

Spatio-Temporal Compressive Sensing and Internet Traffic Matrices (Extended Version)

Matthew Roughan, *Senior Member, IEEE*, Yin Zhang, *Senior Member, IEEE*,
Walter Willinger, *Fellow, IEEE*, Lili Qiu, *Senior Member, IEEE*

Abstract—Despite advances in measurement technology, it is still challenging to reliably compile large-scale network datasets. For example, because of flaws in the measurement systems or difficulties posed by the measurement problem itself, missing, ambiguous or indirect data are common. In the case where such data have spatio-temporal structure, it is natural to try to leverage this structure to deal with the challenges posed by the problematic nature of the data. Our work involving network datasets draws on ideas from the area of compressive sensing and matrix completion, where sparsity is exploited in estimating quantities of interest. However, the standard results on compressive sensing are (i) reliant on conditions which generally don't hold for network datasets, and (ii) don't allow us to exploit all we know about their spatio-temporal structure. In this paper we overcome these limitations with an algorithm that has at its heart the same ideas espoused in compressive sensing, but adapted to the problem of network datasets. We show how this algorithm can be used in a variety of ways, in particular on traffic data, to solve problems such as simple interpolation of missing values, traffic matrix inference from link data, prediction, and anomaly detection. The elegance of the approach lies in the fact that it unifies all of these tasks, and allows them to be performed even when as much as 98% of the data is missing.

Index Terms—Compressed sensing, Interpolation, Prediction methods, Tomography.

I. INTRODUCTION

Network data can often be arranged in the form of multi-dimensional arrays (*e.g.*, matrices or tensors). For example, traffic measurements that specify the traffic volumes exchanged between origin and destination pairs in a network during a given time period naturally form a 2-dimensional matrix or, when viewed over multiple time periods, a 3-dimensional array or tensor. The resulting objects are known as *traffic matrices (TMs)* and are prime examples of network data. They are a critical input to many network tasks including traffic engineering [1], [2], capacity planning [3], and anomaly detection [4], [5]. Due to their importance, there is now a substantial body of work on TMs, for instance see [6] and the references therein. The thrust of much of this research has been on measurement [7], [8] and inference [9]–[16] of TMs, and more recently on topics such as anomaly detection [4], [5], [17]–[19]. At the same time, TMs serve as important examples that network data, due to the engineered nature of the underlying network, are often highly structured, with pronounced spatial or temporal dependencies.

A key challenge that lies at the heart of many problems involving network data is how to cope with missing values. Missing values are common in measurement of real-world systems and can have a detrimental effect on many network tasks, especially those that require complete information. Using the example of TMs and applying novel interpolation techniques for accurately reconstructing missing values in TMs based on direct and/or indirect measurements, we provide in this paper a unified approach to several common network tasks involving measurement and analysis of network data. Our approach uses the first truly spatio-temporal model of TMs and performs well for network tasks ranging from TM estimation and prediction to anomaly detection. A critical feature of our proposed method is that it combines key ideas and recent results from the active area of compressive sensing with detailed domain knowledge that has accumulated over the last few years of TM research.

Motivation: In practice it is challenging to reliably measure TMs for large networks. First, in many networks the TM is not directly observable, and can only be estimated through link load measurements. Such measurements, while linearly related to the TM itself, are not sufficient to unambiguously identify the true TM. Typically, the problem has been posed as an underconstrained linear-inverse problem, where the solution relies on a prior model of the TM (*e.g.*, the Poisson model of Vardi [9], the gravity model [12], [14], or the independent flow model [16]). Second, although many networks now collect (sampled) flow-level measurements for at least part of their network, there are still serious impediments to *reliable* large-scale or network-wide TM data collection: data collection systems can fail, flow collectors often use an unreliable transport protocol, and legacy network components may not support flow collection or be resource-challenged. Third, scalability requirements may mean that flow-level collection cannot occur at the edge of a network (where we would wish it for true TM recovery [7]), but has to be restricted to just some subset of the routers. Finally, when we find an anomaly in a set of TMs, we often need to know the non-anomaly-related traffic either for other network tasks, or just so that we can infer the cause of the anomaly.

Recognizing this aspect of real-world networks, any large set of TM measurements is bound to have some, and quite often, a significant number of missing values, and recovery of the actual ingress-egress TM from such incomplete data is non-trivial. At the same time, the accurate reconstruction of missing values from partial and/or indirect TM measurements

Matthew Roughan is with the University of Adelaide, Australia.
Yin Zhang is with The University of Texas at Austin.
Walter Willinger is with AT&T Labs – Research.
Lili Qiu is with The University of Texas at Austin.

is important because many network engineering tasks that require TMs are either intolerant or highly sensitive to missing data. *Interpolation* is the mathematical term for filling in these missing values. In the context of matrices, interpolation is sometimes referred to as matrix completion, and there is now a developing set of methodologies for solving such problems using *compressive sensing*, a generic methodology for dealing with missing values that leverages the presence of certain types of structure and redundancy in data collected from many real-world systems. Compressive sensing has recently attracted considerable attention in statistics, approximation theory, information theory, and signal processing. Several effective heuristics have been proposed to explicitly exploit the sparse or low-rank nature of empirically obtained matrices [20]–[24]. Meanwhile, the mathematical theory of compressive sensing has also advanced to the point where the optimality of many of these heuristics has been proven under certain technical conditions on the matrices of interest [20]–[24].

Contributions: Despite much recent progress in the area of compressive sensing, our evaluation of existing compressive sensing algorithms on real TMs has shown that they do not perform well [25], and moreover, are not flexible enough to encompass the range of applications we wish to address. Real TM measurements often violate the mathematical conditions under which existing compressive sensing algorithms are designed to operate and are provably optimal. Our approach (originally proposed in [25]) is SPARSITY REGULARIZED MATRIX FACTORIZATION (SRMF). It finds sparse approximations (in the sense that they are low-rank) to TMs, augmented by spatio-temporal operators and local interpolation to achieve high accuracy. While most of the known approaches are best described as purely spatial (*e.g.*, [4], [12], [14], [16], [17]) or purely temporal (*e.g.*, [9], [26]), to the best of our knowledge, SRMF represents the first genuine spatio-temporal model of TMs, and its performance in problems involving interpolation with real TMs is excellent. Low-rank approximation combined with spatio-temporal modelling works well for high-loss scenarios, while local interpolation allows it to perform equally well for low-loss scenarios.

We present in this paper the results of a detailed study where we apply our SRMF approach to a wide range of TM-related problems including network tomography, traffic prediction, and anomaly detection. Using real TMs from three operational networks to evaluate the effectiveness of our approach in these diverse applications, we show that our SRMF-type algorithms consistently outperform other commonly-used methods and do so across the whole range of missing values scenarios, from purely random to highly structured scenarios where whole columns and/or rows of a TM are missing, and from very low levels of missing values to very high levels (*e.g.*, 90% and more). While it may be surprising that such good performance can be achieved with so little data, our results are an indication of the degree of structure present in real-world TMs and the fact that our approach can find and exploit this structure. Indeed, the main reason for the superior performance of our techniques when compared to their most widely-used competitors (*e.g.*, PCA, Tomo-gravity [12], [14],

or independent flows [16]) is their reliance on truly spatio-temporal models of TMs that are capable of capturing much of the global and local structure inherent in actual TMs.

We have applied our approach to matrices with close to a million entries, and our algorithms can process them in only a few seconds. In fact, the algorithms we developed scale linearly with the size of the data so that in principle, much larger matrices can be analyzed. Such larger matrices arise naturally in the context of TMs (*e.g.*, the AS-level TM that gives the traffic volumes exchanged between pairs of Autonomous Systems over some time period), but also in many non-TM-related applications. For example, measuring the performance between origin and destination nodes in a network during some time interval yields so-called *network matrices*, with delay matrices denoting the special case where the performance measure of interest is delay. While these network matrices are notorious for containing many missing values, they are also known to exhibit a significant degree of structure or redundancy. They are therefore natural candidates for our spatio-temporal compressive-sensing inspired approach, and one of our future research efforts will be to extend our approach to such datasets and beyond. Our method will not work on unstructured data, but its effectiveness and robustness to real-world measurement issues are demonstrated in this paper. The results suggest that the approach can be applied when large amounts of missing data exist in much larger and more diverse datasets than considered in this paper, provided they exhibit equally strong spatio-temporal structure.

Last but not least, the fact that we can perform such a variety of network tasks so well with so little data suggests fundamental changes in the way TM measurements are collected in practice. Much as sampling has enabled large-scale flow collection, the type of SRMF-based techniques considered in this paper can enable truly network-wide collection of TM data that is provably robust to the existing impediments to reliable measurement of TMs. In this sense, our work adds traffic matrices to the set of examples where compressive sensing theory asserts that one can accurately recover certain objects from far fewer measurements than traditional methods use [27]. Exploiting this observation in practice for more efficient, effective, and robust TM data collection will be part of our future work.

II. BACKGROUND

A. Traffic Matrices

A Traffic Matrix (TM) is a non-negative matrix $Z(i, j)$ that describes volumes of traffic (in bytes, packets, or flows) between a source i and a destination j . For a network with N locations the TM is a square $N \times N$ matrix. In practice we need a number of addenda to this simple definition. First, a TM is typically measured over some time interval, and the value reported is an average. So we denote $Z(i, j; t)$ to be the traffic from i to j averaged over $[t, t + \Delta t)$. We call the TM $Z(*, *, t)$ a *snapshot* despite the fact that it really represents an interval. Second, although it is common to speak of “origin-destination” TMs, it is often difficult to accurately map IP addresses present in traffic to the true origin and destination of traffic when we

examine a network or Autonomous System (AS). So typically the matrix is aggregated into a router-level *ingress-egress* TM, where $Z(i, j; t)$ describes the traffic entering a network at router i , and leaving at router j .

The TM may be thought of as a 3-dimensional array $Z \in \mathbb{R}^N \times \mathbb{R}^N \times \mathbb{R}^m$ (where there are m time intervals present). It is common to take a TM snapshot and stack the columns to form a column vector which we denote \mathbf{x}_t . We can compile these vectors into the columns of a larger matrix $X \in \mathbb{R}^n \times \mathbb{R}^m$ (where $n = N^2$), and this form of the TM is often more convenient for algebraic manipulation than a 3-dimensional array. The columns of X represent the TM at different times, while the rows represent the time evolution of a flow.

One example of how this notation is useful is in TM inference (the so-called network tomography problem [9]). In this problem the TM is related to the more easily measured link loads Y by the following linear matrix equation

$$Y = AX, \quad (1)$$

where A is the routing matrix, which expresses which links are used by which routes¹. TM inference involves finding the “best” solution \hat{X} to (1) from link-load measurements Y .

