

# AN SS-SVM APPROACH TO GENERATE SYNTHETIC NETWORK DELAYS

José A. Hernández and Iain W. Phillips  
Research School of Informatics,  
Loughborough University, Leics, LE11 3TU, UK.  
J.A.Hernandez@lboro.ac.uk  
I.W.Phillips@lboro.ac.uk

Javier M. Moguerza  
School of Engineering  
Universidad Rey Juan Carlos,  
28933 Móstoles (Madrid), Spain.  
Javier.Moguerza@urjc.es

## KEYWORDS

Network simulation, Performance modelling, Support Vector Regression, Slice Sampling algorithm, Internet delay synthesis.

## ABSTRACT

The research community has understood the need to measure packet-switched networks under real operation to obtain performance-related quantities that can help in the understanding of the uses, applications and protocols on the run. Empirical network metrics such as packet delay and loss, achievable throughput and traffic volumes traversing monitored points can potentially provide valuable information in the study, development and improvement of the next generation of the Internet.

However, when the cost of measuring is substantially high, it might result appropriate to generate synthetic measurements, similar in properties and behaviour to those collected under real network operation. Furthermore, the generation of synthetic data might be useful in the simulation of various types of network scenarios according to different profiles, and the study of the properties and characteristics of certain network events, along with their impact in both particular and global performance issues, without the risk of launching them on real scenarios.

In particular, the generation of network delays can set the simulation scenario to a first test and assesment on the behaviour of new software, protocols, applications and services, previous to their deployment in more real scenarios. This work describes a suitable technique to firstly capture the essentials of a given data set of network delays, and secondly generate a statistical clone of it for further simulation purposes. This technique is built up by combining the cutting-edge tool of Support Vector Machines and the recently proposed Slice Sampling algorithm.

## 1. MEASURING AND GENERATING SYNTHETIC MEASUREMENTS

In broad terms, the performance of a networked application depends on the cooperation and coordination of an immense number of interacting elements. Disparate applications such as gaming, videoconference, instant messaging or file sharing have coexisted in the same underlying infrastructure reliably. However, in spite of such

robustness, a deeper analysis of the network dynamics yet shows a large number of undesired behaviours, leading to a degradation in network performance. Previous work has acknowledged the harmful impact of many of such unwanted situations and pathologies, ranging TCP chaoticness (Veres & Boda 2000), routing misconfiguration (Paxson 1997) and loops (Hengartner, Moon, Mortier & Diot 2002), packet reordering (Paxson 1999), denial of use and virus attacks (Lad, Zhao, Zhang, Massey & Zhang 2003) among many others.

For these reasons, the Internet research community has understood the need to take measurements of key performance parameters to analyse the causes of such pathologies and their impact in both individual quality of service degradation and global network performance. To this end, a large number of research projects and institutions (see for instance Ripe NCC, NLANR, CAIDA, Surveyor, Abilene websites for further details) are devoted to collecting and providing all types of measurements that can help in the comprehension and characterisation of network performance, with the goal of its improvement, and even optimisation.

However, in some circumstances, it is unfeasible to obtain accurate measurements for various reasons. For instance, when active probing, the overmeasure may cause network overload, hence risking the accuracy of the reading. In other cases, the cost of deploying an adequate measurement infrastructure might result too expensive. Other situations include when recently past measurements are required, and therefore impossible to repeat.

In all cases, when measurements are not easy or possible to obtain, it is important to have the chance to reproduce them in a way that preserves their statistical properties and features as much as possible.

Essentially, the generation of real-like measurements can help in the creation of controlled simulation scenarios and testbeds to explore the emerging Internet. This includes not only the development of testbeds to test new protocols and applications, but also the isolation of particular aspects of an environment in order to analyse its behaviour under certain harmful situations without risking real infrastructures. Other challenges, potential possibilities and applications of network simulation can be reviewed in the literature (Floyd & Paxson 2001).

Among all measurable quantities, packet delay constitute a major network metric under study for various

reasons. Firstly, they are highly representative of network performance. Secondly, their dynamics impact the behaviour of congestion control of TCP, the dominant transport protocol in Internet communications. Thirdly, the popularity and number of delay-sensitive applications, such as online gaming and multimedia streaming, is increasing rapidly. Finally, estimates of network delays are easy to capture, for instance using the *ping* or *traceroute* facilities.

