A Skewness-Aware Matrix Factorization Approach for Mesh-Structured Cloud Services

Yongquan Fu^(D), Dongsheng Li, Pere Barlet-Ros, Chun Huang, Zhen Huang, Siqi Shen, and Huayou Su

Abstract—Online cloud services need to fulfill clients' requests scalably and fast. State-of-the-art cloud services are increasingly deployed as a distributed service mesh. Service to service communication is frequent in the mesh. Unfortunately, problematic events may occur between any pair of nodes in the mesh, therefore, it is vital to maximize the network visibility. A state-of-theart approach is to model pairwise RTTs based on a latent factor model represented as a low-rank matrix factorization. A latent factor corresponds to a rank-1 component in the factorization model, and is shared by all node pairs. However, different node pairs usually experience a skewed set of hidden factors, which should be fully considered in the model. In this paper, we propose a skewness-aware matrix factorization method named SMF. We decompose the matrix factorization into basic units of rank-one latent factors, and progressively combine rank-one factors for different node pairs. We present a unifying framework to automatically and adaptively select the rank-one factors for each node pair, which not only preserves the low rankness of the matrix model, but also adapts to skewed network latency distributions. Over real-world RTT data sets, SMF significantly improves the relative error by a factor of 0.2 x to 10 x, converges fast and stably, and compactly captures fine-grained local and global network latency structures.

Index Terms—Service mesh, matrix factorization, skewness, latent factor model, residual learning.

I. INTRODUCTION

ARGE-scale cloud services are typically organized as a mesh of micro-services deployed over hundreds to thousands of nodes, as illustrated in Figure 1. Service-to-service communication is frequent on the mesh structured service topology. For example, a Web request may traverse thousands of servers to search and aggregate results. A request's service level agreement (SLA) is determined based on the response from the slowest server. As problematic locations are

Manuscript received May 7, 2018; revised September 14, 2018 and April 22, 2019; accepted June 12, 2019; approved by IEEE/ACM TRANSAC-TIONS ON NETWORKING Editor J. Liu. This work was sponsored in part by National Key R&D Program of China under Grant No. 2018YFB0204300, the National Natural Science Foundation of China (NSFC) under grant 61602500 and 61402509, and in part by the Spanish Ministry of Economy and Competitiveness (MINECO) and the European Fund for Economic and Regional Development (EU FEDER) under grant TEC2017-90034-C2-1-R (ALLIANCE project). (*Corresponding author: Dongsheng Li.*)

Y. Fu, D. Li, Z. Huang, S. Shen, and H. Su are with the Science and Technology Laboratory of Parallel and Distributed Processing, College of Computer, National University of Defense Technology, Changsha 410073, China (e-mail: fyongquanf@nudt.edu.cn; dsli@nudt.edu.cn; huangzhen@nudt.edu.cn; shensiqi@nudt.edu.cn; hyoug@nudt.edu.cn).

P. Barlet-Ros is with the Computer Architecture Department, UPC BarcelonaTech, 08034 Barcelona, Spain (e-mail: pbarlet@ac.upc.edu).

C. Huang is with the College of Computer, National University of Defense Technology, Changsha 410073, China (e-mail: chunhuang@nudt.edu.cn).

This paper has supplementary downloadable material available at http:// ieeexplore.ieee.org, provided by the authors. This includes a PDF containing Appendices A–D.

Digital Object Identifier 10.1109/TNET.2019.2923815

node-1 node-N node-2 application Service A Service A Service A traffic Service B Service B Service B Service C Service C Service C data plane data I proxied traffic monitor and control

Fig. 1. An illustration of a service-mesh scheme [3], [4]. A cloud service consists of a collection of loosely-coupled micro-services hosted on a set of distributed servers. The service-to-service communication is managed by a dedicated service-mesh proxy layer for autonomous traffic control and optimization. The monitoring functionality is located at the proxy layer to collect pairwise RTT status for the service mesh layer.

unpredictable a priori, we need to track RTTs by all nodes and for all nodes [23].

Timely response is vital for ensuring the Quality of Experience (QoE) [29], [45], otherwise, the increased delay significantly affects users' experience and providers' revenue [10], [51]. Unfortunately, high latency issues may arise between any node pairs of the service mesh, due to changing routing paths, degraded path conditions, or transient network congestions. To correctly correlate the network issue with the problematic service, it is vital to tell if a service is affected by a network issue by monitoring pairwise RTTs of the service mesh. Further, optimizing the service mesh needs pairwise RTTs. For example, we may choose the detour routing scheme [13], [40], [52] that selects a relay on the direct routing path between a set of hosts and forwards packets via the relay.

A straightforward approach is to directly collect all-pair network latency, which requires a quadratic number of probing packets with respect to the system size. To provide enough network visibility, we need to predict missing measurement results. The measurement may also disturb normal application traffic, as node resources are usually shared among multiple tenants.

To reduce the probing cost, researchers have proposed prediction methods that embed nodes into a low-dimensional vector space and estimate the pairwise RTTs based on the vector distances. This representation model succinctly captures the pairwise network latency matrix, and needs only O(N)probes to predict the full pairwise matrix for N hosts in edge data centers. Further, the vector representation can act as the input for various applications [15], [16], [26].

The vector space follows from a low-rank matrix factorization model [17], [31], [32], [34], which represents each node as a vector of variables. These variables are "latent" in the sense that they are not directly observed, but should be inferred via a mathematical model from observed data.

1063-6692 © 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

A matrix factorization model is well known to be equivalently represented as the sum of a set of rank-one matrices [22]. A rank-one matrix represented as the product of two vectors $F_k G_k^T$ for $k \leq r$ (the constant r denotes the number of rank-one matrices) serves as a "latent factor"; and the vector F_k and G_k serve as the soft memberships of each row and each column of the RTT matrix to this latent factor [12], [43], respectively. As the same set of rank-one matrices are shared by all nodes, all node pairs are implicitly assumed to share the same set of latent factors. However, our analyses in Section III-C show that a real-world network latency matrix is heterogeneous and highly skewed, thus different node pairs are typically correlated with heterogeneous factors. Thus the matrix factorization should be revisited to account for skewed factors.

In this paper, we focus on improving the adaptation to the skewness for the matrix factorization model. We propose a skewness-aware matrix factorization method called SMF, which decomposes the matrix factorization framework into the basic units of rank-one matrices, and selectively combines the rank-one entries for each node pair to account for skewed factors.

The first challenge is how to smartly combine the rank-one matrices. Existing methods assume positive correlations between the rank-one matrices and the approximation results. We relax this assumption and consider generalized scenarios where each rank-one matrix is either positively correlated, negatively correlated or irrelevant. Then, we model these relations within a unifying residual-learning framework, by extending the well-known orthogonal matching pursuit algorithm [37]. In each iteration, we add a new rank-one matrix into the prediction, where each entry is either added, bypassed, or decreased to account for the skewness of the network latency distribution.

The second challenge is how to complete the selection decision for each entry. We treat the selection decision as the rating score and the service mesh as both clients and goods, and establish a collaborative filtering task that predicts missing rating scores between a set of clients and a set of goods. We propose an adaptive method to map the collaborating filtering scores to discrete combination choice by extending the well-known maximum margin matrix factorization method [38], [44].

Finally, extensive experiments using real-world data sets confirm that our approach reduces the relative errors by a factor of 0.2x to 10x compared to state-of-the-art methods. SMF converges fast and stably. Applying SMF to inform the low-latency detour routing [13], [40], [41] achieves close to optimal performance. Further, we have evaluated the parameter regions where SMF obtains good approximations and verified its computational efficiency.

In summary, we make three primary contributions in this paper:

- We quantify the skewness of the RTT metric with respect to a set of local and global metrics.
- We develop a skewness-aware matrix factorization framework SMF that adaptively learns rank-one latent factors.

• We perform extensive experiments on real-world data sets to confirm that that SMF finds a good balance between the low rankness and the adaptation to skewed RTT distributions.

The rest of the paper is organized as follows. Section II summarizes the related literature. Next, Section III presents the background and states the requirements for skewed latent factors. Next, Section IV introduces the basic ideas of the skewness-ware matrix factorization. Section V presents the detailed algorithms and analysis. Section VI reports simulation experiments compared with state-of-the-art methods. Finally, we conclude in Section VII.

II. RELATED WORK

Extensive studies have been made to enable network-latency measurement for large-scale distributed systems and data center networks. We only introduce representative studies that are most related to us.

A. Mesh Network Monitoring

iPlane [33] predicts end to end network latency based on an Internet topology model. iPlane issues active probes from wide-area vantage points to routable network addresses. PingMesh [23] collects all-pair round-trip time (RTT) measurement system in several scales for geo-distributed data centers, which accumulates 24 TBs of probe results each day. [50] approximates network latency in an OpenFlow network environment based on control messages to and from the OpenFlow controller. Mobilyzer [36] provides a controlled and isolated library for mobile network measurement experiments. Our work is complementary to these studies by predicting missing measurements.

B. Network Latency Prediction

Researchers proposed to predict network latency with network coordinate methods for scalability. GNP [35] pioneers this field with an Euclidean coordinate system. Vivaldi [8] combines a 2-d coordinate system with a height model that reflects the first-hop delay of traversing the accessing link. IDES [34] and DMFSGD [31] embed nodes in a two-factor matrix factorization model. DMFSGD [31] trains the low-rank matrix factorization with the SGD optimization technique. Liu et al. [32] decompose the latency metric to a distance component and a network feature component for flexibility. Fu and Xu [17] stabilizes the matrix factorization process under churns via a relative coordinate based matrix model. Zhu et al. [54] propose an adaptive matrix factorization approach by data transformation. The transformed metrics become more symmetric than the raw metrics, however, retransformed metrics become less stable when the observation is incomplete, since the estimator's inaccuracy will be amplified exponentially with respect to the transformation base. A second approach is to estimate the quantiles [56] based on the quantile regression framework. Generally, these methods assume that each pair of nodes shares the same

set of latent factors, which may not hold when node pairs experience diverse hidden factors. Our work addresses this challenge via a skewness-aware matrix factorization model.

