Stefan Schmid

YuriFest

Happy birthday Yuri Gurevich!





Stefan Schmid

YuriFest

Happy birthday Yuri Gurevich!



Collaboration: B. EATCS

Since 2015



Stefan Schmid

YuriFest

Happy birthday Yuri Gurevich!



Collaboration: Visit in B. EATCS

Since 2015



Berlin 2022



Stefan Schmid

INTERACTIVE CLASSICAL ALGORITHMS: PREVIEW

Yuri Gurevich June 2025



The power of interaction

Stefan Schmid

INTERACTIVE CLASSICAL ALGORITHMS: PREVIEW

Yuri Gurevich



The power of interaction

The Bulletin of the EATCS

On logic and generative AI

Yuri Gurevich and Andreas Blass University of Michigan in Ann Arbor, MI, USA

> I think the most beautiful thing about deep learning is that it actually works. —Ilya Sutskever [13, 29:46]

§1 Thinking fast and slow¹

Q: I just learned that Daniel Kahneman, Nobel laureate in economics and the author of "Thinking, fast and slow" [7], passed away on March 27, 2024. I heard a lot about this book but have never read it. What did he mean by thinking fast and thinking slow?

A: Daniel Kahneman and Amos Tversky discovered that human thinking is driven by two distinct systems, System 1 and System 2.

System 1 supports *fast thinking*. It "operates automatically and quickly, with little or no effort and no sense of voluntary control... The capabilities of System 1 include innate skills that we share with other animals" [7, pp. 41-43]. System 1 is good at detecting patterns and reading situations on the fly. It allows us to make snap

ON A MEASURE OF INTELLIGENCE

Yuri Gurevich

The measure of intelligence is the ability to change.

-Albert Einstein

Abstract

The Fall 2024 column is a little discussion on intelligence, measuring intelligence, and related issues, provoked by a fascinating must-read article "On the measure of intelligence" by François Chollet. The discussion includes a modicum of critique of the article.

§1 Cybernetics vs. AI, and podcasts vs. reading

Quisani2 (walking in): What are you reading?

Author: An article "On the measure of intelligence" by François Chollet [3].

Q: Is it about psychology?

A: It is mostly about AI. Chollet is a prominent figure in AI.

Q: We spoke about AI last spring. But you didn't seem to be interested in AI before that.

A: This is largely correct, though I read Norbert Wiener's "Cybernetics" [18], when it was translated to Russian in 1968, and was taken with it. For a while I tried to follow cybernetics developments, at least in the USSR.

WHAT ARE KETS?

Yuri Gurevich Computer Science & Engineering University of Michigan Andreas Blass Mathematics University of Michigan

Abstract

According to Dirac's bra-ket notation, in an inner-product space, the inner product $\langle x | y \rangle$ of vectors x, y can be viewed as an application of the bra $\langle x |$ to the ket $| y \rangle$. Here $\langle x |$ is the linear functional $| y \rangle$ is $\langle x |$ y and $| y \rangle$ is the vector y. But often — though not always — there are advantages in seeing $| y \rangle$ as the function $a \mapsto a \cdot y$ where a ranges over the scalars. For example, the outer product $| y \rangle \langle x |$ becomes simply the composition $| y \rangle = \langle x |$. It would be most convenient to view kets sometimes as vectors and sometimes as functions, depending on the context. This turns out to be possible.

While the bra-ket notation arose in quantum mechanics, this note presupposes no familiarity with quantum mechanics.

1 The question

Q¹: Gentlemen, I have a question for you. But first I need to motivate it and explain where I am coming from.

The question is related to the so-called inner-product spaces which are vector spaces

Stefan Schmid

Many more at:

European Association for Theoretical Computer Science								
About EATCS How to Join	about issues contributions deadlines editors how to submit							
Organization Bulletin Conferences Awards Publications	Issues Since 2003 all issues of the EATCS Bulletin have been produced entirely electronically with colour pictures and made available on the web for members. From Number 101 and on, the Bulletin of EATCS will be published electronically using the open journal system							
Schools	Number 145, February 2025, 195pp							
Home Contact	 Access the table of contents and individual contributions via the Open Journal System Read the whole issue. (pdf 9,2 Mb) 							
Social Media	Number 144, October 2024, 146pp							
Members	 Access the table of contents and individual contributions via the Open Journal System Read the whole issue. (pdf 2,9 Mb) 							
	Number 143, June 2024, 117pp							
	 Access the table of contents and individual contributions via the Open Journal System Read the whole issue. (pdf 4,6 Mb) 							
	Number 142, February 2024, 99pp							
	 Access the table of contents and individual contributions via the Open Journal System Read the whole issue. (pdf 29,2 Mb) 							
	Number 141, October 2023, 190pp							
	 Access the table of contents and individual contributions via the Open Journal System Read the whole issue. (pdf 11,5 Mb) 							
	Number 140, June 2023, 213pp							
	 Access the table of contents and individual contributions via the Open Journal System Read the whole issue. (pdf 11,5 Mb) 							
	Number 139, February 2023, 136pp							
	 Access the table of contents and individual contributions via the Open Journal System Read the whole issue. (pdf 4,2 Mb) 							

Stefan Schmid

Yuri on Alan Turing's 100th birthday:

> Digital

Neue Zürcher Zeitung

<u>Anr</u>

Abonnieren

Geheimgespräche zweier Gründerväter

In wenigen Tagen jährt sich Alan Turings Geburtstag zum hundertsten Mal. Er hat, von seinen Zeitgenossen kaum beachtet, die Informatik grundlegend geprägt.

Stefan Betschon 14.06.2012

🗋 Merken 🛱 Drucken 🏳 Teilen



Alan Turing in einer Ausstellung im Heinz Nixdorf Museumsforum (HNF) in Paderborn im Januar (Bild: Keystone)

Stefan Schmid

"We cannot direct the wind, but we can adjust the sails." (Folklore)

Stefan Schmid

"We cannot direct the wind, but we can adjust the sails." (Folklore) Not just my research, e.g. also Microsoft Research!

Stefan Schmid

"We cannot direct the wind, but we can adjust the sails." (Folklore) Not just my research, e.g. also Microsoft Research!

Acknowledgements:





Stefan Schmid

WNetKAT: A Weighted SDN Programming and Verification Language*

Kim G. Larsen¹, Stefan Schmid², and Bingtian Xue³

- 1 Aalborg University, Aalborg, Denmark kgl@cs.aau.dk
- 2 Aalborg University, Aalborg, Denmark schmiste@cs.aau.dk
- 3 Aalborg University, Aalborg, Denmark bingt@cs.aau.dk

AalWiNes: A Fast and Quantitative What-If Analysis Tool for MPLS Networks

Peter Gjøl Jensen Aalborg University Denmark Dan Kristiansen Aalborg University Denmark Stefan Schmid Faculty of Computer Science University of Vienna Austria

Morten Konggaard Schou Aalborg University Denmark Bernhard Clemens Schrenk Faculty of Computer Science University of Vienna Jiří Srba Aalborg University Denmark

Acknowledgements:



CACM'25 (with Chen Avin)

Check out our **YouTube interviews** on Reconfigurable Datacenter Networks:



Revolutionizing Datacenter Networks via Reconfigurable Topologies Chen Avin and Stefan Schmid. Communications of the ACM (CACM), 2025. Watch here: <u>https://www.youtube.com/@self-adjusting-networks-course</u>



A bit of context

Data-Centric Applications



Interconnecting networks:
a critical infrastructure
of our digital society.



The Problem

Huge Infrastructure, Inefficient Use

- Network equipment reaching capacity limits
 - → Transistor density rates stalling
 - → "End of Moore's Law in networking" [1]
- Hence: more equipment, larger networks
- Resource intensive and:
 inefficient



Annoying for companies, opportunity for researchers

Root Cause

Fixed and Demand-Oblivious Topology

How to interconnect?



Root Cause

Fixed and Demand-Oblivious Topology



Root Cause

Fixed and Demand-Oblivious Topology



©°	© 	© 	© ■ 	© 	∎	© 	∎













Our Motivation

Much Structure in the Demand

Empirical studies:

traffic matrices sparse and skewed



destinations



destinations

traffic bursty over time



Our hypothesis: can be exploited.

Traffic is also clustered: Small Stable Clusters



Opportunity: *exploit* with little reconfigurations!

Förster et al., Analyzing the Communication Clusters in Datacenters. WWW 2023

Sounds Crazy? Emerging Enabling Technology.



H2020:

"Photonics one of only five key enabling technologies for future prosperity."

