



Available at

www.ElsevierMathematics.com

POWERED BY SCIENCE @ DIRECT®

Computational Statistics & Data Analysis 45 (2004) 25–50

**COMPUTATIONAL
STATISTICS
& DATA ANALYSIS**

www.elsevier.com/locate/cstda

A diffusion model of roundtrip time

Stephan Bohacek^{a,*}, Boris Rozovski^b

^a*Department of Electrical and Computer Engineering, University of Delaware, USA*

^b*Department of Mathematics, University of Southern California, USA*

Received 18 January 2002; received in revised form 15 July 2002

Abstract

A diffusion model of the roundtrip time experienced by packets on the Internet is developed. The parameters of the diffusion system are modeled as Markov processes. The transition probabilities of the parameters are modeled with a mixture of Laplace distributions. The stationary distributions of the parameters are modeled as Normal distributions. The results show that the model is able to reproduce the probabilities observed in the Internet.

© 2003 Elsevier B.V. All rights reserved.

1. Introduction

The large majority of Internet communication is acknowledgment based. The sender transmits a data packet to the receiver. If the receiver receives the packet, it responds with an acknowledgement. The time until the acknowledgement arrives (assuming it arrives at all), is known as the *round-trip time*. There has been extensive work focused on the stationary distribution of round-trip time (Pointek et al., 1996; Mukherjee, 1994; Quarterman et al., 1994; Mills, 1983; Acharya and Saltz, 1996). However, little work has focused on the dynamic aspects of round-trip time. This paper develops a simple diffusion process that is able to reflect the behavior round-trip time observed in the Internet.

The importance of accurate models of round-trip time is well established. When a packet is sent and an acknowledgement never arrives, at some point the sender must declare a time-out, that is the sent packet or the acknowledgement has been dropped and an acknowledgement will never arrive. The decision as to when to declare a time-out, that is based on estimates of the round-trip time and is critical to maximizing

* Corresponding author.

E-mail addresses: bohacek@eecis.udel.edu (S. Bohacek), rozovski@cams.usc.edu (B. Rozovski).

TCP's¹ performance (Paxson and Allman, 2000). Congestion control schemes such as TCP-Vegas (Brakmo and Peterson, 1995) use the variation in round-trip time to estimate the congestion. Furthermore, some recently suggested congestion control mechanisms for multipath routing ignore out-of-order packets and only utilize time-out. In this case, the decision as to whether a time-out has occurred is of paramount importance and models of round-trip time play a critical role. Furthermore, in models of TCP found in Misra et al. (1999) and Padhye et al. (1998), the data sending rate depends on the round-trip time. Thus, accurate models of round-trip time are required for both the theoretical development of TCP as well as its implementation.

Additionally, there is hope that the Internet can be used for voice communication. One difficulty with such an approach is that data that transverses the Internet is subjected to time-varying delay, i.e., jitter. In order to fully understand the impact of this jitter and to develop countermeasures, the dynamics of the round-trip time must be understood.

Network tomography (Vardi, 1996; Tebaldi and West, 1998) and network monitoring also use dynamic models of round-trip time. Specifically, end users typically only have access to the edges of the network and are not able to directly assess the state of the network. Even network operators only have access to routers in their control and are unable to determine the state of neighboring ISPs. Furthermore, most routers are only able to forward packets and do not monitor traffic. Thus, there is a need to determine the state of the network based only on external measurements (i.e., measurements made at the end hosts and not at the internal router). Beyond the fact that most routers are not able to closely monitor traffic, another advantage of monitoring connections is that a set of connections can provide information about a far larger set of links. Techniques using connection information to understand link properties are developed in Duffield et al. (2001). In this paper, the state of the network is assessed by determining the parameters of a dynamical model the round-trip time. By observing the model parameters, the network can be monitored and anomalies, either failures or network attacks, can be detected. This application of the work presented has special importance since there is an increase in network attacks that do not attack hosts directly, but instead attack the network by either attacking routers or links.

To elucidate this last possibility, imagine a distributed denial of service attack on a single link. The attack could be realized by coordinating a diverse set of hosts to send large UDP packets to another diverse set of hosts. If these hosts are picked carefully, then the packets would traverse a common link causing considerable congestion along this link. Note that the end hosts themselves would not be under attack and would not necessarily be aware of the on-going attack. The only way that such an attack could be detected is by monitoring the link itself or monitoring connections that utilize this link. Such monitoring would detect the attack by noticing a jump in congestion and change in delay. Furthermore, it is likely that since the congestion is caused by UDP packets and not the typical TCP connections, the dynamics of the latency over this attacked

¹ More than 85% of the traffic on the Internet uses the TCP protocol to determine how fast data packets should be sent.

link would differ from the dynamics typically found over other links. Indeed, [Jonckheere et al. \(2002\)](#) has found some difference between the network dynamics under attack and under normal conditions.

In this paper two related models are considered. First, a diffusion process with time-invariant parameters is considered (Section 3). In this case, the system has six scalar parameters. Techniques to estimate these parameters are developed (Section 3.1). While such time-invariant models are valid over short time intervals, in which the parameter variation is typically very small, the model is too simple to capture substantial changes in the dynamics of the round-trip time. For example, at night, the round-trip time may vary by less than 5 ms, while over the same connection during the afternoon, the round-trip time may vary by more than 30 ms. To account for these varying environments, a second model is developed where the parameters of the diffusion process vary (Section 5). In this case, the dynamics of the parameters are of interest. These dynamics are modeled as Markov processes on a finite space (Section 5.2). The transition probabilities for these parameters are well modeled by a mixture of Laplace distributions, while the stationary distributions of the parameters are modeled as Normal distributions.

2. Measurement procedure

The results presented here are part of a large ongoing data collection and modeling effort. The purpose of this paper is not to discuss the spatial variability of the statistics (i.e., how they vary from one geographical point to the next), but to develop models of the dynamics of round-trip time. To exclude distractions by the sources of the data, all data utilized in this paper is from two connections, Los Angeles to San Jose and Los Angeles to Tampa, Florida. Packets going from Los Angeles to San Jose and back pass through 32 routers, while packets going from Los Angeles to Tampa and back traverse 18 routers. These connections and the data collected represent what we have found to be typical cases. From this data not only the best, but rather the full range of results is shown. Therefore, most plots include the case where the results are not as strong. The rationale is to give an idea of the variation of the results. However, there are some outlier data points such as those shown in Fig. 24 that are not discussed. These occasional abrupt events are worthy of their own study and are not discussed here.

The data analyzed was collected during March 2001. During this time an ICMP echo request packet was sent every 10 ms for 100 s, then no data was sent for 80 s and then the process was repeated. The packets had 64 bytes of payload, making them compatible with many voice applications. Furthermore, when using small packets, the measurement process should not have a large effect on the network that is being measured. However, the size of the packets is clearly smaller than typical TCP data packets. The difference in the behavior of the round-trip time dynamics when sending small packets as compared to when sending larger packets is not known. Some small experiments indicate that the dynamics on these connections did not show much variation when small packets were sent as compared to large packets.

3. A stationary model

The round-trip time experience by a packet sent at time t from source S to destination D and back can be decomposed into a fixed, non-varying component due to propagation time and transmission time and a time-varying component that is principally due to queuing delay. Denote this time-varying component as R_t . A diffusion model for R_t is

$$dR_t = \frac{\sigma^2}{2} R_t^{\rho-1} \left(\frac{(\delta R_t^\delta + \gamma \beta^\gamma R_t^\gamma)}{(R_t^\delta + \beta^\gamma R_t^\gamma)} + \rho - 2\phi \ln(R_t) \right) dt + \sqrt{\sigma^2 R_t^\rho} dW_t. \quad (1)$$

Note that for $\beta = 0$ and $\rho = 1$, the above reduces to

$$dR_t = \frac{\sigma^2}{2} (\delta + 1 - 2\phi \ln(R_t)) dt + \sqrt{\sigma^2 R_t} dW_t,$$

which is similar to the Black–Derman–Toy diffusion model for short-term interest rates (Black et al., 1990).

Applying Theorem 1 in Grigelionis (2000), it can easily be shown that (1) is an ergodic H-diffusion and with unique invariant distribution given by

$$h(r) = C(r^\delta + \beta^\gamma r^\gamma) \exp(-\phi(\ln(r))^2), \quad (2)$$

where C is such that $\int_0^\infty h(r) dr = 1$.

Notice that σ and ρ play no role in the invariant density. These parameters only effect the transition probabilities. Let $p_t(y, x) dx$ be the probability of moving from $R_0 = y$ to $x \leq R_t \leq x + dx$. Then p satisfies the Kolgomorov's forward equation

$$\frac{\partial}{\partial t} p_t(y, x) = \frac{\sigma^2}{2} \left(-\frac{\partial}{\partial x} (\mu(x) p_t(y, x)) + \frac{\partial^2}{\partial x^2} (x^\rho p_t(y, x)) \right), \quad (3)$$

where $\mu(x) = x^{\rho-1}((\delta x^\delta + \gamma \beta^\gamma x^\gamma)/(x^\delta + \beta^\gamma x^\gamma) + \rho - \phi 2 \ln(x))$. Since the process is ergodic, $p_t(y, x) \rightarrow h(x)$ where h is given by (2). Thus, σ controls how fast the conditional density converges to the invariant density. The role of ρ is not so easily seen. This separation between the roles of the parameters $(\beta, \delta, \gamma, \phi)$ and (σ, ρ) is helpful in parameter estimation. In particular, we are able to first estimate $(\beta, \delta, \gamma, \phi)$ from the stationary density. These parameters can be found using an optimization technique such as quasi-Newton. Then, once these parameters are determined, σ and ρ can be found by examining the transition probabilities. This second optimization is far more difficult than the first as it requires repeatedly solving Kolgomorov's forward equation. Since this second optimization has only two unknown parameters, as opposed to six, this separation reduces the number of times equation (3) must be solved.

