# **Graphing Your Way Through the Cosmos**

## Chad Green

CodeMash January 9, 2020

#### Who is Chad Green

Director of Software Development ScholarRx



chadgreen@chadgreen.com

whadgreen.com

ChadGreen

in ChadwickEGreen



# What are Graph Databases

Graphing Your Way Through the Cosmos

- Discrete mathematics
- Structure amounting to a set of objects in which some pairs of the objects are in some sense related
- Objects correspond to mathematical abstractions called vertices and each of the related pairs of vertices is called an edge
- Graph Theory is the study of graphs

• Depicted in diagrammatic form as a set of dots or circles for the vertices,

joined by lines or curves for the edges





## History of Graph Theory



#### History of Graph Theory



#### History of Graph Theory



# Leonard Euler



History of Graph Theory



Solutio problematis ad geometriam situs pertinentis

The solution of a problem relating to the geometry of position

## Applications of Graph Theory

- Linguistics
- Physics and Chemistry
- Social Sciences
- Biology
- Computer Science





- Collection of *vertices* and *edges*
- Represent entities as vertices and the ways in which those entities relate to the world as relationships
- Allow us to model all kinds of scenarios



What is a Graph Database

## A graph database is a database

## that uses graph structures to

represent and store data.

#### What is a Graph Database

- Represents data as it exists in the real world that are naturally connected
- Does not try to change them in any way to define them as entities
- Graphs are composed of *vertices* and *edges*
- Vertices represent specific objects
- Edge is a relation between vertices
- Both vertices and edges can have any number of properties

#### Property Graph Model

Contains **nodes** (vertices) and **relationships** (edges) Nodes and relationships contain **properties** 

Relationships are **named** and **directed** with a **start** and **end** node

#### Employee

Name: Chad Green Location: Louisville, KY Title: Director of Software Development

Works For

Date of Employment: 2/28/2019

Company

Name: ScholarRx Location: Elizabethtown, KY

#### The Power of Graph Databases

• Performance

• Graph database performance tends to remain relatively constant, even as the dataset grows

• Flexibility

• Graph data model better accommodates changing business needs

• Agility

• Equip us to perform frictionless development and graceful system maintenance

• Governance is typically applied in a programmatic fashion

#### Common Graph Use Cases

- Internet of Things
- Customer 360
- Asset management
- Recommendations
- Fraud detection
- Data Integration
- Identity and access management

- Social networks
- Communication networks
- Genomics
- Epidemiology
- Semantic Web
- Search

# Graph vs Relational

Graphing Your Way Through the Cosmos

#### Graph Databases vs Relational Databases

Relational

Tables

Schema with nullables

Relations with foreign keys

Related data fetched with joins

Graph

Vertices (Nodes)

No schema

Relation is first class citizen

Related data fetched with a pattern

EmployeeId	EmployeeName	EmployeeGroup
1	Willis B. Hawkins	Sales
2	Neil S. Vega	Sales
3	Ada C. Lavigne	Engineering

Em	ployees
8	EmployeeID
	EmployeeName
	EmployeeGroup



```
-- Create the Employee Table

CREATE TABLE Employees

(

EmployeeID INT IDENTITY(1,1),

EmployeeName VARCHAR(64),

EmployeeGroup VARCHAR(32),

CONSTRAINT pkcEmployees PRIMARY KEY CLUSTERED (EmployeeId)

)

GO
```

```
-- Populate the Employee Table

INSERT INTO Employees (EmployeeName, EmployeeGroup)

VALUES ('Willis B. Hawkins', 'Sales'),

('Neil S. Vega', 'Sales'),

('Ada C. Lavigne', 'Engineering');
```

Graph Databases vs Relational Databases

```
// Create group nodes
g.addV('group').property('id', 'Sales')
g.addV('group').property('id', 'Engineering')
```

# // Create employee nodes g.addV('employee').property('id', 'Willis B. Hawkins') g.addV('employee').property('id', 'Neil S. Vega') g.addV('employee').property('id', 'Ada C. Lavigne')

// Create relationships between groups and employees
g.V('Sales').addE('member').to(g.V('Willis B. Hawkins'))
g.V('Sales').addE('member').to(g.V('Neil S. Vega'))
g.V('Engineering').addE('member').to(g.V('Ada C. Lavignee'))

