Graphs for all and everything

Biennial Workshop VI - panel #4
ICES Foundation

It all started with a stroll…

Is there a walk through the city that would cross each of those bridges once and only once?

In 1736, Euler demonstrated this is not possible, using a *graph* abstraction.

In 285+ years, *graph theory* provided many more theoretical results and graph metrics.

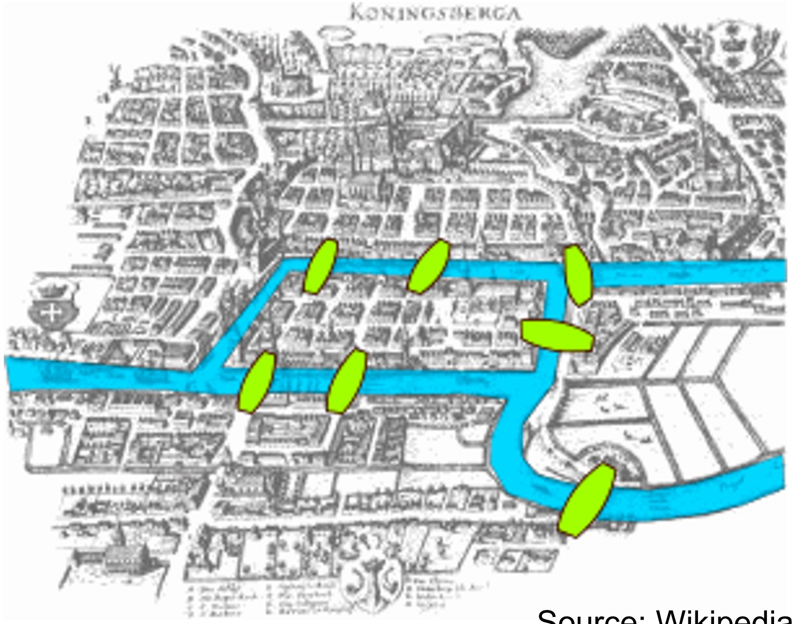

Source: Wikipedia

Graphs are everywhere …

Social media

Logistics

Bioinformatics

Text analysis

Brain modeling

Business processes

Fraud detection

…

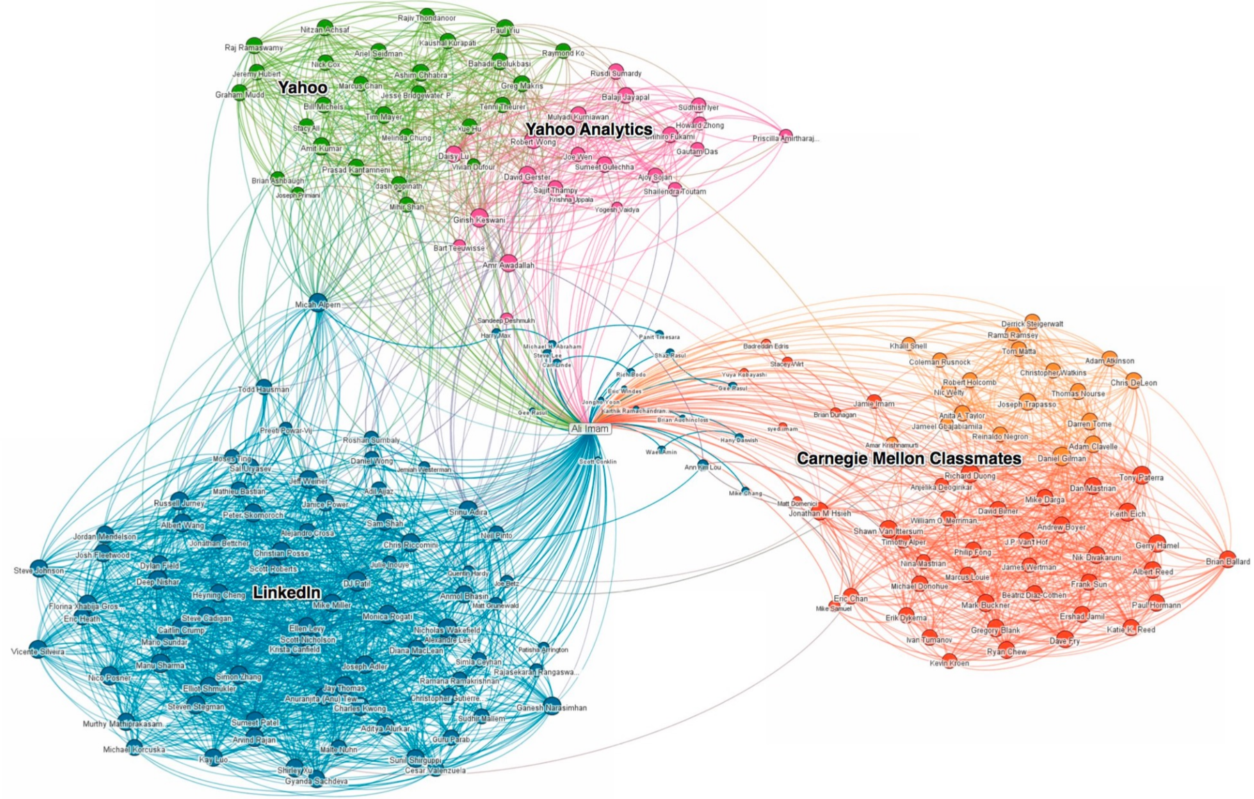

Why should we care?

Graphs are powerful computational models, fundamental to reason about and generate new knowledge from inter-connected entities.

Graphs are challenging…

- To build
- To analyse at scale
- To predict

… but bring **wisdom** to (big) data.

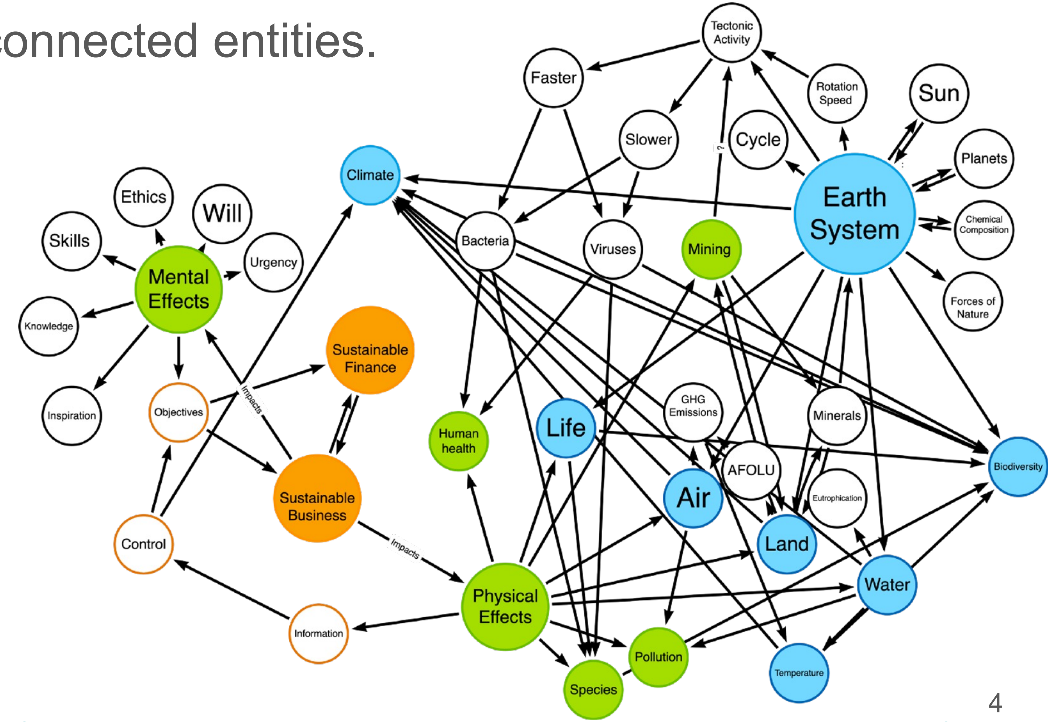

[Sustainable Finance – using knowledge graphs to model impacts on the Earth System]

Our vision [CACM 21]

COMMUNICATIONS OF THE ACM

Search

HOME | CURRENT ISSUE | NEWS | BLOGS | OPINION | RESEARCH | PRACTICE | CAREERS | ARCHIVE | VID

Home / Magazine Archive / September 2021 (Vol. 64, No. 9) / The Future Is Big Graphs: A Community View on Graph... / Full Text

CONTRIBUTED ARTICLES

The Future Is Big Graphs: A Community View on Graph Processing Systems

By Sherif Sakr, Angela Bonifati, Hannes Voigt, Alexandru Iosup, Khaled Ammar, Renzo Angles, Walid Aref, Marcelo Arenas, Maciej Besta, Peter A. Boncz, Khuzaima Daudjee, Emanuele Della Valle, Stefania Dumbrava, Olaf Hartig, Bernhard Haslhofer, Tim Hegeman, Jan Hidders, Katja Hose, Adriana Iamnitchi, Vasiliki Kalavri, Hugo Kapp, Wim Martens, M. Tamer Özsu, Eric Peukert, Stefan Plantikow, Mohamed Ragab, Matei R. Ripeanu, Semih Salihoglu, Christian Schulz, Petra Selmer, Juan F. Sequeda, Joshua Shinavier

Communications of the ACM, September 2021, Vol. 64 No. 9, Pages 62-71
10.1145/3434642

Comments

https://cacm.acm.org/magazines/2021/9/255040-the-future-is-big-graphs/

Meet the panel

Prof. Ana-Lucia Varbanescu
University of Twente, NL

Prof. Angela Bonifati
University of Lyon1

Prof. Torsten Hoefler
ETH Zurich

Prof. Adriana Iamnitchi,
University of Maastricht, NL

What's next?