An implicit assumption in (1) is that R_t is a Markov process. Specifically, it is assumed that for $0 < \tau_1 < \tau_2 < \dots$,

$$P(R_t \in A | R_{t-\tau_1}, R_{t-\tau_2}, \dots, R_{t-\tau_n}) = P(R_t \in A | R_{t-\tau_1}). \quad (4)$$

To a great extent, the assumption that round-trip time is Markov is supported by the data. Fig. 1 shows the conditional distribution of the latency. The figure indicates that

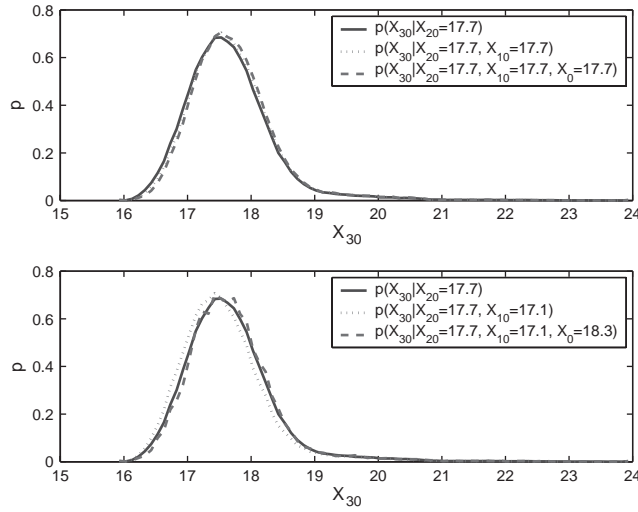


Fig. 1. The above plots show the conditional distribution of the round-trip time. The distributions are conditioned on either one, two or three past values.

the distribution depends little on the past given the present, i.e., the Markov condition (4) appears to be a reasonable approximation.

3.1. Parameter estimation for the stationary model

There are many possible approaches to estimating the parameters. A very popular approach to parameter estimation is the method of maximum likelihood. While the maximum likelihood parameter estimators are tractable in this setting, they do not yield very good estimates. In particular, these estimates appear to overestimate all parameters. The reason for this poor performance is that the model (1) is not the actual system and the method of maximum likelihood performs poorly in the face of even a few data points that are not explained by the model. Instead of maximizing the likelihood, we are interested in the parameters that best fit the data. To this end we seek to minimize the L^1 distance between the observed probabilities and the model. Note that this norm has the desirable property (Devroye, 1987) that when f and g are densities,

$$\int |f - g| = 2 \sup_A \left| \int_A f - \int_A g \right|. \quad (5)$$

In this paper, we are able to make use of large amounts of data. Hence, in most cases, the histograms provide accurate estimates of the actual density. Thus, the L^1 error between the fitted density and the observed histogram can be easily translated into the size of the error in terms of probability.

Define

$$h_{\beta, \delta, \gamma, \phi}(x_i) = C(r^\delta + \beta^\gamma r^\gamma) \exp(-\phi(\ln(r))^2).$$

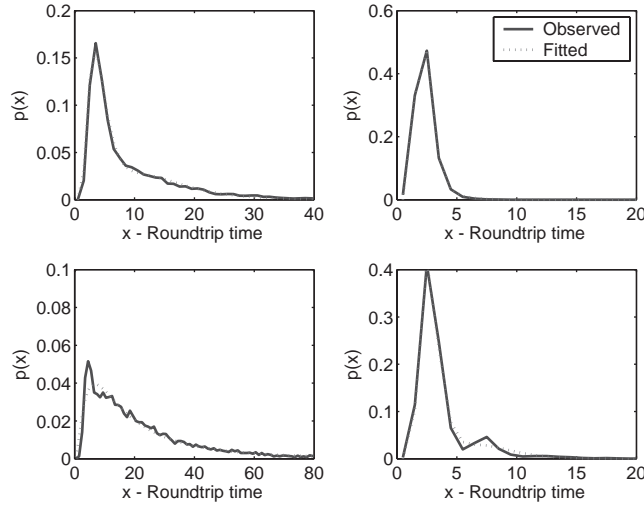


Fig. 2. Observed and fitted stationary distributions of the round-trip time. The top left-hand plot shows the round-trip time between Los Angeles and San Jose at 3:30 p.m. on March 26, 2001. The right-hand plot show the round-trip time for the same connection at 1 a.m. on March 27, 2001. The lower two plots show the round-trip times at the same times as the upper plots but for the Los Angeles to Tampa connection.

Define equally spaced bins centered at $x_i = \Delta i$. The observed density is

$$p_e(x_i) := \frac{1}{\Delta \cdot N} \sum_{k=1}^N 1_{\{x_i - \Delta/2 \leq R_k < x_i + \Delta/2\}},$$

where $1_{\{x_i - \Delta/2 \leq R_k < x_i + \Delta/2\}}$ is one if R_k is in the i th bin and zero otherwise and where N is the number of observations. Fitting cost is defined as

$$J_1(\beta, \delta, \gamma, \phi) := \sum_{i=1}^{\infty} |p_e(x_i) - h_{\beta, \delta, \gamma, \phi}(x_i)| \Delta. \quad (6)$$

The objective is to solve $\min_{\beta, \delta, \gamma, \phi} J_1(\beta, \delta, \gamma, \phi)$. This four-dimensional minimization can be solved using numerical methods such as quasi-Newton (Luenberger, 1984). While this optimization appears to be non-convex, with carefully chosen initial parameters the optimization usually converged quickly. Note that if $N \rightarrow \infty$ and if $\Delta \rightarrow 0$, the minimization becomes

$$\min_{\beta, \delta, \gamma, \phi} \int_0^{\infty} |p_e(x) - h_{\beta, \delta, \gamma, \phi}(x)| dx.$$

As discussed in the introduction, round-trip time measurements were made for a connection between Los Angeles and San Jose and between Los Angeles and Tampa. Fig. 2 shows the observed and the minimum distance fit probability density of the variable part of the latency. The fixed delay was subtracted and will continue to be for the remainder of the paper. Hence, a round-trip time of zero means that the packet experienced the shortest round-trip time observed. Fig. 2 shows the round-trip time for

Table 1
Estimated parameter values for the stationary distribution

Connection	β	δ	γ	ϕ	L^1
Los Angeles–San Jose (afternoon)	0.33	6.3	13	2.6	0.09
Los Angeles–San Jose (night)	0.44	7.8	16	5.4	0.03
Los Angeles–Tampa (afternoon)	0.01	2.1	369	0.5	0.126
Los Angeles–Tampa (night)	0.41	9.8	21	5.4	0.115

The right-hand column is the L^1 norm of the difference between the fitted density and the observed histogram.

both connections during the day and night. These plots are typical. Note that larger round-trip times are experienced during the afternoon. This increase in round-trip time is due to increased congestion during the afternoon, resulting in longer queueing delays. The parameter values for these connections and the L^1 fitting errors are shown in Table 1.

3.2. Estimating the diffusion parameters

As in the previous section, a minimum distance fit is used to estimate the parameters σ and ρ . As discussed above, σ and ρ control the distribution of the increments $R_{t+T} - R_t$. Specifically, if β, δ, γ , and ϕ are given, then σ and ρ are uniquely identified by the functions $p(R_T | R_0 = r_0)$ for different r_0 . Thus, σ and ρ can be found by observing R_T whenever $|R_0 - r_0|$ is small where T is the sampling period. In this paper, $T = 10$ ms. We approximate the density $p(R_T | R_0 = r_0)$ by subdividing the positive real line into bins as above and observing the distribution $P_e(x_i - \frac{\Delta}{2} \leq R_T < x_i + \frac{\Delta}{2} | R_0 = r_0)$. Once P_e is determined, the following minimization can be solved:

$$\min_{\sigma, \rho} \sum_{i=1}^{\infty} \left| P_e \left(x_i - \frac{\Delta}{2} \leq R_T < x_i + \frac{\Delta}{2} \middle| R_0 = r_0 \right) - P \left(x_i - \frac{\Delta}{2} \leq R_T < x_i + \frac{\Delta}{2} \middle| R_0 = r_0, \beta, \delta, \gamma, \phi, \sigma, \rho \right) \right|,$$

where

$$\begin{aligned} P \left(x_i - \frac{\Delta}{2} \leq R_T < x_i + \frac{\Delta}{2} \middle| R_0 = r_0, \beta, \delta, \gamma, \phi, \sigma, \rho \right) \\ = \int_{x_i - \frac{\Delta}{2}}^{x_i + \frac{\Delta}{2}} p_T(x | r_0, \beta, \delta, \gamma, \phi, \sigma, \rho) dx \end{aligned}$$

with $p_T(x | r_0, \beta, \delta, \gamma, \phi, \sigma, \rho)$ found by solving (3).

