



NORTHWESTERN UNIVERSITY

Computer Science Department

Technical Report

NWU-CS-02-13

October 15, 2002

Revised: February 29th, 2004

Multiscale Predictability of Network Traffic

Yi Qiao

Jason Skicewicz

Peter Dinda

Abstract

Distributed applications use predictions of network traffic to sustain their performance by adapting their behavior. The timescale of interest is application-dependent and thus it is natural to ask how predictability depends on the resolution, or degree of smoothing, of the network traffic signal. To help answer this question we empirically study the one-step-ahead predictability, measured by the ratio of mean squared error to signal variance, of network traffic at different resolutions. A one-step-ahead prediction at a coarse resolution is a prediction of the average behavior over a long interval. We apply a wide range of linear and nonlinear time series models to a large number of packet traces, generating different resolution views of the traces through two methods: the simple binning approach used by several extant network measurement tools, and by wavelet-based approximations. The wavelet-based approach is a natural way to provide multiscale prediction to applications. We find that predictability seems to be highly situational in practice—it varies widely from trace to trace. Unexpectedly, predictability does not always increase as the signal is smoothed. Half of the time there is a sweet spot at which the ratio is minimized and predictability is clearly the best. Also surprisingly, predictors that can capture nonstationarity and nonlinearity provide benefits only at very coarse resolutions. We conclude by describing plans for an online wavelet-based prediction system.

Effort sponsored by the National Science Foundation under Grants ANI-0093221, ACI-0112891, and EIA-0130869. The NLANR PMA traces are provided to the community by the National Laboratory for Applied Network Research under NSF Cooperative Agreement ANI-9807479. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation (NSF).}

Keywords: network traffic characterization, network traffic prediction, wavelet methods, multiscale methods

An Empirical Study of the Multiscale Predictability of Network Traffic

Yi Qiao
yqiao@cs.northwestern.edu

Jason Skicewicz
jskitz@cs.northwestern.edu

Peter Dinda
pdinda@cs.northwestern.edu

Department of Computer Science, Northwestern University

Abstract—Distributed applications use predictions of network traffic to sustain their performance by adapting their behavior. The timescale of interest is application-dependent and thus it is natural to ask how predictability depends on the resolution, or degree of smoothing, of the network traffic signal. To help answer this question we empirically study the one-step-ahead predictability, measured by the ratio of mean squared error to signal variance, of network traffic at different resolutions. A one-step-ahead prediction at a coarse resolution is a prediction of the average behavior over a long interval. We apply a wide range of linear and nonlinear time series models to a large number of packet traces, generating different resolution views of the traces through two methods: the simple binning approach used by several extant network measurement tools, and by wavelet-based approximations. The wavelet-based approach is a natural way to provide multiscale prediction to applications. We find that predictability seems to be highly situational in practice—it varies widely from trace to trace. Unexpectedly, predictability does not always increase as the signal is smoothed. Half of the time there is a *sweet spot* at which the ratio is minimized and predictability is clearly the best. Also surprisingly, predictors that can capture non-stationarity and nonlinearity provide benefits only at very coarse resolutions. We conclude by describing plans for an online wavelet-based prediction system.

I. INTRODUCTION

The predictability of network traffic is of significant interest in many domains, including adaptive applications [6], [37], congestion control [23], [8], admission control [24], [11], [10], wireless [25], and network management [9]. Our own focus is on providing application-level performance queries to adaptive applications, ranging from fine-grain interactive applications such as immer-

sive audio [27] in local environments to coarse-grain scientific applications on computational grids [18]. For example, an application can ask the Running Time Advisor (RTA) system to predict, as a confidence interval, the running time of a given size task on a particular host [14]. We are trying to develop an analogous Message Transfer Time Advisor (MTTA) that, given two endpoints on an IP network, a message size, and a transport protocol, will return a confidence interval for the transfer time of the message. A key component of such a system is predicting the aggregate background traffic with which the message will have to compete. We model this traffic as a discrete-time *resource signal* representing bandwidth utilization. For example, a router might periodically announce the bandwidth utilization on a link.

The timescale for prediction that a tool like the MTTA is interested in depends on the query posed to it. If the application wants to send a small message, the MTTA requires a short-range prediction of the signal, while for a large message the prediction must be long-range (as is often the case with wide area data transfers [39], [29]). However, the appropriate resolution of the signal varies with the query. A short-range query demands a fine grain resolution while a long-range query can make do with a coarse resolution. Note that a one-step-ahead prediction of a coarse grain resolution signal corresponds to a long-range prediction in time.

To easily support this need for multi-resolution views of resource signals, we have proposed disseminating them using a wavelet domain representation [36]. A sensor would capture a one-dimensional signal at high resolution and apply an N -level streaming wavelet transform to it, generating N signals with exponentially decreasing resolutions and sample rates. Tools like the MTTA would then reconstruct the signal at the resolution they require by using a subset of the signals, consuming a minimal amount of network bandwidth to get an appropriate resolution view of the resource signal. We call this view a *wavelet approximation signal*.

Effort sponsored by the National Science Foundation under Grants ANI-0093221, ACI-0112891, ANI-0301108, EIA-0130869, and EIA-0224449. The NLANR PMA traces are provided to the community by the National Laboratory for Applied Network Research under NSF Cooperative Agreement ANI-9807479. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation (NSF).

In our scheme, the final signal an application receives is in effect an appropriately low-pass filtered version of the original signal. How close the filter is to an ideal low-pass filter depends on the nature of the wavelet basis function that is used. Interestingly, in currently available network monitoring systems like Remos [13] and the Network Weather Service [41] an analogous filtering step occurs in the form of binning. The signal is smoothed by reducing it to averages over non-overlapping bins, producing what we refer to as a *binning approximation signal*. For example, Remos’s SNMP collector periodically queries a router about the number of bytes transferred on an interface and uses the difference between consecutive queries divided by the period as a measurement of the consumed bandwidth.

