

Spatiotemporal Traffic Matrix Synthesis

Paul Tune and Matthew Roughan
ARC Centre of Excellence for Mathematical and Statistical Frontiers
School of Mathematical Sciences
University of Adelaide, Australia
{paul.tune,matthew.roughan}@adelaide.edu.au

ABSTRACT

Traffic matrices describe the volume of traffic between a set of sources and destinations within a network. These matrices are used in a variety of tasks in network planning and traffic engineering, such as the design of network topologies. Traffic matrices naturally possess complex spatiotemporal characteristics, but their proprietary nature means that little data about them is available publicly, and this situation is unlikely to change.

Our goal is to develop techniques to synthesize traffic matrices for researchers who wish to test new network applications or protocols. The paucity of available data, and the desire to build a general framework for synthesis that could work in various settings requires a new look at this problem. We show how the principle of maximum entropy can be used to generate a wide variety of traffic matrices constrained by the needs of a particular task, and the available information, but otherwise avoiding hidden assumptions about the data. We demonstrate how the framework encompasses existing models and measurements, and we apply it in a simple case study to illustrate the value.

Categories and Subject Descriptors

C.2.5 [Computer Communications]: Local and Wide Area Networks—*Internet*; C.4 [Performance of Systems]: Modeling Techniques

Keywords

Maximum entropy; network design; spatiotemporal modeling; traffic engineering; traffic matrix synthesis

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGCOMM '15, August 17 - 21, 2015, London, United Kingdom

© 2015 ACM. ISBN 978-1-4503-3542-3/15/08...\$15.00

DOI: <http://dx.doi.org/10.1145/2785956.2787471>

1. INTRODUCTION

A Traffic Matrix (TM) describes the traffic volume between a set of sources and destinations in a network. TMs are used as inputs in a variety of network planning and traffic engineering tasks. For instance, new power aware routing and traffic engineering algorithms [9] need TMs for testing and validation. TMs are also required for network traffic forecasting and planning [32] and topology generation [5]. The lack of public TM data remains an obstacle to researchers in these areas.

The main reason for this lack is the propriety nature of TMs: network operators are reluctant to release such data as they fear it may lead to the erosion of their competitive advantage. Consequently, researchers resort to testing network designs on limited public data, such as from the Abilene [23] or GÉANT [41] networks. These, however, are research and education networks, leading to the possibility that they are not representative of commercial networks. Even the release of a single commercial operator's TM data would not quell the question of representativeness – so we really need many such datasets. Hence, the current deficit is unlikely to be remedied in the foreseeable future.

Researchers also need multiple TMs to test ideas statistically and in order to have ground truth when testing inference techniques. To do this, large ensembles of TMs must be generated, so that confidence intervals on the performance of these techniques be meaningful.

We can solve these issues by synthesizing TMs, but how can we create a good model for synthesis without first having data? The answer is to abandon the typical network modeling approach – measure then model – and think about the problem from first principles.

Additionally, we are not aiming here at the typical target of past TM research: better inference of TMs from limited data [14, 22, 35, 45, 48]. When we synthesize a TM, we have different requirements:

- *Control*: A particular network protocol we aim to test will make assumptions. We should be able to generate TMs that match those assumptions to test its performance. For completeness, we also need to generate TMs that violate the assumptions, to test what happens to the protocol in unfriendly situations.

The required assumptions and constraints on TMs will vary from setting to setting. However, in creating TMs that satisfy some set of assumptions, we should not be forced to introduce additional (often hidden) assumptions about the traffic. We want to include the desired constraints on the TMs but no others.

- *Efficiency*: We want to be able to generate potentially large numbers of TMs computationally efficiently to obtain clear statistical measures of performance.
- *Consistency*: We want to make apples to apples comparisons: two studies using the same model should reach the same conclusions (regardless of specific random values), so we need to be able to draw instances from a well-defined, consistent ensemble of TMs.
- *Simplicity*: A model should be “as simple as it can be – but not simpler” [7]. Simplicity has many virtues: it improves our intuitive understanding, reduces the complexity of parameter estimation, and prevents over fitting. Part of simplicity is *parsimony* of parameters, and part is having *meaningful* parameters.

Finally, we might like our TMs to be “realistic”, but that might vary tremendously from setting to setting: for instance it seems likely that Wide Area Network (WAN) TMs are nothing like data center TMs. So realism is context dependent, and our goal here is to create a framework that can encompass and be adapted easily to any such setting. Indeed, if some data is available to populate a model in a particular context, we would like our framework to be able to include that data.

Our approach is based on the *principle of maximum entropy (MaxEnt)* [15]. Information theory and its darling *entropy* have been applied to inferring TMs in the past with success [31, 45, 46], but there are some subtle differences between creating models for inference and synthesis. Most notably, we don’t necessarily have *any* data, though if we do, we would like to incorporate it.

The resulting MaxEnt framework encompasses and extends most of the existing models for TMs, *e.g.*, [31, 36]. The models it produces are conservative in the sense that they do not assume anything about the TMs beyond the given constraints. The framework also links the constraints and assumptions directly to models, so that we can learn from an existing model exactly what minimal set of assumptions led to it.

A list of our contributions are as follows:

- a systematic methodology for generating *controlled synthetic TMs with prescribed properties*, achieved via the principle of maximum entropy,
- *closed-forms* of example models that would be useful in testing various network algorithms (*e.g.*, routing protocols) and network designs,
- *fast generation* of ensembles of TMs with spatiotemporal properties, and
- a *unification* of previous TM models under the family of maximum entropy models.

We examine the utility of our approach through its use in topology generation [5]. Here, the topology is designed around a set of TMs as input. Our case study’s two chief findings are 1. although TMs have an impact on the topology design, the impact of some TM features is small, consistent with the findings of [5], and 2. optimizing a network around the peak TM is almost as good an approach as optimizing with respect to a whole-of-week pattern, supporting the common practice of basing designs around a peak traffic measure such as a *busy hour*.

Our code is at <https://github.com/ptuls/MaxEntTM>

2. BACKGROUND

2.1 The principle of maximum entropy

The principle of insufficient reason, often credited to Pierre Simon Laplace in the late 18th century, states that given an observation one should remain undecided about all the potential events explaining the observation. In other words, one assigns equal probability to each event. Laplace and Jacob Bernoulli considered the concept to be intuitively obvious, seeing little need to formalize it [37].

The principle of maximum entropy formally defines and generalizes the principle of indifference. Proposed by Jaynes [15] in the context of statistical mechanics the principle advocates that, given data, we should choose the distribution providing a plausible explanation of the observations while making the fewest assumptions. Jaynes showed that several models in statistical mechanics arise as a natural consequence.

The principle is defined in terms of Shannon *information entropy* [10]. If X is a discrete random variable taking values in set \mathcal{X} with probability mass function $p(x)$, then Shannon’s entropy is defined as [10]

$$H(X) = - \sum_{x \in \mathcal{X}} p(x) \log p(x), \quad (1)$$

with the convention $0 \log 0 = 0$. Entropy measures the average uncertainty of a random source, for instance, if X is completely deterministic, $H(X) = 0$. Shannon’s entropy is used because it is the only measure that nicely separates independent components in a joint distribution, and has a natural axiomatic derivation as a measure of uncertainty [34].

The principle of indifference now arises as a special case: when the only constraint is that $p(x)$ be a probability distribution, the entropy is maximized by the uniform distribution. Thus, maximizing entropy is consistent with the intuitive notion of being least committed to any one outcome, *i.e.*, indifference.

If additional observations exist, we describe them as constraints on the data. MaxEnt states that the best strategy is to choose a distribution that maximizes entropy subject to these constraints. The solution is one

that conforms to known observations about the data and no more than that.

Jaynes [15] used the Shannon entropy to construct the MaxEnt framework. Since we consider continuous Random Variables (RVs) as an approximation of the traffic volumes, we use the *differential entropy* [10],

$$h(f) = - \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} f(\mathbf{x}) \log f(\mathbf{x}) d\mathbf{x}, \quad (2)$$

of a random TM process \mathbf{x} with distribution $f(\mathbf{x})$.

