

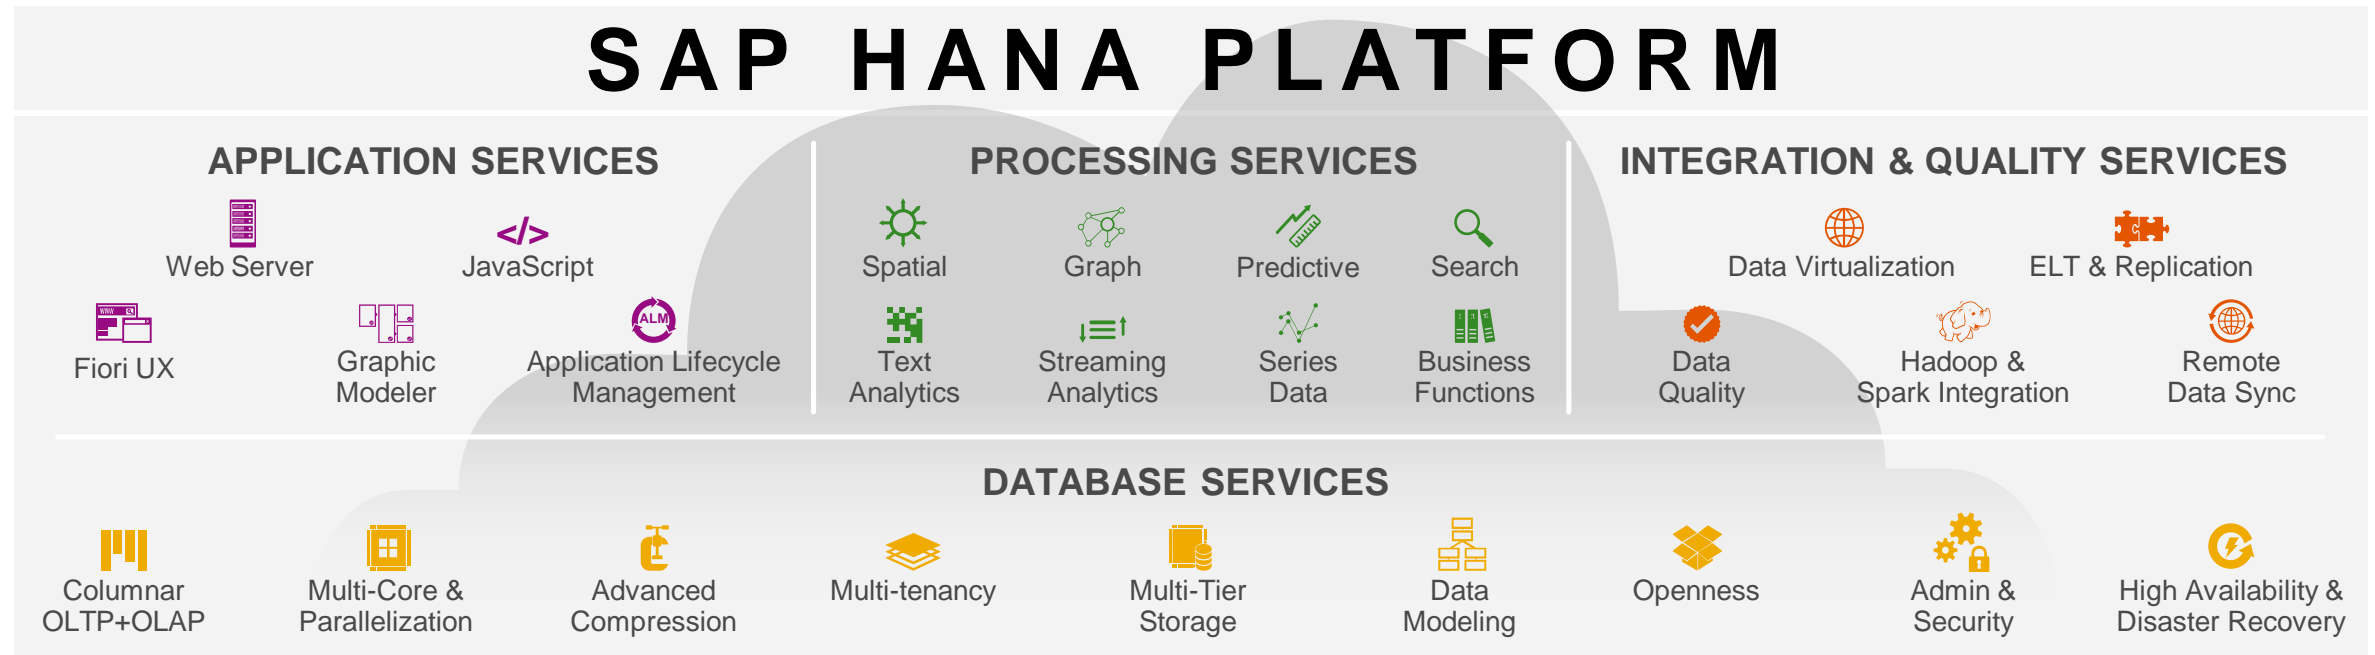


GraphScript: Implementing Complex Graph Algorithms in SAP HANA

Marcus Paradies, SAP SE
September 1, 2017

PUBLIC

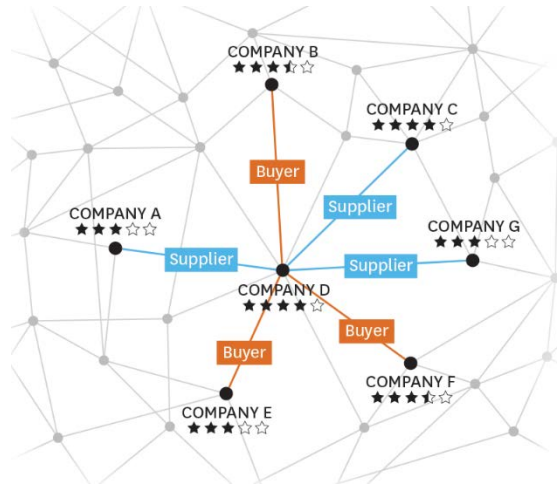
SAP HANA Overview



- Offers advanced analytics features for graph, text, geospatial, and machine learning directly on business data

Graph Querying Paradigms in SAP HANA

Graph Pattern Matching




Example Query

"Retrieve all suppliers of Company D"

openCypher*

Graph Analysis



"Compute all communities in the graph"

GraphScript

Design Principles

Expressiveness & Simplicity

- Easy-to-use for graph algorithm implementers
- Support for a large variety of graph algorithm classes and workflows

Minimality & Orthogonality

- Limited but effective set of types and operations thereon
- Extensibility of built-in graph operators

Native Graph Abstraction

- Native exposure of graph-specific types
- Full exposure of graph data model
- Relational only for returning complex results

