

Trinity Graph Engine

Bin Shao, Yatao Li, Wei-Ying Ma

Microsoft Research Asia

Why do we need a graph system?

Existing Systems

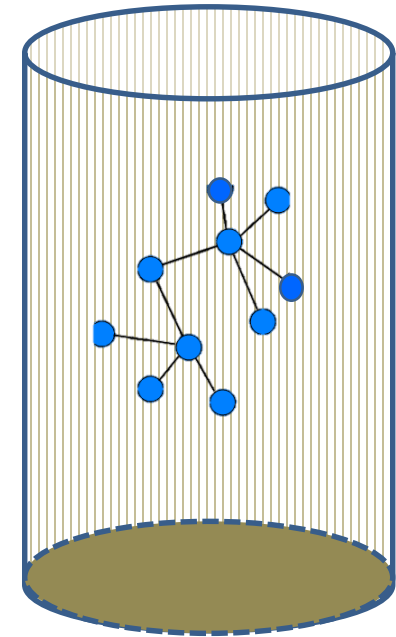
- Mature data processing systems
 - RDBMS
 - Map Reduce Systems, e.g. cosmos
- Systems specialized for certain graph operations:
 - PageRank

Graph Data is “Special” ...

- Random access (Poor Locality)
 - For a node, its adjacent nodes’ content cannot be accessed without “jumping” no matter how you represent a graph
 - Not cache-friendly, data reuse is hard
- Unstructured nature of graph
 - Difficult to extract parallelism by partitioning data
 - Hard to get an efficient “Divide and Conquer” solution


Graph in the Jail of Storage

- RDBMS/cosmos, mature but not for graphs
- The commonest graph operation “traversal” incurs excessive amount of table joins

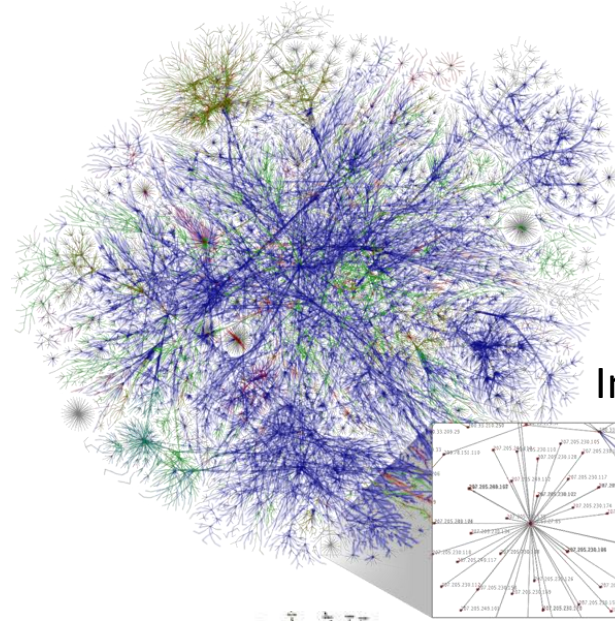


**Graph in the
Jail of the storage**

Challenge I: Diversity of Graphs




Satori Schema Graph



Internet Web Graph



Social Network



Protein Interaction Network

Do we need to design algorithms for each type of graphs?

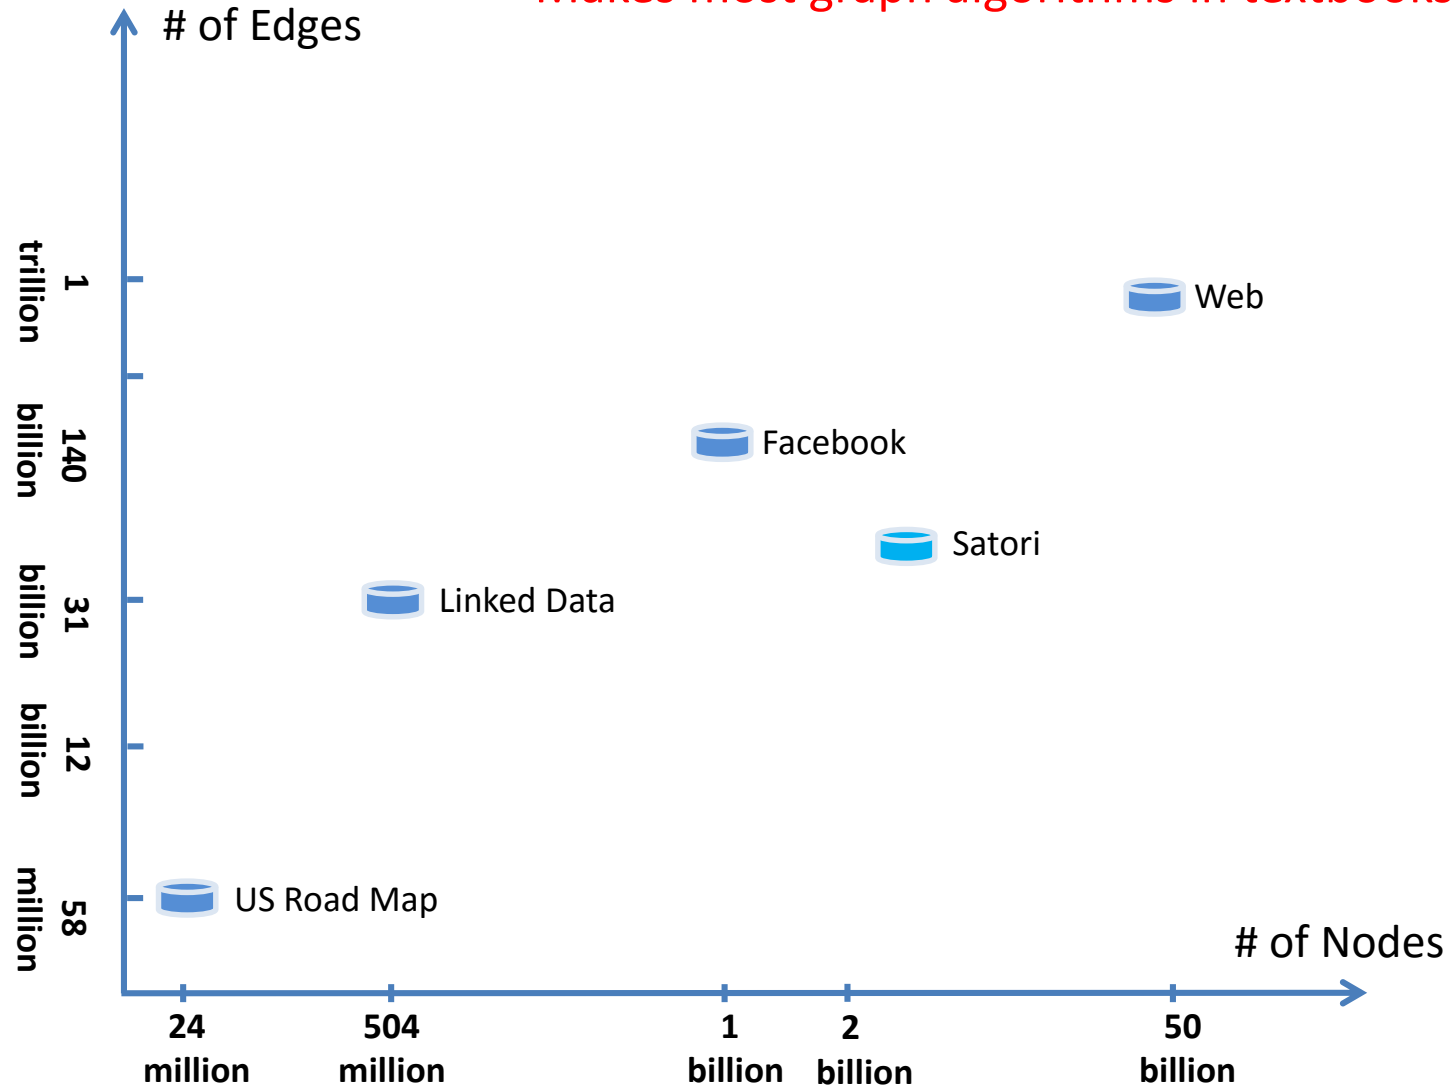
Challenge II: Diversity of Computations

- Online query processing
 - Shortest path query
 - Subgraph matching query
 - SPARQL query
 - ...
- Offline graph analytics
 - PageRank
 - Community detection
 - ...
- Other graph operations
 - Graph generation, visualization, interactive exploration, etc.