C. Matrix Completion

Our study is related with the matrix-completion theory [5], which recovers an incomplete matrix via a subset of observed entries. For a rank- $r \ m \times n$ matrix ($r \ll (m, n)$) that meets an incoherent condition¹, a unique rank-r matrix can be recovered with a high probability. Minimizing the matrix rank exactly is NP-hard [5]. OR1MP [43] iteratively finds a rank-one matrix out of the approximation residual with the SVD. However, it is generally impossible to exactly recover the SVD result for a partially observed matrix. Further, different node pairs are likely to be correlated with heterogeneous latent factors.

D. Traffic Matrix Interpolation

Real-world traffic matrices are usually incomplete. Consequently, interpolating missing entries becomes important. Traffic matrix interpolation is a related, but different problem, with different properties. Xie *et al.* [46]–[48] exploit hidden spatial and temporal structures with three-dimensional lowrank tensors, which effectively reduces the estimation error. Zhang *et al.* [53] interpolate incomplete traffic matrices with structure regularized low-rank matrix factorization and local interpolation procedures. LENS [7] models the traffic matrices as the sum of multiple matrices that are positively correlated with the traffic matrix. Our work proposes a unifying model that keeps the low-rank interpretation and adapts well to skewed latent factors.

III. PROBLEM STATEMENT

We first present the measurement environment for the mesh-structured cloud services, then introduce the matrix-factorization results, and discuss the open questions.

A. Measurement Architecture

A service mesh typically consists of a set of nodes located in mega data center networks or edge data-center networks. Each node hosts a set of networked micro-services, as discussed in the introduction. Service-to-service communication is frequent, while the latency between sending service requests and obtaining responses should meet network SLAs. We assume that, the service mesh should have synchronized their clocks, as otherwise we could not correlate the network problems in different locations. The synchronization protocols such as Network Time Protocol (NTP) [2] or the IEEE 1588 Precise Time Protocol (PTP) [1] can provide millisecond-level precision for geo-distributed nodes.

The measurement system is comprised of two main components inspired by the software defined networks [23], [33], [50]: a data plane that consists of service-mesh nodes and a control plane on a logically centralized server.

(i) At the control plane, the logically centralized controller schedules the Round-Trip Time (RTT) measurement process

in the data plane. The controller randomly samples a small list of nodes as *probing targets* for each service-mesh node. The choices of probing targets are randomized for different nodes for load balancing. The number of probing targets depends on the node's measurement capability. For a scale of hundreds of nodes, selecting tens of probing targets suffices to obtain a good level of accuracy.

Further, the controller handles the churns of nodes, since an offline node is useless and should be detected and filtered. Accordingly, the controller keeps the online status of the data plane as volatile states in the main memory. Each online node periodically sends a heart-beating message to the centralized controller to notify its online status; as a response, the controller piggybacks a list of sampled online nodes. The frequency of the heart-beat messages is platform-dependent, where stable platforms could choose a long period, while edge platforms should choose a relatively short period (e.g., one minute) to reflect system churns.

(ii) At the data plane, each service-mesh node performs a number of measurements towards other nodes in the same platform. It downloads the list of probing targets from the controller, and measures the RTT status towards these probing targets in a periodical approach. After collecting the RTT samples in an interval, each node uploads the RTT results to the persistent storage that is accessed by the controller.

The data plane could use any kinds of measurement methods. For example, at the network or transport level, the data plane may choose ICMP or TCP protocol based measurement methods; at the application level, the data plane could use RPC or HTTP protocol based methods. Generally, the RTT value amounts to the absolute difference between the time of sending a request message to the probing target and that of receiving the response message from this probing target. The unit of a measurement interval determines the granularity of the monitoring process. Increasing the sampling interval towards a probing target yields a coarser granularity.

B. Challenges for RTT Matrix Completion

For a set of N nodes, the pairwise RTTs between N nodes in an interval can be represented as a N-by-N matrix **D**. The state-of-the-art approaches predict pairwise RTT values based on the matrix factorization approach, which factorizes a matrix $\mathbf{D} \in \mathbb{R}^{N \times N}$ as a product of two low-dimensional factor matrices $\mathbf{F} \in \mathbb{R}^{N \times r}$ and $\mathbf{G} \in \mathbb{R}^{N \times r}$, i.e., $\mathbf{D} \approx \mathbf{F}\mathbf{G}^{T}$, where $r \ll N$, and T denotes the transpose of a matrix.

A matrix factorization model is equivalent to a sum of a set of rank-one matrices:

$$\hat{\mathbf{D}} = FG^T = \sum_k F_k G_k^T \tag{1}$$

where F_k , G_k denote the k-th ($k \le r$) column vector of the matrix F and G, respectively. An objective function seeks to minimize the approximation residual between the observed entries and the sum of the rank-one matrices:

$$\min_{F,G} \left\| \mathbf{D} - \sum_{k=1}^{r} F_{*k} G_{*k}^{T} \right\|$$
(2)

¹*Incoherence* [5] states that the singular vectors spread out to help the matrix be loosely aligned with the coordinate axes.

As the matrix factorization approximates the target matrix **D** via the sum of each rank-one matrix $F_k G_k^T$, each of which is assumed to be positively correlated with the approximation results. Consequently, the matrix factorization assumes that all node pairs experience the same set of latent factors. This assumption may not hold for the RTT metric, which additively consists of many factors, e.g., the propagation latency, the queueing latency, the transmission latency. For two node pairs with a number of different routing links, their RTTs are likely to be affected by independent latency components, or hidden factors. Furthermore, the RTT metric only reveals the sum of all factors, not individual factors.

C. Empirical Distributions

Next, we empirically analyze the RTT characteristics of real-world data sets and motivate the design requirements.

1) Data Set: Real-world network connections are heterogeneous, therefore, an ideal network latency prediction algorithm should adapt to different network connections. We choose three publicly available data sets that differ in terms of the node size and RTT distributions to study real-world performance of different prediction algorithms:

- Seattle: The Seattle platform is an open peer-to-peer cloud computing platform that includes donated personal devices like personal computers, laptops and mobile phones [6] This data set was collected in summer 2014 for three hours between 99 nodes. Each interval aggregates pairwise RTT samples within 15.7 seconds, which indicates short-term dynamics between Seattle nodes.
- **PlanetLab**: The PlanetLab platform has been widely used in many previous RTT-prediction studies [17], [31], [33]. This data set was collected in 2013 for a 9-day period between 490 distributed nodes. Each interval aggregates pairwise RTT measurements within 14.7 hours, which represents long-term RTT trends between PlanetLab nodes.
- **RIPE**: The RIPE Atlas measurement platform consists of tiny networked devices that issue measurements to a small number of addresses, most of which are chosen by the platform owner. There are two kinds of devices, i.e., Probes and Anchors, depending on their measurement capability. The Anchor is more powerful than the Probe. As we need dense RTT matrices to train the matrix factorization methods, we choose RIPE Atlas nodes that are powerful enough to probe each other. This data set² was collected in August 19, 2018 between 250 RIPE Atlas nodes.

Table I summarizes several basic statistics about the data sets. We can see that the data sets span wide ranges.

2) Distributions: First, we compare the RTT distributions for each node. We calculate the interquartile ranges for each RTT vector, i.e., the difference between the 25th and the 75th percentiles of the samples in the vector. In Figure 2, the interquartile ranges span wide intervals, thus different node pairs are likely to experience diverse hidden factors.

TABLE I BASIC STATISTICS OF DATA SETS

IEEE/ACM TRANSACTIONS ON NETWORKING

Trace	Interval	Mean	STD	Min	Max
Seattle	15.7s	0.37s	0.90s	0.01s	90.50s
PlanetLab	14.7hr	147.38ms	103.65ms	0.064ms	7892.8ms
RIPE	2hr	118.47ms	106.02ms	0.08ms	10,425.75ms



Fig. 2. The CCDFs of interquartile ranges of each row vector in the dataset.



Fig. 3. Heat maps of the pairwise correlation coefficients. We choose the first interval for illustration purpose. Varying the intervals yields the same conclusions. We randomly select a number of nodes from the data set as the landmarks, calculate the vector of correlation coefficients from this node to the landmarks. Then we compute the K-means clustering with these feature vectors based on the Lloyd method [49]. Finally, we reorganize the correlation coefficient matrix by putting nodes in the same cluster at adjacent positions, and plot the heat map of the matrix where darker pixels correspond to larger correlation values. During the experiments, we set the number of landmarks to 16, the number of clusters to three. The same conclusions hold as we vary the parameters.

Next, we evaluate the correlation between pairs of nodes. We compute the linear correlation coefficient [39] for any pair of nodes i and j, defined as

$$r(i,j) = \frac{\left\langle \vec{D_i} \circ \vec{D_j} \right\rangle - \left\langle \vec{D_i} \right\rangle \left\langle \vec{D_j} \right\rangle}{\sigma\left(\vec{D_i} \right) \sigma\left(\vec{D_j} \right)}$$

where D denotes the pairwise RTT matrix, \vec{D}_j denotes the *i*-th row vector, \circ denotes the hadamard operator

$$\left(\vec{D}_i \circ \vec{D}_j\right)_k = D_{ik} D_{jk}$$

for $i, j, k \in [1, N]$, $\langle \cdot \rangle$ denotes the average of the vector, and

$$\sigma\left(\vec{D_i}\right) = \sqrt{\left\langle \vec{D_i} \circ \vec{D_i} \right\rangle - \left\langle \vec{D_i} \right\rangle^2}.$$

Figure 3 plots the heat map of the linear correlation coefficients of each node pair. Darker pixels correspond to larger correlation coefficients. Two to three groups of nodes are separable from the rest of the plot, where intra-group correlation coefficients are relatively larger than node pairs from different groups. Thus the pairwise RTT distribution is highly skewed, and intra-group nodes are more likely to experience similar latent factors than inter-group node pairs. Thus we need to study more powerful representative models to account for these skewed distributions.

