Machine Learning-based Algorithms to Infer End-to-End Network Performance Matrices

Guy Leduc
with Yongjun Liao, Wei Du and Pierre Geurts

Research Unit in Networking (RUN)
University of Liège, Belgium

INFORMS Telecommunications, Lisbon, March 4, 2014

End-to-End Network Performance

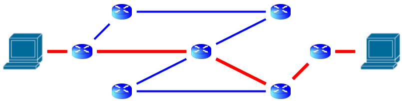

Definition

the performance of a network path linking two end systems

Metrics

- round-trip time (RTT), or one-way delay (OWD)
- available bandwidth (ABW)
- ...

Internet Services rely on End-to-end Performance Measurements

Content Delivery Networks

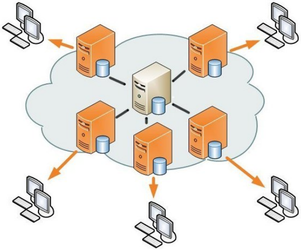

User requests are directed to nearby and/or well-connected servers.

P2P Overlay Networks

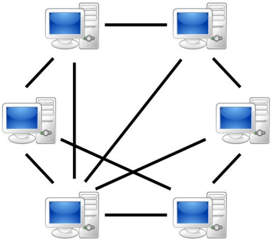

Peers fetch data from nearby and/or well-connected peers.

Intelligent Peer Selection

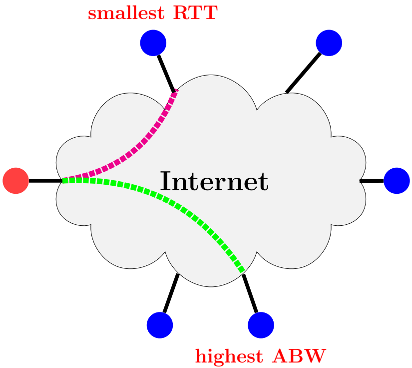

P2P Applications

- reduce cross-domain traffic
- improve download rate

Network Performance Acquisition

How to acquire network performance on large networks?

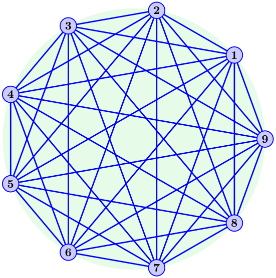

full-mesh active measurements

n nodes $\Rightarrow o(n^2)$ measurements

Network Performance Acquisition

How to acquire network performance on large networks?

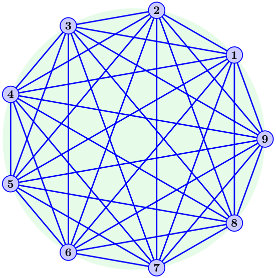

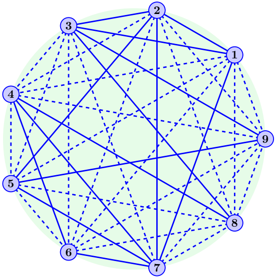

full-mesh active measurements

network performance prediction

n nodes $\Rightarrow o(n^2)$ measurements

n nodes $\Rightarrow$ measurements $\ll o(n^2)$

Learning to Predict Network Performance

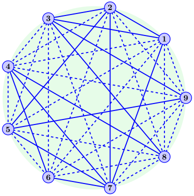

Statistical Inference by Machine Learning

Questions and Answers

- Q: Which model is suitable?
- A: matrix completion by matrix factorization
- Q: Which and how many paths have to be monitored?
- A: a few randomly selected paths ($1 \sim 2\%$ for 2500 nodes)

Université de Liège

Learning to Predict Network Performance

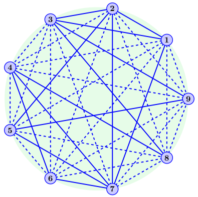

Statistical Inference by Machine Learning

Questions and Answers

- Q: Which model is suitable?
- A: matrix completion by matrix factorization
- Q: Which and how many paths have to be monitored?
- A: a few randomly selected paths ($1 \sim 2\%$ for 2500 nodes)

Learning to Predict Network Performance

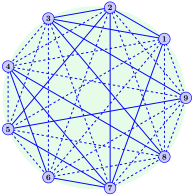

Statistical Inference by Machine Learning

Questions and Answers

- Q: Which model is suitable?
- A: matrix completion by matrix factorization
- Q: Which and how many paths have to be monitored?
- A: a few randomly selected paths ($1 \sim 2\%$ for 2500 nodes)

Université de Liège

Related Work

Tomography-based Approaches

- TOM Chen et al. SIGCOMM 2004
- Network Kriging Chua et al. JSAC 2006
- NetQuest Song et al. SIGMETRICS 2006

Model-based Approaches

- Euclidean Embedding
  - GNP Ng et al. TON 2002
  - Vivaldi Dadeck et al. SIGCOMM 2004
- Matrix Factorization
  - IDES Mao et al. JSAC 2005

Related Work: Tomography-Based Approaches

Idea

Infer link performance from a few path measurements.

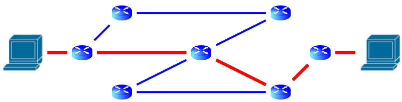

Related Work: Tomography-Based Approaches

Idea

Infer link performance from a few path measurements.

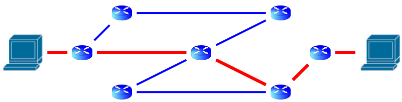

Limitations

- require routing information
- only applicable to additive metrics (RTT, packet loss rate)

Related Work: Model-Based Approaches

Related Work: Model-Based Approaches

Euclidean Embedding

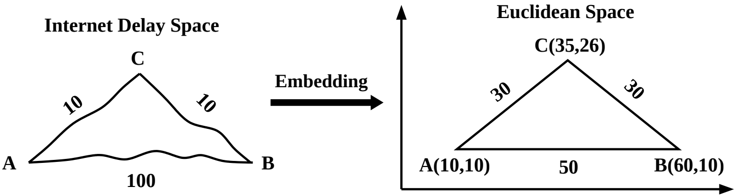

Related Work: Model-Based Approaches

Euclidean Embedding

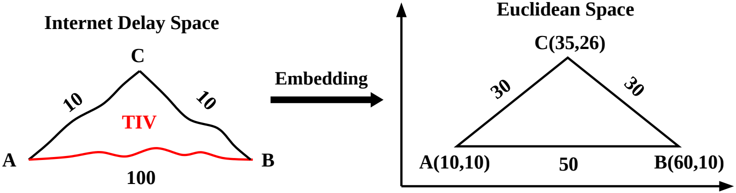

$\overline{AB} > \overline{AC} + \overline{CB} \Rightarrow$ Triangle Inequality Violations

Limitations:

- subject to geometric constraints (symmetry, triangle inequalities)

Related Work: Model-Based Approaches

Euclidean Embedding

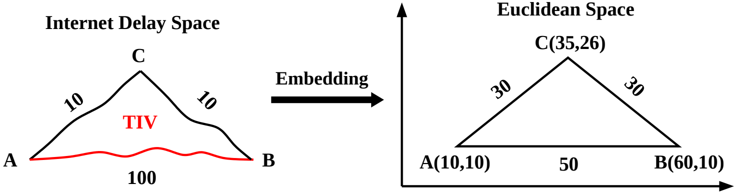

$\overline{AB} > \overline{AC} + \overline{CB} \Rightarrow$ Triangle Inequality Violations

Limitations:

- subject to geometric constraints (symmetry, triangle inequalities)
- only applicable to additive metrics (typically RTT)

Related Work: Model-Based Approaches

Matrix Factorization

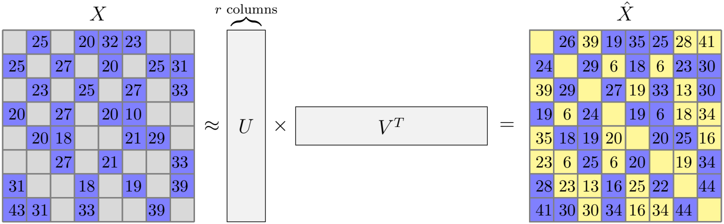

Related Work: Model-Based Approaches

Matrix Factorization

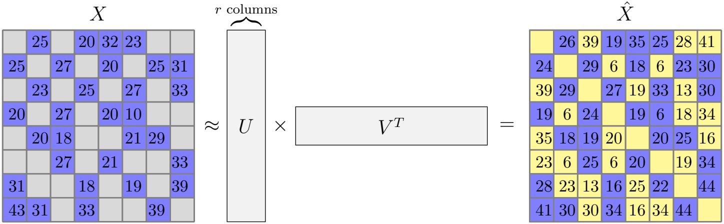

Advantages:

- no routing information
- no geometric constraints
- also applicable to non-additive metrics (available bandwidth)

Contributions

Learning to predict end-to-end network performance

1. Formulation as Matrix Completion

network performance prediction as matrix completion

2. Decentralized Solution

decentralized matrix factorization by stochastic gradient descent

3. Qualitative Representations of Network Performance

- represent network performance by binary classes
- represent network performance by ordinal ratings

Université de Liège

Contributions

Learning to predict end-to-end network performance

1. Formulation as Matrix Completion

network performance prediction as matrix completion

2. Decentralized Solution

decentralized matrix factorization by stochastic gradient descent

3. Qualitative Representations of Network Performance

- represent network performance by binary classes
- represent network performance by ordinal ratings

Université de Liège

Contributions

Learning to predict end-to-end network performance

1. Formulation as Matrix Completion

network performance prediction as matrix completion

2. Decentralized Solution

decentralized matrix factorization by stochastic gradient descent

3. Qualitative Representations of Network Performance

- represent network performance by binary classes
- represent network performance by ordinal ratings

Université de Liège

Contributions

Learning to predict end-to-end network performance

1. Formulation as Matrix Completion

network performance prediction as matrix completion

2. Decentralized Solution

decentralized matrix factorization by stochastic gradient descent

3. Qualitative Representations of Network Performance

- represent network performance by binary classes
- represent network performance by ordinal ratings

Université de Liège

Contributions

Learning to predict end-to-end network performance

1. Formulation as Matrix Completion

network performance prediction as matrix completion

2. Decentralized Solution

decentralized matrix factorization by stochastic gradient descent

3. Qualitative Representations of Network Performance
   - represent network performance by binary classes
   - represent network performance by ordinal ratings

Université de Liège

A Matrix Completion View

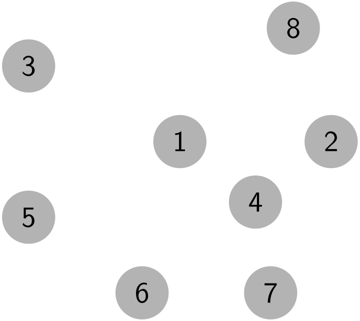

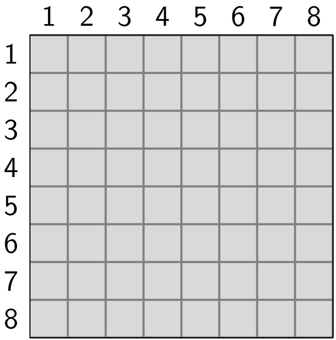

A Matrix Completion View

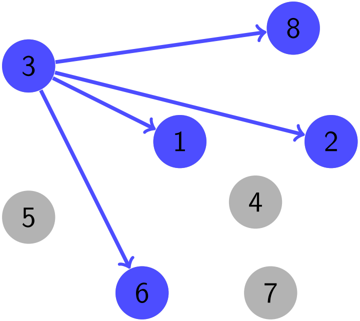

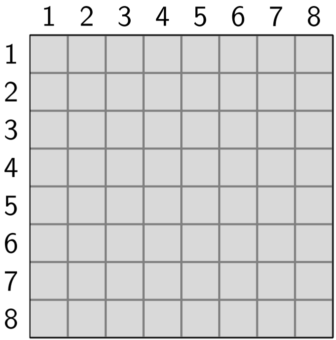

A Matrix Completion View

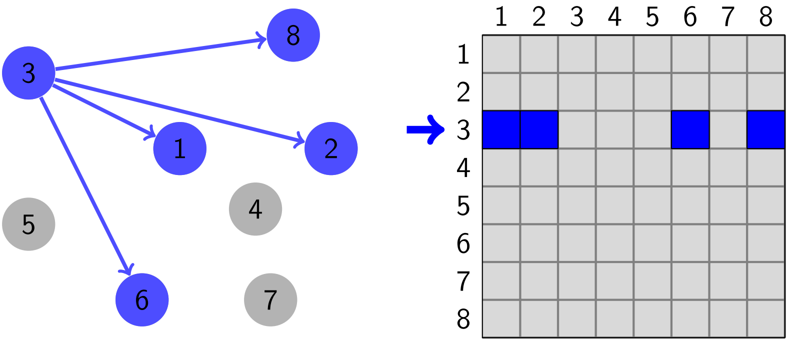

A Matrix Completion View

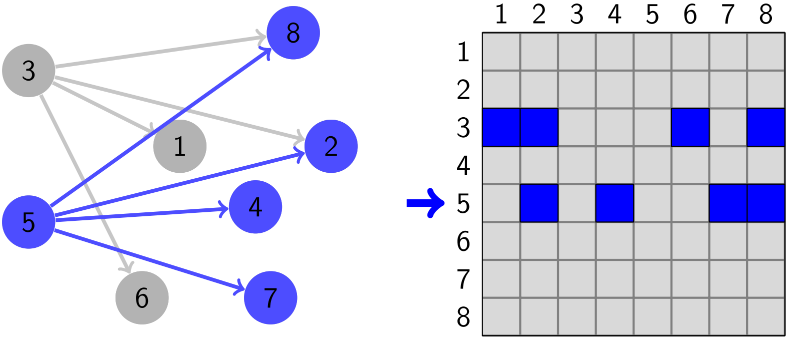

A Matrix Completion View

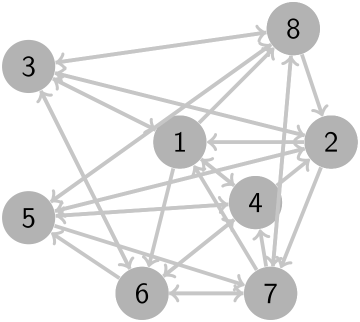

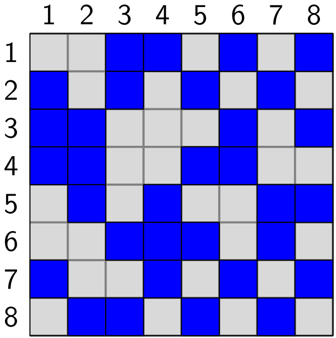

A Matrix Completion View

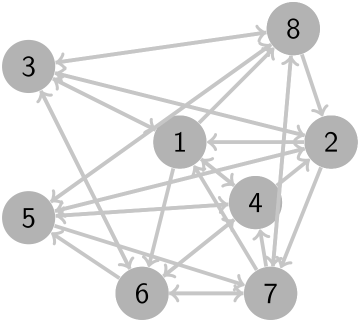

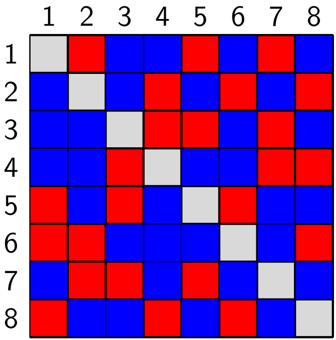

Connections to Recommender Systems

	*movie*1	*movie*2	*movie*3	*movie*4	*movie*5	*movie*6
*user*1	5	3	4	1	?	2
*user*2	5	3	4	1	5	?
*user*3	5	?	4	1	5	3
*user*4	1	3	2	5	1	4
*user*5	4	?	4	4	4	?