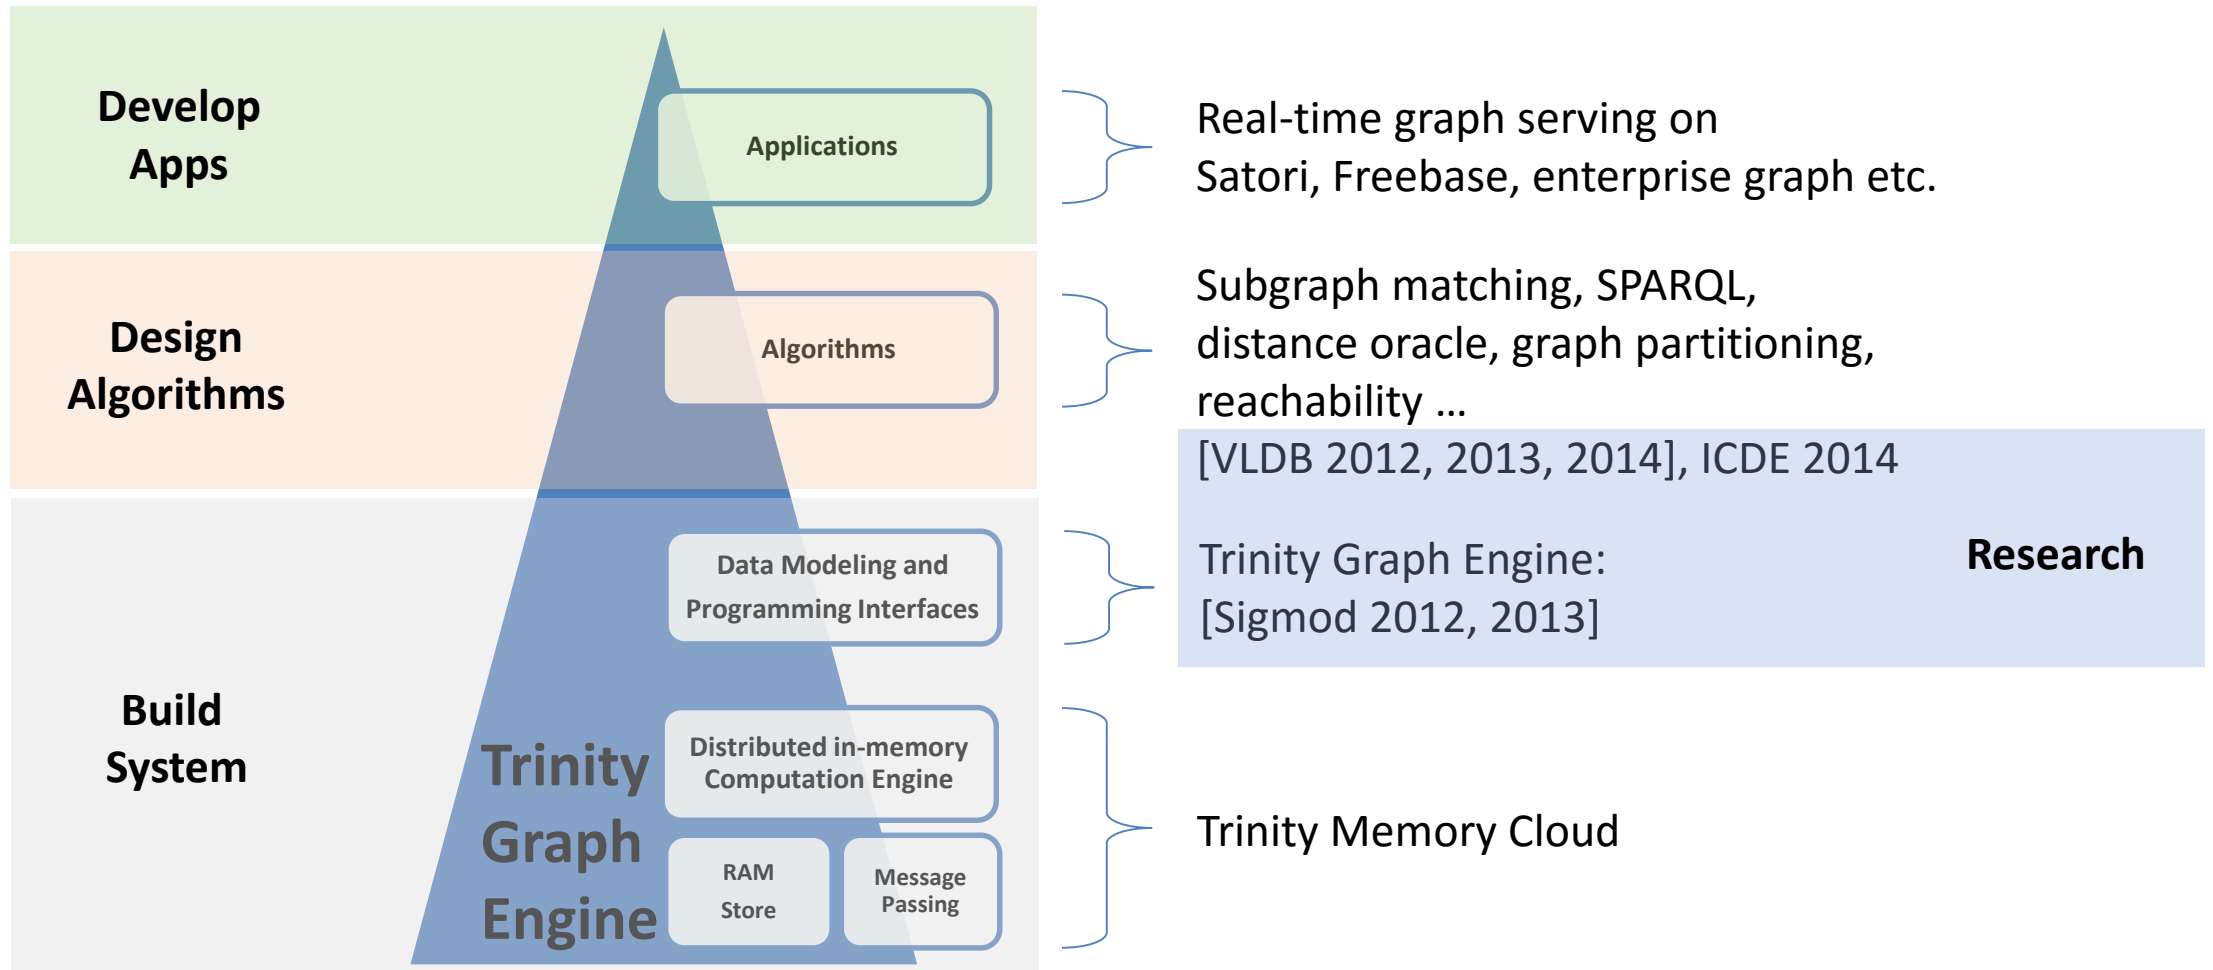
Do we need to implement systems for each graph operation?

Challenge III: The Scale of Graphs

Makes most graph algorithms in textbooks ineffective!



Roadmap of the Graph Engine project



Design Philosophy

Not a one-size-fits-all graph system, but a graph engine

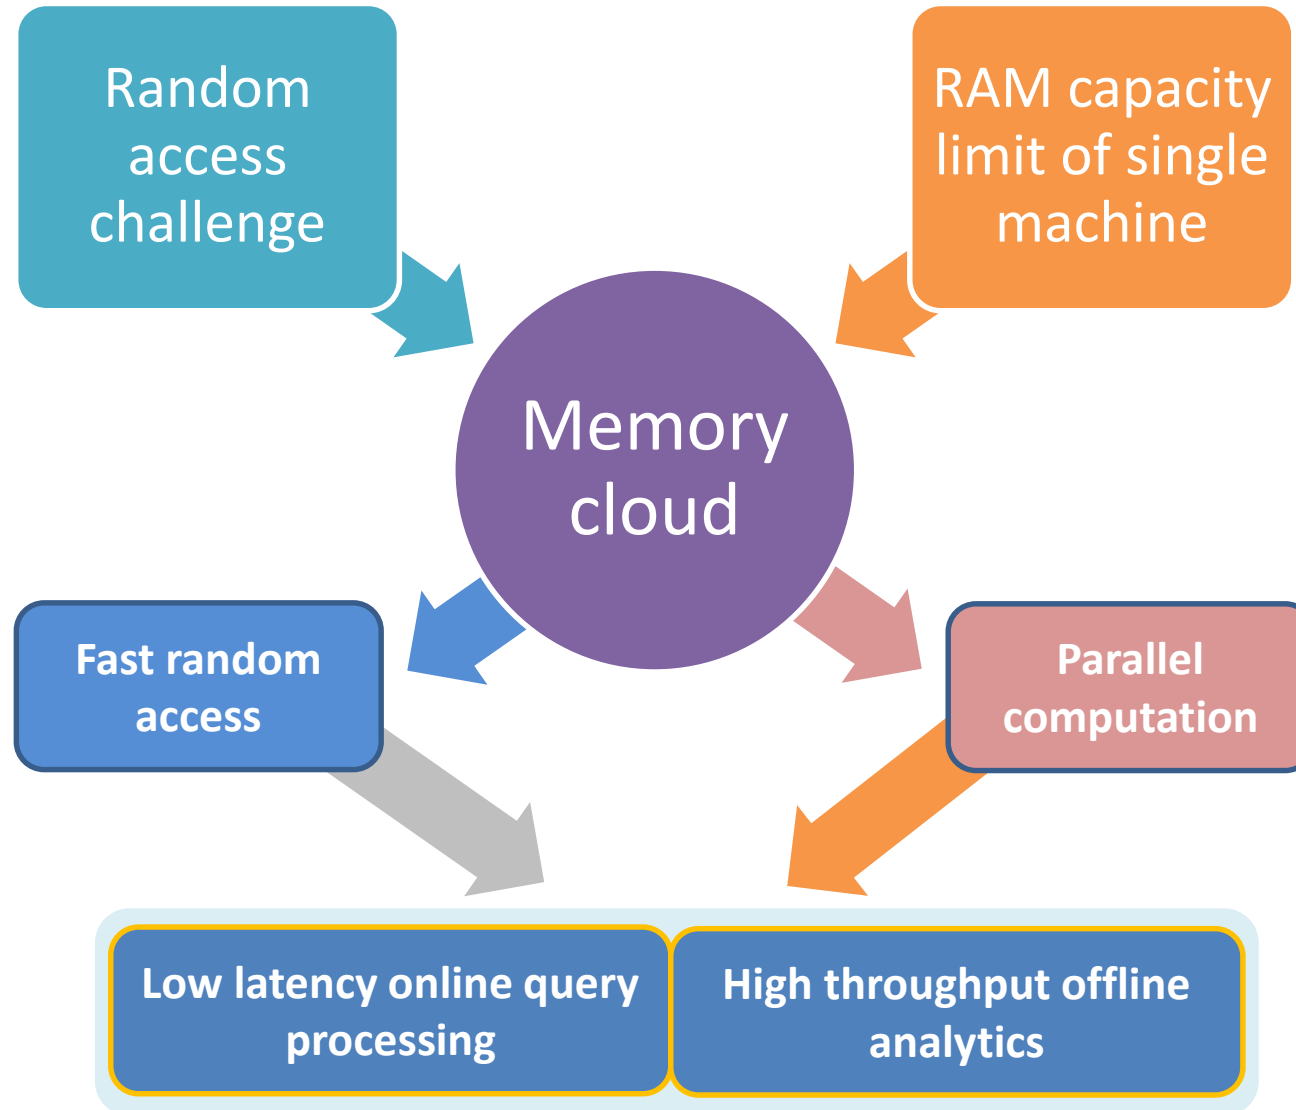
Flexible data and computation modeling capability



