Big Graph Processing Systems

Organisation and **Motivation**

Christopher Spinrath

CNRS – LIRIS – Lyon 1 Université

DISS Master 2025

This presentation is an adaption of slides from Angela Bonifati

Outline

- Motivation
- Part I: Graph Query Paradigms and their Semantics
- Part II: Property Graphs
  - Chapter 1: An introduction to a concrete query language
  - Chapter 2: PG Schema Languages and PG Constraints
  - Chapter 3: Schema Discovery and Property Graph Transformations
- Part III: RDF and Sparql

Organisation

Tentative Schedule

January	06/01	CM + CM	Part I
	13/01	CM + TP	Part II
	20/01	CM + TP	Part II
	27/01	CM + TP	Part II
February	03/02	TP + TP	Part II
	10/02	TP + TP	Part II
	17/02	TP + TP	Part II
	24/02	CM + CM	Part III
	25/02	CM + TP	Part III
March	03/03	CM (Exam)	

- 15 CM + 15 TP
- Please ask (and answer) questions

Organisation

Tentative Schedule

January	06/01	CM + CM	Part I
	13/01	CM + TP	Part II
	20/01	CM + TP	Part II
	27/01	CM + TP	Part II
February	03/02	TP + TP	Part II
	10/02	TP + TP	Part II
	17/02	TP + TP	Part II
	24/02	CM + CM	Part III
	25/02	CM + TP	Part III
March	03/03	CM (Exam)	

- 15 CM + 15 TP
- Please ask (and answer) questions
- Grading
  - 50% Practical lab (TP)
  - 50% Exam (CM)

Hands-on Part – Preparations

- The warm-up project consists of analysing a dataset containing genomic information
- Please install Neo4j and import the dataset before the first hands-on lesson

Hands-on Part – Preparations

- The warm-up project consists of analysing a dataset containing genomic information
- Please install Neo4j and import the dataset before the first hands-on lesson

1. Install Neo4j
   - Make sure you are using version 5.26 or newer
   - Make sure you have enough memory, at least 20GB
   - A simple option is to install Neo4j Desktop by follwowing the official instructions: `https://neo4j.com/docs/desktop-manual/current/installation/`
2. Download and import the database
   - Download the database dump file from `https://partage.liris.cnrs.fr/index.php/s/LoEtp24fk38P6n5`
   - Import the dump file by following the official instructions: `https://neo4j.com/docs/desktop-manual/current/operations/create-from-dump/`

Everything is Data

[https://www.acronis.com/en-us/blog/posts/data-everything-8-noble-truths/]

A New Era of Science

The quest for knowledge used to begin with grand theories. Now it begins with massive amounts of data.

Welcome to the ~~Peta~~Zettabyte Age.

New Realities

- Everything is data
- Rise of data-driven culture
- High-performant data analytics
- Exploit sophisticated statistical methods

How do we structure/implement/live with this trend?

Level of Analytics

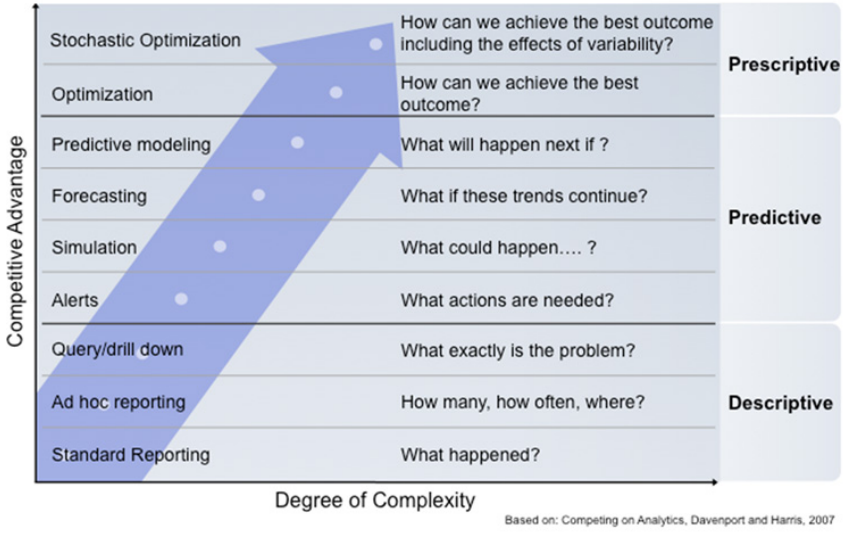

Based on: Competing on Analytics, Davenport and Harris, 2007

Focus of Interest

Properties of Entities

- Captured/measured values
- What are the sales figures/temperatures/etc.?
- Multidimensional data/time series/matrices

