

Measuring the Relationships between Internet Geography and RTT

Raul Landa, Richard G. Clegg, João Taveira Araújo, Eleni Mykoniati, David Griffin, Miguel Rio
 Department of Electronic and Electrical Engineering, University College London
 Email: {rlanda, jaraujo, rclegg, emykonias, dgriffin, mrio}@ee.ucl.ac.uk

Abstract—When designing distributed systems and Internet protocols, designers can benefit from statistical models of the Internet that can be used to estimate their performance. However, it is frequently impossible for these models to include every property of interest. In these cases, model builders have to select a reduced subset of network properties, and the rest will have to be estimated from those available.

In this paper we present a technique for the analysis of Internet round trip times (RTT) and its relationship with other geographic and network properties. This technique is applied on a novel dataset comprising ~ 19 million RTT measurements derived from ~ 200 million RTT samples between ~ 54 thousand DNS servers.

Our main contribution is an information-theoretical analysis that allows us to determine the amount of information that a given subset of geographic or network variables (such as RTT or great circle distance between geolocated hosts) gives about other variables of interest. We then provide bounds on the error that can be expected when using statistical estimators for the variables of interest based on subsets of other variables.

I. INTRODUCTION

One of the difficulties in the design and modelling of large-scale distributed systems and Internet protocols is that their performance can be markedly affected by the *real-life* characteristics of their underlying networks. Unfortunately, the sheer complexity of the Internet precludes a full modelling or simulation of these characteristics. It is however possible to use network measurements to develop scalable models based on specific statistical properties, which can then be used to construct statistical estimators for those properties not explicitly included in the model. However, this presents researchers with a new challenge: selecting the subset of variables to model in order to get the highest quality information.

In this paper we approach this question using information theory, and present quantitative measures of the degree to which observation of a given network property can be used to estimate other related network properties. To achieve this we collected one of the most comprehensive sets of Internet delay measurements available today, comprising more than 200 million individual RTT samples taken between ~ 54 thousand measurement points; this dataset is available to researchers on request. We treat each measurement in this dataset as an observation of a multidimensional random variable \mathbf{X} that also includes geographic (e.g. country pair) and network (e.g. ASN pair) metadata relevant for the measurement points involved. This allows us to build a *multi-dimensional histogram*, i.e. a histogram where each bin is associated not with a single-variable interval, but with the cartesian product of intervals for

all variables of interest.

Our main contribution is an information-theoretical analysis that allows us to determine the amount of information that knowledge of a given subset of variables \mathbf{Y} gives about any given variable of interest X . Then, given that information, we provide bounds on the error that can be expected from an estimator that uses observations of \mathbf{Y} to estimate the value of X . We discuss which variables provide the greatest usefulness to estimate the RTT, geographic variables (e.g. the countries on which communication endpoints are located), and network variables (e.g. the AS to which each one of them belongs).

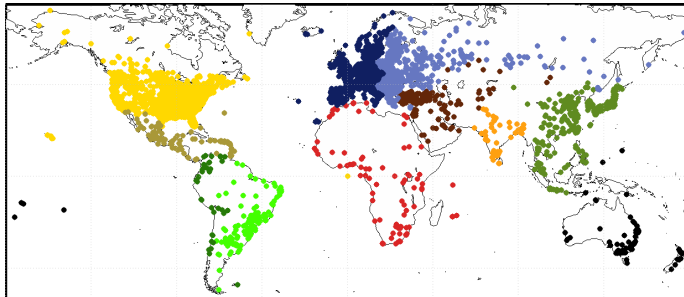
Our work goes beyond previous contributions in the topic by applying information-theoretic techniques on a novel dataset with enough diversity to account for the heterogeneity of Internet infrastructure and geography. Other analyses that have addressed the geographic properties of RTT do so in specific contexts, such as the design of IP address geolocation techniques [9], [11], [23], [20] or the analysis of the relationships between IP prefixes and RTT [7]. There has also been interest in the study of the relationship between geography and IP prefixes [10], and between network locations and RTT [13]. Although all these works share the common premise that by observing certain variables it is possible to gain insights on the behaviour of related variables, a more general analysis of how helpful certain variables are to estimate the values of other variables is still lacking in the literature. Such analysis would be very useful to make informed tradeoffs in protocol design, and to construct statistical estimators with improved cost-benefit. This paper aims to be a step in this direction.