It is important to understand that wavelet approximation signals *include* binning. A binning approximation signal is equivalent to a wavelet approximation signal constructed using a Haar wavelet. We study binning because it is widely used in network monitoring systems, and we study wavelets because they provide a generalization that may prove to be more appropriate. Wavelet approximation signals tend to be much smoother than binning approximation signals, avoiding abrupt artifacts that are not actually a part of the data. Furthermore, the extent of smoothing can be varied by the choice of the underlying wavelet.

Given the context of the MTTA and these two methods, binning and wavelets, for producing approximations to resource signals that represent network traffic, what is the nature of the predictability of the signals, how does it depend on the degree of approximation, and what does this imply for the MTTA? This paper reports on an empirical study that provides answers to these questions. The study is based on a large number of packet traces collected on WANs and LANs. We studied the predictability of these traces, which cover all of the classes with multiple traces per class.

Our methodology is simple. We generate a very high resolution view of a trace by binning the packets into very small bins. Then we produce an approximation using either the binning approach or the wavelet approach. We fit a linear or nonlinear time series model to the first half of the approximated trace and use it to produce one-step-ahead predictions for the second half. The nonlinear models can refit themselves as they traverse the second half of the trace. As we noted earlier, one-step-ahead predictions are sufficient in the context of the MTTA because as the timescale of prediction increases, the resolution can decrease. Our measure of predictability is the ratio of the mean squared error for the predictions to the variance of the second half of the signal. This is basically the “noise

to signal” ratio of the predictor. The smaller the ratio, the better the predictability.

Our conclusions from this study and their implications are listed below to aid the reader in judging what follows.

- Generalizations about the predictability of network traffic are very difficult to make. Network behavior can change considerably over time and space. Prediction should ideally be adaptive and it must present confidence information to the user.
- Aggregation appears to improve predictability. WAN traffic is generally more predictable than LAN traffic. In this we agree with the results of the earlier studies. The implication is that wide area network prediction systems are likely to be more successful than those in the local area. Happily, they are also more necessary.
- Smoothing often does *not* monotonically increase predictability. About 50% of the long traces in our study exhibit a *sweet spot*, a degree of smoothing at which predictability is maximized, contradicting earlier work. This suggests that there is a “natural” timescale for prediction-driven adaptation.
- There are some differences in the predictability of wavelet-approximated and binning-approximated traces, although they are not large. Both approximation approaches are effective. The implication is that concerns other than predictability will drive the choice between these approaches.
- There clearly are differences in the performance of different predictive models. An autoregressive component is clearly indicated, although it is often also helpful to have a moving average component and an integration. Fractional models, which capture long-range dependence, are effective, but do not warrant their high cost for prediction. Happily, this implies that simple models that can be practically deployed in an online system will be effective.
- The nonlinear models we evaluated generally do not perform better than the linear models until the degree of smoothing is considerable, and even then the performance gain is small. This suggests that modeling the nonstationarity and nonlinear behavior of network traffic is only significant for very coarse grain prediction.

In the following, we begin by summarizing related work. Next, we describe the traces on which our study is based. We then describe our results for predicting binning approximation signals. This is followed by a parallel presentation of our results for predicting wavelet approximation signals. Finally, we conclude by describing the structure of a multi-resolution prediction system.

II. RELATED WORK

We use a wide range of predictive models, including the classical AR, MA, ARMA, and ARIMA models [7], fractional ARIMAs [21], [19], [5], a variation of TAR nonlinear models [38], and simple models such as LAST and a windowed average. Our prediction tools, which are used both for offline studies and in online resource signal prediction systems, are currently publicly available as part of our RPS Toolbox [15]. Our wavelet results use our Tsunami wavelet toolbox, which is summarized in Section V, with more detail elsewhere [35]. Tsunami will be publicly available soon.

The earliest work in predicting network traffic of which we are aware is that of Groschwitz and Polyzos who applied ARIMA models to predict the long-term (years) growth of traffic on the NSFNET backbone [20]. Basu, et al produced the first in-depth study of modeling FDDI, Ethernet LAN, and NSFNET entry/exit point traffic using ARIMA models [4]. As in our binning study, they binned packet traces into non-overlapping bins in order to produce a periodic time series to study. Our trace sets overlap slightly in that they also used the Bellcore LAN traces. Leland, et al demonstrated that Ethernet traffic is self-similar [26], while Willinger, et al suggested a mechanism for this phenomenon [40]. The long-range dependence implied by self-similarity suggests that fractional ARIMA models might be appropriate. On the other hand, You and Chandra found that traffic collected from a campus site exhibited nonstationary and nonlinear properties and studied modeling it using threshold autoregressive (TAR) models [42].

In terms of network prediction, closest to the work described in this paper is that of Sang and Li [34], who analyzed the prospects for multi-step prediction of network traffic using ARMA and MMPP models. Their analysis is based on continuous time ARMA and MMPP models driven by Gaussian noise sources. Making the assumption that such models were appropriate, they then developed analytic expressions for how far into the future prediction was possible before errors would exceed a bound, and for how this limit was affected by traffic aggregation and smoothing of measurements. They found that both aggregation and smoothing monotonically increased predictability. They then fitted such models to several traces and evaluated the resulting models' parameters and noise variances to determine the maximum prediction interval for the traces. Only their WAN traces could be predicted significantly into the future and then only after considerable smoothing. Our work differs in several ways. First, we are approaching this problem from the context of a user

tool like the MTTA. Second, we use a much larger set of traces. Third, we applied nonlinear models as well as linear models. Finally, we find that predictability often does *not* increase monotonically with smoothing.

Researchers have applied wavelet-based techniques to understand network traffic and packet traces for some time. The self-similar nature of network traffic was an important discovery in the early 90s. Abry, et al, have developed wavelet-based techniques to estimate the Hurst parameter, the degree of self-similarity [1]. Feldmann, et al have extensively used wavelets to characterize network traffic as multi-fractal [16] and to study the impact of this property on control mechanisms such as TCP congestion control [17]. Riedi, et al have shown how to use wavelets to synthesize network traffic [32], computing results in an efficient manner that appear to match real Ethernet traces visually and statistically. Our work is the first of which we are aware that empirically studies the *predictability* of wavelet approximations of real network traffic. Several wavelet systems also exist. For example, WIND uses wavelet-based scaling analysis to detect network performance problems [22]. Another relevant system estimates the Hurst parameter at the router to make adaptive changes in congestion control, or to provide up to date information about traffic dynamics without storing all of the data for offline analysis [33].