Suppose by experimental observation, there are $L+1$ constraints on the data $\mathbb{E}[\phi_\ell(\mathbf{x})] = b_\ell$ (where $\phi_\ell(\cdot)$ is a convex function and b_ℓ is a scalar value setting the constraint) are known. The constraint $\mathbb{E}[\phi_0(\mathbf{x})] = 1$ is the normalization constraint (that applies to all probability distributions). The MaxEnt optimization problem is

$$\max_{f(\mathbf{x})} h(f), \text{ s.t. } \{\mathbb{E}[\phi_\ell(\mathbf{x})] = b_\ell\}_{\ell=0}^L, \text{ and } \mathbf{x} \succeq \mathbf{0}, \quad (3)$$

where $\mathbf{x} \succeq \mathbf{0}$ denotes non-negativity, a constraint we enforce on all TM models.

In the cases we primarily consider, the problem is convex so any local maximum is the global maximum, the solution being the *maximum entropy model*. Non-convex functions are possible, but more effort is required to search for the globally maximum model.

Boltzmann's theorem [15] states that the optimal solution takes the form of distributions from the exponential family, generically given by

$$f(\mathbf{x}) = \frac{1}{Z} \exp \left(- \sum_{\ell=1}^L \lambda_\ell \phi_\ell(\mathbf{x}) \right), \quad (4)$$

where the $\{\lambda_\ell\}_{\ell=1}^L$ are the Lagrange multipliers obtained via the Calculus of Variations applied to (3), and Z is the normalization factor which incorporates the normalization constraint.

2.2 Entropy and TM models

A general introduction to modeling TMs can be found in [38]. We examine some of these models in detail here so that we can discuss how they relate to our framework.

Zhang *et al.* [45, 46] used a technique referred to as maximizing *relative entropy* to develop TM inference methods. However, this is different from maximum entropy modeling. Zhang *et al.*'s approach was aimed at finding a particular TM that matched data and was close to a prior model (in their case the generalized gravity model). Here we aim to create a model, not estimate a particular TM. The output of our approach is a distribution, not a TM. The part of that work closer to our own is in their prior model: we shall discuss gravity models and their ilk at length here.

Roughan [31] proposed using a random gravity model as a spatial model for TM synthesis, and Oikonomou [25] noted that it is a MaxEnt model for traffic under a certain set of assumptions. In fact, if one examines the

transportation literature, that insight is even older [28]. This is a single case of the framework presented here.

The work listed above considered spatial TMs only. The only work of which we are aware specifically on spatiotemporal synthesis was that of Nucci *et al.* [24], who developed static, dynamic stationary and dynamic cyclo-stationary models of TMs. Their spatial model was simple: a log-normally distributed Independent and Identically Distributed (IID) set of TM elements were generated. However, their models' raw outputs might not conform to pre-specified constraints, such as link capacities, necessitating adjustments to make them admissible TMs. We show here how this IID model fits into the MaxEnt framework, and how the framework allows us to avoid this *ad hoc* mapping step.

Most other works proposing spatiotemporal TM models do so with the aim of improving inference accuracy, rather than addressing the synthesis problem (though occasionally, passing reference has been made to potential alternative uses of their models). Soule *et al.* [36] proposed two models of Origin-Destination (OD) flows, one stationary and the other cyclo-stationary. Roughan *et al.* [33] developed a simple purely temporal model of OD flows traversing backbone routers comprising of a growth trend component, a seasonal component and a component for sudden spikes in traffic due to anomalies. Zhang *et al.* [48], exploiting new developments in signal processing, developed a model under the assumption of spatiotemporal low rankedness, enabling accurate TM inference via convex programming [8, 30].

Another work addressing the synthesis problem is [40], but its focus is entirely different. Our conceptual model in that paper was to help a network operator understand the effect of errors in predictions of TMs. In that context, the idea was to generate an ensemble of TMs centered around the prediction, whereas here, we do not require prior data about the TM, only constraints on its general properties. Interestingly, as the errors become large in that model, the results approach a maximum entropy distribution, which is exactly what we propose here: when little is known about the TM, MaxEnt provides a natural path towards modeling.

Notation: The column vector of N ones is denoted by $\mathbf{1}_N$. Some common statistical distributions are used here:

- An exponentially distributed RV X with mean $1/\lambda$ is denoted $X \sim \text{Exp}(\lambda)$.
- The continuous uniform distribution in the interval $[a, b]$ is denoted by $U(a, b)$.
- $X \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ means that X is distributed as a normal distribution with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$.
- $X \sim \text{TNorm}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ means that X is distributed according to a truncated normal distribution with support $\mathcal{I} := [0, \infty)^N$. Note that $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are the mean and covariance *after* truncation.

We define all other notation when needed.

3. SPATIOTEMPORAL TM SYNTHESIS

Network traffic information is contained in the TM $\mathbf{X}(t)$, which measures traffic between locations in a network. Traffic is typically measured in discrete time intervals, with 5 to 15 minutes. So a “traffic matrix” is actually a series of matrices:

$$\mathbf{X}(t_k), k = 0, 1, \dots, M - 1, \quad (5)$$

where $X_{i,j}(t_k)$ denotes the average volume of traffic entering node i and exiting node j in a measurement interval $[t_k, t_{k+1})$.

Let $\mathcal{X} := \{\mathbf{X}(t_k)\}_{k=0}^{M-1}$ denote the TM process. We are concerned here with the distribution governing the generation of \mathcal{X} , denoted by $f(\mathcal{X})$. In general, multivariate stochastic processes are all governed by such a distribution, which here define an ensemble of TMs. However, all successful models simplify this high dimensional construct in some (usually dramatic) way.

Although network traffic is discrete in nature (measured in bytes or packets), we model the TM entries as continuous RVs. This assumption is reasonable where we consider high volumes of traffic.

Spatial properties of the TMs refer to the statistical properties between the TM entries either at a fixed t , *i.e.*, a *snapshot* $\mathbf{X}(t_k)$, or over time, *e.g.*, the average TM $\frac{1}{M} \sum_{k=0}^{M-1} \mathbf{X}(t_k)$. We consider a network with N locations, so each snapshot is an $N \times N$ matrix.

Temporal properties refer to the statistical properties only varying with t , either with fixed spatial indices, *i.e.*, the process $X_{i,j}(t_k)$, or some summary such as the total traffic $S(t_k) = \sum_{i,j} X_{i,j}(t_k)$ at each time point t_k .

3.1 Maximum entropy models

In the introduction we discussed the qualities that make a good synthesis model. In this section our primary foci are simplicity and controllability.

Testing does not (necessarily) require realistic input, since the goal is to understand the effect a specific input has on an algorithm. For instance, inputs could be extreme cases, such as a TM with its total traffic spread uniformly over all entries (called a *uniform load model* [6, 38]). What we do need is the ability to control the input of an algorithm to perform systematic tests.

Chiefly, we need to incorporate what we believe about the network. Different types of networks have different properties and constraints. For instance, data centers have entirely different traffic mixes from WANs [1, 16] so there exists different constraints on their respective TMs. The constraints we primarily consider here are in the form of constraints on statistics of the TM: for instance, a constraint on the average or expected value $E[\cdot]$ of some aspect of the TM. Deterministic constraints can also be added, but are rather simpler to analyze, so we focus on the more difficult case.

Though we primarily focus on WANs, the general methodology provides tools for researchers to develop synthesis models for alternative networks. Part of the

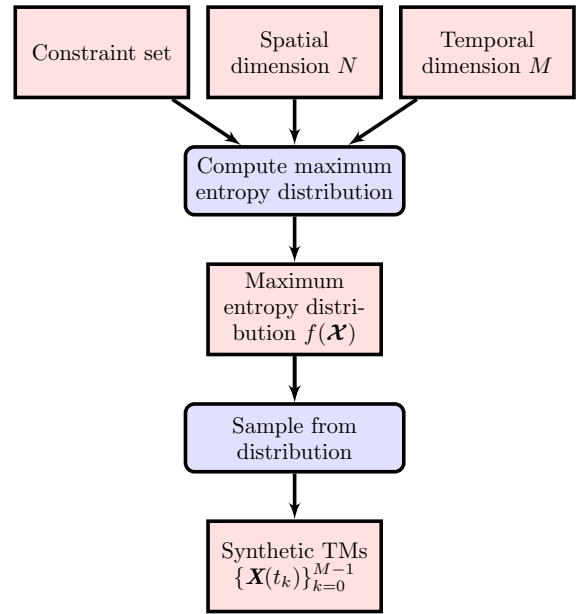


Figure 1: Methodology of the spatiotemporal synthesis of the maximum entropy TMs. The rectangles denote the inputs and outputs at each stage, while the rounded rectangles denote the stages of the process. The steps generate one run of M matrices.

goal here is to provide a modeling framework for researchers even *before* a new technology is deployed, to understand its properties under various hypotheses.