The structure of the paper is as follows. We commence in §III by presenting our modelling and measurement considerations. We then present our analysis technique and its underlying theory in §IV. In §V we present other research contributions that relate to ours, and in §VI, our conclusions.

II. MEASUREMENTS

A comprehensive study of the geography of Internet RTTs requires the measurement of round trips between a large number of measurement points in as many distinct geographical locales as possible. Due to their ubiquitous nature, DNS servers are ideal for this purpose; this led us to select TurboKing [17] as our main measurement technique. Data collection took place from May 2011 to February 2012, and comprises ~ 200 million individual RTT samples between ~ 54 thousand recursive, non-forwarding DNS servers. To allow the system to filter

out any episodic RTT effects, each RTT measurement was calculated as the median value of 10 individual samples spaced 10 seconds apart. Hence, the RTT samples collected yielded ~ 20 million RTT measurements, which were reduced to ~ 19 million after additional processing required to remove spurious results arising from DNS server configuration. Before database storage, each measurement was tagged with the metadata information presented in Table II.



Subcontinental Zone	Servers
Africa	519
Central Asia and the Middle East	1,490
Asia Pacific and China	7,730
Indian Subcontinent	449
North America (North)	21,276
North America (South) and the Caribbean	526
Oceania	1,116
South America (West)	270
South America (East)	1,333
Eastern Europe (including Russia)	6,798
Western Europe	12,953

TABLE I: Geographic distribution of measurement servers

Since we were interested in the relationships between RTT and both geographic and network variables, we made extensive use of IP geolocation databases. Specifically, we used both MaxMind GeoLite and [1] and Neustar IP Geolocation [2] alongside an spatial index [12] derived from the GeoNames city names database [4]. The two latter sources resulted in 122,952 place bindings representing cities with 1,000 inhabitants or more, which we then used along the geolocated latitude and longitude of each IP address to obtain the country associated with it. Since we also use AS information to aggregate RTT measurements, we required a mechanism to resolve IP addresses to their respective AS. We used both MaxMind GeoLite and services offered by Team Cymru [5].

As indicated by our geolocation and AS look-up sources, DNS servers used as measurement endpoints for our dataset were present in 5,455 autonomous systems over 3,384 cities and 189 countries. According to [3], 99.6% of global Internet users are located in a country with at least one measurement DNS server. To verify the IP coverage of our measurement set, we compared it with a daily routing table snapshot obtained from the RouteViews server in WIDE [6]. Of the 216,344 prefixes received at that point, 20,881 included at least one measurement DNS server. Overall, 476 million addresses ($\sim 32\%$ of 1.017 billion routable IP addresses) belong to network prefix that contains at least one measurement server.

Categorical Variables		
Variable	Interpretation	Bins used
X_p	Pairs of /8 prefixes	14,189
X_a	Pairs of Top AS numbers	215,392
X_z	Pairs of subcontinental zones	66
X_c	Pairs of countries	2,648
Quantitative Variables		
X_l	Common prefix length	32
X_d	Great circle distance	300
X_t	Round trip time (RTT)	300

TABLE II: Definition of variables used

III. MODELLING AND PRELIMINARIES

In the following we address the relationships between RTT and other related network and geographic variables shown in Table II. Since RTT measurements are always performed between two distinct hosts, all these variables refer to properties of distinct *pairs* of hosts. We will distinguish two kinds of variables: *categorical* and *quantitative*. The first set includes variables in which each RTT measurement is associated with any one of a set of equivalent labels; the second one includes variables that represent measurements on an underlying numerical scale. Some categorical variables in Table II are self-explanatory, e.g. X_p refers to a pair of octets, one for each host, and X_c refers to the pair countries in which each one of the hosts associated with an RTT measurement are located. We now provide further explanation for the others.

In order to better elucidate the large-scale effects that Internet routing has on RTT we divided our measurement endpoints into 11 *subcontinental zones* consisting of geographically adjacent countries, as shown above Table I. To map an IP address to its subcontinental zone, we first map it to a country, which is then mapped to a subcontinental zone. This leads to a set of 66 distinct values for X_z describing distinct pairs of subcontinental zones. Regarding X_a , we distinguished an AS as a *top AS* if the dataset included at least $\sim 10,000$ measurements associated with that AS. In total, we found 658 *top AS* numbers. If a measurement endpoint did not belong to a *top AS*, we labelled it as belonging to a generic *non-top AS*; otherwise, we labelled it with the AS number to which its IP address belonged. This ensured that there was enough probability mass associated with *top AS* pairs to make an accurate estimation of their relevant distributions. Regarding quantitative variables, the common prefix length X_l between the host addresses is naturally discrete; in our dataset it ranges from 0 (the first bit of the two IP addresses is different) to 31 (only the last bit of the two IP addresses was different). To simplify the analysis, we discretised X_d and X_t to the number of bins reported on Table II, yielding a resolution of 67 km and 3.3 ms respectively.