Graph Databases vs Relational Databases

EmployeeId	EmployeeName	EmployeeGroup
1	Willis B. Hawkins	Sales
2	Neil S. Vega	Sales
3	Ada C. Lavigne	Engineering



8 documents (vertices and edges)

#### 3 rows, 3 columns

Graph Databases vs Relational Databases

EmployeeId	EmployeeName	EmployeeGroup
1	Willis B. Hawkins	Sales
2	Neil S. Vega	Sales
3	Ada C. Lavigne	Engineering



#### SELECT \* FROM Employees;

g.V().hasLabel('employee')

# **Company Reorganization**

```
-- Create the Groups table

CREATE TABLE Groups

(

GroupId INT IDENTITY(1,1),

GroupName VARCHAR(64),

CONSTRAINT pkcGroups PRIMARY KEY CLUSTERED (GroupId)
```



```
-- Create the Employee_Group join table
CREATE TABLE Employee_Group
```

```
GroupId INT,
EmployeeId INT,
CONSTRAINT pkcEmployeeGroup PRIMARY KEY CLUSTERED (GroupId, EmployeeId),
CONSTRAINT fkEmployeeGroup_Groups FOREIGN KEY (GroupId) REFERENCES Groups(GroupId),
CONSTRAINT fkEmployeeGroup_Employees FOREIGN KEY (EmployeeId) REFERENCES Employees(EmployeeId)
```

Graph Databases vs Relational Databases

-- Populate the Employee\_Group table from Employees and Groups
INSERT INTO Employee\_Group (GroupId, EmployeeId)
SELECT Groups.GroupId,
Employees.EmployeeId
FROM Employees,
Groups
WHERE Groups.GroupName = Employees.EmployeeGroup

Graph Databases vs Relational Databases

-- Drop the Employees.EmployeeGroup column that is no longer valid ALTER TABLE Employees DROP COLUMN EmployeeGroup

#### Graph Databases vs Relational Databases

EmployeeId	EmployeeName
1	Willis B. Hawkins
2	Neil S. Vega
3	Ada C. Lavigne
GroupId	GroupName
1	Engineering
2	Sales

GroupId	EmployeeId
1	3
2	1
2	2



Graph Databases vs Relational Databases

// Add link to existing node
g.V('Sales').addE('member').to(g.V('Ada C. Lavigne'))



#### Graph Databases vs Relational Databases

Employeeld	EmployeeName
1	Willis B. Hawkins
2	Neil S. Vega
3	Ada C. Lavigne

GroupId	GroupName
1	Engineering
2	Sales

GroupId	EmployeeId
1	3
2	1
2	2

Added 2 tables; 6 rows; 4 new columns Removed a column



+1 document

Graph Databases vs Relational Databases



g.V('Sales').outE('member').inV()

@chadgreen

SELECT Employees.EmployeeId, Employees.EmployeeName FROM Employees INNER JOIN Employee\_Group ON Employee\_Group.EmployeeId = Employees.EmployeeId INNER JOIN Groups ON Groups.GroupId = Employee\_Group.GroupId WHERE Groups.GroupName = 'Sales'
## Corporate Merger

LAR

Graph Databases vs Relational Databases

#### -- Create the new Product Group INSERT INTO Groups (GroupName) VALUES ('Product Group')



Graph Databases vs Relational Databases

-- Associate everyone to the new Product Group
INSERT INTO Employee\_Group (GroupId, EmployeeId)
SELECT Groups.GroupId,
 Employees.EmployeeId
FROM Groups,
 Employees
WHERE Groups.GroupName = 'Product Group

Graph Databases vs Relational Databases

```
-- Create the Group/Group union table CREATE TABLE Group_Group
```

ParentGroupId INT, ChildGroupId INT, CONSTRAINT pkcGroup\_Group PRIMARY KEY CLUSTERED (ParentGroupId, ParentGroupId), CONSTRAINT fkGroupGroup\_Groups\_Parent FOREIGN KEY (ParentGroupId) REFERNCES Groups(GroupId), CONSTRAINT fkGroupGroup\_Groups\_Child FOREIGN KEY (ChildGroupId) REFERNCES Groups(GroupId)