Connections to Recommender Systems

	*movie*1	*movie*2	*movie*3	*movie*4	*movie*5	*movie*6
*user*1	5	3	4	1	?	2
*user*2	5	3	4	1	5	?
*user*3	5	?	4	1	5	3
*user*4	1	3	2	5	1	4
*user*5	4	?	4	4	4	?

Resemblance, analogy

- Network nodes are "users".
- Network performance is a "preference" measure.
- Peer selection is a "friend" recommendation task!

Université de Liège

Connections to Recommender Systems

	*movie*1	*movie*2	*movie*3	*movie*4	*movie*5	*movie*6
*user*1	5	3	4	1	?	2
*user*2	5	3	4	1	5	?
*user*3	5	?	4	1	5	3
*user*4	1	3	2	5	1	4
*user*5	4	?	4	4	4	?

Resemblance, analogy

- Network nodes are "users".
- Network performance is a "preference" measure.
- Peer selection is a "friend" recommendation task!

Why is matrix completion possible?

Connections to Recommender Systems

	*movie*1	*movie*2	*movie*3	*movie*4	*movie*5	*movie*6
*user*1	5	3	4	1	?	2
*user*2	5	3	4	1	5	?
*user*3	5	?	4	1	5	3
*user*4	1	3	2	5	1	4
*user*5	4	?	4	4	4	?

Resemblance, analogy

- Network nodes are "users".
- Network performance is a "preference" measure.
- Peer selection is a "friend" recommendation task!

Why is matrix completion possible?

Université de Liège

Connections to Recommender Systems

	*movie*1	*movie*2	*movie*3	*movie*4	*movie*5	*movie*6
*user*1	5	3	4	1	?	2
*user*2	5	3	4	1	5	?
*user*3	5	?	4	1	5	3
*user*4	1	3	2	5	1	4
*user*5	4	?	4	4	4	?

Resemblance, analogy

- Network nodes are "users".
- Network performance is a "preference" measure.
- Peer selection is a "friend" recommendation task!

Why is matrix completion possible?

Connections to Recommender Systems

	*movie*1	*movie*2	*movie*3	*movie*4	*movie*5	*movie*6
*user*1	5	3	4	1	?	2
*user*2	5	3	4	1	5	?
*user*3	5	?	4	1	5	3
*user*4	1	3	2	5	1	4
*user*5	4	?	4	4	4	?

Resemblance, analogy

- Network nodes are "users".
- Network performance is a "preference" measure.
- Peer selection is a "friend" recommendation task!

Why is matrix completion possible?

Why is Matrix Completion Possible?

Feasibility

- Matrix entries are correlated.
- The correlations induce low rank.
- $n \times n$ matrix of rank $r < n$
  - only r linearly independent columns or rows

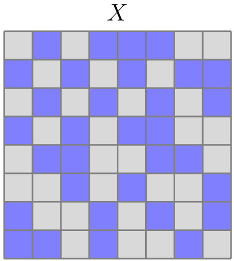

You don't need all $n \times n$ entries!

Theorem

A $n \times n$ matrix of rank $r < n$ can be exactly or accurately recovered from just $O(nr\log n)$ randomly observed entries.

Emmanuel Candes and Benjamin Recht, Exact Matrix Completion Via Convex Optimization, Foundations of Computational Mathematics.9 pp.717--772. 2009.

Why is Matrix Completion Possible?

Feasibility

- Matrix entries are correlated.
- The correlations induce low rank.
- $n \times n$ matrix of rank $r < n$
  - only r linearly independent columns or rows

X

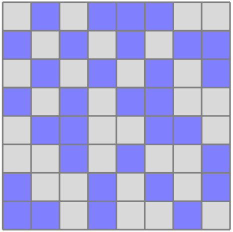

You don't need all $n \times n$ entries!

Theorem

A $n \times n$ matrix of rank $r < n$ can be exactly or accurately recovered from just $O(nr\log n)$ randomly observed entries.

Emmanuel Candes and Benjamin Recht, Exact Matrix Completion Via Convex Optimization, Foundations of Computational Mathematics.9 pp.717--772. 2009.

Université de Liège

Why is Matrix Completion Possible?

Feasibility

- Matrix entries are correlated.
- The correlations induce low rank.
- $n \times n$ matrix of rank $r < n$
  - only r linearly independent columns or rows

You don't need all $n \times n$ entries!

X

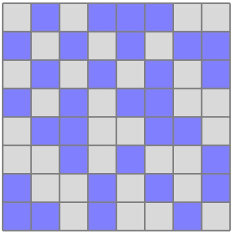

Theorem

A $n \times n$ matrix of rank $r < n$ can be exactly or accurately recovered from just $O(nr\log n)$ randomly observed entries.

Emmanuel Candes and Benjamin Recht, Exact Matrix Completion Via Convex Optimization, Foundations of Computational Mathematics.9 pp.717--772. 2009.

Why is Matrix Completion Possible?

Feasibility

- Matrix entries are correlated.
- The correlations induce low rank.
- $n \times n$ matrix of rank $r < n$
  - only r linearly independent columns or rows

X

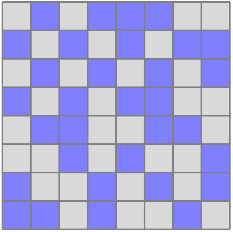

You don't need all $n \times n$ entries!

Theorem

A $n \times n$ matrix of rank $r < n$ can be exactly or accurately recovered from just $O(nr\log n)$ randomly observed entries.

Emmanuel Candes and Benjamin Recht, Exact Matrix Completion Via Convex Optimization, Foundations of Computational Mathematics.9 pp.717--772. 2009.

Why is Matrix Completion Possible?

Feasibility

- Matrix entries are correlated.
- The correlations induce low rank.
- $n \times n$ matrix of rank $r < n$
  - only r linearly independent columns or rows

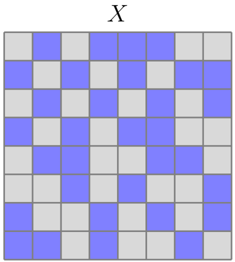

You don't need all $n \times n$ entries!

Theorem

A $n \times n$ matrix of rank $r < n$ can be exactly or accurately recovered from just $O(nr\log n)$ randomly observed entries.

Emmanuel Candes and Benjamin Recht, Exact Matrix Completion Via Convex Optimization, Foundations of Computational Mathematics.9 pp.717--772. 2009.