²https://data-store.ripe.net/datasets/atlas-daily-dumps/

TABLE II Frequently Used Notations in This Paper

Notations	Meaning
N	Number of nodes
$f D, \hat D \ f X$	RTT matrix and the estimated RTT matrix rank-1 matrix
β,\hat{eta}	Sign matrix and the estimated sign matrix
\mathbf{E}	Approximation residual
F, G	rank-1 matrix model
$(ec{u},ec{v},ec{ heta},b)$	Sign matrix model
r	Approximation rank
S_i	Probing targets of node <i>i</i>
n_p	Number of probing targets
r_s	Dimension of the estimated sign matrix
λ	Regularized parameter

IV. SKEWNESS-AWARE MATRIX FACTORIZATION

Having analyzed the skewness with real-world RTT data sets, we see that the RTTs of different node pairs are likely to be affected by divergent hidden factors, which should be accounted by the prediction algorithm. Next, we first present naive approaches and discuss its limitations, then present a new model that keeps the interpretation of the low-rank representation and adapts to the skewness of the RTT distributions.

Table II summarizes frequently used notations used in this paper.

A. Strawman Approaches

1) Increasing Rank: A straightforward approach is to increase the rank of the matrix factorization model, which amounts to the maximal number of columns or rows that are linearly independent with each other. The rank is neatly characterized by the number of positive singular values of the SVD of this matrix [22]. The SVD represents D as D = USV^T , where V^T denotes the transpose of $V, U \in \mathbb{R}^{N \times N}$ is an orthogonal matrix, i.e., $UU^T = U^T U = I, V \in \mathbb{R}^{N \times N}$ is also an orthogonal matrix, i.e., $VV^T = V^T V = I$, S is a squared and diagonal matrix consisting of a vector of descendingly-ordered real numbers $(\sigma_1, \cdots, \sigma_N)$, where $\sigma_i \geq 0$ for $i \leq N$. Further, the SVD amounts to the sum of k rank-1 matrices $\hat{D} = \sum_{i=1}^k \delta_i u_i v_i^T$, where δ_i is the *i*-th singular value, u_i is the *i*-th left singular vector, and v_i is the *i*th right singular vector. Each rank-1 matrix represents a latent factor that introduces a degree of freedom to approximate the matrix.

Figure 4 shows the residual variance and the relative error with an increasing number of top singular values. First, 90% of its variance can be captured with two to four top singular values, therefore, the RTT matrix is approximately low-rank. Second, the relative error is still high, since a longer tail of singular values implies a higher relative error. This is due to the fact that, the residual of the SVD approximation is correlated with the remaining set of singular values, which is represented as $D - \hat{D}_k = \sum_{i=1}^N \delta_i u_i v_i^T - \sum_{i=1}^k \delta_i u_i v_i^T = \sum_{i=k+1}^N \delta_i u_i v_i^T$.

2) Incorporating Weights: A second approach is to regularize the latent factor model for different node pairs. For example, [55] sets $\hat{D}'[i,j] = \hat{D}[i,j] \cdot w[i,j]$, where w



Fig. 4. Energy contained in the top-k singular values VS. the relative error of the rank-k approximation using the SVD [22]. We compute the **relative error** $\frac{|\mathbf{D}_{ij} - \hat{\mathbf{D}}_{ij}^r|}{|\mathbf{D}_{ij}|}$ for each node pair between the rank-k approximation $\hat{\mathbf{D}}^r$ and the RTT matrix **D**. Further, we compute residual fraction of the total variance captured by top-k singular values as $1 - \frac{\sum_{i=1}^{k} \delta_i^2}{\sum_{i=1}^{N} \delta_i^2}$, where δ_i represents the *i*-th largest singular value.

denotes the weight matrix. The estimation of each node pair is scaled independently, which introduces N^2 degrees of freedom to the prediction \hat{D} . Accordingly, the weighted model has enough degrees of freedom to recover any matrix exactly. [55] assigns an N-by-N real-valued weighted matrix to the low-rank matrix, and estimates the weight matrix in the matrix completion framework. However, [55] still assumes that each latent factor is of equal importance for each node pair, which may not hold due to complex routing decisions and varying latency components.

B. Ideal Skewness-Aware Model

Next, we present an ideal model to account for skewed latent factors assuming that we obtain the complete RTT matrix. In the next subsection, we relax this assumption and present a practical approach.

Recall that matrix factorization is equivalent to the sum of rank-1 matrices, while each rank-1 matrix serves as a latent factor. To make the matrix factorization be aware of skewed latent factors, we should incorporate a rank-1 matrix entry for a node pair only if this node pair is correlated with this latent factor.

Let k denote the index of the current rank-1 matrix $(1 \leq k \leq r)$. Let $\hat{\mathbf{D}}_k$ be the approximation result of combining the first k rank-1 matrices, and $\mathbf{E}_k = \mathbf{D} - \hat{\mathbf{D}}_{k-1}$ the current residual, where $\hat{\mathbf{D}}_0$ represents an empty matrix.

1) Correlation Model: We define three kinds of correlation types to combine the latent factor for each node pair (i, j) as follows:

(i) *Positive correlation*: If the approximation residual $\mathbf{E}_{k}[i, j]$ will be smaller by adding the entry $\mathbf{X}_{k}[i, j]$ to the current approximation $\hat{\mathbf{D}}_{k-1}[i, j]$, i.e., $\mathbf{D} - (\hat{\mathbf{D}}_{k-1}[i, j] + \mathbf{X}_{k}[i, j]) < \mathbf{E}_{k}[i, j]$, then the entry $\mathbf{X}_{k}[i, j]$ is said to be **positively correlated** with the current approximation residual. Accordingly, we update the approximation for the node pair i, j as: $\hat{\mathbf{D}}_{k}[i, j] = \hat{\mathbf{D}}_{k-1}[i, j] + \mathbf{X}_{k}[i, j]$.

(ii) *Negative correlation*: Otherwise, if the approximation residual $\mathbf{E}_k[i, j]$ will be smaller if we subtract the entry $\mathbf{X}_k[i, j]$ from the current approximation $\hat{\mathbf{D}}_{k-1}[i, j]$, i.e., $\mathbf{D} - (\hat{\mathbf{D}}_{k-1}[i, j] - \mathbf{X}_k[i, j]) < \mathbf{E}_k[i, j]$, then the entry $\mathbf{X}_k[i, j]$ is claimed to be **negatively correlated** with the approximation residual. To reduce the approximation error, we update the approximation as: $\hat{\mathbf{D}}_k[i, j] = \hat{\mathbf{D}}_{k-1}[i, j] - \mathbf{X}_k[i, j]$.



Fig. 5. The diagram of the ideal correlation model.

(iii) *Irrelevance*: If the above (i) and (ii) do not hold, we should skip this component $\mathbf{X}_k[i, j]$ for node pair i, j, since otherwise, the approximation residual will increase. This entry $\mathbf{X}_k[i, j]$ is said to be *irrelevant* with the current approximation residual.

Accordingly, the approximation residual either decreases monotonically (positive or negative correlation holds) or keeps the current status (irrelevance holds). Thus we not only keep the interpretation of the matrix factorization, but also adapt to the skewed latent factors.

2) Algebraic Representation Model: Having presented the correlation model in pieces, we next represent it in a compact framework. We summarize the correlation model in a pairwise matrix (denoted as a **sign matrix**) $\beta \in \{+1, -1, 0\}^{N \times N}$. The "positive correlation, negative correlation, and irrelevance" choices are transformed to three discrete choices (+1, -1, 0), respectively. Consequently, we transform the correlation model from a combination optimization problem to a linear algebraic problem.

We next represent the k-th approximation residual as follows: $\hat{E}_k = \beta_k \circ \mathbf{X}_k$. The overall estimation amounts to $\hat{D} = \sum_{k=1}^r \hat{E}_k = \sum_{k=1}^r \beta_k \circ X_k$.

The ideal model can be represented as a *recurrent* framework that predicts the residuals layer by layer, as illustrated in Figure 5. Each recurrence aims to find a skewness-aware rank-1 matrix to minimize the current residuals. The model sequentially finds a rank-1 matrix with respect to the residual and a sign matrix according to the correlation types. Then, it combines the rank-1 matrix and the sign matrix to approximate the residual. Finally, the overall estimation is updated and the refreshed residual is forwarded to the next recurrence.

Lemma 1 shows that the ideal model either decreases the approximation residual monotonically or keeps the current status in each layer. The proof is put in the Appendix A.

Lemma 1: Assume that we obtain the perfect sign matrix, for all $k \ge 1$, $\|\mathbf{E}_{k+1}\| \le \|\mathbf{E}_k\|$

Unfortunately, manually deciding the selection choices is only possible for observed RTTs, but impossible for unobserved node pairs. To address this challenge, we next predict the selection choices to make the ideal model practical.

C. Our Work

1) Overview: We realize the ideal model with a residual-learning framework called SMF, which unrolls the recurrent framework to a chained sequence of layers, and trains the model layer by layer similar to the stacking architecture of the deep neural network [25].

Each layer k takes the partially-observed residual as the input, and computes a new rank-1 matrix by maximizing the correlation with the current residual, then combines each entry

of this rank-1 matrix into the current approximation based on the sign matrix, both of which are learnt from partial observations. Next, we calculate the approximation residual at this layer, and finally refresh the residual and forwards that to the next layer until completing the final layer.

Our work improves the matrix factorization in several aspects:

- **Explainability**: SMF progressively finds a rank-1 matrix to best explain the residual in each layer. Further, higher layers' rank-1 matrices are trained based on the residuals to lower layers' rank-1 matrices. Thus there exists no ambiguity for the rank-1 matrices.
- What-if Analysis: SMF optimizes a separate rank-1 matrix for each layer, and sequentially combines them with the sign matrices to produce the residual approximation. Thus SMF enables the operator to test the choice of rank based on the layerwise RTT approximation.
- **Robustness**: SMF formulates the model with respect to the residual in each layer, which is known to be robust to the gradient-vanishing problem, as proved by the ResNet method [24].
- **Modularity**: SMF divides each layer to two conditionally independent optimization problems and combines layerwise estimation into a modular framework.

2) SMF Components: (i) Rank-1 Matrix: As each rank-1 matrix amounts to a product of two vectors [22], [43], we represent the rank-1 matrix \mathbf{X}_k as $\mathbf{X}_k = F_k(G_k)^T$, where F_k and G_k denote two vectors of length N. Moreover, to keep the latent factor be interpretable, we enforce the rank-1 matrix \mathbf{X}_k to be nonnegative. As the observation is incomplete, we formulate a rank-1 matrix factorization problem to estimate the vectors F_k and G_k .