Graph Databases vs Relational Databases

```
-- Relate the child groups to the parent group
INSERT INTO Group_Group (ParentGroupId, ChildGroupId)
SELECT (SELECT GroupId FROM Groups WHERE GroupName = 'Product Group'),
        Groups.GroupId
FROM Groups
WHERE Groups.GroupName <> 'Product Group'
```



2

3

#### Graph Databases vs Relational Databases

EmployeeId	EmployeeName	GroupId	EmployeeId		
1	Willis B. Hawkins	1	3		
2	Neil S. Vega	2	1		
3	Ada C. Lavigne	2	2		
		2	3		
GroupId	GroupName	2	1		
1	Fngineering	3	T		
-		3	2		
2	Sales	3	3		
3	Product Group	3	5		
ParentGroupId ChildGroupId					
3	1				



Graph Databases vs Relational Databases

```
// Add supergroup node
g.addV('group').property('id', 'Product Group')
```

```
// Link to adjacent nodes
g.V('Product Group').addE('contains_subgroup').to(g.V('Engineering'))
g.V('Product Group').addE('contains_subgroup').to(g.V('Sales'))
```



#### Graph Databases vs Relational Databases

EmployeeId	EmployeeName	GroupId	EmployeeId	
1	Willis B. Hawkins	1	3	
2	Neil S. Vega	2	1	
3	Ada C. Lavigne	2	2	
-		2	3	
GroupId	GroupName	3	1	
1	Engineering	3	2	
2	Sales	3	2	
3	Product Group	3	3	
ParentGroup	d ChildGroupId			
3	1			
3	3 2			



#### Added 1 table; 6 rows; 2 new columns

+3 documents

Graph Databases vs Relational Databases

GroupId	GroupName
1	Engineering
2	Sales



SELECT Groups.GroupId, Groups.GroupName FROM Groups INNER JOIN Group\_Group ON Group\_Group.ChildGroupId = Groups.GroupId WHERE Group\_Group.ParentGroupId = (SELECT GroupId FROM Groups WHERE GroupName = 'Product Group') g.V('Product Group')
.outE('contains\_subgroup')
.inV()

## Management

Graph Databases vs Relational Databases

```
-- Create the Employee/Employee join table CREATE TABLE Employee_Employee
```

ParentEmployeeId INT, ChildEmployeeId INT, CONSTRAINT pkcEmployeeEmployee PRIMARY KEY CLUSTERED (ParentEmployeeId, ChildEmployeeId), CONSTRAINT fkEmployeeEmployee\_Employee\_Parent FOREIGN KEY (ParentEmployeeId) REFERENCES Employees(EmployeeId), CONSTRAINT fkEmployeeEmployee\_Employee\_Child FOREIGN KEY (ChildEmployeeId) REFERENCES Employees(EmployeeId)

Graph Databases vs Relational Databases

```
-- Make Ada the boss

INSERT INTO Employee_Employee (ParentEmployeeId, ChildEmployeeId)

SELECT (SELECT EmployeeId FROM Employees WHERE EmployeeName = 'Ada C. Lavigne'),

EmployeeId

FROM Employees

WHERE EmployeeId IN (SELECT EmployeeId

FROM Employee_Group

WHERE Employee_GroupId = (SELECT GroupId

FROM Groups

WHERE GroupName = 'Sales'))
```



#### Graph Databases vs Relational Databases

Employeeld	EmployeeName	GroupId	EmployeeId
1	Willis B. Hawkins	1	3
2	Neil S. Vega	2	1
3	Ada C. Lavigne	2	2
		2	3
GroupId	GroupName	2	1
1	Engineering	5	T
2	Color	3	2
2	Sales	3	3
3	Product Group		
ParentGroup	d ChildGroupId	ParentEmployeel	d ChildEmployeeId
3	1	3	1
3	2	3	2
		3	3



Graph Databases vs Relational Databases

// Add relationships
g.V('Ada C. Lavigne').addE('has\_report').to(g.V('Willis B. Hawkins'))
g.V('Ada C. Lavigne').addE('has\_report').to(g.V('Neil S. Vega'))