Figure 1 presents an overview of our methodology. The inputs are the set of constraints, and the spatial and temporal dimensions N and M respectively. Solving (3) gives the MaxEnt distribution $f(\mathcal{X})$, which entails computing the Lagrange multipliers $\{\lambda_\ell\}_{\ell=1}^L$. Generally, these can be computed using the Newton-Raphson method (see [43]), but it is more instructive where we can derive closed forms for the distributions, and so we spend some effort here to derive a number of such models in a manner that is explanatory and extensible.

Finally, $f(\mathcal{X})$ is sampled to generate the synthetic TM process \mathcal{X} . Again, this can be achieved numerically, say via Markov Chain Monte Carlo (MCMC) techniques, but we demonstrate several models for which sampling can be performed much more efficiently. Constraints that often arise in TM research lead to closed forms, which in turn provide fast sampling methods.

For all models, let $\mathbf{R}(t_k) = \mathbf{X}(t_k)\mathbf{1}_N$ and $\mathbf{C}(t_k) = \mathbf{X}(t_k)^T\mathbf{1}_N$, *i.e.*, the row and column sum processes (total ingress and egress traffic), and $S(t_k)$ denotes the total traffic in each interval t_k .

Model derivations involve (sometimes) lengthy, but standard derivations, and so in the interests of space we omit the details. However, there are some simple facts which should make the results somewhat intuitive:

1. Given the constraints of non-negativity and mean, the MaxEnt distribution is the exponential [27].

2. With both a mean and variance constraint, the MaxEnt distribution is the Normal distribution, and if we add to this the requirement of non-negativity, then the MaxEnt distribution is the truncated normal distribution [39].
3. Many other distributions arise from other moment-based or otherwise constraints on the RVs [27].
4. If we have two (or more) RVs, the MaxEnt model is that the RVs are independent (unless there is a joint constraint on the RVs) [10].

3.2 Purely spatial models

We divide our models into three classes: *purely spatial*, *purely temporal*, and *spatiotemporal models*. First generation TM models such as the gravity model [45] and the discrete choices model [22] are spatial, so we start there, noting that if we have no temporal information, the TM snapshots will be independent (Property 4), so we need only focus on a single interval and thus we drop t_k from our notation.

Independent entries model: In the absence of constraints on the joint distribution of TM entries $X_{i,j}$, by Property 4, these will be independent. If we also have homogenous constraints, *i.e.*, they don't favor any group of entries, then the entries will be IID.

There are many subcases of this type. For instance if the mean or expected value of each entry is fixed, *i.e.*, $\mathbb{E}[X_{i,j}] = \bar{X}$, then the MaxEnt model has entries $X_{i,j} \sim \text{Exp}(\bar{X}^{-1})$. If we further impose control on each entry's variance by setting the variance to σ^2 then the MaxEnt model has entries $X_{i,j} \sim \text{TNorm}(\bar{X}, \sigma^2)$. If instead we only set bounds $[D_L, D_U]$ on each TM entry, the MaxEnt model has IID entries $X_{i,j} \sim U(D_L, D_U)$.

These models are perhaps oversimplified, but they have been used! For instance, the Poisson model of Vardi [42] and the log-normal model of Nucci *et al.* [24].

In the latter case we can derive an interesting insight. The log-normal distribution is the result of constraints on the average and variance of the *log* of entries [27]. The natural conclusion is that TM entries might be better modeled in the log-domain.

However, the key takeaway message is not the log-normal distribution, but rather that MaxEnt leads to a link between the model and its hidden assumptions, and these can provide potential insights into the data and/or the validity of the model. Perhaps another model provides a slightly less tight fit to the data, but is simpler and more intuitive when considered in this way, and therefore has more explanatory power?

Outer product models: IID models are clearly limited. Real TMs have correlations. Two attempts to build such correlations are the Random Gravity Model (RGM) [31] and the mapping phase of [24].

Let us consider here a simple yet powerful method of creating structure, which is to impose constraints on the row and column sums of the matrix, an idea that

recurs through much of the TM literature¹ (*e.g.*, [14, 22, 31, 35, 45, 46]), and even into the realms of oblivious design where TMs are assumed to be unknowable [49] except for these constraints.

If we impose row and column constraints (only), then the MaxEnt model is formed via an outer product: for instance given non-negativity and the constraints

$$\mathbb{E}[\mathbf{R}] = \frac{T}{N} \mathbf{1}_N, \quad \mathbb{E}[\mathbf{C}] = \frac{T}{N} \mathbf{1}_N, \quad \mathbb{E}[S] = T,$$

the MaxEnt model is given by the *outer product*

$$\mathbf{X} = \frac{T}{N^2} \mathbf{U} \mathbf{V}^T, \quad (6)$$

where $U_i, V_j \sim \text{Exp}(1)$. This case is in fact the RGM of [31], and that illustrates how we can find new information about minimal assumptions of a model. The RGM was derived from data, but now we can understand the hidden assumptions of the model.

More generally, given constraints $\mathbb{E}[\mathbf{R}] = \mathbf{r}$ and $\mathbb{E}[\mathbf{C}] = \mathbf{c}$, MaxEnt implies a model that has *outer-product* form

$$\mathbb{E}[\mathbf{X}] = \frac{1}{T} \mathbf{r} \mathbf{c}^T, \quad (7)$$

and from this and the simple distributional arguments above we can derive a set of *rank-1* matrix models.

Much more can be said about outer-product models. There are many generalizations, many appearing in older works in the transportation literature, for instance MaxEnt models that include a distance *friction* term that reduces traffic between distant locations. Such terms have not had a clear justification in most Internet studies, but can be included via a distance-cost constraint within the framework if needed [28].

One point, however, which we must make again is that this type of approach applies to any matrix for which we know row and column properties, and we will reuse this idea in a moment.

3.3 Purely temporal models

An example of a purely temporal traffic model is that proposed in [33], which focused on the long-term trends and cyclical patterns in traffic. There are countless other examples of such models, applied either to single TM entries, link traffic or traffic totals. Our goal here is not to describe the panoply of possible models, but to present how these models fit within the MaxEnt framework.

In the past, temporal models have chiefly been used for traffic prediction². Predictions generally estimate a mean, and sometimes a variance around that mean.

¹Row and column sums are really the ingress and egress traffic volumes, hence their importance.

²It may appear that applications such as anomaly detection are otherwise, but a closer look shows that most such techniques revolve around predicting "typical" traffic, and detecting outliers from these predictions.

As before, if only the mean of a set of (non-negative) random variables is given, then the RVs will be independent and exponentially distributed (Property 1). If variance is also fixed, then the RVs will be independent truncated normals (Property 2). More sophisticated models are also possible, for instance, if the autocovariances are also known [27], but these are rarely reported in Internet studies so we won't discuss these models further except to note that they can be added to the MaxEnt framework.

Thus, to derive a MaxEnt temporal model we might choose a mean and variance as function of time and then simulate from the truncated normal distribution. The choice of mean and variance varies with the desired network properties. Commonly, they would not be constants but may include diurnal and weekly cycles.

It should be obvious that the above provides almost too much flexibility. We can accommodate almost any temporal model, so the question is how should it be restricted to a reasonable set of models in practice. A very simple case could arise from modeling cyclical behavior using Fourier analysis, which is designed specifically for that purpose. In this case we do so by choosing constraints on the Fourier coefficients of the traffic:

$$\mathcal{F}\left(\{\mathbb{E}[S(t_k) - T]\}_{k=0}^{M-1}\right) = \mathbf{a}. \quad (8)$$

The coefficients \mathbf{a} would be chosen to reflect the cyclical nature of the traffic. Then the MaxEnt model is an exponential model where the means at each time step are given by $\mathbb{E}[S(t_k)] = \bar{s}(t_k) = T + w(t_k)$, $\{w(t_k)\}_{k=0}^{M-1} = \mathcal{F}^{-1}(\mathbf{a})$ and \mathcal{F}^{-1} denotes the inverse discrete Fourier transform. Notice that now the model generates values that are only conditionally independent, not IID.

As in other cases, we could also control variances and then would arrive at a truncated normal model.

