


# Kùzu

Graph Database Management System

# What are graphs?

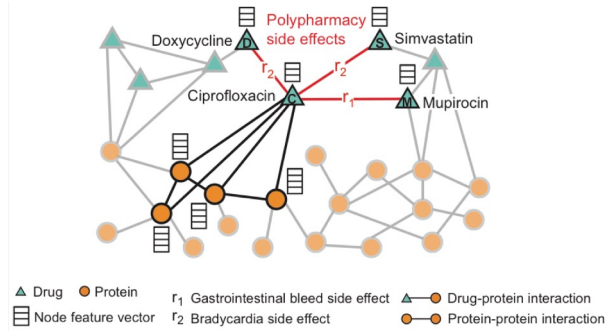
*Abstract representation of entities and relationships*



# Graphs: Natural ways to represent data



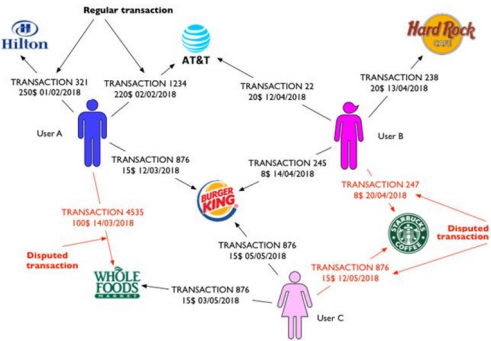
Social network




Drug interactions



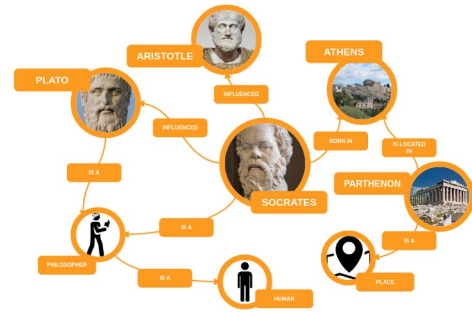
Traffic networks



Transaction graph



Molecular networks



Knowledge graphs

## Pattern Matching

- Pre-defined joins
- Complex pattern

## Recursive Join

- Different algorithms
- Parallel computation
- Skewness handling

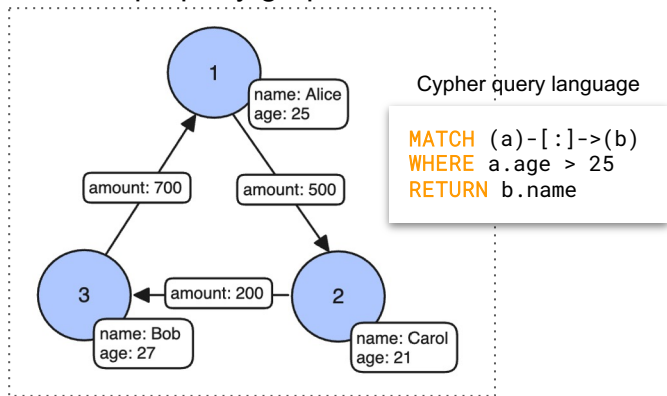
## Relational Operation

- Filter
- Projection
- ...

# What is Kùzu



## Structure property graph data model




## Embeddable (similar to DuckDB/SQLite)

```
import kuzu

db = kuzu.Database("db")
conn = kuzu.Connection(db)
res = conn.execute("MATCH (a)-[:]->(b)")
print(res.get_as_df())
```




## ACID transactions



## Permissively licensed



## Integrations with ML/AI frameworks




## Rich bindings



Learn more at <https://kuzudb.com>

Interoperable with databases  
and data formats

# History of Kùzu



	<b>Name</b>	<b>Age</b>
	Jean-Christophe	20
Node Group 0	Karl	
	...	...
Node Group 1	Alice	29
	Dave	30
Node Group ...	...	...