#### Graph Databases vs Relational Databases

Employee	Id	EmployeeName		GroupId	Empl	oyeeld	
1 W		Willis B. Hawkins		1	3		
2		Neil S. Vega		2	1		
3		Ada C. Lavigne		2	2		
			2	3			
GroupId	GroupName			3	1	1	
1	Engineering			2	2		
2	Sales			3	Z		
-	Sales			3	3		
3	Product Group						
ParentGroupId ChildGroupId			ParentEmployeeId		ChildEmployeeId		
		ChildGroupid	ChildGroupid	3		1	
3		1				2	
3		2		3		2	
-				3		3	





+2 documents

Graph Databases vs Relational Databases

EmployeeName

Ada C. Lavigne

#### Ada C. Lavigne

SELECT DISTINCT EmployeeName
FROM Employees
INNER JOIN Employee\_Group
ON Employee\_Group\_EmployeeId = Employes.EmployeeId
INNER JOIN Employee\_Employee
ON Employee\_Employee.ParentEmployeeId = Employees.EmployeeId
WHERE Employee\_Group.GroupId = (SELECT GroupId
FROM Groups
WHERE GroupName = 'Engineering')

g.V('Engineering')
.outE('member')
.inV()
.outE('has\_report')
.values('id')

### **Challenges of Relational Databases**

Graph Databases vs Relational Databases

- Schema management
- Table alterations
- Costly writes against multiple tables
- Multiple JOIN operations
- Complex read queries



## What is Gremlin

Graphing Your Way Through the Cosmos

#### What is a TinkerPop

- Open source, vendor-agnostic, graph computing framework
- Apace2 license
- Allows users to model their domain as graph and analyze using Gremlin
- TinkerPop-enable systems integrate with one another



### What is a TinkerPop

- Gremlin
- Gremlin Console
- Gremlin Server
- TinkerGraph

- Programming Interfaces
- Documentation
- Useful Recipes



#### What is a Gremlin

- Graph traversal language and virtual machine
- Supports OLTP and OLAP
- Supports imperative and declarative querying
- Supports user-defined domain specified languages



# What is Cosmos DB

Graphing Your Way Through the Cosmos

A globally distributed, massively scalable, multi-model database service

Turnkey global distribution



A globally distributed, massively scalable, multi-model database service

Compegingensiale distablesion



A globally distributed, massively scalable, multi-model database service

## Elastipretalesivet of storageS&Athroughput



A globally distributed, massively scalable, multi-model database service

## Gua Etantie date datency of att three 98 pleroegtile ut



A globally distributed, massively scalable, multi-model database service

## Guariantered - defilatedncy atothsis 90 th cpencettelle



A globally distributed, massively scalable, multi-model database service

#### No schema or index management



A globally distributed, massively scalable, multi-model database service

Battle tested database service



A globally distributed, massively scalable, multi-model database service

#### Battle tested database service



Turnkey global

distribution

A globally distributed, massively scalable, multi-model database service

#### Ubiquitous regional presence



A globally distributed, massively scalable, multi-model database service

#### Secure by default and enterprise ready



A globally distributed, massively scalable, multi-model database service



# Exploring Graph Traversals

Graphing Your Way Through the Cosmos

#### Requirements for Speaking Engagement Management

- Where has a presentation been submitted?
- What presentations are tagged with a particular tag?
- Where have presentations tagged with a particular tag been accepted?
- What events has a speaker been accepted at?
- Where were the events a speaker has been accepted to?




### "Schema"



@chadgreen



### g.V()

Displaying 246 of 246 vertices and 500 of 1212 edges

#### @chadgreen

### View all of the edges

×) F	ile Edit Selection View Go Debug Terminal Help SpeakingEngagements - Visual Studio Code	
C	≣ SpeakingEngagements ×	Ш
Q	SpeakingEngagements / SpeakingEngagements	
ဖို	g.E() Execute	
逯	Graph JSON	
	<pre>{     "id": "49739ebc-fd37-47bf-a379-80c2d2009968",     "label": "partOf",     "type": "edg",     "inViabel": "state",     "uitWiabel": "state",     "uitWiabel": "state",     "uitWiabel": "Treformer the state is a part of"     "description": "The country the state is a part of"     "label": "countryLocation",     "type": "edg",     "inViabel": "country,     "uitWiabel": "event",     "inNotel": "sobcoff.ffs9-acc-9192-78b9654ac882",     "uitWiabel": "countryLocation",     "type": "edg",     "inViabel": "country.,     "uitWiabel": "country,     "uitWiabel": "country,     "uitWiabel": "country,     "uitWiabel": "country,     "uitWiabel": "country,     "inViabel": "sobcoff.ffs9-2cc-9192-78b9654ac882",     "uitWiabel": "sobcoff.ffs9-2cc-9192-78b9654ac882",     "uitWiabel": "stocotion",     "type": "edge",     "inViabel": "stocotion",     "uitWiabel": "stocotion",</pre>	
563		