Why is Matrix Completion Possible?

Feasibility

- Matrix entries are correlated.
- The correlations induce low rank.
- $n \times n$ matrix of rank $r < n$
  - only r linearly independent columns or rows

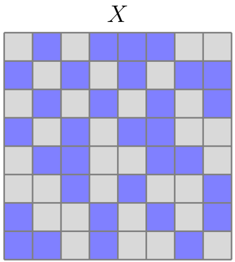

You don't need all $n \times n$ entries!

Theorem

A $n \times n$ matrix of rank $r < n$ can be exactly or accurately recovered from just $O(nr\log n)$ randomly observed entries.

Emmanuel Candes and Benjamin Recht, Exact Matrix Completion Via Convex Optimization, Foundations of Computational Mathematics.9 pp.717--772. 2009.

Correlations between Network Measurements

Link sharing across network paths

Topology: simple core

Abilene network

Routing

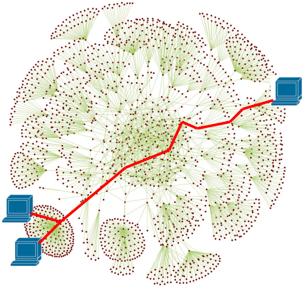

generated by Orbis

Low Rank of the Internet

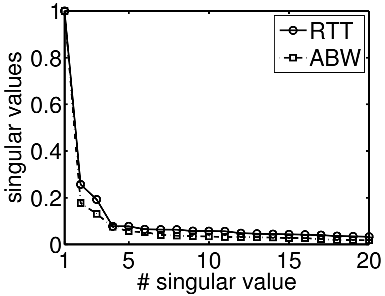

Empirical Justification

- Meridian RTT matrix of 2255×2255
- PlanetLab ABW matrix of 201×201

Université de Liège

Low Rank of the Internet

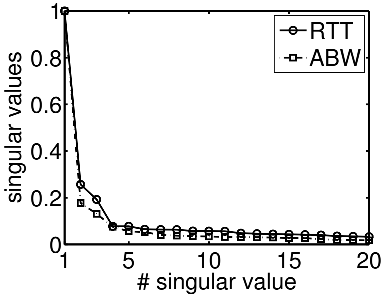

Observation

- Performance matrices are approximately low rank.
  - A perfect recovery is impossible.

Low Rank of the Internet

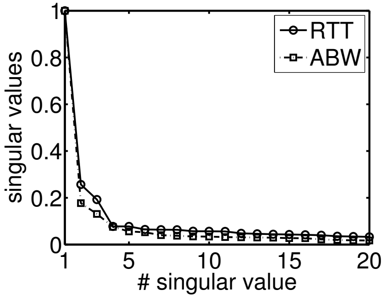

Observation

- A rank-r dominant component exists.
  - It is a fairly accurate approximation to the original matrix.

Low Rank of the Internet

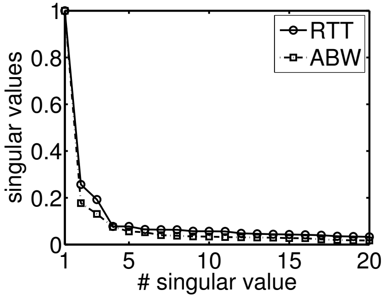

Observation

- Rank r cannot be determined a priori.
  - r is treated as a parameter and tuned for a given dataset.

Université de Liège

Low-Rank Matrix Factorization

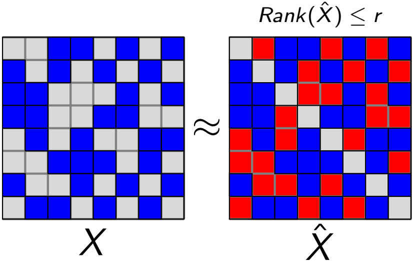

Low-Rank Matrix Factorization

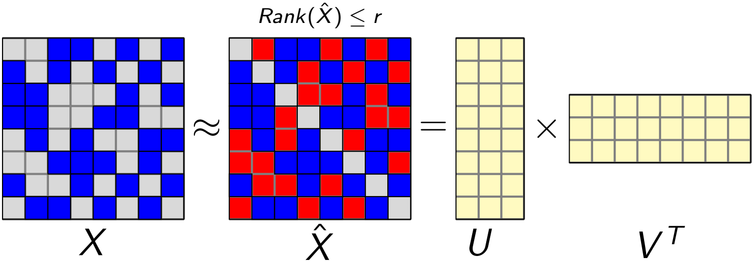

Low-Rank Matrix Factorization

Look for (U, V), instead of $\hat{X}$

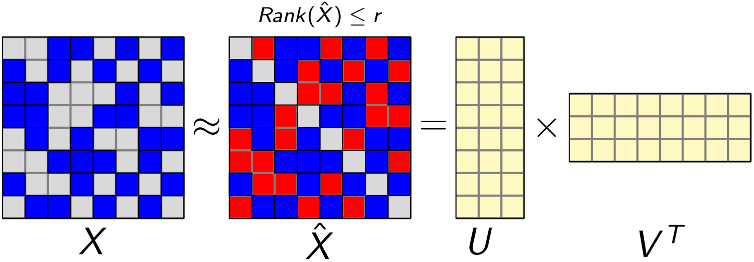

Low-Rank Matrix Factorization

Look for (U, V), instead of $\hat{X}$

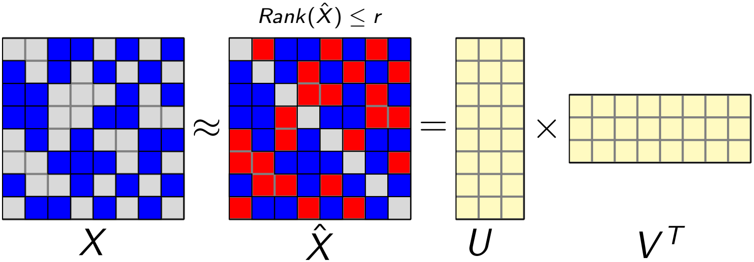

Stochastic Gradient Descent

Low-Rank Matrix Factorization

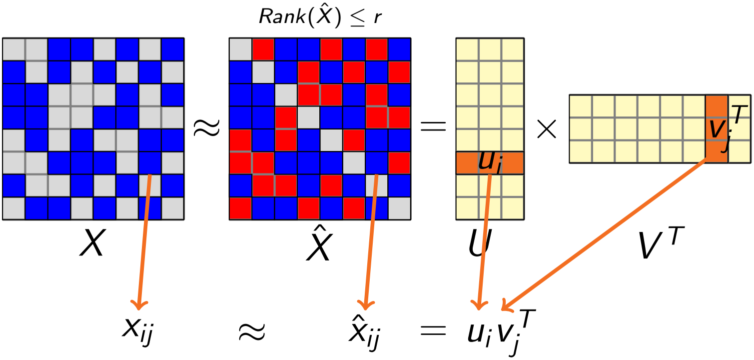

Université de Liège

Stochastic Gradient Descent (SGD)

$$\min \sum_{ij \in \Omega} l(x_{ij}, u_i v_j^T)$$

l : loss function

Ω : the set of observed entries

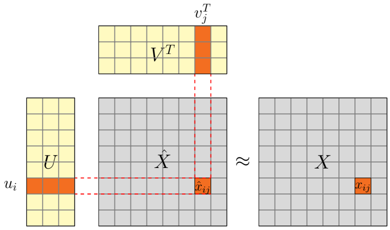

$$u_i v_j^T = \hat{x}_{ij} \approx x_{ij}$$

Stochastic Gradient Descent (SGD)

$$\min \sum_{ij \in \Omega} l(x_{ij}, u_i v_j^T)$$

l : loss function

Ω : the set of observed entries

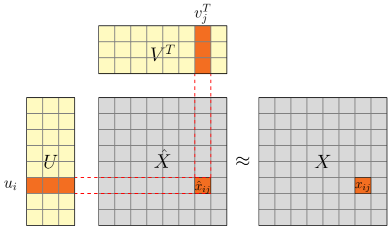

$u_i v_j^T = \hat{x}_{ij} \approx x_{ij}$

1. pick x_{ij} at random
2. update u_i, v_j by gradient descent

$$u_i = u_i - \eta \frac{\partial l(x_{ij}, u_i v_j^T)}{\partial u_i}$$

$$v_j = v_j - \eta \frac{\partial l(x_{ij}, u_i v_j^T)}{\partial v_j}$$

(η is the **learning rate**.)

