

Contribution to the Analysis of the Design-Space of a Distributed Transformation Engine

Jolan PHILIPPE

PhD Defense, speciality: Computer Science

Referees:	Jesús SÁNCHEZ CUADRADO Matthias TICHY	Associate professor, Universidad de Murcia, Spain Professor, Ulm University, Germany
Examiners:	Thomas LEDOUX Leen LAMBERS Antonio VALLECILLO	Professor, IMT Atlantique, France Professor, Brandenburg University of Technology, Germany Professor, University of Málaga, Spain
Ph.D. director:	Gerson SUNYE	Associate professor, University of Nantes, France
Ph.D. advisors:	Hélène COULLON Massimo TISI	Associate professor, Institut Mines-Telecom Atlantique, France Associate professor, Institut Mines-Telecom Atlantique, France



lowcomote



19th December 2022



naomod

THÈSE DE DOCTORAT DE

L'ÉCOLE NATIONALE SUPÉRIEURE
MINES-TÉLÉCOM ATLANTIQUE BRETAGNE
PAYS-DE-LA-LOIRE - IMT ATLANTIQUE

ÉCOLE DOCTORALE N° 601
*Mathématiques et Sciences et Technologies
de l'Information et de la Communication*
Spécialité : *Informatique*

Par
Jolan PHILIPPE

**Contribution to the Analysis of the Design-Space of a Distributed
Transformation Engine**

Thèse présentée et soutenue à Nantes, le tbd
Unité de recherche : Laboratoire des Sciences du Numérique de Nantes
Thèse N° : tbd

Rapporteurs avant soutenance :

Jesus SANCHEZ CUADRADO Associate professor, Universidad de Murcia, Spain
Mathias TICHY Professor, Brandenburg University of Technology, Germany

Composition du Jury :

Président : Thomas LÉDOUX Professor, IMT Atlantique, France
Examinateurs : Leen LAMBERS Professor, University of Ulm, Germany
Antonio VALLECILLO Professor, University of Málaga, Spain
Dir. de thèse : Gerson SUNYE Associate professor, University of Nantes (France)
Co-dir. de thèse : Massimo TISI Associate Professor, Institut Mines-Télécom Atlantique (France)
Hélène COULLON Associate Professor, Institut Mines-Télécom Atlantique (France)

1 CONTEXT & MOTIVATION

2 CONTRIBUTIONS

2.1 SPARKTE: A DISTRIBUTED TRANSFORMATION ENGINE


2.2 DISTRIBUTED QUERY EVALUATION STRATEGIES

2.3 FEATURE ANALYSIS

3 CONCLUSION

1

CONTEXT & PROBLEMATIC





lowcomote



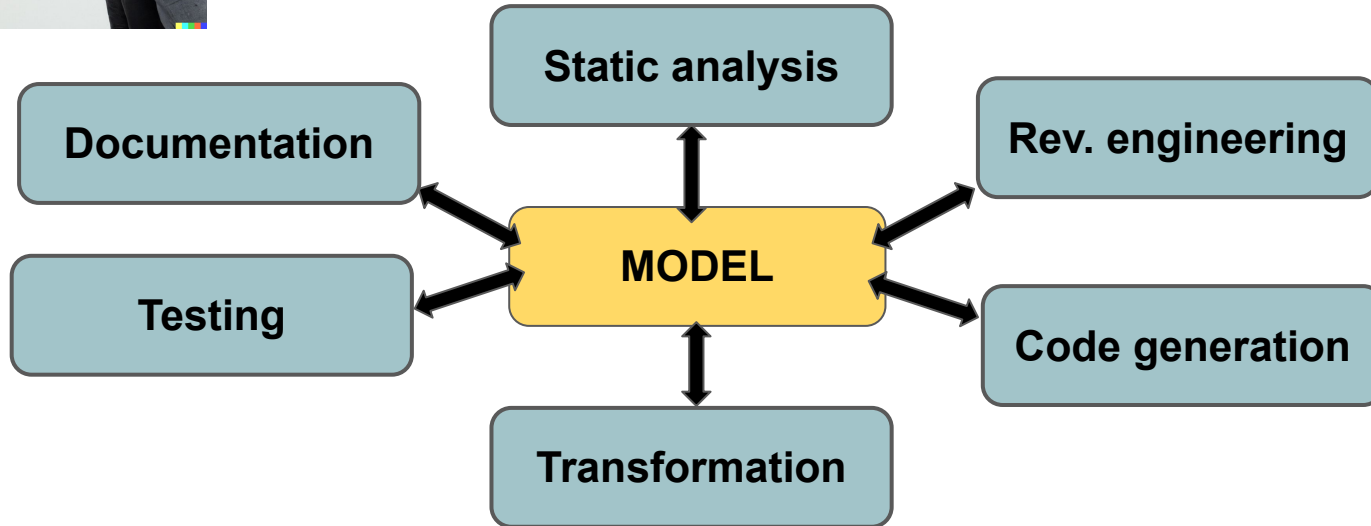
Lowcomote is a H2020-ITN project aiming at training 15 PhD students, and build a low-code development platforms based on

- **Model-Driven Engineering**
- **Cloud Computing**
- **Machine Learning**







- **Software engineering** approach
- Models as the **central artifact** to represent systems






- **Software engineering** approach
- Models as the **central artifact** to represent systems

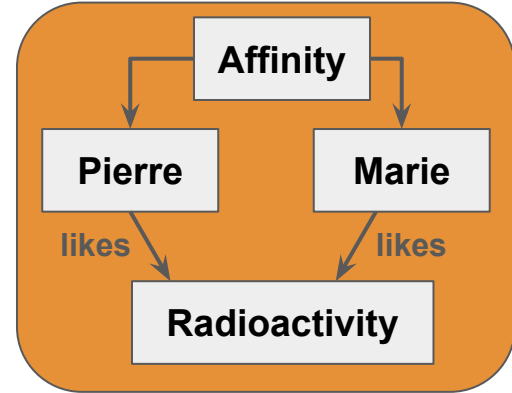





SOURCE MODEL



TARGET MODEL



SOURCE MODEL





TRANSFORMATION

rule copy (e: Element)
output:
new Element (content ← e.content)


rule affinity (e₁: Element,
e₂: Element)
matching:
 $(e_1.likes) \cap (e_2.likes) \neq \emptyset$
output:
new Affinity (from ← e₁, to ← e₂)

TARGET MODEL



Many transformation languages: ATL, ETL, QVT, Henshin, Viatra, ...

SOURCE MODEL



TRANSFORMATION

rule copy (e: Element)

output:

new Element (content \leftarrow e.content)

rule affinity (e_1 : Element,
 e_2 : Element)


matching:

$(e_1.\text{likes}) \cap (e_2.\text{likes}) \neq \emptyset$

output:

new Affinity (from $\leftarrow e_1$, to $\leftarrow e_2$)

TARGET MODEL



Many transformation languages: ATL, ETL, QVT, Henshin, Viatra, ...

The expression $e_i.\text{likes}$ can be expressed as a **query**

System

Marie Curie
May 4th, 1898
I just discovered radioactivity!

3 comments

Comment

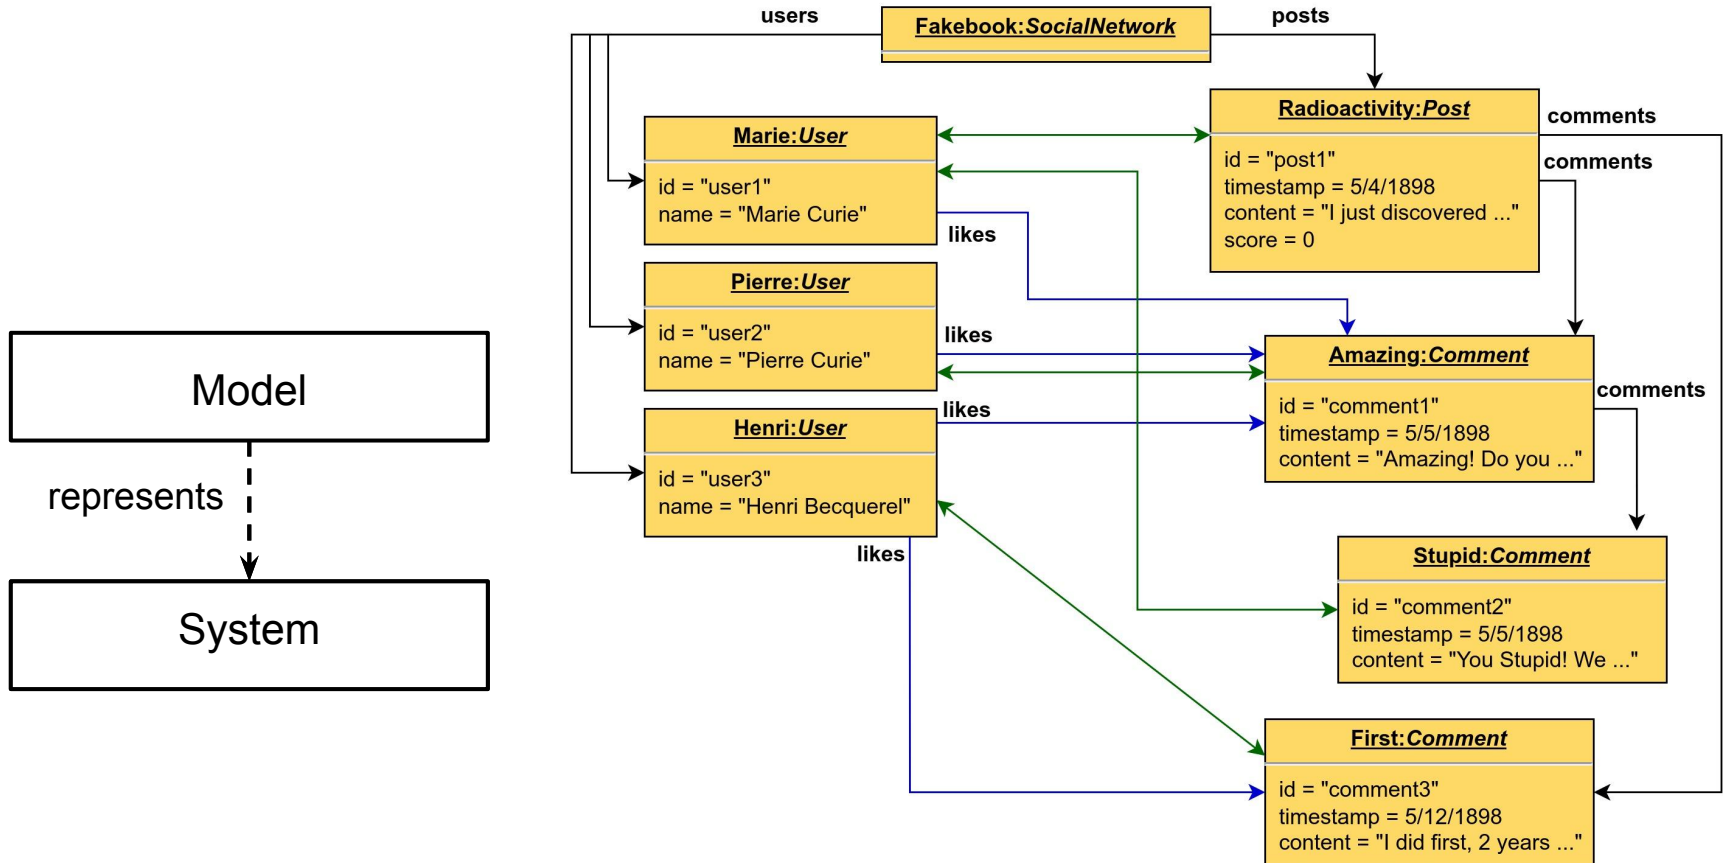
Pierre Curie
Amazing! Do you want to marry me?
Like Comment May 5th, 1898 3

