



Self-Organization in Distributed Semantic Storage System

Anne Augustin, Sebastian Koske, Marko Harasic, Hannes Mühleisen, Philipp
Obermeier, Kia Teymourian, Robert Tolksdorf

{firstname.lastname}@inf.fu-berlin.de

Presenter: Kia Teymourian

Freie Universität Berlin

Department of Mathematics and Computer Science

Networked Information Systems

<http://digipolis.ag-nbi.de>



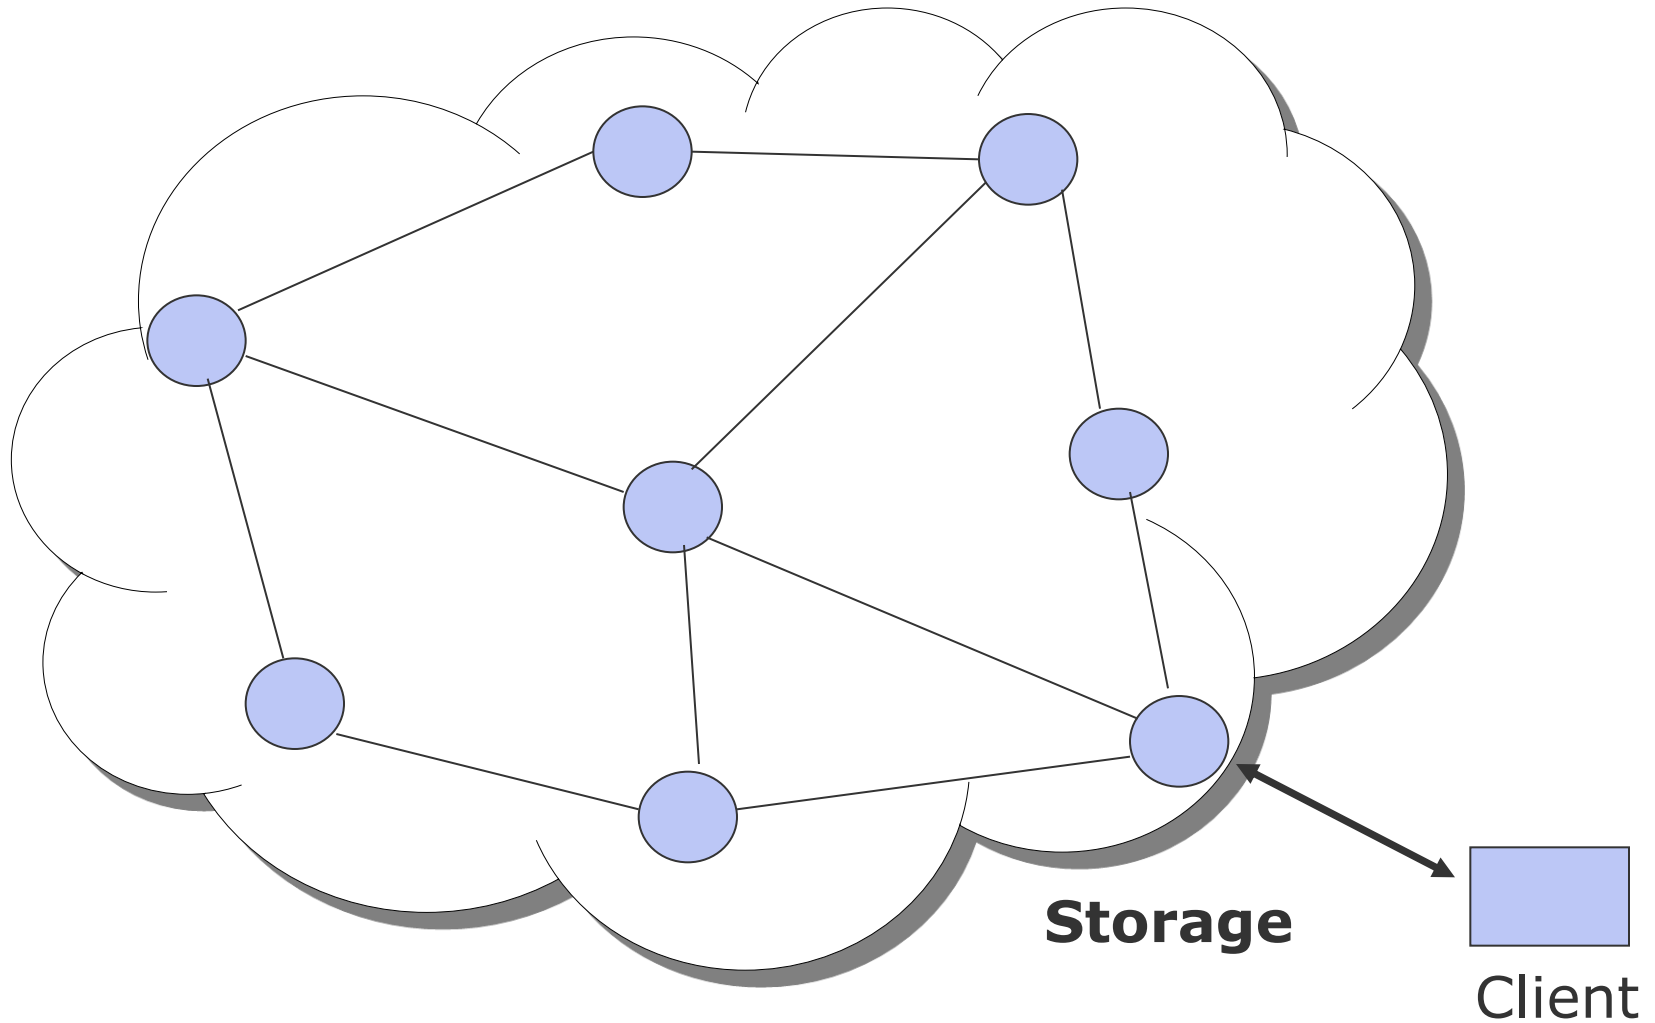
- Semantic Storage Service

- General Requirements
 - **High Scalability,**
 - **Reliability,**
 - **Resilience,**
 - **Inexpensiveness ,**
 - **Simplicity,**
 - **Security,**
 - **etc ...**
- Semantic requirements
 - **Reasoning**
 - **Optimized Metadata Storage**





Storage Service





State of the Art

	<i>Store</i>	<i>Description</i>		<i>#Triples</i>
centralized	Virtuoso	No reasoning, proprietary		1.0 B [9]
	AllegroGraph	RDFS supported, proprietary		1.1 B [16]
	Jena TDB	No reasoning, open source		1.7 B [31]
	OWLIM	RDFS / OWL-Lite supported, proprietary		12 B [26]
		<i>Overlay Network</i>	<i>Routing Cost</i>	
distributed	Edutella	Flooding	$O(N)$	-
	S-RDF	Restricted Flooding	$O(N)$	-
	RDFPeers	CHORD DHT	$O(\log(N))$	-
	YARS2	unspecified DHT	$O(N)$	6.8 B [14]
	GridVine	P-Grid	$O(\log(N))$	-

- Existing Distributed Storage Systems have some difficulties to scale out
- Billion triple challenge
 - DBpedia consists about 0.5 billion RDF triples
 - In future, billions , trillions, quadrillions, etc ...

The Self-Organized Semantic Storage Service (S4)



Swarms:


- Simplicity
- Dynamism
- Locality





Idea

- Operations are represented by ants
- Nodes can be ant colonies, ant hills
- Ants carry data chunks/queries




Data are considered
as brood or food




1. Ant Foraging

Ants are collectively able to find the shortest path between a food source and the nest.





2. Ant Brood Sorting



Lumer, E. D. and Faieta, B. 1994. Diversity and adaptation in populations of clustering ants. In *From Animals To Animats 3*, pp. 501-508.




Algorithms



- RDF (Resource Description Framework)
 - Language to represent knowledge in a formal way
 - RDF-statements (RDF-triples) are primitive which build a directed Graph

- RDF-triple <**S**ubject, **P**redicate, **O**bject>
 - Resource (**S**ubject) has a property (**P**redicate) with a value (**O**bject)
 - S and P have to be URIs; O can be an URI or a literal (e.g. integer)

<http://www.w3.org/1999/02/22-rdf-syntax-ns#type>



<http://birds.org/description/onto.rdf#sparrow>

<http://birds.org/description/onto.rdf#bird>

- Basic Triple Pattern e.g. <?, ?, **O**>
 - Primitive for lookup operations
 - Finds all triples with same value in the field (e.g. all triples with **O** as object)



Write Triples

- Create three ants, one for each triple (S,P,O)
- Each ant carrying the triple and the corresponding field as ant-template



- These ants search the cluster by following the pheromone which matches the triple
- After storing the triple, the ant returns home and amplifies the pheromone on its way.



