

Modeling of Multi-resolution Active Network Measurement Time-series

Prasad Calyam, Ananth Devulapalli

Ohio Supercomputer Center, 1224 Kinnear Road, Columbus, OH 43212

{pcalyam, ananth}@osc.edu

Abstract—Active measurements on network paths provide end-to-end network health status in terms of metrics such as bandwidth, delay, jitter and loss. Hence, they are increasingly being used for various network control and management functions on the Internet. For purposes of network health anomaly detection and forecasting involved in these functions, it is important to accurately model the time-series process of active measurements. In this paper, we describe our time-series analysis of two typical active measurement data sets collected over several months: (i) routine, and (ii) event-laden. Our analysis suggests that active network measurements follow the moving average process. Specifically, they possess $ARIMA(0,1,q)$ model characteristics with low q values, across multi-resolution timescales. We validate our model selection accuracy by comparing how well our predicted values using our model match the actual measurements.

I. INTRODUCTION

Real-time applications such as videoconferencing and Grid computing are being widely deployed on the Internet. To cater to the end-user expectations, Internet Service Providers (ISPs) need to regularly monitor their networks and perform adaptation using various network control and management functions such as path switching [1] and bandwidth-on-demand [2]. These functions require collection of network health measurements using active measurement tools such as Ping, Traceroute, Iperf [3], Pathchar [4], and Pathrate [5]. The collected measurements are analyzed for network-wide performance anomaly detection [6] and forecasting [7].

The degree of success achievable by such analyses greatly depends on a sound understanding of the time-series process of the active measurements. Analysis of active measurement data sets is challenging because of the high variability caused by factors such as end-user behavioral patterns, network fault events and cross-traffic congestion. Studies such as [8] and [9] have attempted to characterize and model the high variability of network measurements. The variations manifest in the form of short spikes, burst spikes and distinct plateaus, which complicate the time series analysis. Network outages due to failure of communication protocols lead to failure of measurement probes, which result in missing data points in the time series data collection. Hence, it is common to find gaps in active network measurement data sets, which compound the time series analysis challenges.

Further, different network control and management functions require the models of measurements at multi-resolution timescales. For example, measurements in time scales on the order of few hours/days need to be modeled to forecast

the performance (e.g. delay or loss) of a network for the subsequent few hours/days. Whereas, measurements in time scales on the order of few months need to be modeled for network planning that involves predicting bandwidth upgrades that will cater to future user-demands. Also, anomaly detection functions need the measurements to be modeled in time resolutions corresponding to notable network events. Here, the goal in anomaly detection is to identify how the model changes due to a network event (e.g. route change). Such an event detection may then lead to an anomaly-alarm for problem resolution.

The work presented in this paper corresponds to the early results of an extensive study we are conducting to address the above challenges and requirements in modeling active network measurements at multi-resolution timescales. Our ultimate goal is to develop new perspectives for the on-going debate about how to best model the high variability of network path performance on the Internet. For the modeling, we use the Auto-Regressive Integrated Moving Average (ARIMA) class that is most commonly used for modeling chaotic time series. Recent works such as [10] and [11] have successfully applied ARIMA models to measure network traffic variability. Results from [10] suggest that prediction accuracy of ARIMA models can be improved when combined with a non-linear time series model. Authors in [11] showed that seasonal ARIMA models can be used to model and predict wireless network traffic characteristics.

There are two main contributions of this paper: (i) we present a systematic methodology to determine the ARIMA model parameters that are most suitable for characterizing active network measurements, and (ii) we evaluate the impact of multi-resolution timescales due to both absence and presence of network events on the ARIMA model parameters.

Our time-series analysis involves two phases. In the first phase called the model training phase, we use the classical decomposition procedure [12] on two typical active measurement data sets: (i) “routine”, and (ii) “event-laden”, that are part of a large dataset collected over several months. The data sets were collected using OSC’s ActiveMon software [13] that orchestrates active measurement probes and collects the measurements into a central database. The routine data set contains jitter measurements on a campus backbone network path that has not experienced any notable network events. The event-laden data set contains delay measurements on a regional backbone network path showing noticeable plateaus caused by




Fig. 1. ActiveMon System Setup

multiple route-change network events. In the second phase called the model validation phase, we evaluate our model selection accuracy by evaluating the forecasting ability of our model. In other words, we compare how well our predicted values using our model match the actual measurements. The actual measurements correspond to a portion of measurements that were held-back in the model training phase.