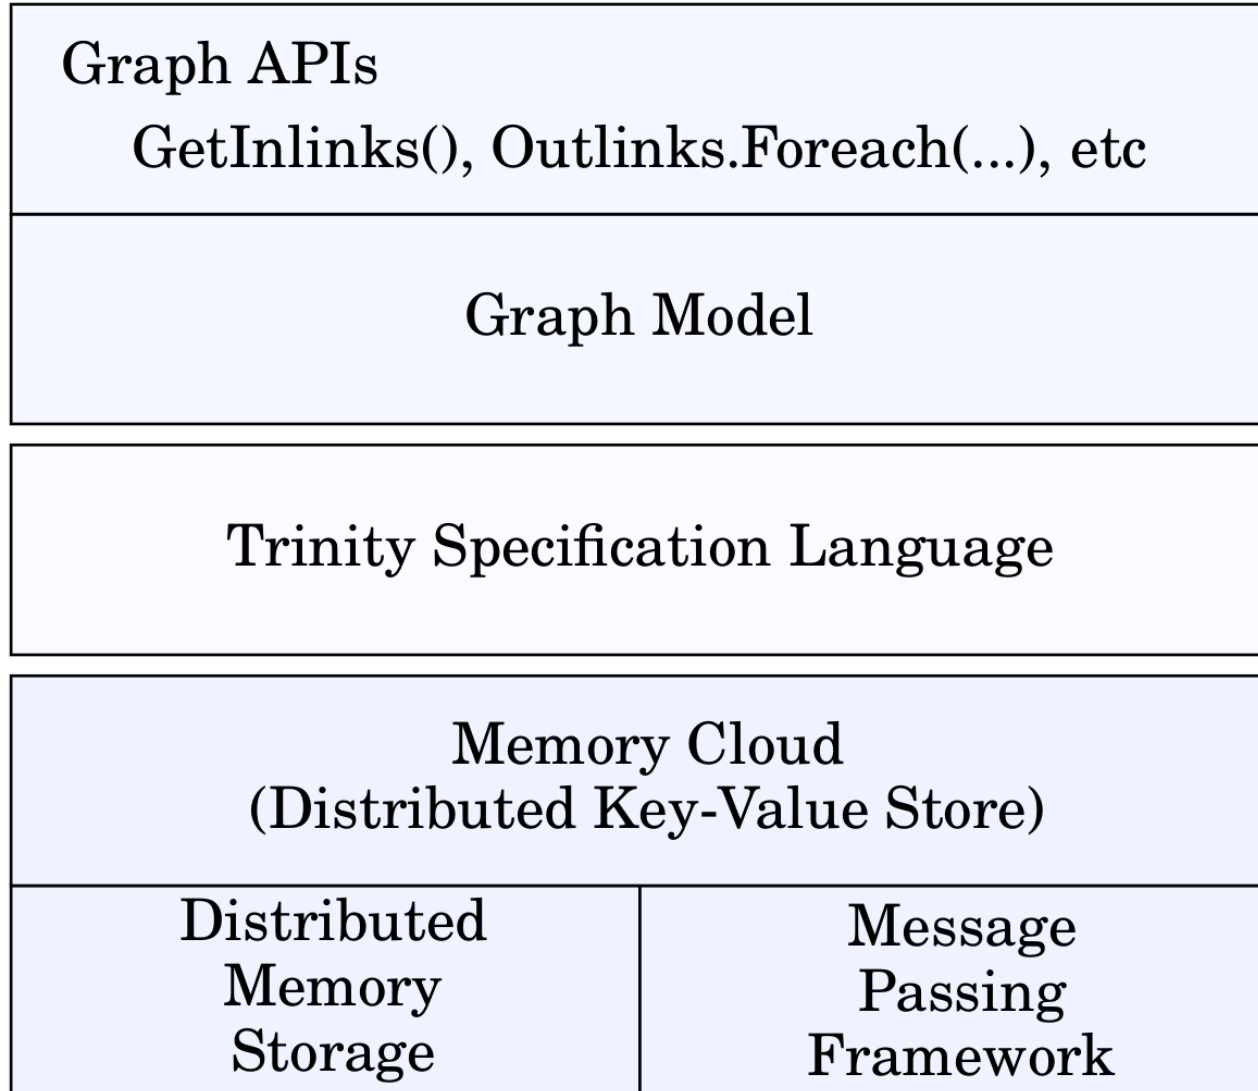
Trinity can morph into
a large variety of graph processing systems

***Trinity* = Graph Modeling Tools +
Distributed In-memory Data Store +
Declarative Programming Model**



Design Rationale of Memory Cloud



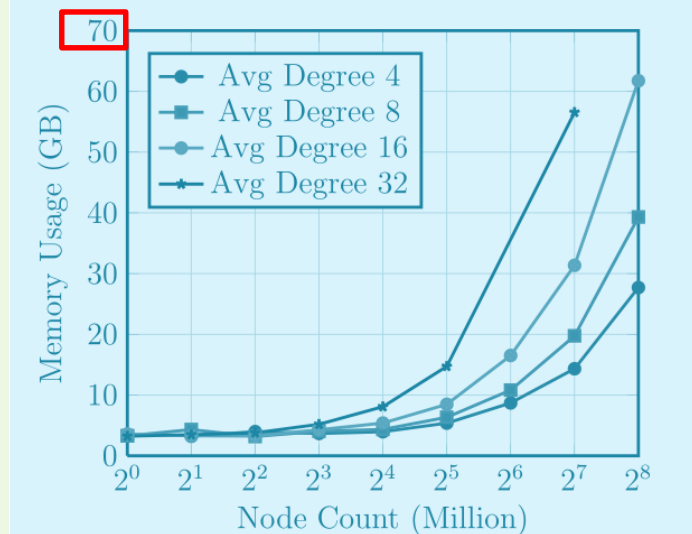

System Stack



One Byte Counts (Trinity vs. PBGL)



Execution Time

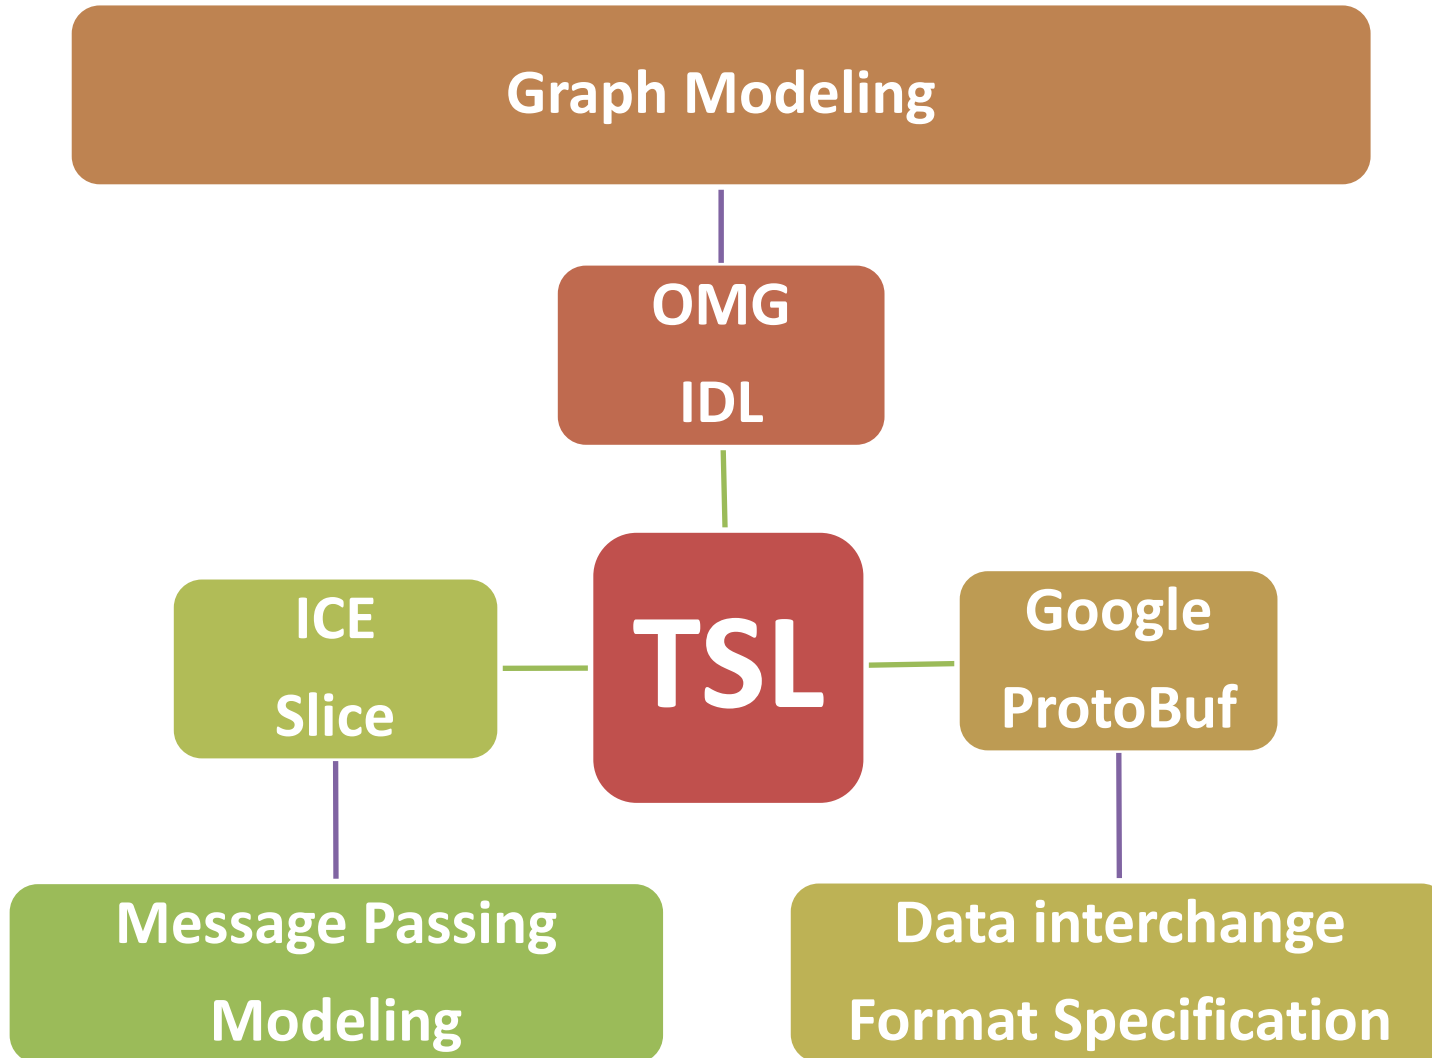


Memory Usage

PBGL

Trinity

Trinity Specification Language



Why TSL?