III. TRACES

Our study is based on the three sets of traces shown in Figure 1. The NLANR set consists of short period packet header traces chosen at random from among those collected by the Passive Measurement and Analysis (PMA) project at the National Laboratory for Applied Network Research (NLANR) [30]. The PMA project consists of a growing number of monitors located at aggregation points within high performance networks such as vBNS and Abilene. Each of the traces is approximately 90 seconds long and consists of WAN traffic packet headers from a particular interface at a particular PMA site. We randomly chose 180 NLANR traces provided by 13 different PMA sites. The traces were collected in the period April 02, 2002 to April 08, 2002. We have developed a hierarchical classification scheme for these traces. The scheme is based largely on the autocorrelative behavior of the traces, which is summarized below. A separate technical report provides much more detail [31]. We identified 12 classes for the NLANR set. For the present study, we worked with 39 of the traces, covering each of the classes we identified.

The AUCKLAND set, which we focus on in detail in this paper, also comes from NLANR's PMA project. These traces are GPS-synchronized IP packet header

Name	Number of			Duration	Range of Resolutions
	Raw Traces	Classes	Studied		
NLANR	180	12	39	90 s	1,2,4,...,1024 ms
AUCKLAND	34	8	34	1 d	0.125, 0.25, 0.5,..., 1024 s
BC	4	n/a	4	1 h, 1 d	7.8125 ms to 16 s
Totals	218	n/a	77	90 s to 1 d	1 ms to 1024 s

Fig. 1. Summary of the trace sets used in the study.




Fig. 2. Signal variance as a function of bin size for the AUCKLAND traces.

traces captured with a DAG3 system at the University of Auckland’s Internet uplink by the WAND research group between February and April 2001. These also represent aggregated WAN traffic, but here the durations for most of the traces are on the order of a whole day (86400 seconds). Our classification approach, also described in the technical report, netted 8 classes here. For the present study, we chose 34 traces, collected from February 20, 2001 to March 10, 2001, which cover the different classes.

The BC set consists of the widely used Bellcore packet traces [26] which are available from the Internet Traffic Archive [3]. There are four traces, which are detailed in the technical report. In summary, two of them are hour-long captures of packets on a LAN on August 29, 1989 and October 5, 1989, while the other two are day-long captures of WAN traffic to/from Bellcore on October 3, 1989 and October 10, 1989.

While the packet traces represent “ground truth” for prediction, the predictors that we study require discrete-time signals. To produce such a signal, we bin the packets into non-overlapping bins of a small size and average the sizes of the packets in a particular bin by the bin size. This result is an estimate of the instantaneous bandwidth usage that becomes more accurate as the bin size declines. It is important to note that as the bin size decreases the variance of the resulting signal increases. It is this variance that we are trying to model with a predictor. Figure 2 shows this effect for the 34 AUCKLAND traces. Notice that the graph




Fig. 3. Autocorrelation structure of an NLANR trace that is not predictable using linear models.

is on a log-log scale. The linear relationship and gradual slope we can see on the graph for bin sizes greater than 125 ms indicate that the traces are likely long-range dependent with a Hurst parameter near one. The more abrupt slope for smaller binsizes indicates that H declines in that region.

The linear time series models that we evaluate attempt to model the autocorrelation function (ACF) of a discrete-time signal in a small number of parameters. It is important to understand that the ACF has limited meaning if the signal is nonstationary. However, the integration of a stationary signal (modeled by ARIMA models), which is one form of nonstationarity, does show up as an ACF effect. Furthermore, piecewise stationarity (modeled by TAR models), another form of nonstationarity, is very likely to show up as an ACF effect.

The bottom line is that if there is no autocorrelation function, there is nothing to model, a linear approach is bound to fail, a nonlinear approach is likely to fail, and the best predictor is probably the mean of the signal. For this reason, we studied the autocorrelation functions of our traces in considerable detail at different bin sizes. For space reasons, we can not go into detail about this study here. Instead, we shall show representative ACFs from our three different trace sets to explain our choice of presenting detailed results for the AUCKLAND set. We show ACFs at a bin size of 125 ms for each trace.




Fig. 4. Autocorrelation structure of an AUCKLAND trace that is likely to be very predictable using linear models.

Figure 3 shows the ACF of a representative NLANR trace. For any lag greater than zero, the ACF effectively disappears. This signal is clearly white noise and the prospects for predicting it using linear models are very dim. Only 2% of the autocorrelation coefficients rise to significance at a significance level (p -value) of 0.05. 80% of our NLANR traces exhibit this sort of behavior. For the other 20%, more than 5% of the autocorrelation coefficients are significant, but none are very strong. In these cases we can not claim that the signals are white noise, but it is likely that linear models will not do very well.

Figure 4 shows the ACF of a typical AUCKLAND trace. Obviously, the plot is quite different from what we have seen in the previous figure. Over 97% of the autocorrelation coefficients are not only significant, but quite strong. We can also see a low frequency oscillation, which is likely the diurnal pattern. We expect that such a trace will be quite predictable using linear models. 80% of the AUCKLAND traces have similar strong ACFs.

Figure 5 shows the ACF of a BC LAN trace. It is clearly not white noise, and yet it does not have the strong behavior of the AUCKLAND traces. We would expect that such a trace is predictable to some extent using linear models. All of the BC traces have similar ACFs that are suggestive of predictability.

IV. PREDICTABILITY OF BINNING APPROXIMATIONS

To create binning approximation signals in general, we simply bin the packet traces according to the chosen bin size. If we want to support a variety of bin sizes that are integer multiples, we can simply start with the smallest bin size and produce coarser approximations recursively.