This work is concerned with the generation of synthetic, but statistically equivalent to properly measured data, delays using a two-stage procedure named SS-SVM. Given a seed of real measurements, the proposed method proceeds first capturing the statistical dynamics of the sample, and then randomly generating new samples from such inferred statistical features.

Section 2 reviews previous meaningful results in network-delay analysis and characterisation. Section 3 introduces the mathematical framework of the SS-SVM algorithm and describes its main blocks in some detail. Section 4 outlines the experiments carried out and explores the results and findings. Finally, section 5 summarises the conclusions and possible further lines of investigation.

## 2. PREVIOUS WORK

Internet delays have traditionally been subject of study in network performance characterisation and modelling (Bolot 1993, Paxson 1994, Papagiannaki, Moon, Fraleigh, Thiran & Diot 2003, Mochalski, Micheel & Donnelly 2002). Network delay and loss measurements are often referred to as low-level performance-related quantities, since they can easily and accurately capture operation and performance status of a monitored network or a part of it. Indeed, heavily loaded networks make packets experience increased variation in delay and increased loss rates, whereas uncongested networks shows bounded delays and low loss rates.

But also, the analysis and characterisation of packet delay and loss have proven utility in other network-performance related applications. For instance: the detection of network pathologies and routing exchange protocol misbehaving (Paxson 1997), the identification of undesired packet dynamics such as packet reordering and replication (Paxson 1999), and the estimation of router buffer usage and available link throughput (Prasad, Murray & Claffy 2003), amongst others.

For these reasons altogether, the analysis and characterisation of network delays have traditionally been a major subject of study among the network research community. However, some parts of its modelling remain yet challenging. Previous work has reported on the difficulties in finding single probability distributions to accurately match delay histograms (Mukherjee 1994, Bovy, Mertodimedjo, Hooghiemstra, Uijterwaal & Van Mieghem 2002, Papagiannaki et al. 2003, Choi, Moon, Zhang, Papagiannaki & Diot 2004, Lelarge, Liu & Xia 2004). Essentially, this is due to the fact that packet-

switched network delays exhibit extreme variability, with outlying data points occurring with non-negligible frequencies, leading to moment estimates often enormous and uninformative (Willinger, Paxson & Taquq 1997).

Such difficulties identified in empirical data suggests the use of alternative techniques than parametric modelling to characterise network delay samples. Accordingly, non-linear regression models provide a good means to construct a non-linear curve that best fit a pair sample according to the minimisation of a cost function. In light of this, Support Vector Regression provides a batch procedure to construct a nonlinear function that matches a given set of data, easily and optimally. A large number of SVR software packages are publically available, for instance, SVMlight, cSVM, GiniSVM, LS-SVMlab, SVMtorch, WinSVM, etc.

Once the nonlinear function is obtained, random samples of it can be effectively generated using a wide range of sampling methods (Gelfand & Smith 1990, Gilks & Wild 1992, Metropolis, Rosenbluth, Rosenbluth & Teller 1953, Hastings 1970). However, the Slice-Sampling method (Neal 2003) outstands for its simplicity and efficiency.

The forthcoming section briefly reviews the mathematical foundations of the two methods.

## 3. MODEL INSIDES

The model proposed combines two fairly recent and powerful statistical techniques in the cutting edge of decision theory and simulation methods: Support Vector Machines and Slice Sampling. It is aimed to combine the benefits of both in a two-phase procedure, as shown in figure 1.

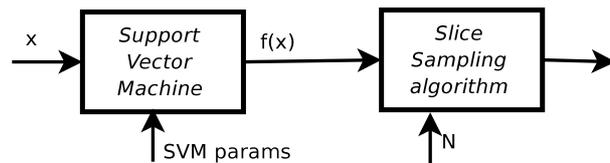


Figure 1: SS-SVM block

In the SVR stage, a  $K$ -bin histogram that corresponds to a  $M$ -sized set of either real data measurements, name  $\mathbf{x}$ , inputs a Support Vector Machine in nonlinear regression operation. The SVM block generates a continuous function  $f(x)$ , which is assumed to be proportional to the probability density function of the delays, i.e.  $f(x) = Cp(x)$ . Such function is built on the basis of a combination of Gaussian kernel operations of the data with the support vectors obtained, according to a set of SVM parameters a priori defined. Such function inputs to the second stage, at which the Slice Sampling algorithm shall produce a synthetic and statistically equivalent  $N$ -sized sample.