We model our data as a 7-dimensional discrete random variable \mathbf{X} , whose joint probability distribution will be denoted as $\Phi(\mathbf{x})$. Each measurement in our dataset is then is a *realisation* of \mathbf{X} , so that $\mathbf{X} = \{X_p, X_a, X_c, X_z, X_l, X_d, X_t\}$. We construct an empirical probability density for $\Phi(\mathbf{x})$ simply by considering a multidimensional histogram with bins defined by the values and/or value ranges achieved by each component of \mathbf{X} . Since we choose the DNS server endpoints for each one of

our measurements randomly, we expect this multidimensional histogram to converge to the appropriate density $\Phi(\mathbf{x})$. A naive implementation would require the maintenance of a histogram with in the order of 10^{21} bins, leading to a memory requirement in the order of 10^{11} GB. This is clearly unfeasible. Fortunately, only a few of these bins have nonzero entries. Thus, by relying on data structures based on hash maps, it becomes feasible to perform all needed calculations on $\Phi(\mathbf{x})$ in minutes with no sampling or approximations required.

Since we will be interested in several marginal distributions of $\Phi(\mathbf{x})$, it will be advantageous to define a shorthand for the marginal distribution of $\Phi(\mathbf{x})$ with respect to a subset of variables $\mathbf{Y} \subset \mathbf{X}$. As an example, let \mathbf{Y} be the set of geographic variables in the dataset, so that $\mathbf{Y} = \{X_z, X_c, X_d\}$. Then, we have that $\mathbf{y} = \{x_z, x_c, x_d\}$ and that $\Phi(\mathbf{y})$, the marginal distribution containing only variables in \mathbf{y} , can be defined as

$$\Phi(\mathbf{y}) = \Phi(x_z, x_c, x_d) = \sum_{\substack{x_p, x_a, \\ x_l, x_t}} \Phi(x_z, x_c, x_d, x_p, x_a, x_l, x_t).$$

Of course, for other subsets \mathbf{Y} of \mathbf{X} , $\Phi(\mathbf{y})$ can be defined in an analogous manner by summing over the remaining variables. We continue by presenting the mathematical tools that we will use to explore the relationships between variables in Table II.

A. Conditional Entropy

In §IV we will be interested in estimating how useful are variables in Table II in estimating the values of other variables. To this end, we rely on information-theoretical arguments based on the *conditional entropy* between these variables. In the following, we use X_\bullet to denote any variable selected from \mathbf{X} , and $\mathbf{Y} \in \{\mathbf{X} \setminus X_\bullet\}$ to denote a subset of observed variables.

The conditional entropy of X_\bullet given \mathbf{Y} , denoted as $\mathcal{H}(X_\bullet|\mathbf{Y})$, quantifies the remaining entropy (uncertainty) in X_\bullet given that variables in \mathbf{Y} are known. Equivalently, $\mathcal{H}(X_\bullet|\mathbf{Y})$ is a measure for the information in X_\bullet which is not present in \mathbf{Y} , and it can be used to bound the performance of approximations of X_\bullet based on \mathbf{Y} . Formally,

$$\mathcal{H}(X_\bullet|\mathbf{Y}) = \sum_{\mathbf{x}_\bullet, \mathbf{y}} \Phi(\mathbf{x}_\bullet, \mathbf{y}) \log \left(\frac{\Phi(\mathbf{y})}{\Phi(\mathbf{x}_\bullet, \mathbf{y})} \right).$$

This can be shown to be equivalent to

$$\mathcal{H}(X_\bullet|\mathbf{Y}) = \mathcal{H}(X_\bullet, \mathbf{Y}) - \mathcal{H}(\mathbf{Y}),$$

where $\mathcal{H}(\mathbf{Y})$ corresponds to the *entropy* of \mathbf{Y} , so that

$$\mathcal{H}(\mathbf{Y}) = - \sum_{\mathbf{y}} \Phi(\mathbf{y}) \log (\Phi(\mathbf{y})).$$