Prof. Angela Bonifati
University of Lyon1

Prof. Ana-Lucia Varbanescu
University of Twente, NL

Panel will present their views on graph processing and its challenges.
…followed by a (lively!) discussion with the audience.

, NL

Prof. Angela Bonifati
University of Lyon1

Prof. Ana-Lucia Varbanescu
University of Twente, NL

Prof. Torsten Hoefler
ETH Zurich

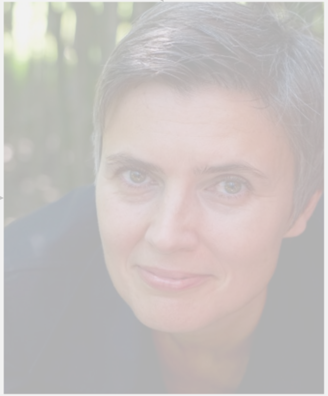

Prof. Adriana Iamnitchi,
University of Maastricht, NL

About myself

- Professor in Computer Science at Lyon 1 University (France)
- Leader of the Database group at LIRIS CNRS lab (France)
- Adjunct Professor at the University of Waterloo (Canada)
- Expertise on **Big data**, graph **querying** and indexing, **property graphs**, graph **schemas** and constraints, schema discovery, graph transformations, graph streaming, **distributed** graph databases with **performance** guarantees
- Member of WGs on standard graph query languages and property graph schemas (LDBC and ISO/IEC committees)

Angela Bonifati

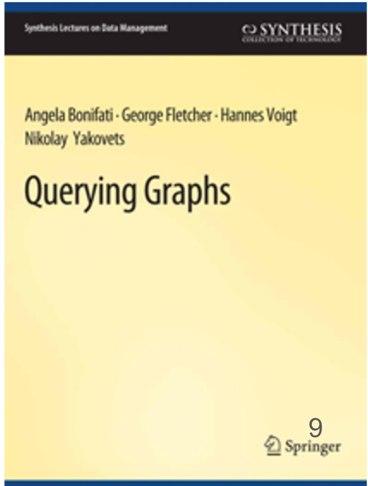

Graphs are universal!

Graphs provide a universal and simple blueprint for how to look at the world and make sense of it.

Everyone* uses graphs!

Tech-driving applications = data science + multi-hop relationships

*not yet :-(

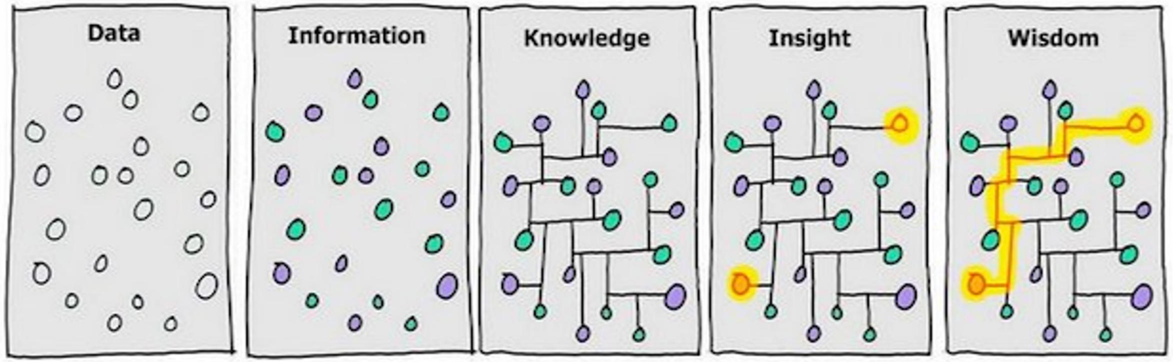

[Cartoon by David Somerville, based on a two pane version by Hugh McLeod.]

A plethora of applications

- Among which, the covidgraph.org initiative aiming at building the Covid19 knowledge graph:
    - Collecting patents, publications about the human coronaviruses
    - Biomedical data (genomics and omics)
    - Experimental data about clinical trials
    - Key demographic indicators

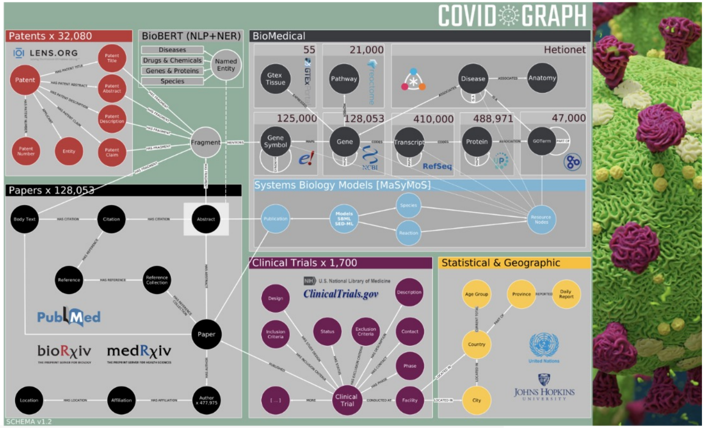

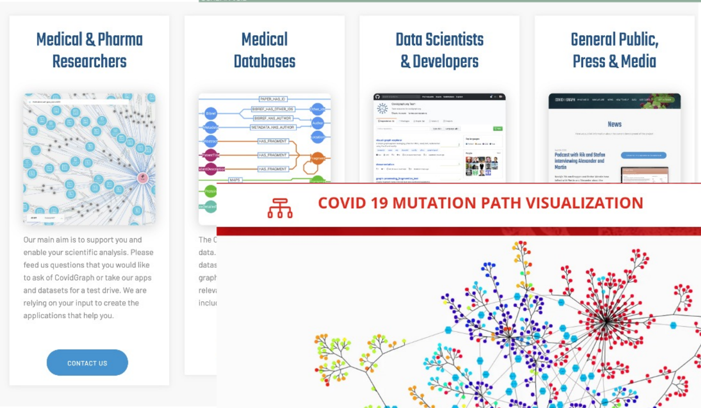

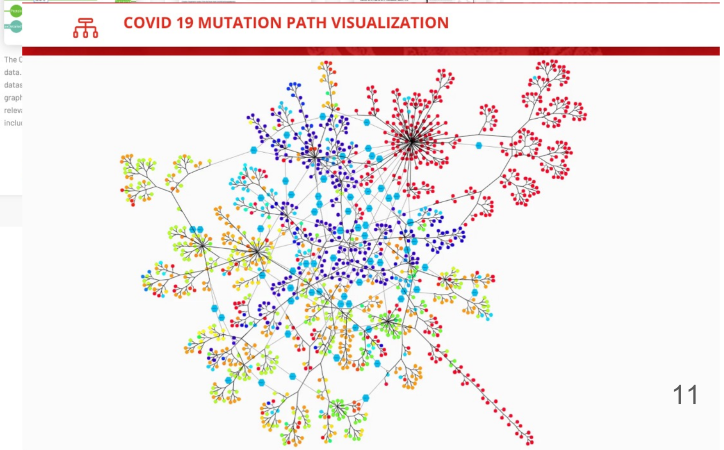

Several graph database engines on the rise

- The number of graph engines is growing over the years as well as their popularity