- TSL allows users to define graph schemata, and communication protocols through declarative interfaces.
- TSL makes Trinity memory cloud beyond a key-value store
 - Users are allowed to freely define the data schema
 - TSL makes message passing programming ever so easy

Modeling a Movie and Actor Graph

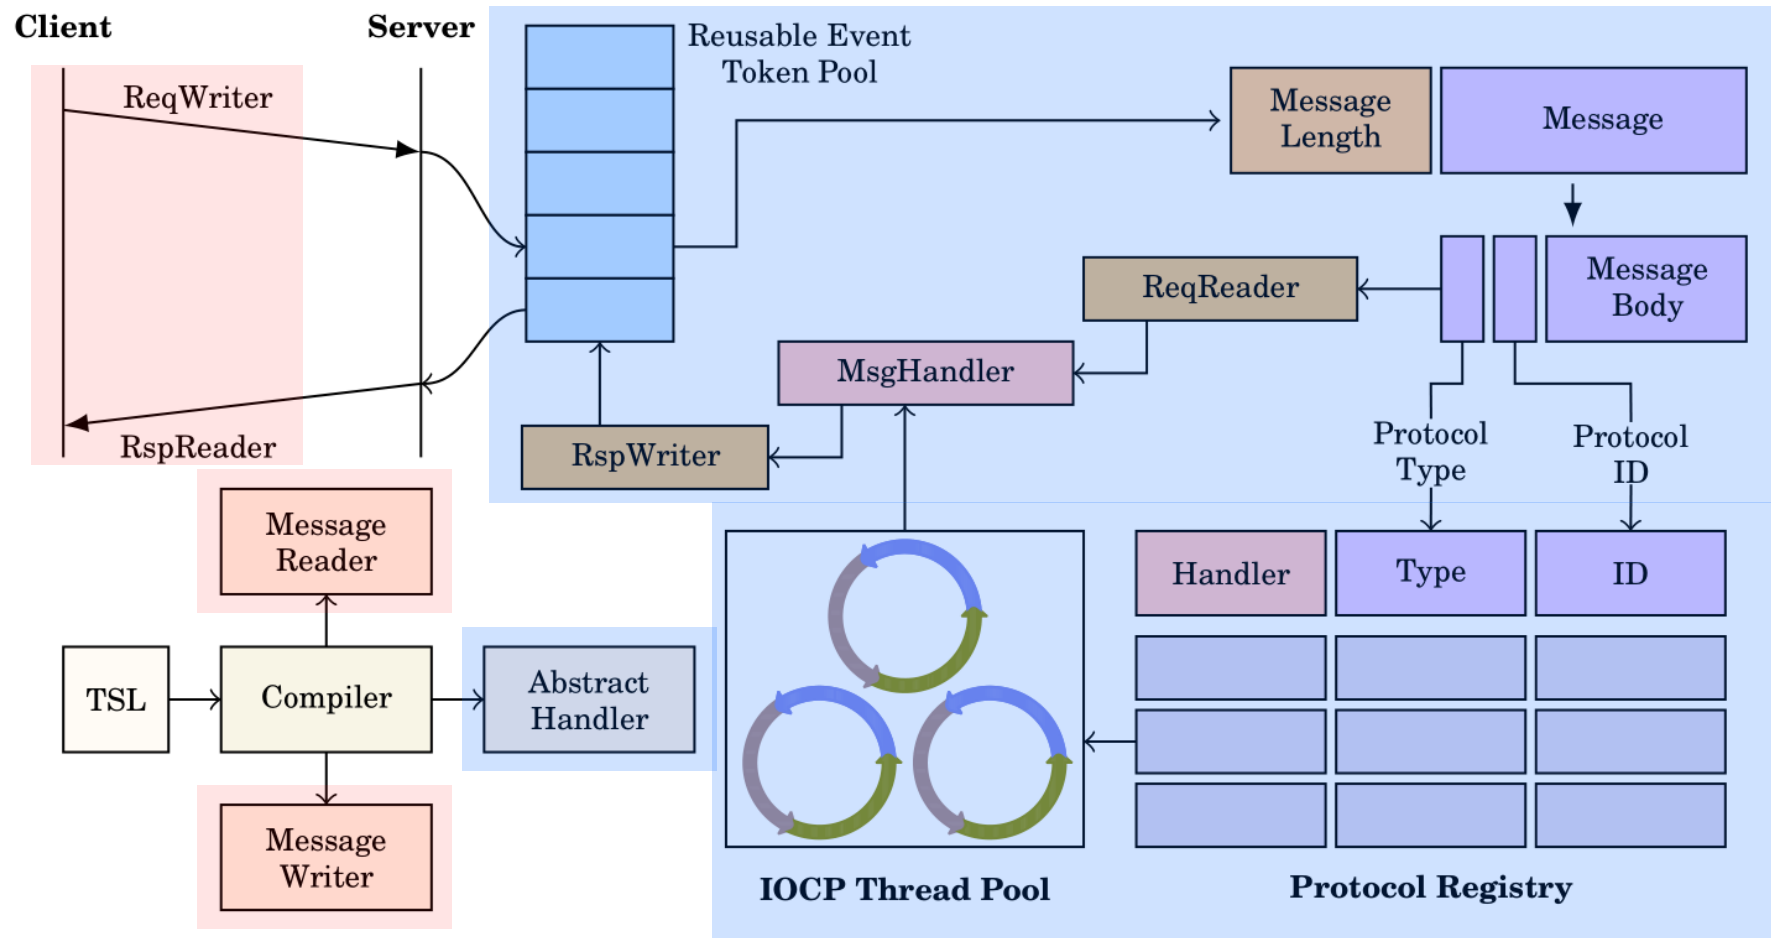
```
cell struct Movie
{
    string Name;
    [GraphEdge]
    List<CellId> Actors;
}
```

```
cell struct Actor
{
    string Name;
    [GraphEdge]
    List<CellId> Movies;
}
```


Modeling Message Passing

```
struct MyMessage
{
    string Text;
}
protocol Echo
{
    Type: Syn;
    Request: MyMessage;
    Response: MyMessage;
}
```

TSL-Powered Message Passing



Trinity-enabled Graph Computation Paradigms

- **Vertex-centric graph analytics**
 - Prosperous since Pregel, e.g. Giraph, GraphChi
- **Approximate graph computation based on local sampling**
 - Enabled by randomly partitioned in-memory graph
 - Fast approximate computation with minimum communication costs
 - Application: distance oracle [VLDB 2014]
- **Index-free real-time online query processing**
 - Enabled by fast in-memory distributed graph exploration
 - Examples, subgraph match (vldb 2012) and Trinity.RDF (vldb 2013)

Query Index Examples

Algorithms	Index Size	Index Time	Update Cost
Ullmann [Ullmann76], VF2 [CordellaFSV04]	-	-	-
RDF-3X [NeumannW10]	$O(m)$	$O(m)$	$O(d)$
BitMat [AtreCZH10]	$O(m)$	$O(m)$	$O(m)$
Subdue [HolderCD94]	-	Exponential	$O(m)$
SpiderMine [ZhuQLYHY11]	-	Exponential	$O(m)$
R-Join [ChengYDYW08]	$O(nm^{1/2})$	$O(n^4)$	$O(n)$
Distance-Join [ZouCO09]	$O(nm^{1/2})$	$O(n^4)$	$O(n)$
GraphQL [HeS08]	$O(m + nd^r)$	$O(m + nd^r)$	$O(d^r)$
Zhao [ZhaoH10]	$O(nd^r)$	$O(nd^r)$	$O(d^L)$
GADDI [ZhangLY09]	$O(nd^L)$	$O(nd^L)$	$O(d^L)$

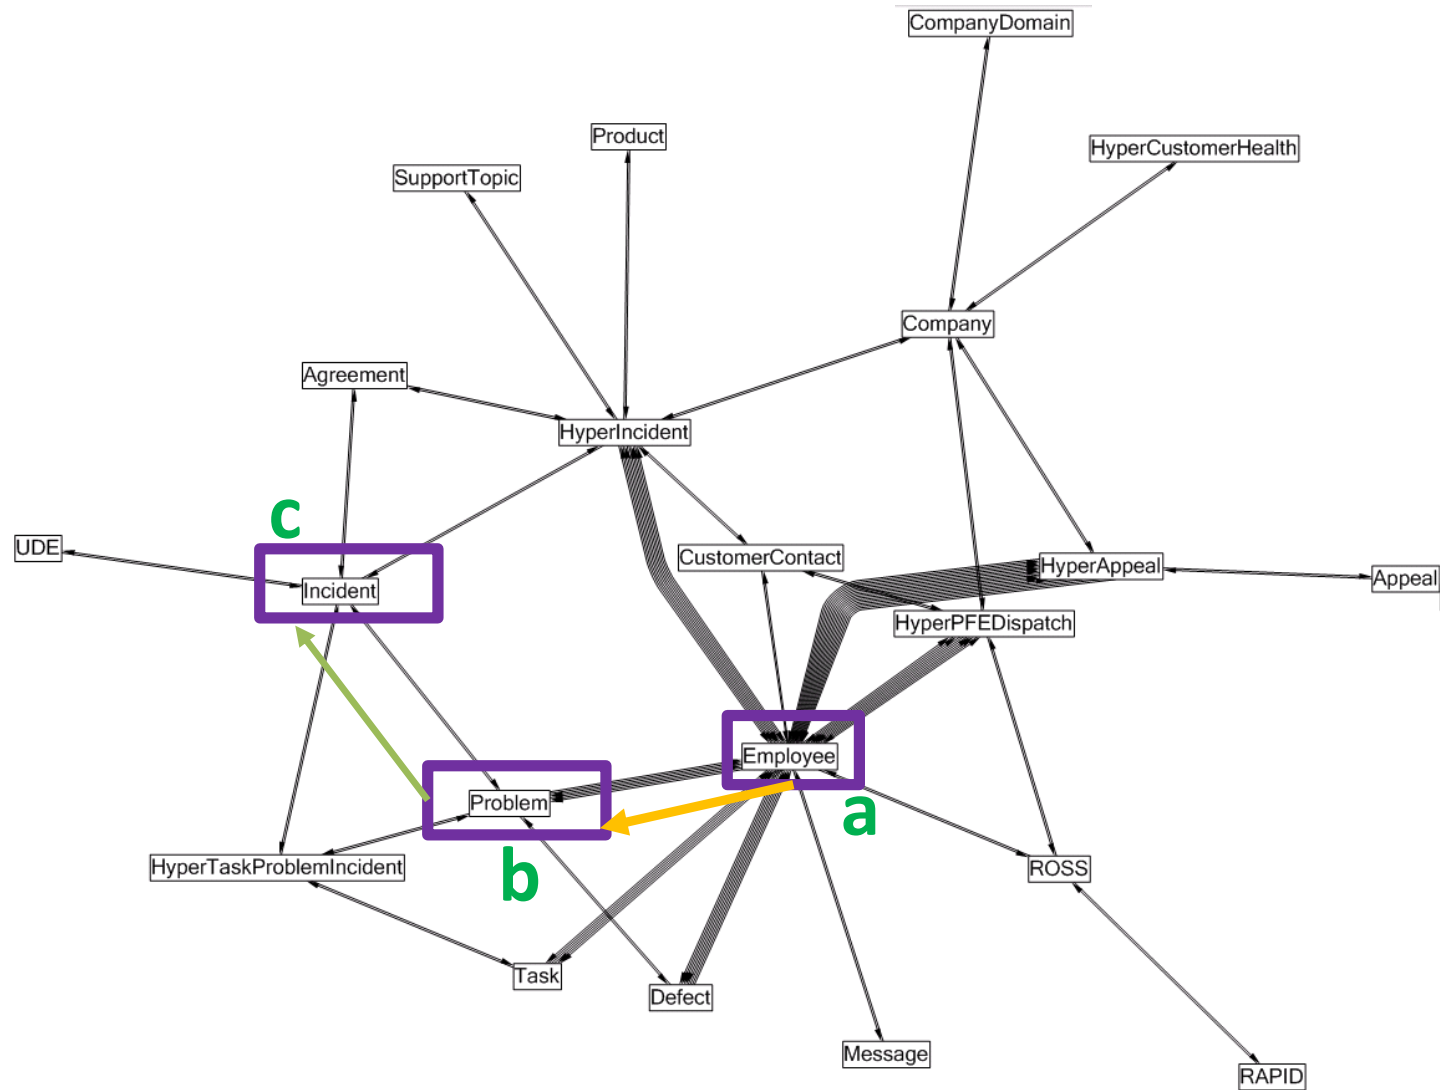
Index-based subgraph matching [Sun VLDB 2012]