Connections Between Entities

- Network structure
- What do the friends of your customers buy?
- Graph data

Connectedness

Connections Manifest in Many Different Ways, Facets, Values, Scopes, and Scales

Cause: interaction, proximity, information, relationship, affiliation, ownership, allegiance, force, repulsion, ...

Nature: social, cultural, economical, physical, chemical, ...

Quality: existential, essential, inconsequential, or indiscernible; persistent or temporary; positive, neutral, or negative

Scale: from interacting cells, such as neurons in our brain, to high-scale entities, such as families, groups, clubs, companies, organizations, states, nations, ... or stars and galaxies

```
(graphs)-[:ARE]->(everywhere)
```

Graph Building Blocks

Nodes (Dots, Circles)

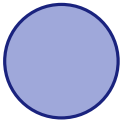

- Like an entity in conceptual models
- Exist on their own
- Have an object identity

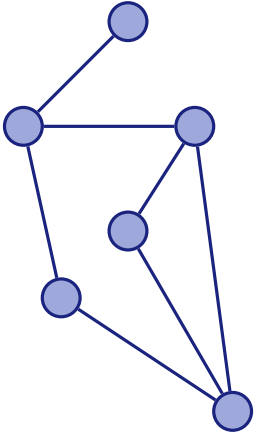

Edges (Lines)

- Like a relationship in conceptual models
- Exist only between nodes
- Identity depends on the nodes they connect

Graphs as Unifying Abstractions

- Graphs are natural abstractions
  for representing interconnected objects when encoding, explaining and predicting real-world and digital-world phenomena.
- Graphs are underpinning several data management ecosystems,
  in societal, scientific, RDF, product and digital domains.
- There is no unique killer application for graphs,
  but several exist.
- Nevertheless, the data models, query languages and system requirements
  needed for graphs are constantly evolving.

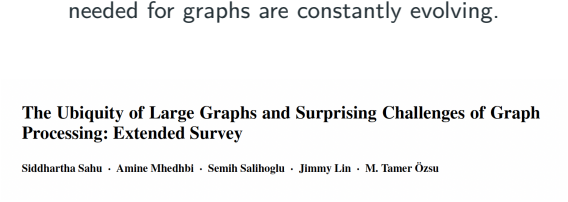
The Ubiquity of Large Graphs and Surprising Challenges of Graph Processing: Extended Survey

Siddhartha Sahu · Amine Mhedhbi · Semih Salihoglu · Jimmy Lin · M. Tamer Özsu


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:
Introduction
Key Insights
Abstractions
Ecosystems
Performance