The following sections briefly revises the very foundations of Support Vector Regression and Slice Sampling. The authors refer to (Smola & Schölkopf 2004)

for further details on Support Vector Regression and to (Neal 2003) for a more extensive analysis on the Slice Sampling technique.

### Support Vector Regression overview

Since the first proposal of Support Vector Machines (often referred as SVMs) in 1992 (Boser, Guyon & Vapnik 1992), extensive research has been conducted to explain their features, show their benefits and use them in a large range of applications in classification, regression, clustering, mainly. Support Vector methods are nonlinear generalisation algorithms, firmly grounded in the framework of statistical learning theory, or VC theory, which enable them to generalise well to unseen data.

Essentially,  $\varepsilon$ -SV Regression (Smola & Schölkopf 2004) aims to find a continuous-valued function,  $f(x)$ , that best fits a set of  $N$  training data pairs, i.e.  $\{(x_i, y_i)\}_{i=1}^N$  with  $(x_i, y_i) \in \mathbf{R}$ , in a way that shares many of the benefits of Support Vector Machines. Such function  $f(x)$  takes the expression:

$$f(x) = w^T \phi(x) + b \quad (1)$$

which achieves non-linearity via the  $\phi$ -mapping of the input vector into the feature space.

The choice of parameters  $w$  and  $b$  employs the formulation of a constrained optimisation problem with two characteristics: (1)  $f(x)$  must be as flat as possible, that is, the norm  $\|w\|^2$  must be minimum; and (2) the training data points must be as close as possible to the subsequent predicted value. Large deviations will be penalised according the so-called  $\varepsilon$ -insensitive loss function  $|\xi|_\varepsilon$  defined as follows:

$$|\xi|_\varepsilon = \begin{cases} 0 & \text{if } |\xi| < \varepsilon \\ |\xi| - \varepsilon & \text{otherwise} \end{cases} \quad (2)$$

where  $\xi = y - f(x)$ . As shown, the  $\varepsilon$ -insensitive cost function linearly penalises deviations larger than  $\varepsilon$ , yet permitting errors below it with no cost.

Vapnik's formulation (Vapnik 1995) defines a constant  $C > 0$  which determines the trade-off between the flatness of the objective function  $f(x)$  and the total deviation larger than  $\varepsilon$  permitted, which leads to the following constrained optimisation problem:

$$\begin{aligned} \text{minimise:} & \quad \frac{1}{2} \|w\|^2 + C \sum_{i=1}^N (\xi_i + \xi_i^*) & (3) \\ \text{subject to:} & \quad \begin{cases} y_i - w^T \phi(x_i) - b \leq \varepsilon + \xi_i \\ w^T \phi(x_i) + b - y_i \leq \varepsilon + \xi_i^* \\ \xi_i, \xi_i^* \geq 0 \end{cases} & (4) \end{aligned}$$

This problem is often solved in its dual form using Quadratic Programming. Faster-than-QP algorithms have also been proposed (Musicant & Feinberg 2004, Flake & Lawrence 2002, Navia-Vázquez, Pérez-Cruz, Artés-Rodríguez & Figueiras-Vidal 2001).

### Slice sampling review

Traditional Markov Chain Monte Carlo techniques such as the Gibbs sampler (Gelfand & Smith 1990), the Acceptation Rejection Sampling (Gilks & Wild 1992) and the Metropolis-Hastings algorithms (Metropolis et al. 1953, Hastings 1970) have been extensively used to sample from both univariate and multivariate probability distributions. However, these methods require extra knowledge on how to sample from the conditional distributions either directly or by finding an appropriate 'auxiliar' distribution, which is often hard to find or leads to unefficient approaches.

Slice sampling (Neal 2003) arises as an easy and more efficient alternative to sample from any univariate or multivariate distribution, just by alternating uniform sampling in the vertical direction with uniform sampling from the horizontal slice defined by the current vertical position.