One drawback to this approach is that there is no clear choice of r_0 . Indeed, for a single r_0 , there may be many pairs (σ, ρ) that achieve the above minimum. To resolve

this ambiguity, a set $\{r_{o_j} : 1 \leq j \leq M\}$ of r_o is chosen and a single pair (σ, ρ) is found for all of these r_o . Specifically, define the following loss function:

$$J_2(\sigma, \rho) := \sum_{j=1}^{M-1} (r_{o_{j+1}} - r_{o_j}) p(r_{o_j} | \beta, \delta, \gamma, \phi) \\ \times \sum_{i=1}^{\infty} \left| P_e \left(x_i - \frac{\Delta}{2} \leq R_T < x_i + \frac{\Delta}{2} \middle| R_0 = r_{o_j} \right) \right. \\ \left. - P \left(x_i - \frac{\Delta}{2} \leq R_T < x_i + \frac{\Delta}{2} \middle| R_0 = R_{o_j}, \beta, \delta, \gamma, \phi, \sigma \rho \right) \right|.$$

The best σ, ρ are found by solving $\min_{\sigma, \rho} J_2(\sigma, \rho)$. Note that the weighting $(r_{o_{j+1}} - r_{o_j}) p(r_{o_j} | \beta, \delta, \gamma, \phi)$ places more weight the more likely r_o . Furthermore, as $M \rightarrow \infty$, $\sup_j (r_{o_{j+1}} - r_{o_j}) \rightarrow 0$, and the number of observation $N \rightarrow \infty$, we have,

$$J_2(\sigma, \rho) \rightarrow E \left(\left| P_e \left(x_i - \frac{\Delta}{2} \leq R_T < x_i + \frac{\Delta}{2} \middle| R_0 = r_{o_j} \right) \right. \right. \\ \left. \left. - P \left(x_i - \frac{\Delta}{2} \leq R_T < x_i + \frac{\Delta}{2} \middle| R_0 = r_{o_j}, \beta, \delta, \gamma, \phi, \sigma \rho \right) \right| \right).$$

However, we have not found a significant benefit in taking M large. The results presented here use $M = 5$ and r_{o_j} is spaced at the $j \times 16\frac{2}{3}$ percentile. This minimization was carried out by gridding the parameter space and evaluating J_2 at each point. We found that J_2 is non-convex. However, as can be seen by examining (3), $p(x_t | y_0, \rho, 2\sigma) = p(x_{4t} | y_0, \rho, \sigma)$. Furthermore, $p(x_t | y_0, \rho, \sigma)$ can be approximately represented as $p(x_t | y_0, \rho, \sigma) = A p(y_0)$ and $p(x_{2t} | y_0, \rho, \sigma) = A^2 p(y_0)$ where A is a matrix. Thus, if the dimension of A is small, then evaluating J_2 for a fixed ρ and many σ is not significantly more difficult than evaluating J_2 for a fixed ρ and σ .

Figs. 3–6 show some observed and fitted transition probabilities. Different σ and ρ were found for each time frame but, within a specific time frame, the same σ and ρ were used for all initial conditions. Table 2 shows the parameters found for the times shown in the figures.

While the quality of fit shown in Figs. 3–6 is typical, the quality is not always so good. Fig. 7 shows the observed and fitted transition probabilities for March 26 at 6:40 p.m. for the Los Angeles to San Jose connection. Note that the observed and fitted stationary density appear to match. Indeed, the L^1 error is 0.21. However, the observed and fitted transition probabilities appear to differ with the average L^1 cost for the transition probabilities was found to be 0.42. This difference is most apparent when the initial conditions, R_o , is large. Note the multiple modes in the transition probabilities. If the parameters are fixed, the model (1) cannot yield such a transition probability. On the other hand, if the parameters are allowed to rapidly switch, then it might be possible to yield such behavior. However, in this paper we only consider slowly varying parameters. The possibility of rapidly varying parameters is left for future investigations. It is also possible that the observed transition probabilities are not very accurate. This inaccuracy could be due to the fact that for this observation,

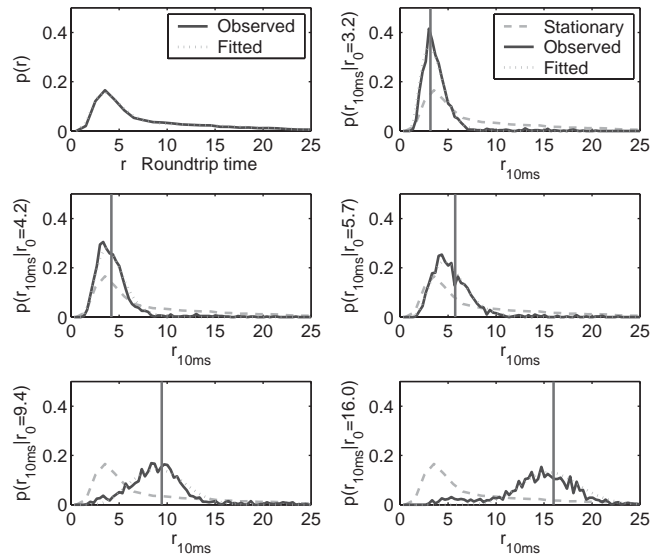


Fig. 3. Observed and fitted transition probabilities for the Los Angeles to San Jose connection on March, 26 at 3:30 p.m. The *upper left* shows the stationary density. The remaining plots show the observed transition probability (solid line), the fitted transition probability (dotted line) and, for reference, the stationary density. The vertical solid line indicates the initial condition.

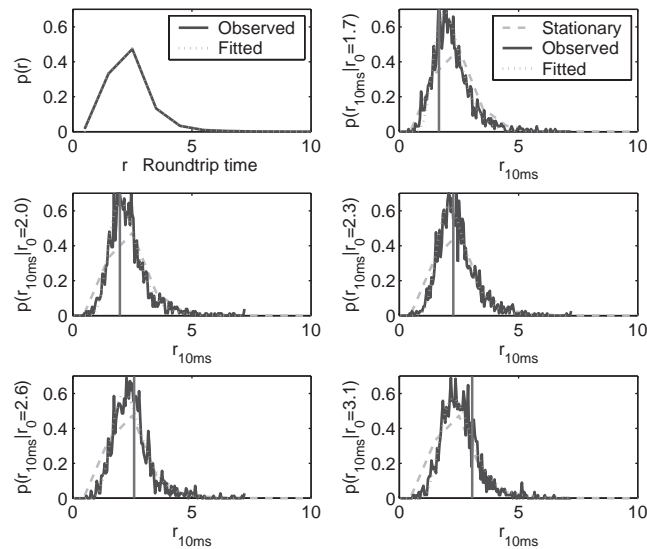


Fig. 4. Observed and fitted transition probabilities for the Los Angeles to San Jose connection on March, 26 at 1 a.m.

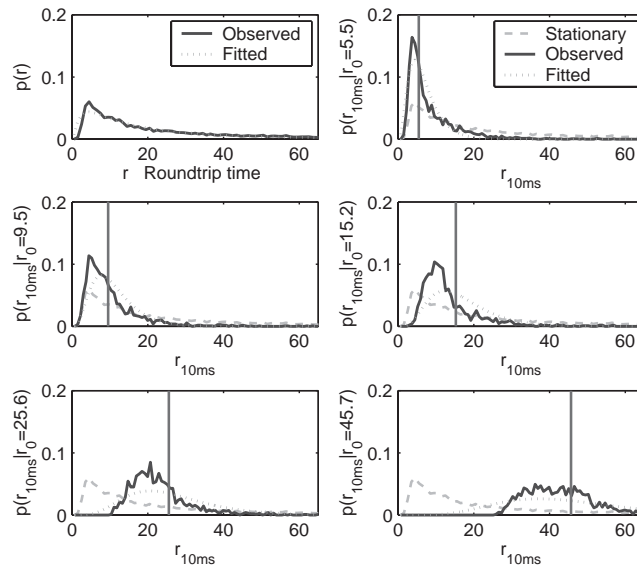


Fig. 5. Observed and fitted transition probabilities for the Los Angeles to Tampa connection on March, 26 at 3:30 p.m.

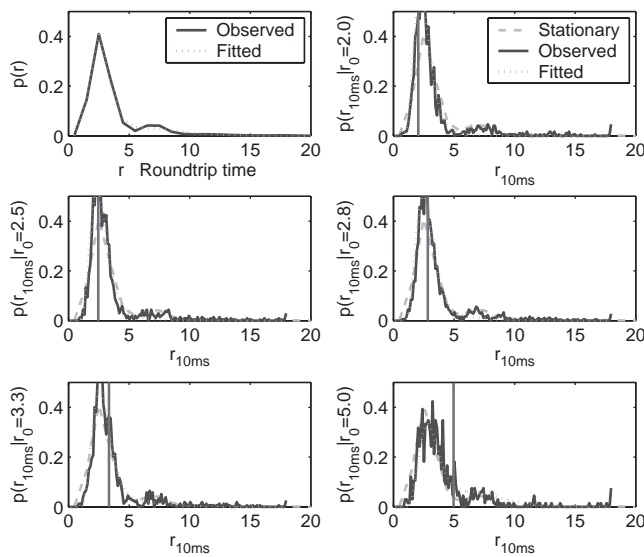


Fig. 6. Observed and fitted transition probabilities for the Los Angeles to Tampa connection on March, 26 at 1 a.m.