For *quantitative* variables, the conditional entropy can be used to bound the mean squared error (MSE) of estimating the value of X_\bullet given the values of variables in \mathbf{Y} . Let the MSE of an estimator $\hat{X}_\bullet(\mathbf{Y})$ be defined as

$$\text{MSE} = \mathbb{E} \left\{ (\hat{X}_\bullet(\mathbf{Y}) - X_\bullet)^2 \right\},$$

where the statistical average is taken over the joint distribution of X_\bullet and \mathbf{Y} . Then, we have that the \hat{X}_\bullet that minimises the MSE is $\hat{X}_\bullet = \mathbb{E}\{X_\bullet|\mathbf{Y}\}$. In this case, the MSE becomes

the *minimum* MSE (MMSE), and it corresponds to the MSE provided by an optimal estimation of X_\bullet given knowledge of the values of variables in \mathbf{Y} . It can be shown [18] that a lower bound for the MMSE can be provided by $\mathcal{H}(X_\bullet|\mathbf{Y})$ so that

$$\text{MMSE} \geq \frac{1}{2\pi e} \exp(2\mathcal{H}(X_\bullet|\mathbf{Y})).$$

Hence, by considering the conditional entropies of variables in Table II with respect to subsets of related variables, it becomes possible to estimate the minimum reconstruction error that an estimator could provide. This will allow us to explore which variables give the better estimation results, hence informing the design decisions for RTT and IP geolocation estimators. Clearly, lower values for $\mathcal{H}(X_\bullet|\mathbf{Y})$ will lead to lower best-case values for the MMSE. As is standard when making error comparisons with the underlying variable, the *root-mean square error* $\text{RMSE} = \sqrt{\text{MSE}}$ is used in this paper.

For *categorical* variables the problem of using known values in \mathbf{Y} to guess the corresponding value of X_\bullet is not a problem of estimation, but of *classification*. In this case, we can use Fano's inequality to relate the probability of error in determining X_\bullet given \mathbf{Y} to the conditional entropy $\mathcal{H}(X_\bullet|\mathbf{Y})$. In particular, if we define an estimator $\hat{X}_\bullet = g(\mathbf{Y})$ for an arbitrary function g , we are interested in the probability of error $P_e(X_\bullet) = \Pr\{X_\bullet \neq \hat{X}_\bullet\}$ that the classification is incorrect. Then, it can be shown [8] that

$$P_e(X_\bullet) \geq \frac{\mathcal{H}(X_\bullet|\mathbf{Y}) - 1}{\log N},$$

where N is the total number of values that the categorical variable X_\bullet can take. As before, lower values for the conditional entropy will lead to reduced lower bounds for P_e .

B. Correlation between Network and Geographic Variables

Before proceeding to quantify the amount of information that various sets of variables give about one another, it is useful to provide some intuition on the nature of the correlations present in our dataset. To this end, in Fig. 1 we present a subset of the measurements taken between DNS servers in four distinct subcontinental zones: South America (East), denoted as **SAE**; North America (North), denoted as **NAN**; Eastern Europe, denoted as **EUE**; and Asia-Pacific and China, denoted as **ASP**. Clearly, when both measurement DNS servers are located in the same zone, RTT and geolocation distances tend to be low and exhibit high positive correlation. However, when different zones are concerned, correlations tend to be lower and can range from positive (e.g. **EUE/SAE**, **ASP/NAN**) to negative (e.g. **EUE/ASP**). Some pairs of zones exhibit very low degrees of linear correlation (e.g. **ASP/SAE**). As expected, when the measurement DNS servers are located in different subcontinental zones, both the RTT and geolocation distance between them tends to be exhibit higher mean and variance. As shown in a related paper [16], these differences arise from Internet routing at the subcontinental zone scale. For the present paper, the most important property exemplified by Fig. 1 is that geographic and network properties exhibit rich correlations, and that a given set of network/geographic variables can hence provide information on others.




Fig. 1: Relationship between great circle distance and RTT. Each panel corresponds to measurements taken between DNS servers located in a given pair of subcontinental zones, as indicated by the panel headings.

IV. CONDITIONAL ENTROPY AND ESTIMATION ERROR

There have been numerous proposals to use certain geographic and network variables to accurately estimate the values of others [9], [11], [23], [20], [14], [21]. As explained in §III-A, we contribute to the research in this area by using conditional entropy to quantify the degree to which various network and geographic variables give us information regarding one another. This can be useful to bound estimation errors and has direct applicability for many problems that have been addressed by the research community.