US National Research Council: "Photons are the new Electrons."

Enabler

Novel Reconfigurable Optical Switches

---> **Spectrum** of prototypes

- \rightarrow Different sizes, different reconfiguration times
- → From our last year's ACM **SIGCOMM** workshop OptSys



Example Optical Circuit Switch

---> Optical Circuit Switch rapid adaption of physical layer



 \rightharpoonup Based on rotating mirrors

Optical Circuit Switch

By Nathan Farrington, SIGCOMM 2010

Another Example

Tunable Lasers

---> Depending on wavelength, forwarded differently

---> Optical switch is passive



Electrical switch with tunable laser

Optical switch Passive

Ballani et al., Sirius, ACM SIGCOMM 2020.

Another Example

Tunable Lasers

---> Depending on wavelength, forwarded differently

---> Optical switch is passive



Electrical switch with tunable laser

Optical switch Passive

Ballani et al., Sirius, ACM SIGCOMM 2020.

Another Example

Tunable Lasers

---> Depending on wavelength, forwarded differently

---> Optical switch is passive



Electrical switch with tunable laser

Optical switch Passive

Ballani et al., Sirius, ACM SIGCOMM 2020.

The Big Picture



Now is the time!

The Big Picture



Now is the time!

Our goal: Develop the theoretical **foundations** of demand-aware, self-adjusting networks.

Analogy



Golden Gate Zipper

Unique Position

Demand-Aware, Self-Adjusting Systems


Question 1:

How to Quantify such "Structure" in the Demand?

Which demand has more structure?

Which demand has more structure?

--> Traffic matrices of two different distributed ML applications -- GPU-to-GPU

More uniform

More structure

Spatial vs temporal structure

- ---> Two different ways to generate same traffic matrix:
 - \rightarrow Same non-temporal structure
- ---> Which one has more structure?



Spatial vs temporal structure

- ---> Two different ways to generate same traffic matrix:
 - \rightarrow Same non-temporal structure
- ---> Which one has more structure?



Systematically?



Information-Theoretic Approach
"Shuffle&Compress"



Increasing complexity (systematically randomized)

More structure (compresses better)







Our Methodology

Complexity Map



temporal complexity



Our Methodology

Complexity Map



Our Methodology

Complexity Map



14

Further Reading

ACM SIGMETRICS 2020

On the Complexity of Traffic Traces and Implications CHEN AVIN, School of Electrical and Computer Engineering, Ben Gurion University of the Negev, Israel MANYA GHOBADI, Computer Science and Artificial Intelligence Laboratory, MIT, USA CHEN GRINER, School of Electrical and Computer Engineering, Ben Gurion University of the Negev, Israel STEFAN SCHMID, Faculty of Computer Science, University of Vienna, Austria This paper presents a systematic approach to identify and quantify the types of structures featured by packet traces in communication networks. Our approach leverages an information-theoretic methodology, based on iterative randomization and compression of the packet trace, which allows us to systematically remove and measure dimensions of structure in the trace. In particular, we introduce the notion of trace complexity which approximates the entropy rate of a packet trace. Considering several real-world traces, we show that trace complexity can provide unique insights into the characteristics of various applications. Based on our approach, we also propose a traffic generator model able to produce a synthetic trace that matches the complexity levels of its corresponding real-world trace. Using a case study in the context of datacenters, we show that insights into the structure of packet traces can lead to improved demand-aware network designs: datacenter topologies that are optimized for specific traffic patterns. CCS Concepts: • Networks \rightarrow Network performance evaluation; Network algorithms; Data center **networks**; • Mathematics of computing \rightarrow Information theory; Additional Key Words and Phrases: trace complexity, self-adjusting networks, entropy rate, compress, complexity map, data centers **ACM Reference Format:** Chen Avin, Manya Ghobadi, Chen Griner, and Stefan Schmid. 2020. On the Complexity of Traffic Traces and

Implications. Proc. ACM Meas. Anal. Comput. Syst. 4, 1, Article 20 (March 2020), 29 pages. https://doi.org/10. 1145/3379486

1 INTRODUCTION

Packet traces collected from networking applications, such as datacenter traffic, have been shown to feature much *structure*: datacenter traffic matrices are sparse and skewed [16, 39], exhibit

20

Question 2:

Given This Structure, What Can Be Achieved? Metrics and Algorithms?