Once again, note that the goal of this paper is to show how knowledge can be exploited to construct a temporal model, not to suggest that the above is the only sensible approach. This approach, however, does have the advantage that both Fourier transforms and expectation are linear operators.

3.4 Spatiotemporal models

Now that we have both spatial and temporal models, we can combine ideas in those models to form spatiotemporal models as follows.

Time modulated traffic: In most situations, we perhaps know or can assume information about the spatial and/or the temporal behavior. Many studies have examined details of these two facets of network traffic [19, 20]. Many more will no doubt follow, but also remember that we wish to consider potential hypothetical future traffic sources, *e.g.*, traffic in hosting centers. It is often possible to reason about the temporal and spatial properties of such hypothetical cases.

Importantly though, we rarely have information correlating spatial to temporal behavior (though we will note an important exception resulting from timezones in the following model).

We proceed in this case using a common trick in TM modeling (*e.g.*, see [44, 48]). We *stack* the columns of a single TM snapshot and then construct a new matrix from the columns at each time interval, that is,

$$\mathbf{Y} = \begin{bmatrix} \mathbf{X}_{*,1}(t_0) & \mathbf{X}_{*,1}(t_1) & \cdots & \mathbf{X}_{*,1}(t_{M-1}) \\ \mathbf{X}_{*,2}(t_0) & \mathbf{X}_{*,2}(t_1) & \cdots & \mathbf{X}_{*,2}(t_{M-1}) \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{X}_{*,N}(t_0) & \mathbf{X}_{*,N}(t_1) & \cdots & \mathbf{X}_{*,N}(t_{M-1}) \end{bmatrix} \quad (9)$$

The new matrix \mathbf{Y} is also sometimes called the traffic matrix, but now it is a $N^2 \times M$ spatiotemporal matrix.

We are then in the same situation as in generating spatial models: we want to generate a matrix about which we only have row and column constraints. A natural building block towards more complex spatiotemporal models is to consider a model with separate independent spatial and temporal components. Not surprisingly, this has been studied by statisticians, and is called a *separable spatiotemporal model* [11].

In our context, for instance, column-sum constraints correspond to a model for the total traffic $S(t_k)$ at each time t_k , and row-sum constraints correspond to constraints on the average traffic per source/destination $\bar{X}_{i,j}$. The logical construction then is to use the same outer-product MaxEnt suggested for spatial models, *i.e.*,

$$\mathbb{E}[X_{i,j}(t_k)] = \bar{X}_{i,j} \bar{s}(t_k) / T, \quad (10)$$

where T is the time average of $\bar{s}(\cdot)$. When seen as a matrix operation, this is just an outer-product of the stacked vector version of $\bar{X}_{i,j}$ and $\bar{s}(t_k)$.

The MaxEnt model then takes the form $X_{i,j}(t_k) = G_{i,j} P(t_k)$, with $\mathbb{E}[G_{i,j}] = \bar{X}_{i,j}$ and $\mathbb{E}[P(t_k)] = \bar{s}(t_k) / T$, *i.e.*, separate spatial and temporal components. The exact form of $X_{i,j}(t_k)$ depends on the constraints, again with a constraint on means resulting in an exponential distribution, and constraints on mean and variance resulting in a truncated normal. The decomposable nature of the MaxEnt model allows mixing the two, using one for space and the other for time.

We can also include detailed spatial structure. For instance, we could have a model we refer to as the *Modulated Gravity Model (MGM)* where spatially we constrain the traffic by gravity-model-like constraints

$$\frac{1}{M} \sum_{k=0}^{M-1} \mathbb{E}[\mathbf{R}(t_k)] = \bar{\mathbf{r}}, \quad \frac{1}{M} \sum_{k=0}^{M-1} \mathbb{E}[\mathbf{C}(t_k)] = \bar{\mathbf{c}}, \quad (11)$$

and temporally we require (8). The MaxEnt model is

$$\mathbf{X}(t_k) = T \mathbf{U} \mathbf{V}^T P(t_k), \quad (12)$$

$$U_i \sim \text{Exp}\left(\frac{T}{\bar{r}_i}\right), V_j \sim \text{Exp}\left(\frac{T}{\bar{c}_j}\right), P(t_k) \sim \text{Exp}\left(\frac{T}{\bar{s}(t_k)}\right),$$

where as before $\bar{s}(t_k) := T + w(t_k)$. Once again, if we aim to control the variances, the relevant components U_i , V_j and/or $P(\cdot)$ would follow appropriate truncated normal distributions. The simplest choice is to set $\Sigma_{\mathbf{r}} = \Sigma_{\mathbf{c}} = \sigma^2 \mathbf{I}_N$. Likewise, we can also control $P(\cdot)$ by introducing a single variance parameter σ_p .

Note that these choices result in *simple* models, with minimal parameters, and meaningful parameters that can be directly related to network properties.

The model's name comes from the fact that its conditional mean is the RGM modulated by $\bar{s}(t_k)/T$, *i.e.*,

$$\mathbb{E}[\mathbf{X}(t_k)] = \frac{\bar{\mathbf{r}}\bar{\mathbf{c}}^T}{T} \frac{\bar{s}(t_k)}{T}. \quad (13)$$

Non-stationary conditionally independent model

(NCIM): Suppose the mean total traffic $\bar{s}(t_k)$ is known, but no temporal correlation information is available. The MaxEnt model is *non-stationary and conditionally independent*, with $\mathbf{X}(t_k)$ independent $\forall k$, and

$$X_{i,j}(t_k) \sim \text{Exp}\left(\frac{N^2}{\bar{s}(t_k)}\right), \forall i, j, k. \quad (14)$$

This model assumes no information on the row and column sums in each interval. If we also knew the mean row and column sums, *i.e.*, $\mathbf{r}(t_k)$ and $\mathbf{c}(t_k)$, $\forall k$ then,

$$\mathbf{X}(t_k) = \bar{s}(t_k) \mathbf{U}(t_k) \mathbf{V}(t_k)^T, \quad (15)$$

with $U_i(t_k) \sim \text{Exp}\left(\frac{\bar{s}(t_k)}{r_i(t_k)}\right)$ and $V_j(t_k) \sim \text{Exp}\left(\frac{\bar{s}(t_k)}{c_j(t_k)}\right)$. Including covariance information of $\mathbf{r}(t_k)$ and $\mathbf{c}(t_k)$ leads once again to the equivalent truncated normal forms.

One interesting use of this model is for TMs of networks large enough that differing timezones distort traffic patterns. Traffic cycles are determined by user behavior, which is in turn determined by timezones. In a globe-spanning network, for example, users on the side of the world experiencing daytime might be transmitting traffic, while users on the other side might not.

For instance, the Abilene network [23] spans North America, with Points-of-Presence (PoPs) concentrated on the East and West Coast of the United States. A simple way to model the 3 hour difference in the traffic cycle between the coasts is to set $r_i(t_k)$ as a delayed, or phase shifted, version of a base constraint $\tilde{r}(t_k)$ (and similarly for $c_j(t_k)$). Once $r_i(t_k)$ and $c_j(t_k)$ have been generated, we generate the NCIM using (15).

In §5, we study the spatiotemporal characteristics of the above two models using Principle Component Analysis (PCA), which allows us to isolate traffic components. We will find that the models emulate some characteristics of real traffic, despite their simplicity.

In particular, these models are low rank, a property shared by real TMs [3, 19, 20, 48]. The MGM actually has a very good rank-1 approximation, and is in fact a rank-1 (three dimensional) tensor [18]. The NCIM model is hardly more complex, but the simple timezone shift described above increases its rank. However, the rank remains low.

Neither of these models can capture the spikes also seen in real traffic, leading us to the next model.

Spike model: Anomalies occur in networks due to worm propagation, distributed Denial-of-Service attacks, or flash crowds, causing a sudden spike in demand of a specific resource in the network [38]. These spikes have been demonstrated to be a feature of actual traffic in various studies [20, 33]. Here, we present a simple model for anomalies.

Under the constraint $\|\mathbf{X}\|_0 = K$ (and non-negativity), *i.e.*, exactly K entries of \mathbf{X} are non-zero and all others are zero, and setting the mean magnitude of the non-zero entries to A , the MaxEnt model is

$$X_{i,j}(t_k) \sim \begin{cases} \text{Exp}(A^{-1}), & \text{if } (i, j, k) \in \mathcal{S}, \\ 0, & \text{otherwise,} \end{cases} \quad (16)$$

where \mathcal{S} is the support of non-zero entries of size K sampled uniformly without replacement from the set $\{1, 2, \dots, N\} \times \{1, 2, \dots, N\} \times \{0, 1, \dots, M-1\}$. When $K = 1$, the model is an example of the *peak load* model [6] which is used to test network designs.