The main results in this section are presented in Fig. 2, which shows the remaining uncertainty after estimators are constructed that take as input the variables shown. Each of the subgraphs shown (Figs. 2a to 2e) depicts the conditional entropy for a given variable as different sets of additional variables are considered. It is important to note that, due to synergy and redundancy effects between variables, it is impossible to simply associate a given amount of information with knowledge of a given variable. Instead, different subsets of variables must be considered, and their conditional entropies analysed and compared individually.

We proceed as follows. First, we select a variable of interest X_{\bullet} , and consider its conditional entropy $\mathcal{H}(X_{\bullet}|\mathbf{Y})$ for a number of sets \mathbf{Y} selected from the set of all subsets of variables in Table II that do not contain X_{\bullet} . Then, we consider the implications that these values of $\mathcal{H}(X_{\bullet}|\mathbf{Y})$ have for the design of an estimator \hat{X}_{\bullet} .

A. Estimating RTT

First, we consider the estimation of RTT from the other variables in Table II. The values of $\mathcal{H}(X_t|\mathbf{Y})$ for various sets \mathbf{Y} can be found in Fig. 2a. First, we note that the single indicator that gives the most information regarding the RTT X_t is X_c , the pair of countries in which the communicating hosts are. Knowledge of this reduces the uncertainty in RTT from its maximum value ~ 7 bits to ~ 5.3 bits. In fact, X_c gives more information about X_t than X_p , X_z and X_l taken together. Since X_z is a function of X_c , the conditional entropy $\mathcal{H}(X_z|X_c)$ is zero; considering the subcontinental zones in which the endpoints of an Internet flow are located brings no new information once their countries have been considered. By considering the geolocation distance between the two endpoints in addition to their countries, uncertainty can be decreased from ~ 5.3 bits to ~ 4.7 bits; a similar improvement can be achieved by considering the top AS pair X_a to which the communication endpoints belong, along with the subcontinental zone X_z in which the endpoints are located ($\mathcal{H}(X_t|X_a, X_z) \approx 4.8$ bits). This points to X_a being a better estimator for X_t than geolocation distance X_d , and suggests that considering them simultaneously could lead to better performance. That is indeed the case, with $\mathcal{H}(X_t|X_a, X_d) \approx 3.99$ bits being lower than many three-variable combinations. If in addition we consider X_c , the pair of countries in which the hosts are located, the remaining uncertainty decreases to ~ 3.3 bits. Conversely, if we consider not the AS pair but the first octet of the host IP addresses X_p , we have that




Fig. 2: Conditional entropy $\mathcal{H}(X_{\bullet}|Y)$ of some network and geographic variables X_{\bullet} , given subsets of related variables Y

$\mathcal{H}(X_t|X_p, X_c, X_d) \approx 2.7$ bits; this is close to the minimum conditional entropy that can be achieved by considering four variables or less. It is interesting to note that longest common prefix length X_l between the IP addresses of the endpoints yields very limited information to estimate the RTT. This is consistent with the findings of [10], which show that IP prefixes have very limited geographic locality. Therefore, due to IP address space fragmentation, X_l is of little value to estimate X_t .

By casting the synthesis of RTT information from other variables as an estimation problem, it becomes possible to use the observed values of $\mathcal{H}(X_t|Y)$ to give bounds on the MMSE, and thus, on the maximum reconstruction imprecision that a given estimator can provide. If no knowledge of other variables is assumed, the minimum RMSE is simply the variability of $\Phi(\mathbf{x})$, as measured by its standard deviation. In this case, we have that the uncertainty in predicting RTT is ~ 900 ms, which is close to the full range of X_t . By

incorporating X_c , the best-case RMSE can be improved to ~ 162 ms, and by incorporating both X_c and X_d , further reduced to ~ 89 ms. It is possible to achieve higher accuracy on RTT estimation by using X_p , X_c and X_d ; in this case, the lower bound for the RMSE decreases to ~ 12 ms. Hence, an optimal estimator can predict RTTs based only on X_p , X_c and X_d with an RMSE of ~ 12 ms or higher.

B. Estimating Host Geolocation

We now investigate the usefulness of geographic and network variables for host geolocation. We approach this question by using the conditional entropy to bound the errors involved in computing estimates for the geographical properties X_c , X_z and X_d from knowledge of various subsets of the network variables X_p , X_a , X_l , and X_t .