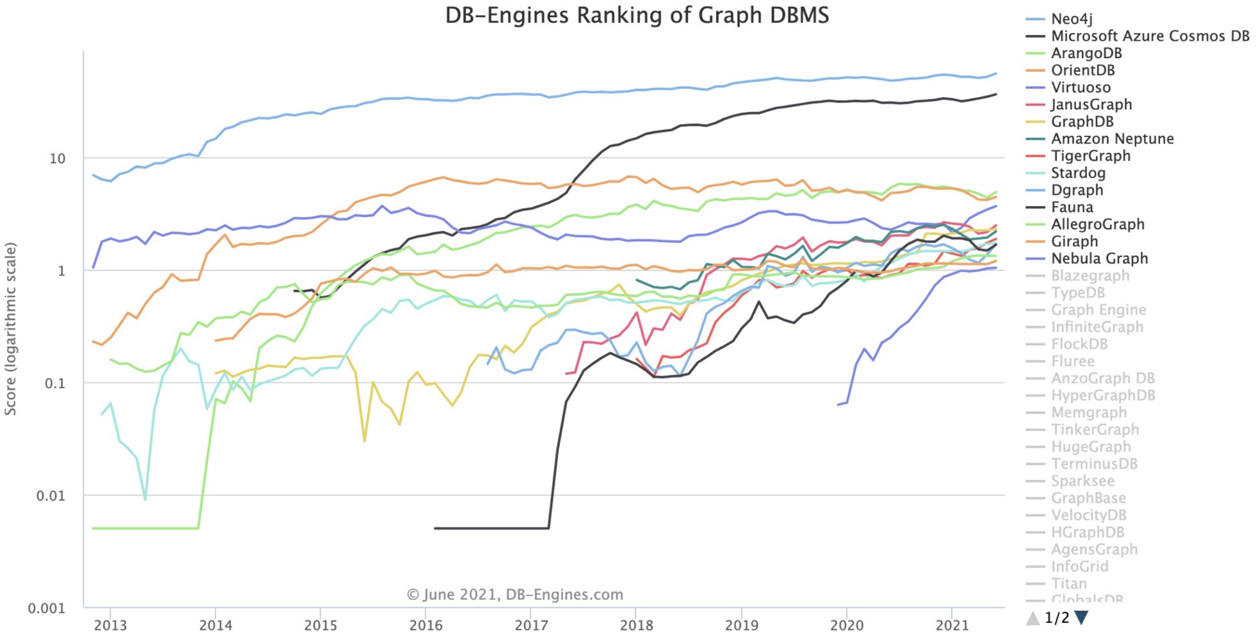

Dynamic and Streaming Graphs

- **Dynamic graphs** are graphs that can accommodate updates (insertions, deletions, changes) and allow querying on the new/old state
- **Streaming graphs** are graphs that are unbound as new data arrives at high-speed.
- Current systems and libraries focus on aggregates/projections and disregard complex analytics (recursion, paths as results of graph queries)

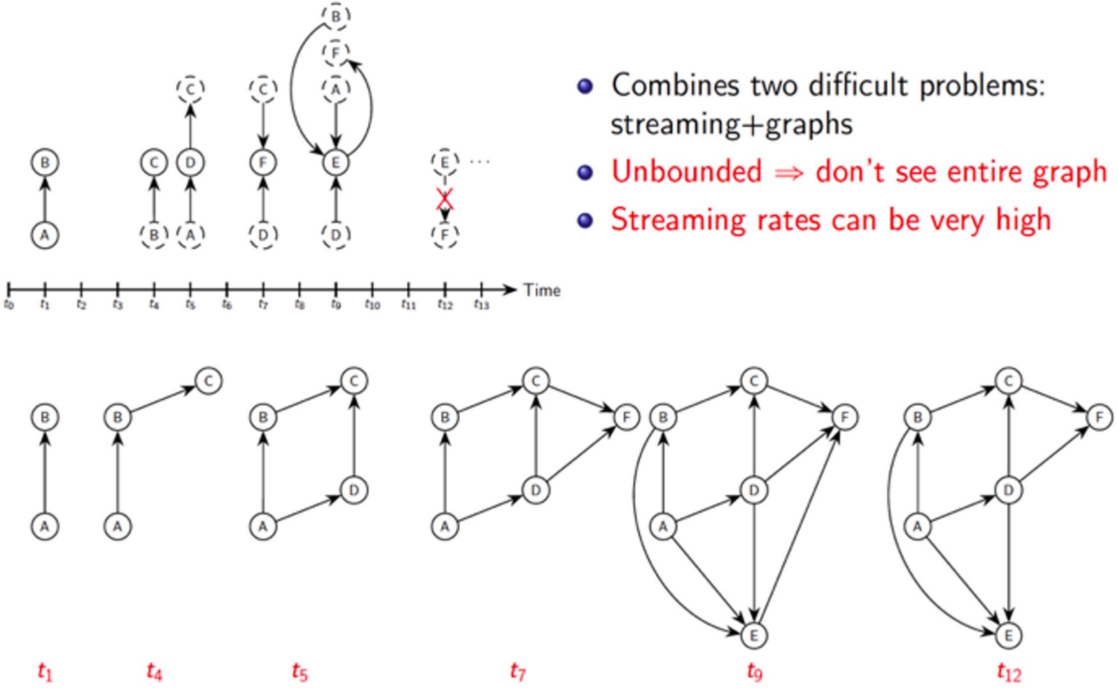

[Pa20] A. Pacaci, A. Bonifati and T. Ozsu.: Regular Path Query Evaluation on Streaming Graphs. SIGMOD Conference 2020: 1415-1430
[Pa22] A. Pacaci, A. Bonifati and T. Ozsu. Evaluating Complex Queries on Streaming Graphs. In IEEE ICDE 2022 (**Best Paper Award**)

PG-Keys: keys for property graphs

- Declaratively specify the scope of the key and its values in your favourite PG query language (a parameter of PG-Keys). Here we use Cypher-like syntax.
- For instance
    - FOR p WITHIN (p:Person) IDENTIFIER p.login; says that “**each person is identified by their login**”, and
    - FOR f WITHIN (f:Forum)<-[:joined]-(:Person) IDENTIFIER f.name, p WITHIN (f)<-[:moderates]-(p:Person); says that “**each forum with a member is identified by its name and moderator**”.

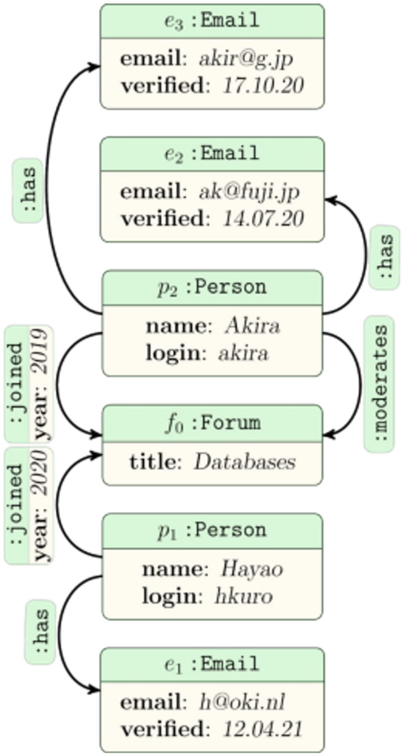

[An21] Renzo Angles, Angela Bonifati, Stefania Dumbrava, George Fletcher, Keith W. Hare, Jan Hidders, Victor E. Lee, Bei Li, Leonid Libkin, Wim Martens, Filip Murlak, Josh Perryman, Ognjen Savkovic, Michael Schmidt, Juan F. Sequeda, Slawek Staworko, Dominik Tomaszuk:PG-Keys: Keys for Property Graphs. SIGMOD Conference 2021: 2423-2436

Our vision [CACM 21]

COMMUNICATIONS OF THE ACM

Search

HOME | CURRENT ISSUE | NEWS | BLOGS | OPINION | RESEARCH | PRACTICE | CAREERS | ARCHIVE | VID

Home / Magazine Archive / September 2021 (Vol. 64, No. 9) / The Future Is Big Graphs: A Community View on Graph... / Full Text

CONTRIBUTED ARTICLES

The Future Is Big Graphs: A Community View on Graph Processing Systems

By Sherif Sakr, Angela Bonifati, Hannes Voigt, Alexandru Iosup, Khaled Ammar, Renzo Angles, Walid Aref, Marcelo Arenas, Maciej Besta, Peter A. Boncz, Khuzaima Daudjee, Emanuele Della Valle, Stefania Dumbrava, Olaf Hartig, Bernhard Haslhofer, Tim Hegeman, Jan Hidders, Katja Hose, Adriana Iamnitchi, Vasiliki Kalavri, Hugo Kapp, Wim Martens, M. Tamer Özsu, Eric Peukert, Stefan Plantikow, Mohamed Ragab, Matei R. Ripeanu, Semih Salihoglu, Christian Schulz, Petra Selmer, Juan F. Sequeda, Joshua Shinavier
Communications of the ACM, September 2021, Vol. 64 No. 9, Pages 62-71
10.1145/3434642
Comments

ARTICLE CONTENTS:

https://cacm.acm.org/magazines/2021/9/255040-the-future-is-big-graphs/

Graph processing ecosystems

- Complex workflows combining OLTP and OLAP processing are needed in order to handle heterogeneous data and heterogeneous queries and algorithms in full-fledged graph ecosystems

