

Trace checking of Metric Temporal Logic with Aggregating Modalities using MapReduce

Srdan Krstić

with

Domenico Bianculli and Carlo Ghezzi



Trace checking of Metric Temporal Logic with Aggregating Modalities using MapReduce

Srdan Krstić

with


Domenico Bianculli and Carlo Ghezzi



Execution traces


The image shows a screenshot of a network execution trace. At the top left, it displays the date and time: "Wednesday, April 22, 2009 18:52:56". Below this, the source IP address "dialup.xtra.co.nz" is visible. The trace consists of numerous lines of text, each representing a network packet. These lines are color-coded: blue for standard data packets, red for error packets, and green for other types of traffic. On the right side of the screenshot, there is a vertical navigation menu with four items: "CSS", "Script", "Images", and "Misc". The "Images" item is currently selected, indicated by a blue highlight. At the bottom of the screenshot, the URL "logstalgia.googlecode.com" is displayed in white text.

Trace checking



```
64.242.88.10 - - [07/Mar/2004:16:05:49 -0800] "GET /twiki/bin/edit/Main/Double_bounce_send
64.242.88.10 - - [07/Mar/2004:16:06:51 -0800] "GET /twiki/bin/rdiff/TWiki/NewUserTemplate?
64.242.88.10 - - [07/Mar/2004:16:10:02 -0800] "GET /mailman/listinfo/hsdivision HTTP/1.1"
64.242.88.10 - - [07/Mar/2004:16:11:58 -0800] "GET /twiki/bin/view/TWiki/WikiSyntax HTTP/1
64.242.88.10 - - [07/Mar/2004:16:20:55 -0800] "GET /twiki/bin/view/Main/DCCAndPostFix HTTP
64.242.88.10 - - [07/Mar/2004:16:23:12 -0800] "GET /twiki/bin/oops/TWiki/AppendixFileSyste
64.242.88.10 - - [07/Mar/2004:16:24:16 -0800] "GET /twiki/bin/view/Main/PeterThoeny HTTP/1
64.242.88.10 - - [07/Mar/2004:16:29:16 -0800] "GET /twiki/bin/edit/Main/Header_checks?topi
64.242.88.10 - - [07/Mar/2004:16:30:29 -0800] "GET /twiki/bin/attach/Main/OfficeLocations
64.242.88.10 - - [07/Mar/2004:16:31:48 -0800] "GET /twiki/bin/view/TWiki/WebTopicEditTemp
64.242.88.10 - - [07/Mar/2004:16:32:50 -0800] "GET /twiki/bin/view/Main/WebChanges HTTP/1.
64.242.88.10 - - [07/Mar/2004:16:33:53 -0800] "GET /twiki/bin/edit/Main/Smtpd_etrn_restric
64.242.88.10 - - [07/Mar/2004:16:35:19 -0800] "GET /mailman/listinfo/business HTTP/1.1" 20
64.242.88.10 - - [07/Mar/2004:16:36:22 -0800] "GET /twiki/bin/rdiff/Main/WebIndex?rev1=1.2
64.242.88.10 - - [07/Mar/2004:16:37:27 -0800] "GET /twiki/bin/view/TWiki/DontNotify HTTP/1
64.242.88.10 - - [07/Mar/2004:16:40:24 -0800] "GET /twiki/bin/view/Main/TokyoOffice HTTP/1
64.242.88.10 - - [07/Mar/2004:16:43:54 -0800] "GET /twiki/bin/view/Main/MikeMannix HTTP/1.
64.242.88.10 - - [07/Mar/2004:16:45:56 -0800] "GET /twiki/bin/attach/Main/PostfixCommands
64.242.88.10 - - [07/Mar/2004:16:47:12 -0800] "GET /twiki/bin/view/Main/PostfixCommands.txt HTTP/1.1" 200 68
64.242.88.10 - - [07/Mar/2004:16:47:46 -0800] "GET /twiki/bin/rdiff/Know/ReadmeFirst?rev1=
64.242.88.10 - - [07/Mar/2004:16:49:04 -0800] "GET /twiki/bin/view/Main/TWikiGroups?rev=1.
64.242.88.10 - - [07/Mar/2004:16:50:54 -0800] "GET /twiki/bin/rdiff/Main/ConfigurationVari
64.242.88.10 - - [07/Mar/2004:16:52:35 -0800] "GET /twiki/bin/edit/Main/Flush_service_name
64.242.88.10 - - [07/Mar/2004:16:53:46 -0800] "GET /twiki/bin/rdiff/TWiki/TWikiRegistratio
64.242.88.10 - - [07/Mar/2004:16:54:55 -0800] "GET /twiki/bin/rdiff/Main/NicholasLee HTTP/
64.242.88.10 - - [07/Mar/2004:16:56:39 -0800] "GET /twiki/bin/view/Main/Sandbox/WebHome?rev=1.6
64.242.88.10 - - [07/Mar/2004:16:58:54 -0800] "GET /mailman/listinfo/administration HTTP/1
64.242.88.10 - - [07/Mar/2004:17:09:01 -0800] "GET /twiki/bin/view/Main/SearchResult?sc
64.242.88.10 - - [07/Mar/2004:17:10:20 -0800] "GET /twiki/bin/rdiff/TWiki/TextFormattingRu
64.242.88.10 - - [07/Mar/2004:17:13:50 -0800] "GET /twiki/bin/edit/Main/DefaultPlugin?t=1
64.242.88.10 - - [07/Mar/2004:17:16:00 -0800] "GET /twiki/bin/search/Main/?scope=topic&re
64.242.88.10 - - [07/Mar/2004:17:17:27 -0800] "GET /twiki/bin/search/TWiki/?scope=topic&re
```