Graph Data – Transportation Network

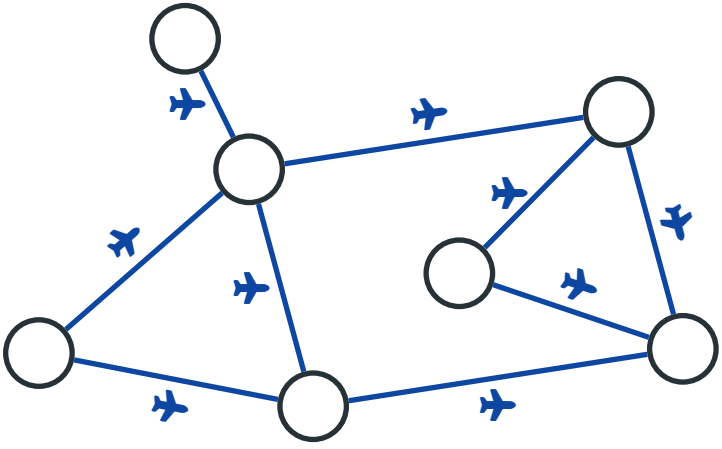

Graph Data – Transportation Network

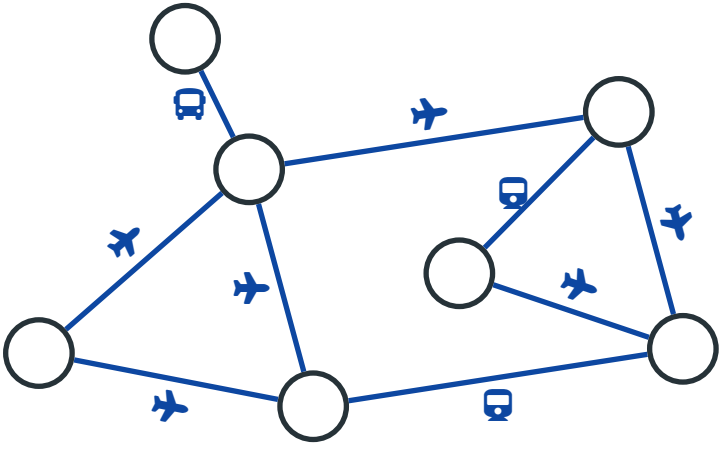

Graph Data – Transportation Network

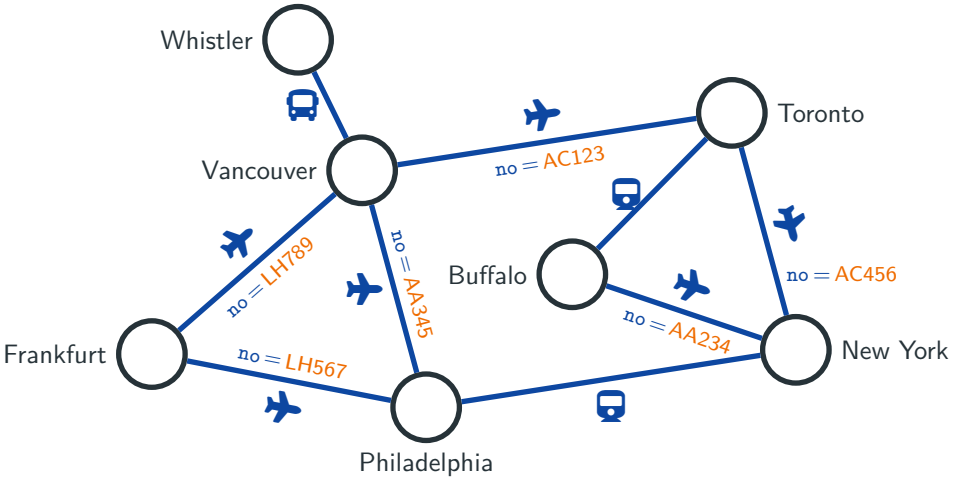

Graph Data – Social Network