3. repeat until convergence

Stochastic Gradient Descent (SGD)

$$\min \sum_{ij \in \Omega} l(x_{ij}, u_i v_j^T)$$

l : loss function

Ω : the set of observed entries

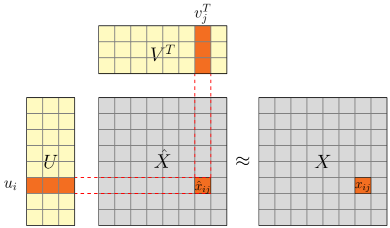

$$u_i v_j^T = \hat{x}_{ij} \approx x_{ij}$$

1. pick x_{ij} at random
2. update u_i, v_j by gradient descent

$$u_i = u_i - \eta \frac{\partial l(x_{ij}, u_i v_j^T)}{\partial u_i}$$

$$v_j = v_j - \eta \frac{\partial l(x_{ij}, u_i v_j^T)}{\partial v_j}$$

(η is the **learning rate**.)

3. repeat until convergence

Université de Liège

ReCap: Formulation as Matrix Completion

Network Performance Prediction

- A matrix completion view
- Connections to recommender systems
- Feasibility and low-rank characteristic
- Matrix factorization by Stochastic Gradient Descent (SGD)

ReCap: Formulation as Matrix Completion

Network Performance Prediction

- A matrix completion view
- Connections to recommender systems
- Feasibility and low-rank characteristic
- Matrix factorization by Stochastic Gradient Descent (SGD)

Question

- How should this problem be solved on networks?

Contributions

Learning to predict end-to-end network performance

1. Formulation as Matrix Completion

network performance prediction as matrix completion

2. Decentralized Solution

decentralized matrix factorization by stochastic gradient descent

3. Qualitative Representations of Network Performance

- represent network performance by binary classes
- represent network performance by ordinal ratings

Université de Liège

Decentralized Network Performance Prediction

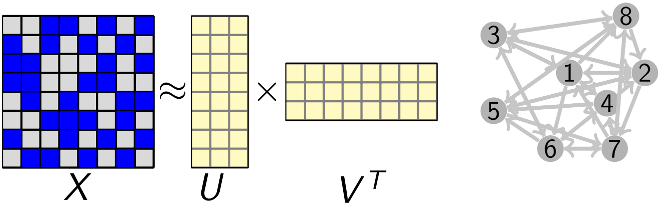

Constraints

- No central server collects and processes the measurements.
- No matrices (X, U, V) are explicitly constructed.

Decentralized Network Performance Prediction

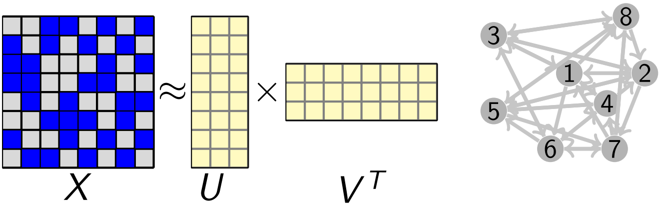

Design

Decentralized Network Performance Prediction

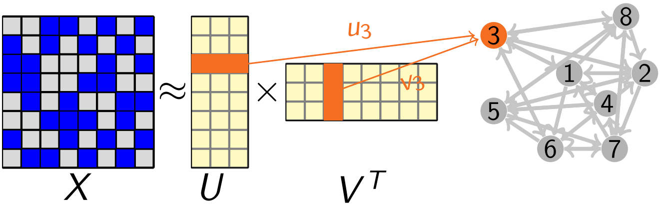

Design

- Row vectors (u_i, v_i) are stored at node i.

Decentralized Network Performance Prediction

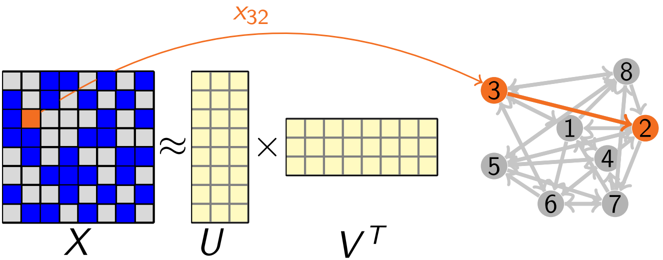

Design

- Row vectors (u_i, v_i) are stored at node i.
- Measurements are collected and processed locally.

Decentralized Network Performance Prediction

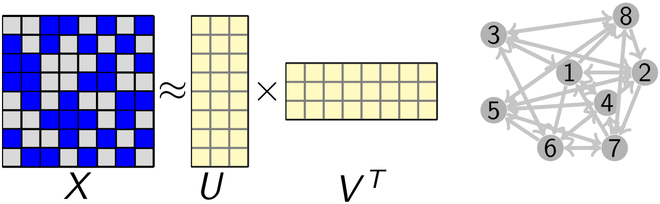

Design

- Row vectors (u_i, v_i) are stored at node i.
- Measurements are collected and processed locally.
- Network nodes exchange messages.

Metrics of Interest

round-trip time (RTT)	available bandwidth (ABW)
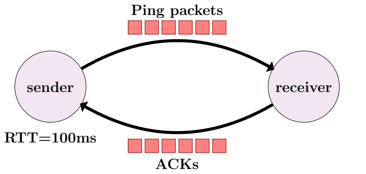	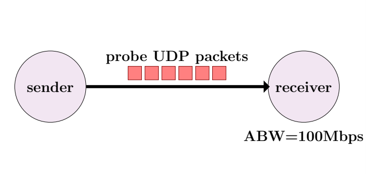
inferred by sender symmetric	inferred by receiver asymmetric

Université de Liège

Metrics of Interest

round-trip time (RTT)	available bandwidth (ABW)

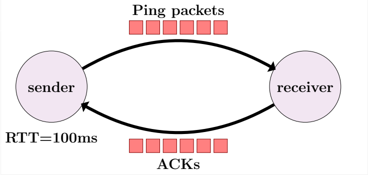

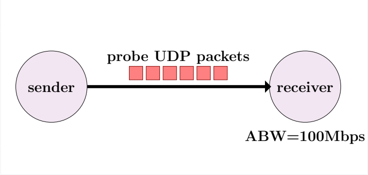

inferred by sender	inferred by receiver
symmetric	asymmetric

Message exchange is adapted for RTT and for ABW.