The remainder of this paper is organized as follows: Section II describes the data sources and modeling methodology. Section III presents the time-series analysis results and related discussions. Finally, Section IV concludes the paper and suggests future work.

II. MODELING METHODOLOGY

A. Data Sources

The “routine jitter measurement” and “event-laden delay measurement” data sets used in this study are part of a large data set that was collected by active measurement probes over several months. The active measurements were collected over three hierarchically different Internet backbone paths: campus, regional and national paths [13]. The campus path refers to a network path between a Computer Science Lab (OSUL) at The Ohio State University and Ohio Supercomputer Center’s (OSC) border router for OSU campus (OSUB); only OSU campus network backbone routers were present in between, along this path. The regional path refers to a network path between OSUB and border router of University of Cincinnati (UOCB); only OSCnet backbone routers were present in between, along this path. Note that OSCnet is the ISP for both OSU and UOC. The national path refers to a network path between OSUB and a Computer Science Lab (NCSL) at North Carolina State University; backbone routers of two regional and one national provider were present in between, along this path.

The measurement probes used OSC’s “ActiveMon” [13] system setup shown in Figure 1 for measurement data col-

lection and database storage. ActiveMon consists of two main components: (i) *node*, and (ii) *root*. The node refers to one or more measurement servers co-located with a core router and associated switch at a strategic point within a network being monitored. High-precision one-way delay measurements are obtained using CDMA devices attached to a measurement server. The node servers host various active measurement tools for network path measurements between any two nodes. The network-wide measurements are centrally collected at a root server that contains a database along with a webservice. Various scripts at the root server analyze and visualize the collected data as required by the network administrators. The root server also hosts a scheduler that is responsible for scheduling the network-wide active measurements without measurement conflicts and with measurement regulation [14].

The routine jitter measurement data set shown in Figure 2 was obtained on the OSUB to UOCB network path using the Iperf tool [3] over a two-month period. The UDP measurement probing feature was used in Iperf tool with a peak bandwidth of 768 Kbps for the probe traffic. The network operation’s logs indicate that there were no major network events during the two-month period of jitter measurement collection. The event-laden delay measurement data set shown in Figure 3 was obtained on the OSUB to OSUL network path using the Ping tool over a six-month period. The network operation’s logs indicate that four route-change events occurred due to various network management activities involved in transitioning of the campus traffic from an old ATM network to a new IP network.




Fig. 2. Plot of Jitter Measurements Data

B. Classical Decomposition Procedure

Figure 4 shows the classical decomposition procedure [12], specifically the *Box-Jenkins modeling methodology* followed in this paper. The first step in modeling is to visually inspect the data set to verify presence of any seasonal, or time-based trends. Such an inspection is helpful in reducing the model search space. If necessary, the data after the inspection is




Fig. 3. Plot of delay data

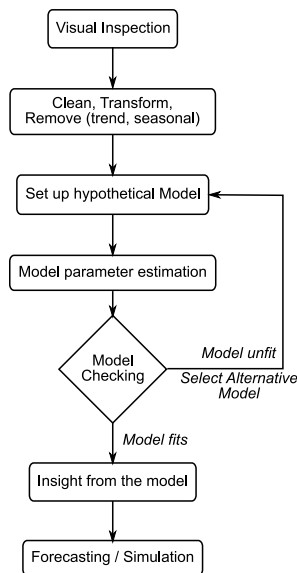


Fig. 4. Flowchart for Box-Jenkins model identification

cleaned and transformed. Further, any notable trends, seasonality and cycles in the data need to be eliminated to obtain stationary data. One of the popular techniques to eliminate trend effects and achieve data stationarity is the *differencing* technique, which we use in this paper. A *one-lag* differencing transforms the data set by getting the differences between consecutive data points. Note that differencing recursively applied N times is referred to as N -lag differencing.