Table 2
Estimated parameters for the transition probabilities

Connection	σ	ρ	L^1 Cost
Los Angeles–San Jose (afternoon)	0.0125	0.8	0.14
Los Angeles–San Jose (night)	0.0045	2.2	0.15
Los Angeles–Tampa (afternoon)	0.067	1.5	0.33
Los Angeles–Tampa (night)	0.043	2.3	0.18

The right-hand column shows the L^1 norm of the difference between the fitted density and the observed histograms.

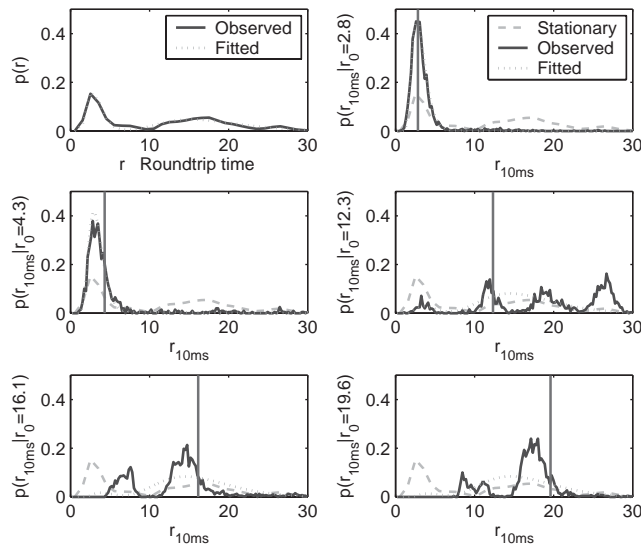


Fig. 7. Observed and fitted transition probabilities for the Los Angeles to San Jose connection on March, 26 at 6:40 p.m.

the support of the stationary density is large. Hence, each particular round-trip time was not observed very often. This is particularly true for large x_0 . Therefore, the variance of the estimate of $P(R_{10} - R_0 | R_0)$ is large for large R_0 .

4. Goodness of fit

The above techniques were applied every 3 min during the month of March, 2001. The data collected over a 100 s period was analyzed and new model parameters were found. Figs. 8 and 9 show an estimate of the L^1 error for the Los Angeles–San Jose and Los Angeles–Tampa connections, respectively. These figures also show the average round-trip time during the period. The L^1 error in modeling the stationary distribution

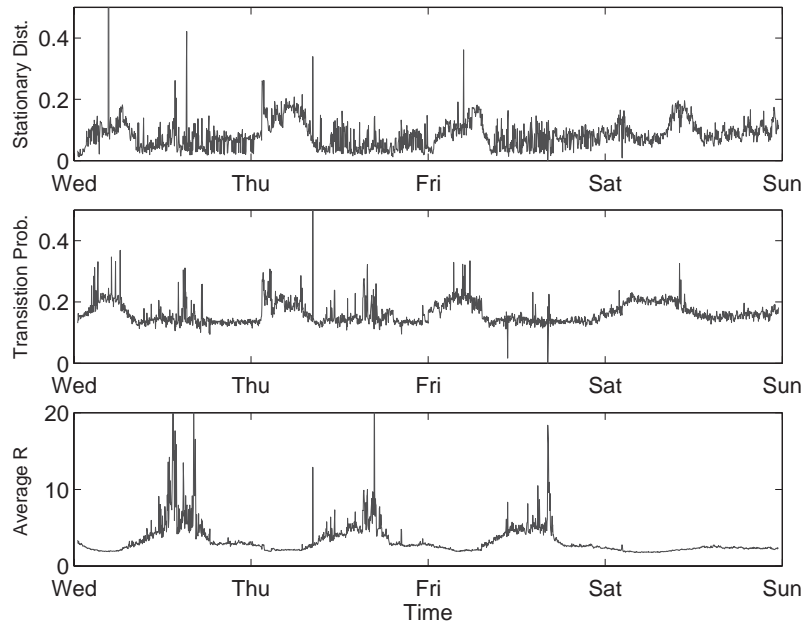


Fig. 8. Parameter goodness of fit for the Los Angeles–San Jose connection. The upper plot shows a typical time series of J_1 , the estimated error in modeling the stationary distribution, while the middle plot shows a typical time series of J_2 , the estimated error in modeling the transition probabilities. The lower plot shows the average round-trip time over 3 minute windows.

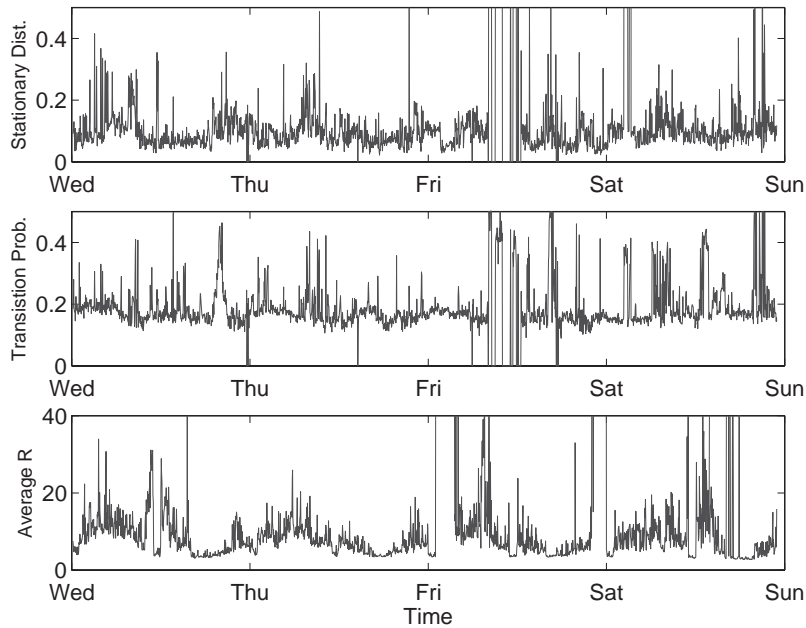


Fig. 9. Parameter goodness of fit for the Los Angeles–Tampa connection.

Table 3
Quality of fit for different connections and different times of day

Connection	Stationary Distribution				Transition Probabilities			
	Afternoon	Evening	Night	Morning	Afternoon	Evening	Night	Morning
LA–SJ	0.073	0.067	0.11	0.093	0.15	0.14	0.20	0.18
LA–Tampa	0.10	0.12	0.17	0.12	0.20	0.20	0.20	0.17

The measure of fit is the L^1 norm of the difference between the fitted density and the observed histogram.

is given by the minimum value of J_1 above. While the L^1 error in modeling the transition probabilities is given by the minimum value of J_2 . Note that for the Los Angeles–San Jose connection, the diurnal variation in the round-trip time is clearly visible. Furthermore, it can be seen that the quality of fit is better during the more congested daytime than the lightly congested nighttime. Table 3 shows the mean quality of fit for different times of the day and different connections.

5. Dynamic model and parameter dynamics

In Section 3.1, it was assumed that the model parameters remain constant. However, examining Figs. 3–6, it is clear that the parameters vary with time. Figs. 10 and 11 show a time series of the estimated parameters β, δ, γ , and ϕ along with the average round-trip time over the measured period. In this section, models are developed of parameter variation. Since the parameters are not directly observable, the variation in the parameters cannot be directly determined. Instead, we make the assumption that the parameters slowly vary and that the parameters can be taken to be nearly constant over a short time interval. In this case, the methods described in Section 3 can be used to estimate the parameters over short time intervals. These estimates are then used to infer the dynamics of the parameters. To make clear the difference between observed distributions and inferred distributions, we denote P_e as the observed quantity (e for ensemble) and \hat{P} as the inferred probability.

Models of parameter dynamics are developed and justified in the next two subsections. In Section 5.1, the parameters are divided into three groups. It is assumed that the parameters in one group are independent of the parameters in the other groups. In Section 5.2, the models of parameter variation are developed for each parameter group. Each model consists of a model of the stationary distribution and a model of the parameter transition.