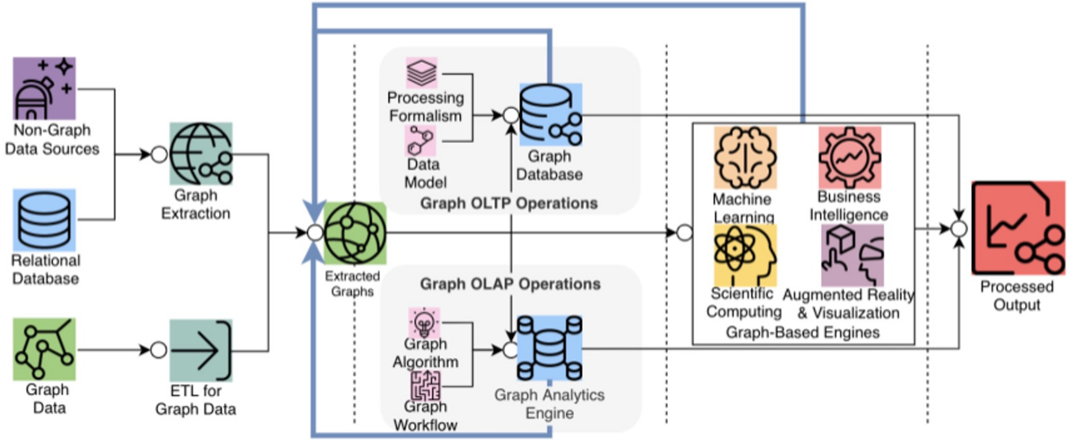

[Sa21] S.Sakr et al. The Future is Big Graphs! a community view on graph processing systems. Commun. ACM 64(9): 62-71 (2021)

Graph analytics at scale

Multi-hop analysis faces combinatorial scaling problem: Every step deeper into the graph multiplies the number of choices and cases to consider

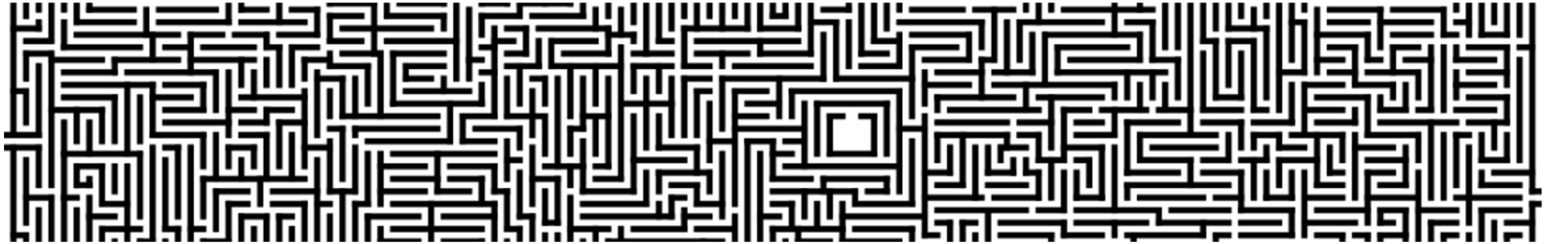

Dealing with this technical challenge is not the typical business interest of a user.

Which challenges are ahead of us to ready graph processing systems for the future?

Challenges to overcome: **Abstractions, Ecosystems, Performance**

Graph Analytics challenges require expertise of many different fields

- Computer systems
- Data management systems
- Data management theory
- Data analytics
- Visualization
- Human computer interaction
- ML/Artificial Intelligence
- …

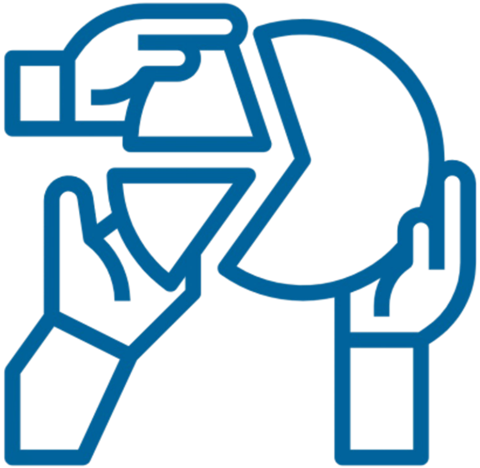

[collaborate by ArmOkay from the Noun Project]

Prof. Angela Bonifati
University of Lyon1

Prof. Ana-Lucia Varbanescu
University of Twente, NL

Prof. Torsten Hoefler
ETH Zurich

Prof. Adriana Iamnitchi,
University of Maastricht, NL

Adriana Iamnitchi

- Professor/Chair of Computational Social Sciences, Maastricht University, NL
    - Until 2021 Professor of Computer Science at University of South Florida, USA
- Expertise in large-scale networked systems, network science, social media forensics, social media modeling/forecasting

2022

[j31] Sameera Horawalavithana, Nazim Choudhury, John Skvoretz, Adriana Iamnitchi:
Online discussion threads as conversation pools: predicting the growth of discussion threads on reddit. Comput. Math. Organ. Theory 28(2): 112-140 (2022)

[j30] Kin Wai Ng, Sameera Horawalavithana, Adriana Iamnitchi:
Social media activity forecasting with exogenous and endogenous signals. Soc. Netw. Anal. Min. 12(1): 102 (2022)

[c54] Catalina Goanta, Thales Bertaglia, Adriana Iamnitchi:
The Case for a Legal Compliance API for the Enforcement of the EU's Digital Services Act on Social Media Platforms. FAccT 2022: 1341-1349

Stories with graphs: an information campaign

The **White Helmets**: a Syrian volunteer organization known for:

- Humanitarian actions
- Efforts to rescue civilians in war zones during the Syrian civil conflict
- Refusal to align with groups or military factions

They also provided:

- Video footage documentation of search and rescue operations
- Videos showing the human impact of the conflict

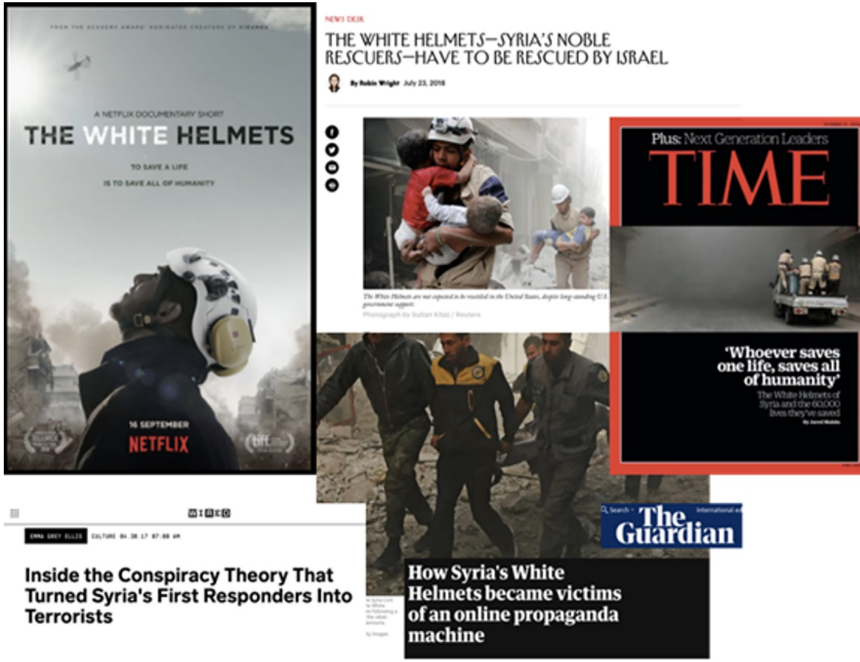

Data Forensic Questions

- Are there signs of coordinated actions in promoting videos on YouTube? (single platform) (Choudhury et al., 2020 and NG et al., 2021)
  - We discovered the promotion of near-identical videos posted in different channels
- How are YouTube videos publicized on Twitter and Facebook? (multiple platforms) (NG et al., 2021)
  - We discovered unusual patterns of synchronized behavior between users from multiple platforms

Check for updates

Strategic Information Operation in YouTube: The Case of the White Helmets

Nazim Choudhury(✉), Kin Wai Ng, and Adriana Iamnitchi

University of South Florida, Tampa, USA
{nachoudhury,kinwaing,aii}@usf.edu

Abstract. Strategic information operations (e.g. disinformation, political propaganda, and other forms of online manipulation) are critical concerns for researchers in social cyber security. Two strategies, spoofing and astroturfing, are often employed in disinformation campaigns to

Multi-platform Information Operations: Twitter, Facebook and YouTube against the White Helmets

Kin Wai NG, Sameera Horawalavithana, Adriana Iamnitchi
University of South Florida
kinwaing@usf.edu, sameera1@usf.edu, anda@cse.usf.edu

Abstract

Social media platforms are often used as a tool for hosting strategic information operations (e.g., disinformation, influence campaigns, or political propaganda). Understanding how these operations span across multiple social media plat-

paign motivates the need to study how these operations are deployed across multiple platforms as opposed to a single platform.

This study is a new look at the information campaign against the White Helmets by analyzing the activity on three

Social Media Datasets

- YouTube
  - Data collected using YouTube API Keywords
  - 666 videos posted between June 19, 2014 and April 30, 2019
  - For each video: date published, channel, and English translation of title and captions
- Twitter
  - Data collected using GNIP API from April 1, 2018 to April 30, 2019
  - Selected only Twitter posts containing links to videos in YouTube dataset
  - 14,776 tweets
- Facebook
  - Data collected using CrowdTangle's URL endpoint query API
  - Public Facebook posts with links to YouTube videos present in our dataset.
  - 961 posts by 611 users between April 1, 2018 and April 30, 2019
  - Out of 666 videos, only 236 were present in this dataset.