Figure 6 illustrates our methodology for evaluating the




Fig. 5. Autocorrelation structure of a BC LAN trace.

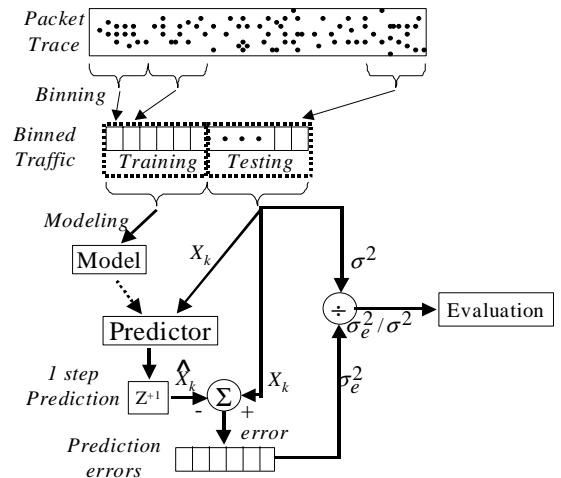


Fig. 6. Methodology of binning prediction.

predictability of a given packet trace at a given bin size. We slice the discrete-time signal produced from binning (X_k) in half. We then fit a predictive model to the first half and create a prediction filter from it. The data from the second half of the trace is streamed through the prediction filter to generate one-step-ahead predictions. Next, we difference these predictions and the values they predict to produce an error signal. We then compute the ratio of the variance of this error signal (the MSE, σ_e^2) to the variance of the second half of the binning approximation signal (σ^2). The smaller the ratio, the better the predictability.

We evaluated the performance of the following models: MEAN, LAST, BM(32), MA(8), AR(8), AR(32), ARMA(4,4), ARIMA(4,1,4), ARIMA(4,2,4), ARFIMA(4,-1,4), and MANAGED AR(32). MEAN uses the long-term mean of the signal as a prediction. Its predictability ratio is typically 1.0 for obvious reasons. LAST simply uses the last observed value as the prediction for the next value. BM(32) predicts that the next value will be the

average of some window of up to the 32 previous values. The size of the window is chosen to provide the best fit to the first half of the signal. MA(8) is a moving average model of order 8. AR(8) and AR(32) are autoregressive models of orders 8 and 32, respectively. ARMA(4,4) is a model with 4 autoregressive parameters and 4 moving average parameters. ARIMA(4,1,4) and ARIMA(4,2,4) are once and twice integrated ARMA(4,4) models. Unlike the other models, they can capture a simple form of non-stationarity. The ARFIMA(4,-1,4) model is a “fractionally integrated” ARMA model that can capture the long-range dependence of self-similar signals. The MANAGED AR(32) model is an AR(32) whose predictor continuously evaluates its prediction error and refits the model when error limits are exceeded. The error limits and the interval of data which the model uses when it is refit are additional parameters. In our presentation, we typically show the best performing MANAGED AR(32). Generally, the sensitivity to the additional parameters is small. MANAGED AR(32) models are variants of threshold autoregressive (TAR) models.

In choosing the number of parameters for our models, we did not employ the Box-Jenkins methodology or AIC. Box-Jenkins is unlikely to be deployable in an online prediction system. AIC trades off between model parsimony and fit, but disregards the computational cost of fitting and using the model in an online system. Our choice of number of parameters for these models was a-priori. We provided a relatively large number of parameters, looking for little sensitivity to a change in the number.

In the following, and for wavelet prediction (Section V), we focus our *presentation* on AUCKLAND for three reasons. First, many of the NLANR traces show minimal predictability. Second, the strength of the ACFs in the AUCKLAND traces allow us to focus on how predictability is affected by the resolution of the signal. Third, unlike the BC traces, the AUCKLAND traces are very long and we have many of them. This lets us consider a wide range of resolutions. It is important to note that our *conclusions* are drawn from studying all of the traces noted in Figure 1.

A. AUCKLAND traces

For each of the 34 AUCKLAND traces, we performed the analysis described above with each of the different predictors. We studied 14 different bin sizes: 0.125 s, 0.25 s, 0.5 s, 1 s, 2 s, 4 s, 8 s, 16 s, 32 s, 64 s, 128 s, 256 s, 512 s, and 1024 seconds. In the discussion that follows, we plot the predictability ratio versus bin size for all the predictors except MEAN. We have elided MEAN since the other predictors typically do much better and including MEAN makes it difficult to see how the other predictors compare.




Fig. 7. Predictability ratio versus bin size of AUCKLAND trace 31 (20010309-020000-0). 44% of traces.




Fig. 8. Predictability ratio versus bin size of AUCKLAND trace 23 (20010305-020000-0). 42% of traces.

It is also important to note that some data points in the graphs are missing. Given enough data it is always possible to fit a model to a trace and glean an estimate of its predictive power from the quality of the fit. However, producing a predictor from the model and sending new data through it as we do often reveals that the model is not as good as the fit might imply. We have elided points in two cases. The first case is when the predictor became unstable as evidenced by a gigantic prediction error. This is sometimes the case with the ARIMA models, which are inherently unstable because they include integration. The second case is when there are insufficient points available to fit the model. This happens at large bin sizes for large models like the AR(32) and the ARFIMA(4,-1,4). Fewer than 5% of points have been elided and we have tried to make it obvious where this happens.

The characteristics of prediction on the AUCKLAND




Fig. 9. Predictability ratio versus bin size of AUCKLAND trace 20 (20010303-020000-1). 14% of traces.

traces fall into three classes, representatives of which are shown in Figures 7 through 9.

The behavior of Figure 7 occurs in 15 of the 34 traces. The most interesting feature here is that the graph shows concavity for all predictors: we can clearly see a *sweet spot* for the traffic prediction. In other words, there is an optimal bin size around 32 seconds at which the trace is most predictable. As we noted earlier, this *contradicts* the conclusions of earlier papers. Because it occurs in half of the AUCKLAND traces, we do not believe that it is a coincidence. The location of the sweet spot varies from trace to trace. In some traces it occurs at quite small bin sizes, which suggests that it is not an artifact of the fact that we are fitting and predicting on smaller amounts of data as we increase bin size. It is clearly an artifact of the data itself.