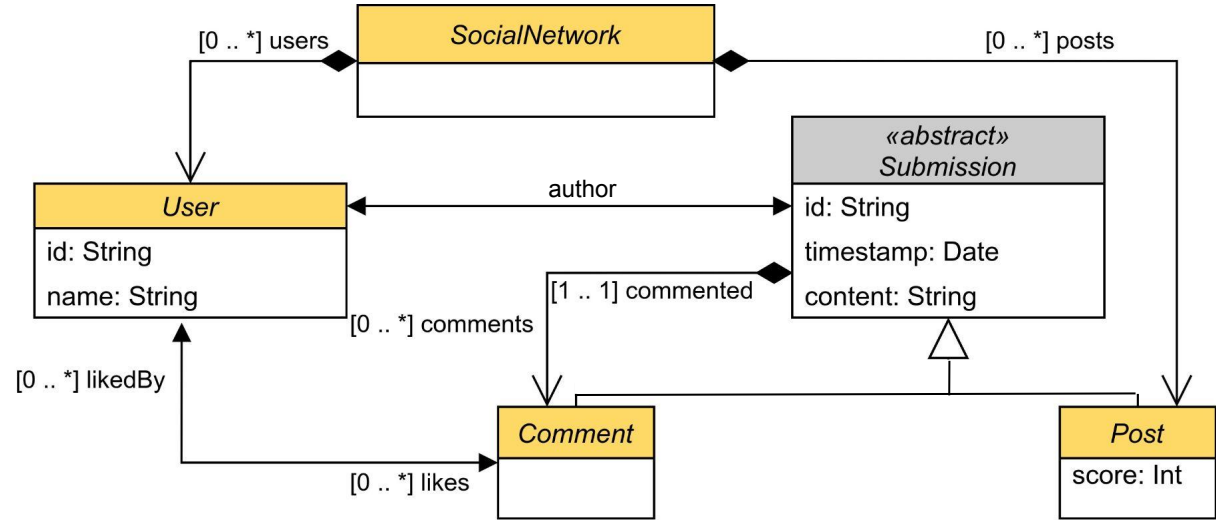
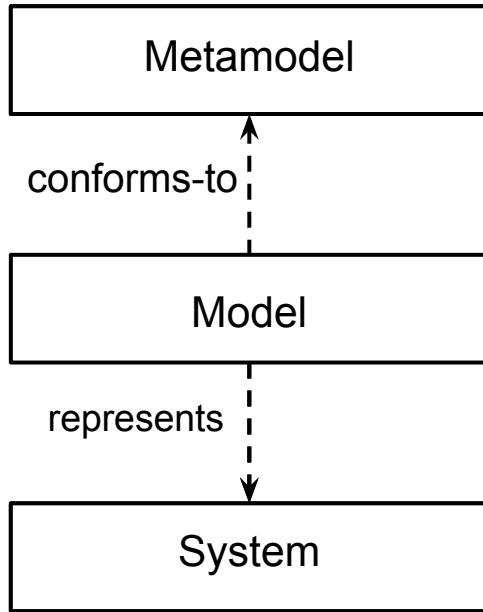
Marie Curie
You stupid! We already are
Like Comment May 5th, 1898

Henri Becquerel
I did first, 2 years ago.
Like Comment May 12th, 1898 1

Use case: A platform for analysing a social network


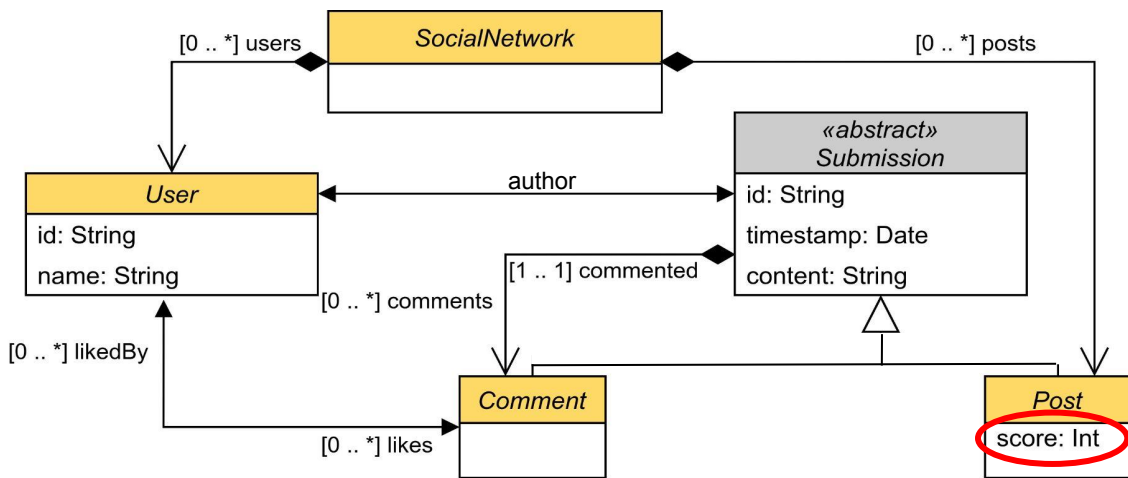
1 CONTEXT & MOTIVATION





Example 1: Give an activity score for posts in a social network

1 CONTEXT & MOTIVATION


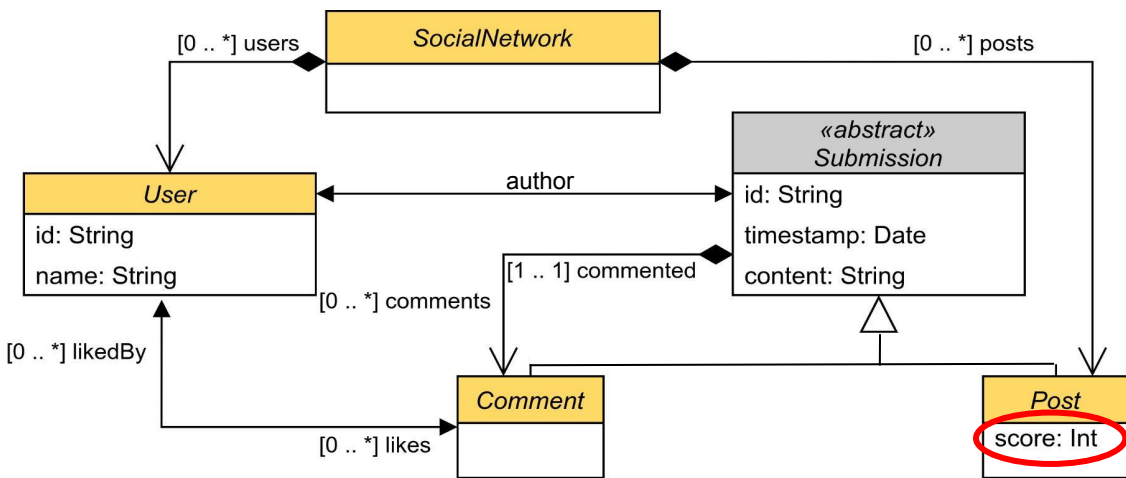


Example: $\text{score}(p: \text{Post}) = \# \text{ comments} \times 10 + \# \text{ likes}$

$\text{score}(\text{Radioactivity}) = 3 \times 10 + 4 = 34$

Example 1: Give an activity score for posts in a social network

1 CONTEXT & MOTIVATION



Example: $\text{score}(p: \text{Post}) = \# \text{ comments} \times 10 + \# \text{ likes}$

rule Post2ScoredPost (p:Post)

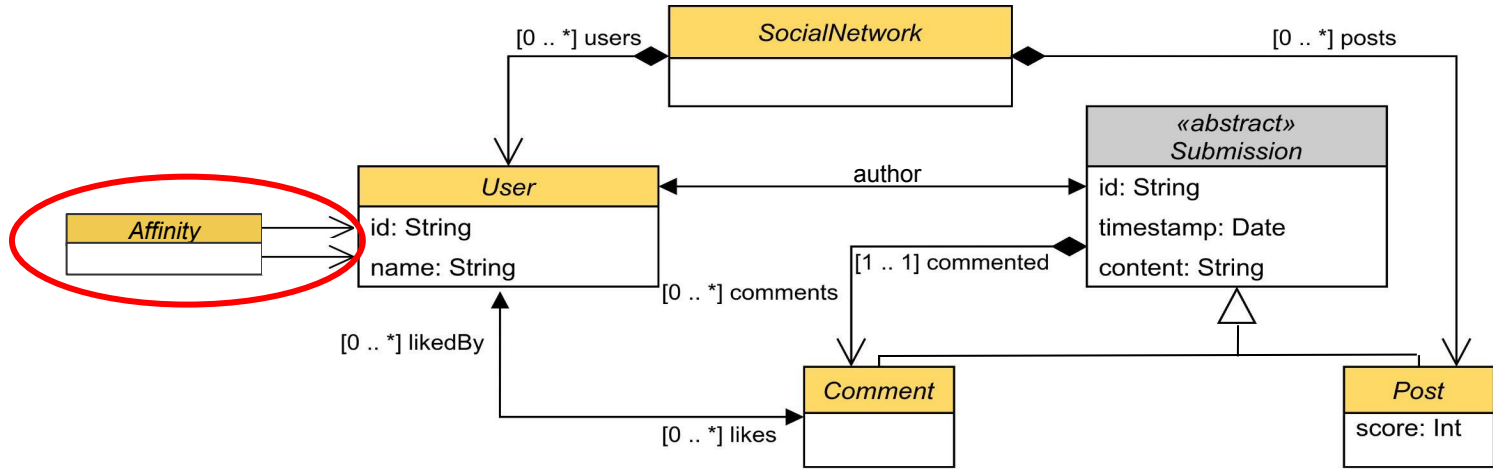
output:

```
new Post (id ← p.id,
          timestamp ← p.timestamp
          content ← p.content,
          score ← score(p))
```

score as a query

Example 2: Look for user affinities in a social network

1 CONTEXT & MOTIVATION



Example: Comment at least 3 same posts




Pierre Curie



Marie Curie

Example 2: Look for user affinities in a social network

1 CONTEXT & MOTIVATION



Example: Comment at least 3 same posts


rule FindAffinity (u_1 :User, u_2 :User)

matching:

$\text{commentedPosts}(u_1) \cap \text{commentedPosts}(u_2) \geq 3$

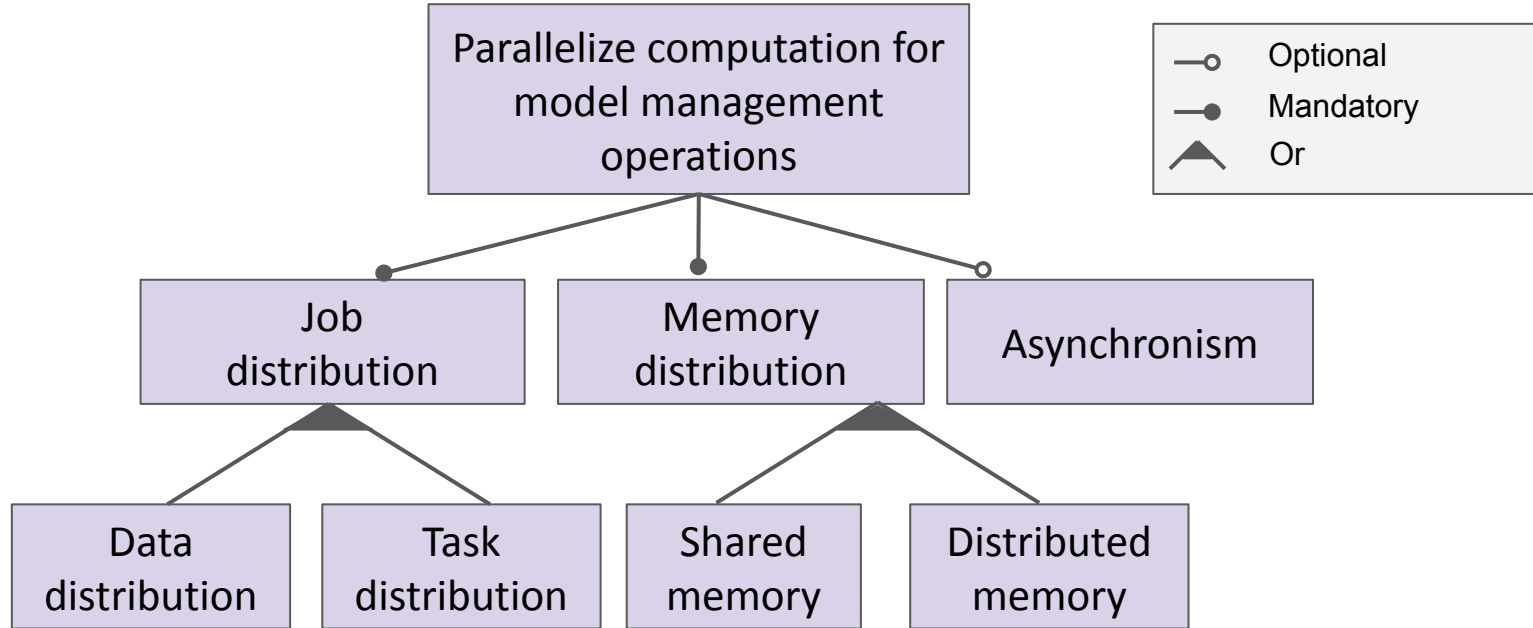
output:

$\text{new Affinity}(\text{user}_1 \leftarrow u_1, \text{user}_2 \leftarrow u_2)$



- Computational complexity
 - Size of the model
 - Storage and memory constraints
 - Scalability with increasing resources
 - Implicit optimization
- Two main approaches
- Avoid computation
 - Parallelize computation

[1] Dimitrios S. Kolovos, Louis M. Rose, Nicholas Drivalos Matragkas, Richard F. Paige, Esther Guerra, Jesús Sánchez Cuadrado, Juan de Lara, István Ráth, Dániel Varró, Massimo Tisi, Jordi Cabot. **A research roadmap towards achieving scalability in model driven engineering.** *BigMDE@STAF 2013*

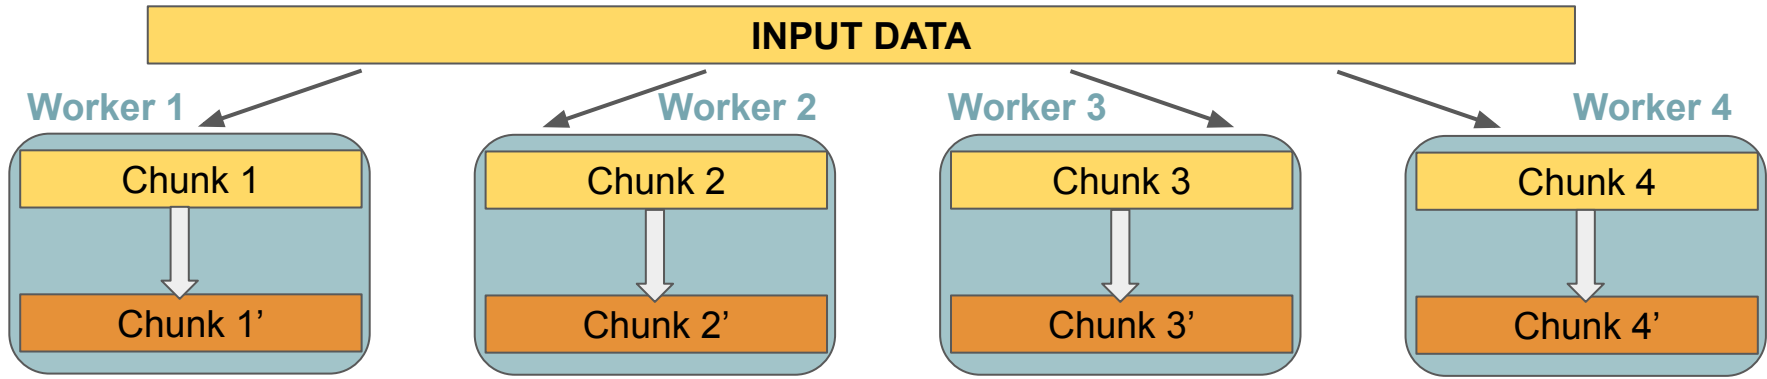




Parallelization in model transformation

	<i>Model query</i>	<i>Model transfo.</i>	<i>Pattern match.</i>	<i>Optimization</i>	<i>Shared mem.</i>	<i>Distrib. mem.</i>	<i>Task-parallel</i>	<i>Data-parallel</i>	<i>Asynchronism</i>
Amine Benelallam et al. « Efficient model partitioning for distributed model ... » SLE 2016		X		X		X		X	
Amine Benelallam et al. « ATL-MR: model transformation on MapReduce » SPLASH 2015		X				X		X	
Loli Burgueño et al. « A Linda-Based platform for the parallel execution ... » IST 2016		X			X			X	X
Loli Burgueño et al. « Towards distributed model transformations with LinTra » JISBD 2016		X		X		X		X	X
Loli Burgueño et al. « Parallel in-place model transformations with LinTra » CEUR-WS 2015		X			X		X		X
Jesús S. Cuadrado et al. « Efficient execution of ATL model transformations ... » TSE 2020		X			X			X	
Gábor Imre et al. « Parallel graph transformations on multicore systems » MSEPT 2012		X			X		X		
Christian Krause et al. « Implementing graph transformations in the BSP model » FASE 2014			X			X		X	
Sina Madani et al. « Distributed model validation with Epsilon » SSM 2021	X				X	X		X	
Sina Madani et al. « Towards optimisation of model queries: a parallel ... » ECMFA 2019	X			X	X		X		
Gergely Mezei et al. « Towards truly parallel model transformations: a ... » EURCON 2019			X			X	X		
Massimo Tisi et al. « Parallel execution of ATL transformation rules » MODELS 2013		X			X		X		
Le-Duc Tung et al. « Towards systematic parallelization of graph transfo. ... » IJPP 2017		X				X		X	
Tamás Vajk et al. « Runtime model validation with parallel object ... » MoDeVva 2011	X				X		X		


Parallelization in model transformation

	<i>Model query</i>	<i>Model transfo.</i>	<i>Pattern match.</i>	<i>Optimization</i>	<i>Shared mem.</i>	<i>Distrib. mem.</i>	<i>Task-parallel</i>	<i>Data-parallel</i>	<i>Asynchronism</i>
Amine Benelallam et al. « Efficient model partitioning for distributed model ... » SLE 2016		X		X		X		X	
Amine Benelallam et al. « ATL-MR: model transformation on MapReduce » SPLASH 2015		X				X		X	
Loli Burgueño et al. « A Linda-Based platform for the parallel execution ... » IST 2016		X			X			X	X
Loli Burgueño et al. « Towards distributed model transformations with LinTra » JISBD 2016		X		X		X		X	X
Loli Burgueño et al. « Parallel in-place model transformations with LinTra » CEUR-WS 2015		X			X		X		X
Jesús S. Cuadrado et al. « Efficient execution of ATL model transformations ... » TSE 2020		X			X			X	
Gábor Imre et al. « Parallel graph transformations on multicore systems » MSEPT 2012		X			X		X		
Christian Krause et al. « Implementing graph transformations in the BSP model » FASE 2014			X			X		X	
Sina Madani et al. « Distributed model validation with Epsilon » SSM 2021	X				X	X		X	
Sina Madani et al. « Towards optimisation of model queries: a parallel ... » ECMFA 2019	X			X	X		X		
Gergely Mezei et al. « Towards truly parallel model transformations: a ... » EURCON 2019			X			X	X		
Massimo Tisi et al. « Parallel execution of ATL transformation rules » MODELS 2013		X			X		X		
Le-Duc Tung et al. « Towards systematic parallelization of graph transfo. ... » IJPP 2017		X				X		X	
Tamás Vajk et al. « Runtime model validation with parallel object ... » MoDeVva 2011	X				X		X		






MapReduce



**Pregel
(vertex-centric)**



Blackboard

- Large number of distributed engines
 - Designed with **≠ purposes**
 - Following **≠ design choices**
 - Implemented on **≠ languages** for **≠ infrastructures**

⇒ What are the optimal design choices for a given case?



- Automatic adapted strategy
 - Pattern matching (Bergman et al.)
- Classification of features of MDE solution
 - For languages (Tamura et al., M Rose et al.)
 - Transformation approaches (Czarnecki et al., Kahani et al.)
 - Performance oriented (Groner et al.)
 - Specific topic: bi-directionality (Hidaka et al.)

Optimization in model transformation

	<i>Model query</i>	<i>Model transfo.</i>	<i>Pattern match.</i>	<i>Optimization</i>	<i>Shared mem.</i>	<i>Distrib. mem.</i>	<i>Task-parallel</i>	<i>Data-parallel</i>	<i>Asynchronism</i>
Amine Benelallam et al. « Efficient model partitioning for distributed model ... » SLE 2016		X		X		X		X	
Amine Benelallam et al. « ATL-MR: model transformation on MapReduce » SPLASH 2015		X				X		X	
Loli Burgueño et al. « A Linda-Based platform for the parallel execution ... » IST 2016		X			X			X	X
Loli Burgueño et al. « Towards distributed model transformations with LinTra » JISBD 2016		X		X		X		X	X
Loli Burgueño et al. « Parallel in-place model transformations with LinTra » CEUR-WS 2015		X			X		X		X
Jesús S. Cuadrado et al. « Efficient execution of ATL model transformations ... » TSE 2020		X			X			X	
Gábor Imre et al. « Parallel graph transformations on multicore systems » MSEPT 2012		X			X		X		
Christian Krause et al. « Implementing graph transformations in the BSP model » FASE 2014			X			X		X	
Sina Madani et al. « Distributed model validation with Epsilon » SSM 2021	X				X	X		X	
Sina Madani et al. « Towards optimisation of model queries: a parallel ... » ECMFA 2019	X			X	X		X		
Gergely Mezei et al. « Towards truly parallel model transformations: a ... » EURCON 2019			X			X	X		
Massimo Tisi et al. « Parallel execution of ATL transformation rules » MODELS 2013		X			X		X		
Le-Duc Tung et al. « Towards systematic parallelization of graph transfo. ... » IJPP 2017		X				X		X	
Tamás Vajk et al. « Runtime model validation with parallel object ... » MoDeVva 2011	X				X		X		

- What solution to use?
- How to optimally configure a solution?

Problem 1:
Many solutions for
executing
rules distributively

Problem 2:
Many solutions for
executing
queries distributively

Problem 3:
Lack of unified proposition
for comparing design
choices

- **Goal:** Getting an insight of how design choices impact scalability of a distributed transformation

Problem 1:

Many solutions for executing rules distributively

Evaluation of distributed design choices for **rule execution**

- Building a new distributed transformation engine: SparkTE

Problem 1:

Many solutions for executing rules distributively

Evaluation of distributed design choices for **rule execution**

- Building a new distributed transformation engine: SparkTE

Problem 2:

Many solutions for executing queries distributively

Evaluation of distributed design choices for **query execution**

- Analysing different distributed execution strategies for a query

Problem 1:

Many solutions for executing rules distributively

Evaluation of distributed design choices for **rule execution**

- Building a new distributed transformation engine: SparkTE

Problem 2:

Many solutions for executing queries distributively

Evaluation of distributed design choices for **query execution**

- Analysing different distributed execution strategies for a query

Problem 3:

Lack of unified proposition for comparing design choices

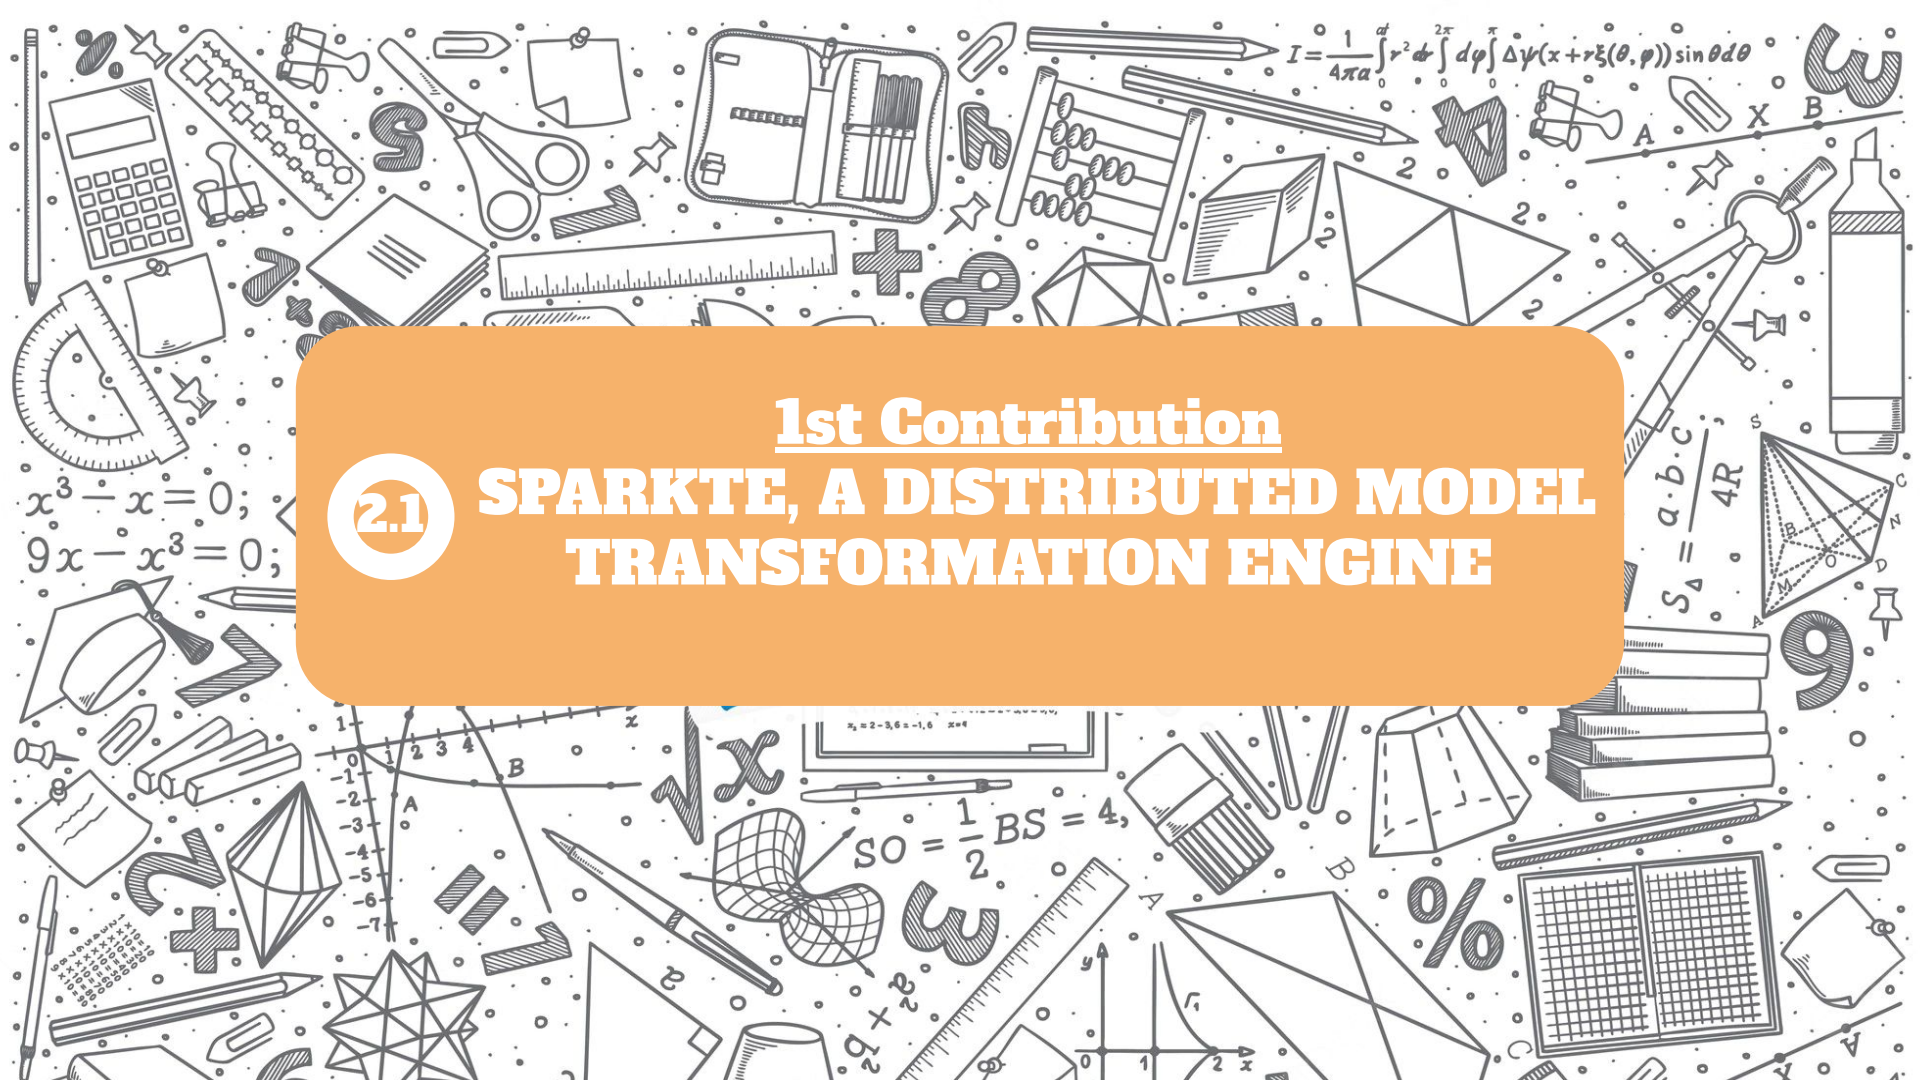
Make possible configurable distributed transformation

- Modeling the design space
- Making the configurable engine: Configurable SparkTE

2.1

1st Contribution

SPARKTE, A DISTRIBUTED MODEL TRANSFORMATION ENGINE



$$I = \frac{1}{4\pi\alpha_0} \int_0^{2\pi} \int_0^{\pi} r^2 dr \int_0^{\pi} \Delta\psi(x+r\xi(\theta, \varphi)) \sin\theta d\theta$$

$$x^3 - x = 0;$$
$$9x - x^3 = 0;$$

$$S_A = \frac{a \cdot b \cdot c}{4R}$$

$$SO = \frac{1}{2} BS = 4,$$

$$a^2 + b^2 = c^2$$

Many solutions for executing rules distributively

- Evaluation of distributed design choices for **rule execution**
 - An engine with design choices for rule execution: SparkTE
 - Prove design choices have no impact on the result
 - Evaluate the scalability of a such engine



```
CoqIDE
Coq_examples.v
theorem and_comm : A A B -> B A A.
Proof.
  prove_imp.
1 subgoals
A : Prop
B : Prop
H : A A B
----- (1/1)
B A A
Ready in Propositional_Logic, provingand_comm
Line: 20 Char: 1 CoqIDE started
```

- Designed for specifying semantics
- A proof assistant based on **Hoare logic**
- Extraction mechanism (to ML lang)

- **DSL** for rule-based model transformation
- Made for **reasoning on transformations**
- Can **reason on the semantic** of the transformation


atlanmod/**coqtl**




CoqTL allows users to write model transformations and prove engine/transformation correctness in Coq


4 Contributors 12 Issues 10 Stars 14 Forks





- Popular distributed computing for **large-scale data** processing
- Support for **many paradigms**
 - MapReduce, vertex-based (Pregel), ...
- **Open-source**

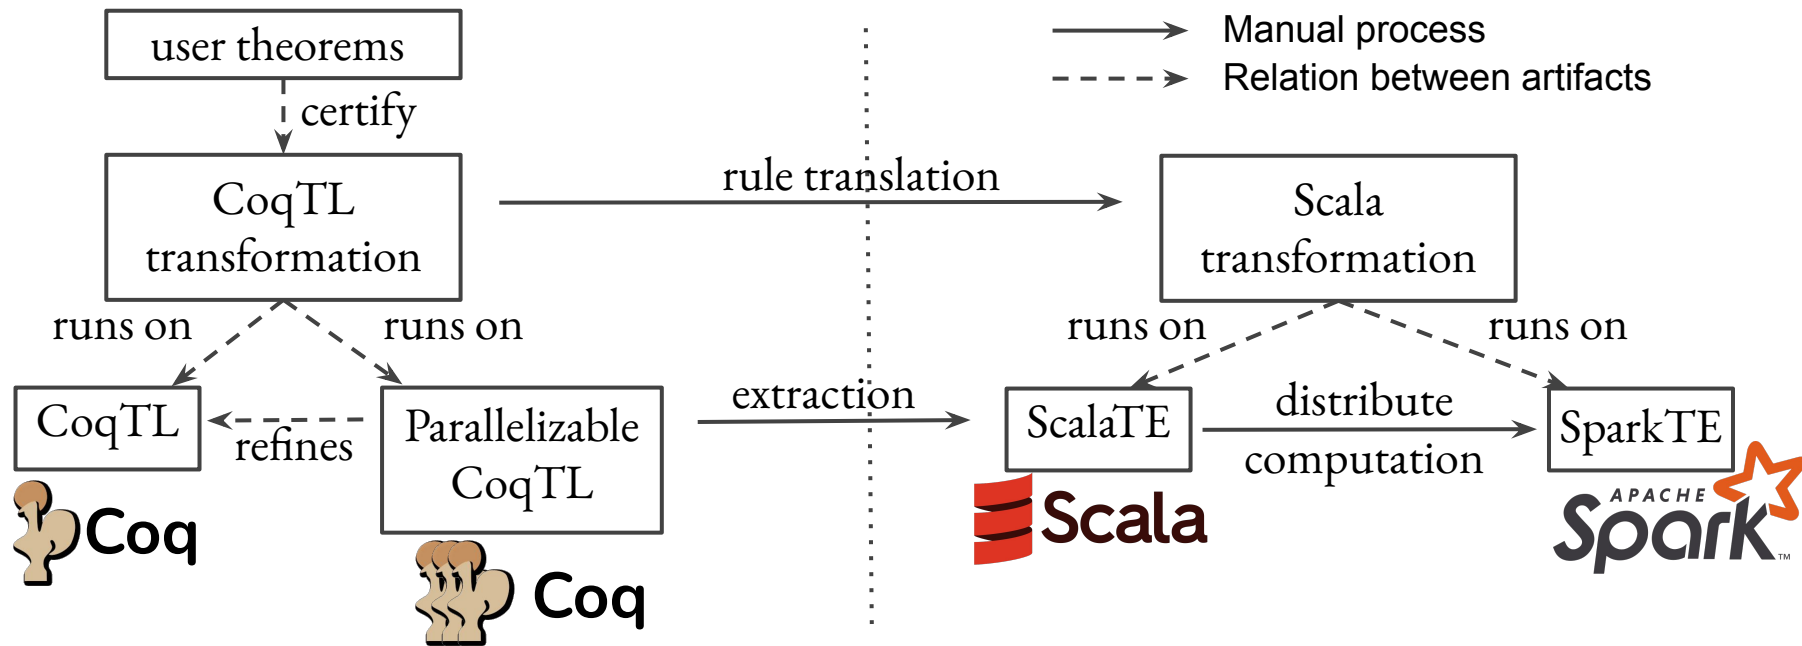


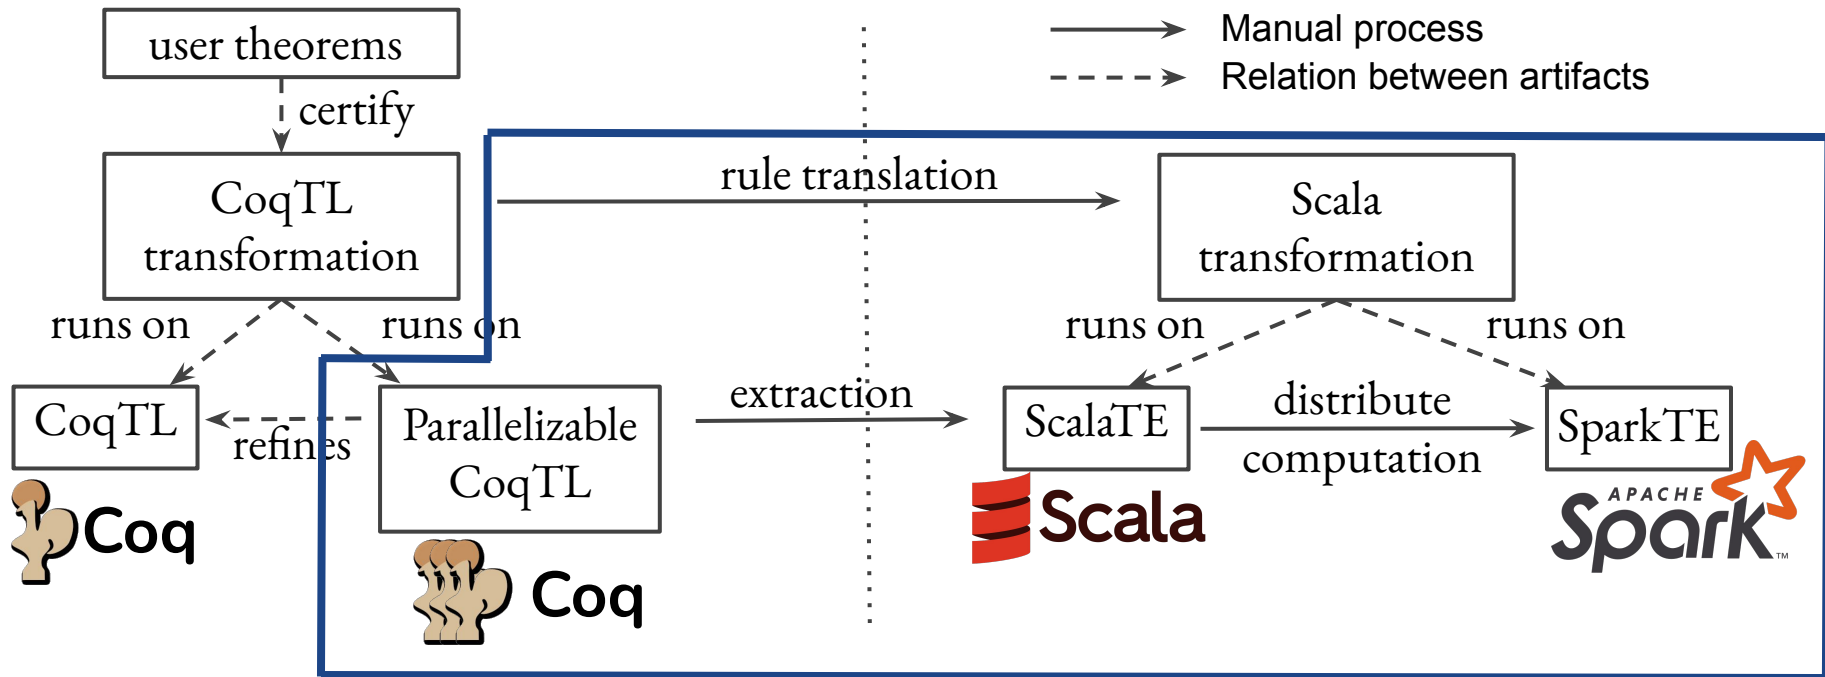


- French cluster for experimentation
- Library for benchmarking
- Support for distributed computing
- More than 15,000 cores; 800 nodes