Data Collection Keywords
'white helmets', 'cascos blancos', 'capacetes brancos', 'caschi bianchi', 'casques blancs', 'elmetti bianchi', 'weisshelme', 'weiß helme', 'syrian civil defence', 'белые каски', 'الدفاع المدني السوري '

Message promotion on YouTube

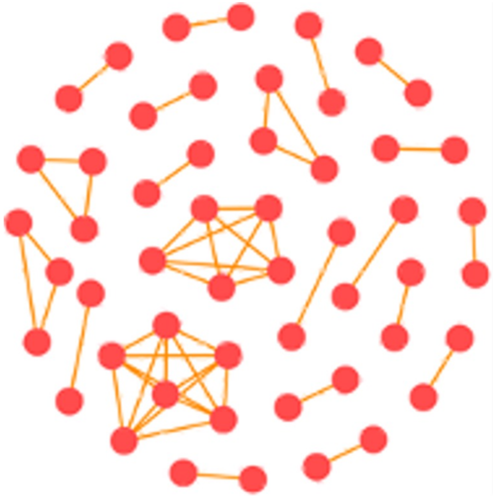

53 videos out of 666 with near-identical content uploaded to 35 different channels

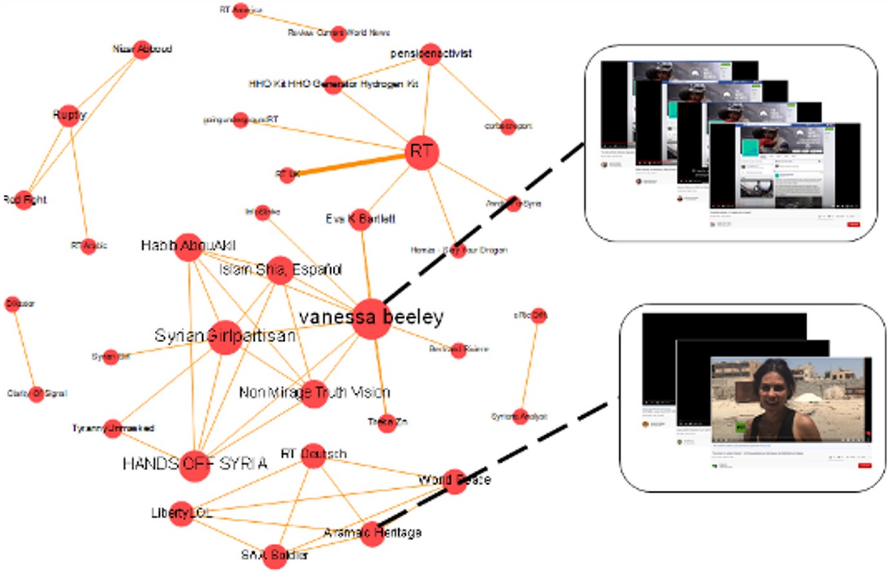

Channels connected by near-identical videos. Heavy presence of **Russian media, Western journalists, and information activists** involved in content coordination.

Inorganic Activity from Top-level Comments on YouTube

- 62 out of 666 videos had near-duplicate comments
- Out of 14K comments, 241 have at least one near duplicate

Videos-Comments Network

- 62 videos (red nodes)
- 241 comments (blue nodes)
- Green edges represent a comment to a video
- Orange edges represent a pair of near-duplicate comments

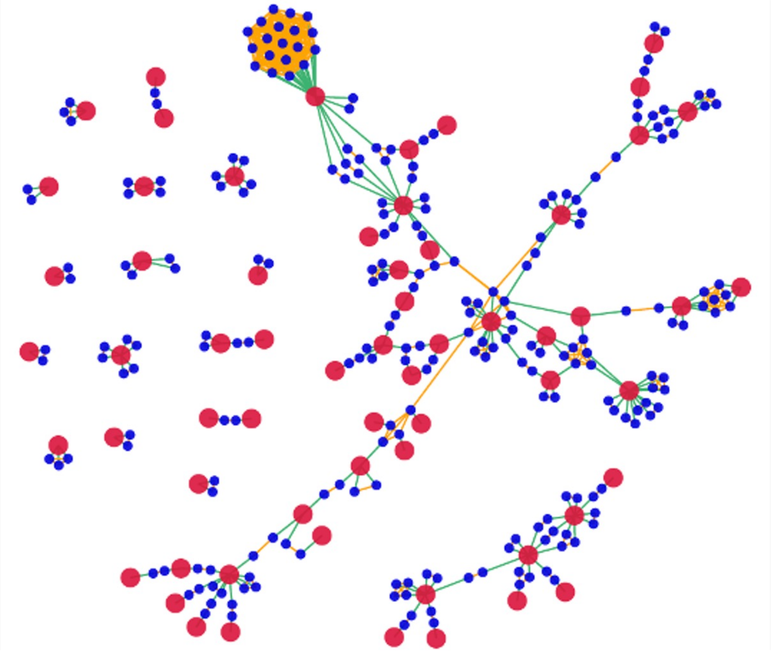

How are YouTube videos posted on Twitter and Facebook?

- Nodes are:
  - Red: YouTube videos (144)
  - Blue: Twitter users (471)
  - Green: Facebook users (161)
- Time threshold of 52 seconds computed based on an inter-arrival analysis between posts to the same video
- We connect social media user accounts that post the same YouTube videos within 52 seconds
  - 450 edges