Query Index Examples

Algorithms	Index Size for Facebook	Index Time for Facebook	Query Time on Facebook (s)
Ullmann [Ullmann76], VF2 [CordellaFSV04]	-	-	>1000
RDF-3X [NeumannW10]	1T	>20 days	>48
BitMat [AtreCZH10]	2.4T	>20 days	>269
Subdue [HolderCD94]	-	> 67 years	-
SpiderMine [ZhuQLYHY11]	-	> 3 years	-
R-Join [ChengYDYW08]	>175T	> 10^{15} years	>200
Distance-Join [ZouCO09]	>175T	> 10^{15} years	>4000
GraphQL [HeS08]	>13T($r=2$)	> 600 years	>2000
Zhao [ZhaoH10]	>12T($r=2$)	> 600 years	>600
GADDI [ZhangLY09]	> 2×10^5 T ($L=4$)	> 4×10^5 years	>400

Index-based subgraph matching [Sun VLDB 2012]

Index-free Query Processing



Trinity Applications

Source Code Graph (Visual Studio)

Welcome to Progression Web

[About Progression Web](#)

Find node by (partial) name:
DrawPreview
 Partial match
In container:
GATetris
Find node(s)
Reset zoom
Remove nodes w/o links
Clear graph
Run force layout

Select an exploration design:
Design 7
Select a subgraph design:
Design 7

Quick user guide
[Expand](#)

The graph shows a central node **GATetris** (UserControl) with 37 links. It is contained by **container**, **gaTetris**, **Options**, and **tetris**. It contains **DrawPreview**, **controlPanel_Paint**, and **controlPanel_Resize**. It also has a link to **GATetris** (parent class). The **UserControl** class is highlighted in blue, and the **GATetris** class is highlighted in yellow.

Inc	Out	Color	Link category
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Green	Attribute use
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Dark Green	Method calls
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Light Purple	Element type
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Light Green	Function pointer
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Red	Generic member use
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Light Orange	Generic parameter use
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Orange	Generic type use
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Dark Purple	Return type
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Light Blue	Implements
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Blue	Inherits from
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Yellow	References
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Gold	Contains
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Yellow	Delayed Cross Group Link
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Grey	Inactive links
<input type="checkbox"/>	<input type="checkbox"/>		Select all
<input type="checkbox"/>	<input type="checkbox"/>		Deselect all

```
12 [DesignTypeof(TetrisGridContainerDesigner)]
13 public sealed class GATetris : UserControl
14 {
15
16     public GATetris()
17     {
18         InitializeComponent();
19
20         SetStyle(ControlStyles.AllPaintingInWmPaint, true);
21         SetStyle(ControlStyles.DoubleBuffer, true);
22         SetStyle(ControlStyles.UserPaint, true);
23     }
24 }
```

Show declarations only
GATetris
NB: Changing the assembly clears the current graph!
About.cs
BestPlayers.cs
BestPlayersCollection.cs
EnterName.cs

Source Code Graph (Visual Studio)

The screenshot displays the Source Code Graph tool in Visual Studio. The main window shows a search for the method 'DrawPreview' in the 'GATetris' assembly. A single node labeled 'DrawPreview' is visible in the center of the graph area. The left sidebar contains search filters and navigation options. The right sidebar shows a list of link categories with checkboxes for 'Inc' and 'Out' colors. At the bottom right, there is a dropdown menu for the assembly and a list of files.

Welcome to Progression Web

[About Progression Web](#)

Find node by (partial) name:
DrawPreview
 Partial match
In container:
GATetris
Find node(s)
Reset zoom
Remove nodes w/o links
Clear graph
Run force layout

Select an exploration design:
Design 7
Select a subgraph design:
Design 7

Quick user guide
[Expand](#)

DrawPreview (Method)

DrawPreview

Inc	Out	Color	Link category
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Green	Attribute use
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Dark Green	Method calls
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Light Blue	Element type
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Light Green	Function pointer
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Red	Generic member use
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Orange	Generic parameter use
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Dark Orange	Generic type use
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Blue	Return type
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Light Blue	Implements
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Dark Blue	Inherits from
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Yellow	References
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Gold	Contains
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Light Yellow	Delayed Cross
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Yellow	Group Link
		Grey	Inactive links
<input type="checkbox"/>	<input type="checkbox"/>		Select all
<input type="checkbox"/>	<input type="checkbox"/>		Deselect all

Show declarations only
GATetris
NB: Changing the assembly clears the current graph!
About.cs
BestPlayers.cs
BestPlayersCollection.cs

ACADEMIC SEARCH

Demo

Example

```
FROM a in {"Author.FullName='Leslie Lamport'"} MATCH a-->b(PaperAuthorOrganization)-->c(Paper) SELECT a.FullName,c.Title
```