While much care was put into the selection of the model for the round-trip time, the models for the round-trip time model parameters are more crude. One reason for this crudeness is that the round-trip time model is already somewhat complex and complex models for the parameters would result in very complex model overall. Another reason for the simplified models is that we anticipate that parameters models will be used in a Bayesian framework. Thus, parameter initial distributions and transition distributions are

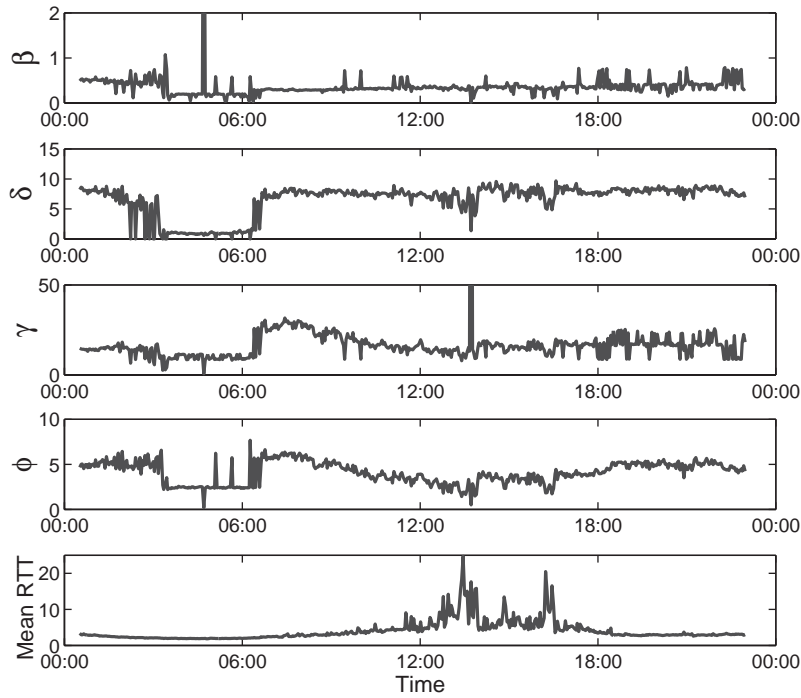


Fig. 10. Time series of inferred parameter values for the Los Angeles to San Jose connection from March, 28 to March 31, 2001.

required. However, the observed dynamics of the round-trip time has a strong impact on the parameter estimates and tends to compensate for errors in the distributions.

5.1. Parameters correlation

The parameters dynamics will be modeled as Markov processes. Thus, it is assumed that the parameters obey

$$p(\phi_t | \phi_{t_0}) = p(\phi_t | \phi_{t_0}, \phi_{t-1}, \dots, \phi_{t-n}).$$

This Markov assumption is only an approximation and is unlikely to exactly hold. Fig. 12 shows the autocorrelation of the parameter increments for the Los Angeles–Tampa connection. The Los Angeles–San Jose connection gives nearly identical results. Note that there is little correlation after one time step. The correlation at one step indicates that a low order AR model might be appropriate. However, in order to develop a simple model, the Markov assumption is made.

Another modeling assumption is that some parameters are dependent while others are independent. The upper plot in Fig. 13 shows the observed pairs of the parameters (ϕ, δ) . There is a clear affine relationship between ϕ and δ . Fig. 14 is similar to 13

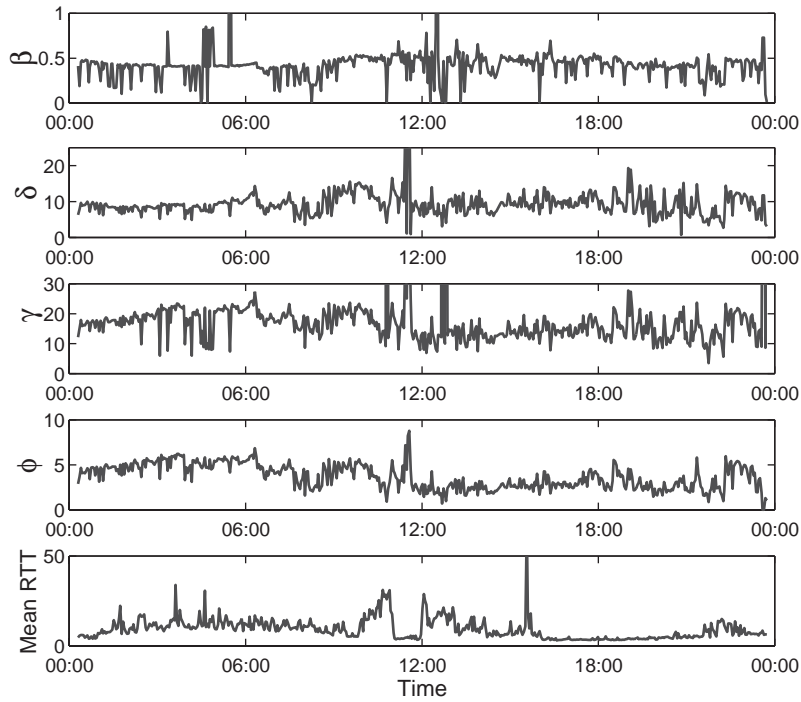


Fig. 11. Time series of inferred parameter values for the Los Angeles to Tampa, Florida connection from March, 28 to March 31, 2001.

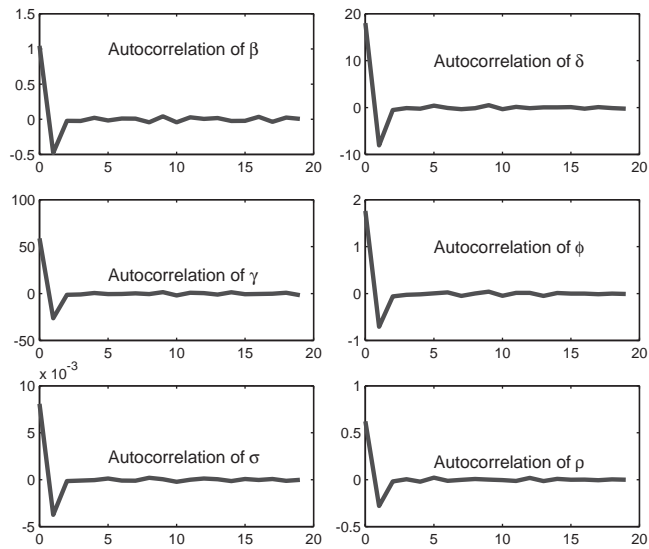


Fig. 12. Autocorrelations of the parameters.

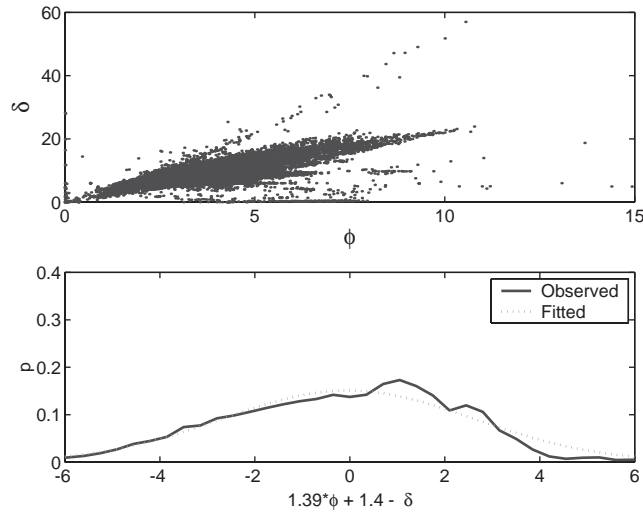


Fig. 13. *Upper*: Observed parameter pairs (ϕ, δ) . *Lower*: Residual error of the best least square affine predictor $\delta = a + b\phi$ (solid line), and a normal distribution fit of the residual error (dashed line).

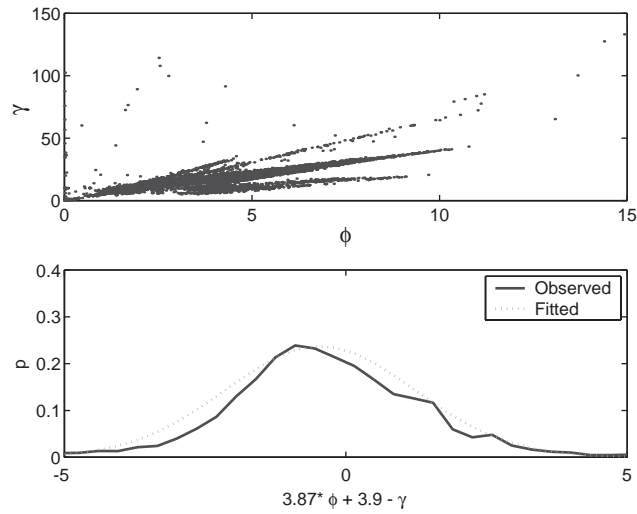


Fig. 14. *Upper*: Observed parameter pairs (ϕ, γ) . *Lower*: Residual error of the best least square affine predictor $\gamma = c + d\phi$ (solid line), and a normal distribution fit of the residual error (dashed line).

but for the parameters (ϕ, γ) . The parameter dependence is model as

$$\begin{aligned}\delta &= a + b\phi + n_\delta, \\ \gamma &= c + d\phi + n_\gamma,\end{aligned}\tag{7}$$

Table 4

The least squares parameter estimates of the parameters that related ϕ to δ and ϕ to γ , the parameters of the Gaussian approximation of the residual error and the L^1 difference between the histogram of the residual error and the Gaussian approximation

Connection	a	b	μ	σ	L^1	c	d	μ	σ	L^1
Los Angeles–San Jose	1.4	1.39	−0.14	2.4	0.16	3.9	3.0	−0.45	1.68	0.31
Los Angeles–Tampa	2.1	2.1	0.3	2.3	0.43	3.7	3.7	1.3	3.7	0.37

where a , b , c , and d are constants and n_δ and n_γ are normally distributed. The scalar parameters are found by minimizing the least squares error and the parameters of the normal distributions are found by minimizing the L^1 norm of the difference between the residual error distribution and the fitted normal distribution. Table 4 shows the parameters found for the two connections analyzed as well as the L^1 distance between the residual error and the fitted normal densities. Because of this dependence, the dynamics of δ and γ are not investigated. Only the dynamics of ϕ is modeled and (7) is utilized to find δ and γ .