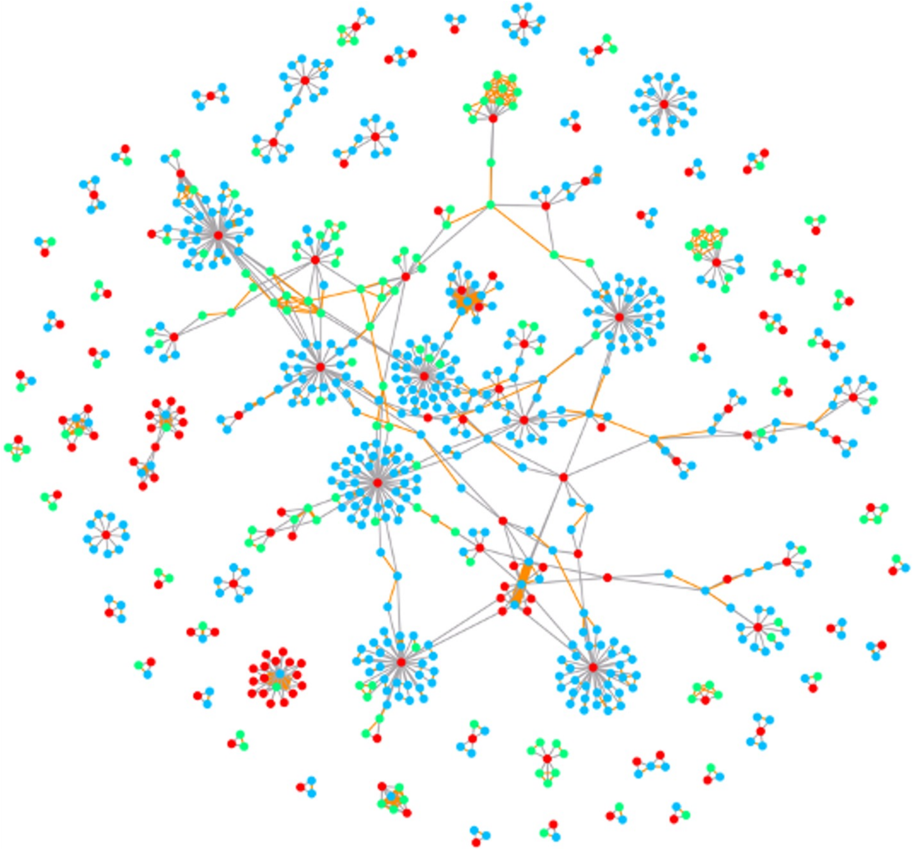

Agenda Broadcasting

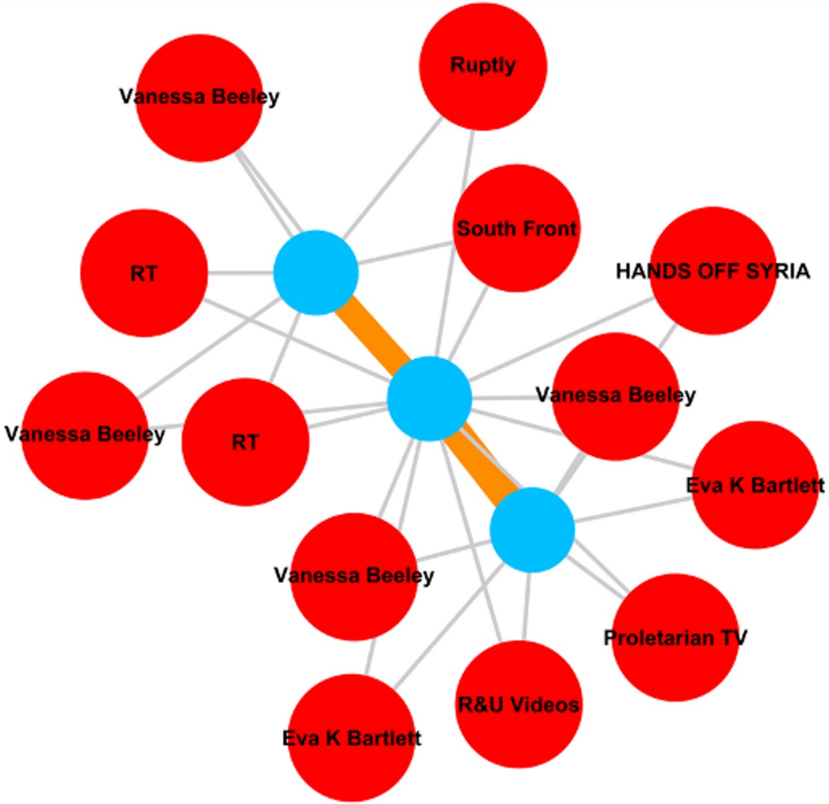

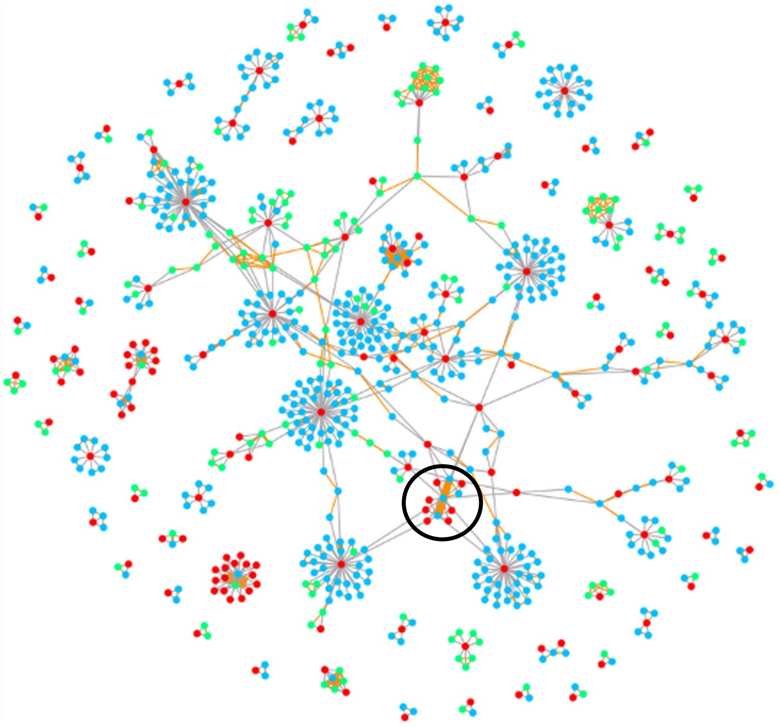

The node at the center shared all videos in detail network, whereas the other two coordinately (within 52 seconds) posted 6 and 7 videos, respectively.

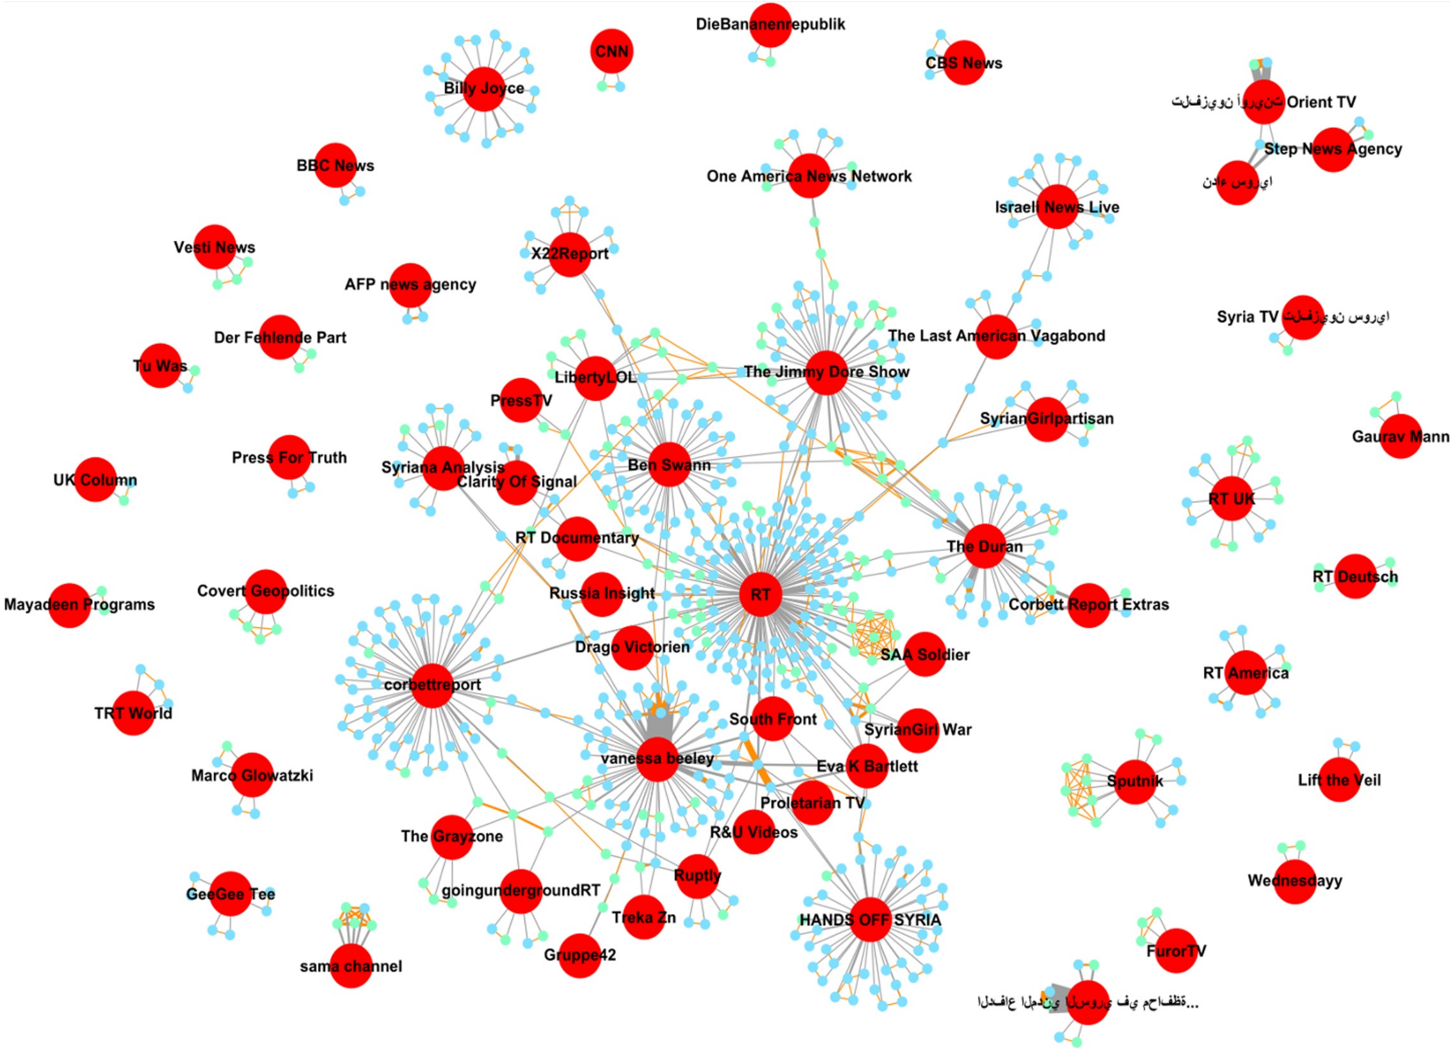
DieBananenrepublik
CNN
Billy Joyce
CBS News
تلفزيون أورينت Orient TV
Step News Agency
نداء سوريا
BBC News
One America News Network
Israeli News Live
Vesti News
X22Report
AFP news agency
Der Fehlende Part
Tu Was
The Last American Vagabond
Syria TV تلفزيون سوريا
LibertyLOL
The Jimmy Dore Show
PressTV
SyrianGirlpartisan
Gaurav Mann
Press For Truth
Syriana Analysis
Clarity Of Signal
Ben Swann
UK Column
RT UK
RT Documentary
The Duran
RT Deutsch
Covert Geopolitics
Russia Insight
RT
Mayadeen Programs
Corbett Report Extras
Drago Victorien
SAA Soldier
corbettreport
RT America
TRT World
South Front
SyrianGirl War
vanessa beeley
Eva K Bartlett
Marco Glowatzki
Sputnik
Lift the Veil
Proletarian TV
R&U Videos
The Grayzone
Ruptly
Wednesdayy
GeeGee Tee
goingundergroundRT
Treka Zn
HANDS OFF SYRIA
sama channel
Gruppe42
FurorTV
الدفاع المدني السوري في محافظة...