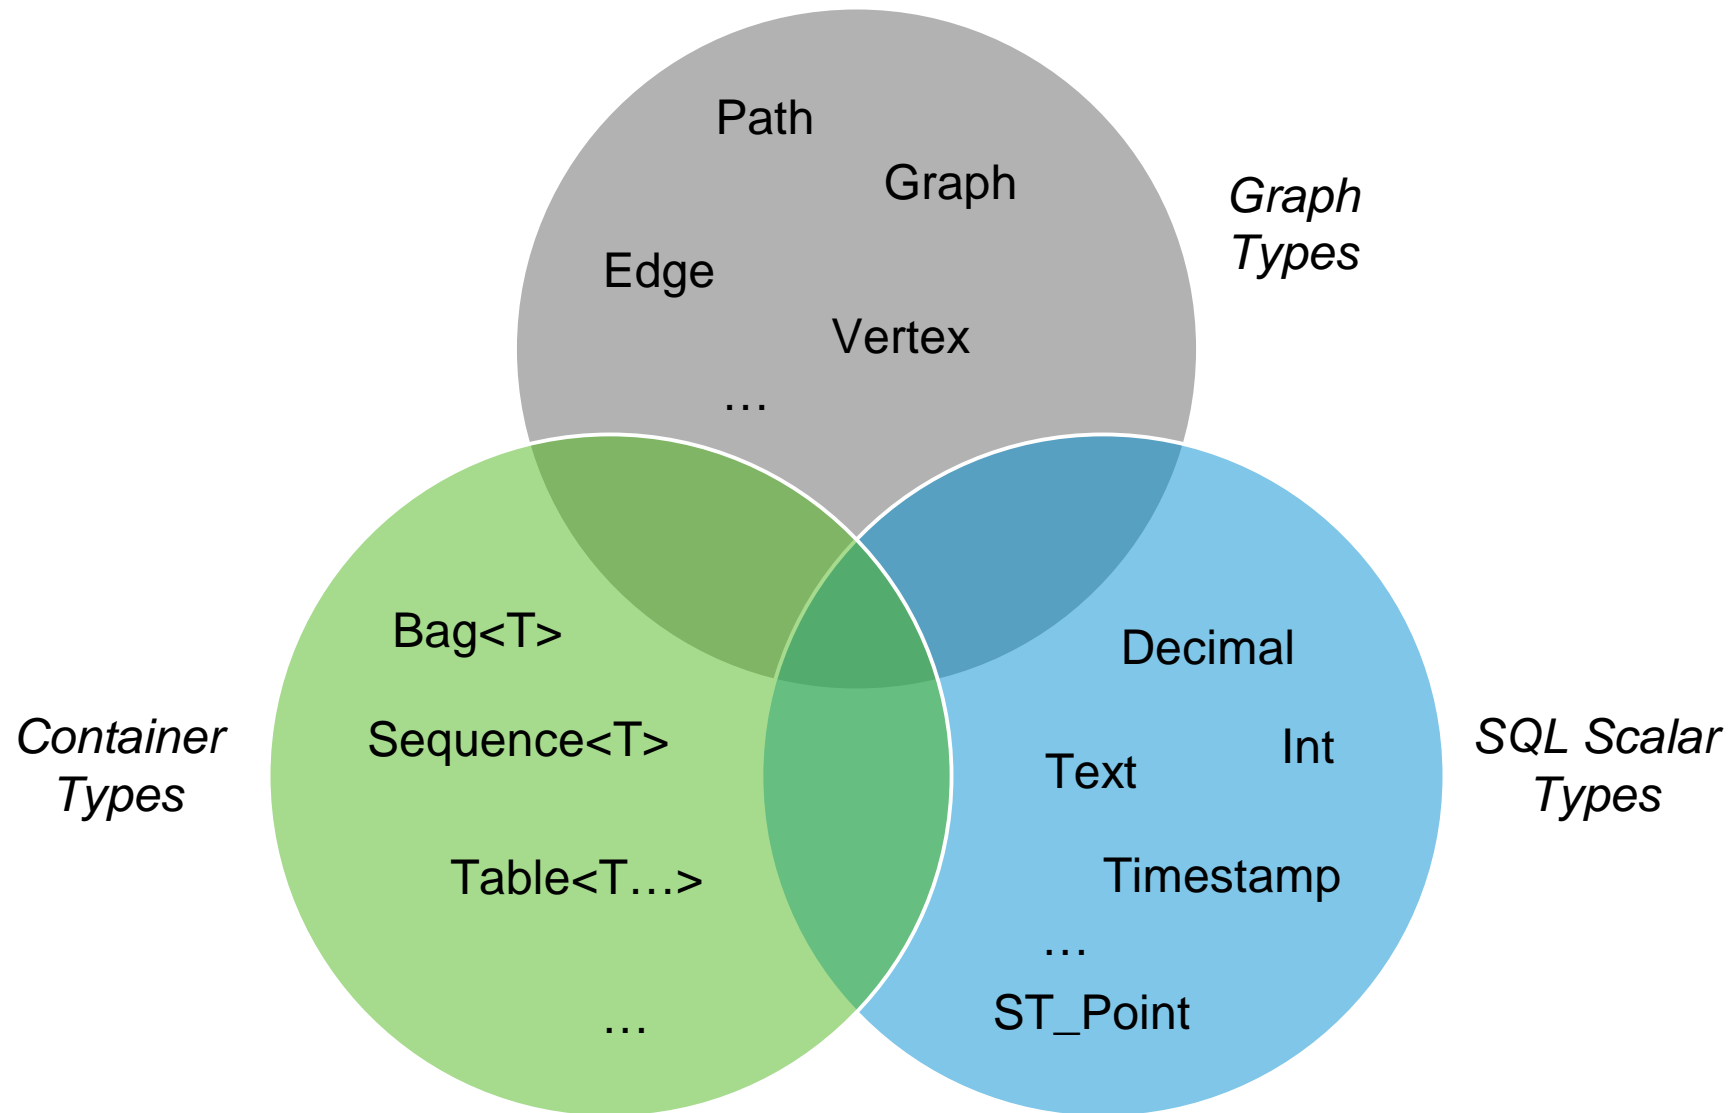
Tight Integration

- Pushdown of operations to relational store
- Reuse of dependency management
- Reuse of resource management

High Performance

- Desired performance close to hand-written code
- Explicit parallelization
- Effective Program Rewritings

GraphScript Type System



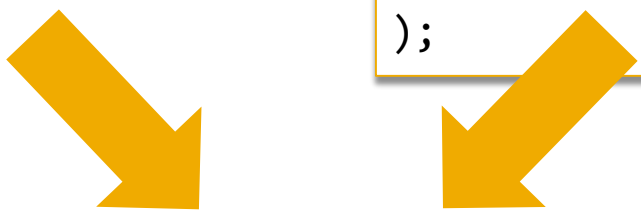
Graph Data Exposure in GraphScript

Vertex Table

```
CREATE COLUMN TABLE MYSCHEMA.VERTICES (  
  ID VARCHAR(100) PRIMARY KEY,  
  TYPE VARCHAR(100),  
  NAME VARCHAR(100),  
  TITLE VARCHAR(100)  
);
```

Edge Table

```
CREATE COLUMN TABLE MYSCHEMA.EDGES (  
  ID INTEGER PRIMARY KEY,  
  SRC VARCHAR(100) NOT NULL  
    REFERENCES MYSCHEMA.VERTICES (ID)  
  TRGT VARCHAR(100) NOT NULL  