### g.E()

### View the schema definition



### g.V() .has('ownerEmailAddress', 'schemaDefinition')

⊗ 0 ▲ 0 Azure: chadgreen@chadgreen.com

Displaying all 6 vertices and 6 edges

@chadgreen

### View the schema definition



### g.V() .has('ownerEmailAddress', 'schemaDefinition')

@chadgreen

⊗ 0 ▲ 0 Azure: chadgreen@chadgreen.com

### What presentations are in my repertoire?



g.V() .has('about (dep Eeseit & ations)s', 'chadgreen@chadgreen.com') .hasLabel('presentation')

### What presentations are in my repertoire?

×	File	Edit	Selection	View	Go Deb	oug Terminal	Help	SpeakingEr	ngagements - Vis	ual Studio Code			Ľ	
Ð	•	Speal	kingEngag	ements >										
Q	SpeakingEngagements / SpeakingEngagements													
ီစ	L	g.V().	has('owr	nerEma	ilAddres	s', 'chadgree	en@chadgreen.c	om').has	Execute					
逯	Graph JSON													
₿	[ "Which Microsoft Framework Am I Supposed to Use?", "Building Great Libraries with .NET Standard", "Black Box Magic".													
Δ		"Froi "Froi "Sof "Sof "Ser" "Get "Tim "Get "Azu "Azu "Azu "Pev "Azu "Pev "Azu "Cev "Azu "Cev "Cev "Sof "Cev "Cev "Cev "Cev "Cev "Cev "Cev "Cev	ck Box J m Zero <sup>2</sup> to be : rets of tware C. tware C. verless ting St. lding A: e Trave ting St. phing Y. re Dural re Serv. ellop a 0 ting Gr nt-Driv. ellop a 0 ting Gr mt-Driv. lding H Hitchh ivering Architec lding a ch-ch-cl thing Y re Does	Magic", to Server a Leade Confli raftsmæ raftsmæ raftsmæ raftsmæ raftsmæ n Act with a Server States of Real-1 cts Gui Cultur hanges: ou Can	verless er", cct Res f anship f ion", with Azu vith Co Serve vith Co Serve vith Co Serve vith Co Serve vith Co Serve vith Co Serve vith Co Serve vith Co Serve vith Co Serve vith Co Serve vith Serve	, , , , , , , , , , , , , ,	relopers", relopers", mg Azure Funct mporal Tables" , oss, eds to Know", tion in a Day" lata", cloud", sing Solutions s", in Azure Cosmo er)", Native?"	ions", ", estration" , in Azure" uny", os DB",						

g.V() .hasLabel('presentation') .values('name')

### How many presentations are in my repertoire?

× F	ile Edit :	Selection	View G	o Debug	Terminal	Help	Spea	kingEngagements - V	isual Studio Code		
Q	≣ Speaki	ingEngage	ments $ imes$								
Q	<u> </u>	Sp	beaki	ingEr	ngag	emen	ts / Sp	eakingE	ngager	nents	
ဠ၀	g.V().h	nas('own	erEmailA	ddress', '	chadgree	n@chadgre	en.com').has	Execute			
逯	Graph	JSON									
₿	[ 27 ]										
Δ											
573											
563											

# g.V() .hasLabel('presentation') .count()

### What are the events that I have submitted to?



g.V() .hasLabel('event') .has('year', '2020')

Displaying all 49 vertices and 0 edges

### Where has Graphing Your Way Through the Cosmos been submitted to?



g.V() .hasLabel('presentation') .has('name', 'Graphing Your Way Through the Cosmos') .outE('submittedTo') .inV()

@chadgreen

<u>⊗ 0 ≜ 0 Azure: chadgreen@chadgreen.com</u>

### Where has Graphing Your Way Through the Cosmos been submitted to?



g.V() .hasLabel('presentation') .has('name', 'Graphing Your Way Through the Cosmos') <u>.outE('submittedTo')</u> .inV() .has('year', '2020')