Remark 1. Note that Figs. 14 and 13 indicate that a simple linear relationship may not be appropriate. It appears that a model that switches between linear relationships might perform better. This approach will be left for future work.

The data indicates that there is a deterministic relationship between σ and ρ . Specifically, the data indicates that for the large majority of the parameters found obeys $\sigma\rho = c$, where $c = 0.01$ for the Los Angeles–San Jose connection and $c = 0.1$ for the Los Angeles–Tampa connection. Fig. 15 shows the inferred relationship between σ and ρ . Note that except for some points along the line $\rho = 0$, the fitted relationship matches exactly the inferred relationship. Because of this dependence, only the dynamics of σ will be investigated.

In light of the above results, the parameters are divided into three groups; (δ, γ, ϕ) , β and (σ, ρ) . The variation of the parameters in distinct groups is assumed to be independent. However, this is a simplifying assumption and is not entirely true. For example, the inferred correlation coefficients between the parameters ϕ , β and σ were found to be $\rho_{\beta, \phi} = 0.44$, $\rho_{\beta, \sigma} = 0.35$, $\rho_{\phi, \sigma} = 0.55$.

5.2. Models of parameter variation

To account for the variation in parameters, the model for latency (1) can be extended to

$$dR_t = \frac{\sigma^2}{2} R_t^{\rho-1} \left(\frac{(\delta R_t^\delta + \gamma \beta^\gamma R_t^\gamma)}{(R_t^\delta + \beta^\gamma R_t^\gamma)} + \rho - \phi 2 \ln(R_t) \right) dt + \sqrt{\sigma^2 R_t^\rho} dW_t.$$

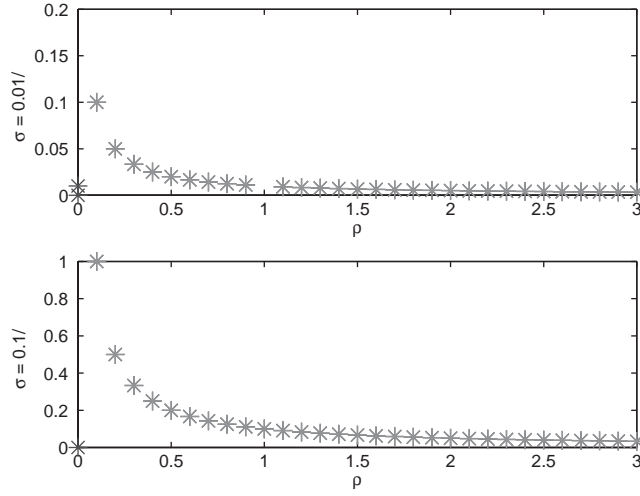


Fig. 15. The Inferred Parameter Pairs (σ, ρ) for the Los Angeles–San Jose connection (upper plot) and the Los Angeles–Tampa connection (lower plot). The plot only shows one set of points because the relation is exact for $\rho \neq 0$.

We model the parameter as Markov process and define a similar model

$$\begin{aligned} dR_t = & \frac{\sigma(\omega_t)^2}{2} R_t^{\rho(\omega_t)-1} \left(\frac{(\delta(\theta_t)R_t^{\delta(\theta_t)} + \gamma(\theta_t)\beta(v_t)^{\gamma(\theta_t)}R_t^{\gamma(\theta_t)})}{(R_t^{\delta(\theta_t)} + \beta(v_t)^{\gamma(\theta_t)}R_t^{\gamma(\theta_t)})} \right. \\ & \left. + \rho(\omega_t) - \phi(\theta_t)2 \ln(R_t) \right) dt + \sqrt{\sigma(\omega_t)^2 R_t^{\rho(\omega_t)}} dW_t, \end{aligned} \quad (8)$$

where θ_t , ω_t and v_t are independent Markov process. The θ process is associated with δ, γ, ϕ , the ω process is associated with σ and ρ , and the v process is associated with β . Thus we can interchange $\phi(\theta_t)$ and ϕ_t as well as $\sigma(\omega_t)$ and σ_t , etc. We assume that the processes θ , ω , and v take integers values $1, 2, \dots, M$. Set $\bar{\phi} := \max(\phi_t)$. The interval $[0, \bar{\phi}]$ is divided into M equally sized intervals. Thus

$$P(\theta_t = j | \theta_0 = i) \equiv P \left((j-1) \frac{\bar{\sigma}}{M} \leq \phi_t < j \frac{\bar{\sigma}}{M} \mid (i-1) \frac{\bar{\sigma}}{M} \leq \phi_0 < i \frac{\bar{\sigma}}{M} \right).$$

Defining $\bar{\sigma} := \max(\sigma_t)$ and $\bar{\beta} := \max(\beta_t)$, $P(\omega_t = j | \omega_0 = i)$ and $P(v_t = j | v_0 = i)$ can be defined similarly.

The transition probabilities of the Markov processes can be written as

$$p(\theta_t = j | \theta_0 = i) = (\exp(-Q^\theta t))_{i,j},$$

where $(\exp(-Q^\theta t))_{i,j}$ is the (i, j) component of the matrix $\exp(-Q^\theta t)$. We assume that the parameters are slowly varying. Specifically, we assume that $(-Q^\theta t)_{i,j}$ are small for

t on the order of a few minutes. Thus, for t in this range,

$$p(\theta_t = j | \theta_0 = i) = (\exp(Q^\theta t))_{i,j} \approx (I + Q^\theta t)_{i,j},$$

$$p(\omega_t = j | \omega_0 = i) = (\exp(Q^\omega t))_{i,j} \approx (I + Q^\omega t)_{i,j},$$

$$p(v_t = j | v_0 = i) = (\exp(Q^v t))_{i,j} \approx (I + Q^v t)_{i,j},$$

where I is the identity matrix. Thus, by estimating the parameter ϕ every T seconds (here, T was taken to be 180), it is possible to estimate $p(\theta_T = j | \theta_0 = i)$. Then Q^θ can be estimated via

$$Q_{i,j}^\theta \approx \frac{1}{T} (p(\theta_T = j | \theta_0 = i) - 1_{\{i=j\}}),$$

where

$$1_{\{i=j\}} = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases}$$

Likewise

$$Q_{i,j}^\omega \approx \frac{1}{T} (p(\omega_T = j | \omega_0 = i) - 1_{\{i=j\}}),$$

$$Q_{i,j}^v \approx \frac{1}{T} (p(v_T = j | v_0 = i) - 1_{\{i=j\}}).$$

A simplification can be made by assuming that

$$p_{a,b,c}(\phi_T | \phi_0) = (1 - c)\alpha \exp(-2\alpha|\phi_T - \phi_0|) + cb \exp(-2b|\phi_T - \phi_0|). \quad (9)$$

This model is the mixture of Laplace distributions and gives a reasonable fit of the data with great reduction in the complexity of the model. Note that the transition $\phi_T - \phi_0$ is taken to be independent of the initial value ϕ_0 . While the data supports this class of mixture of Laplace distributions, it appears that the parameters a, b, p should depend on ϕ_0 . However, we found this dependence to be minor and, hence, make the simplifying assumption that the distribution of the increment of the parameter is independent of the initial value. We make the same assumption for the other parameters β and σ .

In order to estimate the parameters a, b and c , define the cost

$$\begin{aligned} J_3(a, b, c) := & \sum_{i=1}^M \hat{p} \left(\phi_0 = i \frac{\bar{\phi}}{M} \right) \times \sum_{k=1}^M \left| \left((1 - c)a \exp \left(-2a \frac{\bar{\phi}}{M} |k - i| \right) \right. \right. \\ & \left. \left. + cb \exp \left(-2b \frac{\bar{\phi}}{M} |k - i| \right) \right) - \hat{p} \left(\phi_T = k \frac{\bar{\phi}}{M} \middle| \phi_0 = i \frac{\bar{\phi}}{M} \right) \right|, \end{aligned} \quad (10)$$

where

$$\begin{aligned} \hat{p} \left(\phi_T = k \frac{\bar{\phi}}{M} \middle| \phi_0 = i \frac{\bar{\phi}}{M} \right) := & \hat{P} \left(k \frac{\bar{\phi}}{M} - \frac{1}{2} \frac{\bar{\phi}}{M} < \phi_T \leq k \frac{\bar{\phi}}{M} \right. \\ & \left. + \frac{1}{2} \frac{\bar{\phi}}{M} \middle| i \frac{\bar{\phi}}{M} - \frac{1}{2} \frac{\bar{\phi}}{M} < \phi_0 \leq i \frac{\bar{\phi}}{M} + \frac{1}{2} \frac{\bar{\phi}}{M} \right) \end{aligned}$$

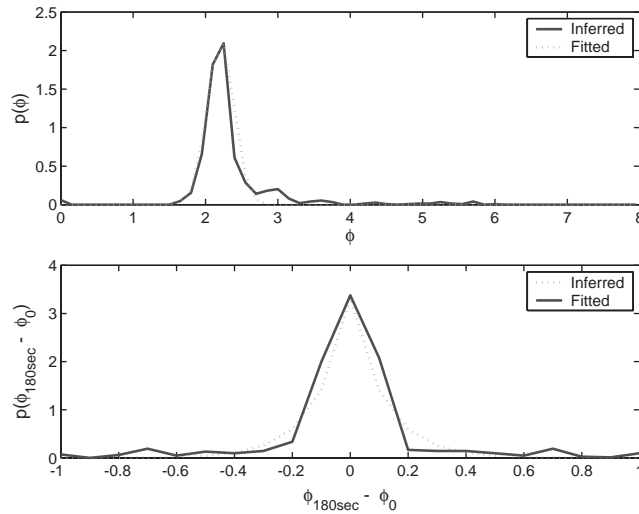


Fig. 16. *Upper*: The stationary distribution of ϕ for the Los Angeles–San Jose connection during times of very low congestion. *Lower*: The transition probabilities of ϕ for the Los Angeles–San Jose connection during times of very low congestion.

is the inferred transition probability and $\hat{p}(\phi_0 = i(\bar{\phi}/M)) := \hat{P}(i(\bar{\phi}/M) - \frac{1}{2}(\bar{\phi}/M) < \phi \leq i(\bar{\phi}/M) + \frac{1}{2}(\bar{\phi}/M))$ is the inferred probability. The parameters were found by solving $\min_{a,b,p} J_3(a,b,c)$.