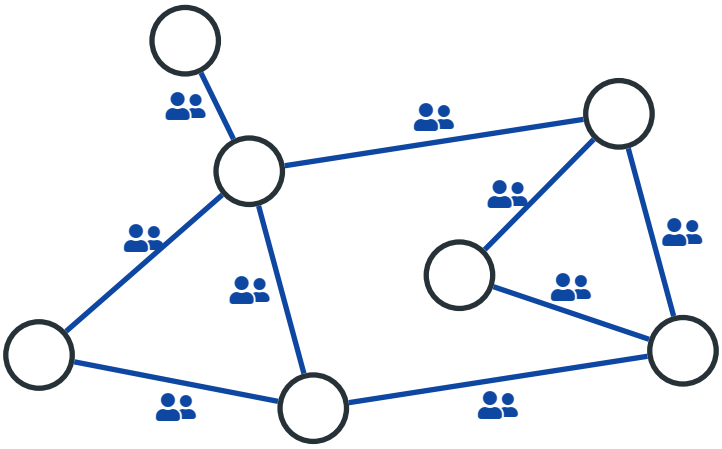

Graph Data – Social Network

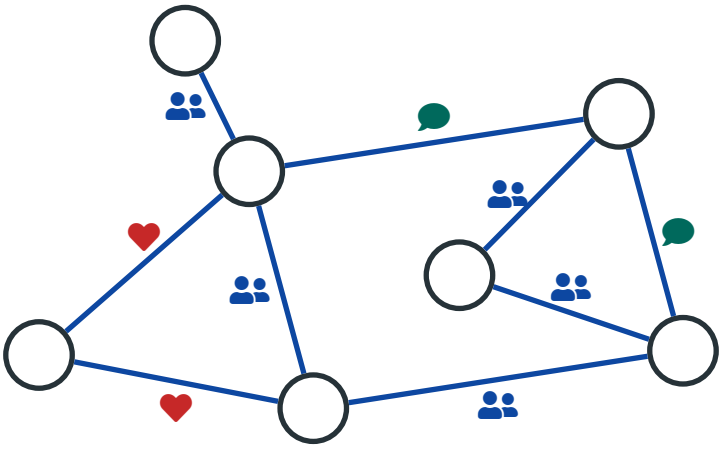

Graph Data – Social Network

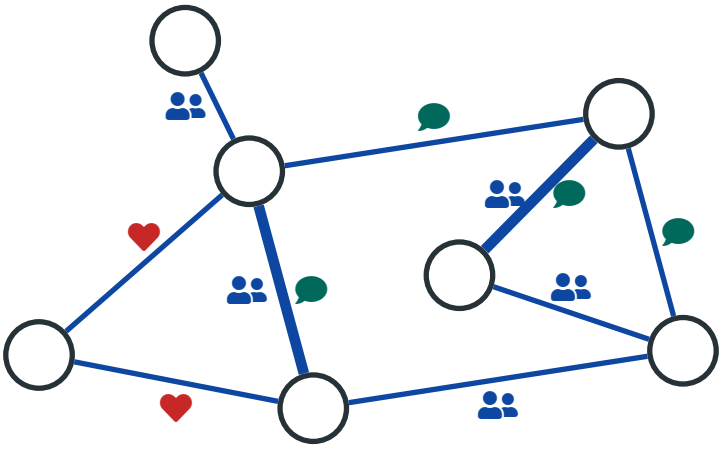

Graph Data – Heterogeneity

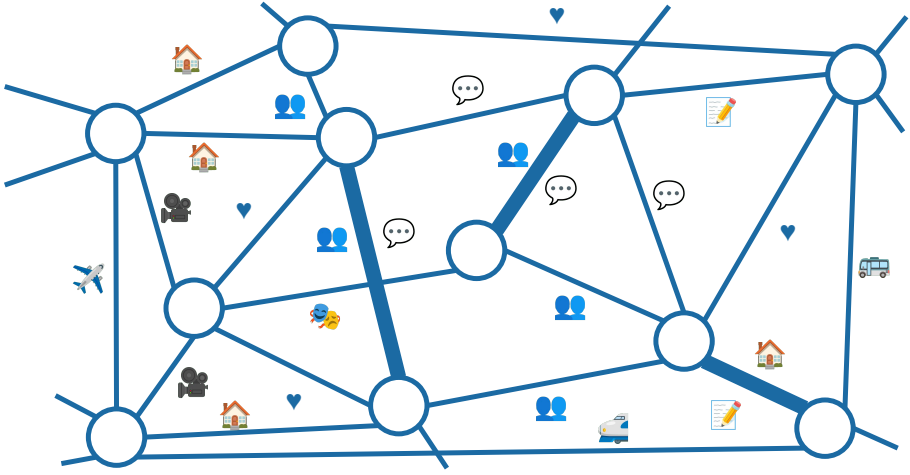

Example (Facebook Graph Search)