Traffic class constraints: Network traffic can often be decomposed into several classes. For instance, different network services or applications possess different characteristics, so naturally, there are separate information about each. Assuming no dependencies, the prior information on each of them are to be accounted for separately by MaxEnt.

Several works provide strong evidence of distinct traffic behaviors. Both [20, 33] showed that OD flows can be divided into three groups: large flows with a strong diurnal cycle, spike-like flows corresponding to anomalies, and small, IID white noise-like flows. Suppose we know that the TM process $\mathbf{X} = \mathbf{X}_D + \mathbf{X}_S + \mathbf{X}_N$, where \mathbf{X}_D comprises of large traffic flows with a distinct diurnal cycle, while \mathbf{X}_S and \mathbf{X}_N accounts for anomalies (spikes) and small noise-like flows respectively. With no correlation information between these components, MaxEnt suggests that the processes \mathbf{X}_D , \mathbf{X}_S and \mathbf{X}_N are independent. This means that we are free to solve for the MaxEnt distributions of these processes separately. To model large flows, we take \mathbf{X}_D to be the MGM. The spiky traffic process \mathbf{X}_S is modeled by the spike model and IID white noise process \mathbf{X}_N by the independent entries model.

Traffic classes also apply to traffic traversing peering or access links in a network. For instance, let the total traffic $T(t_k)$ be known for all k . Denote the access and peer node sets as \mathcal{A} and \mathcal{P} respectively. Now there are 4 classes of traffic: access-to-access (\mathcal{A}, \mathcal{A}), access-to-peer (\mathcal{A}, \mathcal{P}), peer-to-access (\mathcal{P}, \mathcal{A}) and peer-to-peer (\mathcal{P}, \mathcal{P}). Given the average total traffic constraints $T_{\mathcal{A}, \mathcal{A}}(t_k)$, $T_{\mathcal{A}, \mathcal{P}}(t_k)$, $T_{\mathcal{P}, \mathcal{A}}(t_k)$ and $T_{\mathcal{P}, \mathcal{P}}(t_k)$, with

$\sum_{m,n \in \{\mathcal{A}, \mathcal{P}\}} T_{m,n}(t_k) = T(t_k)$ [47], the model is

$$\mathbf{X}(t_k) = T(t_k) \sum_{m,n \in \{\mathcal{A}, \mathcal{P}\}} q_{m,n}(t_k) \mathbf{U}_m(t_k) \mathbf{V}_n(t_k)^T, \quad (17)$$

such that $\sum_{m,n \in \{\mathcal{A}, \mathcal{P}\}} q_{m,n}(t_k) = 1$.

Here, source and destination nodes are independent only within their traffic class. $\mathbf{U}_{\mathcal{A}}(t_k)$ and $\mathbf{U}_{\mathcal{P}}(t_k)$ are equivalent to source independence stratified according to access and peering nodes. There is independence in the destination nodes through $\mathbf{V}_{\mathcal{A}}(t_k)$ and $\mathbf{V}_{\mathcal{P}}(t_k)$. This form is equivalent to the generalized gravity model [46, 47], a successor to the gravity model [45], where independence is only assumed within each class of traffic, leading to the notion of *conditional independence*.

Providing information on traffic classes is a natural method to include additional data, and increase the rank of the matrices if such is seen to be too low in some problems.

General constraints: The constraints we have considered only scratch the surface of all possible constraints on TMs, for instance, nonlinear constraints. In the examples above, it is assumed that the information is given without error. If one wasn't confident about the mean information of the TM, say, a mean squared error (Frobenius norm) constraint could be used, *i.e.*,

$$\mathbb{E} \|\mathbf{X}(t_k) - \bar{\mathbf{X}}\|_F \leq \epsilon, \quad \forall k, \quad (18)$$

where $\epsilon > 0$ is a scalar value. Essentially, this expands the set of constraints from just linear to the more general set of convex constraints.

Convex constraints guarantee that (3) has a unique solution. Non-convex constraints may possibly arise in certain applications. In this case, admissible solutions have to be tested to find the global maximum of (3).

In practice, non-convex constraints are unlikely for WANs. In fact, *a priori* information from measurements can often be expressed as simple linear constraints, since linear network measurements are easier to implement (*e.g.*, internal router byte counts or SNMP measurements). Linear constraints, besides guaranteeing the existence of the MaxEnt model, typically result in simple models with closed forms, so generating these models are easy despite the high dimensionality of \mathbf{X} . However, in data centers, non-convex constraints might arise due to entry sparsity and clustering constraints of the server-to-server TMs [1, 2, 16].

Nonlinear constraints are unlikely to lead to analytically tractable closed forms of the distribution, but numerical optimization methods can be used to compute the Lagrange multipliers of (4) instead.

4. PROPERTIES

4.1 Model properties

We briefly remark about the models' connection to the axioms of MaxEnt. The seminal work of Shore and

Johnson [34] viewed MaxEnt as a natural consequence of four (desirable) inference axioms:

- A1** *uniqueness*: the solution of the inference problem must be unique,
- A2** *consistent*: either the optimal solution from a finer set of constraints must also be an admissible solution in a coarser set of constraints, or there is no solution at all in the finer set,
- A3** *system independence*: information on several independent systems can be accounted either jointly, or considered for each system separately, and
- A4** *subset independence*: information of independent subsystems (*i.e.*, in a system) can be accounted either jointly over the whole system, or separately for each subsystem.

These axioms explicitly define the assumptions our synthesis models satisfy. Generally, modeling relies heavily on two principles: *parsimony* of parameters and *independence* between components. Both are inherent in MaxEnt: parsimony is achieved by eliminating extraneous assumptions, and information of components (and subcomponents) are considered separately, in the absence of evidence that proves otherwise.

Thus, given convex constraints, **Axiom A1** guarantees a unique model. Models only get more refined as new information is given, by **Axiom A2**. The independence between the spatial and temporal processes arise due to **Axiom A3**, since no correlation information between the spatial and temporal processes are known.

As an example, consider model (17). By **Axiom A3**, each traffic class is an independent system. Since the ingress and egress traffic of each traffic class is known, by **Axiom A4**, the source and destination in each traffic class are separated into subsystems. We then have the notion of a hierarchical independence in the model: independence between traffic classes (independent systems) and conditional independence within a class (independent subsystems).

4.2 Computational issues

Generating ensembles of TMs could be surprisingly computationally expensive. For instance, given an arbitrary correlation structure, generating even simple normally distributed RVs would require specification of an $N^4 M^2$ covariance matrix, which might then need to be inverted (depending on the method of generation).

Decomposing the correlation structure, and recomposing via simple matrix and vector products leads to efficient methods for building the complete spatiotemporal correlation structure without ever having to represent this in memory.

Moreover, in our models, the distributions follow that of (4), *i.e.*, the exponential family of distributions. With simple constraints, we need only generate classical distributions such as the exponential distribution, which can be simply generated via von Neumann's inverse method [12], with a very low computational cost.




Figure 2: Total traffic for one day of the MGM with $\sigma^2 = \{1, 10\}$, $\sigma_p^2 = \{1, 100\}$, and PM ratios 1.5 and 2.

For more complicated distributions, we advocate the use of the Hamiltonian Monte Carlo (HMC) [29], as the method has the closest relationship to the exponential family (4). For instance, the multivariate truncated normal distribution is more efficiently generated via HMC [26] compared to the standard Metropolis-Hastings technique (further supported by our numerical experiments). HMC also avoids random walks, so sampling from the distribution is far more efficient than other canonical MCMC methods.

In addition, the independence property (**Axioms A3** and **A4**), enables parallelization: independent model components can be generated separately. For instance, each independent traffic class can be synthesized in parallel, then combined to form the model.

4.3 An Example

Here, we demonstrate the synthesis of a series of TMs from the MGM. To generate MGM, recall that we need three components: the average incoming and outgoing PoP traffic \mathbf{c} and \mathbf{r} respectively and the total traffic $T(t_k)$ for times $k = 0, 1, \dots, M - 1$. We set $N = 20$ and $M = 672$, which corresponds to 7 days with one hour intervals.