Next, a preliminary analysis of the data is performed using the *sample auto-correlation* (ACF), and *sample partial auto-correlation* (PACF). This analysis allows for selection of a set of hypothetical models. Sample ACF measures the correlation between different elements at a lag and is useful in determining if the data follows a *moving average* (MA) process. Sample

PACF is more complex, and is useful in determining if the data follows an *auto-regressive* (AR) process. Based on ACF, PACF and other methods, a promising set of hypothetical models corresponding to $ARIMA(p, d, q)$ is selected, where p , d , and q are integers (≥ 0), and refer to the order of MA, differencing, and AR parts of the model respectively.

For each of the selected hypothetical models, the model parameters are estimated. Next they are checked using a “goodness-of-fit” test. If a model is deemed to be unfit, an alternative model is selected. Ultimately the models that show statistically significant fitness are selected. Amongst these short-listed models, one best-fitting model needs to be selected. For this, the typical approach is to use the “Akaike’s information criterion” (AIC) as a quantitative metric for model selection. The preferred model will have the lowest AIC amongst all the short-listed models. Once we have selected the model, the data sets are related to the selected model to obtain insights into the underlying process involved in the data generation. Finally, the selected model is used for various purposes such as forecasting, and simulation in the domain of the data sets.

III. ANALYSIS RESULTS AND DISCUSSION

In this section we first discuss our time-series analysis of the routine jitter measurement data set. Next, we explain our time-series analysis of the event-laden delay measurement data set. Lastly, we present a parts versus whole time-series analysis of the two data sets.

A. Analysis of Routine Jitter Measurements

1) *Preliminary Data Examination:* On visual inspection of the jitter measurements data shown in Figure 2, we can observe that the data exhibits no apparent trends or seasonality. There are frequent spikes and dips in the data without any specific patterns. The variance of the data seems to be bounded within 0.1 ms to 0.5 ms during the entire data collection. Hence, stabilization of variance using transformation of the data is not necessary.

Inspecting the actual data file, there were a total of 493 data points. All the data points had a constant two-hour inter-sampling period. However, there was an 18 hour gap in the data collection. It is common for network measurements data to have gaps due to factors such as downtimes of measurement servers or network outages that stall measurements collection. However, given that this gap corresponds to less than 10 data points, we can expect the missing data to have minimal influence on the model.

We split the jitter measurements data set into two parts. The first part contains $n - m = 469$ ($n = 493$; $m = 24$) observations, which are used as part of the “training data set” for estimation and model building. The second part with $m = 24$ observations is used as a “test data set” for computing measures of predictive accuracy of the selected model. Note that the following model selection, diagnostics and other analyses are done on the training data set. The test data set is held-back to verify the

power of the model by comparing the predicted measurements and their confidence bounds with the actual measurements.




Fig. 5. ACF (a) and PACF (b) of the jitter data set

2) *ACF and PACF*: Figures 5(a) and (b) show the ACF and PACF of the jitter measurements, respectively. From the ACF plot, we can notice that there is no indication of $MA(q)$ since there is no clear cut-off at any lag. Also since the ACF is not decaying exponentially, there is no indication of $AR(p)$. Looking at the PACF plot, again it is not possible to conclude whether an $MA(q)$ or $AR(p)$ model would be suitable. This suggests that there are inherent trends in the data that are not visually noticeable. We proceed with differencing to remove any inherent trends in the data.




Fig. 6. ACFs for First-order, Second-order and Third-order Differencing

3) *Differencing*: One-lag differencing of the jitter measurements shows indications of $MA(2)$. Next, we verify whether higher order differencing i.e., second-order and third-order differencing is required on the data set. Figure 6 shows the

differenced data and their corresponding ACFs. The autocorrelation values are greater than -0.5 at lag1 for second-order and third-order differencing, which indicates effect of *over-differencing*. Hence, we conclude that first-order differencing is sufficient on the data set.




Fig. 7. ACF and PACF of jitter dataset after differencing, (a) and (b) respectively

Figures 7(a) and (b) show the ACF and PACF plots of one-lag differenced jitter data, respectively. The ACF plot clearly indicates that the process might be $MA(2)$ since there is a sharp cut-off after lag-2. Further, the PACF is decaying exponentially that eliminates AR process possibility.