# History of Kùzu




## RDF support

```
CREATE RDFGRAPH KuzuKG;  
  
COPY KuzuKG FROM 'kuzu_kg.ttl';  
  
WITH "http://kuzu.io/rdf-ex#" as kz  
MATCH (s {iri: kz+"Waterloo"})-[p]->(o)  
RETURN s.iri, p.iri, o.iri;
```

Initial commit  
Sep 2020

Incorporated  
Mar 2023

0.2.0  
Feb 2024



Nov 2022

Nov 2023

Open sourced

0.1.0

## NodeGroup-based storage

	Name	Age
Node Group 0	Jean-Christophe	20
	Karl	
	...	...
Node Group 1	Alice	29
	Dave	30
Node Group ...	...	...

# History of Kùzu



## RDF support

```
CREATE RDFGRAPH KuzuKG;  
  
COPY KuzuKG FROM 'kuzu_kg.ttl';  
  
WITH "http://kuzu.io/rdf-ex#" as kz  
MATCH (s {iri: kz+"Waterloo"})-[p]->(o)  
RETURN s.iri, p.iri, o.iri;
```

Initial commit  
Sep 2020

Incorporated  
Mar 2023

0.2.0  
Feb 2024



Nov 2022

Nov 2023

May 2024

Open sourced

0.1.0

0.4.0

NodeGroup-based storage

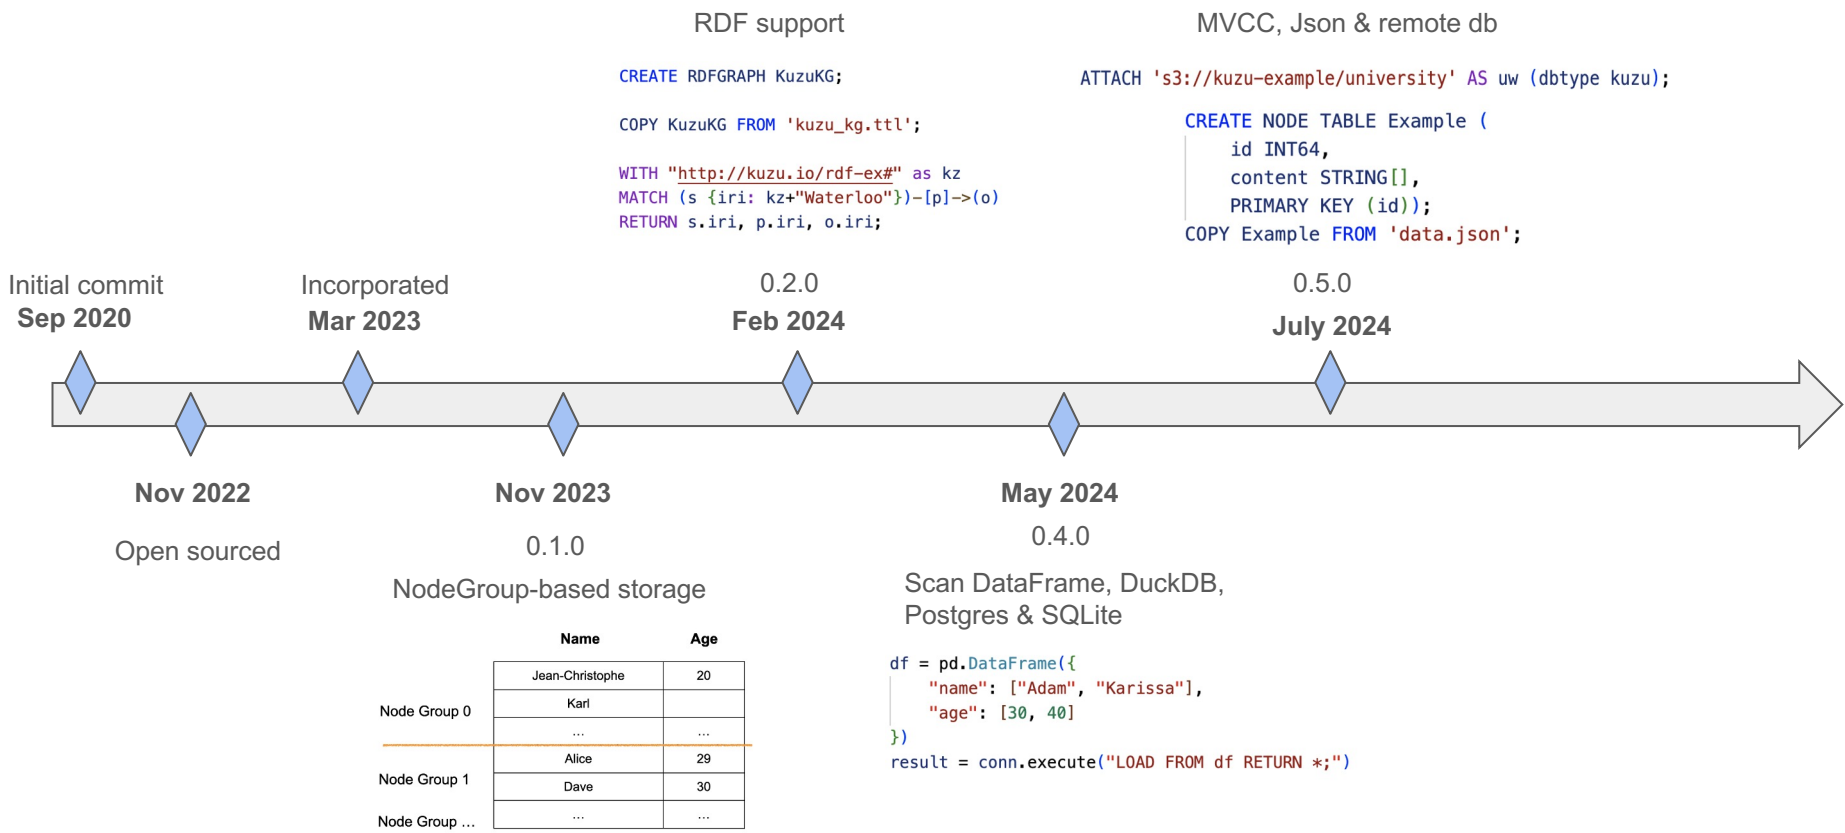
	Name	Age
Node Group 0	Jean-Christophe	20
	Karl	
	...	...
Node Group 1	Alice	29
	Dave	30
Node Group ...	...	...