The behavior of Figure 8 occurs in 14 of the 34 AUCKLAND traces and is commensurate with conclusions from earlier papers. There is no sweet spot here and it is clear that predictability converges to a high level with increasing bin size.

Both of these figures also show significant differences between the performance of the predictors. In general, it is important to have an autoregressive component to the prediction. Fractional models do quite well, but the performance of classical models such as large ARs is close enough to suggest that the extra costs of the fractional models are probably not warranted.

Figure 9 shows an uncommon behavior, as seen in 5 of the 34 AUCKLAND traces. Unlike the two previous kinds of traces, here we have a strong impression of disorder: there are multiple peaks and valleys at different bin sizes. The relative performance of the different predictive models

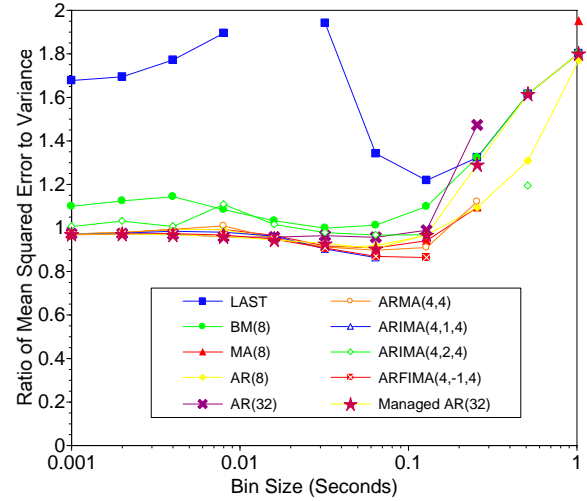


Fig. 10. Predictability ratio versus bin size of a representative NLNLR trace (ANL-1018064471-1-1). 80% of traces.

remains much the same, however.

Our general conclusions about the 34 AUCKLAND traces are the following:

- All of the traces are predictable in the sense that their predictability ratio is less than one. Furthermore, 80% of the traces show strong divergences from one, indicating high predictability. Figures 7 and 8 are examples of traces that are highly predictable. In each of these examples, the predictability ratios are less than 0.4 for all of the predictors at all of the bin sizes. In many cases the ratios are less than 0.1, meaning that the predictor explains 90% of the variation of the signal. This confirms our initial judgment of the predictability of the AUCKLAND traces based on the ACFs from Section III.
- There is considerable variation among the predictors. In almost all cases, LAST, BM, and MA predictors will perform considerably worse. The other six predictors have similar performance except with very large bin sizes where LAST or MA often gives the best results. This is probably due to the fact that there are insufficient data points to produce good fits for some of the predictors at such bin sizes.
- The predictability of a trace varies considerably with bin size. There is often a sweet spot at which predictability is maximized. The location of the sweet spot varies from trace to trace and so is most likely a property of the data. We are trying to understand the origins of this phenomenon. Equally often, predictability increases with bin size, approaching a limit.
- The nonlinear MANAGED AR(32) model provides only marginal benefits, and only at very coarse granularities.




Fig. 11. Predictability ratio versus bin size of a representative BC trace (BC-pOct89).

B. NLANR traces

Because the NLANR traces are only 90 seconds long, we can not use the same range of bin sizes as we did for the AUCKLAND traces. Instead, we chose the following bin sizes: 1, 2, 4, 8, 16, 32, 64, 128, 256, 512, and 1024 ms. Figure 10 shows the predictability ratio for a representative NLANR trace. As we might expect given the ACF behavior described in Section III, this trace is basically unpredictable, turning in predictability ratios around 1.0 or worse for most of the predictors at all the different bin sizes. About 80% of the NLANR traces display similar unpredictability. For the 20% of the traces with non-vanishing ACFs, we see some modicum of predictability, but it is very weak. At coarser granularities, predictability actually declines. The nonlinear MANAGED AR(32) provides no benefits here.

C. BC traces

In Figure 11 we show the performance of the predictors on a BC LAN trace. As the trace is only 1700 seconds long, we have chosen 12 different bin sizes, ranging from 0.0078125 second to 16 seconds, doubling at each step. The predictability here is not as good as for the AUCKLAND traces, although it is much better than for the NLANR traces. All of the BC traces behave similarly. ARIMA models are the clear winners for these traces. Again, we do not necessarily see a monotonic increase in predictability with increasing smoothing. The nonlinear MANAGED AR(32) works much better than its linear AR(32) counterpart at coarse granularities, but other linear models do just as well.




Fig. 12. Methodology of wavelet prediction.

V. PREDICTABILITY OF WAVELET APPROXIMATIONS

Binning is an intuitive mechanism for producing multi-resolution views of network behavior. Wavelet-based mechanisms are a more powerful approach to providing such views because they are parameterized by a wavelet basis function which can be chosen appropriately to optimize for different properties. In fact, the wavelet approach we describe here, when parameterized with the simplest wavelet basis function, the Haar wavelet, is equivalent to the binning approach of the last section. Abry, et al provide a very nice discussion of this equivalence [2]. In the following, we use a D8 wavelet [12], which provides a much smoother multi-resolution analysis than binning/Haar. Wavelet analysis is a rich and multifaceted area of mathematics and signal processing. We will focus our discussion of it on aspects relevant to this study. Interested readers can learn more via the work of Mallat [28], Daubechies [12], Abry, et al [2], or our earlier work [36].

Intuitively, a wavelet transform, in our case a discrete wavelet transform parameterized by a D8 wavelet, splits a 1-dimensional time-domain signal into a 2-dimensional signal representing time and scale (or frequency) information. The output can be thought of as a tree, such that as we move level-by-level toward the root, we see coarser and coarser versions of the signal. Each level of the tree provides both a low-pass filtered version of the signal (the approximations) and a high-pass output (the details). The original signal can be reconstructed using any approximation and the details of all the levels further from the root. We can also reconstruct any finer grain approximation by choosing just the levels we need. In the following, we simply use successive approximations (i.e., low-passing with increasingly lower cutoff frequencies) for successive levels of smoothness, corresponding to larger bin sizes.