More generally, we can combine link measurements with additional TM measurement strategies, which often yields a better estimate of the TM than using each individual type of measurements by itself [15]. For example, flow-records are typically collected at ingress routers [7]. In this case, each router sees one row of a TM snapshot, so over time, router i sees $Z(i, *, *)$. Missing data from a single router means we will be missing a row of Z , or a group of rows of X . Flow-records could also be collected at egress or backbone routers. In this case, although it is difficult to unambiguously determine the ingress router for the observed traffic, we can still form a set of linear constraints on where the traffic could have originated. An alternative measurement strategy [8], [14] is to collect local TMs at each router, which can again be represented as linear constraints on the global TM. In combination we have a set of linear constraints on the TM,

$$A(X) = B, \quad (2)$$

where $A(\cdot)$ is a linear operator, and the matrix B contains the measurements. The operator expresses the information available in our measurements. Note that the presence of missing data is implicit in (2); for instance, the operator \mathcal{A} could include TM measurements at ingress routers with no measurement errors (but with missing data), by writing (2) as

$$M.*X = M.*D, \quad (3)$$

where $D(i, j)$ contains the direct measurements (where available) and M is a $N^2 \times m$ matrix given by

$$M(i, j) = \begin{cases} 0, & \text{if } X(i, j) \text{ is missing.} \\ 1, & \text{otherwise.} \end{cases} \quad (4)$$

and $.*$ denotes an element-wise product, *i.e.*, $A = C.*B$ means $A(i, j) = B(i, j)C(i, j)$. When both link measurements and

direct measurements are available, then constraints (3) will (typically) be incorporated into (2) to simplify notation.

In addition to the above concerns we note that all data sources contain errors. Flow-level collection usually involves sampling, often at quite high rates, and the Simple Network Management Protocol (SNMP) used for collecting link measurements is often noisy [28].

We seek an estimated TM \hat{X} that satisfies the conditions imposed by the set of measurements. However, as is the case in many such linear-inverse problems, there may not be enough information to unambiguously determine X . We call these *underconstrained* linear-inverse problems (in the case of TM estimation from link data, the problem is very highly underconstrained). To solve such problems, we can use side information about the nature of the TM being considered, for instance the gravity model of [12], [14] or independent-flows model of [16]. Regularization is a process used to solve such problems in which we “regularize” towards some prior model of the data in question. The low-rank model we will propose here is motivated by the literature on compressive sensing.

B. Compressive Sensing

Compressive sensing is a relatively new idea in signal processing [20], [21]. The main idea is that many real-world signals or datasets exhibit some structure or redundancy, and we can utilize this prior knowledge for both acquisition and reconstruction of the signal.

Structure and redundancy in data are often synonymous with *sparsity*. A *sparse* vector is simply a vector that has only a few non-zero elements. Often our vectors of interest might have only a few large elements, and many small elements. We call such a vector *compressible*, in the sense that most of its information is carried in the larger elements. Note that the majority of work on compressive sensing has concerned vectors of data, so a naive approach to TMs might be to compile these into vectors and then apply vector techniques. However, some of the structure of a TM is inherent in the matrix itself, so there is value in treating our matrix X as a genuine matrix. In the context of matrices, low rank is analogous to sparsity, because the spectrum formed by the singular values of a low-rank matrix is sparse (see below). It is now well known that TMs may be approximated by matrices of low rank [4], [17], [29], and so this concept fits well here. We explicitly use this type of *structural sparsity* as our approach to resolve the underconstrained nature of the measurement problems we face. In the following section we draw on the recent matrix compressive-sensing literature [22]–[24] to explain how such “sparsity regularized” algorithms can be constructed.

However, we cannot apply existing compressive sensing algorithms. We observed that they performed poorly on real traffic matrices, especially under real-world (*i.e.*, structured, non-random) data loss scenarios. This observation suggests that real TMs often exhibit characteristics that violate the mathematical conditions under which existing compressive-sensing algorithms are designed to operate. Specifically, the optimality results for existing compressive sensing algorithms

¹Issues such as changing network topology or routing have been dealt with in practical network tomography algorithms [5], [14].

often assume that (i) the matrix elements are drawn from a Gaussian-like distribution, (ii) the matrix is exactly low-rank, (iii) data loss is independent for different matrix elements, and (iv) the measurement constraints on the matrix satisfy certain technical conditions (*e.g.*, the restricted isometry property [23]). Unfortunately, none of these conditions hold for real TMs. Real TM elements are non-negative and often exhibit a highly skewed distribution, where the largest and smallest elements often differ in size by several orders of magnitude. Moreover, real TMs are only approximately low-rank, and data loss in real network data tends to be highly structured [28]. Finally, the available measurements are limited by business and technological constraints, not chosen according to a set of desirable properties, and so there is no guarantee that the constraints on real-world TM measurements satisfy technical conditions such as restricted isometry. In fact we have observed that this condition only occurs as a rare exception.

C. Singular Value Decomposition

One tool for creating low-rank matrix approximations is the Singular Value Decomposition (SVD). Simply stated, any $n \times m$ real matrix X can be decomposed such that

$$X = U\Sigma V^T, \quad (5)$$

where V^T is the transpose of V , and U is a $n \times n$ unitary matrix (*i.e.*, $U^T U = U U^T = I$), and V is a $m \times m$ unitary matrix (*i.e.*, $V^T V = V V^T = I$), and Σ is a $n \times m$ diagonal matrix containing the singular values σ_i of X . Typically the singular values are arranged so that $\sigma_i \geq \sigma_{i+1}$. The rank of a matrix is the number of linearly independent rows or columns, which equals the number of non-zero singular values.

To understand SVD's use in matrix approximations, consider the following interpretation of the SVD. The matrix Σ is diagonal, so the SVD of a matrix X can be rewritten as:

$$X = U\Sigma V^T = \sum_{i=1}^{\min(n,m)} \sigma_i \mathbf{u}_i \mathbf{v}_i^T = \sum_{i=1}^{\min(n,m)} \sigma_i A_i, \quad (6)$$

where \mathbf{u}_i and \mathbf{v}_i are the i th columns of U and V respectively, and the matrices A_i are (by construction) rank-1. We can then create a rank- r approximation \tilde{X} from the SVD by keeping only the r largest singular values in the summation and dropping the others:

$$\tilde{X} = \sum_{i=1}^r \sigma_i A_i. \quad (7)$$

The above \tilde{X} is known to be the best rank- r approximation with respect to the Frobenius norm $\|\cdot\|_F$ of the approximation errors, where $\|Z\|_F \triangleq \sqrt{\sum_{i,j} Z(i,j)^2}$ for any matrix Z . That is, truncation of the SVD provides the natural solution to:

$$\begin{aligned} & \text{minimize} && \|X - \tilde{X}\|_F, \\ & \text{subject to} && \text{rank}(\tilde{X}) \leq r. \end{aligned} \quad (8)$$

In Internet measurement, the SVD has most commonly appeared in the form of the Principal Components Analysis (PCA), which has been used, for instance in anomaly detection [4], [17]. PCA is directly related to SVD by the fact

that the columns of U form the principal axes of the PCA coordinate transform.

Sparsity Regularized SVD: Many matrix interpolation techniques try to create a SVD-like factorization of a matrix X , though it is common to write this in the equivalent form

$$X = U\Sigma V^T = LR^T, \quad (9)$$

where $L = U\Sigma^{1/2}$ and $R = V\Sigma^{1/2}$, and so we will subsequently use this form for consistency.

Now SVD by itself is not an interpolation algorithm. Typical algorithms for calculating the SVD assume that X is completely known. Where there is missing data we can instead look for a factorization that satisfies the measurement equations, *i.e.*, $\mathcal{A}(LR^T) = B$. Previous studies have suggested that typical TMs inhabit a relatively low-dimensional subspace [4], [17], so an intuitive approach for finding such a factorization is to solve the following rank minimization problem:

$$\begin{aligned} & \text{minimize} && \text{rank}(LR^T), \\ & \text{subject to} && \mathcal{A}(LR^T) = B. \end{aligned} \quad (10)$$

Rank minimization has a non-convex objective and is thus difficult to solve. A key insight from the matrix compressive-sensing literature [22]–[24] is that under certain conditions, we can solve a simpler problem instead and obtain equivalent results. Specifically, when a certain technical condition (the restricted isometry property [23]) holds on $\mathcal{A}(\cdot)$, then a heuristic — minimizing the nuclear norm — can perform rank minimization exactly for a matrix of low enough rank. Further, if the rank of X is less than the rank of LR^T then (10) is equivalent to

$$\begin{aligned} & \text{minimize} && \|L\|_F^2 + \|R\|_F^2, \\ & \text{subject to} && \mathcal{A}(LR^T) = B. \end{aligned} \quad (11)$$

In TM interpolation, looking for a low-rank solution that strictly satisfies the measurement equations is likely to fail, because (i) the real TM X is often only approximately low-rank, and (ii) the measurements often contain errors. So instead we solve the following

$$\text{minimize } \|\mathcal{A}(LR^T) - B\|_F^2 + \lambda (\|L\|_F^2 + \|R\|_F^2). \quad (12)$$

This solution aims for a low-rank approximation without strictly enforcing the measurement equations. The regularization parameter λ allows a tunable tradeoff between a precise fit to the measured data and the goal of achieving low rank.

We derive L and R from (12) using an alternating least squares procedure. We initialize L and R randomly. We then solve the above optimization taking one of L and R to be fixed and the other to be the optimization variable (which is a standard linear least squares problem). We then swap their roles, and continue alternating towards a solution till convergence. Our implementation of the alternating least squares procedure in Matlab further utilizes sparse matrix operations to minimize memory requirement and maximize speed (details are omitted due to space restriction, but we will supply Matlab code on request). The above approach will be referred to as *Sparsity Regularized SVD (SRSVD)* interpolation.

D. Other Interpolation Algorithms

There are a number of approaches that have been proposed in the recent literature for matrix interpolation besides SVD. These algorithms can be classified as either low-rank approximation algorithms or local interpolation algorithms, depending on whether they exploit the global low-rank structure or the local structure and redundancy. We describe them here for comparison with our approach.