Large execution traces



Complex Specifications

Complex Specifications



Timing information

Complex Specifications



Timing information

$$f_n : \{|Q|\}_n \rightarrow Q$$
$$n \in \mathbb{N} \cup \infty$$

Aggregate operations

Specification

Specification Language for service compositions interactions

SOLOIST

$\neg p$

propositional

$pU_{[0,+\infty)}q$

metric temporal logic

$e_{<n}^K(\phi)$

additional aggregating modalities

Service Provisioning Patterns

625 Industrial requirements
specification



290 Academic requirements
specification



Service Provisioning Patterns

Service Provisioning Patterns

Counting the number of events

Service Provisioning Patterns

Counting the number of events

Average number of events

Service Provisioning Patterns

Counting the number of events

Average number of events

Maximum number of events

Service Provisioning Patterns

Counting the number of events

Average number of events

Maximum number of events

Average response time

Service Provisioning Patterns

Counting the number of events

$$e_{\bowtie n}^K(\phi)$$

Average number of events

$$\mathcal{U}_{\bowtie n}^{K,h}(\phi)$$

Maximum number of events

$$\mathcal{M}_{\bowtie n}^{K,h}(\phi)$$


Average response time

$$\mathcal{D}_{\bowtie n}^K(\phi, \psi)$$


Counting Modality

$$C_n^K(\phi)$$


Counting Modality Semantics




Counting Modality Semantics



Counting Modality Semantics



Counting Modality Semantics



$$4 < 7$$

Counting Modality Example


“If operation **A** or operation **B** executed **exactly 4 times** in the **past 5 minutes**, than operations **C** and **D** must be executed simultaneously **within 32 seconds**.”