To evaluate the predictability of wavelet approximation signals, we use the methodology shown in Figure 12. As

Binsize in seconds	Approximation scale	Number of points	Bandlimit frequency
0.125	Input = 0.125 binsize	n	$f_s/2$
0.25	0	$n/2$	$f_s/4$
0.5	1	$n/4$	$f_s/8$
1	2	$n/8$	$f_s/16$
2	3	$n/16$	$f_s/32$
4	4	$n/32$	$f_s/64$
8	5	$n/64$	$f_s/128$
16	6	$n/128$	$f_s/256$
32	7	$n/256$	$f_s/512$
64	8	$n/512$	$f_s/1024$
128	9	$n/1024$	$f_s/2048$
256	10	$n/2048$	$f_s/4096$
512	11	$n/4096$	$f_s/8192$
1024	12	$n/8192$	$f_s/16384$

Fig. 13. Scale comparison between binning and multi-resolution analysis based on the number of bins and scales used in the AUCKLAND study (n = number of points at 0.125 second binning).

with the binning study, we begin with the packet header trace. We perform a fine-grain binning that produces a highly dynamic discrete time signal, which we denote X_k . We can think of this signal as sampled at a rate f_s and bandlimited to a frequency of $f_s/2$. This signal is sent through a multi-resolution analysis using D8 wavelets and the successive approximations ($approx_i$) are only using the approximations for analysis. For each approximation we run a prediction test identical to the one we used for each bin size in the previous Section.

The D8-based analysis produces different, smoother approximations than the binning approach. For this reason, we reasonably expect (and see) different predictability from the traces. In most cases the behavior is similar, but there are some clear distinctions that we will call out. In our analysis, we have matched the time scale of binsize to that of the approximation subspace. In other words, there are the same number of points in a wavelet approximation signal as in its corresponding binning approximation signal. This is shown in Figure 13. The figure also indicates the bandlimited frequency at each scale. In addition, we plot the predictability of the input signal to match with the binning study. Our input signal is equivalent to the smallest binsize from the binning study.

Our analysis is implemented using Tsunami, our toolkit for wavelet analysis in distributed systems. Tsunami is a C++ template library that can be used to instantiate many different kinds of wavelet transforms (forward or reverse, streaming or block, parameterizable data types, arbitrary frequency decompositions, etc).




Fig. 14. Predictability ratio versus approximation scale AUCKLAND trace 31 (20010309-020000-0). 38% of traces.

A. AUCKLAND traces

For each of the 34 AUCKLAND traces, we studied the predictability of 13 scales of wavelet approximations. As with the binning study we have elided the MEAN predictor and data points that resulted from unstable predictors or having insufficient data to fit a model. There are two principle differences between the wavelet and binning results. The first is that we found four classes of behavior instead of three. The second is that monotonically increasing predictability with increasing approximation is much less common with the wavelet-based approach. We are exploring reasons for this difference.

The behavior of Figure 14 occurs in 13 of the 34 AUCKLAND traces. The figure uses the same trace as Figure 7 from the binning study. As before, we can clearly see that there is a sweet spot, the approximation scale at which predictability is maximized—there is concavity in the figure for all predictors. As before, this behavior does not appear to be a coincidence since it shows up in a number of traces at different levels of approximation. As with binning, this behavior contradicts earlier work.

Figure 15 shows behavior that occurs in 11 of the 34 AUCKLAND traces. It is similar to the behavior we saw in five traces in the binning study and represented in Figure 9. However, here it is far more common. Again, there is a non-monotonic relationship between the approximation scale and the predictability.

Figure 16 shows behavior that occurs in 7 of the 34 AUCKLAND traces. Except for the outliers, this shows the monotonic relationship that was conjectured in earlier work. Note that it is an uncommon behavior in our study.

Figure 17 shows the final class of behavior in the AUCKLAND traces, which occurs in 3 of the 34. Here




Fig. 15. Predictability ratio versus approximation scale for AUCKLAND trace 11 (20010225-020000-0). 32% of traces.




Fig. 17. Predictability ratio versus approximation scale for AUCKLAND trace 4 (20010221-020000-1). 9% of traces.

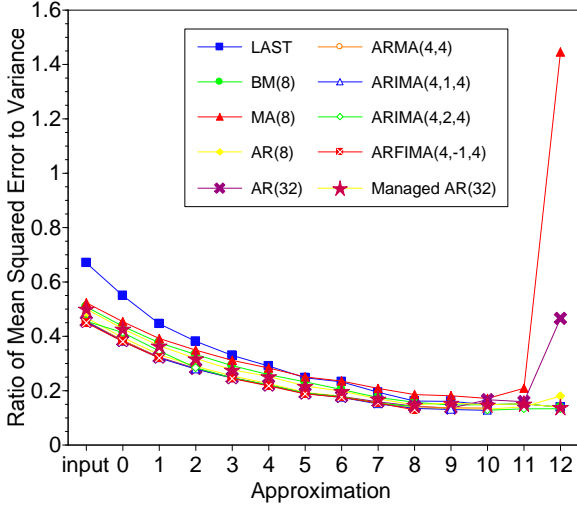


Fig. 16. Predictability ratio versus approximation scale for AUCKLAND trace 32 (20010309-020000-1). 21% of traces.




Fig. 18. Predictability ratio versus approximation scale of a representative NLANR trace (ANL-1018064471-1-1). 80% of traces.

the predictability ratio reaches a plateau and then becomes even more predictable at the coarsest resolutions. Interestingly, this is a kind of behavior that we did not see in the binning study.