A first insight: entropy of the demand.

Connection to Datastructures

Traditional BST



Connection to Datastructures



Connection to Datastructures & Coding



Connection to Datastructures & Coding



Connection to Datastructures & Coding



More than an analogy!

Connection to Datastructures & Coding





 $ERL(\mathcal{D},N) = \sum_{(u,v)\in\mathcal{D}} p(u,v) \cdot d_N(u,v)$

Sources





$$ERL(\mathcal{D},N) = \sum_{(u,v)\in\mathcal{D}} p(u,v) \cdot d_N(u,v)$$



 $ERL(\mathcal{D},N) = \sum_{(u,v)\in\mathcal{D}} p(u,v) \cdot d_N(u,v)$

Sources

Examples

→ DAN for △=3 → E.g., complete binary ´ tree would be log n → Can we do better?



 \rightarrow DAN for $\triangle = 2$ \rightarrow Set of **lines** and **cycles**



Examples

→ DAN for △=3 → E.g., complete binary ´ tree would be log n → Can we do better?



 \rightarrow DAN for $\triangle = 2$ \rightarrow Set of **lines** and **cycles**





Example △=2: A Minium Linear Arrangement (MLA) Problem → Minimizes sum of virtual edges



Example △=2: A Minium Linear Arrangement (MLA) Problem → Minimizes sum of virtual edges



Example △=2: A Minium Linear Arrangement (MLA) Problem → Minimizes sum of virtual edges



Related Problem

Virtual Network Embedding Problem (VNEP)

Example △=2: A Minium Linear Arrangement (MLA) Problem → Minimizes sum of virtual edges

MLA is NP-hard → … and so is our problem!



Example △=2: A Minium Linear Arrangement (MLA) Problem → Minimizes sum of virtual addees

edges

MLA is **NP-hard**

 \rightarrow ... and so is our problem!

But what about $\triangle>2$?

- \rightarrow Embedding problem still hard
- → But we have a new degree of freedom!

Example △=2: A Minium Linear Arrangement (MLA) Problem → Minimizes sum of virtual edges

MLA is NP-hard

 \rightarrow ... and so is our problem!

But what about $\triangle > 2$?

- \rightarrow Embedding problem still hard
- → But we have a new degree of freedom!



Simplifies problem?!

Algorithm: Idea



Entropy Upper Bound

- ---> Idea for algorithm:
 - \rightarrow union of trees



Entropy Upper Bound

- ---> Idea for algorithm:
 - \rightarrow union of trees
 - \rightarrow reduce degree
 - → but keep distances


Entropy Upper Bound

- ---> Idea for algorithm:
 - \rightarrow union of trees
 - \rightarrow reduce degree
 - → but keep distances
- ---> Ok for sparse demands
- \rightarrow not everyone gets tree
- → helper nodes



Entropy Upper Bound



Intuition of Algorithm



Entropy Lower Bound



Dynamic Algorithm



Much Research Ahead

- ---> So far: tip of the iceberg
- → Many challenges
 - → Demand-aware graphs for *dense matrices*?
 - → *Online* competitive algorithms?
 - → *Scalable control* (e.g., Google Jupiter)
 - → Impact on routing, congestion control, buffer management?



Thank you! Questions?



Diverse topology components:

→ demand-oblivious and demand-aware

> Demandoblivious Demandaware









Static





... On Traffic Diversity

Diverse patterns:

- → Shuffling/Hadoop: all-to-all
- → All-reduce/ML: ring or tree traffic patterns → Elephant flows
- → Query traffic: skewed → Mice flows
- → Control traffic: does not evolve but has non-temporal structure

Diverse requirements:

→ ML is bandwidth hungry, small flows are latencysensitive







Topology





Topology



Topology



Demand Matrix



Metric: throughput of a demand matrix...

Demand Matrix





Metric:			throughput	
of	а	der	nand	matrix…

... is the maximal scale down factor by which traffic is feasible.

Demand Matrix



 $\times \theta(T) =$



Metric: throughput
of a demand matrix...

... is the maximal scale down factor by which traffic is feasible.

Throughput of network θ^* : worst case T

Demand Matrix





Worst demand matrix for static and rotor: permutation. Best case for demand-aware!





Completion Time

Demand completion time: How long does
it take to serve a demand matrix?



Also useful in analysis: throughput can be computed more easily via demand completion time.