- Finding subgraph structures
- Very natural way of formulating queries

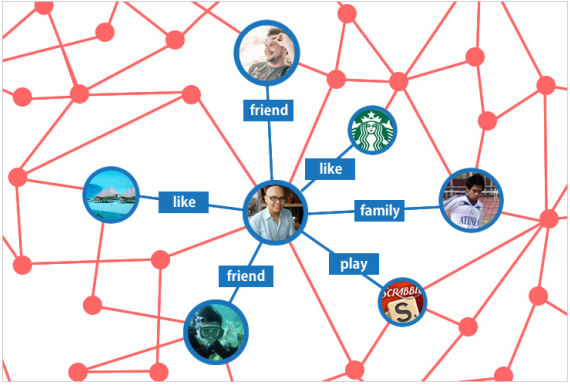

Web Graph

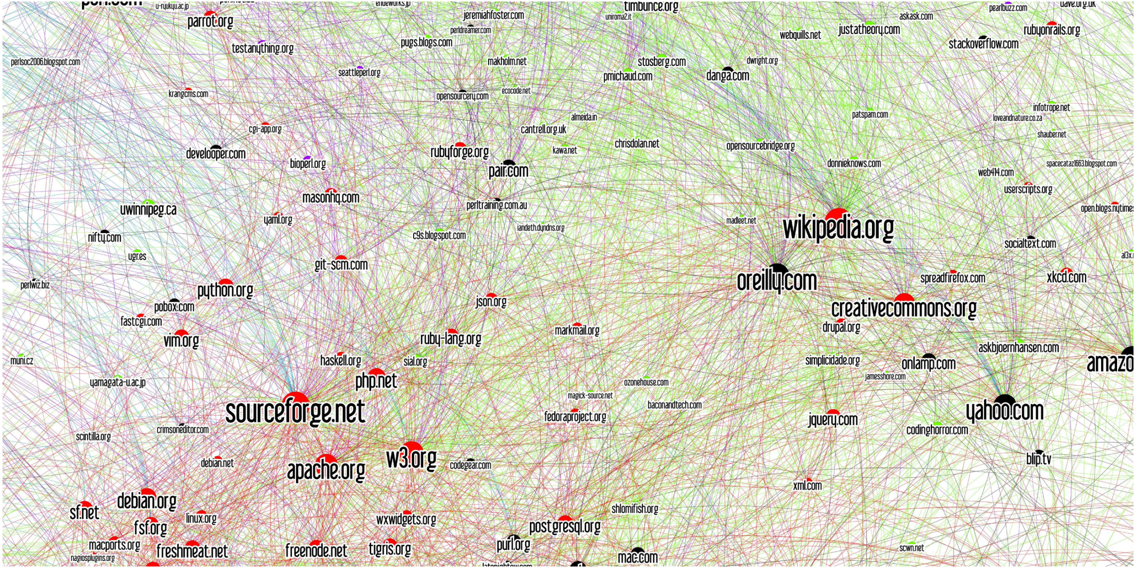
parrot.org
testanything.org
pugs.blogs.com
jeremiahfoster.com
timbunce.org
webquills.net
justatheory.com
askask.com
rubyonrails.org
stackoverflow.com
perlsoc2006.blogspot.com
seattleperl.org
makholm.net
stosberg.com
pmichaud.com
danga.com
dwright.org
krangcms.com
opensourcery.com
ecocode.net
cgi-app.org
almeida.in
cantrell.org.uk
patspam.com
infotrope.net
loveandnature.co.za
shaiber.net
chrisdolan.net
opensourcebridge.org
developer.com
bioperl.org
rubyforge.org
kawa.net
pair.com
donnieknows.com
web414.com
userscripts.org
masonhq.com
perltraining.com.au
uwinnipeg.ca
yaml.org
madleet.net
wikipedia.org
iandeth.dyndns.org
c9s.blogspot.com
nifty.com
socialtext.com
ugr.es
git-scm.com
oreilly.com
spreadfirefox.com
xkcd.com
perlwiz.biz
python.org
json.org
creativecommons.org
pobox.com
fastcgi.com
drupal.org
vim.org
ruby-lang.org
markmail.org
askbjoernhansen.com
muni.cz
haskell.org
sial.org
simplicidade.org
onlamp.com
php.net
jamesshore.com
yamagata-u.ac.jp
ozonehouse.com
magick-source.net
sourceforge.net
baconandtech.com
fedoraproject.org
jquery.com
yahoo.com
scintilla.org
crimsoneditor.com
codinghorror.com
debian.net
w3.org
apache.org
codegear.com
blip.tv
xml.com
debian.org
sf.net
shlomifish.org
fsf.org
linux.org
wxwidgets.org
postgresql.org
macports.org
freshmeat.net
freenode.net
tigris.org
purl.org
mac.com
scwn.net
nagiosplugins.org