Write Algorithm

initialize *age* to a given integer value > 0

while *age* > 0 **do**

if (drop probability currently high) **then**

Drop triple, drop the scent of cluster-resource and
in weaker quantity on the neighbor nodes and die

else

move to neighborhood node based on cluster-resource

age = *age* - 1

end if

end while

Drop triple, drop the scent of cluster-resource and
in weaker quantity on the neighbor nodes and die



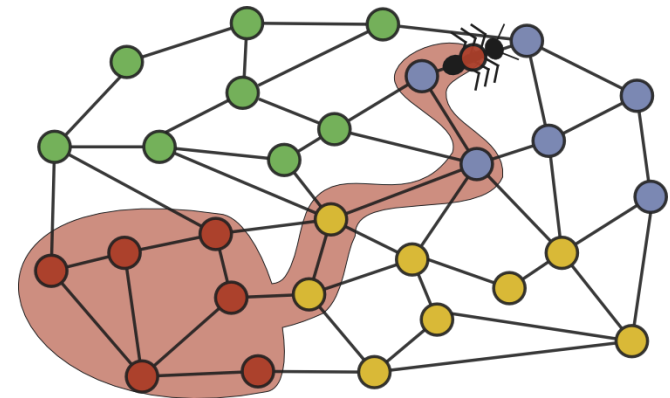

Read Algorithm

- Create an ant carrying the fixed part of the Basic Triple Pattern as ant-template



- The ant follows the pheromone which matches the template and searches in the appropriate layer
- When it finds a result, returns home carrying the result and amplifies the pheromone

- Ants choose their next node
 - Only with local knowledge
 - At random
 - With the probability of each edge depending on the pheromone strength matching the carried template
- Ants leave pheromones on their way back, which evaporate over the time



Shortest paths to clusters emerge



Read Algorithm

initialize *age* to a given integer value > 0

while *age* > 0 **do**

add the current node to memory

if (match found on current node) **then**

make a copy and use memory to return to the origin. Leave scent of cluster-resource on each node on the way back with each hop in a weaker quantity.

return the copy as result and die.

else

Move to neighborhood node based on cluster-resource

$age = age - 1$

end if


end while

die


Algorithms

Data Clusters, Brood Sorting

- Data is stored on a node
 - only with local knowledge
 - with a probability depending on the similarity of the data to the neighborhood
 - only if the similarity to the current node is over a threshold value



Creates clusters of similar data




- Clusters have to be maintained by moving misplaced data to its cluster

Algorithms

Brood Sorting (Cleaning)

- Misplaced data („corpses“)
- Special ants without template roam the network randomly and clean it from corpses
- If a corpse is found, it will be moved by a spawned *write-ant* carrying the corpse