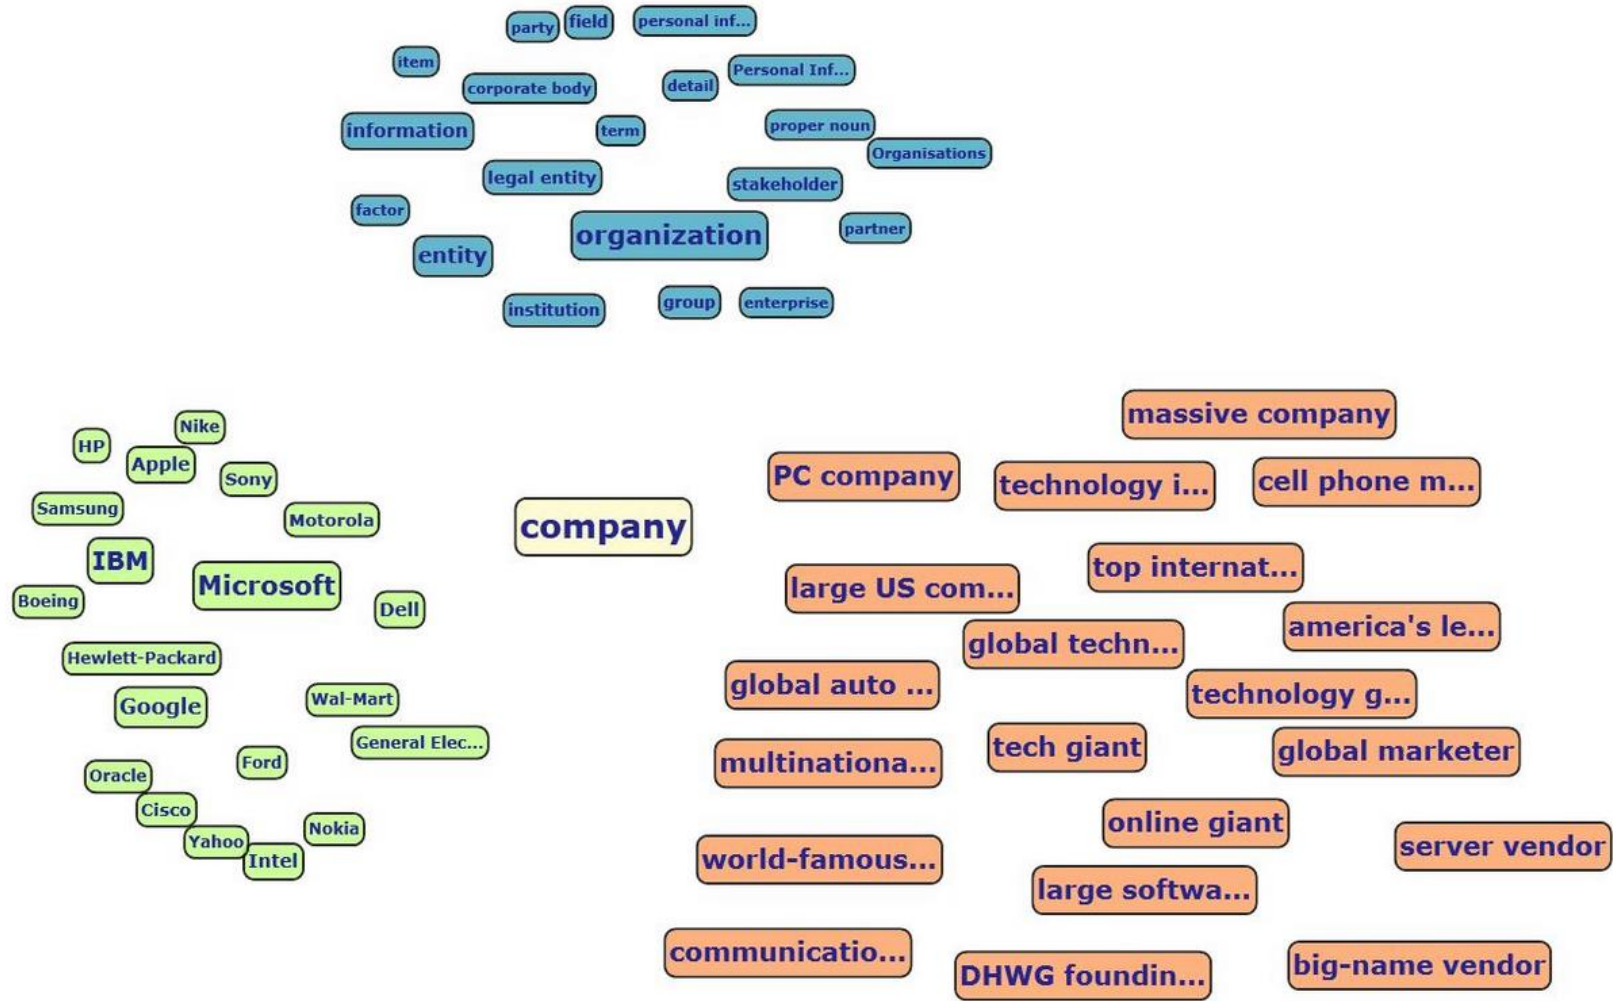
Search

STOP

Query Result

a.FullName	c.Title
Leslie Lamport	Composition: A Way to Make Proofs Harder
Leslie Lamport	A Formal Basis for the Specification of Concurrent Systems
Leslie Lamport	The Operators of TLAC
Leslie Lamport	The Synchronization of Independent Processes
Leslie Lamport	Corrigendum: "A New Approach to Proving the Correctness of Multiprocess Programs"
Leslie Lamport	Comment on Bell's quadratic quotient method for hash coded searching
Leslie Lamport	SIFT: Design and analysis of a fault-tolerant computer for aircraft control
Leslie Lamport	Latex: a document preparation system
Leslie Lamport	Constructing digital signatures from a one-v~ray function
Leslie Lamport	Specifying
1 2 3 4 5 6 7 8 9 10 ... >>	

Knowledge Graph



**Satori Knowledge Graph
Powered by Trinity**

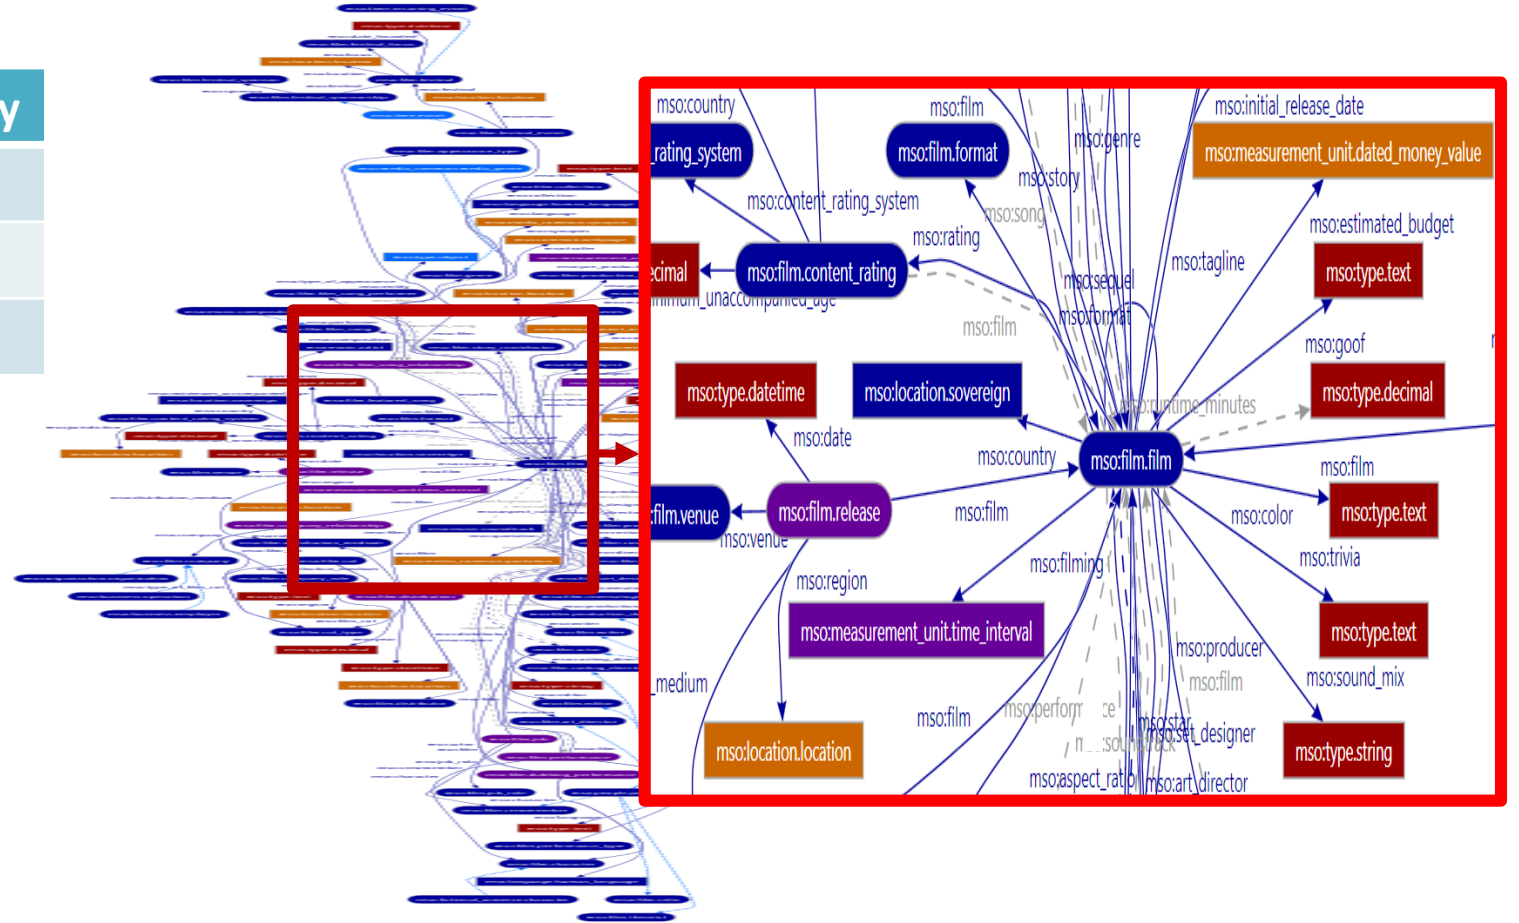
Challenges of Serving Satori

悟 SATORI

Satori: An ever-growing knowledge repository

Raw RDF data	4T+
Entities	2.4B+
Triple Facts	20B+

- Complex data schema
 - Rich relations



Challenges of Serving Satori

悟 SATORI

Satori: An ever-growing knowledge repository

Raw RDF data	4T+
Entities	2.4B+
Triple Facts	20B+

- Complex data schema
 - Rich relations
 - Multi-typed entities

123 mso/type.object.name "Pal"

*123 mso/type.object.type
mso/organism.dog*

123 mso/organism.dog.breeds "Collie Rough"

123 mso/type.object.type mso/film.actor

123 mso/film.actor.film 789

789 mso/type.object.type mso/film.film

789 mso/type.object.name "Lassie Come Home"

"Pal"