The stationary distribution of the parameter values was also estimated. A normal distribution was chosen to model the stationary distribution. Again, the L^1 norm was used to estimate the parameters of the normal distribution.

By examining Fig. 10 it is clear that the average round-trip time is highly correlated to the parameters. In particular, for the Los Angeles to San Jose connection, very low average round-trip times coincide with small values of ϕ . To account for this correlation, two models are developed; one for typical values of average round-trip time and one for very small values of round-trip time. We considered any round-trip time less than 2.5 ms to be very small. These two model scenarios are only required for the Los Angeles–San Jose connection. The probability of transitioning from a state of very low round-trip time to a state of higher round-trip time is 0.051, while the probability of going from higher round-trip time to very low round-trip time is 0.056. Furthermore, the probability of being in the low round-trip time state is 0.47.

Figs. 16–18 show the inferred stationary density and transition probabilities for ϕ . Table 5 provides the parameter values found.

A similar model for β is also developed. Specifically, the transition probabilities were modeled using (9) with parameters a, b, c and modeling the stationary distribution with a Normal distribution with parameters μ and σ . Table 6 shows the parameter values and the minimum cost for the transition probabilities and the stationary distribution and Figs. 19–21. Note that the transitions densities are well modeled by a mixture of Laplace

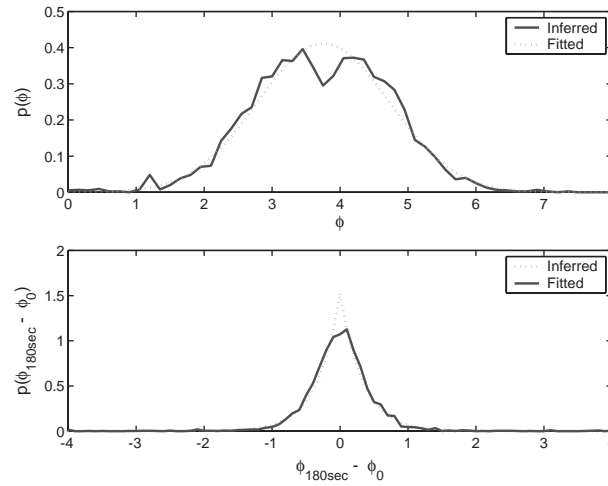


Fig. 17. *Upper*: The stationary distribution of ϕ for the Los Angeles–San Jose connection during times of normal congestion. *Lower*: The transition probabilities of ϕ for the Los Angeles–San Jose connection during times of normal congestion.

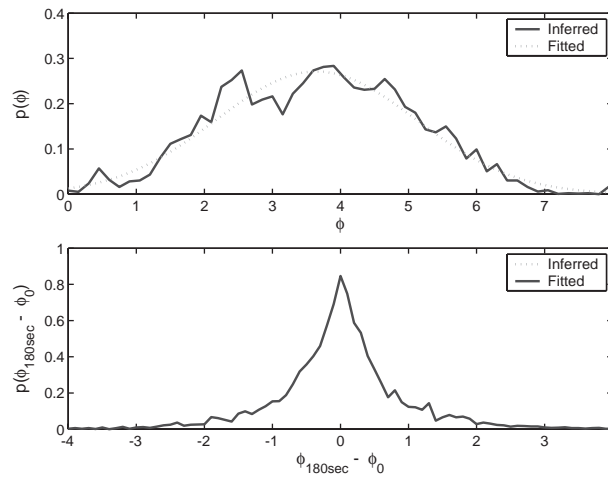


Fig. 18. *Upper*: The stationary distribution of ϕ for the Los Angeles–Tampa connection. *Lower*: The transition probabilities of ϕ for the Los Angeles–Tampa connection.

densities. However, the Normal distribution only provides a rough approximation of the stationary distribution.

The final parameter to be modeled is σ . Again, the transitions are modeled as a mixture of Laplace random variables and the stationary distribution was modeled as a

Table 5
Parameters of the stationary and transition probabilities of ϕ

Connection	a	b	c	L^1	μ	σ	L^1
Los Angeles–San Jose (Very Low Congestion)	0.005	4.3	0.77	0.37	2.2	0.19	0.28
Los Angeles–San Jose (Normal Congestion)	1.5	1.1	0.003	0.16	3.8	1.0	0.12
Los Angeles–Tampa	0.26	0.94	0.21	0.16	3.7	1.45	0.090

Table 6
Parameters of the stationary and transition probabilities of β

Connection	a	b	c	L^1	μ	σ	L^1
Los Angeles–San Jose (Very Low Congestion)	75.	3.5	0.56	0.35	0.17	0.03	0.75
Los Angeles–San Jose (Not Low Congestion)	26.	3.1	0.45	0.19	0.32	0.07	0.46
Los Angeles–Tampa	30.	2.7	0.57	0.21	0.40	0.082	0.37

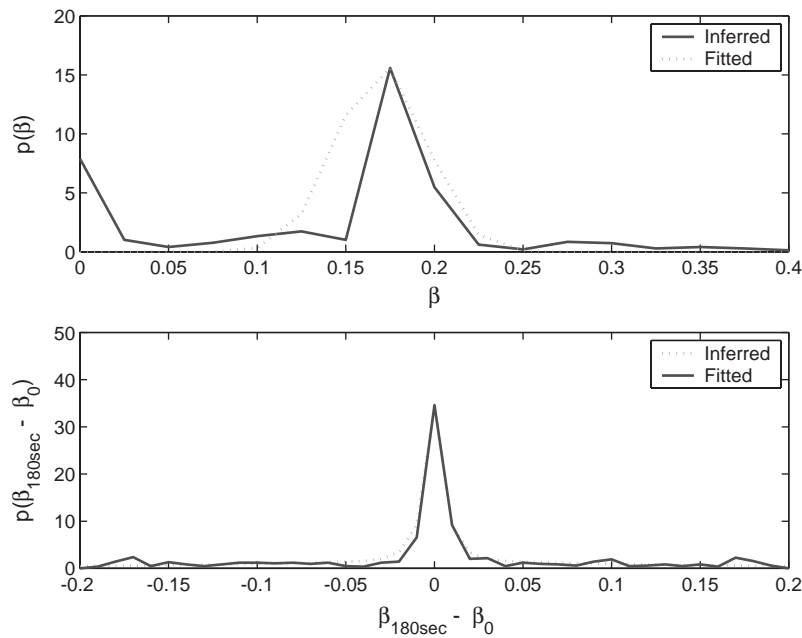


Fig. 19. For the Los Angeles–San Jose connection during periods of very low congestion the inferred and fitted stationary density of β (*upper*) and the inferred and fitted density of the increment of β (*lower*).

Normal distribution. Unlike the other parameters, there is no benefit to model the case of very low round-trip time differently from normal round-trip times. Table 7 shows the inferred parameters and the measure of fit. Figs. 22 and 23 show the inferred and fitted distributions of σ .

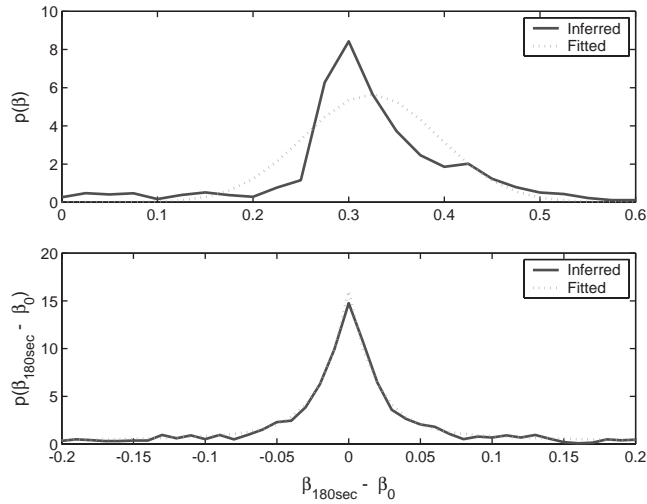


Fig. 20. For the Los Angeles–San Jose connection during periods of not very low congestion the inferred and fitted stationary density of β (*upper*) and the inferred and fitted density of the increment of β (*lower*).