The generalizations we draw are much the same as for the binning study:

- Most of the traces show a high degree of predictability, which confirms what we concluded from the ACFs of Section III. On a trace-by-trace basis, the predictability ratio of the binning study is similar to that of the wavelet study when we have similar classes of behavior.
- While there is considerable variation in the performance of the predictors, it is clearly a good idea to have an autoregressive component to the prediction filter. An integrative

component is also useful.

- There is often a sweet spot, the approximation scale at which predictability is maximized.
- There is an additional class of behavior with wavelets compared to binning.
- The nonlinear MANAGED AR(32) model works slightly better than its linear AR(32) counterpart at coarse granularities, but its performance can usually be matched by other linear models.

B. NLANR traces

As we noted earlier, most of the the NLANR traces typically have empty autocorrelation structures, which suggests that there is little predictability. As we might expect, wavelet approximations do not change this. Figure 18




Fig. 19. Predictability ratio versus approximation scale of a representative BC trace (BC-pOct89).

shows typical results using the same trace as Figure 10. As before, the prediction error variance is essentially the same as the signal variance. As with binning, we see that predictability does not increase monotonically with smoothing, and that the nonlinear models provide little benefit.

C. BC traces

Figure 19 shows prediction results for wavelet approximations of the same BC LAN trace studied using binning in Figure 11. We see very similar performance using wavelet approximation signals and binning approximation signals.

VI. CONCLUSIONS

We have presented an empirical study of the predictability of network traffic at different resolutions using linear and nonlinear models. The goal was to assess the prospects for a Message Transfer Time Advisor (MTTA), a tool that could predict, for end-users, the transfer time of application-level messages over an IP network. The feasibility of an MTTA depends on the multiscale predictability of network traffic.

Our results are similar for two methods of producing different resolutions: binning and wavelet approximations. We found that making generalizations about predictability is very difficult in practice because networks tend to vary in behavior considerably over time and space. The behavior at many aggregation points appears to have little predictability using linear models. On the other hand, in situations where predictability exists, it is the case that increased traffic aggregation is correlated with enhanced predictability. This agrees with earlier results. Our study contradicts earlier work in that we find that predictability does

not necessarily monotonically increase with smoothing. About half of the predictable traces we studied have degrees of smoothing at which predictability is maximized. We found that having an autoregressive component to the predictive models is important. Nonlinear models can provide a slight boost in performance over their purely linear analogs at high degrees of smoothing, but are typically matched by other forms of linear models.

Our results imply several things. First, an online multiresolution prediction system to support the MTTA is feasible, but will likely be more accurate on wide area and at coarser timescales. Second, for many wide area environments, there is a natural timescale at which adaptation that relies on network prediction would perform best. Third, while simple predictive models work well, the prediction system should itself be adaptive because network behavior can change.

We are currently working on building online resource signal dissemination and prediction systems based on the Tsunami wavelet toolbox and the RPS system. Tsunami will be incorporated into the next public release of RPS and it will support the study of other techniques, such as wavelet denoising, for further compression of the signal could be applied before it is sent over the network. We also hope to decouple sensor measurement rates and application queries using Tsunami. Once we have a functioning online prediction system, we expect to study the effects of using adaptive prediction filters, such as Kalman filters, in multi-resolution prediction. This work will be the next step toward the Message Transfer Time Advisor.

REFERENCES

- [1] ABRY, P., FLANDRIN, P., TAQUU, M. S., AND VEITCH, D. *Long-Range Dependence: Theory and Applications*. Birkhauser, 2002, ch. Self-similarity and long range dependence through the wavelet lens.
- [2] ABRY, P., VEITCH, D., AND FLANDRIN, P. Long-range dependence: Revisiting aggregation with wavelets. *Journal of Time Series Analysis* 19, 3 (May 1998), 253–266.
- [3] ACM SIGCOMM. Internet Traffic Archive. <http://ita.ee.lbl.gov>.
- [4] BASU, S., MUKHERJEE, A., AND KLIVANSKY, S. Time series models for internet traffic. In *Proceedings of Infocom 1996* (March 1996), pp. 611–620.
- [5] BERAN, J. Statistical methods for data with long-range dependence. *Statistical Science* 7, 4 (1992), 404–427.
- [6] BERMAN, F., AND WOLSKI, R. Scheduling from the perspective of the application. In *Proceedings of the Fifth IEEE Symposium on High Performance Distributed Computing HPDC96* (August 1996), pp. 100–111.
- [7] BOX, G. E. P., JENKINS, G. M., AND REINSEL, G. *Time Series Analysis: Forecasting and Control*, 3rd ed. Prentice Hall, 1994.
- [8] BRAKMO, L. S., AND PETERSON, L. L. TCP vegas: End to end congestion avoidance on a global internet. *IEEE Journal on Selected Areas in Communications* 13, 8 (1995), 1465–1480.