**Leads to more homogeneous clusters und
increases recall**



Similarity Measurement



Similarity

- All creatures are equal...
... but some are more equal than others



<http://birds.org/description/onto.rdf#sparrow>

<http://birds.org/description/onto.rdf#duck>



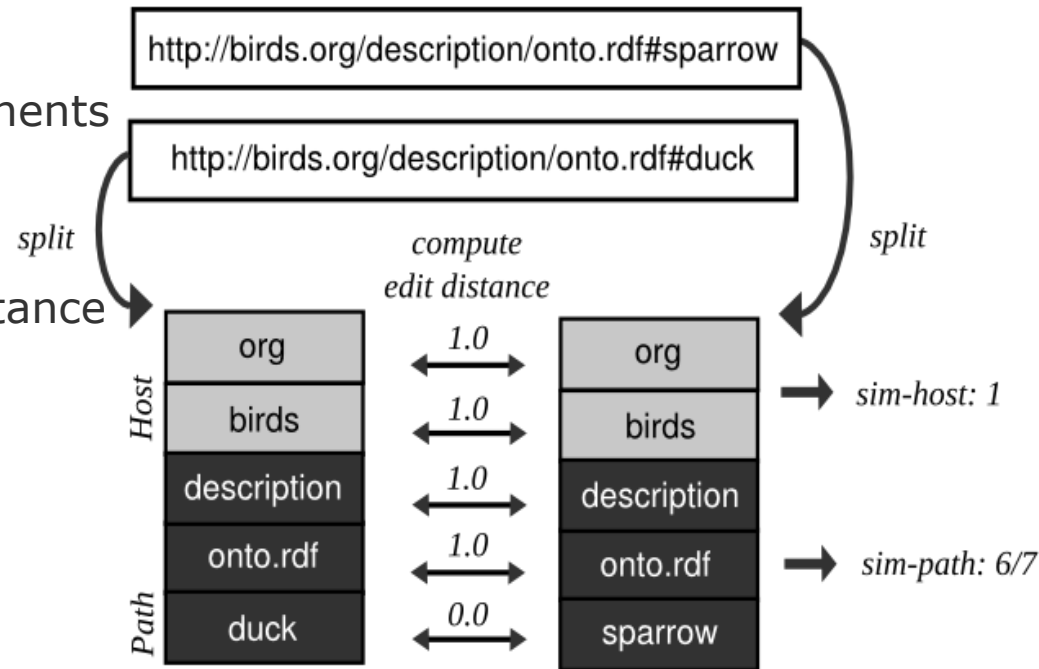
Similarity Measures

URI Similarity

- Idea: **Similar Things might have similar URIs**

- Algorithm :

1. Split path and domain components of the URIs
2. Compare parts of both URIs pairwise with Levenshtein distance
3. Sum up the weighted results
4. Similarity between the URIs is the sum



$$sim-URI: 9/10 \cdot 1 + 1/10 \cdot 6/7 \approx 0.986$$

- No Ontological Semantic
- Clusters based on namespace scheme



Similarity Measures

Fingerprint Similarity

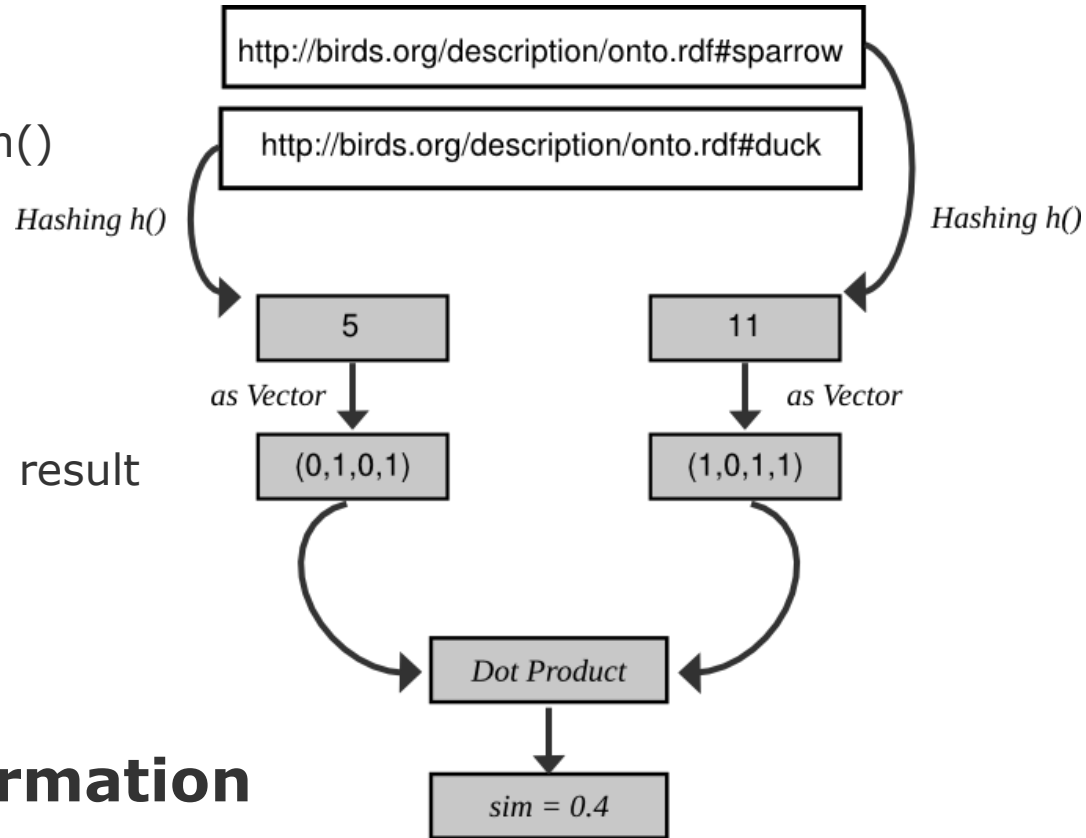
- Idea: String Hashing

- Algorithm :