Fig. 8. AIC versus order of $MA(q)$

4) *Model Fitting*: To confirm our $MA(2)$ observation, we calculate AIC as a function of model sizes i.e., $MA(q)$ with q from 1 to 10. Figure 8 shows the plot of AIC versus the order of MA . The AIC values do not show a “decrease and increase” pattern that clearly indicates the q -value of MA that results in the minimum AIC. Given that the AIC dips with the increase in q -values, we can note that there is no significant minimization of AIC as the q value increases. Hence, we can conclude that $MA(1)$ is a suitable model for the jitter measurements data.

To double-check if our model selection of $MA(1)$ is correct, we calculate AIC for $ARMA(1,1)$, $ARMA(1,2)$, $ARMA(2,1)$ and $ARMA(2,2)$ models. Table I shows the obtained AIC values. We can see that the AIC values of $ARMA(1,1)$ are comparable to AIC value of $MA(1)$ and the AIC values do

not significantly decrease with the increase in the model complexity. Hence, we can further conclude that MA(1) is a suitable model for the jitter measurements data.

ARMA Order	AIC
11	-1611.728
12	-1609.730
21	-1609.755
22	-1609.794

TABLE I
AIC VALUES FOR ARMA(P,Q) WHERE P=1,2 AND Q=1,2

To further verify the correctness of the MA(1) model selection, we compare the significance of the parameter values of MA(1) and other higher orders. For this, we inspect whether the 95% CI values ($\theta_x \pm 1.96 \times \sigma_{\theta_x}$) of MA(1), MA(2) and MA(3) contain zero. Table II shows the 95% CI values for three θ parameters of the MA model. We can note that the 95% CI values of θ_1 are significant because they do not contain zero. However, the 95% CI values of θ_2 and θ_3 contain zero, suggesting that we cannot reject the null hypothesis that MA(1) is not the suitable model. Hence, we can confirm that MA(1) with one-lag differencing or equivalently ARIMA(0, 1, 1) is the suitable model for the jitter measurements data.

Parameter	Value	95% CI
θ_1	-0.9440	(-0.85286, -1.03514)
θ_2	-0.0123	(-0.135584, 0.110984)
θ_3	-0.0114	(-0.10744, 0.08464)

TABLE II
TABLE OF 95% CI VALUES FOR MA MODEL PARAMETERS

5) *Diagnostic Checking*: Figure 9 shows the diagnostic checking using the selected MA(1) model. From the diagnostics of the fitted model we can observe that the residuals look like noise values. The ACF of the residuals indicate that the residuals are uncorrelated and resemble a white noise process. The ‘‘Ljung-Box plot’’ shows that the model is significant at all lags and that the residuals are not just correlated, they are also independent. Above observations conclude that the MA(1) model has captured the structure of the jitter measurement data, and the residuals are just left-over noise.

Based on above analysis, the model for the jitter measurement data is as follows:

$$X_t = Z_t + (-0.9440)Z_{t-1}$$

where $Z_i \sim \text{White Noise}(0, 0.01028)$ and $X_i = \text{diff}(jitter_i)$

6) *Prediction Based on MA(1) Model Fitting*: In this section, we validate our model selection by evaluating the prediction accuracy of our selected MA(1) model. For this, we compare the prediction estimates with the actual measurements




Fig. 9. Diagnostics of fitted model

that were held back i.e., $m = 24$ data points, in the original training data set.




Fig. 10. Prediction using the selected MA(1) Model (a) and actual measurement data overlaid with predicted values and prediction confidence intervals (b)

Figure 10(a) shows the prediction values and the prediction confidence intervals based on the training data set. Figure 10(b) shows the actual measurement or test data overlaid with the predicted values and prediction confidence intervals. Comparing the predicted values and the actual measured values, we can note that the prediction is working well. Except for a couple of measurements, most of the actual measurements lie within the prediction confidence bounds. The outlier measurements can be attributed to unusual fluctuations in the jitter measurements that are not uncommon given the nature of the data set.