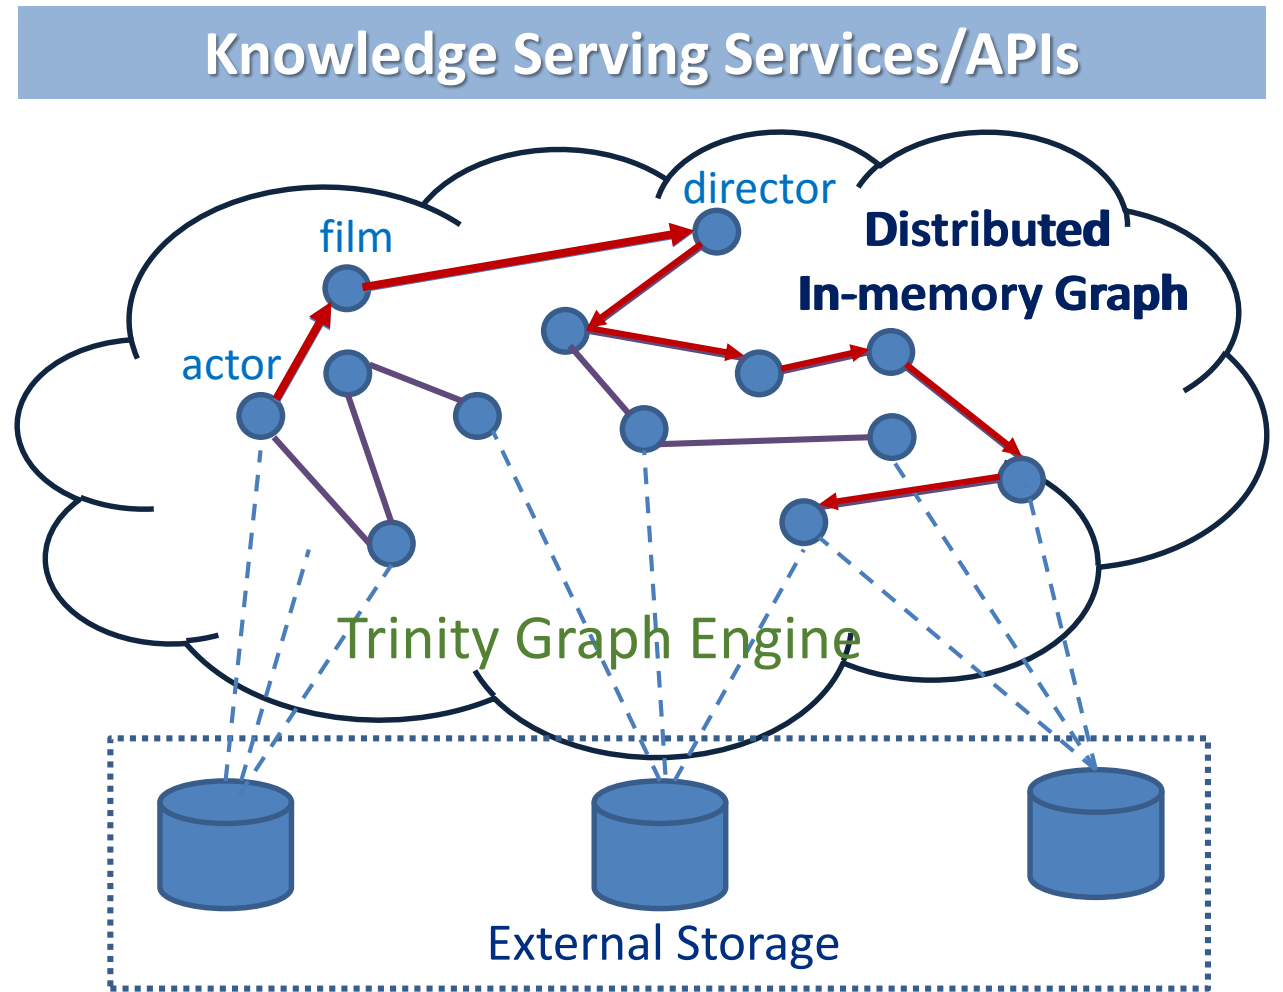
Challenges of Serving Satori

悟 SATORI

Satori: An ever-growing knowledge repository

Raw RDF data	4T+
Entities	2.4B+
Triple Facts	20B+

- Complex data schema
 - Rich relations
 - Multi-typed entities
- Distributed in-memory knowledge graph



Satori in Native Graph Database for Real-time Knowledge Serving

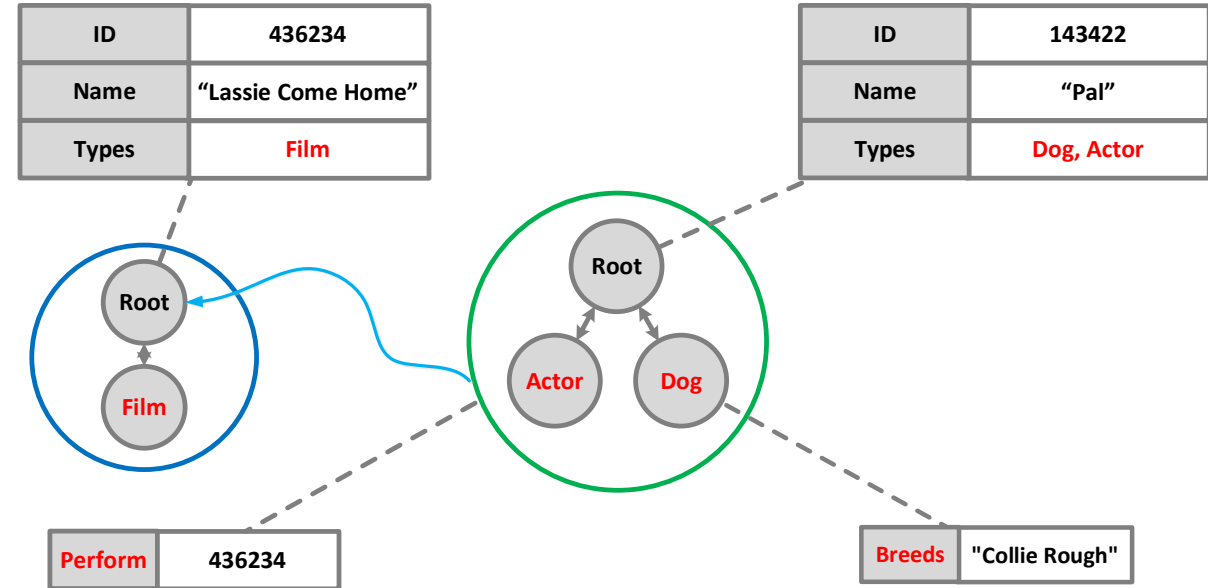
Challenges of Serving Satori

悟 SATORI

Satori: An ever-growing knowledge repository

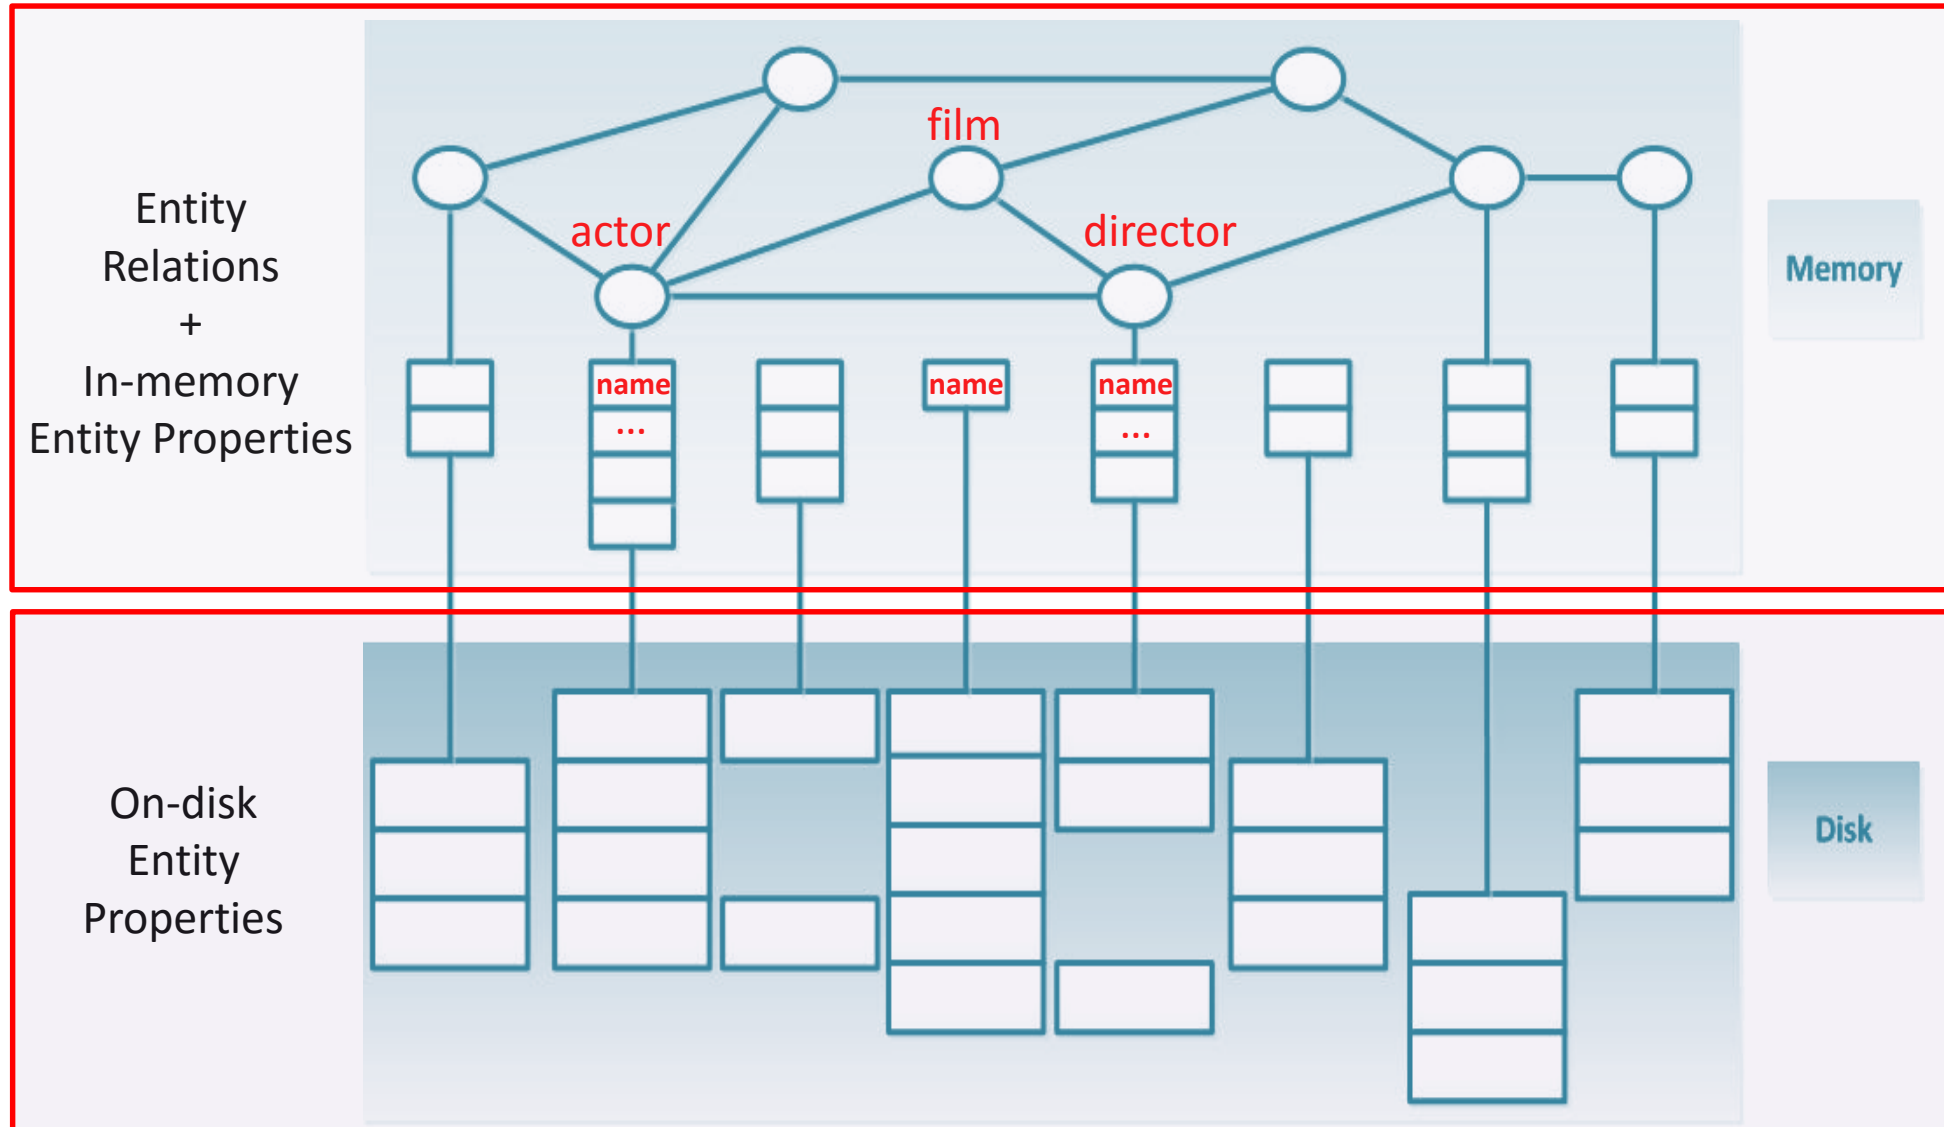
Raw RDF data	4T+
Entities	2.4B+
Triple Facts	20B+