@chadgreen

Displaying all 81 vertices and 0 edges

### Where has Graphing Your Way Through the Cosmos been submitted to?



g.V() .hasLabel('presentation') .has('name', 'Graphing Your Way Through the Cosmos') .outE('submittedTo') .inV()

@chadgreen

<u>⊗ 0 ≜ 0 Azure: chadgreen@chadgreen.com</u>

### Where has Graphing Your Way Through the Cosmos been scheduled?



g.V() .hasLabel('presentation') .has('name', 'Graphing Your Way Through the Cosmos') .outE('submittedTo') .has('status', 'Confirmed') .inV()

⊗ 0 ▲ 0 Azure: chadgreen@chadgreen.com

### View all of the tags



## g.V() .hasLabel('tag')

#### @chadgreen

### Focus on the Graph Data tag



## g.V() .hasLabel('tag') .has('name', 'Graph Data')



### What presentations are tagged with Graph Data?

× 1	File Edit Selection View Go Debug Terminal Help SpeakingEngagements - Visual Studio Code	
¢	≣ SpeakingEngagements ×	
Q	SpeakingEngagements / SpeakingEngagements	
ၟႄၜ	g.V().hasLabel('tag').has('name', 'Graph Data').lnE('tagged') Execute	
逯	Graph JSON	
	<pre>{     "id": "f4d6f410-6bfe-42b1-b3bb-a8035c6c49ec",     "label": "tagged",     "utVIabel": "tag",     "outVIabel": "presentation",     "utVIabel": "presentation",     "utV": "50624330-d2bc-4555-84cf-62185e2263ab",     "outV": "60e69af2-0b2e-43ab-8b79-76482ae10b7b"     ,     "di": "3fe31c1a-f83e-424f-84a5-2a7dc57a2e83",         "label": "tagged",         "ytyPer: "edge",         "utVIabel": "tagged",         "utVIabel": "tagged",         "utVIabel": "tagged",         "utVIabel": "tagged",         "utVIabel": "tagged",         "utVIabel": "tagged",         "utVIabel": "presentation",         "utVIabel": "tagged",         "utVIabel": "presentation",         "utVIabel": "sf624330-d2bc-4555-84cf-62185e2263ab",         "utVIabel": "sf624390-d2bc-4555-84cf-62185e2263ab",         "utVIabel": "tagged",         "utVIabel": "tagged",         "utVIabel": "sf58257-fba1-4fe2-9dfe-5ecb2e7b6414"         /,         "outV": "9a1dbc06-01c9-49f3-87cf-0ad356ed21d5",         "label": "tagged",         "utVIabel": "colls8bad-669e-4075-9046-17fa6fc3840d"         ]         } } </pre>	

g.V() .hasLabel('tag') .has('name', 'Graph Data') .inE('tagged')

### What presentations are tagged with Graph Data?



g.V() .hasLabel('tag') .has('name', 'Graph Data') .inE('tagged') .outV()

### Where have the presentations tagged Graph Data been scheduled?



g.V() .hasLabel('tag') .has('name', 'Graph Data') .inE('tagged') .outV() .outE('submittedTo') .has('status', 'Confirmed') .inV()

Displaying all 10 vertices and 0 edges

### What events have I been scheduled for?



g.E() .hasLabel('submittedTo') .has('status', 'Confirmed') .inV()

@chadgreen

Displaying all 39 vertices and 0 edges

### What events have I been scheduled for?



g.E() .hasLabel('submittedTo') .has('status', 'Confirmed') .inV()

### g.V() .hasLabel('event') .has('presentedAt', 'True')

### What states have I been scheduled to speak in?



⊗ 0 ▲ 0 Azure: chadgreen@chadgreen.com

g.E() .hasLabel('submittedTo') .has('status', 'Confirmed') .inV() .outE('stateLocation') .inV()

### Wrapping Up

- Graphs set of objects in which pairs are in some sense related
- Graph Theory Starts with the 7 bridges of Königsberg
- Graph databases use graph structure to represent and store data
- Azure Cosmos DB globally distributed, multi-model database service
- Graph vs Relational lots of benefits that make graph database worth a look
- Graph Traversal Navigating graph data using patterns

# Thank You