Example (DBLP)

- Open bibliographic information on major computer science journals and proceedings
- >6.5 million publication
- >46000 new publication per month
- >1.7 million authors

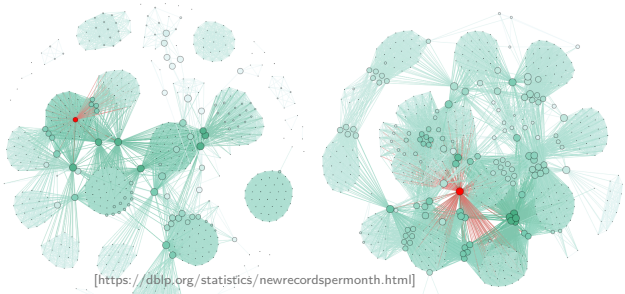

[https://dblp.org/statistics/newrecordspermonth.html]

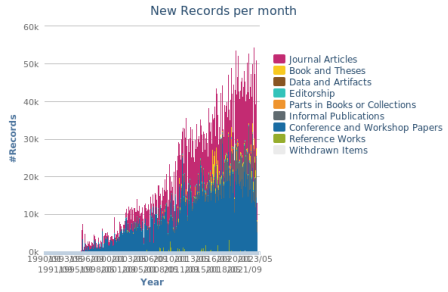

Example (Graph to capture world-knowledge)

- Open knowledge base that can be read and edited by humans and machines
- Structured data of Wikipedia, Wikivoyage, Wikisource, etc.

WIKIDATA

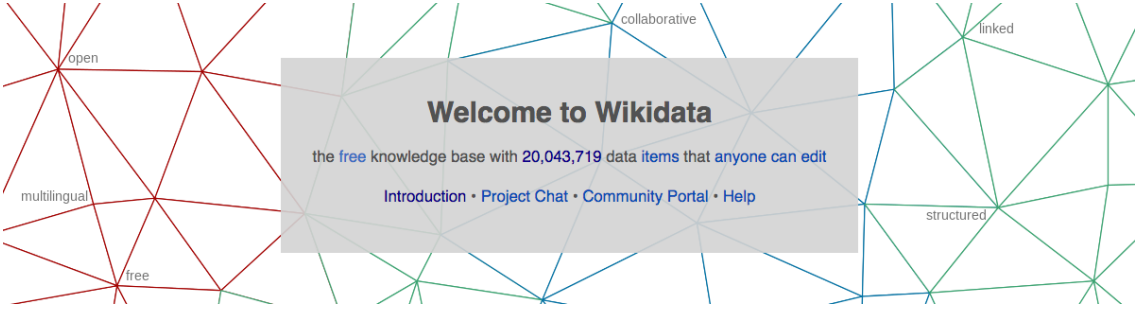

Shortest Paths

Example (Transportation Network)

- Only hop-distances (#edges)
  - in friendship network
- With weights
  (distance, travel time, etc.)
  - in road networks, transportation connections