Basic Process for RTT

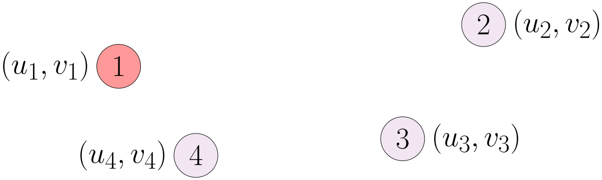

Node 1 joins the network and initializes (u_1, v_1) randomly.

Basic Process for RTT

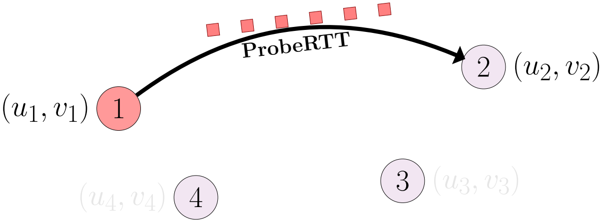

Node 1 probes node 2 for the RTT.

Basic Process for RTT

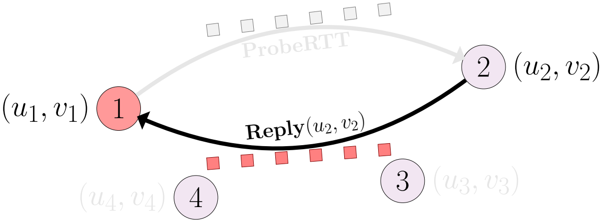

Node 2 replies and sends (u_2, v_2) to node 1.

Basic Process for RTT

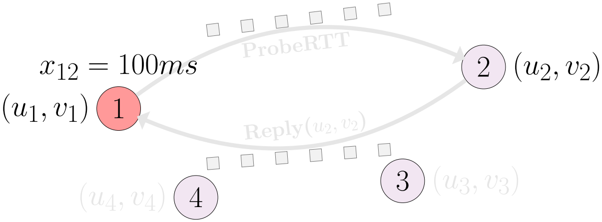

Node 1 computes RTT_{12}.

Basic Process for RTT

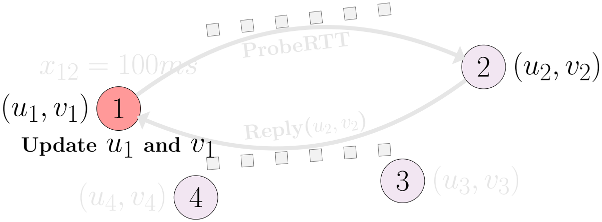

Node 1 updates (u_1, v_1).

Université de Liège

Basic Process for ABW

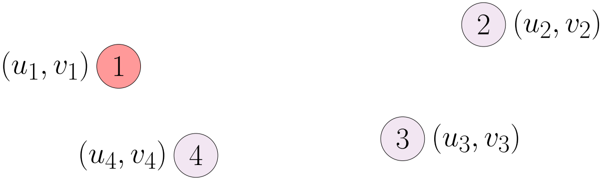

Node 1 joins the network and initializes (u_1, v_1) randomly.

Basic Process for ABW

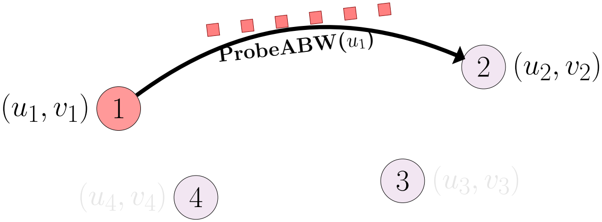

Node 1 probes node 2 for the ABW and sends u_1.

Basic Process for ABW

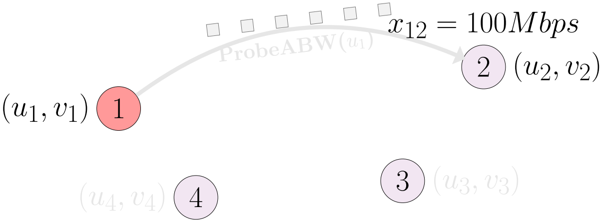

Node 2 computes ABW_{12}.

Basic Process for ABW

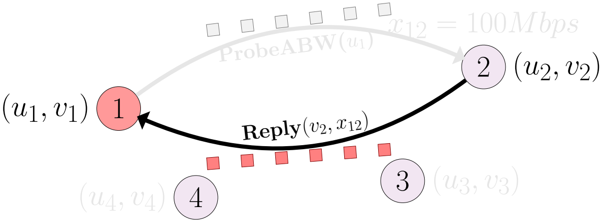

Node 2 replies and sends (ABW_{12}, v_2) to node 1.

Basic Process for ABW

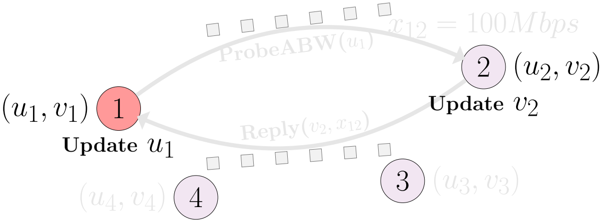

Node 1 updates u_1 and node 2 updates v_2.

Prediction of Network Performance

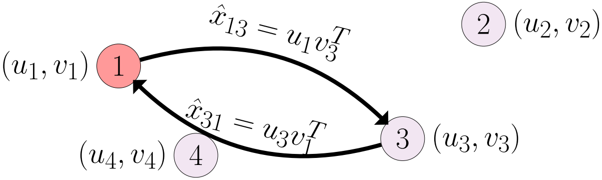

DMFSGD

- Each node selects k neighbors to communicate with.
  - $k = 10$ for small networks (as many as 200 nodes);
  - $k = 32$ for large networks (as many as 2500 nodes, $< 2\%$).
- Each node collaborates with one neighbor at a time.
  - probe measurements;
  - perform SGD updates.
- Rank-10 factorization
  - $r = 10$ in all experiments due to sparse measurements

Advantages

- simple and computationally lightweight
- suitable for large-scale dynamic measurements
- adaptable to various metrics

DMFSGD

- Each node selects k neighbors to communicate with.
  - $k = 10$ for small networks (as many as 200 nodes);
  - $k = 32$ for large networks (as many as 2500 nodes, $< 2\%$).
- Each node collaborates with one neighbor at a time.
  - probe measurements;
  - perform SGD updates.
- Rank-10 factorization
  - $r = 10$ in all experiments due to sparse measurements

Advantages

- simple and computationally lightweight
- suitable for large-scale dynamic measurements
- adaptable to various metrics

DMFSGD

- Each node selects k neighbors to communicate with.
  - $k = 10$ for small networks (as many as 200 nodes);
  - $k = 32$ for large networks (as many as 2500 nodes, $< 2\%$).
- Each node collaborates with one neighbor at a time.
  - probe measurements;
  - perform SGD updates.
- Rank-10 factorization
  - $r = 10$ in all experiments due to sparse measurements

Advantages

- simple and computationally lightweight
- suitable for large-scale dynamic measurements
- adaptable to various metrics

DMFSGD

- Each node selects k neighbors to communicate with.
  - $k = 10$ for small networks (as many as 200 nodes);
  - $k = 32$ for large networks (as many as 2500 nodes, $< 2\%$).
- Each node collaborates with one neighbor at a time.
  - probe measurements;
  - perform SGD updates.
- Rank-10 factorization
  - $r = 10$ in all experiments due to sparse measurements