    REFERENCES MYSCHEMA.VERTICES (ID)  
  TYPE VARCHAR(50)  
);
```



```
CREATE GRAPH WORKSPACE MYSCHEMA.MY_GRAPH  
  EDGE TABLE MYSCHEMA.EDGES  
    SOURCE COLUMN SRC  
    TARGET COLUMN TRGT  
    KEY COLUMN ID  
  VERTEX TABLE MYSCHEMA.VERTICES  
    KEY COLUMN ID;
```

**Graph Workspace
Metadata Object**

Graph Data Exposure in GraphScript /2

Vertex Table View

```
CREATE VIEW MYSCHEMA.VERTEX_VIEW AS
  SELECT * FROM MYSCHEMA.VERTICES
  WHERE TYPE = 'Person';
```

Edge Table View

```
CREATE VIEW MYSCHEMA.EDGE_VIEW AS
  SELECT * FROM MYSCHEMA.EDGES
  WHERE TYPE = 'knows';
```



```
CREATE GRAPH WORKSPACE MYSCHEMA.MY_SUBGRAPH
  EDGE TABLE MYSCHEMA.EDGE_VIEW
  SOURCE COLUMN SRC
  TARGET COLUMN TRGT
  KEY COLUMN ID
  VERTEX TABLE MYSCHEMA.VERTEX_VIEW
  KEY COLUMN ID;
```