1) *Low-Rank Approximation Algorithms: Baseline Approximation:* Many techniques (for instance PCA) implicitly assume that the data has zero mean. So our first step for dealing with network matrices might be to “center” them. However, centering the matrices where we do not have all the data also requires interpolation. Our baseline approximation algorithm implicitly constructs such an interpolation matrix X_{base} to compute row and column means of the matrix. For instance, if we knew all elements of the input X , then the row and column sums of $X - X_{\text{base}}$ would all equal zero. Apart from its use in zeroing the mean, it also forms an interpolation in its own right, and so we will compare its performance below.

variable	description
\bar{X}	an estimate of the mean of X over all i and j .
X_{row}	a vector of length m such that $X_{\text{row}}(i) = \text{an estimate of } \sum_j (X(i, j) - \bar{X})/n$.
X_{col}	a vector of length n such that $X_{\text{col}}(j) = \text{an estimate of } \sum_i (X(i, j) - \bar{X})/m$.
X_{base}	the baseline estimate for X given by $X_{\text{base}}(i, j) = \bar{X} + X_{\text{row}}(i) + X_{\text{col}}(j)$.

TABLE I
OUTPUTS OF BASELINE ESTIMATION.

To compute X_{base} , we use the variables described in Table I. In matrix form, we can represent X_{base} as a rank-2 approximation to X : $X_{\text{base}} = \bar{X} + X_{\text{row}}\mathbf{1}^T + \mathbf{1}X_{\text{col}}^T$, where $\mathbf{1}$ is a column vector consisting of all ones. We use the regularized least-squares algorithm from [30] to compute \bar{X} , X_{row} , X_{col} from input $\mathcal{A}(\cdot)$ and B . That is, we solve the following

$$\begin{aligned} \text{minimize} \quad & \|\mathcal{A}(\bar{X} + X_{\text{row}}\mathbf{1}^T + \mathbf{1}X_{\text{col}}^T) - B\|_F^2 \\ & + \lambda (\bar{X}^2 + \|X_{\text{row}}\|_F^2 + \|X_{\text{col}}\|_F^2), \end{aligned} \quad (13)$$

where λ is a regularization parameter. The first term in this formulation minimizes the Frobenius norm of the difference $\mathcal{A}(X_{\text{base}}) - B$, and the second regularization term helps avoid over-fitting.

SRSVD-base: Techniques like PCA implicitly assume that the data has zero mean, but in TM interpolation we do not know the true mean. Instead we use X_{base} as an estimate. It is not obvious whether such centering is necessary or desirable in interpolation, so we include results for both SRSVD applied to X and SRSVD applied to $(X - X_{\text{base}})$. We refer to the latter as SRSVD-base.

Nonnegative Matrix Factorization: Nonnegative Matrix Factorization (NMF) [31], [32] tries to find nonnegative factor

matrices L and R that minimize the Frobenius norm² of the difference $\mathcal{A}(LR^T) - B$ over the observations. The approach is very similar to the SVD, except for insistence on non-negative factor matrices. We avoid over-fitting by regularizing in the same manner that we do for SVD, *i.e.*, we solve (12) but with the additional constraint of non-negativity. We implement the two most common algorithms for NMF: multiplicative update [32] and alternating nonnegative least squares. Both algorithms are designed for the case where the matrix X is completely known. So we extend them to further cope with measurement equations (2). The two algorithms give similar interpolation performance, but multiplicative update is more efficient. So our results are based on this algorithm.

2) *Local Interpolation Algorithms:* We also test one completely different approach: k -Nearest Neighbors (KNN). Simple nearest neighbors directly uses the nearest neighbor of a missing value for interpolation. KNN extends this by using a weighted average of the k nearest-neighbors’ values. For TMs (and many other datasets) it is difficult to apply KNN because the rows are ordered arbitrarily (for instance based on the names of routers). So the nearest elements in the matrix X may have little correspondence. Hence we need to derive a good distance metric between matrix elements.

We use the approach described in [30]. We can perform the algorithm on either rows or columns of X , but let us start with rows. If two rows are similar (*i.e.*, two TM elements exhibit similar behavior), then it is natural to assume that one might provide a good interpolant of the other. Hence, we base our distance metric on the similarity between rows, *i.e.*, the more similar two rows are, the closer together we consider them. Following [30], we measure the similarity by an approximation to the correlation coefficient of the two rows based on only those directly observed TM elements. To form this coefficient, we would ideally first subtract the mean, but as the mean is unknown we use our proxy X_{base} . The weights used in the k averaged neighbors are proportional to the similarities [30].

Formally, the algorithm proceeds by

- step 1 Form $X' = X - X_{\text{base}}$ where data is available, and zero otherwise.
- step 2 Compute similarities S between rows of X'

$$C(u, v) = \sum_k X'(u, k)X'(v, k).$$

$$S(u, v) = C(u, v) / \sqrt{C(u, u)C(v, v)}.$$
- step 3 Find the k -nearest neighbors $N(i, j)$ for each missing point (i, j) ignoring rows k missing points (k, j) .
- step 4 estimate missing value $X(i, j)$ using weighted average

$$X_{\text{knn}}(i, j) = X_{\text{base}} + \frac{\sum_{k \in N(i, j)} S(k, j)X'(k, j)}{\sum_{k \in N(i, j)} S(k, j)}.$$

In our results we found that temporal similarity was stronger than spatial similarity, so we perform the above algorithms

²There is nothing intrinsically special about the Frobenius norm for this approach. The Kullback-Leibler divergence [32] has also been suggested but our experiments found that the performance of this approach was much worse, and it is not presented here.

on columns first, and then on rows if enough neighbors with positive similarity are not found.

An advantage of KNN is that it is computationally straight forward. It does not rely on convergence of some non-convex optimization problem. However, the choice of k is problematic. When applying this method, we found that in some cases larger values were needed, but often the reverse was true.

III. OUR SOLUTION: SPATIO-TEMPORAL COMPRESSIVE SENSING

The methods above are non-adaptive in the sense that they do not account for the special properties of a particular traffic matrix. Particular TMs have been shown to have strong cyclical behavior [33] (often linked to the diurnal or weekly cycle of network consumers), and strong spatial structure [34]. However, the spatial structure, in particular, is not known *ab initio*, and so we need to be able to estimate the type of spatial structure operations to use for a particular traffic matrix. The result is an algorithm that can adapt to the structures in the data, other than the pure structural sparsity exploited by the majority of the prior algorithms.

The KNN approach is intrinsically different from the other methods described above. It explicitly targets local structure in a TM, whereas the low-rank methods look for global structure, and it estimates this structure as part of the algorithm. This is one of the key motivations for developing a method that seeks to capture both global and local structure. It consists of two key components: (i) SPARSITY REGULARIZED MATRIX FACTORIZATION (SRMF) for incorporating global spatio-temporal properties, and (ii) a mechanism for incorporating local interpolation.

A. Sparsity Regularized Matrix Factorization

The SRSVD approach starts with (12) to find global low-rank structure in the TM. On the other hand, we may *a priori* know that the matrix has additional spatio-temporal structure, *e.g.*, TM rows or columns close to each other (in some sense) are often close in value. We seek to exploit this insight in the new technique we propose here. We propose to solve the following

$$\begin{aligned} \text{minimize} \quad & \|A(LR^T) - B\|_F^2 + \lambda (\|L\|_F^2 + \|R\|_F^2) \\ & + \|S(LR^T)\|_F^2 + \|(LR^T)T^T\|_F^2, \end{aligned} \quad (14)$$

where S and T are the spatial and temporal constraint matrices, respectively. Matrices S and T express our knowledge about the spatio-temporal structure of the TM (*e.g.*, temporally nearby TM elements have similar values). We solve the above optimization problem again using alternating least squares. We call the resulting algorithm *Sparsity Regularized Matrix Factorization (SRMF)*. It has the advantages of SRSVD, but is more general, allowing us to express other objectives in our TM approximation/interpolation algorithm through different choices of S and T .

Below we discuss how to choose S and T . To better illustrate the idea and benefit of SRMF, we intentionally use relatively simple choices of S and T . Both SRSVD and SRMF

also require specification of the input rank of L and R . Our evaluation in Section IV shows that SRMF is not sensitive to the input rank parameter.

Choice of T : The temporal constraint matrix T captures the temporal smoothness of the TM. A simple choice for the temporal constraint matrix is $T = \text{Toeplitz}(0, 1, -1)$, which denotes the Toeplitz matrix with central diagonal given by ones, and the first upper diagonal given by negative ones, *i.e.*,

$$T = \begin{bmatrix} 1 & -1 & 0 & \dots \\ 0 & 1 & -1 & \ddots \\ 0 & 0 & 1 & \ddots \\ \vdots & \ddots & \ddots & \ddots \end{bmatrix} \quad (15)$$

This temporal constraint matrix intuitively expresses the fact that TMs at adjacent points in time are often similar. For instance XT^T is just the matrix of differences between temporally adjacent elements of X . By minimizing $\|(LR^T)T^T\|_F^2$, we seek an approximation that also has the property of having similar temporally adjacent values. A more sophisticated choice taking into account domain knowledge (say knowledge of the periodicity in traffic data) can result in some improvements, and allows almost arbitrary temporal structure into the approach. In general, it is not difficult to develop such temporal models of TMs, as they show strong signs of periodicity over 24 hours and 7 day intervals (for obvious reasons), however, in more difficult cases a more adaptive approach such as used for S below could be used.

Choice of S : The spatial constraint matrix S can be used to express which rows of a TM are close to each other, but due to the arbitrary ordering of rows in the TM, a simple matrix of the above form is not appropriate. We find S by first obtaining an initial TM estimate \hat{X} using a simple interpolation algorithm, and then choosing S based on the similarity between rows of \hat{X} (which approximates the similarity between rows of X).

1. *Computing \hat{X} .* In our current implementation, we take $\hat{X} = X_{\text{base}} * (1 - M) + D * M$, where $*$ means element-wise product, M is defined in (4) and specifies which TM elements are directly measured, and D contains the direct measurements. That is, we use direct measurements where available, and interpolate using X_{base} at other points.
2. *Choosing S based on \hat{X} .* There are many possible methods for choosing S based on \hat{X} . For example, one general method is to (i) construct a weighted graph G , where each node represents a row of \hat{X} and each edge weight represents certain similarity measure between two rows of \hat{X} , and (ii) set S to be the normalized Laplacian matrix [35] of graph G , which acts as a differencing operator on G and induces sparsity by eliminating redundancy between similar nodes of G (*i.e.*, rows of \hat{X}).

We have experimented with several such methods. The following method for choosing S based on KNN and linear regression consistently yields good performance in our tests, and we use it in our evaluation. For each row i of \hat{X} , we find the K most similar rows $j_k \neq i$ ($k = 1, \dots, K$). We perform linear regression to find a set of weights $w(k)$

such that the linear combination of rows j_k best approximates row i : $\hat{X}(i, *) \approx \sum_{k=1}^K w(k) \hat{X}(j_k, *)$. Assuming that $\sum_{k=1}^K w(k) X(j_k, *)$ approximates $X(i, *)$ well, we then set $S(i, i) = 1$ and $S(i, j_k) = -w(k)$.