Advantages

- simple and computationally lightweight
- suitable for large-scale dynamic measurements
- adaptable to various metrics

DMFSGD

- Each node selects k neighbors to communicate with.
  - $k = 10$ for small networks (as many as 200 nodes);
  - $k = 32$ for large networks (as many as 2500 nodes, $< 2\%$).
- Each node collaborates with one neighbor at a time.
  - probe measurements;
  - perform SGD updates.
- Rank-10 factorization
  - $r = 10$ in all experiments due to sparse measurements

Advantages

- simple and computationally lightweight
- suitable for large-scale dynamic measurements
- adaptable to various metrics

Experiments and Evaluations

RTT Datasets

	Harvard	P2psim	Meridian
nodes	226	1740	2500
source	Ledlie et al. NSDI 2007	Dabek et al. SIGCOMM 2004	Wong et al. SIGCOMM 2005

Vivaldi: Competitor for RTT Prediction

Euclidean embedding

- simulation of a spring system
- energy minimization by Hooke's law
- the same architecture as DMFSGD
- adopted in Azureus (now Vuze)

Dabek et al., Vivaldi: a decentralized network coordinate system, SIGCOMM 2004.

Comparison of Prediction Accuracy

Median Absolute Error (MAE) $= median_{ij}(|d_{ij} - \hat{d}_{ij}|)$

	P2PSim	Meridian	Harvard
Vivaldi	13.4ms	9.2ms	5.8ms
DMFSGD	11.5ms	9.0ms	1.1ms

DMFSGD outperforms Vivaldi almost always.

Comparison of Prediction Accuracy

Median Absolute Error (MAE) $= median_{ij}(|d_{ij} - \hat{d}_{ij}|)$

	P2PSim	Meridian	Harvard
Vivaldi	13.4ms	9.2ms	5.8ms
DMFSGD	11.5ms	9.0ms	1.1ms

DMFSGD outperforms Vivaldi almost always.

Université de Liège

Comparison of Prediction Accuracy

Median Absolute Error (MAE) $= median_{ij}(|d_{ij} - \hat{d}_{ij}|)$

	P2PSim	Meridian	Harvard
Vivaldi	13.4ms	9.2ms	5.8ms
DMFSGD	11.5ms	9.0ms	1.1ms

DMFSGD outperforms Vivaldi almost always.

ReCap: Decentralized Solution

DMFSGD

- Decentralized Architecture
- Basic Processing for RTT and ABW
- Experiments on RTT Prediction

ReCap: Decentralized Solution

DMFSGD

- Decentralized Architecture
- Basic Processing for RTT and ABW
- Experiments on RTT Prediction

- DMFSGD works well for available bandwidth prediction.

DMFSGD for Available Bandwidth Prediction

ABW measurement based on self-induced congestion

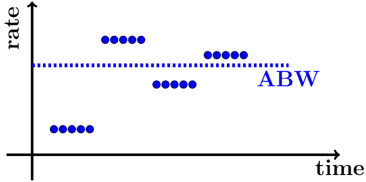

Pathload

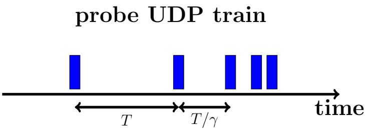

PathChirp

High cost

Problems even if we only measure 1% paths on large networks.

Contributions

Learning to predict end-to-end network performance

1. Formulation as Matrix Completion

network performance prediction as matrix completion

2. Decentralized Solution

decentralized matrix factorization by stochastic gradient descent

3. Qualitative Representations of Network Performance
    - represent network performance by binary classes
    - represent network performance by ordinal ratings

Qualitative Representations of Network Performance

Binary classes

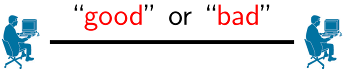

Ordinal ratings

Characteristics of Qualitative Representations

✗ Classes/Ratings are less fine-grained.

✓ Class/Rating information is sufficient to applications.

- Peer selection: "good enough" is often enough.

✓ Classes/Ratings unify different metrics.

- Performance takes a few discrete values.

✓ Class/Rating information can be encoded in a few bits.

✓ Class/Rating measurements are cheap.

- Classes/Ratings are coarse and stable.

Characteristics of Qualitative Representations

✗ Classes/Ratings are less fine-grained.

✓ Class/Rating information is sufficient to applications.

- Peer selection: "good enough" is often enough.

✓ Classes/Ratings unify different metrics.

- Performance takes a few discrete values.

✓ Class/Rating information can be encoded in a few bits.

✓ Class/Rating measurements are cheap.

- Classes/Ratings are coarse and stable.

Characteristics of Qualitative Representations

✗ Classes/Ratings are less fine-grained.

✓ Class/Rating information is sufficient to applications.

- Peer selection: "good enough" is often enough.

✓ Classes/Ratings unify different metrics.

- Performance takes a few discrete values.

✓ Class/Rating information can be encoded in a few bits.

✓ Class/Rating measurements are cheap.

- Classes/Ratings are coarse and stable.

Characteristics of Qualitative Representations

✗ Classes/Ratings are less fine-grained.

✓ Class/Rating information is sufficient to applications.

- Peer selection: “good enough” is often enough.

✓ Classes/Ratings unify different metrics.

- Performance takes a few discrete values.

✓ Class/Rating information can be encoded in a few bits.

✓ Class/Rating measurements are cheap.

- Classes/Ratings are coarse and stable.

Université de Liège

Characteristics of Qualitative Representations

✗ Classes/Ratings are less fine-grained.

✓ Class/Rating information is sufficient to applications.
- Peer selection: “good enough” is often enough.

✓ Classes/Ratings unify different metrics.
- Performance takes a few discrete values.

✓ Class/Rating information can be encoded in a few bits.

✓ Class/Rating measurements are cheap.
- Classes/Ratings are coarse and stable.

Université de Liège

Characteristics of Qualitative Representations

✗ Classes/Ratings are less fine-grained.

✓ Class/Rating information is sufficient to applications.

- Peer selection: "good enough" is often enough.

✓ Classes/Ratings unify different metrics.

- Performance takes a few discrete values.

✓ Class/Rating information can be encoded in a few bits.

✓ Class/Rating measurements are cheap.

- Classes/Ratings are coarse and stable.

Université de Liège

Class/Rating Measurement

Thresholding

- partition the range into several bins;
- map a metric value to a bin.

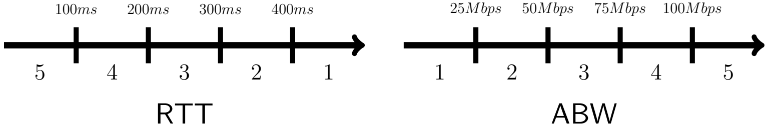

Class/Rating Measurement

Thresholding

- partition the range into several bins;
- map a metric value to a bin.