(ii) **Sign Matrix**: We utilize the analogy between the sign matrix estimation and the collaborative filtering problem that recovers missing discrete rating scores (e.g., one to five stars) between a set of clients and a set of goods, and predict the sign matrix β_k with a Maximum Margin Matrix Factorization (MMMF) [15], [38].

We represent pairwise signs with a low-dimensional model. We assign each node a r_s -dimensioned coordinate (\vec{u}_i, \vec{v}_i) and a threshold vector $\vec{\theta}_i$ that serves as boundaries to obtain the discrete signs. We tolerate skewed distributions by incorporating a bias parameter to each node [44]. Each node *i* keeps a bias variable b_i , the coordinate distance Y_{ij} from node *i* to node *j* is defined as the sum of the dot-product result plus the sum of biases:

$$Y_{ij} = \vec{u}_i \vec{v}_j + b_i + b_j \tag{3}$$

We compute the sign from node *i* to node *j* by mapping the coordinate distance $\vec{u}_i \vec{v}_j$ to "-1", 0, "+1" using node *i*'s threshold values $\vec{\theta}_i$:

$$\widehat{\beta}_{ij} = \begin{cases} -1, \vec{u}_i \vec{v}_j \le \vec{\theta}_i(1) \\ 0, \vec{\theta}_i(1) \le \vec{u}_i \vec{v}_j \le \vec{\theta}_i(2) \\ +1, else \end{cases}$$
(4)

3) Example: We provide a simple example to illustrate the SMF algorithm. Note that our purpose here is not to characterize the performance.

FU et al.: SKEWNESS-AWARE MATRIX FACTORIZATION APPROACH FOR MESH-STRUCTURED CLOUD SERVICES



Fig. 6. Four nodes organized in a hierarchical topology.

Given four nodes in Figure 6, the RTT matrix is defined as:

$$D = \begin{bmatrix} 0 & 20 & 80 & 80\\ 20 & 0 & 80 & 80\\ 80 & 80 & 0 & 40\\ 80 & 80 & 40 & 0 \end{bmatrix}$$

We build a sampled RTT matrix to model the probes in the data plane:

$$D' = \begin{bmatrix} 0 & 20 & 0 & 80\\ 20 & 0 & 80 & 0\\ 0 & 80 & 0 & 40\\ 80 & 0 & 40 & 0 \end{bmatrix}$$

Next, we compute a rank-1 matrix with respect to the matrix D': $F_1 = (-6.6791, -6.6791, -7.1096, -7.1096)$, $G_1 = (-6.6791, -6.6791, -7.1096, -7.1096)$. The first two nodes and the last two nodes are identical with each other, respectively, which implies that the rank-1 components captures the global proximity index. Further, F_1 and G_1 are identical to each other, as the RTT matrix is symmetric.

Multiplying F_1 and G_1 yields the rank-1 matrix as follows:

$$X_1 = \begin{bmatrix} 0 & 44.6102 & 47.4854 & 47.4854 \\ 44.6102 & 0 & 47.4854 & 47.4854 \\ 47.4854 & 47.4854 & 0 & 50.5459 \\ 47.4854 & 47.4854 & 50.5459 & 0 \end{bmatrix}$$

Next, we calculate the residual E_1 to the matrix D as follows:

$$E_1 = \begin{bmatrix} 0 & -24.6102 & 32.5146 & 32.5146 \\ -24.6102 & 0 & 32.5146 & 32.5146 \\ 32.5146 & 32.5146 & 0 & -10.5459 \\ 32.5146 & 32.5146 & -10.5459 & 0 \end{bmatrix}$$

Given the partially observed residual matrix E'_1 :

$$E_1' = \begin{bmatrix} 0 & -24.6102 & 0 & 32.5146 \\ -24.6102 & 0 & 32.5146 & 0 \\ 0 & 32.5146 & 0 & -10.5459 \\ 32.5146 & 0 & -10.5459 & 0 \end{bmatrix},$$

we find F_2 and G_2 to the non-negative absolute numbers of E'_1 as follows: $F_2 = (-4.7934, -4.7934, -4.3030, -4.3030)$, $G_2 = (-4.7934, -4.7934, -4.3030, -4.3030)$. The new rank-1 matrix is calculated as follows:

$$X_2 = \begin{bmatrix} 0 & 22.9771 & 20.6263 & 20.6263 \\ 22.9771 & 0 & 20.6263 & 20.6263 \\ 20.6263 & 20.6263 & 0 & 18.5161 \\ 20.6263 & 20.6263 & 18.5161 & 0 \end{bmatrix}$$

We extract the sign matrix to the residual matrix E'_1 :

$$\beta = \begin{bmatrix} 0 & -1 & 0 & 1 \\ -1 & 0 & 1 & 0 \\ 0 & 1 & 0 & -1 \\ 1 & 0 & -1 & 0 \end{bmatrix}$$

based on Algorithm 1, and predict the sign matrix as follows:

$$U:\begin{bmatrix} 0.6074 & -0.3540\\ 0.8816 & -0.1533\\ -1.0691 & -0.2674\\ -0.5347 & 0.5162 \end{bmatrix}, V:\begin{bmatrix} -0.8211 & -0.1208\\ -0.6933 & 0.1867\\ 0.8124 & -0.9372\\ 0.3635 & 0.0986 \end{bmatrix}, \\ \theta:\begin{bmatrix} 0.3265 & 0.3686\\ 0.2530 & 0.2870\\ 0.3905 & 0.5373\\ 0.0705 & 0.0771 \end{bmatrix}$$

and the bias vector (0.1988, 0.1186, -0.0130, 0.6004).

The coordinate distance matrix \hat{Y} can be derived by Eq (3) as:

$$\hat{Y} = \begin{bmatrix} 0 & -0.1698 & 1.0110 & 0.9851 \\ -0.3880 & 0 & 0.9655 & 1.0243 \\ 1.0959 & 0.7969 & 0 & 0.1724 \\ 1.1759 & 1.1861 & -0.3308 & 0 \end{bmatrix}$$



Fig. 7. The diagram of the SMF method.

Next, we infer the sign matrix to \hat{Y} as follows:

$$\hat{\beta} = \begin{bmatrix} 0 & -1 & 1 & 1\\ -1 & 0 & 1 & 1\\ 1 & 1 & 0 & -1\\ 1 & 1 & -1 & 0 \end{bmatrix}$$

Based on the estimated rank-1 matrix and the sign matrix, we obtain the selective combination $X_2 \circ \beta$ as:

$$X_2 \circ \beta = \begin{bmatrix} 0 & -22.9771 & 20.6263 & 20.6263 \\ -22.9771 & 0 & 20.6263 & 20.6263 \\ 20.6263 & 20.6263 & 0 & -18.5161 \\ 20.6263 & 20.6263 & -18.5161 & 0 \end{bmatrix}$$

Plugging this rank-1 approximation to the rank-1 approximation yields a new estimation matrix:

$$\hat{D}_2 = \begin{bmatrix} 0 & 21.6331 & 68.1117 & 68.1117 \\ 21.6331 & 0 & 68.1117 & 68.1117 \\ 68.1117 & 68.1117 & 0 & 32.0298 \\ 68.1117 & 68.1117 & 32.0298 & 0 \end{bmatrix}$$

The estimation matrix \hat{D}_2 outperforms the rank-1 estimation X_1 , which completes the process of a two-layer residualdriven optimization framework.

V. OPTIMIZATION METHODS

Having presented the residual-driven layerwise representation framework, we next present optimization techniques to find the skewness-aware matrix factorization model.

As stated in Section III-A, the monitoring framework consists of a logically centralized controller in the control plane and distributed service-mesh nodes in the data plane. The controller predicts missing measurements in each interval: It aggregates the measurements of an interval as a partially observed matrix, and runs the SMF to estimate missing matrix entries.

A. Optimization Workflow

We summarize the layerwise optimization workflow in Figure 7:

(i) S1: For the first layer, let E_1 denote the partiallyobserved RTT matrix **D**, which consists of aggregated RTT samples of the current interval.

(ii) S2: We compute a rank-1 matrix \mathbf{X}_k (k is set to one for the first layer) to maximize the correlation with the residual \mathbf{E}_k . \mathbf{E}_1 is initialized to the matrix \mathbf{D} for the first layer.

(iii) S3: Next, we compute a sign matrix β_k to the current rank-1 matrix \mathbf{X}_k , by minimizing the approximation error of the residual \mathbf{E}_k .

(iv) S4: We approximate the residual as the dot product of these two matrices: $\hat{\mathbf{E}}_k = \mathbf{X}_k \circ \beta_k$, and update the RTT estimation as follows: $\hat{\mathbf{D}}_k = \hat{\mathbf{D}}_{k-1} + \hat{\mathbf{E}}_k = \sum_{l=1}^k \mathbf{X}_l \circ \beta_l$, where \circ represents the element-wise product operator.

Next, we update the residual for each observed entry (i, j): $\mathbf{E}_{k+1}[i, j] = \mathbf{D}[i, j] - \hat{\mathbf{D}}_k[i, j].$

(v) Let $k \leftarrow k+1$. If k < r, we move to the next layer until completing at the final layer; otherwise, we stop and output the approximation as follows: $\hat{\mathbf{D}} = \sum_{l=1}^{r} \mathbf{X}_{l} \circ \beta_{l}$.



Fig. 8. Graphical model of key variables in the optimization workflow.

B. Loss Function

Based on the optimization workflow, the rank-1 matrix is optimized with respect to the residual from the last layer, while the sign matrix is optimized with respect to the observed sign samples calculated between the residual and the approximated rank-1 matrix. Thus, each layer optimizes two conditionally independent optimization problem, as clearly shown in Figure 8. Thus we separate the optimization problem to decomposed subproblems and design modular optimization workflow, due to the conditional independence between the rank-1 matrix and the sign matrix.

1) Rank-1 Matrix: The input to the rank-1 matrix completion is the residual. We compute the residual \mathbf{E}_k based on the prediction $\hat{\mathbf{D}}_{k-1}$ that is recursively built from the first to the k-1-th layer:

$$\mathbf{E}_{k}\left[i,j\right] = \mathbf{D}\left[i,j\right] - \mathbf{D}_{k-1}\left[i,j\right]$$
(5)

for $(i, j) \in Observed$ samples and $k \geq 2$.