Let us assume we are given a well defined continuous function  $y = f(x)$  from which, we want to take random samples. Essentially, slice sampling proceeds in a three step procedure, say:

- (1) Sample uniformly from the range  $[x_{min}, x_{max}]$  giving  $x_0$ . That is  $x_0 \sim U(x_{min}, x_{max})$  (figure 2, top-left).
- (2) Define the vertical range  $[0, f(x_0)]$ , from which we take a uniformly random value  $y_0$  (figure 2, top-right).
- (3) Define the horizontal slice containing the  $x$  values for which  $f(x)$  lay underneath  $y_0$ . That is  $S_0 = \{x : f(x) > y_0\}$  (figure 2, bottom-left).
- (4) Sample uniformly from the slice  $S_i$ ,  $x_i \in S_i$  (figure 2, bottom-right).
- (5) Go to step 2 and repeat.

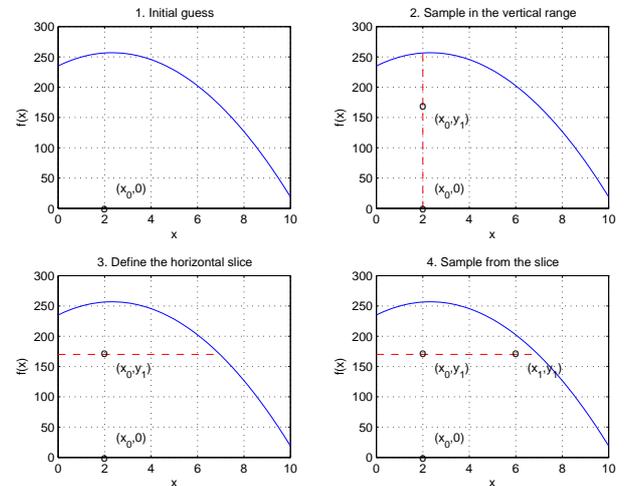


Figure 2: Slice sampling example

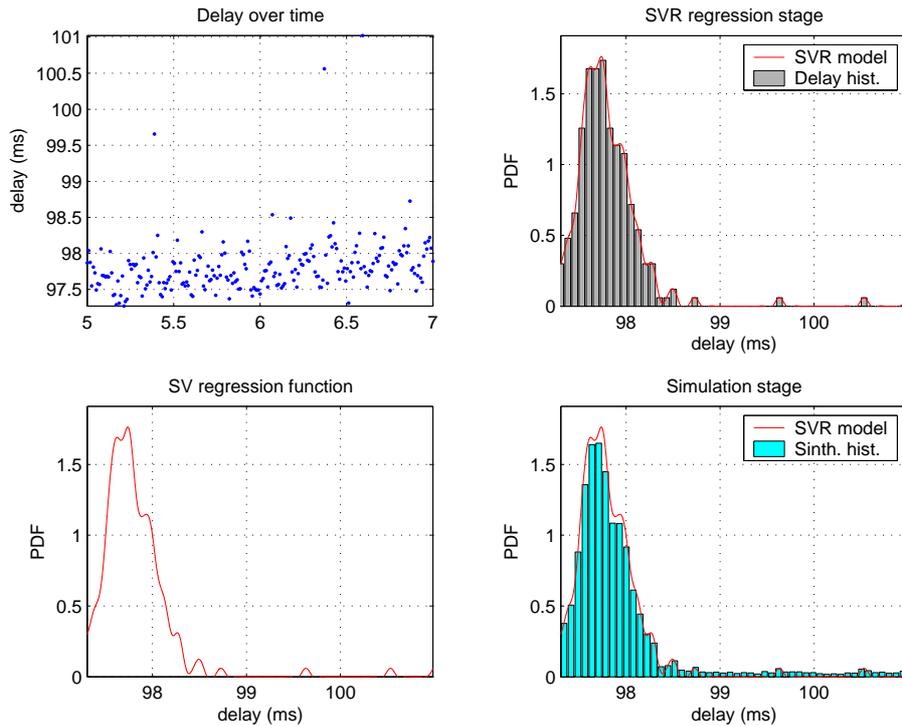


Figure 3: SV regression and synthetic generation for the link Amsterdam - New Zealand. Delay over time of day (top-left); delay histogram and SV regression function (top-right); SV regression function (bottom-left); SVR function and synthetic delay histogram (bottom-right)

After a relatively large number of iterations, the procedure outputs a set of random samples from which, the last say  $N$  samples can be considered as being drawn from the distribution  $f(x)$ .