Scaling of S and T : Finally, we need to scale S and T properly so that $\|S(LR^T)\|_F$, $\|(LR^T)T^T\|_F$, and $\|A(LR^T) - B\|_F$ are of similar order of magnitude — otherwise they may overshadow each other during the optimization of (14). In our experiments, we simply scale S and T such that $\|S\hat{X}\|_F = 0.1\sqrt{\lambda}\|B\|_F$ and $\|\hat{X}T^T\|_F = \sqrt{\lambda}\|B\|_F$, where $\sqrt{\lambda}\|B\|_F$ reflects the level of approximation error $\|A(LR^T) - B\|_F$ that we are willing to tolerate. Our results show that such scaling yields good performance over a wide range of scenarios and that the performance is not sensitive to the choice of λ . Note that we intentionally make $\|S\hat{X}\|_F$ smaller than $\|\hat{X}T^T\|_F$ because we expect the temporal model obtained through domain knowledge is more reliable.

B. Combining Global and Local Methods

A quick preview of the (non-hybrid) performance results that follow shows that for small amounts of missing data, KNN is the best performer (in most cases). On the other hand, for large amounts of loss, SRMF outperforms KNN. The intuition behind this result is obvious. When only a few data points are missing, the k -nearest neighbors of a missing data point will be close by. There are strong temporal and spatial correlations in our data, so the nearest neighbors provide good interpolants for the missing data. However, when there is substantial missing data, the nearest neighbors will come from further away. As the correlations between data points drop, the low-rank global model of the data expressed by the matrix factorization becomes superior.

To take advantage of local structure and redundancy present in the TM, we use the low-rank approximation obtained by SRMF as a prior and augment it with a local interpolation procedure. In this way, we obtain a TM estimate that is close to the low-rank prior yet can account for constraints imposed by the local interpolation procedure. Note that such an approach generalizes the Tomo-gravity method for TM estimation [12], [14], which uses a rank-1 approximation (*i.e.*, gravity model) as the prior solution.

The choice of the local interpolation procedure is application dependent, and we present below two hybrid algorithms, *SRMF+KNN* and *SRSVD-base+KNN*, both of which incorporate KNN.

SRMF+KNN: We first compute the SRMF interpolation of X . Call this X_{SRMF} . For each missing data point (i, j) we then examine its row to see if any of the elements $X(i, j-3), \dots, X(i, j+3)$ are present. If we cannot observe any of these neighbors, then we simply use the value $X_{\text{SRMF}}(i, j)$, but if we do have any of these values, we will use them to better approximate $X(i, j)$. We do so by forming a local model for the temporal process using all of the other rows of the TM. We perform a regression to find a set of weights $w(k)$ that best approximates $X_{\text{SRMF}}(p, j) = \sum_{k \in \text{nbrs}} w(k) X_{\text{SRMF}}(p, k)$ for all $p = 1, 2, \dots, n$. Then we apply a weighted linear interpolation of the nearest neighbors, using the weights derived

above, *i.e.*,

$$X_{\text{SRMF+KNN}}(i, j) = \sum_{k \in \text{nbrs}} w(k) X(i, k). \quad (16)$$

SRSVD-base+KNN: We will show that the above approach is superior, but to understand the importance of incorporating the spatio-temporal constraints (given by S and T), we also consider an algorithm that uses SRSVD-base as the prior in the same procedure. We call the resulting algorithm *SRSVD-base+KNN*.

IV. INTERPOLATION PERFORMANCE ON TM DATA

For the basic network task of TM interpolation, we illustrate in the following the performance of our approach with real-world TM data where we generate a whole range of missing values scenarios, from purely random to highly structured scenarios, and from low levels of missing values to very high levels. Comparisons to commonly-used alternative techniques attest to the effectiveness and robustness of our approach to real-world measurement issues.

A. Data

The main dataset we use here is real TM data. The first two are the Abilene (Internet2) [36] dataset used previously in various studies [4], [5], [17], and the GÉANT TM dataset provided in [37], and previously examined in [6]. Although these are now older datasets, we use them because they are valuable for comparisons with other work. In addition, we use one longer and more recent commercial TM dataset from a large Internet service provider.

In addition, we use one set of SNMP link load data (also from Abilene) collected using the RRD tool [38]. Although this is still traffic data, it has rather different characteristics to the raw traffic matrix data. As with all SNMP data we only see link loads, and so there are only 28 rows in the data, but such data is easier to collect, process and analyze, and so we have a longer dataset with finer resolution, *i.e.*, the matrix we study in this case has rather different size from our traffic matrices (which are not so far from square).

The properties of the data are summarized in Table II.

Dataset	Date	Duration	Resolution	Size
Abilene TM	Apr. 2003	1 week	10 min.	121×1008
Commercial TM	Oct. 2006	3 weeks	1 hour	400×504
GÉANT TM	Apr. 2005	1 week	15 min.	529×672
Abilene SNMP	Jan 4th-Mar. 28th, 2006	12 weeks	5 min.	$28 \times 23,498$

TABLE II
DATASETS UNDER STUDY.

B. Methodology

The methodology we use here is to drop some data from existing measurements, and then apply the interpolation algorithm. This provides us with ground truth for comparison. The pseudo-missing data is not used in the interpolation algorithms in any way.

The typical approach when measuring the performance of an interpolation algorithm is to drop data at random. We will start our experiments with this case. However, in real measurements of TMs there are different mechanisms that result in missing data, and these result in the missing data having structure. Such structure is obviously important for interpolation, so we will explore several structured models of missing data in Section 4.5 below.

We measure performance using the Normalized Mean Absolute Error (NMAE) in the interpolated values. That is, we calculate

$$NMAE = \frac{\sum_{i,j:M(i,j)=0} |X(i,j) - \hat{X}(i,j)|}{\sum_{i,j:M(i,j)=0} |X(i,j)|}, \quad (17)$$

where \hat{X} is the estimated matrix. Note that we *only measure errors on the missing values*. So the NMAE is defined only when there is at least one missing value, and will not approach zero as the proportion of missing data decreases. We computed three other performance metrics (root mean squared error, normalized root mean squared error, and the correlation-coefficient) but the results are substantively the same. In each case we perform the process of randomly dropping data and reconstructing the matrix 10 times. The results presented show the mean NMAE.

C. Initial Comparisons

Figure 1 shows a comparison of algorithms for independent random loss for data loss rates ranging from 0.02 to 0.98. We perform these algorithms using the same regularization and input rank parameters $\lambda = 0.1$ and $r = 8$ for each global algorithm, and $k = 4$ in KNN (we defer justification of these choices to the section below).

First consider the three TM datasets shown in Figure 1 (a)-(c). For low loss probabilities KNN achieves better performance than SRMF. For high loss probabilities we see that SRMF's performance exceeds KNN. However, the hybrid SRMF+KNN outperforms all algorithms over the whole range of loss values. Interestingly, the hybrid is noticeably better than either method individually.

Meanwhile, the hybrid SRSVD-base+KNN also performs well, though not as well as SRMF+KNN. The performance gap typically widens for large amounts of loss. This is because under independent random loss, when the loss rate is not too high, it is likely that the near neighbors of a missing value are directly observed, making KNN an effective recovery strategy. However, when loss is large or when the loss is highly structured (see Section IV-E), the performance gap between SRSVD-base+KNN and SRMF+KNN widens.

The other methods all have worse performance. For low loss, the baseline method is the worst (as we might expect given it is only a rank-2 approximation). However, for high loss, the baseline performs surprisingly well, certainly better than SRSVD, whose performance is very bad for high loss. However, the SRSVD applied after baseline removal achieves reasonable performance over the whole loss range, in some cases almost as good as the simple SRMF. NMF performs poorly for all loss probabilities.

We have examined many such graphs. Several approaches are uniformly poor, so we do not examine them in further detail, in order to simplify our presentation.

Figure 1 (d) shows the algorithms applied to SNMP link data. The results are substantively similar, though the difference between the various algorithms is much smaller, and for small loss, SRSVD-base+KNN performs marginally better than SRMF+KNN. The size of the matrix in this dataset is very different (it has only 28 rows, but many more columns) than our TM datasets. We hypothesize that the smaller number of rows allows a lower-rank approximation to fit the data better, and that the larger number of columns provides more data, so most techniques can perform better in this case, but nevertheless, SRMF+KNN still shows the best performance over the widest range of parameters.

D. Parameter Sensitivity and Settings

The algorithms we consider have several input parameters. The performance of these algorithms in relation to these parameters is (in most cases) dependent on the dataset in question. In practice, when interpolating a real dataset, we would not be able to precisely optimize λ and r for the dataset in question, so it is desirable to have algorithms that are not very sensitive to their values. In fact, all algorithms display some dependence on the parameter settings, and no single parameter setting is optimal for all datasets. However, we found rough parameter settings that are never too far from optimal.

The first input parameter is the rank. Given our motivation from the compressive sensing literature, *i.e.*, that we will aim to minimize matrix sparsity or rank, it may seem strange that we input a rank when performing the algorithm. However, although they seek to minimize the rank of the decomposition, the algorithms work by optimizing an L and R that have a fixed number of columns r (the input rank). The final rank of the solution might be smaller.

In theory, as long as the input rank is greater than the real rank of X , the various algorithms will converge to the correct matrix [22]–[24]. However, note that the theoretical results that inform our intuition here concern matrices with exact ranks, whereas our matrices typically have a number of small, but non-zero singular values. Moreover, there are measurement errors in our data, so we cannot expect to get zero error reconstructions.

Figure 2 shows a sample of performance results with respect to rank (note that the baseline algorithm is excluded here because it is a fixed rank-2 approximation). We find that most of the rank-dependent methods have better performance as the input rank increases. Although this is not always the case, the deviations are minor. However, note the logarithmic x -axis, so that the results suggest a decrease in the marginal improvement with increasing rank. There is also an additional computational cost for higher ranks, and we find that an input rank of $r = 8$ is a reasonable operating point for use on TM data. Going to $r = 16$ yields only a very small incremental improvement at the expense of extra computation. On the SNMP link data, with fewer rows, we found that an even lower rank ($r = 4$) performed as well, or sometimes better.