Engine based on a formal semantic: from CoqTL to SparkTE






Contribution


- **Increase parallelization**
 1. Two distinct phases: **instantiate & apply**
 - Define map-reduce phases
 2. Iterate on rules instead of src patterns
 - Avoid unnecessary computations
 3. Iterate on trace links instead of src patterns
 - Reuse of intermediate results
- **Formal proof of equivalence with CoqTL**



- **Increase parallelization**
 1. Two distinct phases: **instantiate & apply**
 - Define map-reduce phases
 2. Iterate on rules instead of src patterns
 - Avoid unnecessary computations
 3. Iterate on trace links instead of src patterns
 - Reuse of intermediate results
- **Formal proof of equivalence with CoqTL**

	Spec. size (LoC)	Cert. size (LoC)	Proof effort (man-days)
1.	69	484	10
2.	42	487	7
3.	69	520	4







1. Produce executable and maintainable code

- Object-oriented approach
- Pure Scala functions (correctness)

2. Distribute the computation

- Distribute data-structures
- Explicit communication operations
 - Take advantage of scatter/gather operations
 - Broadcast global knowledge



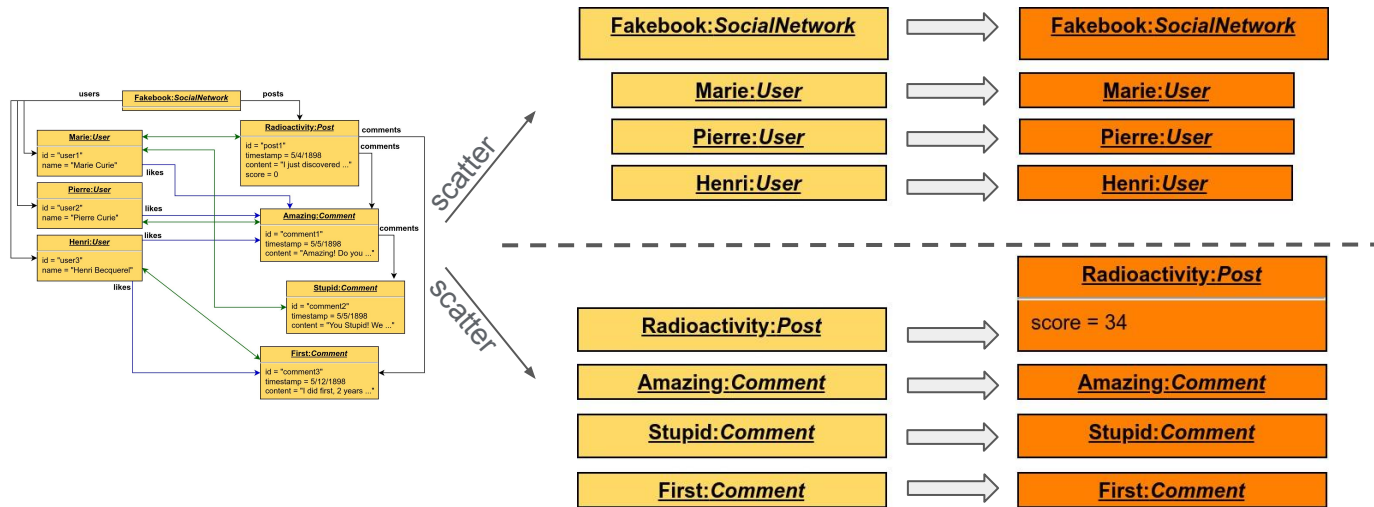


Data-distributed strategy: (*Map-Reduce* phase)

- Input **elements** are **distributed**
- Input **model** is **broadcasted**

As output:

- Instantiated **output model elements**
- **Trace-links** (mapping input-output)

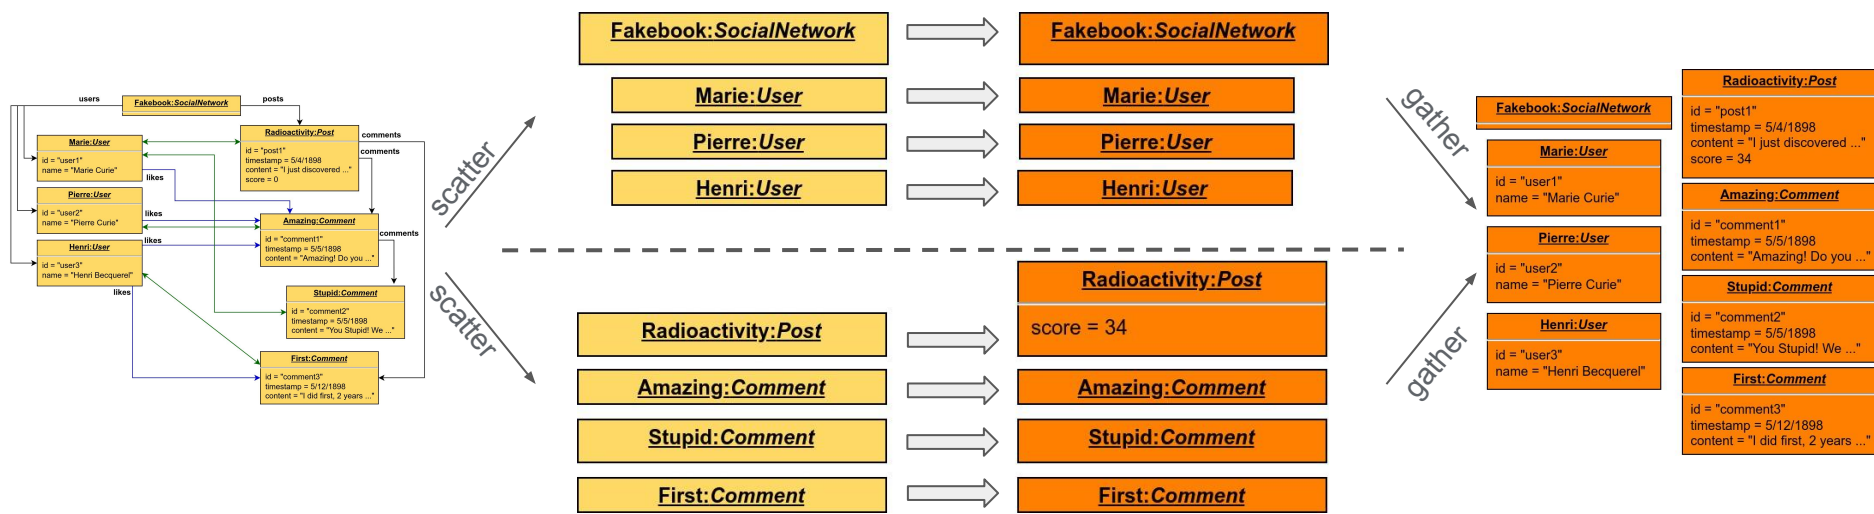


Data-distributed strategy: (*Map-Reduce* phase)

- Input **elements** are **distributed**
- Input **model** is **broadcasted**

As output:

- Instantiated **output model elements**
- **Trace-links** (mapping input-output)




Data-distributed strategy: (Map-Reduce phase)

- Input elements are distributed
- Input model is broadcasted

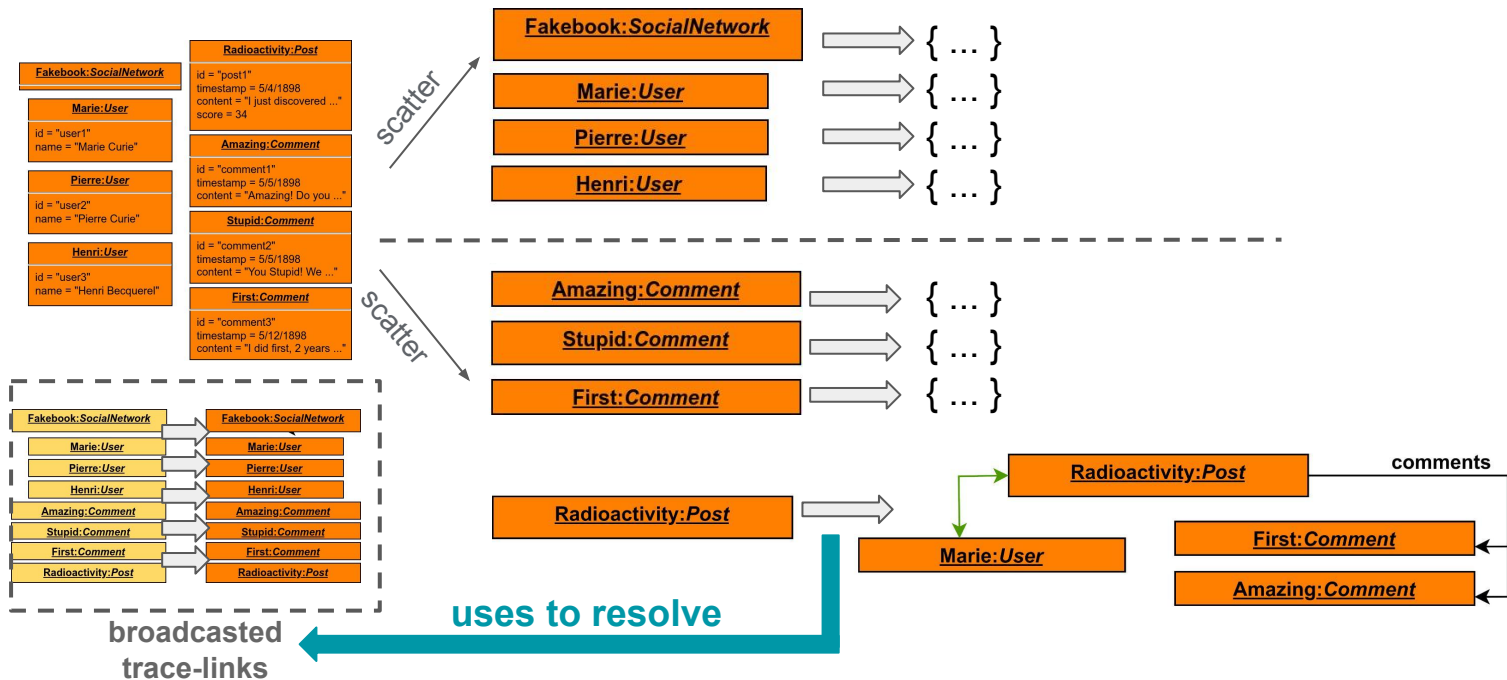
As output:

- Instantiated output model elements
- Trace-links (mapping input-output)