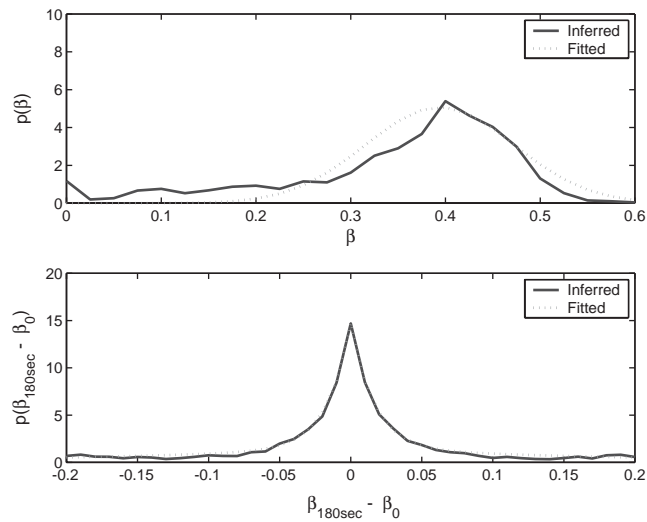


Fig. 21. For the Los Angeles–Tampa connection the inferred and fitted stationary density of β (*upper*) and the inferred and fitted density of the increment of β (*lower*).

6. Conclusion

A diffusion model of the round-trip time has been presented. Although this model has few parameters, it is able to reproduce the probabilities distributions observed.

Table 7
Parameters of the stationary and transition probabilities of σ

Connection	a	b	c	L^1	μ	σ	L^1
Los Angeles–San Jose	1452	152	0.35	0.23	0.0047	0.0012	0.35
Los Angeles–Tampa	243.	3.7	0.43	0.32	0.042	0.0040	0.28

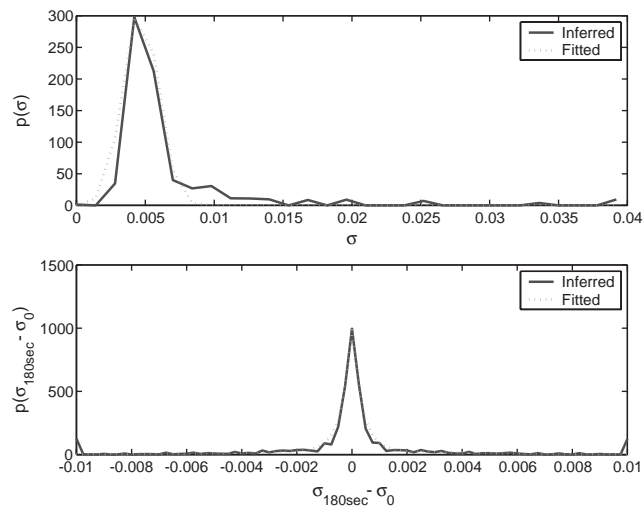


Fig. 22. For the Los Angeles–San Jose connection the inferred and fitted stationary density of σ (*upper*) and the inferred and fitted density of the increment of σ (*lower*).

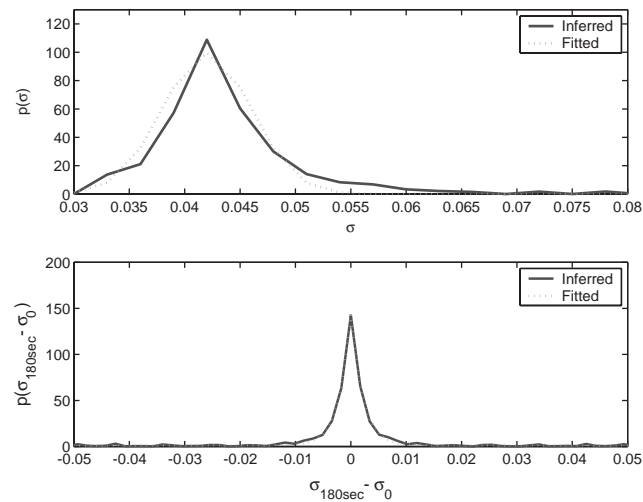


Fig. 23. For the Los Angeles–Tampa connection the inferred and fitted stationary density of σ (*upper*) and the inferred and fitted density of the increment of σ (*lower*).

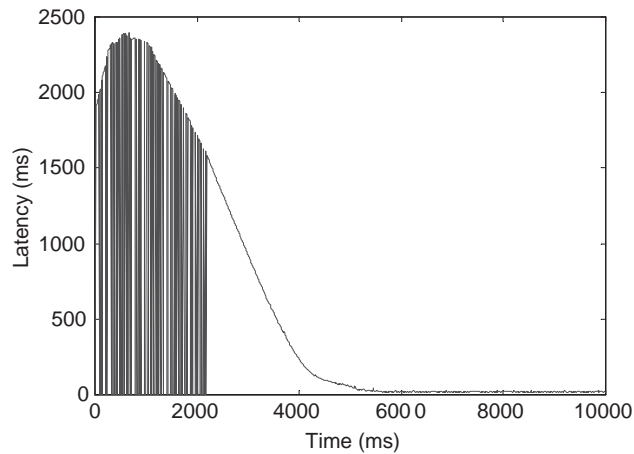


Fig. 24. Round-trip time observations not explained by the model. The round-trip time is set to 0 if a drop occurs.

This paper not only provides useful models for round-trip time, but shows that it is possible to model round-trip time. This result, to some degree, contradicts the popular opinion that the Internet is unmodelable. With such models the normal and abnormal behavior of the Internet can be detected. Such monitoring is the first step to providing network attack detection implementable from the edge hosts. However, many issues require further investigation. For example, these models do not accurately model the occasional large variation in latency. For example, Fig. 24 shows a time series of the latency. Clearly, some anomaly occurred some time before $t = 0$ and ended at around $t = 1$ s. In this figure the round-trip time is set to zero if a packet is dropped. Note that a large number of drops occurred until about 2 s after the trial commenced. After this time, the latency slowly decreased back to its ambient level. While such events do occur, a modeling approach different from the one presented here is required to capture such events. Future work will concentrate in two areas. The first will collect more data from different locations and determine if the models developed here are applicable for other connections. If the models are applicable, then the spatial variation of the model parameters can be determined. A second direction is to develop a more theoretical understanding of round-trip time. For example, the long tail of the round-trip time density is not predicted by simple queue models such as M/M/1.

References

- Acharya, A., Saltz, J., 1996. A study of internet round-trip delay. Technical Report CS-TR-3736, Department of Computer Science, University of Maryland.
- Black, F., Derman, E., Toy, W., 1990. A one-factor model of interest rates and its application to treasury bond options. *J. Financial Analysts* 46, 33–39.
- Brakmo, L., Peterson, L., 1995. TCP vegas: End to end congestion avoidance on a global internet. *IEEE J. Selected Areas Commun.* 13, 1465–1480.

- Devroye, L., 1987. *A Course in Density Estimation*. Birkhauser, Boston.
- Duffield, N.G., Horowitz, J., Presti, F.L., Towsley, D.F., 2001. Network delay tomography from end-to-end unicast measurements. In: IWDC, Taormina, Italy, pp. 576–595.
- Grigelionis, B., 2000. On generalized z-diffusions, Tech. Rep. Preprint 2000-24, Institute of Mathematics and Informatics.
- Jonckheere, E., Shah, K., Bohacek, S., 2002. Dynamic modeling of internet traffic for intrusion detection. In: American Control Conference, Anchorage, Alaska.
- Luenberger, D.G., 1984. *Linear and nonlinear programming*. Addison-Wesley, Reading, MA.
- Mills, D., 1983. Internet delay experiments.
- Misra, V., Gong, W., Towsley, D., 1999. Stochastic differential equation modeling and analysis of TCP-window size behavior. In: Performance '99, Istanbul, Turkey.
- Mukherjee, A., 1994. On the dynamics and significance of low frequency components of Internet load. *Internetworking: Res. Experience* 5 (4), 163–205.
- Paxson, V., Allman, M., 2000. Computing TCP's retransmission time. Request for Comments (RFC) 2988.
- Padhye, J., Firoiu, V., Towsley, D., Kurose, J., 1998. Modeling TCP throughput: A simple model and its empirical validation. In: ACM Sigcomm, September, Vancouver, British Columbia, Canada.
- Pointek, J., Shull, F., Tesoriero, R., Agrawala, A., 1996. Netdyn revisited: A replicated study of network dynamics. Technical Report CS-TR-3696, Department of Computer Science, University of Maryland.
- Quartermann, J., Carl-Mitchell, S., Phillips, G., 1994. Internet interaction pinged and mapped. In: Proceedings of INET'94, June, Prague, Czech Republic.
- Tebaldi, C., West, M., 1998. Bayesian inference on network traffic using link count data. *J. Amer. Statist. Assoc.* 93 (442), 557–572.
- Vardi, Y., 1996. Network tomography: Estimating source-destination traffic intensities. *J. Amer. Statist. Assoc.* 91 (433), 365–377.