#### 4. SIMULATION EXPERIMENTS AND RESULTS

The measurements utilised for the experiments have been gently donated by the Test Traffic Measurements project of the RIPE NCC institution. Among its many current activities, the RIPE NCC institution continuously collects around 3000 one-way delay GPS-synchronised measurements per day from more than 40 monitored points spread over Europe, the United States, Australia and New Zealand. This results in an average of 2 measurements per minute per route, thereby giving a coarse-time estimation of the load status of the network by means of the delay measured.

Two numerical examples have been carried out, both with the same monitoring site. The former (experiment A) explores the characteristics of an inter-continental link and the latter investigates the delay properties for a trans-European link (experiment B).

It is worth pointing out that delay properties do not persist over the time of the day. In other words, we should expect to see different histograms at different times of day. This is rather obvious, since the network at peak time in the morning is expected to be more loaded, than the same routers and links at midnight. Thereby, we can not consider statistical stationarity for large periods of time,

since the network status is continuously changing. In the forthcoming experiments, we have considered local stationarity, assuming that the statistical properties do not change for periods of time less than two hours.

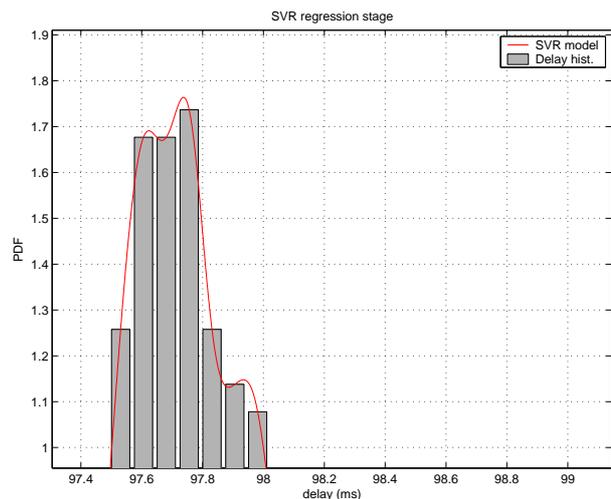


Figure 4: SV regression function and delay histogram zoomed in.

#### Experiment A: The Inter-continental link

In this experiment, we have considered an input of 227 real measurements collected between 5 – 7am the 1<sup>st</sup> of January 2002. A plot of these measurements over

time can be seen in fig. 3 top-left. From these measurements, we have computer and plotted their SV regression function (fig. 3 top-right and bottom-left) and generated 20000 synthetic delays using the SS-SVM approach. A plot of the histogram of the synthetic delays is given in fig. 3 bottom-right. As shown, the histogram of the synthetic measurements is quite close to the SV regression function, and therefore similar to the original data set.

A zoom of fig. 3 top-right is given in fig. 4. This shows the non-linearity nature of the SV regression function and also its capability of matching the centre points of all the histogram bins.

The next experiment shall briefly discuss on the flatness issues of Support Vector Regression and a methodology to quantify the similarity between the real and synthetic data sets.

### Experiment B: The European link

In this experiment, we have considered an input of 2715 samples that corresponds to a 24-hour set of measurements period of the Amsterdam-London link. With this information, and for various SVM parameters, we have generated 9000 synthetic measurements using the SS-SVM method, more than three times the original set.

For testing the results, we have made use of the Kolmogorov-Smirnov goodness-of-fit test for assessing on whether the real measured sample is statistically equivalent to the synthetic data set. The Kolmogorov-Smirnov method is a statistical test used to provide a means to decide whether a data sample may have come from a particular probability distribution or not. It is also very much used to determine whether two independent samples come from the same statistical distribution or not. It proceeds computing the so-called KS statistic and looking up in a table. A good summary on the K-S test can be found on (González, Sahni & Franta 1977).