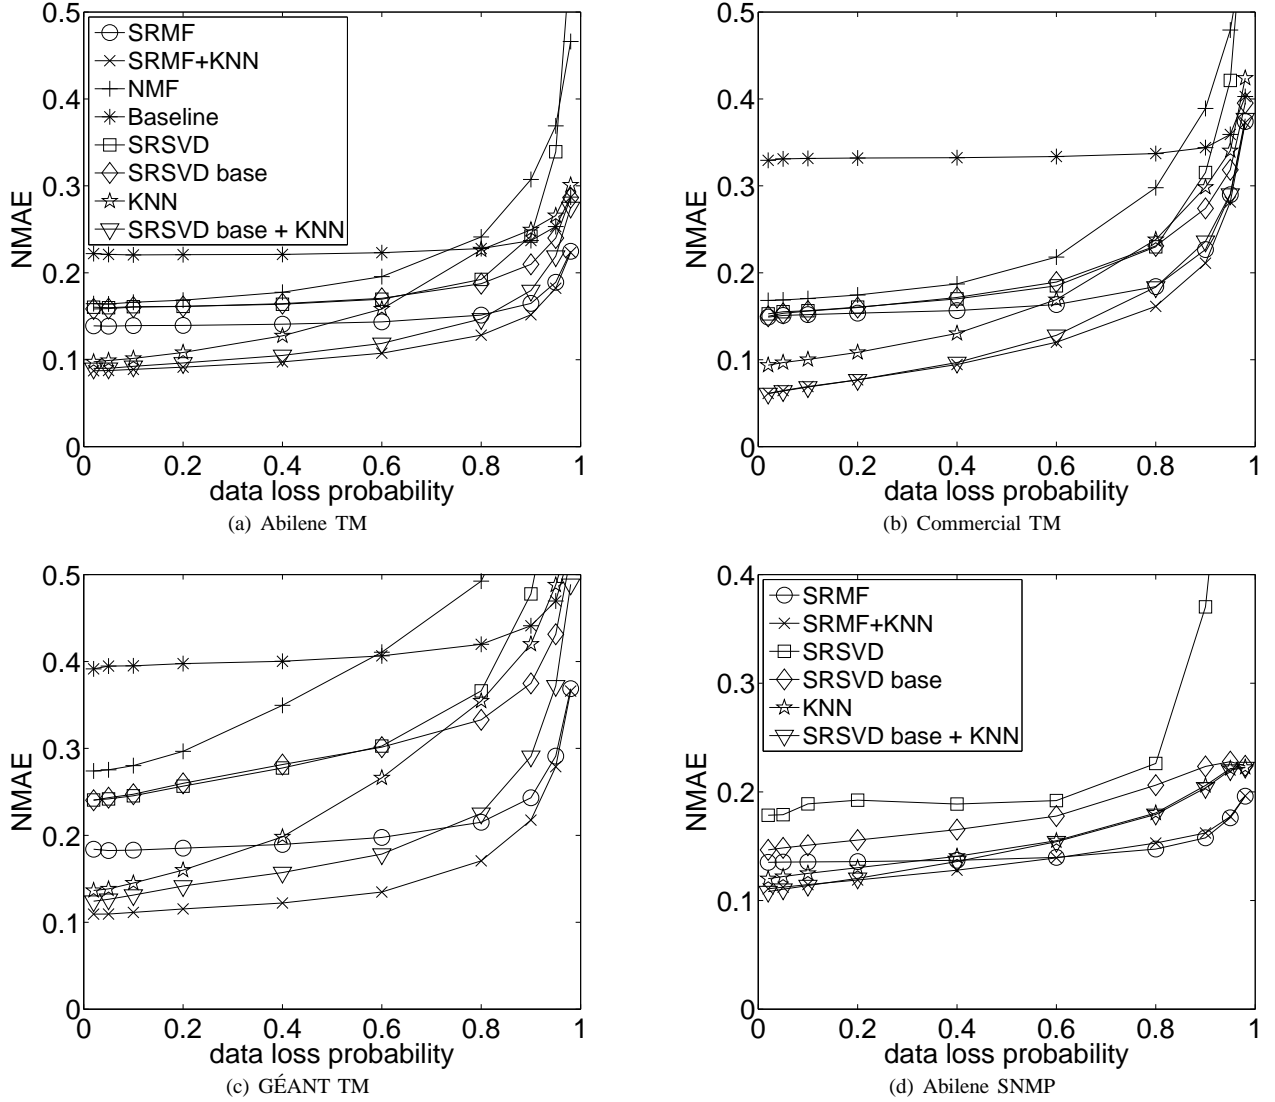


Fig. 1. Interpolation performance for random loss (note that the legend is the same for plots (a-c) but we omit poorly performing approaches from plot (d) to make the plot more legible).

The most important finding in these results, however, is the relative insensitivity of the hybrid algorithm, SRMF+KNN. In general it is the least dependent on the input rank of all the algorithms. There is some improvement for higher ranks, but it is typically smaller than those of other algorithms.

KNN does not use input rank, but rather k , the size of neighborhood. Figure 2 also shows the effect of k on the performance of KNN. We choose to use $k = 4$ for our experiments, since it consistently avoids the worst results.

The final parameter of importance is the regularization parameter λ , which determines the tradeoff (in the optimization) between the measurement constraints and the importance of rank. Larger λ leads to lower rank approximations, whereas smaller values lead to approximations that are a better fit to the data. Figure 3 presents three examples showing the type of variability we encounter over a range of values of λ , for three different loss rates and networks. KNN is omitted because it does not use regularization. Once again note the logarithmic x -axis – we are looking for order of magnitude effects here, not

fine tuning. None of the techniques is too sensitive. Among them, SRSVD is the most sensitive (overall). Larger values of λ typically perform better although again sometimes this trend is reversed, and there are a number of cases where the optimal case is around $\lambda = 0.1$. So we use this value in our experiments.

Note again that SRMF+KNN is the most insensitive algorithm with Figure 3 (c) showing the most extreme case of parameter sensitivity that we observed for this algorithm.

E. Comparison: Other Loss Models

As earlier noted, not all data loss is random [28]. Losses are often highly structured, and in this section we examine the effect this has on the results. The boldface name denotes the label used in our datasets, where **xx** is replaced by the percentage of rows (or columns) effected.

1. **PureRandLoss**: This is the simple random loss model. Data points in the matrix X are dropped independently at random with probability q .

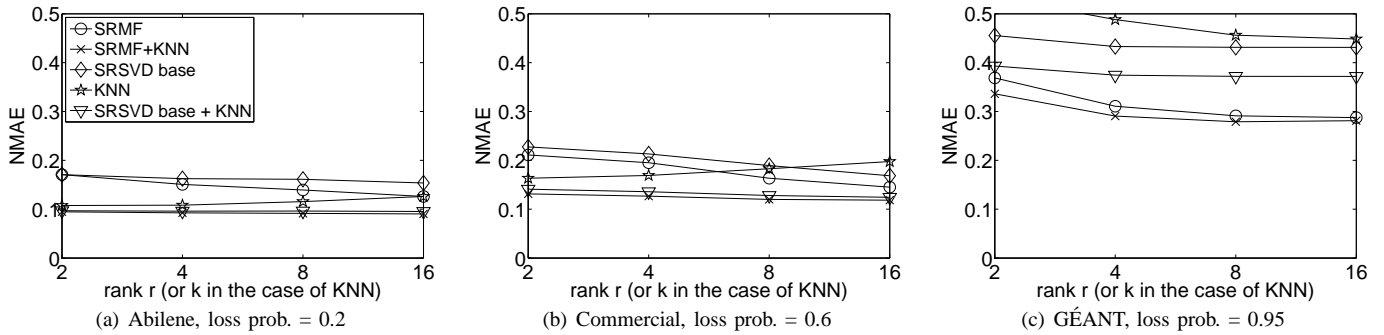


Fig. 2. Sensitivity with respect to the input rank r (or k in the case of KNN).




Fig. 3. Sensitivity with respect to λ .

2. **xxTimeRandLoss:** This simulates a structured loss event at certain times if, for example, our monitoring equipment became overloaded. In these cases, we may lose some random proportion of the data at a particular point in time. We simulate this loss by choosing, at random, $xx\%$ of the columns of X , and dropping data from these at random with probability q .
3. **xxElemRandLoss:** This simulates a structured loss event where a set of randomly chosen TM elements suffers from lost data. This type of loss might occur where unreliable transport mechanisms are used to transport measurements. Often the problems with such transport depend on the locations where measurements are made (*e.g.*, locations close to the management station are less likely to suffer congestion based losses). We randomly select $xx\%$ of the rows of X to be effected. Note that the case 100ElemRandLoss corresponds to PureRandLoss.
4. **xxElemSyncLoss:** This simulates a structured loss event where a group of TM elements all suffer from missing data from the same cause. Hence, the losses on each element are synchronized. We do so by selecting $xx\%$ of rows of X to be effected, and a set of times with probability q . Lost data comes from the intersection of the selected rows and columns.
5. **RowRandLoss:** Random element loss, as presented above, is not a particular realistic model for data loss. With flow level measurements, data are collected by a router. If that router cannot collect data, then an entire row of each TM snapshot Z will be missing. The effect on X is to remove a set of structurally associated rows. We simulate this by dropping rows from the original TM Z (before it is formed

into the matrix X).

6. **ColRandLoss:** It is perhaps less likely that a column of the original TM Z is dropped from measurement. One can construct scenarios where a software bug causes such an error, but in fact we primarily consider the random column loss scenario for completeness.

In this section we examine the impact of the loss model on the performance of the interpolation algorithms. Obviously there are many ways of viewing this data. Due to space limitations, we present here only two summary figures (more details appear in [25]). Figure 4 shows bar charts of the performance of the key algorithms for two different loss levels, across all loss models. The key observations are that for low-to moderate loss, SRMF+KNN performs significantly better across all loss models. When loss is higher, there are some cases where the performance of SRSVD-base and KNN is similar to SRMF+KNN, and occasionally slightly better, but where losses are highly structured (*e.g.*, AllElemSyncLoss) SRMF+KNN is always clearly superior.

F. Computational Times

We measure the computation times of SRMF (KNN takes a small amount of additional time) on matrices of various sizes. The computation time is linear in the number of matrix elements, and quite reasonable. A 200×1000 matrix (with rank $r = 10$ used in the decomposition) can be approximated in less than 3.5 seconds (using Matlab on a 2.66 GHz Linux machine). Meanwhile, the computation time with respect to r is reasonably modeled by $O(r^2)$. Figure 5 shows computation times with respect to the number of rows and columns of the matrix, and the input rank r .