- [9] BUSH, S. F. Active virtual network management prediction. In *Proceedings of the 13th Workshop on Parallel and Distributed Simulation (PADS '99)* (May 1999), pp. 182–192.
- [10] CASSETTI, C., KUROSE, J. F., AND TOWSLEY, D. F. A new algorithm for measurement-based admission control in integrated services packet networks. In *Proceedings of the International Workshop on Protocols for High-Speed Networks* (October 1996), pp. 13–28.
- [11] CHONG, S., LI, S., AND GHOSH, J. Predictive dynamic bandwidth allocation for efficient transport of real-time VBR video over ATM. *IEEE Journal of Selected Areas in Communications* 13, 1 (1995), 12–23.
- [12] DAUBECHIES, I. *Ten Lectures on Wavelets*. Society for Industrial and Applied Mathematics (SIAM), 1999.
- [13] DINDA, P., GROSS, T., KARRER, R., LOWEKAMP, B., MILLER, N., STEENKISTE, P., AND MILLER, N. The architecture of the remos system. In *Proceedings of the 10th IEEE International Symposium on High Performance Distributed Computing (HPDC 2001)* (August 2001), pp. 252–265.
- [14] DINDA, P. A. Online prediction of the running time of tasks. *Cluster Computing* (2002). To appear, earlier version in HPDC 2001, summary in SIGMETRICS 2001.
- [15] DINDA, P. A., AND O'HALLARON, D. R. An extensible toolkit for resource prediction in distributed systems. Tech. Rep. CMU-CS-99-138, School of Computer Science, Carnegie Mellon University, July 1999.
- [16] FELDMAN, A., GILBERT, A. C., AND WILLINGER, W. Data networks as cascades: Investigating the multifractal nature of internet WAN traffic. In *Proceedings of ACM SIGCOMM '98* (1998), pp. 25–38.
- [17] FELDMANN, A., GILBERT, A., HUANG, P., AND WILLINGER, W. Dynamics of ip traffic: a study of the role of variability and the impact of control. In *Proceedings of the ACM SIGCOMM 1999* (Cambridge, MA, August 29 - September 1 1999).
- [18] FOSTER, I., AND KESSELMAN, C., Eds. *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann, 1999.
- [19] GRANGER, C. W. J., AND JOYEUX, R. An introduction to long-memory time series models and fractional differencing. *Journal of Time Series Analysis* 1, 1 (1980), 15–29.
- [20] GROSCHWITZ, N. C., AND POLYZOS, G. C. A time series model of long-term NSFNET backbone traffic. In *Proceedings of the IEEE International Conference on Communications (ICC'94)* (May 1994), vol. 3, pp. 1400–4.
- [21] HOSKING, J. R. M. Fractional differencing. *Biometrika* 68, 1 (1981), 165–176.
- [22] HUANG, P., FELDMANN, A., AND WILLINGER, W. A non-intrusive, wavelet based approach to detecting network performance problems. In *Proceeding of ACM SIGCOMM Internet Measurement Workshop 2001* (San Francisco, CA, November 2001).
- [23] JACOBSON, V. Congestion avoidance and control. In *Proceedings of the ACM SIGCOMM 1988 Symposium* (August 1988), pp. 314–329.
- [24] JAMIN, S., DANZIG, P., SCHENKER, S., AND ZHANG, L. A measurement-based admission control algorithm for integrated services packet networks. In *Proceedings of ACM SIGCOMM '95* (February 1995), pp. 56–70.
- [25] KIM, M., AND NOBLE, B. Mobile network estimation. In *Proceedings of the Seventh Annual International Conference on Mobile Computing and Networking* (July 2001), pp. 298–309.
- [26] LELAND, W. E., TAQQU, M. S., WILLINGER, W., AND WILSON, D. V. On the self-similar nature of ethernet traffic. In *Proceedings of ACM SIGCOMM '93* (September 1993).
- [27] LU, D., AND DINDA, P. A. Virtualized audio: A highly adaptive interactive high performance computing application. In *Proceedings of the 6th Workshop on Languages, Compilers, and Run-time Systems for Scalable Computers* (March 2002).
- [28] MALLAT, S. Multiresolution approximation and wavelets. *Transactions American Mathematics Society* (1989), 69–88.
- [29] MILLER, N., AND STEENKISTE, P. Network status information for network-aware applications. In *Proceedings of IEEE Infocom 2000* (March 2000). To Appear.
- [30] NATIONAL LABORATORY FOR APPLIED NETWORKING RESEARCH. Nlanr network analysis infrastructure. <http://moat.nlanr.net>. NLANR PMA and AMP datasets are provided by the National Laboratory for Applied Networking Research under NSF Cooperative Agreement ANI-9807579 .
- [31] QIAO, Y., AND DINDA, P. Network traffic analysis, classification, and prediction. Tech. Rep. NWU-CS-02-11, Department of Computer Science, Northwestern University, January 2003.
- [32] RIEDI, R., CROUSE, M., RIBEIRO, V., AND BARANIUK, R. A multifractal wavelet model with application to network traffic. *IEEE Transactions on Information Theory* 45, 3 (April 1999), 992–1019.
- [33] ROUGHAN, M., VEITCH, D., AND ABRY, P. On-line estimation of the parameters of long-range dependence. In *Proceedings Globecom 1998* (November 1998), vol. 6, pp. 3716–3721.
- [34] SANG, A., AND LI, S. Predictability analysis of network traffic. In *Proceedings of INFOCOM 2000* (2000), pp. 342–351.
- [35] SKICEWICZ, J., AND DINDA, P. Tsunami: A wavelet based toolkit for the analysis and dissemination of resource information. Tech. Rep. NWU-CS-03-16, Department of Computer Science, Northwestern University, April 2003.
- [36] SKICEWICZ, J., DINDA, P., AND SCHOPF, J. Multi-resolution resource behavior queries using wavelets. In *Proceedings of the 10th IEEE International Symposium on High Performance Distributed Computing (HPDC 2001)* (August 2001), pp. 395–405.
- [37] STEMM, M., SESHAN, S., AND KATZ, R. H. A network measurement architecture for adaptive applications. In *Proceedings of INFOCOM 2000* (March 2000), vol. 1, pp. 285–294.
- [38] TONG, H. *Threshold Models in Non-linear Time Series Analysis*. No. 21 in Lecture Notes in Statistics. Springer-Verlag, 1983.
- [39] VAZKHUAI, S., SCHOPF, J., AND FOSTER, I. Predicting the performance of wide area data transfers. In *Proceedings of the 16th International Parallel and Distributed Processing Symposium (IPDPS 2002)* (April 2002).
- [40] WILLINGER, W., TAQQU, M. S., SHERMAN, R., AND WILSON, D. V. Self-similarity through high-variability: Statistical analysis of ethernet lan traffic at the source level. In *Proceedings of ACM SIGCOMM '95* (1995), pp. 100–113.
- [41] WOLSKI, R., SPRING, N. T., AND HAYES, J. The network weather service: A distributed resource performance forecasting system. *Journal of Future Generation Computing Systems* (1999). To appear. A version is also available as UC-San Diego technical report number TR-CS98-599.
- [42] YOU, C., AND CHANDRA, K. Time series models for internet data traffic. In *Proceedings of the 24th Conference on Local Computer Networks LCN 99* (1999), pp. 164–171.