Table 1 shows the numerical results obtained in this experiment. The column on the left gives values of the flatness measures. According to equation 3, such flatness quantity has been computed as follows:

$$\text{flatness} = \frac{\frac{1}{2} \|w\|^2}{\sum_{i=1}^N (\xi_i + \xi_i^*)} \quad (5)$$

Essentially, the smaller value, the smoother the SV regression function is. The second column in the table shows the previously described Kolmogorov-Smirnov statistic. The third column should be read as either accept or reject the null hypothesis. In other words,  $H_\alpha = 0$  implies the two samples are statistically equivalent and  $H_\alpha = 1$ , they are drawn from different probability distributions, according to the Kolmogorov-Smirnov test for different levels of significance  $\alpha$ .

Table 1: K-S statistic and acceptance or rejection results of synthetic samples for different flatness values in the SVR stage

Flatness	K-S stat	$H_{0.1}$	$H_{0.05}$	$H_{0.01}$
0.3065	$9.18 \times 10^{-51}$	1	1	1
0.0476	$1.35 \times 10^{-5}$	1	1	1
0.0491	0.0502	1	0	0
0.0103	0.4415	0	0	0
0.0091	0.3488	0	0	0
0.0060	0.2918	0	0	0
0.0367	0.4279	0	0	0
0.0441	0.0169	1	1	0
0.0015	0.0499	1	1	0
$2.08 \times 10^{-9}$	0.0045	1	1	1

As shown, when the flatness value is too small or too large, the K-S test rejects the synthetic sample. Figure 5 shows the different regression functions achieved for the second, forth and ninth experiments.

Essentially, the effect shown is that, when the SV regression function is too smooth (fig. 5 top), then it cannot actually capture the dynamics of the sample histogram, giving a too generic function which does not quite represent the original sample. Hence, in this case, the synthetic sample will be far different from the original one. On the contrary, when the flatness value is too large (fig. 5 bottom), the effect shown is that the regression function is too specific of the sample, therefore it does not generalise well. This impacts in the slice sampling algorithm as it cannot reproduce an accurate synthetic data set of such an specific function. In between the two cases (fig. 5 middle), when the regression function is not too uniform nor too coarse, it captures the histogram curve acceptably and generalises to produce a good representative synthetic sample.

## 5. SUMMARY AND DISCUSSION

This work has introduced a novel technique to generate synthetic Internet delays with the same statistical properties and characteristics than measurements collected under real network activity. Its methodology is based on the non-linear regression application of Support Vector Machines, combined with an efficient method of random sampling known as Slice Sampling. This two-stage procedure has been visually and quantitatively tested and validated with delay measurements collected and donated by the RIPE NCC institution. However, it is worth remarking that this two-stage procedure is not restricted to the generation of synthetic network delays. Instead, this method can be applied in generating synthetic data sets in a wide range of networking simulation scenarios, with various applications. For instance, this includes the simulation of:

- Packet losses for testing application resilience in mobile computing scenarios. Under wireless environments, lost packets are more frequent due to