We define a loss function that minimizes the difference between the rank-1 matrix and the absolute-valued residual \mathbf{E}_k as follows:

$$\min \left\| F_k(G_k)^T - |\mathbf{E}_k| \right\|_F^2 \tag{6}$$

We extend this loss function to a regularized objective function so as to avoid the overfitting issue [19], [57]:

$$\min J_{k} = \left\| F_{k}(G_{k})^{T} - |\mathbf{E}_{k}| \right\|_{F}^{2} + \lambda \left\| F_{k}^{T}F_{k} - G_{k}^{T}G_{k} \right\|_{F}^{2}$$
(7)

where λ denotes a regularized factor. The regularization $\|F_k^T F_k - G_k^T G_k\|_F^2$ addresses the scaling ambiguity as $F_k(G_k)^T = (F_k R)(G_k R)^T$ holds for any orthogonal matrix R [57]: Suppose J_k is defined in Eq (7), any stationary point (F_k, G_k) of J_k with $\nabla J(F_k, G_k) = 0$ implies that $F_k^T F_k = G_k^T G$.

2) *Sign Matrix:* The sign matrix is vital to keep the combined prediction monotonically decreasing. The input to find the sign matrix is the signs of observed node pairs.

We first present a simple algorithm to obtain the input β that is consistent with the selective combination rules (presented in Algorithm 1).

Let the sign $\beta[i, j]$ of each observed node pair (i, j) be shifted to the interval $\{1, 2, 3\}$ for ease of training. We formulate a soft-margin classification error $L(Y, \vec{\theta})$ for each

Algorithm	1:	SignRule($\mathbf{E}_k[i, j]$, \mathbf{X}_{ij}): The Rule of	
Calculating	the	Sign for an Observed Node Pair (<i>i</i> , <i>i</i>).	

1	SignRule ($\mathbf{E}_{k}[i, j], \mathbf{X}_{ij}$)
	input : $\mathbf{E}_{k}[i, j]$: Current approximation residual, \mathbf{X}_{ij} :
	rank-one esimtation for pair (i, j)
	output : Sign $\beta_k[i, j]$.
2	if $ \mathbf{E}_{k}[i, j] + \mathbf{X}_{ij}] < \mathbf{E}_{k}[i, j] $ then
3	Set $\beta_k[i,j] = "+1";$
4	else if $ \mathbf{E}_k[i,j] - \mathbf{X}_{ij} < \mathbf{E}_k[i,j] $ then
5	L Set $\beta_k[i, j] = (-1)^n$;
6	else
7	
8	return $\beta_k[i,j]$;

observed node pair (*i*, *j*) [11], [15], [38]:

$$L\left(\vec{u}_{i}, \vec{v}_{j}, b_{i}, b_{j}, \vec{\theta}_{i}\right)$$

$$= \sum_{c=1}^{\beta_{ij}-1} h\left(\vec{u}_{i} \cdot \vec{v}_{j} + b_{i} + b_{j} - \vec{\theta}_{i}(c)\right)$$

$$+ \sum_{c=\beta_{ij}}^{2} h\left(\vec{\theta}_{i}(c) - (\vec{u}_{i} \cdot \vec{v}_{j} + b_{i} + b_{j})\right)$$

$$= \sum_{c=1}^{2} h\left(T_{ij}^{c} [\beta_{ij}] \cdot (\theta_{ic} - (\vec{u}_{i} \cdot \vec{v}_{j} + b_{i} + b_{j}))\right)$$
(8)

where $\theta_{ic} = \vec{\theta}_i(c)$, $T_{ij}^c[\beta_{ij}] = \begin{cases} +1 c \ge \beta_{ij} \\ -1 c < \beta_{ij} \end{cases}$ serves as an indicator function for the sign value β_{ij} , and $h(z) = \begin{cases} \frac{1}{2}-z & z < 0 \\ 0 & z > 1 \end{cases}$ represents a smoothed hinge loss with the $\frac{1}{2}(1-z)^2 \text{ otheriwse}$ derivative as follows: $h'(z) = \begin{cases} -1 & z < 0 \\ 0 & z > 1 \end{cases}$ Thus the loss function allows for gradient based optimization techniques.

Further, We present a regularized loss function with respect to Eq (8) for each node i for robustness against outliers and missing items [11], [15], [38]:

$$g_k(x_i) = \sum_{e \in S_i} L\left(\vec{u}_i, \vec{v}_j, b_i, b_j, \vec{\theta}_i\right) + \lambda\left(\|\vec{u}_i\|_F^2 + \|\vec{v}_i\|_F^2 + \|b_i\|_F^2\right)$$
(9)

where $x_i = \left[\vec{u}_i; \vec{v}_i; b_i; \vec{\theta}_i\right]$, and λ denotes the regularization constant.

C. Algorithms

Each layer has two key optimization building blocks, including the optimization of the rank-1 matrix and that of the sign matrix, both of which belong to the non-convex matrixfactorization problem.

1) Rank-1 Matrix Optimization: First, we seek a rank-1 matrix \mathbf{X}_k to be maximally correlated with the approximation residual. We compute the rank-1 matrix with the stochastic gradient descendent (SGD) method. The SGD converges fast to the vicinity of the optimum regions and helps escape the local minimum for distributed optimization [19]–[21], [27], [28], [42].

FU et al.: SKEWNESS-AWARE MATRIX FACTORIZATION APPROACH FOR MESH-STRUCTURED CLOUD SERVICES

Algorithm	2:	RankOne($ E_k $):	SGD	Based	Rank-One
Matrix Opti	miz	ation			

1 RankOne ($|E_k|$) input : Partially observed residual $|E_k|$, learning rate η_1 , optimization rounds Toutput: F_k , G_k 2 for $t = 1, 2, \dots, T$ do 3 for $i \in 1, 2, \dots, N$ do 4 Compute the partial gradients of $\frac{\partial J}{\partial F_k(i)}$ and $\frac{\partial J}{\partial G_k(i)}$

$$\frac{\partial J}{\partial F_k(i)} = \lambda F_k(i) \left(F_k^T F_k - G_k^T G_k \right) - \sum_{\substack{j \in S_i \\ \partial G_k(i)}} \left(|E_{ij}| - F_k(i)G_k(j) \right) G_k(j) - \lambda G_k(i) \left(F_k^T F_k - G_k^T G_k \right) - \sum_{\substack{j \in S_i \\ j \in S_i}} \left(|E_{ji}| - F_k(j)G_k(i) \right) F_k(j)$$
(10)

Update the coordinate as:

$$F_k(i) = F_k(i) - \eta_1 \frac{\partial J}{\partial F_k(i)}$$

$$G_k(i) = G_k(i) - \eta_1 \frac{\partial J}{\partial G_k}(i)$$
(11)

6 return F_k , G_k ;

Algorithm 3: Optimization Procedures of the SMF Algorithm

1	SMF	(D .	r
	O T T T		

- **input** : Partially observed RTT matrix \mathbf{D} , rank r**output**: Estimated RTT matrix $\hat{\mathbf{D}}$
- 2 Compute a rank-1 matrix $\mathbf{X}_1 = F_1(G_1)^T$ that is maximally correlated with the RTT matrix **D** by calling RankOne(**D**) in Algorithm 2;
- 3 Set the RTT prediction as $\hat{\mathbf{D}}_1 = \mathbf{X}_1$;
- 4 for Integer $k \in [2, r]$ do
- 5 Compute the residual $\mathbf{E}_k = \mathbf{D} \hat{\mathbf{D}}_{k-1}$;
- 6 Find a rank-1 matrix $\mathbf{X}_k = F_k(G_k)^T$ with respect to \mathbf{E}_k by calling RankOne ($|E_k|$) in Algorithm 2;
- 7 Calculate the sign matrix $\beta_{\mathbf{k}}$ by calling SignRule ($\mathbf{E}_{k}[i, j], \mathbf{X}_{ij}$) in Algorithm 1;
- 8 Find a sign matrix $\hat{\beta}_k$ by calling SignMatrix (β_k) (Appendix B);
- 9 Update the RTT prediction $\hat{\mathbf{D}}_k = \hat{\mathbf{D}}_{k-1} + \mathbf{X}_k \circ \hat{\beta}_k$; 10 Set $k \leftarrow k+1$;

11 return
$$\hat{\mathbf{D}} = \sum_{k=1}^{r} F_k (G_k)^T \circ \hat{\beta}_k;$$

Let $F_k(i)$ and $G_k(i)$ denote the row vectors of node i, respectively. We optimize these two row vectors in T rounds. In each round t, we compute the partial gradients of each row vectors with respect to probed targets of node i, and then adjust each row vector towards the approximated negative gradient direction scaled by a learning-rate parameter η_1 , which is automatically chosen using the line search method [31]. We summarize the optimization steps in Algorithm 2.

2) Sign Matrix Optimization: Second, we seek a sign matrix to optimize the correlation types for each node pair between the residual and the current rank-1 matrix. We compute the sign matrix based on the MMMF framework that adaptively penalizes the difference between the estimated signs and the ground-truth ones [38].

We form a vector x_i as the concatenation of the coordinate components of i, i.e., $x_i(s) = \begin{bmatrix} \vec{u_i} ; \vec{v_i} ; \vec{\theta_i} ; b_i \end{bmatrix}$, and adjust each node's vector x_i in T rounds. In each round s ($s \ge 1$), we compute the conjugate direction of the vector x_i with respect to probed targets, and move the vector x_i a small distance in the direction of the conjugate direction scaled by a learning-rate parameter η_2 , which is automatically computed to meet the Strong wolfe line search conditions [9]. The optimization steps are summarized in Appendix B.

3) Putting It All: Algorithm 3 summarizes the overall workflow. First, we find a rank-1 matrix that is maximally correlated with the partially-observed RTT matrix **D** (In *line two*). Next, we set the current approximation based on the rank-1 matrix X_1 (in *line three*). Then, we perform a layerwise training process (from *lines four to ten*), by finding the rank-1 matrix and the sign matrix, and forward the residual to the next loop cycle.