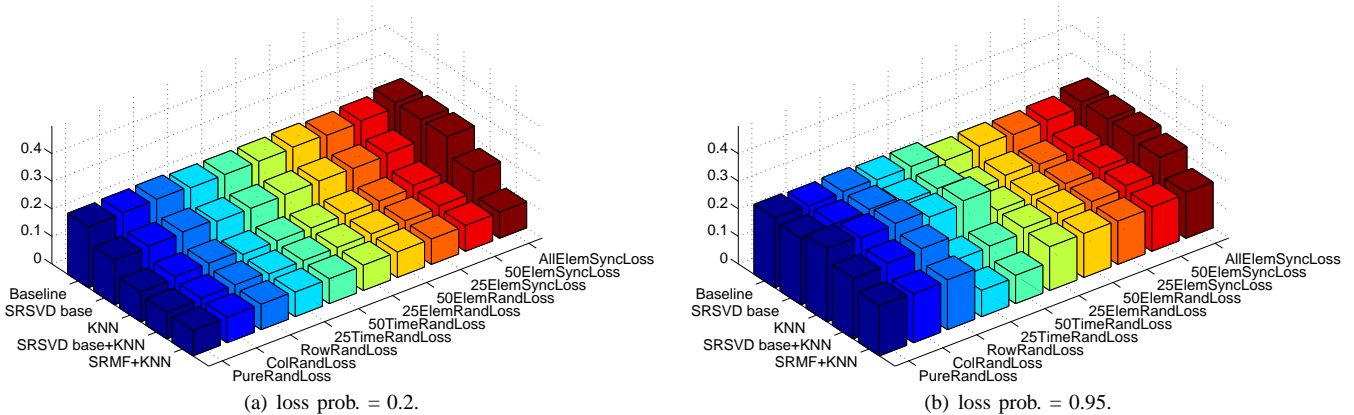


Fig. 4. Comparison between algorithms for the different loss models.




Fig. 5. Computation times (seconds) for a $n \times m$ matrix using a rank r decomposition.

V. APPLICATIONS

A key contribution of this paper is demonstrate the effectiveness and robustness of our approach beyond the simple problem of interpolation to a much wider set of network tasks. To this end, we consider the following different applications of matrix interpolation, where the meaning or importance of the missing values are determined by the application in question. The general approach we use here is to remove some data from existing measurements and then examine how well we can perform a particular application (*e.g.*, tomography, prediction, anomaly detection) without this data. For instance, in the case of prediction, we drop future data points and try to estimate these using the remaining data. This provides us with the ground truth for comparison. The pseudo-missing data is not used in the interpolation algorithms so that we can perform fair comparisons between algorithms.

A. Tomography

A special case of the our approach is the network tomography problem of inferring a TM from link-load measurements. In the interpolation examples discussed earlier, the constraints come from direct measurements. In network tomography, the constraints are given by (1). However, it is common to have some combination of these sets of measurements. So it is desirable to combine them to obtain the best possible approximation to the TM. In this case, we can simply define \mathcal{A} to incorporate both (1) and (3), resulting in a combined penalty term of the form $\|A(LR^T) - Y\|_F^2 + \|(LR^T - D) * M\|_F^2$.

We do not compare all possible algorithms for TM estimation. For instance, most recently Bharti *et al.* [29], used similar techniques to solve a problem with similar goals: discovery of un-observed traffic. However, the goals of [29] are somewhat broader – estimating a traffic matrix for the whole Internet from the limited vantage that any one researcher or network manager is likely to have, from a somewhat different dataset (netflow) which we do not assume available in this section, and moreover, the work assumes temporal information is not available. Further performance improvements might be obtained by using more recently developed algorithms (*e.g.*, [16], which also includes a temporal component). However, we concentrate here on two existing algorithms – the gravity model and Tomo-gravity – because they use a consistent set of data, and have been used in real network operations. The gravity model [14] is a simple rank-1 approximation to a single TM. It is known to be a poor estimator of real TMs, but it has been successfully used as the first step in the Tomo-gravity [14] algorithm. The latter is yet another regularization based on the Kullback-Leibler divergence between the gravity model and the measurements.

In this section we compare the gravity model and Tomo-gravity against three alternatives: the baseline approximation, SRMF, and Tomo-SRMF. In Figure 6 we show the performance of the algorithms with respect to the proportion of the TM elements that are missing, but note that in addition to direct measurement of the matrices, we assume we can measure all of the link loads on the networks. So in this figure, 100% data loss corresponds to the standard network tomog-

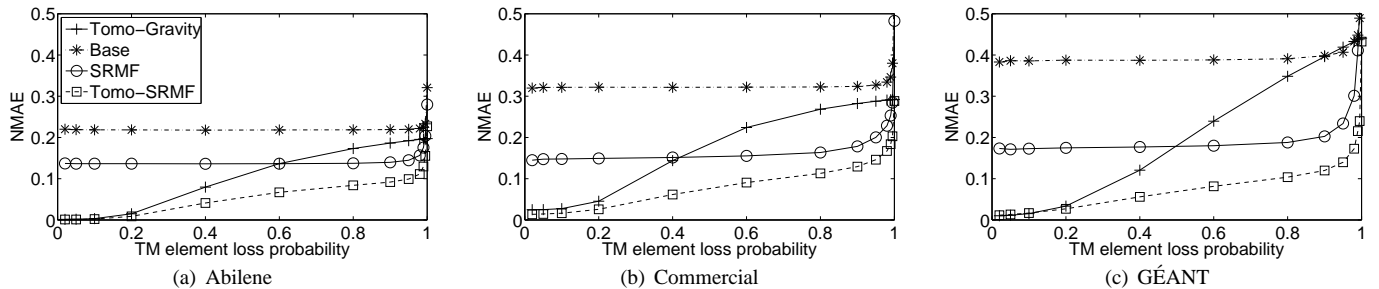


Fig. 6. Network tomography performance. Note, loss probability=1 corresponds to the traditional network tomography problem.

	Abilene	Commercial	GÉANT
Tomo-gravity	0.197 / 0.197	0.292 / 0.292	0.441 / 0.439
Base	0.321 / 0.233	0.566 / 0.380	1.198 / 0.489
SRMF	0.280 / 0.204	0.483 / 0.285	1.185 / 0.516
Tomo-SRMF	0.227 / 0.155	0.288 / 0.203	0.433 / 0.240

TABLE III

NETWORK TOMOGRAPHY PERFORMANCE: THE FIRST NUMBER IS THE PERFORMANCE WHERE WE HAVE NO DIRECT TM MEASUREMENTS, THE SECOND SHOWS WHERE WE MEASURE ONLY 0.5% OF THE ELEMENTS.

raphy problem. As this part of the figure is important, but relatively hard to read, we have duplicated key performance metrics in Table III.

First, note that the gravity model is so bad that its results lie off the scale. The baseline technique is the second worst in most cases, but is still much better than the gravity model. Second, SRMF performs poorly at the pure network tomography task where no direct measurements are available. However, if even a few (as few as 0.5%) of the TM elements are directly observed, then SRMF’s performance improves dramatically, whereas Tomo-gravity’s performance improves roughly linearly with respect to the increase in information. Finally, by combining SRMF and Tomo-gravity, Tomo-SRMF gets the best of both worlds and significantly outperforms each individual method by itself. Figure 6 and Table III show the improvements.

Note that Soule *et al.* [11] also propose to incorporate flow-level measurements in TM estimation. Compared with their “third generation” TM estimation methods, Tomo-SRMF has two key advantages: (i) it does not require any expensive calibration phase in which *entire* TMs are directly measured, and (ii) it is highly accurate and can reduce the error of Tomo-gravity by half with only 0.5-5% observed TM elements (whereas 10-20% directly observed TM elements are required according to [11]).

B. Prediction

In this section we consider the behavior of SRMF with respect to TM prediction. We do so by dividing our data into two segments, an initial training segment up to some time t , and then a test segment over which we try to predict the TM.

Prediction is rather different from the general problem of interpolation. Several techniques (SRSVD and NMF) just fail. KNN does not work well because there are no temporally “near” neighbors, and no spatial neighbors at all. However, if

we can use the temporal pattern in the data more creatively we can make progress. For instance, rather than using a simple nearest neighbors technique, we use seasonal nearest neighbors. TMs show strong diurnal patterns, and so it is not surprising that offsetting neighbors by the 24 hours period has benefits. In essence, the seasonal nearest neighbor approach assumes that today’s traffic has the same pattern as yesterdays.

Likewise for SRMF, we do not need to use the spatial constraint matrix, as an entire slab of the data is missing (the future data we are trying to predict). However, to allow a fair comparison with seasonal nearest neighbors, we also use seasonality in constructing our T matrix. We construct a difference matrix, but where the interval between the differences is 24 hours.

Figure 7 shows the results with respect to the proportion of data being predicted. Note that SRMF outperforms the other techniques, and further that SRMF’s performance degrades very slowly as the length of data being predicted increases (and the training data gets correspondingly smaller). This shows that typical TMs exhibit temporal regularity and SRMF can effectively take advantage of it.

C. Anomaly Detection

A common task in network operations is finding problems. There are specific tools for finding some problems (*e.g.*, SNMP is commonly used to find link failures), and other problems such as specific attacks can be characterized by a *signature*, which signals the attack. However, both of the above approaches rely on pre-knowledge of the problems that we will encounter. There is a complementary need to find unanticipated problems in networks.

Such problems cannot be characterized before-hand, and so the method commonly used to detect such *anomalies* is to find significant differences from historical observations. Most approaches involve some transformation of the data followed by outlier detection [5]. Common examples include simple filtering of the data, Fourier transform, wavelets, or PCA. The transform is aimed at separating the “normal” space of historical events, from the anomalous space. Techniques such as PCA do this explicitly, while others rely on commonly observed properties. For example, Fourier techniques rely on the normal data primarily inhabiting a low- to mid-frequency space, so that anomalies involve high-frequencies such as those incurred by a rapid change. Outlier detection can be performed by taking the normal model of the data, and comparing its




Fig. 7. Network prediction performance.

values at each time point with the real data, and then seeking points where the difference exceeds some threshold T .