Scan DataFrame, DuckDB,  
Postgres & SQLite

```
df = pd.DataFrame({  
    "name": ["Adam", "Karissa"],  
    "age": [30, 40]  
})  
result = conn.execute("LOAD FROM df RETURN *;")
```



# History of Kùzu



# History of Kùzu




RDF support

```
CREATE RDFGRAPH KuzuKG;  
  
COPY KuzuKG FROM 'kuzu_kg.ttl';  
  
WITH "http://kuzu.io/rdf-ex#" as kz  
MATCH (s {iri: kz+"Waterloo"})-[p]->(o)  
RETURN s.iri, p.iri, o.iri;
```

MVCC, Json & remote db

```
ATTACH 's3://kuzu-example/university' AS uw (dbtype kuzu);  
  
CREATE NODE TABLE Example (  
  id INT64,  
  content STRING[],  
  PRIMARY KEY (id));  
COPY Example FROM 'data.json';
```



NodeGroup-based storage

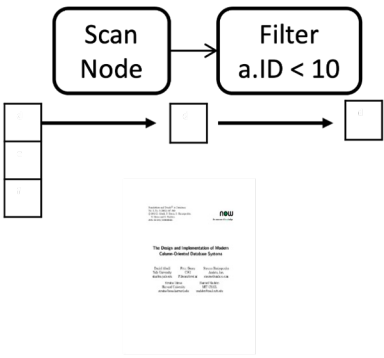
	Name	Age
Node Group 0	Jean-Christophe	20
	Karl	
	...	...
Node Group 1	Alice	29
	Dave	30
Node Group ...	...	...