The constraint for the temporal component is a simple one: a sinusoid with the frequency corresponding to a daily cycle. This is a rough approximation of actual traffic, where peak traffic occur around mid-day and troughs at night. The component has two parameters: the *mean traffic* over the synthesis interval and the *peak-to-mean (PM) ratio*, *i.e.*, how high is the sinusoid’s peak relative to the mean traffic. While simple, this approach has been used to test protocols for power-aware routers [9].

The constraints $\bar{\mathbf{r}}$ and $\bar{\mathbf{c}}$ are set as proportions of the mean traffic $\bar{T} = \frac{1}{M} \sum_{k=0}^{M-1} T(t_k)$, *i.e.*, $\bar{r}_i = f_i \bar{T}$ and $\bar{c}_j = f_j \bar{T}$. The fractions f_i are non-negative and $\sum_{i=1}^N f_i = 1$, and they are generated randomly. We use the truncated normal distribution for the spatial and temporal components so as to allow control of their

variations. Independence between the components allows them to be generated separately by the HMC.

Figure 2 presents the total traffic cycle for a day. Note that the total traffic of the synthetic TMs vary around the shape of the sinusoid. In Figure 2(b), increasing the PM ratio from 1.5 to 2 results in a much larger total traffic, as expected. The larger spatial variation ($\sigma^2 = 10$) will cause large “swings” in traffic due to a larger variation in the starting point *i.e.*, the random gravity model, for each run. As we vary σ_p^2 from 1 to 100, the generated traffic becomes “noisier”, so the total traffic has more variations (Figure 2(c)). Note that the constraints will be satisfied over whole ensembles of the generated TMs, not necessarily in a single run.

Here, and in the examples to follow the TMs were generated on a machine with a 3.06 GHz processor and 4 GB of RAM with both models coded in `Matlab`. Generating the TMs for this example took on average less than a second. At $M = 2,016$ in the next section, both MGM and NCIM took about a second.

4.4 Model Analysis

Here, we want to know if our models replicate real traffic properties, such as low-rank behavior. The models we test are the MGM and the NCIM with timezone information (see §3.1).

We isolate the traffic components of the generated TMs via PCA. PCA was used to analyze traffic in [19, 20] to show that there are three types of components of OD flows, described in §3.4.

MGM: To construct constraint (8), we extracted the 23 largest Fourier coefficients from Abilene’s total traffic [23], over a week starting 1st March 2003, measured in 5-minute intervals, *i.e.*, $M = 2,016$. Besides the coefficients, all other inputs, such as the mean row sums, of the model were computed from the Abilene data.

NCIM: Base constraints $\tilde{r}(t_k)$ and $\tilde{c}(t_k) \forall k$ (time averages \tilde{r} and \tilde{c}) follow cycles derived from 23 Fourier coefficients from the traffic of a PoP belonging to the Abilene network located in New York (NY) over the




Figure 3: Largest and smallest PCA components of the MGM and NCIM where 23 Fourier coefficients were extracted from Abilene’s traffic. Both models were generated from multivariate truncated normal distributions with covariance matrix $\Sigma_{\mathbf{r}} = \Sigma_{\mathbf{c}} = \sigma^2 \mathbf{I}_N$, and $\sigma_p^2 = \sigma^2$. Control of the spatial and temporal variations is achieved by tuning σ^2 . Here, two examples are shown with $\sigma^2 = \{10^{-5}, 10\}$.

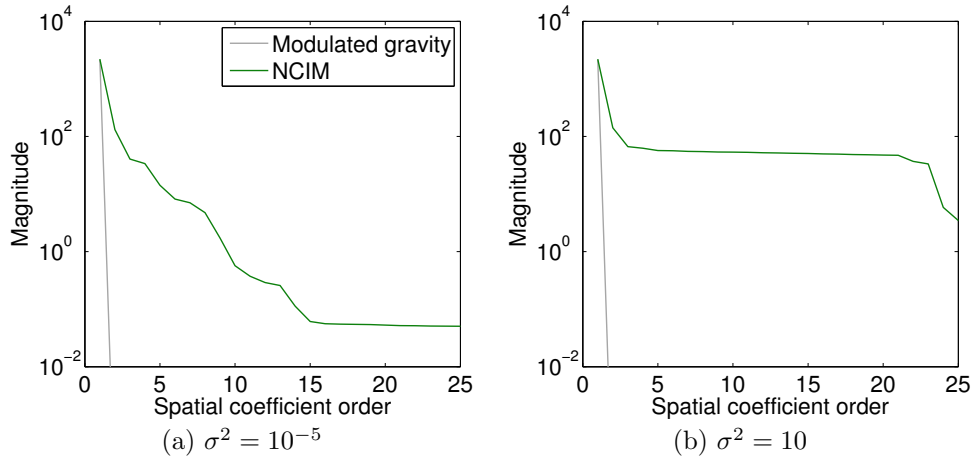


Figure 4: Spatial PCA coefficients (25 out of $N^2 = 144$) of the MGM and NCIM with variation control provided by $\sigma^2 = \{10^{-5}, 10\}$. Only coefficient values above 10^{-2} are shown for clarity. Note the log scale of the y-axis. The MGM is single rank, while the NCIM is almost full rank, with a handful of coefficients possessing large magnitudes.

same duration as before. All other constraints are delayed relative to the NY PoP. For instance, the traffic of the Los Angeles (LA) PoP, is delayed by 3 hours, since LA time is 3 hours behind NY time. The row and column sums are scaled by r_i/\bar{r} and c_j/\bar{c} , $\forall i, j$, since time averages r_i and c_j are known. Conservation of traffic implies $\bar{s}(t_k) = \sum_i r_i(t_k) = \sum_j c_j(t_k)$.

Both models’ spatial and temporal variations are controlled through σ^2 , by setting $\Sigma_{\mathbf{r}} = \Sigma_{\mathbf{c}} = \sigma^2 \mathbf{I}_N$ and $\sigma_p^2 = \sigma^2$.

Intuitively, we expect the MGM to be rank-1, so it would be too simple to replicate the properties of real traffic. The NCIM would then be a better choice since it would still be low-rank, but with enough variations to approximate a real TM process, since real TMs are unlikely to be rank-1.

Figure 3 shows the coefficients of the PCA components of the MGM and the NCIM with $\sigma^2 = \{10^{-5}, 10\}$.

For the MGM, we find that only the first component (with the largest score) has a strong cyclic behavior. Though not shown here, the cyclic behavior persisted to more than one component for the NCIM, which is closer to real traffic (such as Abilene’s) compared to the MGM. Smaller components of both models behave similar to white noise, e.g., the smallest component in Figure 3(c).

In Figure 3(b), once $\sigma^2 = 10$, we now find larger variation in the coefficient magnitudes of both models, though the diurnal cycles are still identifiable. Smaller components are also much noisier. Clearly, the models’ components have parallels to the cyclic and noise-like components of OD flows [19, 20].

Figure 4 plots the spatial PCA coefficients of models. The contrast between both models is evident: the MGM is rank-1, as expected, while the NCIM has a larger rank. Incidentally, the largest coefficient of both

models coincide with the contribution from the average gravity component $\bar{\mathbf{r}}\mathbf{c}^T/T$. The NCIM’s rank increases as σ^2 increases because σ^2 increases the variety of generated TMs. We omitted showing the temporal PCA coefficients as they match the spatial coefficients ³.

What’s interesting here is that just *differences in phase* between the row and column sums lead to a larger rank, rather than the addition of new traffic classes. This may explain the low-rank (but not single rank) behavior of real TMs observed by [19, 20, 48]. This relatively simple model reproduces complex properties such as cycles, noise-like components and low-rank behavior.

Overall, the NCIM’s properties are closer to that of real TMs than the MGM, but clearly require more detailed information, *i.e.*, the PoP locations.

5. CASE STUDY

In our case study, we apply our TM synthesis methodology to PoP-level network topology generation.

Real world PoP-level topologies are far from random, as these topologies are *designed* depending on various factors [21] such as demographic and geographical constraints, traffic demand, the designer’s experience, and financial costs. Real designs have large variability, with some as simple as a hub-and-spoke network, to more complex and meshy ones [17]. They generally do not reflect the properties of random graphs, such as the Erdős-Rényi [13] model, so a realistic topology generator must account for the above-mentioned factors.

Topology design requires the TM as an input as it determines the capacity of the links of a network. A popular choice is the gravity model [45], with several works employing a single instance of it [5, 21, 49].