In this section we will compare several approaches to anomaly detection. To keep things simple so that we can gain an intuitive understanding of the various properties of different approaches, we will consider only three algorithms: one temporal, one spatial, and our spatio-temporal approach. The three approaches we use are

1. **Differencing:** Differencing is a standard time-series technique to remove linear trends (typical traffic data are non-stationary, and over periods of minutes to hours can often be reasonably approximated as having a linear trend). Differencing also highlights sudden changes, such as we would see in a traffic *spike* or a level shift [5]. Implicitly, differencing is using the data from the previous time step as a model for the current time, and so it has not received a great deal of consideration in the networking literature, but it provides a simple temporal benchmark against which we can gain some intuition. We can write the differencing operator as post-multiplication of X by $T = \text{Toeplitz}(0, 1, -1)$, a purely temporal operation that makes no use of spatial correlations between TM elements.
2. **PCA/SVD:** PCA/SVD has received much interest for network anomaly detection in recent years [4], [5], [17]–[19], [39], and is the most common spatial technique for anomaly detection. As noted earlier, PCA/SVD is applied by choosing the rank r of the normal subspace (based on the power contained in the first r singular values), and projecting the input data X into the abnormal subspace, where artifacts are then tested for size. Implicitly, we are looking at the difference between the normal model of the data created by the low-rank SVD approximation and the data itself. Intuitively, the process builds a (simple) model from the historical relationships between TM elements. New time points are compared to see if they satisfy this relationship. If not, they are declared to be anomalies. It is a purely spatial technique, since reordering of the data in time (the columns of X) has no effect on the results. Interestingly, compressive sensing ideas have already appeared in the context of PCA based anomaly detection [39], though in that context the goal was to reduce the volume of data transmitted to a NOC, and the missing data could be controlled, whereas in our context the missing data are out of our control.
3. **SRMF:** SRMF is used for interpolation of missing data

throughout most of the paper, but it can equally be thought of as an approximation algorithm. Here we apply SRMF directly to the traffic data *including* the anomalies, much as one would with SVD, to find a model against which we compare the actual traffic. Large differences indicate anomalies. Our technique, however, is truly spatio-temporal as the model that we create involves both the spatial and temporal properties of the underlying TM. The low-rank approximation is then used as a model for the normal data, and the differences between this and the real traffic are used to signal anomalies. Once again, we use the standard method of thresholding these differences to detect outliers.

We will compare each of these algorithms using simulations. Ringberg *et al.* [19] explain in detail why simulation should be used for accurate comparisons of anomaly detection techniques. In brief their reasons are: (i) accurate and complete ground truth information is needed to form both false-alarm and detection probability estimates (both are needed for comparisons, as one by itself can be entirely misleading); (ii) many more results are needed (than one can obtain from data) to form accurate estimates of probabilities, and (iii) simulation allows one to vary parameters (say the anomaly size) to study their effects. Simulation is necessary, but not sufficient for validation, so we expect that further work is needed on this type of anomaly detection before it is used by network operators.

Our approach to simulation is intended to highlight the features of the different techniques. We make no claim that the simulation is completely realistic, only that it clearly illustrates the properties of the different anomaly detection techniques. We simulate in two steps: we first create the normal traffic, and then inject anomalies. We create the TM by an orthogonal composition of a synthetic gravity model TM [34] in the spatial domain, and a periodic Norros model [40], [41] in the temporal domain. Both models have arguments in their favor but principally we need to create a TM with low rank, but some local spatio-temporal structure that we might find in a real TM.

We use this model to generate 100 instances of the TM X consisting of one weeks worth of measurements at 15-minute intervals. In each instance we inject one anomaly. The anomaly is a spike added to the TM at a single randomly chosen time point, so that one anomaly cannot interfere with the detection of another. The value of the spike is a vector




Fig. 8. Detection probability (for fixed false alarm probability of 10^{-5}).

of Gaussian random variables (in each element of the TM) but we normalize the total size of the spike (measured by its L_2 norm) to be a fixed size, which we vary from 0.1 to 100. Spikes of size 0.1 (in our dataset) are almost indistinguishable (by the naked eye) from the standard random variations in the traffic. Spikes of size 100 are much larger than the typical TM elements, and so are easily observed. We then apply each of the three techniques above to create a “normal” traffic vector, and detect anomalies by thresholding on differences between the normal and measured vector. Note that we do not have missing data in the inputs (it is not obvious how to fairly compare the three algorithms when there are missing data, given the better interpolation performance of SRMF). However the anomalies *are* included in the inputs, so that both SVD and SRMF can be compared fairly.

An important detail is the choice of thresholds for outlier detection. Non-stationarity in our data makes setting thresholds more difficult than in some problems. For instance, in PCA/SVD the anomaly pollutes the data used to create the normal subspace, and so invalidates the standard statistical assumptions used to determine thresholds [19]. So it is hard to theoretically derive thresholds that produce the same false-alarm probabilities for large and small anomalies. Instead, we choose the threshold empirically to fix the false alarm probability. We perform a number of simulations, and choose a threshold such that the false alarm probability falls at a suitably low value 10^{-5} (in practice even smaller values may be desirable, however, estimation of very small probabilities becomes problematic). We then plot the detection probability and compare these.

Figure 8 shows the detection probabilities for the three techniques. First, let us compare PCA/SVD and Differencing for small anomalies. When anomalies are small (in value) and hence the differences they create will be small, the probability of detection of these differences will be small. Note that, even though we allow a larger false-alarm probability as the anomalies get smaller, the detection probability for Differencing drops dramatically. On the other hand, PCA/SVD exploits the spatial relationships between the data elements. It uses all the information available at a particular time point instead of processing the information from each time sequence independently. While the performance of PCA/SVD also declines for smaller anomalies, it is much better than

Differencing.

For moderate sized anomalies, both techniques have perfect detection records. These anomalies are not particularly hard to detect whichever method one uses. On the other hand, when we consider large anomalies, a different story emerges. The larger anomalies pollute the data used to create the normal subspace in PCA/SVD, and this has a tangible effect in that the detection probability drops dramatically. Note that although the false-alarm probabilities of the Differencing and PCA/SVD methods are slightly different, the detection probability of PCA/SVD drops almost to zero, whereas Differencing maintains its perfect detection record. So we can discount the minor difference in false-alarm probability as causing this drop. The result is consistent with those observed in data [19].

A solution might be to temporally pre-filter the data to remove large anomalies, and then apply PCA/SVD. This introduces the problem of interpolating data, which we have noted before. A preferable approach would be to use an approach that deals well with both ends of the spectrum. SRMF provides such a method. Figure 8 shows its performance. We can see that its detection probability is almost indistinguishable from the better of the two other methods for each anomaly size. So SRMF provides a method that deals well with the complete range of anomalies.

Another alternative comes from Brauckhoff *et al.* [42] where a standard extension of PCA was used to improve the performance of PCA for traffic anomaly detection. The approach copes with correlations between measurement intervals, along with the standard spatial approximation used in PCA. In that sense they also propose a spatio-temporal solution to the problem, however, there is no attempt to deal with the pollution of subspace estimates from large anomalies.

SRMF does this through its use of a spatio-temporal model. In the case where the spatial model is more effective (small anomalies) this is the part of the algorithm that is “biting”, whereas when the anomalies are large, they are being detected in the temporal domain, essentially by differencing. What we see here is that by imposing temporal smoothness constraint on LR^T , the effect of contamination is much smaller. Intuitively, if too much energy leaks into the normal subspace (as in PCA), then the projection of X into the normal subspace is no longer smooth, which would then result in a too big penalty in the smoothness term. Thus the smoothness term helps to limit the damage of contamination — the problem seen in PCA/SVD.

Moreover, the technique of [42] cannot be easily extended to deal with missing data. Although we use SRMF here purely to derive an approximate model for the typical traffic, it can be used to both interpolate and approximate, and hence to detect anomalies where data is missing, something that most anomaly detection algorithms cannot do without a preprocessing step of interpolation.

Note that we do not argue that with the naive choices of temporal operator T that we use here that SRMF is the best prediction or anomaly technique for TMs. Given the wealth of methods available for these applications (*e.g.*, see [5]), one can undoubtedly do better by more careful choice of T . However, there is a lesson to be learned here. First, our regularization approach can be generalized to apply to any linear

prediction/anomaly detection technique through appropriate choice of T . In each case we would hope for performance improvements as well, but the more important aspect of this work comes from the features we have demonstrated above: (i) our approach naturally deals with missing data, (ii) it can flexibly include additional data sources (*e.g.*, link data), and (iii) anomaly detection (and exclusion) are an inherent part of the algorithm. We argue that these are ideal features for any set of algorithms based on TMs.

VI. CONCLUSIONS AND FUTURE WORK

By drawing on recent developments in compressive sensing and relying on readily available domain knowledge in the area of TMs, we present in this paper a unified approach to measurement and analysis of TMs. We achieve this by applying our recently developed spatio-temporal compressive sensing framework that exploits the presence of both global structure (*e.g.*, low rank) and local structure (*e.g.*, spatio-temporal properties) in real-world TMs. Whether applied to TM estimation (*i.e.*, tomography), TM prediction, or anomaly detection, our algorithms consistently outperform other commonly-used methods and do so across the whole range of missing values scenarios, from purely random to highly structured scenarios where whole columns and/or rows of a TM are missing, and from very low levels of missing values to very high levels (*e.g.*, 90% and more). The main reason for the superior performance of our proposed technique when compared to its most widely-used competitors is its reliance on truly spatio-temporal models of TMs that capture much of the localized structure inherent in actual TMs. The reported results confirm the effectiveness and robustness of our approach to many real-world measurement issues and highlight the practical relevance of our algorithms for dealing with important network engineering tasks.

There are a number of avenues for future work. From an algorithmic perspective, it is worthwhile to note that many of the techniques described here (including SRMF) naturally extend to tensors (*i.e.*, multi-dimensional arrays), so that the original (unvectorized) TMs can be analyzed directly, *i.e.*, as true 3-d objects with traffic source, traffic destination, and time as the three axis. Such a tensor treatment of TMs has great potential and presents an opportunity to build more sophisticated spatio-temporal descriptions of the TM, including a more detailed modeling of the spatial and temporal constraint matrices S and T .

From an applied view, our work to date has focused exclusively on intra-domain TMs. However, applications of the proposed techniques to other types of TMs (*e.g.*, arising in the context of data centers [43]) and, more generally, to the more diverse set of network matrices loom as an intriguing open problem. In particular, inter-domain TMs that describe the volume of traffic exchanged between pairs of ASes in the Internet per time period are of great interest to researchers and network operators alike, but essentially nothing is known about them (for an initial attempt, see [44]). Our results provide a promising new direction for estimating such matrices. In the context of inter-domain TMs, the main challenge is that

there will be no measurements for the majority of the TM elements, mainly because most ASes view this as sensitive information and are generally not willing to share it with third parties. However, in view of our results, just a few measurements of a few TM elements (say from a small number of contributing ASes), maybe combined with some inter-AS link measurements or some other linear combinations of available TM data could provide enough information for an algorithm such as SRMF to obtain an approximate inter-domain TM. An important note is that not all of this hard-to-come-by information would be needed from the same time interval, and that even different measurements intervals might be accommodated through interpolation.

Acknowledgments

We thank the many people who have provided feedback on this work for their valuable input. This work was supported in part by NSF grants CNS-0546720, CNS-0615104, and CNS-0627020, and ARC grants DP0665427 and DP110103505. We would also like to thank IPAM at UCLA for providing a forum for many valuable discussions related to this work, and the Abilene and GÉANT networks for providing data.