1. Hash the URI Strings with $h()$
2. Interpret the values as bitvectors
3. Compare the vectors with the dot-product
4. Similarity is the normalized result

+ **Very fast to compute**

- **Lost of semantic information**






- Similarity based on ontological knowledge
- Requirements and assumptions:
 - Only local decisions
 - No global ontology
 - Similarity measure operating on a fragment of an ontology



- Extended behavior of out ants:
 - Carry triple and local type hierarchy
 - Type concentration on nodes determines responsibility for a certain type
 - Ants learn on their way and merge T-Boxes
 - Drop probability determined by similarity of type carried with type dominant on node

<cat,is-a,animal>
<dog,colored,white>

<bird,is-a,animal>
<cat,is-a,animal>



<cat,colored,grey>
<cat,is-a,type>
<cat,is-a,animal>
<bird,is-a,animal>
<dog,is-a,animal>

<dog,is-a,animal>
<dog,colored,brown>
<bird,colored,yellow>
<cat,is-a,animal>
<cat,colored,grey>




Simulation & Implementation



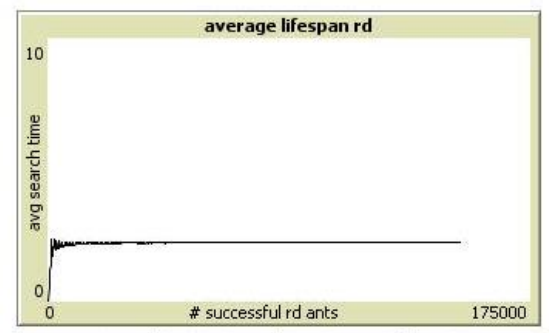
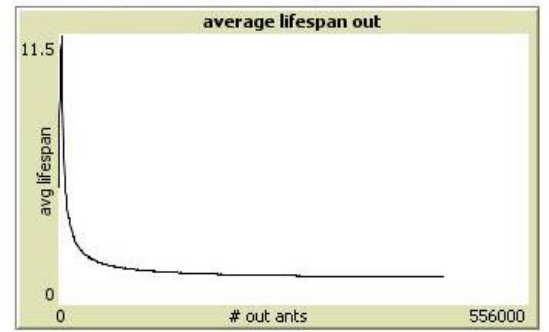
Simulation

Import RDFS-XML	reset import	Statements 15293
::::: Test Data Generation :::::::::::::::		
generate testdata	reset test data	test data 8000
generate templates	reset templates	test templates 9000
generate testdata and templates		reset both
num-testdata 10000	<input type="checkbox"/> On <input checked="" type="checkbox"/> Off random-order	
template-type one-field	import network	
set up	set up all	
::::: Ant Decision Settings ::::		
max-cluster-size 200	::::: Layout Settings ::::::	
<input type="checkbox"/> On <input checked="" type="checkbox"/> Off anti-overclustering	<input type="checkbox"/> On <input checked="" type="checkbox"/> Off node-information	
p-drop mod2	<input type="checkbox"/> On <input checked="" type="checkbox"/> Off show-distribution	
pick-up cleaner 1	<input type="checkbox"/> On <input checked="" type="checkbox"/> Off resize-nodes?	
tolerance 0.8	background-color black	
export output	avg-sim-testdata	
clear output	avg-sim-network	



OUT	run R	clean up	birth-node random
RD	step S	kill cleaners	node-selection 368
num-triples 1000	set-age 30	num-ants 1	distribute randomly

nodes 513	triples 456000	out ants 0	rd ants 117	cleaning ants 0
--------------	-------------------	---------------	----------------	--------------------