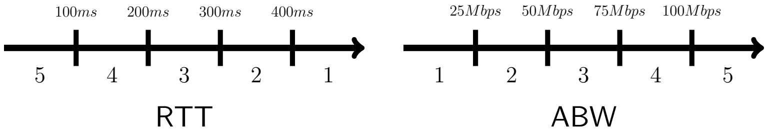

Choice of Threshold(s)

data analysis

- 50% percentile for binary classification
- $\{20\%, 40\%, 60\%, 80\%\}$ percentiles for ordinal rating

Performance Class/Rating Prediction by DMFSGD

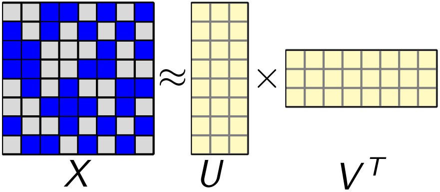

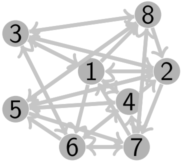

Performance Class/Rating Prediction by DMFSGD

$$x_{32} = \begin{cases} -1 \; or \; 1 \\ 1, ..., 5 \end{cases}$$

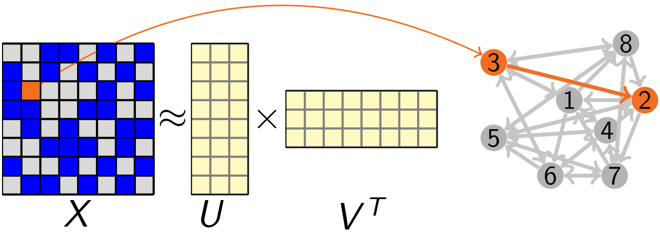

Experiments and Evaluations

Datasets

	Harvard	Meridian	HP-S3
nodes	226	2500	231
metric	RTT	RTT	ABW
source	Ledlie et al. NSDI 2007	Wong et al. SIGCOMM 2005	Ramasubramanian et al. SIGMETRICS 2009

Université de Liège

Experiments and Evaluations

Datasets

	Harvard	Meridian	HP-S3
nodes	226	2500	231
metric	RTT	RTT	ABW
source	Ledlie et al. NSDI 2007	Wong et al. SIGCOMM 2005	Ramasubramanian et al. SIGMETRICS 2009

Confusion Matrix

- Diagonal entries are correct predictions.
- Off-diagonal entries are confusions or wrong predictions.

$Accuracy = \frac{\#\ correct}{\#\ data}$		Predicted	
		"Good"	"Bad"
Actual	"Good"	correct%	confusion%
	"Bad"	confusion%	correct%

Experiments and Evaluations

Datasets

	Harvard	Meridian	HP-S3
nodes	226	2500	231
metric	RTT	RTT	ABW
source	Ledlie et al. NSDI 2007	Wong et al. SIGCOMM 2005	Ramasubramanian et al. SIGMETRICS 2009

Confusion Matrix

- Diagonal entries are correct predictions.
- Off-diagonal entries are confusions or wrong predictions.

$Accuracy = \frac{\#\ correct}{\#\ data}$		Predicted	
		"Good"	"Bad"
Actual	"Good"	correct%	confusion%
	"Bad"	confusion%	correct%

Experiments and Evaluations

Datasets

	Harvard	Meridian	HP-S3
nodes	226	2500	231
metric	RTT	RTT	ABW
source	Ledlie et al. NSDI 2007	Wong et al. SIGCOMM 2005	Ramasubramanian et al. SIGMETRICS 2009

Confusion Matrix

- Diagonal entries are correct predictions.
- Off-diagonal entries are confusions or wrong predictions.

$Accuracy = \frac{\#\ correct}{\#\ data}$		Predicted	
		"Good"	"Bad"
Actual	"Good"	correct%	confusion%
	"Bad"	confusion%	correct%

Experiments and Evaluations

Datasets

	Harvard	Meridian	HP-S3
nodes	226	2500	231
metric	RTT	RTT	ABW
source	Ledlie et al. NSDI 2007	Wong et al. SIGCOMM 2005	Ramasubramanian et al. SIGMETRICS 2009

Confusion Matrix

- Diagonal entries are correct predictions.
- Off-diagonal entries are confusions or wrong predictions.

$Accuracy = \frac{\#\ correct}{\#\ data}$		Predicted	
		"Good"	"Bad"
Actual	"Good"	correct%	confusion%
	"Bad"	confusion%	correct%

Confusion Matrix for Meridian

Binary Classification

Accuracy=85%		Predicted	
		"Good"	"Bad"
Actual	"Good"	88%	12%
	"Bad"	18%	82%

Ordinal Rating

Accuracy=56%		Predicted				
		1	2	3	4	5
Actual	1	78%	18%	3%	1%	0
	2	8%	59%	30%	5%	0
	3	1%	18%	60%	20%	1%
	4	1%	3%	33%	59%	4%
	5	1%	1%	12%	59%	27%

Confusion Matrix for Meridian

Binary Classification

Accuracy=85%		Predicted	
		"Good"	"Bad"
Actual	"Good"	88%	12%
	"Bad"	18%	82%

Ordinal Rating

Accuracy=56%		Predicted				
		1	2	3	4	5
Actual	1	78%	18%	3%	1%	0
	2	8%	59%	30%	5%	0
	3	1%	18%	60%	20%	1%
	4	1%	3%	33%	59%	4%
	5	1%	1%	12%	59%	27%

Confusion Matrix for Meridian

Binary Classification

Accuracy=85%		Predicted	
		"Good"	"Bad"
Actual	"Good"	88%	12%
	"Bad"	18%	82%

Ordinal Rating

Accuracy=56%		Predicted				
		1	2	3	4	5
Actual	1	78%	18%	3%	1%	0
	2	8%	59%	30%	5%	0
	3	1%	18%	60%	20%	1%
	4	1%	3%	33%	59%	4%
	5	1%	1%	12%	59%	27%

Confusion Matrix for Meridian

Binary Classification

Accuracy=85%		Predicted "Good"	Predicted "Bad"
Actual	"Good"	88%	12%
	"Bad"	18%	82%

Ordinal Rating

Accuracy=56%		Predicted 1	2	3	4	5
Actual	1	78%	18%	3%	1%	0
	2	8%	59%	30%	5%	0
	3	1%	18%	60%	20%	1%
	4	1%	3%	33%	59%	4%
	5	1%	1%	12%	59%	27%

Confusion Matrix for Meridian

Binary Classification

Accuracy=85%		Predicted	
		"Good"	"Bad"
Actual	"Good"	88%	12%
	"Bad"	18%	82%

Ordinal Rating

Accuracy=56%		Predicted				
		1	2	3	4	5
Actual	1	78%	18%	3%	1%	0
	2	8%	59%	30%	5%	0
	3	1%	18%	60%	20%	1%
	4	1%	3%	33%	59%	4%
	5	1%	1%	12%	59%	27%

ersité
Liège

ReCap: Qualitative Representations of Network Performance

- Binary classification and ordinal rating
- Advantages
- Class/Rating measurement
- Experiments on binary classification and ordinal rating

Université de Liège

Conclusion: Features of DMFSGD

A general framework for network performance prediction

- Soundness
  - founded on recent advances in matrix completion
- Flexibility
  - deal with various metrics (RTT, ABW, ...)
  - deal with metric values, performance classes and ratings

  A unique feature of DMFSGD!
- Simplicity
  - easy to implement
  - no infrastructure
  - computationally lightweight

Thank you!

Research supported by

- FP7 ECODE project
- FP7 mPlane project

Related Publications

- Liao et al, DMFSGD: A Decentralized Matrix Factorization Algorithm for Network Distance Prediction, IEEE/ACM Transactions on Networking, vol. 21, nb. 5, Oct. 2013, pp. 1511-1524.
- Liao et al, Decentralized Prediction of End-to-End Network Performance Classes, ACM CoNEXT 2011, Tokyo, Japan.
- Liao et al, Network Distance Prediction Based on Decentralized Matrix Factorization, IFIP Networking 2010, **best paper award**, Chennai, India.

ersité
Liège