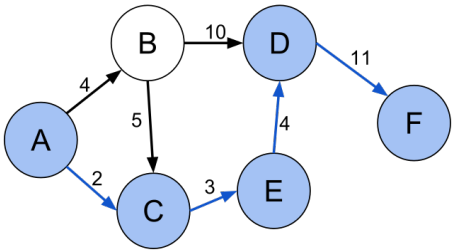

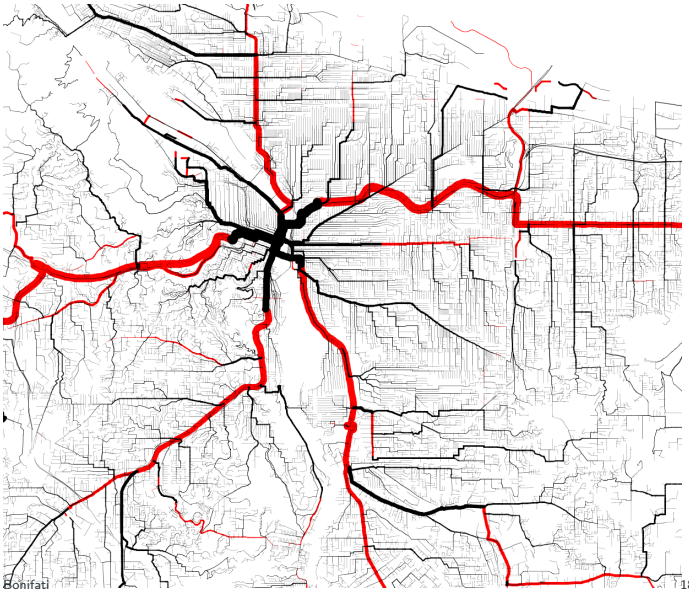

A Plethora of Applications

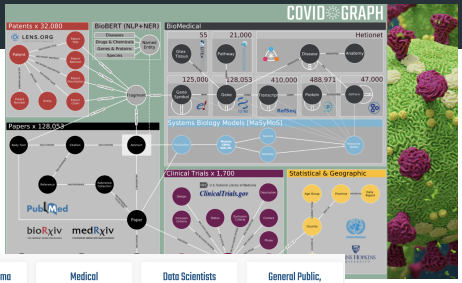

- Among which, the `covidgraph.org` initiative aiming at building the Covid19 knowledge graph
  - Collecting patents, publications about the human coronaviruses
  - Biomedical data (genomics and omics)
  - Experimental data about clinical trials
  - Key demographic indicators
- Practical use case in many data-oriented tasks
  - Property graph schema discovery
  - Threshold queries in Theory and in the Wild

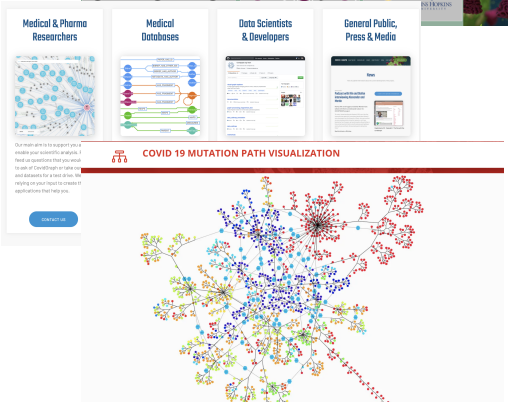

Threshold graph queries on the Covid19 graph

- Find each country that does not have three reports for some age group

```
MATCH (c:Country)
         -[e:CURRENT_FEMALE | CURRENT_MALE | CURRENT_TOTAL]->
      (a:AgeGroup)
WITH c, a, COUNT(type (e)) AS ecount
WHERE ecount < 3 RETURN c, a
```

Adapted from

[Bon22a] Angela Bonifati, Stefania Dumbrava, George Fletcher, Jan Hidders, Matthias Hofer, Wim Martens, Filip Murlak, Joshua Shinavier, Slawek Staworko, Dominik Tomaszuk: Threshold Queries in Theory and in the Wild. Proc. VLDB Endow. 15(5): 1105-1118 (2022)

Threshold graph queries on the Covid19 graph

- Find each country that does not have three reports for some age group

```
MATCH (c:Country)
        -[e:CURRENT_FEMALE | CURRENT_MALE | CURRENT_TOTAL]->
      (a:AgeGroup)
WITH c, a, COUNT(type (e)) AS ecount
WHERE ecount < 3 RETURN c, a
```

- Find each protein that has more than 43917 associated geneontology terms

```
MATCH (p:Protein )-[:MAPS]->*()-[:HAS_ASSOCIATION]->()
      (()-[:IS_A]->*()|()-[:PART_OF]->*())(t:GOTerm)
WITH p, COUNT (DISTINCT t) AS count_go
WHERE count_go > 43917 RETURN p
```

Adapted from

[Bon22a] Angela Bonifati, Stefania Dumbrava, George Fletcher, Jan Hidders, Matthias Hofer, Wim Martens, Filip Murlak, Joshua Shinavier, Slawek Staworko, Dominik Tomaszuk: Threshold Queries in Theory and in the Wild. Proc. VLDB Endow. 15(5): 1105-1118 (2022)