Scan DataFrame, DuckDB, Postgres & SQLite

```
df = pd.DataFrame({  
  "name": ["Adam", "Karissa"],  
  "age": [30, 40]  
})  
result = conn.execute("LOAD FROM df RETURN *;")
```

In-memory mode

### Vectorization



### Factorization

<u>b</u> , name		<u>a</u>		<u>c</u>
{L <sub>1</sub> , Liz}	X	{U <sub>1</sub> , ..., U <sub>100</sub> }	X	{C <sub>1</sub> , ..., C <sub>100</sub> }
U				
{L <sub>2</sub> , Liz}	X	{U <sub>100</sub> , ..., U <sub>199</sub> }	X	{C <sub>100</sub> , ..., C <sub>199</sub> }



### Morsel-driven Parallelism

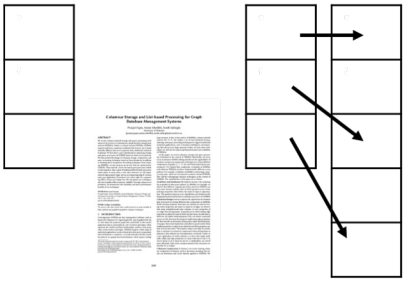
➤ Scans are "morselized" across threads



### Novel Join Algorithms

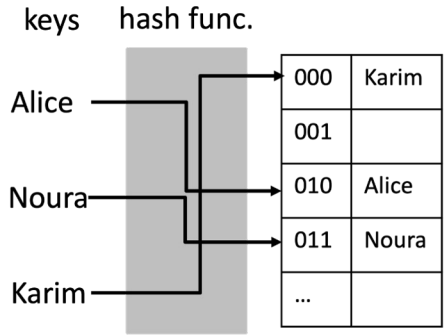


### Disk-based Columnar Storage




node properties      CSR-based join indices & edge props

### Disk-based Hash Index



# Factorized query processing



a1	b1	c1
...	...	...
a1	b1	c10
a10	b1	c1
...	...	...
a10	b1	c10


100 tuples

Flat representation

a1, ..., a10	b1	c1, ..., c10
--------------	----	--------------


Factorized representation

# CSR index




### FWD CSR

Node Group 0




Node Group 1




### BWD CSR

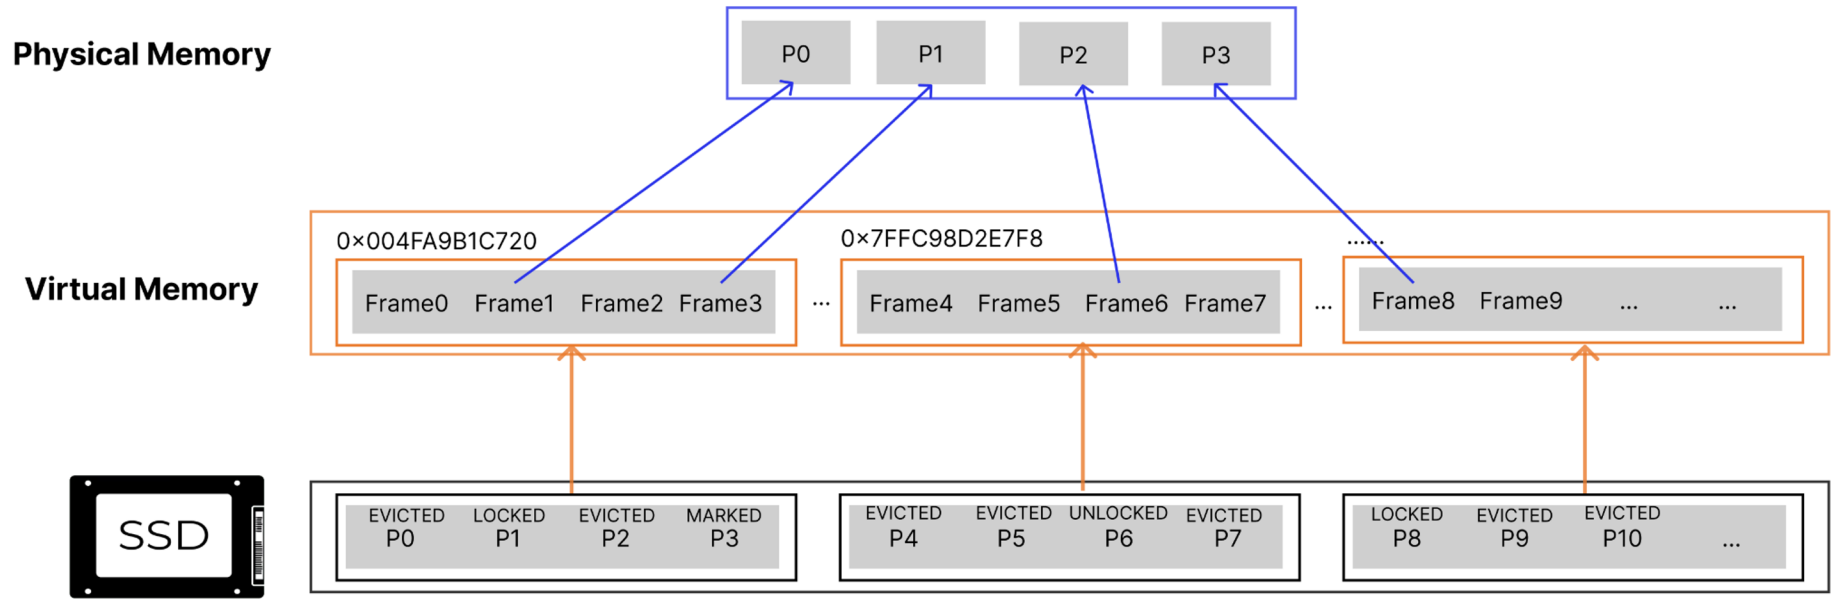
Node Group 0



Node Group 1



# VM-backed buffer manager




# Roadmap - Direct Query over RDBMS



Account		
ID	name	country