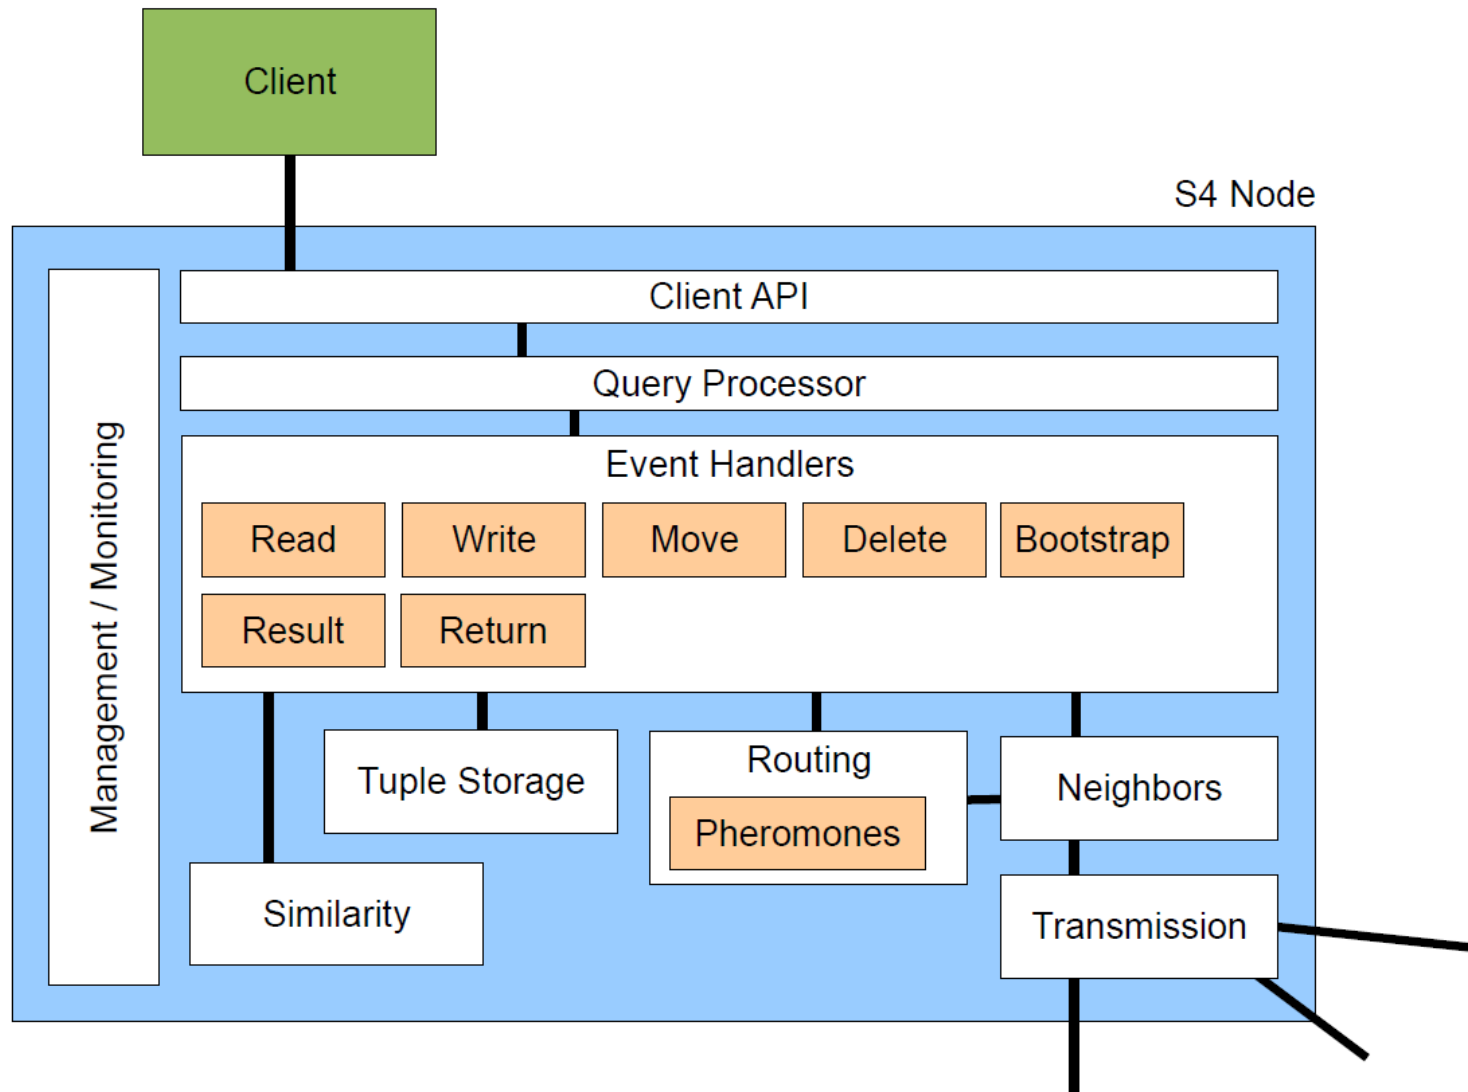


success rd ants 149936	failed rd ants 947	search time rd 2.22289	avg lifespan out 1.15584
---------------------------	-----------------------	---------------------------	-----------------------------

sim-measure Hashed	comp-type Integer
number-digits 20	max-hash-constant 20
run test write	calculate constant
run test read	generate testrun data