B. Analysis of Event-laden Delay Measurements

1) *Preliminary Data Examination:* On visual inspection of the delay measurements data shown in Figure 3, we can clearly see four distinct plateaus. Recall that these plateaus are the result of network route change events. Inspecting the actual data file, there were a total of 2164 data points. The data was periodic for the most part, however, there were minor gaps of data i.e., data was missing for a day or less in a few cases. Since data size is quite large, we can expect that such gaps have minimal influence on the model.

We split the delay measurements data into two parts. The first part contains $n - m = 2100$ ($n = 2164$; $m = 64$) observations, which are used as part of the “training data set” for estimation and model building. The second part with $m = 64$ observations is used as a “test data set” for computing measures of predictive accuracy of the selected model. Note that the following model selection, diagnostics and other analyses are done on the training data set. The test data set is held-back to verify the power of the model by comparing the predicted values and their confidence bounds with the actual measurements.




Fig. 11. ACF (a) and PACF (b) of the delay dataset

2) *ACF and PACF:* Figures 11(a) and (b) show ACF and PACF of the delay measurements, respectively. It is clear from the ACF plot that the auto-correlations are not damping suggesting that the underlying process may not an MA process. PACF of the untransformed data shows damping of auto-correlations, which suggest that the AR process might be a possibility. Further, by observing the PACF damping till lag 7, there is indication of an AR(7) process. Given that the data is showing inherent trends, we proceed with differencing.

3) *Differencing:* Figure 12 shows the plot of one-lag differencing on the delay measurements data to remove any trend in the series. Second-order and third-order differencing showed effects of over-differencing. Consequently, we limit ourselves to first-order differencing on the data set. From the one-lag differencing, we observe that there is neither a trend nor seasonality in the time-series. Hence, we can infer that the process is more or less stationary. To confirm our intuition, we plot the ACF of one-lag differencing of the delay measurements data.

Figures 13(a) and (b) show the ACF and PACF of one-lag




Fig. 12. Plot of 1-lag differencing of delay series




Fig. 13. ACF (a) and PACF (b) of the 1-lag differenced delay dataset

difference of the delay measurements data, respectively. The damping in the ACF plot clearly indicates that the process is an MA(2) process since there is a sharp cut-off after lag-2. Further, PACF is decaying exponentially that eliminates AR process possibility.

4) *Model Fitting:* To confirm our MA(2) observation, we calculate AIC as a function of model sizes i.e., MA (q) with q from 1 to 6. Figure 14 shows the plot of AIC versus order of MA. The AIC values show a clear “decrease and increase” pattern that clearly indicates the q -value of 3 has the minimum AIC. Consequently, the MA(3) process explains the structure of the process without added complexity of higher order models. Hence, we can conclude that MA(3) is a suitable model for the delay measurements data.

Parameter	Value	95% Conf. Interval
θ_1	-0.4876	(-0.5303, -0.4449)
θ_2	-0.0064	(-0.0552, +0.0424)
θ_3	-0.0564	(-0.0983, -0.0145)

TABLE III
TABLE OF CONFIDENCE INTERVALS OF PARAMETERS

To further verify the correctness of the MA(3) model selection, we compare the significance of the parameter values




Fig. 14. AIC versus order of MA(q)

of MA(3) by observing their confidence intervals. Table III shows the confidence intervals for the three parameters of the model. From the confidence intervals, it is clear that θ_1 is significant, since the 95% confidence interval does not include 0. But the confidence interval of θ_2 includes 0, which would have suggested MA(1) model instead of MA(3) model fitted. However, if we observe confidence interval of θ_3 it does not contain zero. This indicates that we cannot reject MA(3) model even though we might have rejected the MA(2) model. Hence, we can confirm that MA(3) with one-lag differencing or equivalently ARIMA(0,1,3) is the suitable model for the delay measurements data.