D. Analysis

1) Convergence: SMF generalizes the truncated SVD that recursively finds an approximately optimal rank-1 matrix to the residual. The SVD sums up each rank-1 matrix to form the low-rank prediction, while SMF flexibly combines each rank-1 matrix to be aware of the skewness of node pairs. As a result, the truncated SVD is a special case for SMF.

SMF and ResNet [24] both optimize the residuals in a layerwise approach. ResNet [24] forces each layer to learn the residual feature between the output of this layer and that of the last layer, which efficiently addresses the gradient-vanishing problems for deep neural networks. While SMF not only forces each layer to approximate the residual, but also regularizes the structure of the layer to be aware of skewed latent factors. Accordingly, SMF adapts well to the matrix-factorization context.

Recall that the rank-1 matrix factorization and the sign matrix estimation are two conditionally independent optimization problems. As a result, the problem of proving SMF's convergence is transformed to prove the convergence of the rank-1 matrix factorization and that of the sign matrix in each layer.

(i) Rank-1 Matrix

A point x is a stationary point *iff* its gradient $\nabla J(x)$ is zero. Geometrically, x is a local minimum *iff* x is a stationary point and there exists a neighborhood area NA(x) of the vector x such that $J(z) \ge J(x)$ for any $z \in NA(x)$ [30]. We restate the strict saddle property for twice differentiable functions such as the rank-1 matrix factorization.

Definition 1 ([18], [19]): A twice differentiable function J is (ψ, Ω, ω) -strict saddle, if for any point x at least one condition holds: (i) $\|\nabla J(x)\| \ge \psi$; (ii) $\lambda_{min} (\nabla^2 J(x)) \le -\Omega$; (iii) x is ω -close to the set of local minima. The strict saddle definition implies that for any point x, either it has a large gradient, or it has a negative directional curvature, or it is close to a local minima. Lemma 2 shows that, the loss function J_k defined in Eq (7) satisfies the strict saddle condition.

Lemma 2: Let σ_1^* , σ_r^* denote the largest and the r-th singular values of the residual \mathbf{E}_k . The loss function J_k is $(\epsilon, \Omega(\sigma_r^*), O(\frac{\epsilon}{\sigma_r^*}))$ -strict saddle. The proof of Lemma 2 directly derives from Theorem 4 in

The proof of Lémma 2 directly derives from Theorem 4 in [19]. Accordingly, the rank-1 matrix factorization satisfies that: All local minima are also globally optimal; any saddle point has at least one strictly negative eigenvalue in its Hessian matrix, thus local search methods efficiently find points towards the local minima [18], [19].

Further, the optimization process of the rank-1 matrix with a random initialization converges to or a local minimizer almost surely, as described in Lemma 3 and proved in Appendix C.

Lemma 3: If $\lambda \geq 0$, $\eta < \frac{1}{L_{smooth}}$, where L_{smooth} denotes the smoothness of function J, i.e. $\|\nabla J(x_1) - \nabla J(x_2)\| \leq L \|x_1 - x_2\|$, Algorithm 2 with a random initialization converges to a local minimizer almost surely.

(ii) Sign Matrix

The nonlinear conjugate-gradient optimization converges to zero gradients, as proved in Theorem 3.5 by Liu et al. [9] under the following conditions:

Lemma 4: The gradient of the nonlinear conjugate-gradient method converges to zero, i.e., $\lim_{s\to\infty} \nabla g_k(s) = 0$ given the conditions:

- g(x) is bounded below on the level set = $\{x|g(x) \le g(x_1)\}$, where x_1 denotes the starting point. And in some neighborhood of, g is continuously differentiable, and the gradient is Lipschitz continuous, i.e., for a constant L > 0, such that $||g(x) - g(y)|| \le L ||x-y||$.
- The level set $= \{x | g(x) \le g(x_1)\}$ is bounded.
- Polak-Ribière scalar $\gamma_s \geq 0$.
- Strong wolfe line search conditions: A positive step-length η_2 computed by a line search satisfies that $g(x_s + \eta_2\Lambda x_s) \leq g(x_s) + \rho\eta_2\Delta x_s^T\Lambda x_s$ and $\Delta(x_s + \eta_2\Lambda x_s)^T\Lambda x_s \geq \sigma\Delta x_s^T\Lambda x_s$, for $0 < \rho < \sigma < 1$.
- Descendent condition: $\Delta x_s^T \Lambda x_s < 0.$
- Property (*): There exists constants b > 0 and $\lambda > 0$ for all k, if $|\gamma_s| \le b$, and $||x_{(s)} - x(s-1)|| \le \lambda$, then $|\gamma_s| \le \frac{1}{2b}$.

(iii) **Discussions**

Further, significant efforts such as [18]–[21], [27], [42] have proved that first-order local search optimization algorithms have nice convergence guarantee for general non-convex optimization problems: No spurious local minimum arise in the optimization landscape; further, simple local search methods escape saddle points efficiently and find the exact low-rank matrix from arbitrary starting points with high probability in polynomial time.

2) Robustness: Due to the non-convex nature, some entries of the sign matrix may be incorrectly predicted. Specifically, there exists a small probability that "+1" is flipped to "-1", "-1" to "+1", and "0" to "+1" or "-1". Note that mapping "+1" or "-1" to "0" does not degrade the current

approximation residual. We next bound the failure probability that the sign of any server pair is always incorrect in Theorem 1.

Theorem 1: Let r be the approximation rank and ϵ the expected fraction of incorrectly predicted signs. For any node pair, the expected number of correctly predicted signs amounts to $r(1-\epsilon)$, the variance is $\epsilon r(1-\epsilon)$, and for t > 0, the number of correct estimations is not within $t\sqrt{r(1-\epsilon)}$ from its expectation is at most $\frac{\epsilon}{t^2}$.

The proof is put in the Appendix D. For example, if $\epsilon = 0.2$ and r = 8, the expected number of correct predictions amounts to $\lfloor 8 \times 0.8 \rfloor = 6$, with the variance $0.2 \cdot 8 (1 - 0.2) = 1.28$. Further, note that when a sign is mapped to zero from "+1" or "-1", then the current approximation will be skipped and the residual will not change accordingly.

3) Time Complexity: The time complexity of SMF amounts to the sum of time spent on each layer, finding the rank-1 matrix and the sign matrix.

(i) **Rank-1 matrix**: The time complexity of the SGD algorithm is proportional to the gradient computation. For each node *i*, one gradient calculation takes $O(n_p)$ time, where n_p denotes the set of probing targets. *N* nodes take $O(Nn_p)$ time to compute the gradient for one round. Algorithm 2 needs *T* rounds, requiring $O(TNn_p)$ time.

(ii) **Sign Matrix**: The time complexity of the conjugate-gradient method is linearly proportional to the calculation of the conjugate gradient [15]. Let r_s denote the dimension of the coordinate, it takes $O(r_s n_p)$ time to compute the partial gradient and $O(n_p r_s^2)$ time to derive the conjugate gradient direction for each node *i*. Consequently, N nodes take $O(Nn_p r_s^2)$. SignMatrix algorithm needs T rounds, requiring $O(TNn_p r_s^2)$.

In summary, training a rank-r SMF model needs $O(rTNn_p(1+r_s^2))$ time.

E. Parameter Choices

As the optimization problem is non-convex, it is challenging to choose optimal parameters for the SMF method. For an unknown dataset, we propose an offline approach to decide the number of rank-1 components, the dimension of the sign matrix, and the regularization constant. We fix all but one parameters, and incrementally adjust the parameter until the average relative error does not decrease significantly.

From our empirical experiments, a wide range of parameters achieve similar degrees of the estimation accuracy. As the majority of a RTT matrix is typically captured by several singular values from Figure 4. Therefore, we could select a relatively small numbers of probing targets and rank-1 matrices, as well as the dimension of the sign-matrix approximation.

VI. EVALUATION

Having presented the optimization algorithms, we next systematically evaluate the performance of our method with state-of-the-art methods on real-world data sets.

A. Experimental Setup

We built a trace-driven controller-agent structured simulator. We reuse the data sets introduced in Section III-C.1 for the simulator. The simulator parses the trace for continuous network latency monitoring. During a slice, each agent randomly samples a small set of nodes as probing targets, then submits network latency values to these targets to the controller. The controller collects probed samples from all agents, then initializes parameters for each node in the continuous vector space, and adjusts parameters based on optimization rules until reaching the local minimum. Afterward, the controller predicts network latency values for missing node pairs.

Performance metric: We quantify the performance of unobserved entries using *relative error* that is widely used for network latency prediction studies. It is defined as the ratio between the relative error and the ground-truth value $\frac{|\mathbf{D}_{ij} - \hat{\mathbf{D}}_{ij}|}{\mathbf{D}_{ij}}$ for each unobserved node pair (i, j). For each setting, the simulation runs in ten times. The results reported show the average relative error.

Default parameters: We choose the default parameters for SMF to balance the diminishing returns of the expected relative error and the computational cost, based on the sensitivity analysis in Section VI-D. Specifically, we set the approximation rank r to 8, the number of probing targets n_p to 32, the bias MMMF coordinate dimension r_s to 16, and the regularized parameter λ to one. Our parameter choices yield sparse matrices. For example, for a 99-by-99 matrix, setting the number of probing targets to 32 implies that 32% of matrix entries are observed, while for a 490-by-490 matrix, it implies that only 6.5% of matrix entries are observed.

We run experiments on a MacBook-Pro Intel Core i7 with Quad-core and 16 GB memory.

B. Comparison Results

First, we compare the relative error of our method with state-of-the-art methods. As we mentioned in the related work, there are many methods proposed in the literature. We choose four baseline algorithms covering the recent well-known approaches: (i) Vivaldi [8]: predicts RTTs in an Euclidean coordinate system and trains the coordinates via a spring-force field simulation. Further, Vivaldi incorporates a height constant to model the access-link portion from end hosts to the edge. (ii) DMFSGD [31]: predicts RTTs in a matrix factorization model and trains the factorized matrices based on the SGD method. (iii) Distance feature decomposition (denoted as DFD in the plot) [32]: predicts RTTs via the product of a low-rank matrix and a complete weight matrix. (iv) OR1MP [43]: recovers an incomplete matrix by incrementally generating a rank-one basis matrix by the SVD method and linearly combining this matrix to the current approximation in a closed form.