REFERENCES

- [1] B. Fortz and M. Thorup, "Optimizing OSPF/IS-IS weights in a changing world," *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 4, pp. 756–767, 2002.
- [2] M. Roughan, M. Thorup, and Y. Zhang, "Traffic engineering with estimated traffic matrices," in *ACM SIGCOMM Internet Measurement Conference (IMC)*, Miami Beach, FL, USA, 2003, pp. 248–258.
- [3] R. S. Cahn, *Wide Area Network Design*. Morgan Kaufman, 1998.
- [4] A. Lakhina, K. Papagiannaki, M. Crovella, C. Diot, E. D. Kolaczyk, and N. Taft, "Structural analysis of network traffic flows," in *ACM SIGMETRICS / Performance*, 2004.
- [5] Y. Zhang, Z. Ge, M. Roughan, and A. Greenberg, "Network anomography," in *Proceedings of the Internet Measurement Conference (IMC '05)*, Berkeley, CA, USA, October 2005.
- [6] D. Alderson, H. Chang, M. Roughan, S. Uhlig, and W. Willinger, "The many facets of Internet topology and traffic," *Networks and Heterogeneous Media*, vol. 1, no. 4, pp. 569–600, December 2006.
- [7] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, J. Rexford, and F. True, "Deriving traffic demands for operational IP networks: Methodology and experience," *IEEE/ACM Transactions on Networking*, pp. 265–279, June 2001.
- [8] G. Varghese and C. Estan, "The measurement manifesto," in *2nd Workshop on Hot Topics in Networks (HotNets-II)*, 2003.
- [9] Y. Vardi, "Network tomography," *Journal of the American Statistical Association*, March 1996.
- [10] A. Medina, N. Taft, K. Salamatian, S. Bhattacharyya, and C. Diot, "Traffic matrix estimation: Existing techniques and new directions," in *ACM SIGCOMM*, Pittsburgh, USA, August 2002.
- [11] A. Soule, A. Lakhina, N. Taft, K. Papagiannaki, K. Salamatian, A. Nucci, M. Crovella, and C. Diot, "Traffic matrices: Balancing measurements, inference and modeling," in *ACM SIGMETRICS*, 2005, pp. 362–373.
- [12] Y. Zhang, M. Roughan, N. Duffield, and A. Greenberg, "Fast accurate computation of large-scale IP traffic matrices from link loads," in *ACM SIGMETRICS*, San Diego, California, June 2003, pp. 206–217.
- [13] Y. Zhang, M. Roughan, C. Lund, and D. Donoho, "An information-theoretic approach to traffic matrix estimation," in *ACM SIGCOMM*, Karlsruhe, Germany, August 2003, pp. 301–312.
- [14] —, "Estimating point-to-point and point-to-multipoint traffic matrices: An information-theoretic approach," *IEEE/ACM Transactions on Networking*, vol. 13, no. 5, pp. 947–960, October 2005.
- [15] Q. Zhao, Z. Ge, J. Wang, and J. Xu, "Robust traffic matrix estimation with imperfect information: making use of multiple data sources," *SIGMETRICS Perform. Eval. Rev.*, vol. 34, no. 1, pp. 133–144, 2006.

- [16] V. Erramilli, M. Crovella, and N. Taft, "An independent-connection model for traffic matrices," in *ACM SIGCOMM Internet Measurement Conference (IMC06)*. New York, NY, USA: ACM, 2006, pp. 251–256.
- [17] A. Lakhina, M. Crovella, and C. Diot, "Diagnosing network-wide traffic anomalies," in *ACM SIGCOMM*, 2004.
- [18] K. Xu, J. Chandrashekar, and Z.-L. Zhang, "A first step toward understanding inter-domain routing dynamics," in *MineNet '05: Proceedings of the 2005 ACM SIGCOMM workshop on Mining network data*, 2005, pp. 207–212.
- [19] H. Ringberg, A. Soule, J. Rexford, and C. Diot, "Sensitivity of PCA for traffic anomaly detection," in *ACM SIGMETRICS*, San Diego, CA, June 2007.
- [20] E. Candes and T. Tao, "Near optimal signal recovery from random projections: Universal encoding strategies?" *IEEE Trans. on Information Theory*, vol. 52, no. 12, pp. 5406–5425, 2006.
- [21] D. Donoho, "Compressed sensing," *IEEE Trans. on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [22] E. Candes and B. Recht, "Exact matrix completion via convex optimization," *Foundations of Computational Mathematics*, vol. 9, pp. 717–772, 2009.
- [23] B. Recht, M. Fazel, and P. A. Parrilo, "Guaranteed minimum rank solutions to linear matrix equations via nuclear norm minimization," *SIAM Review*, vol. 52, no. 3, pp. 471–501, 2010.
- [24] B. Recht, W. Xu, and B. Hassibi, "Necessary and sufficient conditions for success of the nuclear norm heuristic for rank minimization," in *47th IEEE Conference on Decision and Control*, December 2008.
- [25] Y. Zhang, M. Roughan, W. Willinger, and L. Qui, "Spatio-temporal compressive sensing and Internet traffic matrices," in *ACM Sigcomm*, Barcelona, Spain, August 2009, pp. 267–278.
- [26] P. Barford, J. Kline, D. Plonka, and A. Ron, "A signal analysis of network traffic anomalies," in *ACM SIGCOMM Internet Measurement Workshop*, Marseilles, France, November 2002.
- [27] E.J.Candes and M.B.Wakin, "An introduction to compressive sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 21–30, 2008.
- [28] M. Roughan, "A case-study of the accuracy of SNMP measurements," *Journal of Electrical and Computer Engineering*, vol. 2010, 2010, article ID 812979, doi:10.1155/2010/812979. <http://www.hindawi.com/journals/jece/2010/812979.html>.
- [29] V. Bharti, P. Kankar, L. Setia, G. Gürsun, A. Lakhina, and M. Crovella, "Inferring invisible traffic," in *Proceedings of the 6th International Conference*, ser. Co-NEXT '10. New York, NY, USA: ACM, 2010, pp. 22:1–22:12. [Online]. Available: <http://doi.acm.org/10.1145/1921168.1921197>
- [30] R. Bell, Y. Koren, and C. Volinsky, "Chasing the \$1,000,000: How we won the Netflix progress prize." *Statistical Computing and Graphics*, vol. 18, no. 2, December 2007.
- [31] Y. Mao and L. K. Saul, "Modeling distances in large-scale networks by matrix factorization," in *ACM SIGCOMM Internet Measurement Conference (IMC04)*, 2004, pp. 278–287.
- [32] D. D. Lee and H. S. Seung, "Algorithms for non-negative matrix factorization," in *Proc. Neural Information Processing Systems (NIPS)*, 2000, pp. 556–562.
- [33] B. Eriksson, P. Barford, R. Bowden, M. Roughan, N. Duffield, and J. Sommers, "BasisDetect : A model-based network event detection framework," in *ACM SIGCOMM Internet Measurement Conference*, Melbourne, Australia, 2010.
- [34] M. Roughan, "Simplifying the synthesis of Internet traffic matrices," *ACM SIGCOMM Computer Communications Review*, vol. 35, no. 5, pp. 93–96, October 2005.
- [35] F. R. K. Chung., *Spectral Graph Theory*. CBMS Lecture Notes, AMS Publications., 1996.
- [36] Abilene/Internet2, <http://www.internet2.edu/>.
- [37] S. Uhlig, B. Quoitin, S. Balon, and J. Leprepre, "Providing public intradomain traffic matrices to the research community," *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 1, pp. 83–86, January 2006.
- [38] T. Oetiker, "RRDtool," <http://oss.oetiker.ch/rrdtool/>.
- [39] L. Huang, X. Nguyen, M. Garofalakis, J. Hellerstein, M. Jordan, M. Joseph, and N. Taft., "Communication-efficient online detection of network-wide anomalies," in *Proc. of IEEE INFOCOM*, 2007.
- [40] I. Norros, "A storage model with self-similar input," *Queueing Systems*, vol. 16, pp. 387–396, 1994.
- [41] M. Roughan and J. Gottlieb, "Large-scale measurement and modeling of backbone Internet traffic," in *SPIE ITCOM*, Boston, MA, USA, 2002.
- [42] D. Brauckhoff, K. Salamatian, and M. May, "Applying PCA for traffic anomaly detection: Problems and solutions," in *IEEE Infocom*, 2009.
- [43] S. Kandula, S. Sengupta, A. Greenberg, and P. Patel, "The nature of datacenter traffic: Measurements and analysis," in *ACM Sigcomm Internet Measurement Conference*, 2009.
- [44] H. Chang, S. Jamin, Z. Mao, and W. Willinger, "An empirical approach to modeling inter-AS traffic matrices," in *Proc. 2005 ACM SIGCOMM Internet Measurement Conference (IMC 2005)*, 2005, pp. 139–152.

PLACE
PHOTO
HERE

D r Matthew Roughan joined the School of Mathematical Sciences at the University of Adelaide in 2004. Prior to that he worked at AT&T in the United States. His research interests lie in measurement and modelling of the Internet, and his background is in stochastic modelling with his PhD being in Applied Probability from the University of Adelaide, awarded in 1994.

PLACE
PHOTO
HERE

D r Yin Zhang is an Associate Professor at The University of Texas at Austin. His research interests lie in computer networks in general, with particular interests in network management, measurement, and security. He received MS and PhD degrees in Computer Science from Cornell University in 1999 and 2001, respectively. Before joining UT in 2005, he spent four years as a researcher at AT&T Labs – Research. He received NSF CAREER Award in 2006 and is an IEEE Senior Member.

PLACE
PHOTO
HERE

D r Walter Willinger received the Diplom (Dipl. Math.) from the ETH Zurich, Switzerland, and the M.S. and Ph.D. degrees from the School of ORIE, Cornell University, Ithaca, NY. He is currently a member of the Information and Software Systems Research Center at AT&T Labs-Research, Florham Park, NJ. Before that, he was a Member of Technical Staff at Bellcore Applied Research (1986-1996). He is a Fellow of ACM, IEEE, SIAM, and AT&T, and for his work on the self-similar ("fractal") nature of Internet traffic, he received the 1994 W.R. Bennett Prize Paper Award from the IEEE Communications Society, the 1996 IEEE W.R.G. Baker Prize Award from the IEEE Board of Directors, and the 2005 "Test-of-Time" Paper Award from ACM SIGCOMM.

PLACE
PHOTO
HERE

D r Lili Qiu is an Associate Professor at the University of Texas at Austin. She received MS and PhD degrees in Computer Science from Cornell University in 1999 and 2001, respectively. Before joining UT in 2005, she spent four years as a researcher at Microsoft Research working on Internet and wireless networking research. She received NSF career award in 2006. She is an IEEE senior member.