Fig. 15. Diagnostics of fitted model

5) *Diagnostic Checking*: Figure 15 shows the diagnostics checking using the selected MA(3) model. From the diagnostics of the fitted model, we make the following observations:

- The residuals do not exhibit any apparent pattern. They resemble a noise process whose mean is 0. We can also note that there are few large residuals near the plateau transition points.
- ACF of the residuals indicates that the residuals are uncorrelated. After the initial spike all correlations are near 0. Since the residuals have zero mean and are uncorrelated, this means that the residuals represent a white noise process.
- Ljung-Box plot shows that the model is significant at all lags and that the residuals are not just correlated, but they are also independent. This suggests that the residuals resemble an iid noise process.

Above observations conclude that the MA(3) model has captured the structure of the delay measurement data, and the residuals are just left-over noise.

Based on above analysis, the model for the delay measurement data is as follows:

$$X_t = Z_t + (-0.4876)Z_{t-1} + (-0.0064)Z_{t-2} + (-0.0564)Z_{t-3}$$

where $Z_i \sim \text{White Noise}(0, 0.01028)$ and $X_i = \text{diff}(\text{delay}_i)$




Fig. 16. Prediction using the selected MA(3) Model (a) and test data overlaid with predicted values and prediction confidence intervals (b)

6) *Prediction*: In this section, we validate our model selection by evaluating the prediction accuracy of our selected MA(3) model. For this, we compare the prediction estimates with the actual measurements that were held back i.e., $m = 64$ data points, in the original training data set.

Figures 16(a) and 16(b) show the original model overlaid by prediction and confidence intervals. Figure 16(a) shows the trailing part of the training data and actual measurements or test data for 10 values with confidence bounds. To see if the prediction is capturing the data, in Figure 16(b) we fit first 24 points of prediction dataset and overlay the prediction and confidence bounds for those 24 points. All the actual measurements lie within the prediction confidence bounds.

C. Parts Versus Whole Analysis

Herein, we present our multi-resolution time-series analysis to determine the impact of multi-resolution timescales due to both absence and presence of network events on the ARIMA

model parameters. Due to space constraints, we do not show the detailed time-series analysis below.

1) *Jitter data analysis*: For the jitter measurements data, we choose a two-part resolution of the data separated by the gap (corresponding to the 18 hours outage). Our aim is to assess whether the two data portions also exhibit MA(1) characteristics.

For the first data part, we observed that the time-series plot does not show any apparent trends or seasonality. The ACF plot clearly indicated an MA(1) process. We plotted the AIC values for higher order MA and found the plot to be similar to Figure 8, i.e., the first data part is an MA(1) process. For the second data part, the time-series plot showed a slightly decreasing trend. The ACF and PACF plots did not clearly indicate MA(q) or AR(p) characteristics. Upon applying one-lag differencing to remove any inherent trends in the data, the ACF of the differenced data set indicated MA(2) with sharp cut-off at lag 2. However, upon plotting the AIC values for higher order MA, we found the plot to be similar to the AIC values plot for the whole data, i.e., the second data part is also an MA(1) process.

2) *Delay data analysis*: For the delay measurements data, we choose a four-part resolution of the data separated by the plateaus viz., $d1$, $d2$, $d3$, $d4$ shown in Figure 3 (corresponding to the route change network events). Our aim is to assess how the process changes in the four data portions.

ACF plots for the four parts of delay measurement data indicated interesting characteristics. For the parts $d1$ and $d3$, ACF damped quickly, which suggests MA(1). On the other hand, for the parts $d2$ and $d4$ ACF does not damp, hence we reject the MA model for them. PACF plots for the four parts also had interesting characteristics i.e., PACF of $d1$ and $d3$ damped after one-lag differencing indicating AR(1) model. Whereas, PACF of $d2$ and $d4$ damped after lag 12 suggesting AR(12) model. The ACF and PACF plots of the non-differenced data also showed the same behavior as that of the non-differenced data. Hence, the plateau parts had different models in comparison with the whole data set.

From the above analysis of jitter and delay measurement data sets, we can observe that the “parts resemble the whole” when there are no network events. However, if network events are present, the underlying process changes, which can be leveraged for anomaly detection purposes.