Implementation, Architecture





Evaluation



Evaluation

Test environment

- First Beta Implementation
- 10 computers simulating 10, 20, 30, 40 nodes
 - Nodes form a random graph
 - No connected nodes on the same computer

- Test system

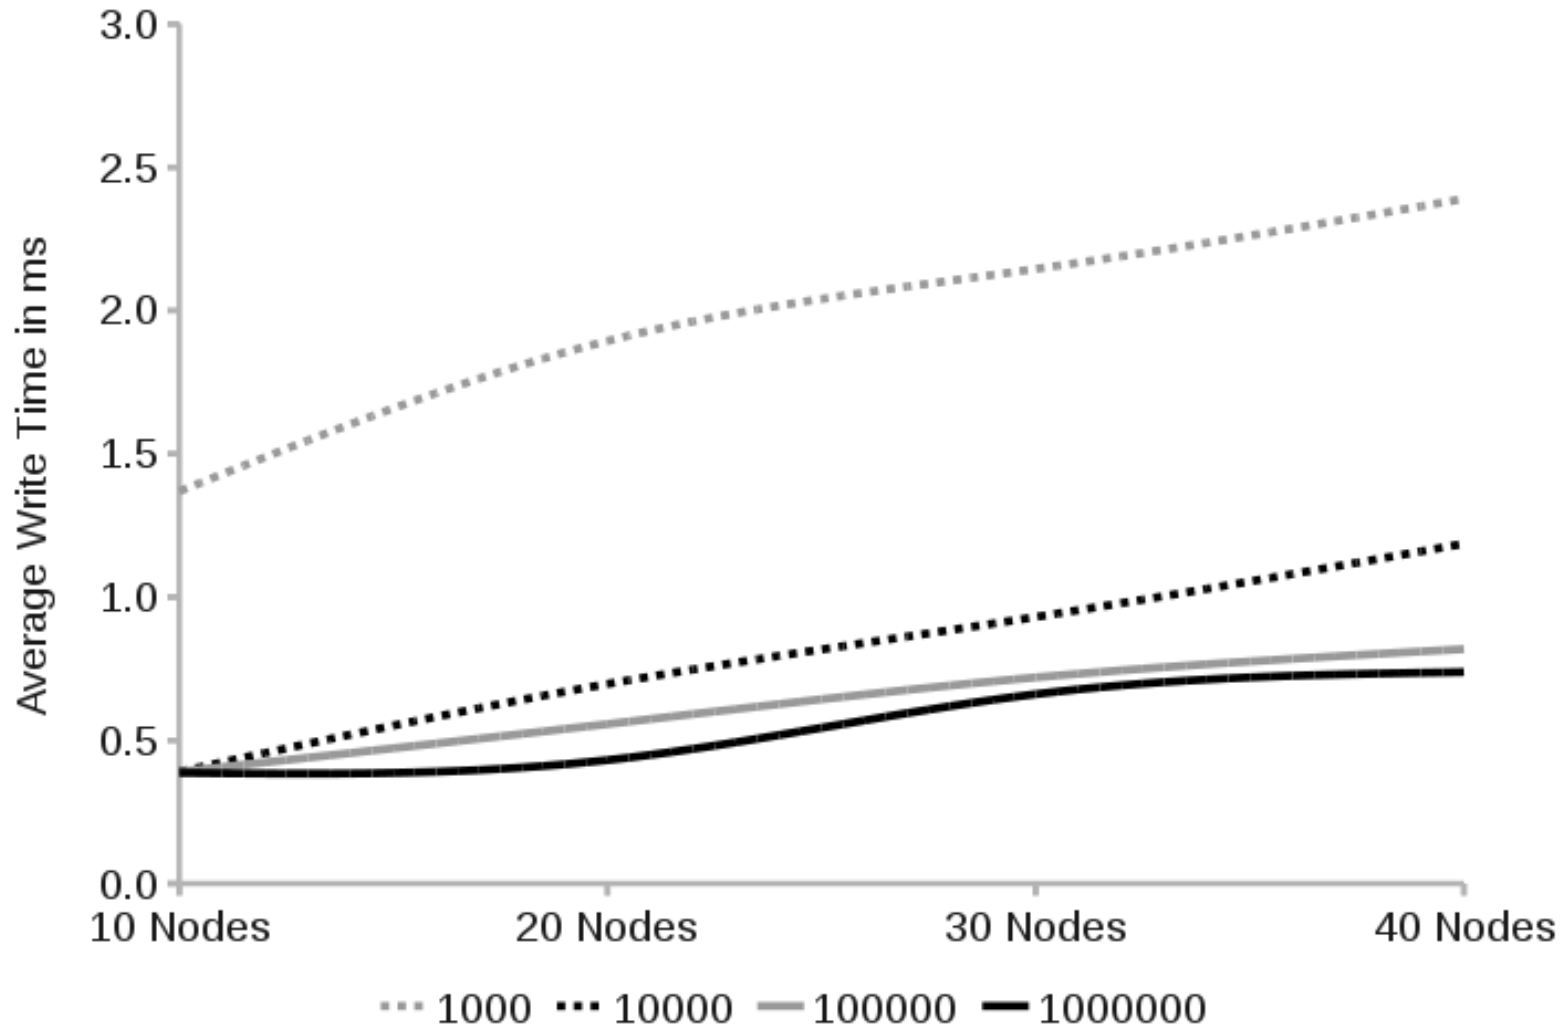
CPU	Pentium 4 2.4 Ghz
RAM	512 MByte
Network	100 MBit/s Ethernet
Operating System	Debian Linux Kernel 2.6.24
Java	1.5.0.17