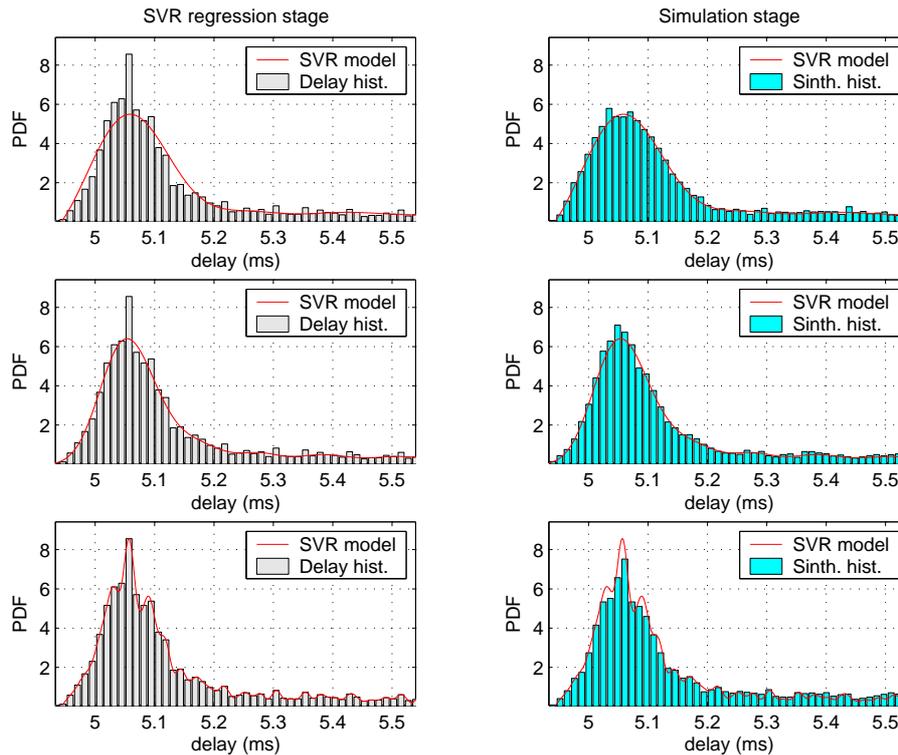


Figure 5: SV regression and synthetic generation for various grades of flatness: flatness= 0.0476 (top), flatness=0.0103 (middle) and flatness=0.0015 (bottom)

channel fading than in wired networks. This can result in the performance and tolerance analysis of certain loss-sensitive applications under such scenarios.

- Flow length and duration for assessing on the grades of self-similarity and long-range dependence obtained in global aggregated traffic volumes. This is useful in forecasting network traffic with further applications in capacity dimensioning and network planning.
- Routing update messages for characterising the robustness and recovery of both inter-domain and intra-domain routing protocols in the events of link-failure, malicious attacks or simply gradual overload.
- Packet arrival rates for determining optimal router buffer sizing at different grades of network use.
- Of course, network delays for examining the feasibility of using real-time (delay-sensitive) applications under different grades of congestion. This includes assessing whether applications such as Internet telephony, multimedia streaming and online gaming would run smoothly at different scenarios.

## ACKNOWLEDGEMENTS

The authors would like to acknowledge the provision of data by RIPE NCC. José A. Hernández is funded by a

Loughborough University PhD scholarship.

## REFERENCES

- Bolot, J. C. (1993), 'Characterizing end-to-end packet delay and loss in the Internet', *Journal of the High Speed Networks* 2(3), 305–323.
- Boser, B. E., Guyon, I. M. & Vapnik, V. (1992), A training algorithm for optimal margin classifiers, in 'Proceedings of the Annual Conference on Computational Learning Theory', ACM Press, Pittsburg, PA, pp. 144–152.
- Bovy, C. J., Mertodimedjo, H. T., Hooghiemstra, G., Uijterwaal, H. & Van Mieghem, P. (2002), Analysis of end-to-end delay measurements in the Internet, in 'Proceedings of Passive and Active Measurements'.
- Choi, B. Y., Moon, S., Zhang, Z. L., Papagiannaki, K. & Diot, C. (2004), Analysis of point-to-point packet delay in an operational network, in 'Proceedings of INFOCOMM 2004'.
- Flake, G. W. & Lawrence, S. (2002), 'Efficient SVM regression training with SMO', *Machine Learning* 46(1-3), 271–290.
- Floyd, S. & Paxson, V. (2001), 'Difficulties in simulating the Internet', *IEEE/ACM Trans on Networking* 9(4), 392–403.
- Gelfand, A. E. & Smith, A. F. M. (1990), 'Sampling-based approaches to calculating marginal densities', *Journal of the American Statistical Association* 85, 398–409.
- Gilks, W. R. & Wild, P. (1992), 'Adaptive rejection sampling for Gibbs sampling', *Applied Statistics* 41, 337–348.
- González, T., Sahni, S. & Franta, W. R. (1977), 'An efficient algorithm for the kolmogorov-smirnov and lilliefors tests',

*ACM Transactions on Mathematical Software* **3**(1), 60–64.