First, we ask if using a series of TMs can help improve a network’s robustness. “Robustness” here refers to *robustness to TM variation*, not network survival under node removals. Second, a key result of [5] was that topology design is largely unaffected by certain facets of TMs. However, their algorithm, Combined Optimized Layered Design (COLD), used the random gravity model [31]. We extend their work by using a wider range of TMs.

We modified COLD to accept a series of TMs by defining a new optimization scheme, though this new scheme is general enough to be applied elsewhere. As far as we are aware, there is no other work on designing topologies using a series of TMs with prescribed spatiotemporal properties.

Topology generation involves combinatorial optimization, so COLD applies a genetic algorithm as a heuristic to solve for the topology with the minimum cost, based on a function with four cost parameters:

- k_0 -cost: cost for the existence of a link,
- k_1 -cost: cost for the physical length of a link,

³Spatial and temporal PCA are singular value decompositions performed on $\mathbf{Y}\mathbf{Y}^T$ and $\mathbf{Y}^T\mathbf{Y}$ respectively.

- k_2 -cost: the bandwidth cost over the length of the link, factoring in operating expenses, initial expenditures on equipment *etc.*, and
- k_3 -cost: complexity cost of a PoP with more than one link, called a non-leaf or *core* PoP.

Let $\mathcal{G}(N, E)$ be the set of all undirected graphs with N nodes and E links with sufficient capacity to carry the traffic. Let \mathcal{N}_C be the set of non-leaf nodes *i.e.*, degree > 1 and $\mathbf{w} := \{w_i | i \in E\}$ be the bandwidth of the links in E . We test the *minimax scheme*, *i.e.*,

$$\min_{\mathcal{G}(N, E)} \max_{\mathbf{w} \in \mathcal{W}(\mathcal{X})} \sum_{i \in E} (k_0 + k_1 \ell_i + k_2 \ell_i w_i) + k_3 |\mathcal{N}_C|,$$

where $\mathcal{W}(\mathcal{X})$ is the space of all possible link loads computed over \mathcal{X} , via OSPF (Open Shortest Path First) routing. The minimax scheme seems like an improvement as it factors in a range of TMs. We test this below.

5.1 Experiments and results

We first fix the locations of the PoPs in a 10 unit by 10 unit square with $N = 12$ PoPs, so randomness only enters via the input TMs. We then generated a week’s worth of 5 minute interval TMs ($M = 2,016$) via the MGM (see §4.3) as input for COLD with the recommended settings [4].

We also tested minimax with the Independent Entries Model (IEM) with the constraint $\mathbb{E}[X_{i,j}(t_k)] = \alpha \bar{X}$ to control the size of each entry. Note that we tested the model with $\alpha = 1, 2, 4$, but all results were similar, so results from $\alpha = 1$ were omitted.

The NCIM discussed in §4.3 was tested too, but we obtained similar results to the MGM, so its results were omitted. We explain the reasons for this below.

We defined the *single peak* scheme, where only the TM with the largest total traffic out of the M TMs is chosen as input in the topology design process, to compete against the minimax scheme.

We measured the Global Clustering Coefficient (GCC) of the topologies generated by the schemes. The GCC is a metric on how well-connected (clustered) a PoP is to its neighbors. The more well-connected, the more choices OSPF has in routing traffic flows. The network would then be more robust to variations in the TM.

Figure 5 show the GCC for the minimax and single peak schemes under the MGM with $\sigma^2 = 10^{-5}$, and the minimax scheme via the IEM with $\alpha = 2$. Cost $k_0 = 5$ and $k_1 = 1$, while $k_3 = \{0.1, 1, 10\}$. Each data point on $k_2 = \{0, 0.1, 1, 10, 10^2, 10^3\}$ is an average of 20 trials.

The GCCs of minimax and the single peak schemes (under the MGM) almost match. As k_2 increases, the difference between minimax and the single peak almost vanishes as both topologies become clique-like. The GCC for the IEM is lower with larger variance, especially at low k_2 . Even when $\sigma^2 = 10$ and $\alpha = 4$ in Figure 6, results were similar. We also computed other statistics such as the *average shortest paths* which decreases for all schemes and models as k_2 increases. How-




Figure 5: *GCC* for the topologies generated by the minimax and single peak schemes via the MGM and minimax with IEM, under three values of k_3 . The MGM has $\sigma^2 = 10^{-5}$, and the IEM has $\alpha = 2$. Each data point is the average of 20 trials. The 95% confidence intervals appear as vertical bars. Confidence intervals for minimax under MGM are very small that they appear unnoticeable. For all models, *GCC* is 0 at $k_2 = 0$.




Figure 6: *GCC* for the topologies generated by the minimax and single peak schemes via the MGM and minimax with IEM, under three values of k_3 . Both models have increased variance, i.e., the MGM has $\sigma^2 = 10$, and the IEM has $\alpha = 4$. Each data point is the average of 20 trials. The 95% confidence intervals appear as vertical bars. Confidence intervals for minimax under MGM are very small that they appear unnoticeable. For all models, *GCC* is 0 at $k_2 = 0$.

ever, topologies under the MGM consistently have lower average shortest paths than the IEM topologies, suggesting better robustness for the former.

The implication here is that peak TMs dominate the overall topology design. The contrast between the MGM and the IEM topologies is due to the larger peak TMs generated by the MGM. Another highlight is that the minimax topologies show very little variation (unnoticeable 95% confidence intervals) compared to the single peak. Interestingly, COLD is using a distribution of peak TMs to converge on a design that is robust to a much wider range of TMs.

We summarize two key lessons. First, PoP-level topology design is largely unaffected by TMs, in line with the result in [5]. TMs have an effect through k_2 but k_2 has relatively minor impact. This is seen through the little variation in the results even though variance in the models were increased by increasing σ^2 and α . More

extreme TMs with entries sampled from the Pareto distribution (which has heavy tails), studied by [5], also had little impact on topology generation.

Second, only peak TMs matter in improving the robustness of a network. This is by no means surprising, as it is standard practice to design networks using peak TMs. As the single peak scheme was competitive with the minimax scheme, average case TMs are not required. Instead, it would be better to use several worse case TMs as input into the design. Models with little variation such as the IEM are not useful here, in contrast to the MGM (and similarly, NCIM) which do generate peak TMs.

6. CONCLUSION

Synthesis of TMs with prescribed spatiotemporal properties is a challenging problem. A big obstacle is the lack of publicly available data for researchers. Another

important issue is controllability, as researchers require a way to generate ensembles of TMs with constraints to test network algorithms and protocols.

Our approach goes back to first principles to avoid the measure-then-model paradigm by applying the principle of maximum entropy. MaxEnt allows for control of the TM properties by setting appropriate constraints on them, while avoiding extraneous assumptions. Our models are controllable and satisfy several desirable characteristics, namely consistency and independence between model components. Despite their simplicity, they were able to replicate characteristics of real TMs such as low rank behavior.

We evaluated our models on an application: PoP-level network topology generation. We modified COLD, a topology generator, to accept a series of TMs with prescribed spatiotemporal characteristics. PoP-level topology generation turns out to be largely unaffected by TMs. We also found that only peak TMs are useful in designing networks that are robust to TM variations. Thus, the common practice of using busy period TMs to design topologies is an excellent strategy.

Future work will focus on developing an algorithm for sampling maximum entropy distributions handling general constraints *e.g.*, correlation constraints and cyclostationarity, and studying the TMs of data centers.

Acknowledgments

The work was supported by the Australian Research Council (ARC) grant DP110103505. We thank our shepherd, Anja Feldmann and the anonymous reviewers for their feedback, eResearch South Australia (eRSA) for providing computational resources used in the experiments, and Rhys Bowden for help with COLD.