The potentially most accurate predictor of the subcontinental zone pair X_z in which the endpoints are located is given by the first octets of their respective IP addresses X_p , as this

knowledge reduces the uncertainty in X_z from its maximum of ~ 3.9 bits to ~ 1.4 bits. This is not surprising, as it follows from the fact that there is a geographic element to the distribution of IP prefixes. If in addition to X_p one considers the RTT X_t , conditional entropy is reduced to $\sim .93$ bits; this value is close to the minimum conditional entropy that can be achieved using purely network-related data. The lowest value of the conditional entropy achievable with 4 variables or less can almost be reached by considering X_p , X_t and X_d , since $\mathcal{H}(X_z|X_p, X_d, X_t) \approx .35$ bits.

Regarding the two countries X_c where the communication parties are situated, the total uncertainty with no further information is ~ 7.87 bits. Knowledge of the the AS pair X_a provides a residual uncertainty of ~ 3.8 bits, which is very similar to the ~ 3.9 bits provided by knowledge of the subcontinental zone pair X_z . Hence, on their own, X_a and X_z are similarly useful to estimate X_c . However, if we consider the geolocation distance X_d , we find that although $\mathcal{H}(X_c|X_z, X_d) \approx 2.7$ bits, $\mathcal{H}(X_c|X_a, X_d)$ has the lower value of ~ 1.7 bits. Unsurprisingly, the information provided by X_d has more in common with that provided by X_z than with that provided by X_a . Hence, X_a provides more information about X_c than X_z if the geolocation distance X_d is known. If all these three variables are known, however, a further reduction in conditional entropy can be achieved, as $\mathcal{H}(X_c|X_a, X_z, X_d) \approx .98$ bits. This is close to the minimum achievable with 4 variables or less. The minimum conditional entropy using only network data is $\mathcal{H}(X_c|X_a, X_l, X_t) \approx 2.35$ bits.

Since X_z and X_c are categorical variables, we evaluate the precision of their estimators by considering the error probability P_e of incorrectly assigning a measurement to an incorrect category \hat{X}_z or \hat{X}_c . For an estimator \hat{X}_z using only X_p , we have that $P_e \geq .09$. Inclusion of additional variables results in negative estimates for the minimum P_e ; these are uninformative, and simply imply that the use of better bounds than that provided by (III-A) is required. For an estimator \hat{X}_c using only X_a , we have that $P_e \geq .35$; an estimator using both X_a and X_d can achieve a lower best-case bound $P_e \geq .09$. An estimator based only on network information could achieve a best-case bound $P_e \geq .17$.

Regarding the geolocation distance X_d , we find that its potentially strongest predictor is the country pair X_c , which can reduce its uncertainty to around ~ 4.6 bits from its maximum value of ~ 7.5 bits. By also considering the AS pair X_a to which the hosts belong, it is possible to reduce this uncertainty to ~ 3.3 bits, and to ~ 2.2 bits by incorporating both X_a and the RTT X_t . The best estimate on X_d using 4 variables or less can be achieved with X_a , X_c , X_l and X_t , leading to a residual uncertainty of ~ 1.76 bits. The best geolocation distance estimates using only network information are $\mathcal{H}(X_d|X_a, X_t) \approx 3.77$ bits if we use two estimation variables, and $\mathcal{H}(X_d|X_a, X_l, X_t) \approx 3.25$ bits if we use three.

Regarding RMSE, considering X_c can potentially reduce the RMSE from its maximum of $\sim 20,000$ km to $\sim 1,600$ km. The precision of the estimator can be further improved to an RMSE of 437 km by incorporating X_a , and to ~ 145 km by including both X_a and X_t . The best RMSE bound for all sets of 4 variables or less is ~ 93.76 km; the best RMSE bound

for network-only information is ~ 416 km. Of course, this does not take into account using multiple RTT measurements for multilateration [9], [11] or using high-resolution spatial information [20].

C. Estimating Host AS Numbers

The determination of network properties such as the AS pair X_a or the network prefix pair X_p from RTT observations and geographical information is of interest in the context of traffic trace anonymisation and dataset privacy, particularly if the dataset includes geographic information. Hence, in this case the estimation difficulty is associated with the capability to reverse engineer un-anonymised information from anonymised, publicly available data sets.