Example: Coordinated Video Recurrence

- Recurrent coordinated attempts among Twitter users in promoting videos from two channels.
- One Twitter account repeatedly posted one video with always the same message.
    - This user posted one video 16 times, and other users coordinated repeatedly by posting simultaneously

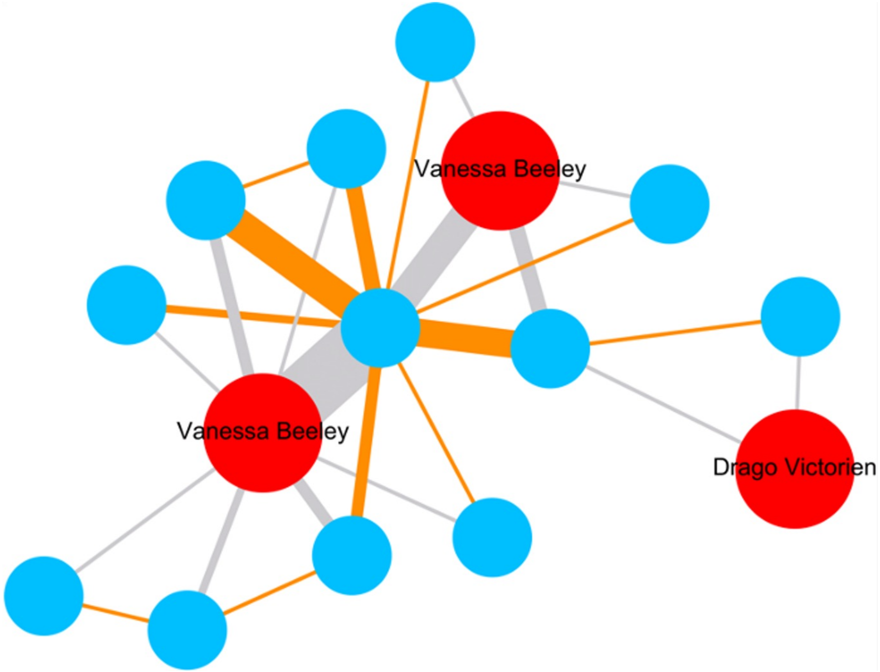

Take-away messages

- Different graph definitions can highlight different issues
- In these examples we reduced datasets to manageable-size graphs
  - Filter datasets by keywords to select only topic-specific content
  - Intersection of datasets (YouTube and (Twitter or Facebook)) reduces graphs further
- Automatic tools that detect/quantify coordination in message promotion can help educate users
  - No need to identify true vs fake news (which is hard to automize)
  - Just signal visibly that message is suspiciously promoted by synchronized user accounts
  - Caveat: sometimes users react organically at the same time due to big breaking news

Prof. Angela Bonifati
University of Lyon1

Prof. Ana-Lucia Varbanescu
University of Twente, NL

Prof. Torsten Hoefler
ETH Zurich

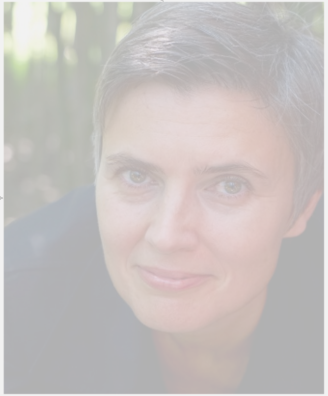

Prof. Adriana Iamnitchi,
University of Maastricht, NL

Torsten Hoefler - background

- **Teaching at ETH Zurich**
  - High-performance Computing, typically very large scale, helped to design top-5 systems
  - Large-scale machine learning
  - All with irregular graphs!
- **View between hardware/software/algorithms**
  - Enjoys fundamental principles and mathematical models

Large-scale Graph Processing

- All about efficiency
  - Compressed representation
  - $\Omega(\log(n))$ bits per vertex
    - Binary distances
  - Additional bits for weights
- Approximate graphs needed

Graph Neural Networks

- Embeddings for vertices
  - Maintain discrete graph structure at high cost
  - Add embeddings to encode structure and semantics
- How to approximate?

Euclidian Embeddings

- Vertices are embedding vectors
  - Approximate graph structure mixed with semantics
  - Quite elegant $O(1)$ encoding per vertex
- Nice tradeoffs for approximation

Processing large graphs

70 trillion edges on 10 million cores

12-trillion edge real-world graph (Internet) in 8.5s per iteration

Gordon Bell Prize finalist with pure graphs!

Still? Largest documented graph job.

ShenTu: Processing Multi-Trillion Edge Graphs on Millions of Cores in Seconds

Heng Lin[1,2], Xiaowei Zhu[1,5], Bowen Yu[1], Xiongchao Tang[1,5], Wei Xue[1], Wenguang Chen[1], Lufei Zhang[3], Torsten Hoefler[4], Xiaosong Ma[5], Xin Liu[6], Weimin Zheng[1], and Jingfang Xu[7]

***Abstract*—Graphs are an important abstraction used in many scientific fields. With the magnitude of graph-structured data constantly increasing, effective data analytics requires efficient and scalable graph processing systems. Although HPC systems have long been used for scientific computing, people have only recently started to assess their potential for graph processing, a workload with inherent load imbalance, lack of locality, and access irregularity. We propose ShenTu[8], the first general-purpose graph processing framework that can efficiently utilize an entire Petascale system to process multi-trillion edge graphs in seconds. ShenTu embodies four key innovations: hardware specialization, supernode routing, on-chip sorting, and degree-aware messaging, which together enable its unprecedented performance and scalability. It can traverse a record-size 70-trillion-edge graph in seconds. Furthermore, ShenTu enables the processing of a spam detection problem on a 12-trillion edge Internet graph, making it possible to identify trustworthy and spam webpages directly at the fine-grained page level.**

***Index Terms*—Application programming interfaces; Big data applications; Data analysis; Graph theory; Supercomputers**

I. Justification for ACM Gordon Bell Prize

ShenTu enables highly efficient general-purpose graph processing with novel use of heterogeneous cores and extremely large networks, scales to the full TaihuLight, and enables graph analytics on 70-trillion-edge graphs. It computes PageRank and TrustRank distributions for an unprecedented 12-trillion-edge real-world web graph in 8.5 seconds per iteration.

III. Overview of the Problem

Graphs are one of the most important tools to model complex systems. Scientific graph structures range from multi-billion-edge graphs (e.g., in protein interactions, genomics, epidemics, and social networks) to trillion-edge ones (e.g., in connectomics and internet connectivity). Timely and efficient processing of such large graphs is not only required to advance scientific progress but also to solve important societal challenges such as detection of fake content or to enable complex data analytics tasks, such as personalized medicine.

Improved scientific data acquisition techniques fuel th rapid growth of large graphs. For example, cheap sequen techniques lead to massive graphs representing millions man individuals as annotated paths, enabling quick advanc medical data analytics [1]. For each individual, human gen researchers currently assemble de Bruijn graphs with ov billion vertices/edges [2]. Similarly, connectomics mod human brain, with over 100 billion neurons and a erage of 7,000 synaptic connections each [3].

Meanwhile, researchers face unprecedented challenges in the study of human interaction graphs. Malicious activities such as the distribution of phishing emails or fake content, as well as massive scraping of private data, are posing threats to human society. It is necessary to scale graph analytics with the

Will graphs survive contact with ML? ☠

ICML'21

ProGraML: A Graph-based Program Representation for Data Flow Analysis and Compiler Optimizations

Chris Cummins [*1] Zacharias V. Fisches [*2] T

Abstract

Machine learning (ML) is increasingly seen
viable approach for building compiler optim
tion heuristics, but many ML methods ca
replicate even the simplest of the data flow an
ses that are critical to making good optimiza
decisions. We posit that if ML cannot do
then it is insufficiently able to reason about
grams. We formulate data flow analyses as su
vised learning tasks and introduce a large
dataset of programs and their corresponding
bels from several analyses. We use this dat

KDD'22

Motif Prediction with Graph Neural Networks

Maciej Besta[1†], Raphael Grob[1], Cesare Miglioli[2], Nicola Bernold[1],
Grzegorz Kwasniewski[1], Gabriel Gjini[1], Raghavendra Kanakagiri[3], Saleh Ashkboos[1],
Lukas Gianinazzi[1], Nikoli Dryden[1], Torsten Hoefler[1†]
[1]ETH Zurich [2]Research Center for Statistics, University of Geneva [3]UIUC [†]Corresponding authors

ABSTRACT

…lems in graph mining. How-
…ance of *higher-order network*
…ed motifs are the first-class
…k prediction schemes fail to
…this, we establish a general
…ose several heuristics that
…tif to appear. To make the
… – among others – *correla-*
…pact of some arriving links
…en motif. Finally, for highest
…etwork (GNN) architecture
… offers vertex features and
…ch structural properties of
…d do not need any training,

example, one could use motif prediction to find probable missing clusters of interactions in biological (e.g., protein) networks, and use the outcomes to limit the number of expensive experiments conducted to find missing connections [65, 67].