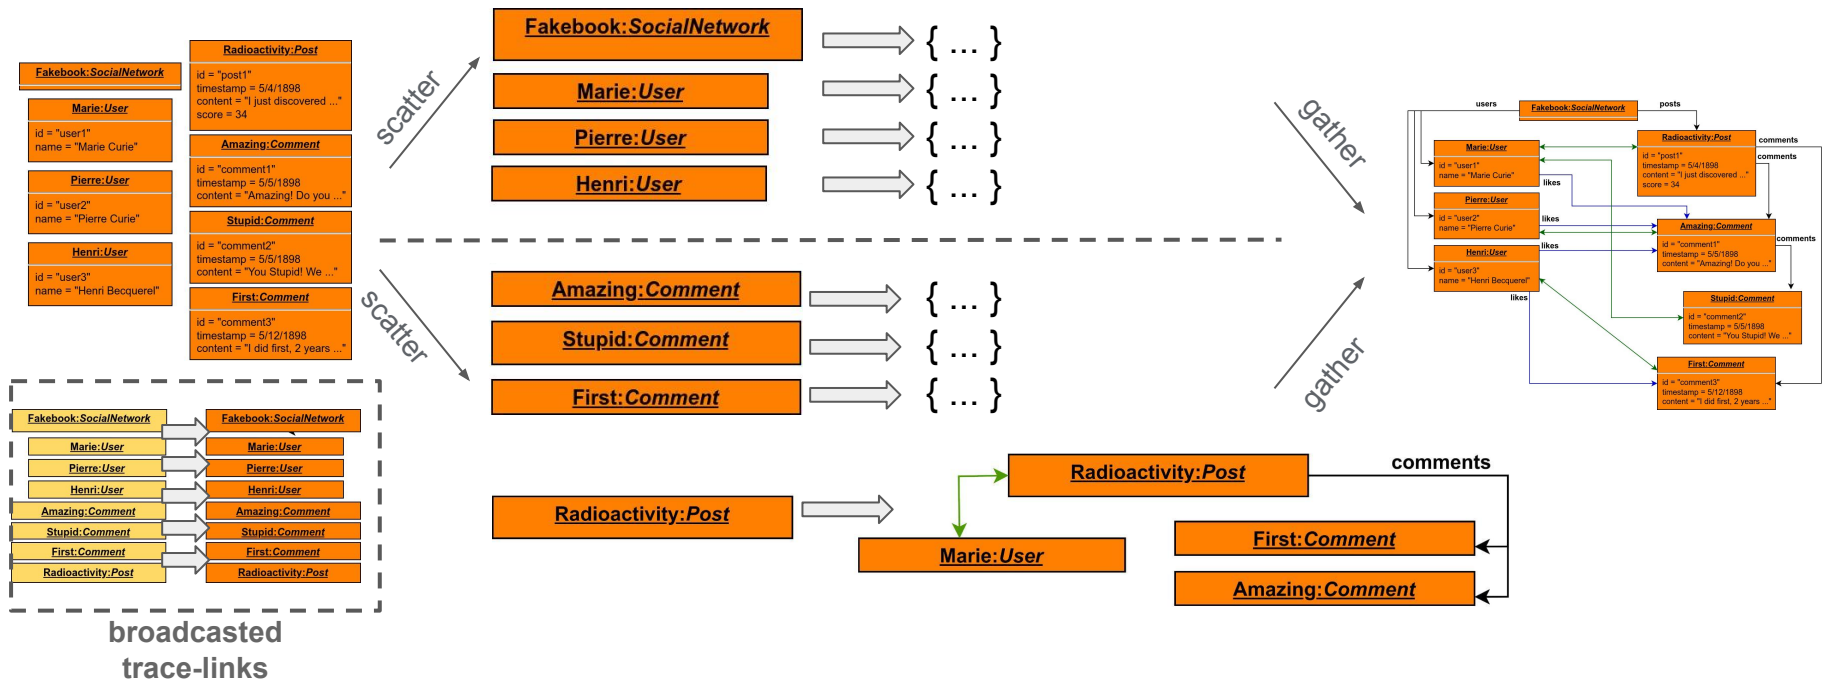
Data-distributed strategy: (*Map-Reduce* phase)

- Output **elements** are **distributed**
- Trace-links are **broadcasted**



Data-distributed strategy: (*Map-Reduce* phase)

- Output **elements** are **distributed**
- Trace-links are **broadcasted**





Data-distributed strategy: (Map-Reduce phase)


- Output **elements** are **distributed**
- Trace-links are **broadcasted**

Vertical scalability of model transformation on Spark

- Simulate a uniform amount of computation on nodes
 - fixed time for each task



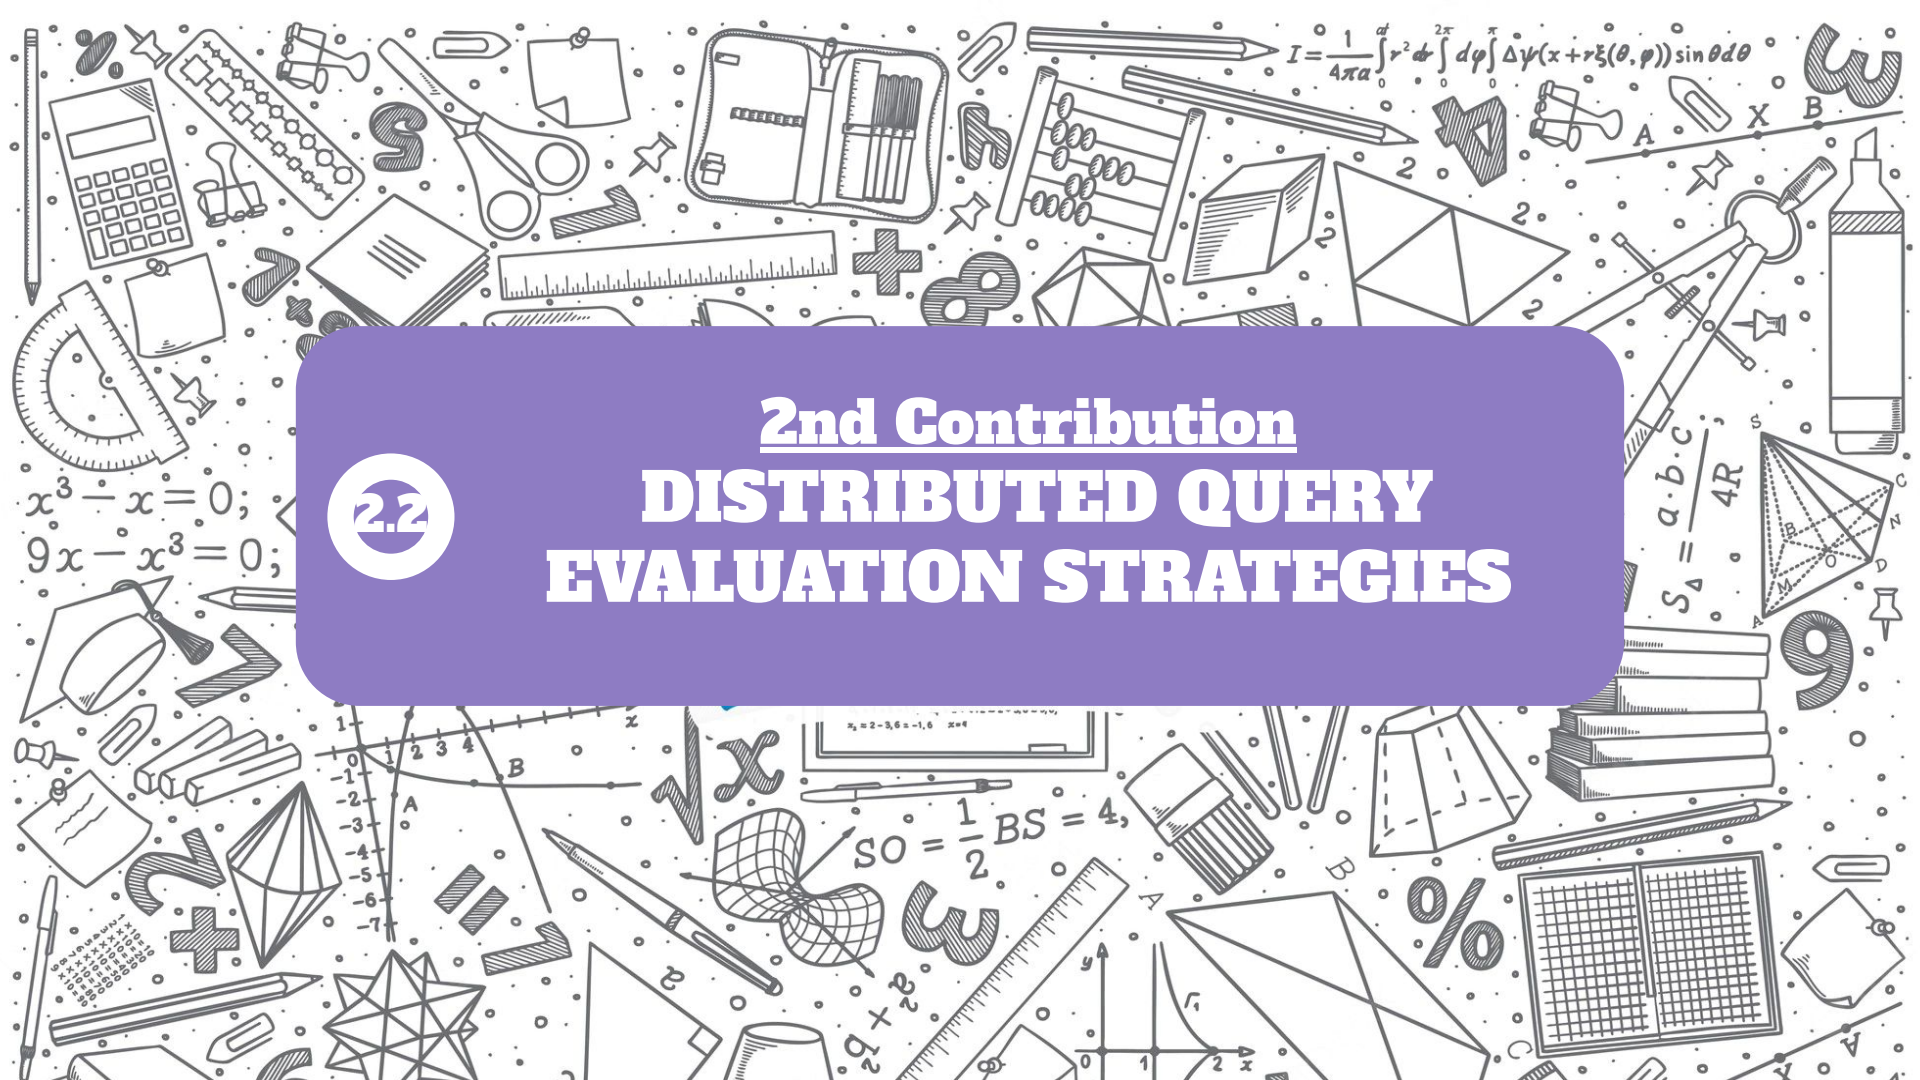
Model of 150 elements and 290 links, on 4 machines



Model of 600 elements and 1060 links, 8 machines

2.2

2nd Contribution DISTRIBUTED QUERY EVALUATION STRATEGIES



Many solutions for executing queries distributively

- Evaluation of distributed design choices for **query execution**
 - Take a query whose evaluation is dependant from input model
 - Implement with several design choices
 - Evaluate them and try to correlate with input

- Query:
What is the score for a post in a social network?
- A score function


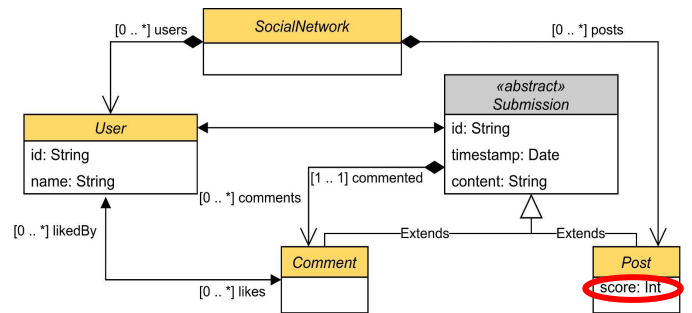
$$\text{score}(p: \text{Post}) := \# \text{ comments} \times 10 + \# \text{ likes}$$

```

score(p: Post) := comments(p).size() * 10
                + likes(p).size()

comments(s: Submission) := [s.comments].
union(c: s.comments.flatMap(
    λc.comments(c))

likes(p: Post) := comments(p).map(λc.likes)
    
```



```
score(p: Post) :=  
  comments(p).size() * 10  
  + likes(p).size()  
  
comments(s: Submission) :=  
  [s.comments].union(  
    c: s.comments.flatMap(  
      λc.comments(c))  
  )  
  
likes(p: Post) :=  
  comments(p).map(λc.likes)
```

- Design-choices for running the query:
 1. **Scala-OCL**
 - No distribution (sequential)
 2. **Spark-OCL** (Spark core API)
 - Delegate distribution to Spark
 3. **MapReduce** (Spark core API)
 - More control of parallelism
 4. **Pregel** from (GraphX)
 - Iterative process
 5. Hybrid approaches
 - **Spark-OCL + Pregel**
 - **MapReduce + Pregel**

- Proposed models from TTC
- Calculate score value
- Cannot really extract relevant metrics about topology