IV. CONCLUSION AND FUTURE WORK

In this paper, we presented a systematic time-series analysis of active network measurement data sets whose inherent characteristics included: sudden fluctuations due to network traffic transients, network events and missing data due to outages. Our analysis presented new insights into the modeling of the high variability in the network path performance on the Internet.

Results from our extensive time-series analysis using diagnostics, 95% CI checking and prediction concluded that it is reasonable to characterize network performance on the

Internet using $ARIMA(0,1,q)$ models. We found that first-order differencing of active network measurement data sets can remove inherent trends (in routine data set) as well as apparent plateau trends (in event-laden data set). Also, MA(q) models with reasonably low q values are suitable for even plateau-laden data. Consequently, network performance data sets have “too much memory” and thus auto-regressive models that are dependent on present and past values may not be well-suited. Another notable observation was that the “parts resemble the whole” in the absence of any major network events i.e., the model selection is comparable if we break down the entire data into parts and analyzed each of the individual parts separately. Whereas, in the presence of network events causing plateaus, the underlying process changes in the parts.

As part of our future investigations, our goal is to apply similar methodology on several other data sets collected using ActiveMon. From this, we hope to determine the underlying process changes across typical short, medium and long time-resolutions - both in the presence and absence of various network events. Ultimately, we believe such analysis can be leveraged in network control and management schemes to obtain better prediction accuracy and lower anomaly detection false-alarms.

REFERENCES

- [1] S. Tao, K. Xu *et al.*, “Improving voip quality through path switching,” in *IEEE INFOCOM*, 2005, pp. 2268–2278.
- [2] M. Yang, Y. Huang *et al.*, “An end-to-end qos framework with on-demand bandwidth reconfiguration,” *Elsevier Computer Communications*, vol. 28, no. 18, pp. 2034–2046, 2005.
- [3] A. Tirumala, L. Cottrell, and T. Dunigan, “Measuring end-to-end bandwidth with iperf using web100,” in *Passive and Active Measurement Workshop*, 2003.
- [4] A. B. Downey, “Using pathchar to estimate internet link characteristics,” *SIGCOMM Computer Communications Review*, vol. 29, no. 4, pp. 241–250, 1999.
- [5] C. Dovrolis, P. Ramanathan, and D. Moore, “Packet-dispersion techniques and a capacity-estimation methodology,” *IEEE/ACM Transactions on Networking*, vol. 12, no. 6, pp. 963–977, 2004.
- [6] M. Thottan and C. Ji, “Anomaly detection in ip networks,” *IEEE Transactions on Signal Processing*, vol. 51, no. 8, pp. 2191–2204, 2003.
- [7] R. Wolski, N. Spring, and J. Hayes, “The network weather service: A distributed resource performance forecasting service for metacomputing,” *Journal of Future Generation Computer Systems*, vol. 15, no. 5, pp. 757–768, 1999.
- [8] W. Willinger, D. Alderson, and L. Li, “A pragmatic approach to dealing with high-variability in network measurements,” in *ACM Internet Measurement Conference*, 2004, pp. 88–100.
- [9] M. Roughan, A. Greenberg *et al.*, “Experience in measuring backbone traffic variability: Models, metrics, measurements and meaning,” in *ACM Internet Measurement Workshop*, 2002.
- [10] B. Zhou, D. He *et al.*, “Network traffic modeling and prediction with arima/garch,” in *HET-NETs*, 2005.
- [11] Y. Shu, M. Yu *et al.*, “Wireless traffic modeling and prediction using seasonal arima models,” *IEICE Transactions on Communications*, vol. E88-B, no. 10, pp. 3992–3999, 2005.
- [12] M. Wheelwright and R. Hyndman, “Forecasting: Methods and applications (3rd edition),” in *Wiley Publication*, 1998.
- [13] P. Calyam, D. Krymskiy, M. Sridharan, and P. Schopis, “Active and passive measurements on campus, regional and national network backbone paths,” in *IEEE ICCCN*, 2005.
- [14] P. Calyam, C.-G. Lee, E. Ekici, M. Haffner, and N. Howes, “Orchestration of network-wide active measurements for supporting distributed computing applications,” *IEEE Transactions on Computers*, vol. 56, no. 12, pp. 1629–1642, 2007.