- Complex data schema
 - Rich relations
 - Multi-typed entities



Modeling Multi-Typed Satori Entities in a Strongly Typed Manner

Trinity's Storage Architecture for Satori




Demo

SATORI

- bill gate
- bill gatekeeper
- bill gately
- bill gatens
- bill gates
- bill gates bill
- bill gates jnr

Satori Service Portal



Harvard University
 award.presenting_organization, award.ranked_item, award.winner, book.author, education.academic_institution, education.educational_institution ...

Harvard University

organization.organization, type.object

harvard university

internet.social_network_user, people.person, type.object

harvard university

internet.social_network_user, people.person, type.object

Harvard University

local.entity, type.object

Harvard University

local.entity, type.object

Harvard University



Harvard University is an American private Ivy League research university located in Cambridge, Massachusetts, United States, established in 1636 by the Massachusetts legislature. Harvard is the oldest institution of higher learning in the United States and the first corporation (officially The President and Fellows of Harvard College) chartered in the country. Harvard's history, influence,...

Types

award.presenting_organization, award.ranked_item, award.winner, book.author, education.academic_institution, education.educational_institution ...

Predicates

- education.educational_institution.total_enrollment
- education.educational_institution.color
- education.educational_institution.subsidiary_or_constituent_schools**
- education.educational_institution.number_of_staff
- education.educational_institution.honorary_degrees_awarded
- education.educational_institution.school_sports_team

[Prev Page](#) [Next Page](#)

Values

- "Harvard Extension School"
- "Harvard Medical School"
- "Harvard Business School"
- "Harvard College"
- "Harvard Division of Continuing Education"
- "John F. Kennedy School of Government"

[Prev Page](#) [Next Page](#)

Entity Explorer

Schema Graph

Meta Graph of Satori

Schema Type:

Go

Schema Path:



Go

Fields:

.bust_measurement mso/type.decimal
.date_of_birth mso/type.datetime
.eye_color mso/type.text
.first_name mso/type.string
.hair_color mso/type.text
.height mso/type.decimal
.hips_measurement mso/type.decimal
.last_name mso/type.string
.waist_measurement mso/type.decimal
.weight mso/type.decimal

Links:

.business_employment_tenure mso/business.employment_tenure
.children mso/people.person
.city_of_birth mso/location.location

mso/people.person	.quotation	mso/media_common.quotation
mso/media_common.quotation	.character	mso/fictional_universe.character
mso/fictional_universe.character	.appears_in_the se_fictional_universes	mso/fictional_universe.universe
mso/fictional_universe.universe	.literary_series	mso/book.literary_series
mso/book.literary_series	.author	mso/book.author

Schema Graph Services

Satori Knowledge Graph Access API

API Names	Availability	Description
GetEntityIdByName	Available	Gets a list of Trinity entity Ids by the specified entity name.
GetPredicatesByEntityId	Available	Gets a list of predicates for the entity with the specified Trinity entity Id.
GetValuesByEntityPredicate	Available	Gets the values of the specified predicates for the specified entity.
GetSubjectsByPredicateObject	Available	Gets the subjects for the given object and a predicate.
GetEntityIdBySatoriId	Available	Gets the corresponding Trinity entity Id for the specified Satori Guid.
GetSatoriIdByEntityId	Available	Gets the corresponding Satori Guid for the specified Trinity entity Id.
GetRankedEntityIdByName	Available	Gets a list of Trinity entity Ids by the specified entity name sorted by their static rank.
GetScoredValuesByEntityPredicate	Available	Gets the values of the specified predicates for the specified entity, sorted by confidence score.
GetSortScoredValuesByEntityPredicate	Available	Gets the values of the specified predicates for the specified entity, sorted by the column index (1 for...
GetEntityDescription	Available	Gets the description of the specified entityid.

Testing: [GetScoredValuesByEntityPredicate](#)

Please input test parameters below:

EntityId

Predicate

PredicateValue	ConfidenceScore	OverallScore
2987469205879	0.71	1.311128
116281907553515	0.71	1.409593
265920831012309	0.71	1.416611
58184534540412	0.71	1.710736
237856925167352	0.71	1.228339
57628320423344	0.71	1.272193

Page 1 of 1

Knowledge Serving APIs

Satori

film

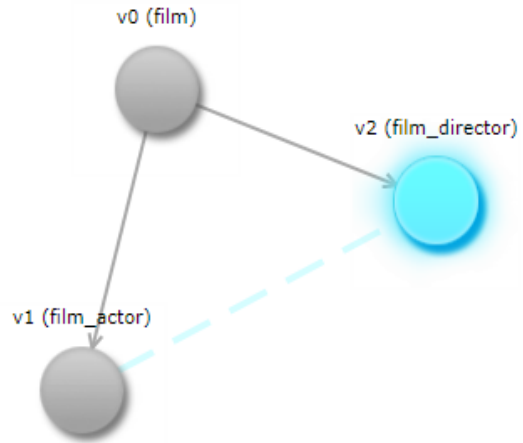
Add

View

TQL

Results

Graphical Query Interface



Node Information

Alias: v2

Type: film_director

URI: mso_film_director

Conditions:

Outputs:

- Steve Saari
- Steve Sacks
- Steve Sale
- Steve Salinaro
- Steve Salisian
- Steve Sanders
- Steve Sanguedolce
- Steve Saporito
- Steve Savage
- Steve Savitz
- Steve Scheffler

You could add some conditions a

name

=

Steve S

Submit

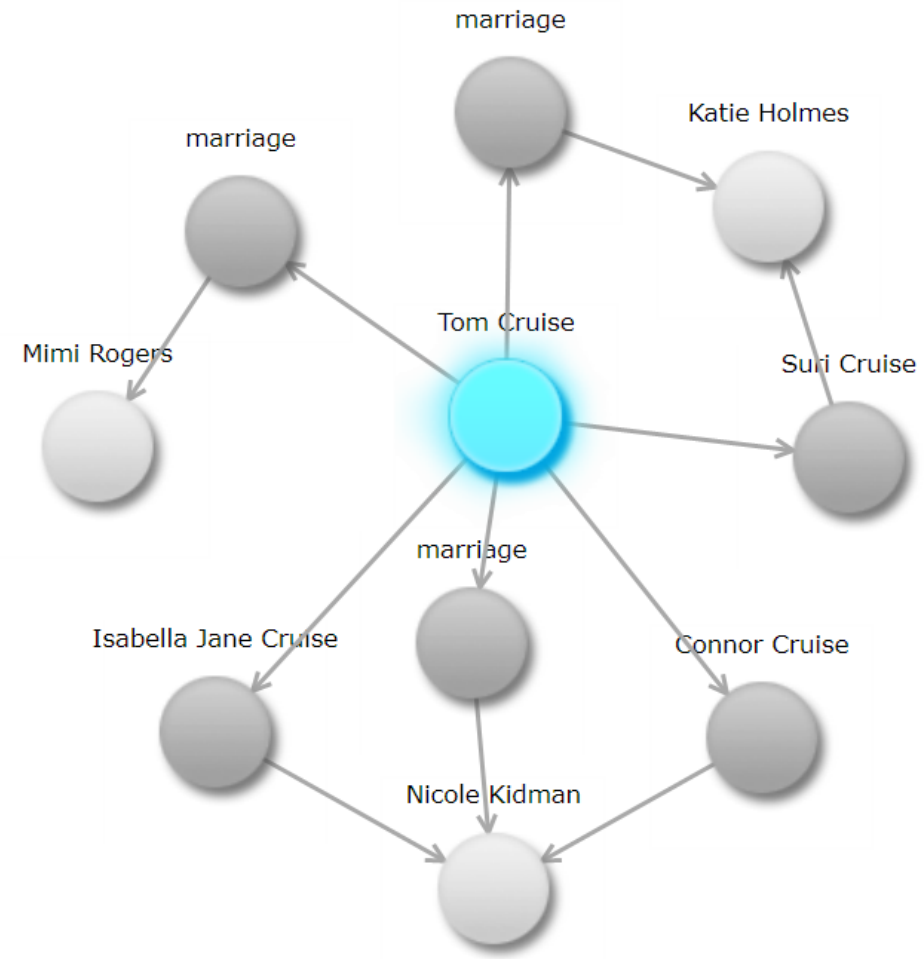
Add

Search

Tom Cruise, Mimi Rogers, Nicole Kidman, Katie Holmes

Results

[View](#)



Tom Cruise



Tom Cruise (born Thomas Cruise Mapother IV; July 3, 1962), is an American film actor and producer. He has been nominated for three Academy Awards and has won three Golden Globe Awards. He started his career at age 19 in the 1981 film Endless Love. After portraying supporting roles in Taps (1981) and The Outsiders (1983), his first leading role was in Risky Business, released in August 1983. Cruise became a full-fledged movie...

Types

award.nominee, award.winner, film.actor, film.director, film.producer, film.story_contributor ...

Relation Search

Thanks!

<http://www.graphengine.io/>