#	Dataset				Speed-up (compared to Sequential Scala-OCL)					
	# users	# posts	# comments	# likes	Scala-OCL	Spark-OCL	Pregel	MapReduce	Spark-OCL + Pregel	MapReduce + Pregel
1	889	1064	118	24	1x	0.39x	0.36x	0.46x	0.44x	0.46x
2	1845	2315	190	66	1x	0.51x	0.68x	0.85x	0.66x	0.71x
3	2270	5056	204	129	1x	0.27x	0.35x	2.34x	0.15x	2.96x
4	5518	9220	394	572	1x	4.25x	5.21x	4.17x	4.68x	4.03x
5	10929	18872	595	1598	1x	4.68x	2.83x	2.39x	1.97x	3.91x
6	18083	39212	781	4770	1x	4.07x	4.12x	4.58x	5.17x	3.27x

Correlation matrix: *input model vs. speed-ups*

Size	Spark- OCL	Pregel	MapReduce	Spark-OCL + Pregel	MapReduce + Pregel
# users	0.78	0.67	0.74	0.76	0.39
# posts	0.71	0.62	0.75	0.75	0.32
# comments	0.86	0.74	0.78	0.79	0.51
# likes	0.62	0.57	0.7	0.73	0.19

Correlation matrix: *input model vs. speed-ups*


Size	Spark- OCL	Pregel	MapReduce	Spark-OCL + Pregel	MapReduce + Pregel
# users	0.78	0.67	0.74	0.76	0.39
# posts	0.71	0.62	0.75	0.75	0.32
# comments	0.86	0.74	0.78	0.79	0.51
# likes	0.62	0.57	0.7	0.73	0.19


Correlation matrix: *ratio in input model vs speed-ups*

	Spark-OCL	Pregel	MapReduce	Spark-OCL + Pregel	MapReduce + Pregel
ratio: #users / #likes	-0.85	-0.79	-0.89	-0.75	-0.82
ratio: #posts / #likes	-0.96	-0.88	-0.82	-0.85	-0.66
ratio: #comments / #likes	-0.8	-0.74	-0.86	-0.69	-0.83


Lack of unified proposition for comparing design choices

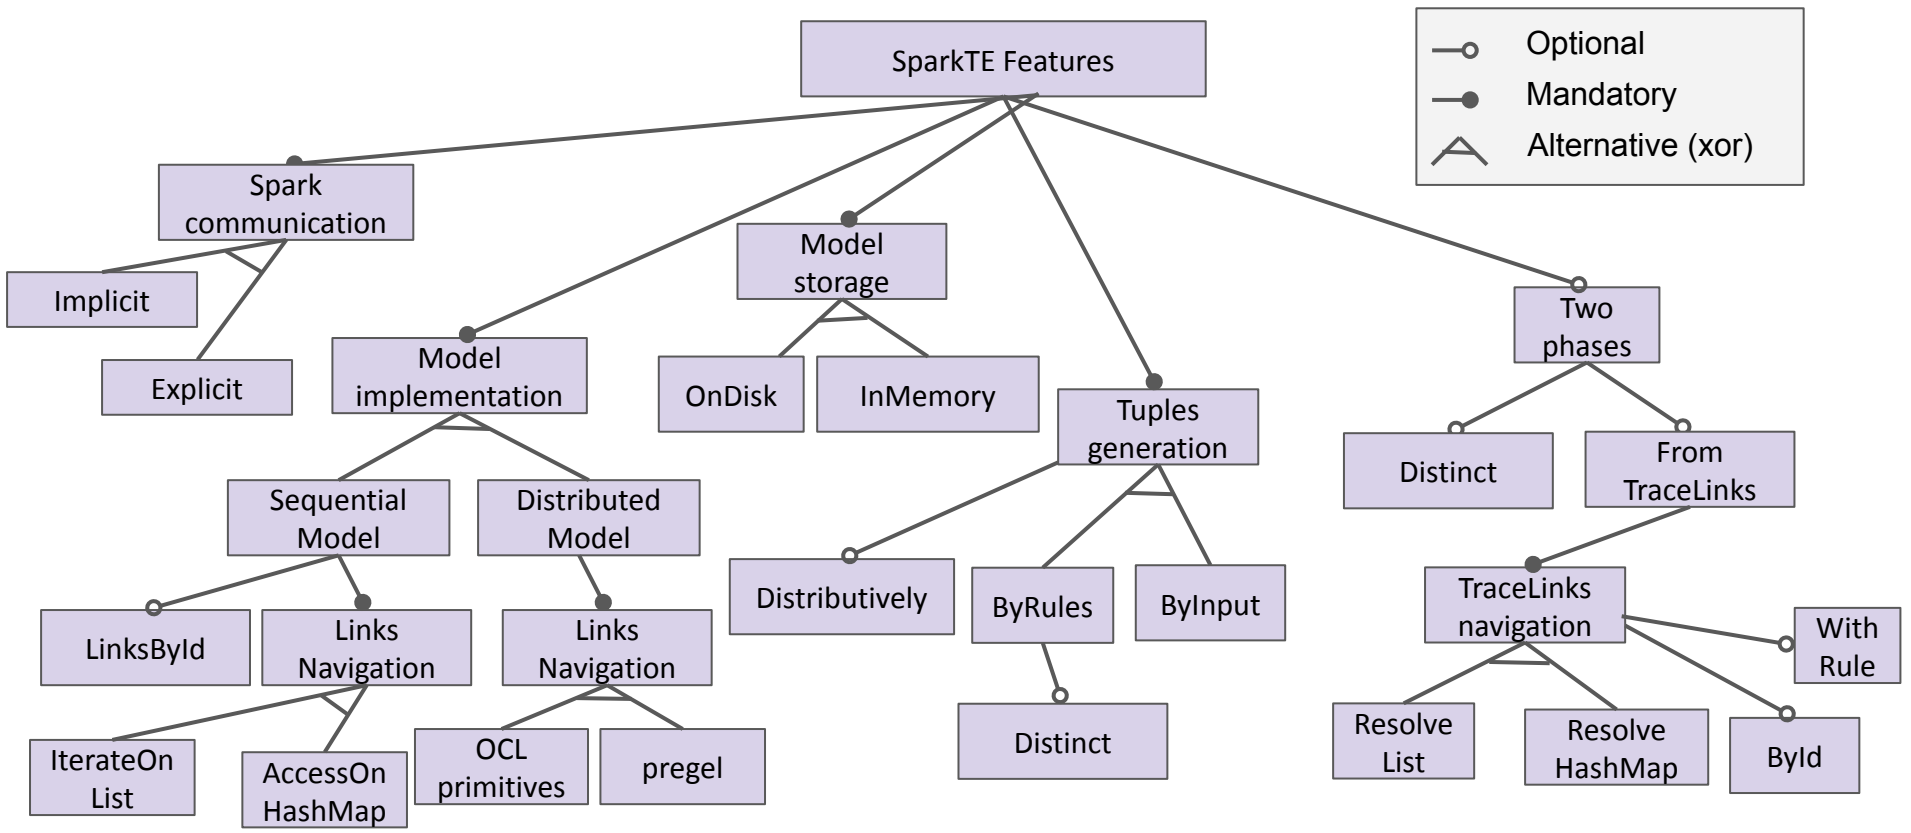
- Make possible configurable distributed transformation
 - Formalized past contributions and additional design choices
 - Design a configurable engine
 - Evaluate them and analyse impact