- Test data

Data Source	WordNet Ontology
Write	Sets of 1K, 10K, 100K, 1M triples
Read	10k random operations for each set




Evaluation Write Operation





Evaluation Read Operation





Future work

- Semantic Similarity
- Swarm-based Reasoning
- Efficient full SPARQL support
- Use geographical information to cluster even better



Conclusion



- We can use self-organization to cluster similar things, and build scalable storage service
- The evaluation of our initial Alpha implementation shows that such swarm-based storage system can scale out.



Thank You

Questions?



Image Reference



- http://newsitemstoday.today.com/files/2009/04/cat_and_dog_toxic_substances.JPG
- http://1.bp.blogspot.com/_TbfuMu7x3_s/SNrOuYUjxTI/AAAAAAAAAKA/f_s8I8bpyLI/s400/antclusterial.jpg
- <http://onionesquereality.wordpress.com/2008/01/31/adaptive-routing-taking-cues-from-stigmergy/>
- <http://www.answers.com/topic/ant-colony-optimization>
- http://upload.wikimedia.org/wikipedia/commons/1/17/Aco_shortpath.svg



Extra Slides



URI Similarity

- URI Similarity:
- Take pairwise Levenshtein-distance in host and path parts eg:

$$sim_{host} = \sum_{i=1}^{\min(k,l)} c_i edit(m_{k-i}, n_{l-i})$$

- Weight components along their hierarchy, eg:

$$c_i = \frac{2^{\max(k,l)-i}}{2^{\max(k,l)} - 1}$$

- Weight host and path importance, eg. 0.9/0.1
- Cats and dogs are quite similar:

$$1 \cdot 0.9 + \left(\frac{2}{3} \cdot 1 + \frac{1}{3} \cdot 0\right) \cdot 0.1 = 0.9\bar{6}$$