We find that the single best predictor of the ASN pair X_a is given by X_c , the pair of countries in which the hosts are located. This can reduce uncertainty in X_a from its maximum of ~ 11.16 bits to ~ 7.17 bits. By incorporating the geolocation distance X_d , uncertainty is further reduced, leading to a $\mathcal{H}(X_a|X_c, X_d) \approx 5.87$ bits; this is the best approximation using only geographic variables. This can be reduced to ~ 4.4 bits by incorporating X_t . The best estimate for this variable is $\mathcal{H}(X_a|X_c, X_l, X_t, X_d) \approx 3.5$ bits.

Regarding the probability of classification error P_e , for \hat{X}_a we have that an estimator using only X_c will have $P_e \geq .5$. Using X_c and X_d as data sources improves this bound to $P_e \geq .4$, and further inclusion of X_t yields a $P_e \geq .27$. The best lower bound for P_e is obtained by constructing an estimator from X_c , X_l , X_d and X_t , and it provides a $P_e \geq .2$. Hence, although significant network information may be recoverable from anonymised traces, in the absence of traffic injection or multilateration attacks this may be difficult even with additional geolocation information.

V. RELATED WORK

An improved modelling of the relationship between RTT and geographic Internet properties has been at the heart of many works in network-centred host geolocation. In [9] Dong et. al. propose a model of the relationship between RTT and geographic distances using segmented polynomial regression and semidefinite programming. This system builds on the multilateration approach presented in [11], which transforms RTT measurements into geographic distance constraints to infer the location of Internet hosts. As shown in Section IV, geolocation accuracy can be improved by assimilating geographic information. By considering properties such as Internet user density, [23] proposes a *demographic* approach to the placement of landmark nodes and RTT probing endpoints for measurement-based IP geolocation. In [20] Wang et. al. escalate this idea by mining numerous publicly available mapping services and associating the IP addresses of web servers with the advertised addresses of their premises. By relying on this high-resolution data and relative RTT measurements, the technique in [20] can achieve median geolocation error distances of ~ 700 m.

The relationship between IP prefixes and geographic locality has also received attention from the research community. In

[10], Freedman et. al. show that IP prefixes express only very coarse geographic information about the networks that comprise them. Although this is trivially explained for shorter IP prefixes that represent a larger portion of the IP address space, it is also a feature of longer prefixes: the authors find that about 1.4% of prefixes with mask length between /21 and /31 span a distance of more than 160 km, and some /24 prefixes span distances of more than 16,000 km. However, this lack of geographic specificity may not be important to all applications. In [7], Beverly et. al. show that, by training support vector machines on a data set of 30,000 Internet latencies, a mean prediction error of 25 ms is achievable using 20% of the samples for training.

Some previous works aim to elucidate the geographical aspects of interdomain routing. In [19] the authors consider, in addition to network paths, the geographic paths traversed by packets. They show that the distance ratio depends on the geographic and network locations of end hosts and tends to be larger when paths traverse multiple ISPs. Another example is [22], where Zheng et. al. show that both intra- and inter-domain routing policies can naturally give rise to TIVs. For inter-domain, the authors show that both *hot potato routing* (which is used by some Tier 1 providers [19]) and *private peering* shortcuts can introduce TIVs. For the intra-domain case, the authors show that link weights which are disproportional to link delay can yield TIVs.

Finally, there have been several efforts to map the geography of Internet resources. The seminal work of Lakhina et. al. [15] mapped routers to their geographical locations using both geolocation registries and DNS-based host naming heuristics. The authors presented an analysis of interface density across regions, with particular emphasis to its relationship with population density. In addition, they study the relationships between geographic distance and link density, and between the size and geographic extent of ASes.

VI. CONCLUSION

In this paper we presented a model for the large-scale analysis of RTT and its related geographic and network properties based on a novel dataset comprising ~ 19 million RTT measurements derived from ~ 200 million samples between ~ 54 thousand recursive, non-forwarding DNS servers.

We approached the modelling of Internet round trip time by treating our measurements as realisations of a multidimensional random variable, whose distribution $\Phi(\mathbf{x})$ we estimated by constructing a histogram where each observation was tagged with metadata as shown in Table II. Then, we used this histogram to analyse the relationship between RTT and both geographical and network properties.

Our main contribution was an information-theoretical analysis of $\Phi(\mathbf{x})$ that allowed us to determine the amount of information that a given subset \mathbf{Y} of geographic or network variables gives about any given variable of interest X_{\bullet} . Then, we provided bounds on the estimation error that can be expected of an estimator $\hat{X}_{\bullet}(\mathbf{Y})$, which allowed us to compare the relative merits of different sets \mathbf{Y} . This can illuminate the relationships between various network and geographical

variables, and help protocol designers choose which variables to use when estimating network properties of interest.

