Mar. 2004.
- [33] M. Wolf, Z. Cai, W. Huang, and K. Schwan. Smart Pointers: Personalized Scientific Data Portals in Your Hand. In *Proc. of IEEE/ACM Supercomputing Conference*, Nov. 2002.
- [34] Y. Zhang, N. Duffield, V. Paxson, and S. Shenker. On the Constancy of Internet Path Properties. In *Proc. of the ACM SIGCOMM Internet Measurement Workshop*, Nov. 2001.