**Graph Workspace
Metadata Object**

A Simple GraphScript Example

```
CREATE PROCEDURE "myGraphProc"(OUT numNeighbors BIGINT)
LANGUAGE GRAPH READS SQL DATA AS
BEGIN
    Graph g = Graph("myGraph");
    ALTER g ADD TEMPORARY VERTEX ATTRIBUTE(BIGINT cnt = 0);
    FOREACH v IN Vertices(:g) {
        v.cnt = Count(Neighbors(:g, :v, 1, 3));
    }
    FOREACH v IN Vertices(:g) {
        numNeighbors += :v.cnt;
    }
END
```


Adjacency List Construction

0	1	2
1	2	5
2		
3	4	2
4	0	2
5	6	
6	2	


Adjacency List

0	C
1	D
2	E
3	A
4	B
5	F
6	G

Vertex Key Dictionary

Key	Color
C	green
D	green
E	yellow
A	green
B	green
F	yellow
G	green

Vertex Attributes




Variants:

- Omit dictionary encoding for dense key domains
- Static/Dynamic/Compressed adjacency list
- Vertex/edge adjacency

Parallel Index Construction with up to 65 Mio. edge insertions/sec

Inducing Subgraphs

Data graph:



<p><i>"Induce a graph over all blue edges"</i></p>	<pre> graph TD 1((1)) -- blue --> 2((2)) 2((2)) -- blue --> 3((3)) 2((2)) -- blue --> 5((5)) 3((3)) -- blue --> 4((4)) 4((4)) -- blue --> 5((5)) </pre>	<pre> Graph g = Subgraph(:g, e IN Edges(:g) WHERE :e.color == 'blue'); </pre>
<p><i>"Induce a graph over all red edges that connect a green and a yellow vertex"</i></p>	<pre> graph TD 1((1)) -- red --> 5((5)) 3((3)) -- red --> 5((5)) 4((4)) -- red --> 6((6)) 5((5)) -- red --> 7((7)) 6((6)) -- red --> 7((7)) 7((7)) -- red --> 5((5)) </pre>	<pre> Graph g = Subgraph(:g, e IN Edges(:g) WHERE Source(:e).color == 'green' AND Target(:e).color == 'yellow' AND :e.color == 'red'); </pre>
<p><i>"Induce a graph over all all vertices that are reachable from vertex 4"</i></p>	<pre> graph TD 4((4)) -- blue --> 5((5)) 4((4)) -- red --> 6((6)) 5((5)) -- red --> 7((7)) 6((6)) -- red --> 7((7)) 7((7)) -- red --> 5((5)) </pre>	<pre> Vertex v1 = Vertex(:g, 4); Graph g = Subgraph(:g, v IN Vertices(:g) WHERE IS_REACHABLE(:g, :v1, :v); </pre>

Integration with other Data Models/Scalar Types

Creation of Relational Output from GraphScript

```
Graph g = Graph("myWorkspace");
ALTER g ADD TEMPORARY VERTEX ATTRIBUTE(DOUBLE length = 0);
FOREACH v IN Vertices(:g) {
    Path p = Shortest_Path(:g, :v, Vertex(:g, 1));
    v.length = Length(:p);
}
outTab = SELECT :v.id, :v.length FOREACH v IN Vertices(:g);
```

Integration with Geospatial Processing

```
Graph g = Graph("myWorkspace");
ST_Geometry area = Vertex(:g, 'Munich').area;
Graph g1 = Subgraph(:g, v IN Vertices(:g) WHERE :v.type == 'Person'
                    AND ST_Within(:v.location , :area));
```

Conclusion

Language Constructs

- Rich type system with native graph types
- Powerful imperative constructs

Code Generation

- Generation of low-level code against internal Graph Storage interface
- Elimination of query processing on external keys
- Pushdown of filter conditions to relational engine

Future Work

- More language extensions towards fast traversals and user-defined function invocations
- More advanced GraphScript program rewritings and optimizations

Thank you.

Contact information:

Marcus Paradies

marcus.paradies@sap.com