For Vivaldi, DMFSGD and OR1MP, we directly downloaded authors' implementations, while for DFD, we implemented the algorithm based on [32]. For fair comparison, all methods use the same number of probes, and set the same coordinate dimension and the identical set of probing targets. We set the regularized parameters based on the recommended configuration of each method.



Fig. 9. The CCDFs of relative errors for different methods. We set the number of probes per node to 32. The x-axis is in log-scale.



Fig. 10. The convergence of the MAE values at each round. We plot the average MAE and the standard deviations.

We plot the distributions of the relative errors. We set the number of probing targets to 32 and the coordinate dimension to eight. We compute the CCDF of relative errors for all methods. From Figure 9, SMF consistently outperforms the other methods in three data sets, while DFD, DMFSGD and Vivaldi are less stable than SMF. OR1MP is less accurate than SMF, since OR1MP combines a set of rank-one matrices with a weight vector, which is insufficient to account for skewed latent factors.

C. Convergence Analysis

Having shown the effectiveness of our method, we next evaluate the convergence of SMF's performance.

1) Sign-Matrix Prediction: First, we evaluate the convergence of the sign-matrix prediction process, as it determines the convergence of the selective combination procedure. We quantify the difference between the estimated sign matrix and the ground-truth result based on the normalized mean absolute error (MAE):

$$\frac{\sum_{(i,j)\in\text{missing entries}} \left| \widehat{Y}_{ij} - Y_{ij} \right|}{\sum_{(i,j)\in\text{missing entries}Y_{ij}}}$$

where Y_{ij} represents the ground-truth sign from *i* to *j*, and \hat{Y}_{ij} denotes the estimated sign.

In Figure 10, we plot the average MAE values of the coordinates at each round. We see that the MAE values monotonically decrease towards the local minimum with increasing numbers of rounds. Further, the MAE values generally increase as we consider more rank-1 matrix components, while most incorrect estimations are either mapping "+1" or "-1" to "0", since most incorrect signs are adjacent to the correct values. Accordingly, the estimation accuracy is not affected by the sign "0", as it implies that we skip the current rank-1 component entry, therefore,

2) Approximation Residual: Having shown the convergence of the sign matrix prediction, we next evaluate the relative error of approximation residuals as we combine more rank-one matrices to the approximation. In Figure 11, we see that



Fig. 11. The average relative error as a function of the number of combined rank-one matrices.



Fig. 12. Sensitivity of the approximation rank as we increase the rank from 4 to 20.



Fig. 13. Sensitivity of the number of probing targets as we increase the number from 16 to 48.

the RTT matrix can be compactly captured with two to three rank-one matrices. The average relative error decreases significantly as we combine the second and the third rank-one matrices, while combining more matrices marginally decreases the average relative error.

D. Sensitivity Results

Having presented the convergence, we next test the sensitivity of our method. We fix all but one parameters to default values, and vary the parameter configuration. We report the average relative errors and the standard deviation.

1) Approximation Rank: First, we compute the distributions of the relative errors as we increase the approximation rank r from four to twenty. In Figure 12, we see that eight to twelve rank-1 components are enough to obtain a reasonably accurate estimation. Further, the median, the 15-th, and the 75-th percentiles of relative errors remain steady on the Planetlab dataset, since the PlanetLab dataset is approximately low-rank. Comparatively, the relative errors on the Seattle and RIPE dataset decrease significantly when we increase the rank from 4 to 8, as these data sets have longer tails of singular values than the PlanetLab dataset.

2) Number of Measurements for Each Node: Next, we evaluate the choice of the number of targets. We compute the relative errors for each setting. From Figure 13, the average relative error decreases progressively, and 32 probing targets suffice to obtain relatively accurate results.

3) Biased-MMMF Dimension: Next, we vary the choice of the coordinate dimension for predicting the sign matrix



Fig. 14. Sensitivity of the bias-MMMF dimension as we increase the coordinate dimension from 4 to 20.



Fig. 15. Sensitivity of the regularization parameter λ as we increase the regularized parameter λ from 0.4 to 2.



Fig. 16. The average relative error and the standard deviation of the found relay as a function of the number of targeting detouring hosts. For a set of nodes (T_1, \ldots, T_L) of size ϖ that need detour routing, we define the optimal relay as the one that minimizes the average network latency: $\frac{1}{\varpi} \sum_{j=1}^{\varpi} \left(d_{T_jR_O} + d_{R_OT_j} \right)$ We quantify the detouring performance for each node pair by comparing the optimal relay selected from the ground-truth network latency matrix and that from the estimated all-pair network latency matrix: $\frac{\left|\sum_{j=1}^{\varpi} \left(d_{T_jR_f} + d_{R_fT_j} \right) - \sum_{j=1}^{\varpi} \left(d_{T_jR_O} + d_{R_OT_j} \right) \right|}{\sum_{j=1}^{\varpi} \left(d_{T_jR_f} + d_{R_fT_j} \right) - \sum_{j=1}^{\varpi} \left(d_{T_jR_O} + d_{R_OT_j} \right)}$, where R_f and

 R_O denote the estimated relay and the optimal relay respectively.

and plot the relative-error distribution. In Figure 14, setting the dimension to 12 to 16 is enough to obtain a relatively high degree of estimation. The median relative error decreases marginally on the PlanetLab, but decrease progressively on the Seattle and RIPE datasets, which also holds for the 15-th and 75-th percentiles of the relative errors.

4) Regularization λ : Next, we study the performance sensitivity with respect to the regularized parameter λ that controls the extent of the regularization. We vary the regularized parameter from 0.4 to 2 and compute the relative error distribution for each setting. From Figure 15, the relative error keeps steady as we change the regularized parameter, thus the prediction is less sensitive to the regularization than other parameters.

Summary: From the sensitivity analysis, a wide range of parameters yield similar accuracy. Therefore, we choose modest parameters in order to trade off well between the accuracy and the computation complexity.

E. Use Case

Having evaluated the performance of SMF, we next evaluate the performance gains for selecting detouring routing nodes that act as proxies to forward packets for end hosts. Given a group of clients that need detour routing via a relay, we choose a server with the minimal average RTT value towards this group of clients. Figure 16 plots the variations of the average relative errors of the found relay as we vary the number of relayed hosts. We can see that on the PlanetLab dataset, our method and DMFSGD obtain the most accurate relays, since the PlanetLab dataset can be well approximated via a low-rank model, while on the Seattle and RIPE datasets, our method significantly outperforms the other methods.

VII. CONCLUSION

Monitoring pairwise RTT status scalably and accurately is vital for network troubleshooting and performance management. Existing matrix factorization based methods overcome the scaling limitations, but could not truthfully capture the skewed distributions. We propose a skewness-aware matrix factorization method named SMF to learn latent factors for different node pairs. We incrementally combine a rank-one matrix weighted by their correlations with the current approximation residual in a scalable residual-learning framework. Extensive experiments over real-world data sets show that SMF significantly improves the relative error by a factor of 0.2x to 10x, converges fast and stably.

Although this paper focuses on a cloud measurement architecture, we can decompose the representation model to decentralized components. First, a rank-1 matrix can be decomposed to separable models with respect to each node. Second, the sign matrix model can be decomposed to separate coordinates, as already proven in previous studies in [14], [15]. After we decompose the representation model, we may optimize tuples a decentralized procedure, which is left as future work.

REFERENCES

- Linux PTP Project. (2016). [Online]. Available: http://nwtime.org/ projects/linuxptp/
- [2] NTP. (2016). [Online]. Available: http://www.ntp.org
- [3] Conduit Service Mesh. (2018). [Online]. Available: https://conduit.io
- [4] Istio Service Mesh. (2018). [Online]. Available: https://istio.io
- [5] E. J. Candès and B. Recht, "Exact matrix completion via convex optimization," *Commun. ACM*, vol. 55, no. 6, pp. 111–119, 2012.
- [6] J. Cappos, I. Beschastnikh, A. Krishnamurthy, and T. Anderson, "Seattle: A platform for educational cloud computing," in *Proc. SIGCSE*, 2009, pp. 111–115.
- [7] Y. Chen, L. Qiu, Y. Zhang, G. Xue, and Z. Hu, "Robust network compressive sensing," in *Proc. MobiCom*, 2014, pp. 545–556.
- [8] F. Dabek, R. Cox, F. Kaashoek, and R. Morris, "Vivaldi: A decentralized network coordinate system," in *Proc. ACM SIGCOMM*, 2004, pp. 15–26.
- [9] Y. Dai et al. "Convergence properties of nonlinear conjugate gradient methods," SIAM J. Optim., vol. 10, no. 2, pp. 345–358, 2000.
- [10] J. Dean and L. A. Barroso, "The tail at scale," Commun. ACM, vol. 56, no. 2, pp. 74–80, Feb. 2013.
- [11] D. DeCoste, "Collaborative prediction using ensembles of maximum margin matrix factorizations," in *Proc. ICML*, 2006, pp. 249–256.
- [12] C. Ding, T. Li, and M. I. Jordan, "Convex and semi-nonnegative matrix factorizations," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 1, pp. 45–55, Jan. 2010.
- [13] Y. Fu, E. Biersack, Y. Fu, and E. Biersack, "MCR: Structure-aware overlay-based latency-optimal greedy relay search," *IEEE/ACM Trans. Netw.*, vol. 25, no. 5, pp. 3016–3029, Oct. 2017.
- [14] Y. Fu and Y. Wang, "iRank: Supporting proximity ranking for peer-topeer applications," in *Proc. IEEE ICPADS*, Dec. 2009, pp. 836–841.
- [15] Y. Fu, Y. Wang, and E. Biersack, "A general scalable and accurate decentralized level monitoring method for large-scale dynamic service provision in hybrid clouds," *Future Gener. Comput. Syst.*, vol. 29, no. 5, pp. 1235–1253, 2013.