In this paper, we first (Section 3) establish and formally describe a general motif prediction problem, going beyond link prediction and showing how to predict higher-order network patterns that will appear in the future (or which may be missing from the data). A key challenge is the appropriate *problem formulation*. Similarly to link prediction, one wants a *score function* that – for a given vertex set V_M – assesses the chances for a given motif to appear. Still, the function must consider the combinatorially increased complexity of the problem (compared to link prediction). In general, contrary to a single link, a motif may be formed by an *arbitrary* set V_M of vertices, and the number of potential edges between these vertices can be large, i.e., $O(|V_M|^2)$. For example, one may be interested in

1

Parallel and Distributed Graph Neural Networks: An In-Depth Concurrency Analysis

Maciej Besta and Torsten Hoefler
Department of Computer Science, ETH Zurich

Abstract—Graph neural networks (GNNs) are among the most powerful tools in deep learning. They routinely solve complex problems on unstructured networks, such as node classification, graph classification, or link prediction, with high accuracy. However, both inference and training of GNNs are complex, and they uniquely combine the features of irregular graph processing with dense and regular computations. This complexity makes it very challenging to execute GNNs efficiently on modern massively parallel architectures. To alleviate this, we first design a taxonomy of parallelism in GNNs, considering data and model parallelism, and different forms of pipelining. Then, we use this taxonomy to investigate the amount of parallelism in numerous GNN models, GNN-driven machine learning tasks, software frameworks, or hardware accelerators. We use the work-depth model, and we also assess communication volume and synchronization. We specifically focus on the sparsity/density of the associated tensors, in order to understand how to effectively apply techniques such as vectorization. We also formally analyze GNN pipelining, and we generalize the established Message-Passing class of GNN models to cover arbitrary pipeline depths, facilitating future optimizations. Finally, we investigate different forms of asynchronicity, navigating the path for future asynchronous parallel GNN pipelines. The outcomes of our analysis are synthesized in a set of insights that help to maximize GNN performance, and a comprehensive list of challenges and opportunities for further research into efficient GNN computations. Our work will help to advance the design of future GNNs.

Index Terms—Parallel Graph Neural Networks, Distributed Graph Neural Networks, Parallel Graph Convolution Networks, Distributed Graph Convolution Networks, Parallel Graph Attention Networks, Distributed Graph Attention Networks, Parallel Message Passing Neural Networks, Distributed Message Passing Neural Networks, Asynchronous Graph Neural Networks.

Learning C

Luka

ABSTRACT

We present a novel neural architect
tion problems where the solution con
allowing us to solve hard problems l
our model using reinforcement learn
ents, which gives us both a greedy a
architecture builds on a graph attenti
inductive biases to improve solution quality. Our learned deterministic heuristics for graph coloring give better solutions than classical degree-based greedy heuristics and only take seconds to apply to graphs with tens of thousands of vertices. Moreover, our probabilistic policies outperform all greedy state-of-the-art coloring baselines and a machine learning baseline. Finally, we show that our approach also generalizes to other problems by evaluating it on

Attention Weights Change

Figure 1: *Spatial locality of the decoding.* After labeling a node, only its neighbors' attention weights change. The example shows how a graph is 2-colored using the vertex order c, e, b, a, d. The nodes whose attention weights change have a box around them. For example, when the first node c is

Neural Graph Databases

…atrick Iff[1] Florian Scheidl[1] Kazuki Osawa[1] Nikoli Dryden[1]
…Podstawski[2,3] Tiancheng Chen[1] Torsten Hoefler[1,†]
[1]Department of Computer Science, ETH Zurich
…Warsaw University of Technology, Warsaw, Poland
[3]TCL Research Europe, Warsaw, Poland
[†]Corresponding authors:
…aciej.besta, torsten.hoefler}@inf.ethz.ch

Abstract

Graph databases (GDBs) enable processing and analysis of unstructured, complex, rich, and usually vast graph datasets. Despite the large significance of GDBs in both academia and industry, little effort has been made into integrating them with the predictive power of graph neural networks (GNNs). In this work, we show how to seamlessly combine nearly any GNN model with the computational capabilities of GDBs. For this, we observe that the majority of these systems

The future may not be big graphs: embeddings to represent relations!

- **E.g., cosine distance as metric for "connectedness"**
  - Equivalent to vector angle for normalized vectors

    orthogonal vectors -> not connected, collinear vectors -> maximally connected, and everything in between ☺
- **Used in Heavily used in ML today**
  - Basis of attention mechanisms – e.g., transformers
  - Nice binary coding possible – get to $O(1)$ bits
  - nice tradeoff between accuracy and overhead

ProbGraph: High-Performance and High-Accuracy Graph Mining with Probabilistic Set Representations

Maciej Besta[1†], Cesare Miglioli[2], Paolo Sylos Labini[3], Jakub Tětek[4], Patrick Iff[1], Raghavendra Kanakagiri[5], Saleh Ashkboos[1], Kacper Janda[6], Michał Podstawski[7,8], Grzegorz Kwaśniewski[1], Niels Gleinig[1], Flavio Vella[9], Onur Mutlu[1], Torsten Hoefler[1†]

[1]ETH Zurich [2]Research Center for Statistics, University of Geneva [3]Free University of Bozen-Bolzano [4]BARC, University of Copenhagen [5]UIUC [6]AGH-UST [7]Warsaw University of Technology [8]TCL Research Europe [9]University of Trento [†]Corresponding authors

Abstract—Important graph mining problems such as Cluster- … Moreover, there are many heuristics for approximating

Slim Graph: Practical Lossy Graph Compression for Approximate Graph Processing, Storage, and Analytics

Maciej Besta, Simon Weber, Lukas Gianinazzi, Robert Gerstenberger, Andrey Ivanov, Yishai Oltchik, Torsten Hoefler
Department of Computer Science; ETH Zurich

ABSTRACT

We propose Slim Graph: the first programming model and framework for practical lossy graph compression that facilitates high-performance approximate graph processing, storage, and analytics. Slim Graph enables the developer to express numerous compression schemes using small and programmable compression kernels that can access and modify require unprecedented amounts of compute power, storage, and energy. For example, running PageRank on the Sogou webgraph using 38,656 compute nodes (10,050,560 cores) on the Sunway TaihuLight supercomputer [71] (nearly the full scale of TaihuLight) takes 8 minutes [101]. The sizes of such datasets will continue to grow; Sogou Corp. expects a ≈60 trillion edge graph dataset with whole-web crawling. *Lower-*

Sparsity in Deep Learning: Pruning and growth for efficient inference and training in neural networks

Torsten Hoefler HTOR@INF.ETHZ.CH
ETH Zürich
Zürich, Switzerland, 8092

Dan Alistarh DAN.ALISTARH@IST.AC.AT
IST Austria
Klosterneuburg, Austria, 3400

Tal Ben-Nun TALBN@INF.ETHZ.CH
Nikoli Dryden NDRYDEN@INF.ETHZ.CH
ETH Zürich
Zürich, Switzerland, 8092

Alexandra Peste ALEXANDRA.PESTE@IST.AC.AT
IST Austria
Klosterneuburg, Austria, 3400

Editor: Samy Bengio

Abstract

The growing energy and performance costs of deep learning have driven the community to reduce the size of neural networks by selectively pruning components. Similarly to their biological counterparts, sparse networks generalize just as well, sometimes even better than, the original dense networks. Sparsity promises to reduce the memory footprint of regular

Big data belongs in (knowledge) graphs!

We briefly explained why and how.

Prof. Ana-Lucia Varbanescu
University of Twente, NL

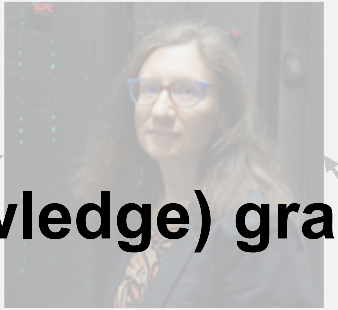

Prof. Angela Bonifati
University of Lyon1

ETH Zurich

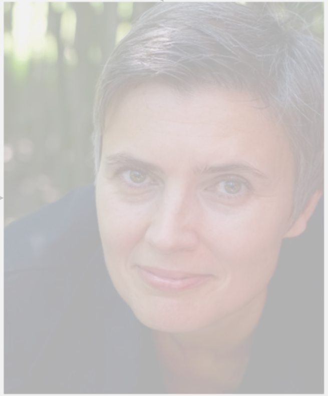

Prof. Adriana Iamnitchi,
University of Maastricht, NL

Time for your questions.