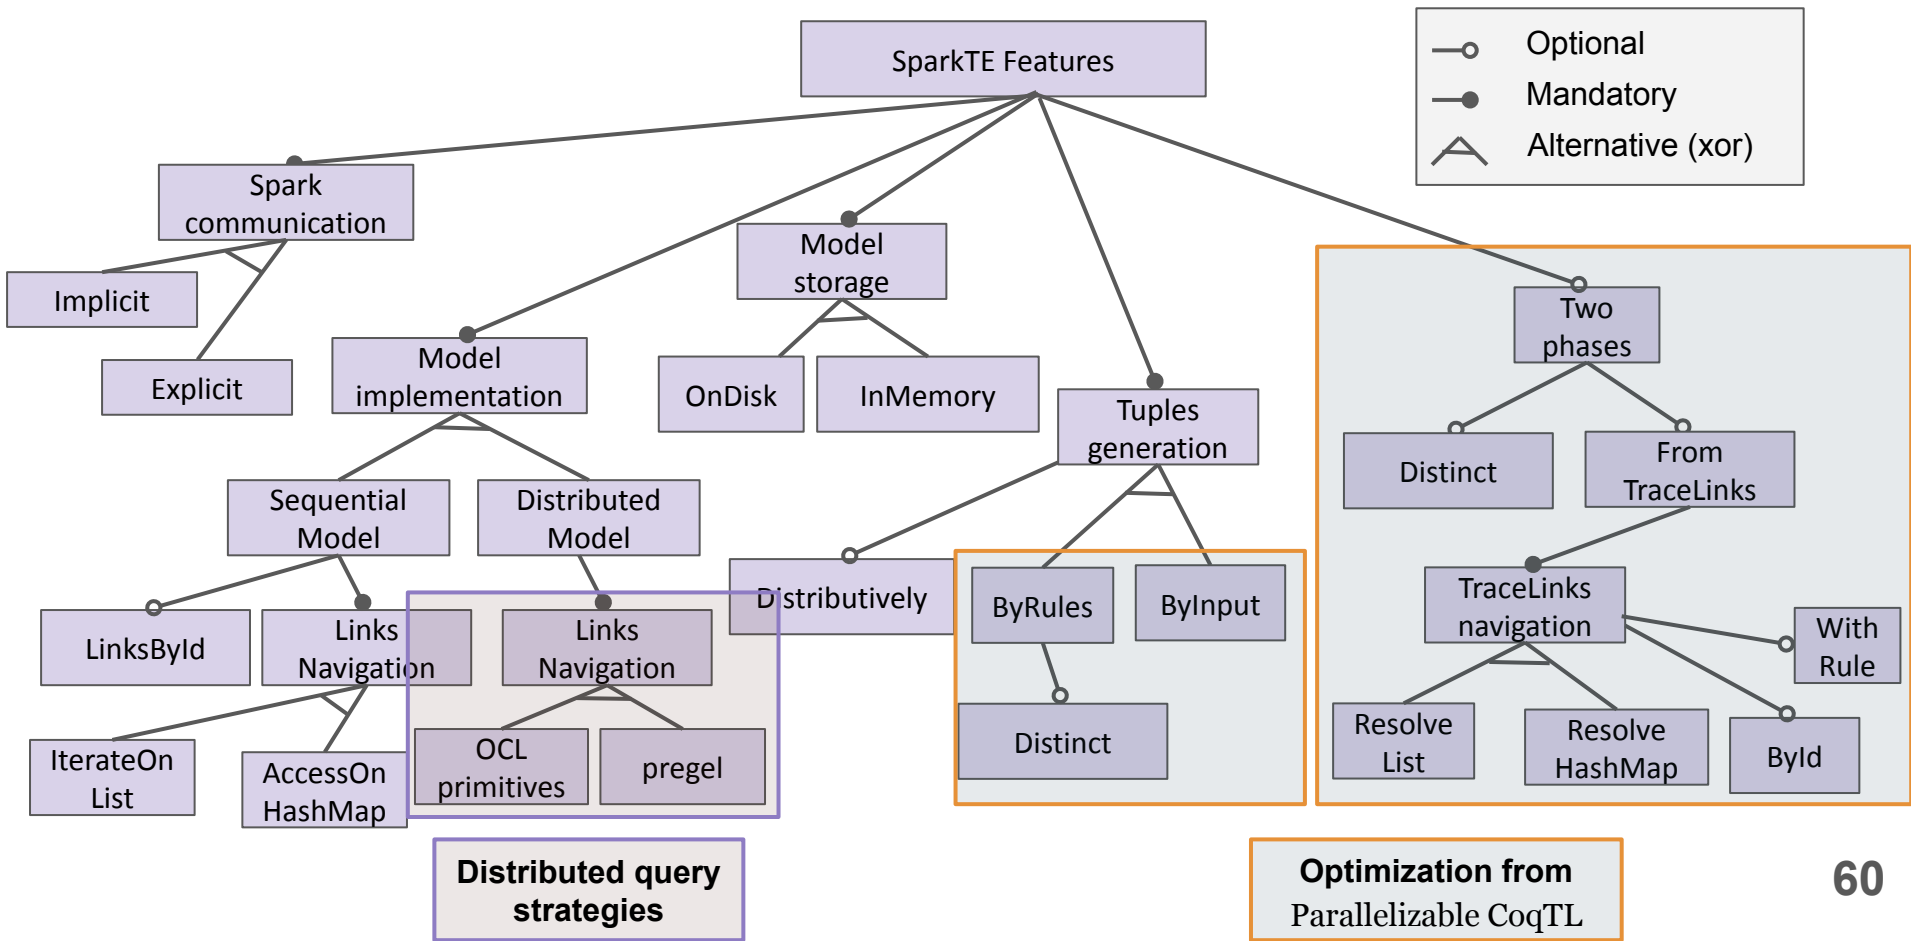


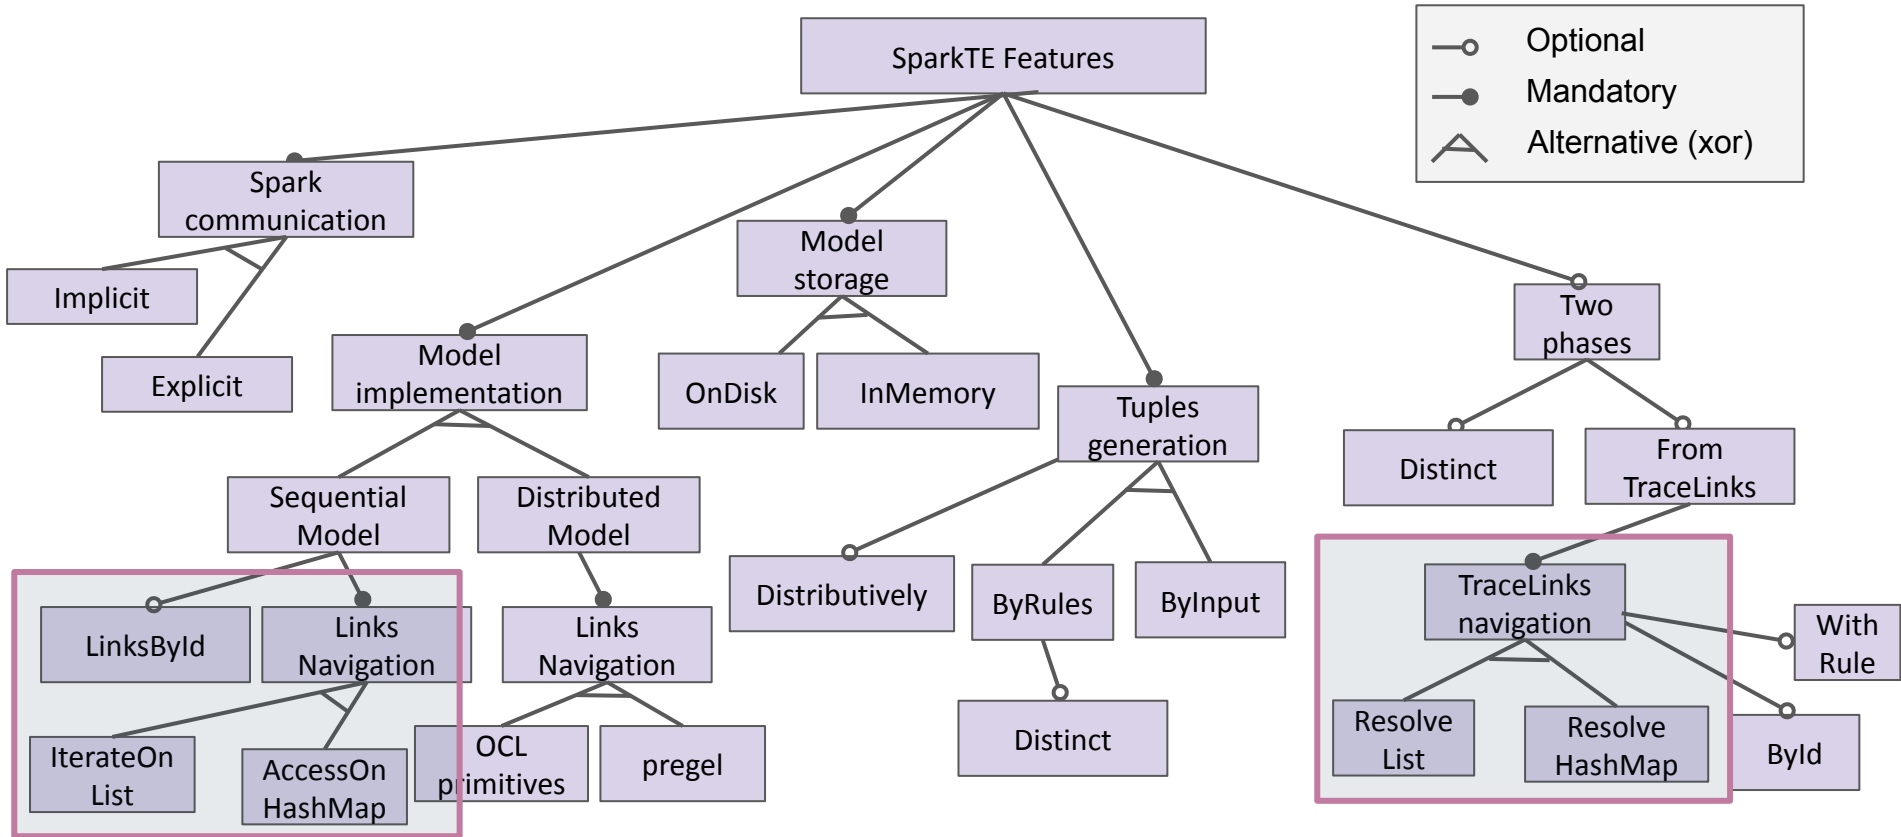


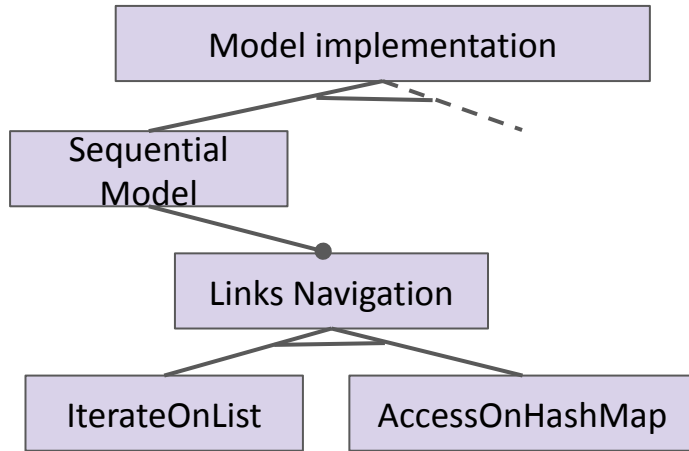
- Take **as input** a configuration conforms to the feature model
- Produce **as output** performance results (computation time)



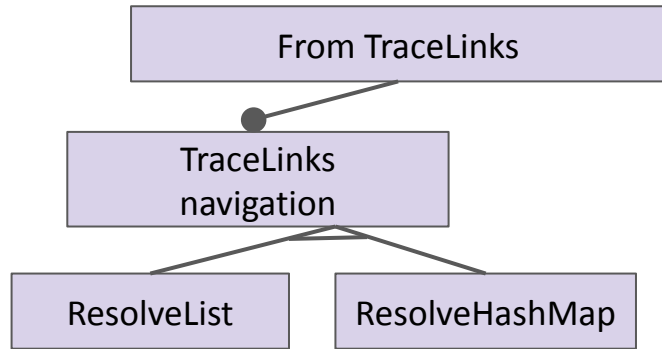








- IterateOnList :
 - Navigation by iteration
 - Simple to set-up
- AccessOnHashMap :
 - Additional computation in model loading
 - Increase memory usage
 - Direct access on links from elements



- **ResolveList** :
 - Resolution by iteration
 - Naturally gathered by master node
- **ResolveHashMap** :
 - Additional computation in instantiate phase
 - Increase memory usage
 - Fastest resolution

Using configurable engine to find features synergie

Execution of **Identity** transformation on a model of 100k elements and 250k links (4 cores)

Configuration 1: Links navigation	Configuration 2: TraceLinks navigation	Computation time (sec)	Instantiate phase (sec)	Apply phase (sec)
IterateOnList	ResolveList	1636 sec	3 sec	1633 sec
IterateOnList	ResolveHashMap	1584 sec	3 sec	1581 sec
AccessOnHashMap	ResolveList	233 sec	6 sec	227 sec
AccessOnHashMap	ResolveHashMap	12 sec	6 sec	6 sec

- **TraceLinks navigation**'s impact
 - on the **whole** computation is **negligible**
 - is **important** when **Links navigation** is processed by **AccessOnHashMap**
- **Links navigation**'s impact
 - decreases the **whole** computation time
 - increases the computation time of the instantiate phase

Design-space exploration for the Find affinity case

Feature label	Parallelizable CoqTL design choices (C1)	Optimal design choices (C2)
Model implementation	Sequential Model	Sequential Model
<ul style="list-style-type: none">linksById	false	false
<ul style="list-style-type: none">Link Navigation	IterateOnList	ResolveHashMap
Model storage	InMemory	InMemory
Spark communication	Implicit	Explicit
Tuples generation	ByRules	ByInput
<ul style="list-style-type: none">Distributively	false	false
<ul style="list-style-type: none">Distinct	false	true
TraceLinks Navigation	ResolveList	ResolveList
<ul style="list-style-type: none">byId	false	false
<ul style="list-style-type: none">withRule	false	true
<ul style="list-style-type: none">Distinct	false	true

Design-space exploration for the Find affinity case

2.3

FEATURE ANALYSIS

Feature label	Parallelizable CoqTL design choices (C1)	Optimal design choices (C2)
Model implementation	Sequential Model	Sequential Model
○ linksById	false	false
● Link Navigation	IterateOnList	ResolveHashMap
Model storage	InMemory	InMemory
Spark communication	Implicit	Explicit
Tuples generation	ByRules	ByInput
○ Distributively	false	false
○ Distinct	false	true
TraceLinks Navigation	ResolveList	ResolveList
○ byId	false	false
○ withRule	false	true
○ Distinct	false	true

#elements	#links	C1 computation time	C2 computation time
1000	3000	9.799 sec	4.978 sec
2500	7300	81.047 sec	7.803 sec
5000	15000	882.708 sec	19.127 sec
7500	22000	> 2h	36.928 sec
10000	45000	Timeout error	65.198 sec

- The feature model is useful for comparing implementations
- Gives useful insights about the engine
- Highlighted correlation between features

Problem 1:

Many solutions for executing rules distributively

Built a **distributed solution** from a **specification**

- Re-designed specification to make it distributable
- Made a proof of equivalence for optimizations
- Shown our solution is scalable

Problem 2:

Many solutions for executing queries distributively

Evaluated distributed execution **strategies** for a query

- Implemented three design-choices
- Proposed hybrid solution
- Performance variation depending on the strategy

Problem 3:

Need an unified proposition for comparing design choices

Formalized features in our distributed solution

- Shown the synergies between them
- Shown the impact on performance

- Jolan Philippe, H el ene Coullon, Massimo Tisi, Gerson Suny e. **Towards Transparent Combination of Model Management Execution Strategies for Low-Code Development Platforms.** *23rd ACM/IEEE International Conference on Model Driven Engineering Languages and Systems (MODELS): Companion Proceedings*, Oct 2020, Montreal (Virtually), Canada. 10.1145/3417990.3420206. Hal-02952952
- Jolan Philippe, Massimo Tisi, H el ene Coullon, Gerson Suny e. **Executing Certified Model Transformations on Apache Spark.** *14th ACM SIGPLAN International Conference on Software Language Engineering (SLE)*, Oct 2021, Chicago IL, United States. 10.1145/3486608.3486901. Hal-03343942
- Ongoing: Jolan Philippe, Massimo Tisi, Gerson Suny e. **Analysis of the Design-Space of a Distributed Transformation Engine.** *Software and Systems Modeling (SoSyM)*
- *Several public Lowcomote deliverables*
 - **Concepts for Multi-paradigm distributed transformation**
 - **Scalable low-code artefact persistence and query**
 - **Multi-paradigm distributed transformation engine**

- **Automated design-space exploration for a given scenario**
 - A model of the input (e.g., topological metrics)
 - A model of the platform (**Spark and ≠**)
 - Constraints and requirements
- **Other parameters to optimize (≠ CPU time)**
 - Network bandwidth
 - Memory consumption
 - Energy consumption/production
- + **Other execution strategies (≠ data-dist)**
 - Take advantage of Spark for task-distribution
 - Combine incrementality and laziness to distribution



Contribution to the Analysis of the Design-Space of a Distributed Transformation Engine

Jolan PHILIPPE

PhD Defense, speciality: Computer Science