- [16] Y. Fu, Y. Wang, and E. Biersack, "HybridNN: An accurate and scalable network location service based on the inframetric model," *Future Gener. Comput. Syst.*, vol. 29, no. 6, pp. 1485–1504, 2013.
- [17] Y. Fu and X. Xiaoping, "Self-stabilized distributed network distance prediction," *IEEE/ACM Trans, Netw.*, vol. 25, no. 1, pp. 451–464, Feb. 2017.
- [18] R. Ge, F. Huang, C. Jin, and Y. Yuan, "Escaping from saddle points— Online stochastic gradient for tensor decomposition," in *Proc. COLT*, 2015, pp. 797–842.
- [19] R. Ge, C. Jin, and Y. Zheng, "No spurious local minima in nonconvex low rank problems: A unified geometric analysis," in *Proc. ICML*, 2017, pp. 1233–1242.
- [20] R. Ge, J. D. Lee, and T. Ma, "Matrix completion has no spurious local minimum," in *Proc. NIPS*, 2016, pp. 2973–2981.
- [21] R. Ge and T. Ma, "On the optimization landscape of tensor decompositions," in *Proc. NIPS*, 2017, pp. 3656–3666.
- [22] G. H. Golub and C. F. V. Loan, *Matrix Computations*, 3rd ed. Baltimore, MD, USA: The Johns Hopkins Univ. Press, 1996.
- [23] C. Guo *et al.*, "Pingmesh: A large-scale system for data center network latency measurement and analysis," in *Proc. SIGCOMM*, 2015, pp. 139–152.
- [24] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. CVPR*, Jun. 2016, pp. 770–778.
- [25] G. E. Hinton and R. R. Salakhutdinov, "Reducing the dimensionality of data with neural networks," *Science*, vol. 313, no. 5786, pp. 504–507, 2006.
- [26] J. Jiang et al., "Via: Improving Internet telephony call quality using predictive relay selection," in Proc. SIGCOMM, 2016, pp. 286–299.
- [27] C. Jin, S. M. Kakade, and P. Netrapalli, "Provable efficient online matrix completion via non-convex stochastic gradient descent," in *Proc. NIPS*, 2016, pp. 4520–4528.
- [28] J. D. Lee *et al.*, "First-order methods almost always avoid strict saddle points," *Math. Program.*, vol. 176, nos. 1–2, pp. 311–337, 2017.
- [29] D. Li, J. Cao, X. Lu, and K. C. C. Chen, "Efficient range Query processing in peer-to-peer systems," *IEEE Trans. Knowl. Data Eng.*, vol. 21, no. 1, pp. 78–91, Jan. 2009.
- [30] X. Li *et al.*, "Symmetry, saddle points, and global optimization landscape of nonconvex matrix factorization," *IEEE Trans. Inf. Theory*, vol. 65, no. 6, pp. 3489–3514, Jun. 2019.
- [31] Y. Liao, W. Du, P. Geurts, and G. Leduc, "DMFSGD: A decentralized matrix factorization algorithm for network distance prediction," *IEEE/ACM Trans. Netw.*, vol. 21, no. 5, pp. 1511–1524, Oct. 2013.
- [32] B. Liu, D. Niu, Z. Li, and H. V. Zhao, "Network latency prediction for personal devices: Distance-feature decomposition from 3D sampling," in *Proc. INFOCOM*, Apr./May 2015, pp. 307–315.
- [33] H. V. Madhyastha et al., "iPlane: An information plane for distributed services," in Proc. USENIX OSDI, 2006, pp. 367–380.
- [34] Y. Mao, L. K. Saul, and J. M. Smith, "IDES: An Internet distance estimation service for large networks," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 12, pp. 2273–2284, Dec. 2006.
- [35] T. S. E. Ng and H. Zhang, "Predicting Internet network distance with coordinates-based approaches," in *Proc. INFOCOM*, Jun. 2002, pp. 170–179.
- [36] A. Nikravesh, H. Yao, S. Xu, D. Choffnes, and Z. M. Mao, "Mobilyzer: An open platform for controllable mobile network measurements," in *Proc. MobiSys*, 2015, pp. 389–404.
- [37] Y. C. Pati, R. Rezaiifar, and P. S. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in *Proc. ACSSC*, vol. 1, Nov. 1993, pp. 40–44.
- [38] J. D. M. Rennie and N. Srebro, "Fast maximum margin matrix factorization for collaborative prediction," in *Proc. ICML*, 2005, pp. 713–719.
- [39] J. L. Rodgers and W. A. Nicewander, "Thirteen ways to look at the correlation coefficient," *Amer. Statistician*, vol. 42, no. 1, pp. 59–66, 1988.
- [40] S. Savage *et al.*, "Detour: A case for informed Internet routing and transport," *IEEE Micro*, vol. 19, no. 1, pp. 50–59, Jan./Feb. 1999.
- [41] A.-J. Su, D. R. Choffnes, A. Kuzmanovic, and F. E. Bustamante, "Drafting behind Akamai (travelocity-based detouring)," in *Proc. SIGCOMM*, 2006, pp. 435–446.
- [42] R. Sun and Z.-Q. Luo, "Guaranteed matrix completion via non-convex factorization," *IEEE Trans. Inf. Theory*, vol. 62, no. 11, pp. 6535–6579, Nov. 2016.
- [43] Z. Wang et al., "Orthogonal rank-one matrix pursuit for low rank matrix completion," SIAM J. Sci. Comput., vol. 37, no. 1, pp. A488–A514, 2015.

14

- [44] M. Weimer, A. Karatzoglou, and A. Smola, "Improving Maximum Margin Matrix Factorization," *Mach. Learn.*, vol. 72, pp. 263–276, Sep. 2008.
- [45] F. M. F. Wong, C. Joe-Wong, S. Ha, Z. Liu, and M. Chiang, "Improving user QoE for residential broadband: Adaptive traffic management at the network edge," in *Proc. IWQoS*, Jun. 2015, pp. 105–114.
- [46] K. Xie *et al.*, "Accurate recovery of Internet traffic data under variable rate measurements," *IEEE/ACM Trans. Netw.*, vol. 26, no. 3, pp. 1137–1150, Jun. 2018.
- [47] K. Xie *et al.*, "Accurate recovery of Internet traffic data: A sequential tensor completion approach," *IEEE/ACM Trans. Netw.*, vol. 26, no. 2, pp. 793–806, Apr. 2018.
- [48] K. Xie *et al.*, "Sequential and adaptive sampling for matrix completion in network monitoring systems," in *Proc. INFOCOM*, Apr./May 2015, pp. 2443–2451.
- [49] R. Xu and D. Wunsch, II, "Survey of clustering algorithms," *IEEE Trans. Neural Netw.*, vol. 16, no. 3, pp. 645–678, May 2005.
- [50] C. Yu *et al.*, "Software-defined latency monitoring in data center networks," in *Proc. PAM*, 2015, pp. 360–372.
- [51] Y. Zhang, D. Li, and L. Liu, "Leveraging glocality for fast failure recovery in distributed RAM storage," *Trans. Storage*, vol. 15, no. 1, pp. 3:1–3:24, 2019.
- [52] Y. Zhang et al., "CSR: Classified source routing in distributed networks," IEEE Trans. Cloud Comput., vol. 6, no. 2, pp. 464–477, Apr./Jun. 2018.
- [53] Y. Zhang, M. Roughan, W. Willinger, and L. Qiu, "Spatiotemporal Compressive Sensing and Internet Traffic Matrices," in *Proc. SIGCOMM*, 2009, pp. 267–278.
- [54] J. Zhu, P. He, Z. Zheng, and M. R. Lyu, "Online QoS prediction for runtime service adaptation via adaptive matrix factorization," *IEEE Trans. Parallel Distrib. Syst.*, vol. 28, no. 10, pp. 2911–2924, Oct. 2017.
- [55] R. Zhu, B. Liu, D. Niu, Z. Li, and H. V. Zhao, "Network latency estimation for personal devices: A matrix completion approach," *IEEE/ACM Trans. Netw.*, vol. 25, no. 2, pp. 724–737, Apr. 2017.
- [56] R. Zhu, D. Niu, and Z. Li, "Robust Web service recommendation via quantile matrix factorization," in *Proc. INFOCOM*, May 2017, pp. 1–9.
- [57] Z. Zhu, Q. Li, G. Tang, and M. B. Wakin, "Global optimality in low-rank matrix optimization," *IEEE Trans. Signal Process.*, vol. 66, no. 13, pp. 3614–3628, Jul. 2018.



Pere Barlet-Ros received the M.Sc. and Ph.D. degrees in computer science from the Universitat Politècnica de Catalunya (UPC) in 2003 and 2008, respectively. He is currently an Associate Professor with the Computer Architecture Department, UPC BarcelonaTech, and alao a Co-Founder of Talaia Networks, a University spin-off that develops innovative network monitoring products. His research interests are in the fields of network monitoring, traffic classification, and anomaly detection.



Chun Huang received the M.S and PhD. degree from the National University of Defense Technology (NUDT), China, in 1997 and 2007, respectively. She is currently a Professor with NUDT. Her research interests include advanced compiler, parallel programming environment, and embedded system.



Zhen Huang received the Ph.D. degree from the National University of Defense Technology, China (NUDT), in 2012. He is currently an Associate Professor with NUDT. His research interests include distributed computing and big data.



Yongquan Fu received the M.S. and Ph.D. degrees in computer science and technology from the National University of Defense Technology (NUDT), in 2007 and 2012, respectively. He is currently an Associate Professor with NUDT. His research interests include network machine learning and distributed systems.



Siqi Shen received the B.S. and M.S. degrees from the National University of Defense Technology (NUDT), China, in 2007 and 2009, respectively, and the Ph.D. degree from the Delft University of Technology, The Netherlands, in 2015. He is currently an Assistant Professor with NUDT. His research interests include computer network and data mining.



Dongsheng Li received the B.Sc. and Ph.D. degrees (Hons.) in computer science from the College of Computer Science, National University of Defense Technology, Changsha, China, in 1999 and 2005, respectively. He is currently a Full Professor with NUDT. His research interests include distributed computing, cloud computing, computer network, and large-scale data management. He received the Prize of the National Excellent Doctoral Dissertation of China by Ministry of Education of China in 2008.



Huayou Su received the Ph.D. degree from the National University of Defense Technology (NUDT), China, in 2014. He is currently an Assistant Professor with NUDT. His research interests include high performance computing, GPU programming, and parallel computing.