7. REFERENCES

- [1] T. Benson, A. Akella, and D. A. Maltz. Network traffic characteristics of data centers in the wild. In *Proc. ACM SIGCOMM Internet Measurement Conference (IMC)*, pages 267–280, November 2010.
- [2] T. Benson, A. Anand, A. Akella, and M. Zhang. MicroTE: Fine grained traffic engineering for data centers. In *Proc. ACM SIGCOMM CoNEXT*, pages 1–12, December 2011.
- [3] V. Bharti, P. Kankar, L. Setia, G. Gürsun, A. Lakhina, and M. Crovella. Inferring invisible traffic. In *Proc. ACM SIGCOMM CoNEXT*, pages 22:1–22:12, 2010.
- [4] R. Bowden. COLD: PoP-level topology generator. <https://github.com/rhysbowden/COLD/>, 2014.
- [5] R. Bowden, M. Roughan, and N. Bean. COLD: PoP-level topology synthesis. In *Proc. ACM SIGCOMM CoNEXT*, pages 173–184, December 2014.
- [6] R. S. Cahn. *Wide Area Network Design*. Morgan Kaufmann, 1998.
- [7] A. Calaprice. *The Ultimate Quotable Einstein*. Princeton University Press, 2010.
- [8] E. Candes and B. Recht. Exact matrix completion via convex optimization. *Found. of Comput. Math.*, 9:717–772, 2008.
- [9] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsang, and S. Wright. Power awareness in network design and routing. In *Proc. IEEE INFOCOM*, April 2008.
- [10] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley and Sons, Inc., 2nd edition, 2006.
- [11] N. Cressie and C. K. Wikle. *Statistics for Spatiotemporal Data*. Wiley, 2011.
- [12] L. Devroye. *Non-Uniform Random Variate Generation*. Springer, 1986.
- [13] P. Erdős and A. Rényi. On random graphs. *Publicationes Mathematicae*, 6:290–297, 1959.
- [14] V. Erramilli, M. Crovella, and N. Taft. An independent-connection model for traffic matrices. In *Proc. ACM SIGCOMM Internet Measurement Conference (IMC)*, pages 251–256, October 2006.
- [15] E. T. Jaynes. Information theory and statistical mechanics. *Phys. Rev.*, 106:620–630, May 1957.
- [16] S. Kandula, S. Sengupta, A. Greenberg, P. Patel, and R. Chalken. The nature of datacenter traffic: Measurements and analysis. In *Proc. ACM SIGCOMM Internet Measurement Conference (IMC)*, pages 202–208, November 2009.
- [17] S. Knight, H. Nguyen, N. Falkner, R. Bowden, and M. Roughan. The Internet Topology Zoo. *IEEE J. Select. Areas Commun.*, 29(9):1765–1775, October 2011.
- [18] T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, September 2009.
- [19] A. Lakhina, M. Crovella, and C. Diot. Diagnosing network-wide traffic anomalies. In *Proc. ACM SIGCOMM*, pages 219–230, September 2004.
- [20] A. Lakhina, K. Papagiannaki, M. Crovella, C. Diot, E. D. Kolaczyk, and N. Taft. Structural analysis of network traffic flows. *SIGMETRICS Perform. Eval. Rev.*, 32(1):61–72, June 2004.
- [21] L. Li, D. Alderson, W. Willinger, and J. Doyle. A first-principles approach to understanding the Internet’s router-level topology. In *Proc. ACM SIGCOMM*, pages 3–14, August 2004.
- [22] A. Medina, N. Taft, K. Salmatian, S. Bhattacharyya, and C. Diot. Traffic matrix estimation: Existing techniques and new directions. In *Proc. ACM SIGCOMM*, 2002.
- [23] NLANR. Abilene Trace Data. <http://pma.nlanr.net/Special/ipls3.html>.
- [24] A. Nucci, A. Sridharan, and N. Taft. The problem of synthetically generating IP traffic matrices:

- Initial recommendations. *ACM Comput. Commun. Rev.*, 35:19–32, July 2005.
- [25] K. N. Oikonomou. Analytic forms for most likely matrices derived from incomplete information. *Int. J. Systems Science*, 43(3):443–458, September 2010.
- [26] A. Pakman and L. Paninski. Exact Hamiltonian Monte Carlo for truncated multivariate Gaussians. *J. Comput. Graph. Stat.*, 23(2):518–542, April 2014.
- [27] S. Y. Park and A. K. Bera. Maximum entropy autoregressive conditional heteroskedasticity model. *J. Econometrics*, 150:219–230, 2009.
- [28] R. B. Potts and R. M. Oliver. *Flows in Transportation Networks*. Academic Press, 1972.
- [29] N. M. Radford. MCMC using Hamiltonian dynamics. In S. Brooks, A. Gelman, G. L. Jones, and X.-L. Meng, editors, *Handbook of Markov Chain Monte Carlo*, volume 31-32, pages 113–162. CRC Press, 2010.
- [30] B. Recht, M. Fazel, and P. A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, 2010.
- [31] M. Roughan. Simplifying the synthesis of Internet traffic matrices. *ACM Comput. Commun. Rev.*, 35(5):93–96, 2005.
- [32] M. Roughan. Robust network planning. In C. R. Kalmanek, S. Misra, and R. Yang, editors, *The Guide to Reliable Internet Services and Applications*, chapter 5, pages 137–177. Springer, 2010.
- [33] M. Roughan, A. Greenberg, C. Kalmanek, M. Rumsewicz, J. Yates, and Y. Zhang. Experience in measuring backbone traffic variability: Models, metrics, measurements and meaning. In *Proc. ACM SIGCOMM Internet Measurement Conference (IMC)*, 2002.
- [34] J. E. Shore and R. W. Johnson. Axiomatic derivation of the principle maximum entropy and the principle of minimum cross-entropy. *IEEE Trans. Info. Theory*, 26(1):26–37, January 1980.
- [35] A. Soule, A. Lakhina, N. Taft, K. Papagiannaki, K. Salmatian, A. Nucci, M. Crovella, and C. Diot. Traffic matrix: Balancing measurements, inference and modeling. In *Proc. ACM SIGMETRICS*, pages 362–373, June 2005.
- [36] A. Soule, A. Nucci, R. Cruz, E. Leonardi, and N. Taft. How to identify and estimate the largest traffic matrix elements in a dynamic environment. *SIGMETRICS Perform. Eval. Rev.*, 32(1):73–84, June 2004.
- [37] S. M. Stigler. *The history of statistics: The measurement of uncertainty before 1900*. Belknap Press of Harvard University Press, 1986.
- [38] P. Tune and M. Roughan. Internet traffic matrices: A primer. In H. Haddadi and O. Bonaventure, editors, *Recent Advances in Networking, Vol. 1*. ACM SIGCOMM, August 2013.
- [39] P. Tune and M. Roughan. Maximum entropy traffic matrix synthesis. In *ACM SIGCOMM MAMA*, June 2014. <http://tinyurl.com/kk7757k>.
- [40] P. Tune and M. Roughan. Network design sensitivity analysis. In *Proc. ACM SIGMETRICS*, pages 449–461, June 2014.
- [41] S. Uhlig, B. Quoitin, J. Lepropre, and S. Balon. Providing public intradomain traffic matrices to the research community. *ACM Comput. Commun. Rev.*, 36(1):83–86, January 2006.
- [42] Y. Vardi. Network Tomography: Estimating source-destination traffic intensities from link data. *J. Am. Statist. Assoc.*, 91:365–377, 1996.
- [43] A. Zellner and R. A. Highfield. Calculation of maximum entropy distributions and approximation of marginal posterior distributions. *J. Econom.*, 37:195–209, 1988.
- [44] Y. Zhang, Z. Ge, A. Greenberg, and M. Roughan. Network anomography. In *Proc. ACM SIGCOMM Internet Measurement Conference (IMC)*, Berkeley, California, USA, October 2005.
- [45] Y. Zhang, M. Roughan, N. Duffield, and A. Greenberg. Fast accurate computation of large-scale IP traffic matrices from link loads. In *Proc. ACM SIGMETRICS*, pages 206–217, 2003.
- [46] Y. Zhang, M. Roughan, C. Lund, and D. Donoho. An information-theoretic approach to traffic matrix estimation. In *Proc. ACM SIGCOMM*, pages 301–312, 2003.
- [47] Y. Zhang, M. Roughan, C. Lund, and D. Donoho. Estimating Point-to-Point and Point-to-Multipoint traffic matrices: An information-theoretic approach. *IEEE/ACM Trans. Netw.*, 13(5):947–960, October 2005.
- [48] Y. Zhang, M. Roughan, W. Willinger, and L. Qiu. Spatio-Temporal compressive sensing and Internet traffic matrices. In *Proc. ACM SIGCOMM*, pages 267–278, August 2009.
- [49] R. Zhang-Shen and M. McKeown. Designing a predictable Internet backbone with Valiant load-balancing. In *Proceedings of the 13th International Conference on Quality of Service (IWQoS)*, pages 178–192, 2005.