1	Alice	CA
2	Carol	USA
3	Bob	CA
...	...	...

eTransfer		
from	to	amount
1	2	700
2	3	800
3	1	900
...	...	...



```
CREATE NODE TABLE Account (  
  ID INT64,  
  name STRING,  
  country STRING,  
  PRIMARY KEY (id)  
);
```

```
CREATE REL TABLE eTransfer (  
  FROM Person TO Person,  
  amount INT64  
);
```

```
ATTACH duck.db AS duck (dbtype duckdb);  
Copy Account FROM duck.Account;  
Copy eTransfer FROM duck.eTransfer;
```

Require graph modeling & data moving

# Roadmap - Direct Query over RDBMS




Account		
ID	name	country
1	Alice	CA
2	Carol	USA
3	Bob	CA
...	...	...

eTransfer		
from	to	amount
1	2	700
2	3	800
3	1	900
...	...	...

```
ATTACH duck.db AS duck (dbtype duckdb);  
Create Node Table Account AS duck.Account;  
Create REL Table eTransfer AS duck.eTransfer;
```

```
MATCH (p:Account)-[:eTransfer]->(p2:Account)  
WHERE a.name = 'Alice'  
RETURN p2.name
```



- Data stored in DuckDB
- Computation happens in Kùzu
- Extendable to Postgres, SQLite, CSV, ...



- Parallelize within single source node
  - Handle skew
- Extendable framework
  - Shortest path
  - Weight shortest path
  - Connected component



Join our Discord!



[github.com/kuzudb/kuzu](https://github.com/kuzudb/kuzu)



[@kuzudb](https://twitter.com/kuzudb)



<https://kuzudb.com/blog/>