The dataset upon which this research is based is available to researchers on request.

This research has received funding from the Seventh Framework Programme (FP7/2007-2013) of the European Union, through the ENVISION project (grant agreement 248565).

REFERENCES

- [1] MaxMind GeoLite City. <http://www.maxmind.com/app/geolitecity>.
- [2] Neustar IP Geolocation. <http://www.neustar.biz/solutions/ip-geolocation>.
- [3] The World Factbook. Country comparison: Internet users. <https://www.cia.gov/library/publications/the-world-factbook/rankorder/2153rank.html>.
- [4] GeoNames. <http://www.geonames.org>, 2012.
- [5] Team Cymru Research: IP to ASN Mapping. <http://www.team-cymru.org/Services/ip-to-asn.html>, 2012.
- [6] WIDE Project. <http://www.wide.ad.jp>, 2012.
- [7] R. Beverly, K. Sollins, and A. Berger. SVM learning of IP address structure for latency prediction. In *Proceedings of ACM MineNet*, pages 299–304. ACM, 2006.
- [8] T. M. Cover and J. Thomas. *Elements of Information Theory*. Wiley-Interscience, 1991.
- [9] Z. Dong, R. D. W. Perera, R. Chandramouli, and K. P. Subbalakshmi. Network measurement based modeling and optimization for IP geolocation. *Comput. Netw.*, 56(1):85–98, Jan. 2012.
- [10] M. J. Freedman, M. Vutukuru, N. Feamster, and H. Balakrishnan. Geographic locality of IP prefixes. In *Proceedings of ACM IMC*, Berkeley, CA, USA, 2005. USENIX Association.
- [11] B. Gueye, A. Ziviani, M. Crovella, and S. Fdida. Constraint-based geolocation of internet hosts. *IEEE/ACM Trans. Netw.*, 14(6):1219–1232, Dec. 2006.
- [12] A. Guttman. R-trees: a dynamic index structure for spatial searching. In *Proc. of the ACM SIGMOD*, 1984.
- [13] N. Hopper, E. Y. Vasserma, and E. Chan-TIN. How much anonymity does network latency leak? *ACM Trans. Inf. Syst. Secur.*, 13(2):13:1–13:28, Mar. 2010.
- [14] S. Kaune, K. Pussep, C. Leng, A. Kovacevic, G. Tyson, and R. Steinmetz. Modelling the Internet Delay Space Based on Geographical Locations. In *Proc. of the EuroMicro PDP*, pages 301–310, feb. 2009.
- [15] A. Lakhina, J. Byers, M. Crovella, and I. Matta. On the geographic location of Internet resources. *Selected Areas in Communications, IEEE Journal on*, 21(6):934–948, aug. 2003.
- [16] R. Landa, J. T. Araújo, R. G. Clegg, E. Mykoniati, D. Griffin, and M. Rio. The Large-Scale Geography of Internet Round Trip Times (*under submission*).
- [17] D. Leonard and D. Loguinov. Turbo King: Framework for Large-Scale Internet Delay Measurements. In *Proc. of IEEE INFOCOM*, 2008.
- [18] S. Prasad. Certain Relations between Mutual Information and Fidelity of Statistical Estimation. 2010.
- [19] L. Subramanian, V. N. Padmanabhan, and R. H. Katz. Geographic Properties of Internet Routing. In *Proceedings of the USENIX Annual Technical Conference (ATEC)*, Berkeley, CA, USA, 2002.
- [20] Y. Wang, D. Burgener, M. Flores, A. Kuzmanovic, and C. Huang. Towards street-level client-independent IP geolocation. In *Proc. of USENIX NSDI*, 2011.
- [21] B. Zhang, T. S. E. Ng, A. Nandi, R. Riedi, P. Druschel, and G. Wang. Measurement based analysis, modeling, and synthesis of the internet delay space. In *Proceedings of ACM IMC*, pages 85–98, 2006.
- [22] H. Zheng, E. K. Lua, M. Pias, and T. G. Griffin. Internet Routing Policies and Round-Trip-Times. In *Proc. of PAM*, pages 236–250, 2005.
- [23] A. Ziviani, S. Fdida, J. F. de Rezende, and O. C. M. B. Duarte. Improving the accuracy of measurement-based geographic location of internet hosts. *Comput. Netw. and ISDN Syst.*, 47(4):503–523, Mar. 2005.