- Hastings, W. K. (1970), 'Monte Carlo sampling methods using Markov chains and their applications', *Biometrika* **57**, 97–109.
- Hengartner, U., Moon, S., Mortier, R. & Diot, C. (2002), Detection and analysis of routing loops in packet traces, in 'Proceedings of IMC'02', Marseille, France.
- Lad, M., Zhao, X., Zhang, B., Massey, D. & Zhang, L. (2003), Analysis of BGP update surge during Slammer worm attack, in 'Proceedings of the International Workshop on Distributed Computing'.
- Lelarge, M., Liu, Z. & Xia, C. H. (2004), Asymptotic tail distribution of end-to-end delay in networks of queues with self-similar cross traffic, in 'Proceedings of INFOCOMM', Atlanta (USA).
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N. & Teller, E. (1953), 'Equation of state calculations by fast computing machines', *Journal of Chemical Physics* **21**, 1087–1092.
- Mochalski, K., Micheel, J. & Donnelly, S. (2002), Packet delay and loss at the Auckland Internet access path, in 'Proceedings of Passive and Active Measurements'.
- Mukherjee, A. (1994), 'On the dynamics and significance of low-frequency components of Internet load', *Internet-working: research and experience* **5**(4), 163–205.
- Musicant, D. R. & Feinberg, A. (2004), 'Active set support vector regression', *IEEE Trans on Neural Networks* **15**(2), 268–275.
- Navia-Vázquez, A., Pérez-Cruz, F., Artés-Rodríguez, A. & Figueiras-Vidal, A. (2001), 'Weighted least squares training of support vector classifiers leading to compact and adaptive schemes', *IEEE Trans on Neural Networks* **12**(5), 1047–1059.
- Neal, R. M. (2003), 'Slice sampling', *Annals of Statistics* **31**(3), 705–767.
- Papagiannaki, K., Moon, S., Fraleigh, C., Thiran, P. & Diot, C. (2003), 'Measurement and analysis of single-hop delay on an IP backbone network', *IEEE Journal on Selected Areas in Communications. Special Issue on Internet and WWW Measurement, Mapping and Modeling* **21**(6).
- Paxson, V. (1994), 'Empirically-derived analytic models of wide-area TCP connections', *IEEE/ACM Trans on Networking* **2**(4), 316–336.
- Paxson, V. (1997), 'End-to-end routing behavior in the Internet', *IEEE/ACM Trans on Networking* **5**(5), 601–615.
- Paxson, V. (1999), 'End-to-end Internet packet dynamics', *IEEE/ACM Trans on Networking* **7**(3), 277–292.
- Prasad, R., Murray, M. & Claffy, K. C. (2003), 'Bandwidth estimation: metrics, measurement techniques, and tools', *IEEE Network* pp. 27–35.
- Smola, A. & Schölkopf, B. (2004), 'A tutorial on support vector regression', *Statistics and Computing* (14), 199–224.
- Vapnik, V. (1995), *The nature of statistical learning theory*, Springer, New York.
- Veres, A. & Boda, M. (2000), The chaotic nature of TCP congestion control, in 'Proceedings of IEEE INFOCOM', Tel Aviv.
- Willinger, W., Paxson, V. & Taqqu, M. S. (1997), 'Self-similarity and heavy-tails: structural modeling of network traffic', *Annals of Statistics* **25**(5), 1805–1869.

## AUTHOR BIOGRAPHIES

### JOSÉ A. HERNÁNDEZ GUTIÉRREZ

received his five-year degree in Telecommunications Engineering from the Carlos III University at Madrid (Spain) in November 2002. After this, he joined Loughborough University's Department of Computer Science to conduct research on network measurement-based bayesian modelling, as part of his PhD. His research interests include the areas at which bayesian statistics, machine learning techniques, nature-based computation and network performance modelling overlap. His email address is J.A.Hernandez@lboro.ac.uk



**DR IAIN W. PHILLIPS** is a lecturer in the Department of Computer Science at Loughborough University in the UK. His research interests are the investigation of performance issues in networking, where performance can refer from the underlying hardware to the human user/operator. He has worked on project funded by Industry, UK Research Councils and Europe, and with many industrial collaborators, including BT and NTL. He is a member of the ACM and the British Computer Society. His email address is I.W.Phillips@lboro.ac.uk



**DR JAVIER M. MOGUERZA** received his Ph. D. in mathematical engineering from Universidad Carlos III de Madrid in 2000. After this, he became associate professor at Universidad Rey Juan Carlos in Madrid (Spain). He has collaborated in both Spanish and European funded projects regarding bayesian statistics and electronic democracy. His research interests are focused on methods for mathematical optimisation, machine learning and pattern recognition using Support Vector Machines. His email address is Javier.Moguerza@urjc.es