Property Graphs – Example

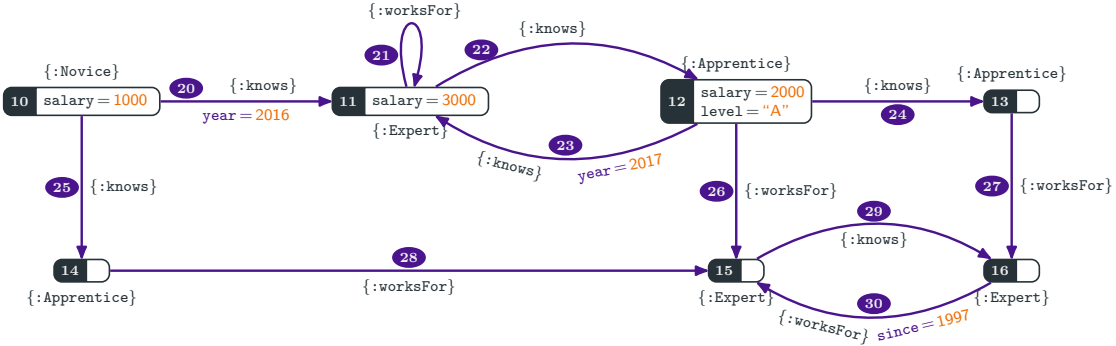

Property Graphs – Example

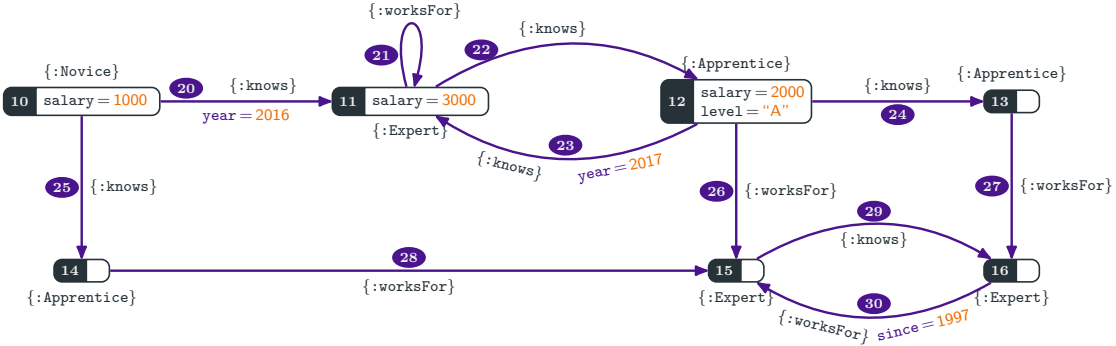

Nodes and edges have

- IDs: 10, 11, ...
- labels, e.g., `:Novice`, `:Apprentice`, `:worksFor`, ...
- properties, e.g. "`salary` = 3000", "`year` = 2016", ...

Property Graphs: A formal definition

Assume pairwise disjoint sets of
$\mathcal{O}$ (objects), $\mathcal{L}$ (labels), $\mathcal{K}$ (property keys), and $\mathcal{N}$ (values)

Definition

A property graph is a structure $(V, E, \eta, \lambda, \vartheta)$ where

- $V \subseteq \mathcal{O}$ is a finite set of vertices,
- $E \subseteq \mathcal{O}$ is a finite set of edges,
- $\eta\colon E \to V \times V$ assigns an ordered pair of vertices to each edge,
- $\lambda\colon V \cup E \to \mathcal{P}(\mathcal{L})$ assigns a finite set of labels to each vertex and edge,
- $\vartheta\colon (V \cup E) \times \mathcal{K} \rightharpoonup \mathcal{N}$ assigns values for properties to vertices and edges.

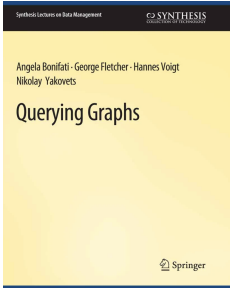

A Lattice of Graph Data Models

- A data model per use case
- How expressive and human-friendly is a data model?
- Need of making different data models interoperable via mappings or direct translations