$$G(\mathcal{C}_{=4}^{300}(A \vee B) \rightarrow F_{[0,32]}(C \wedge D))$$

Terminology

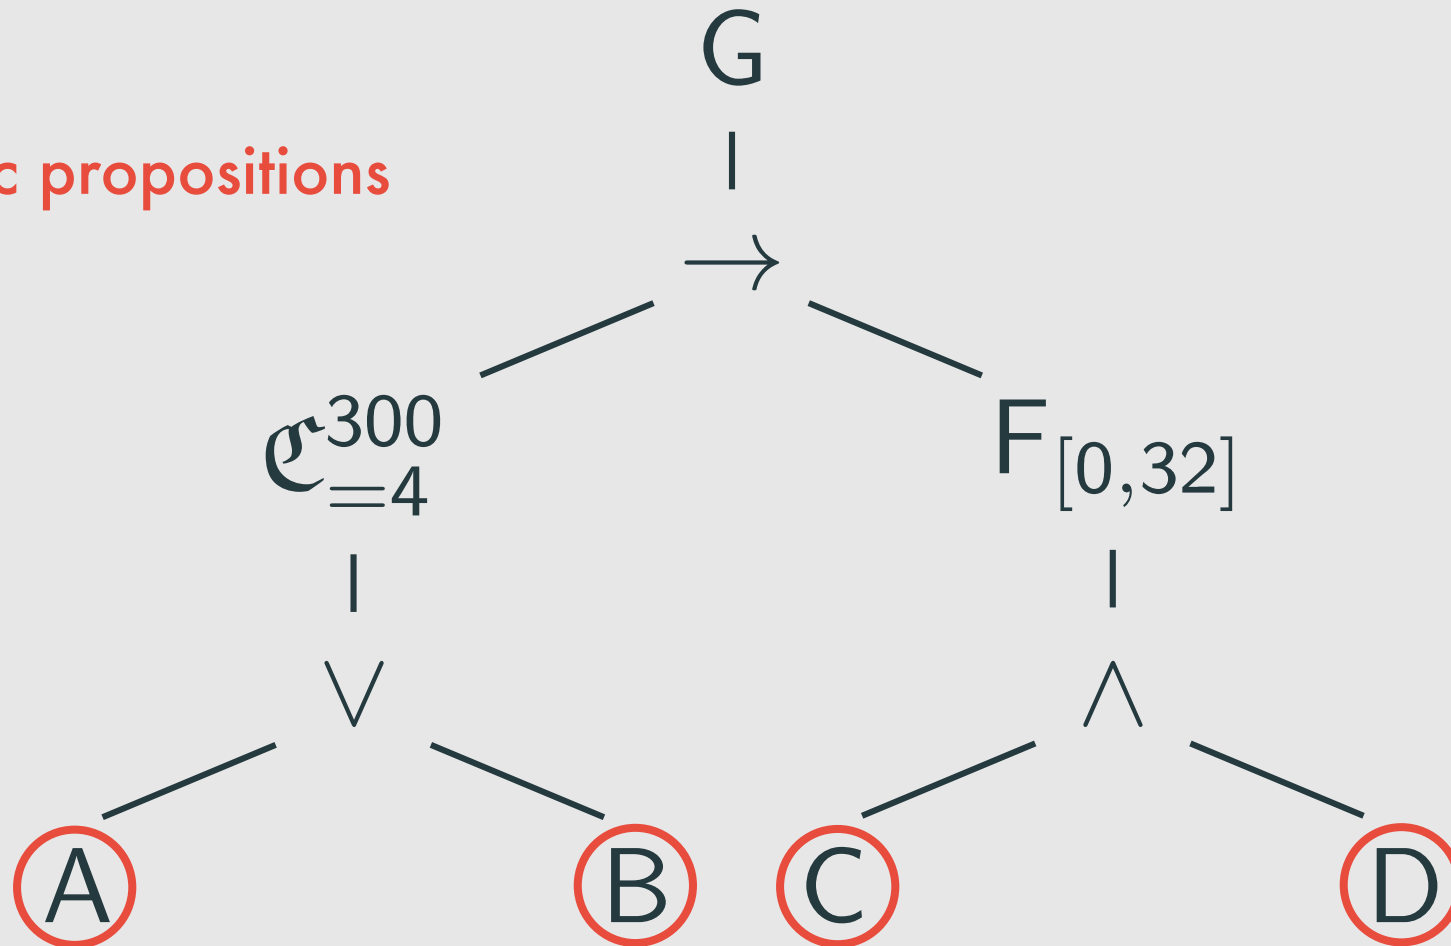
Terminology

Syntax tree

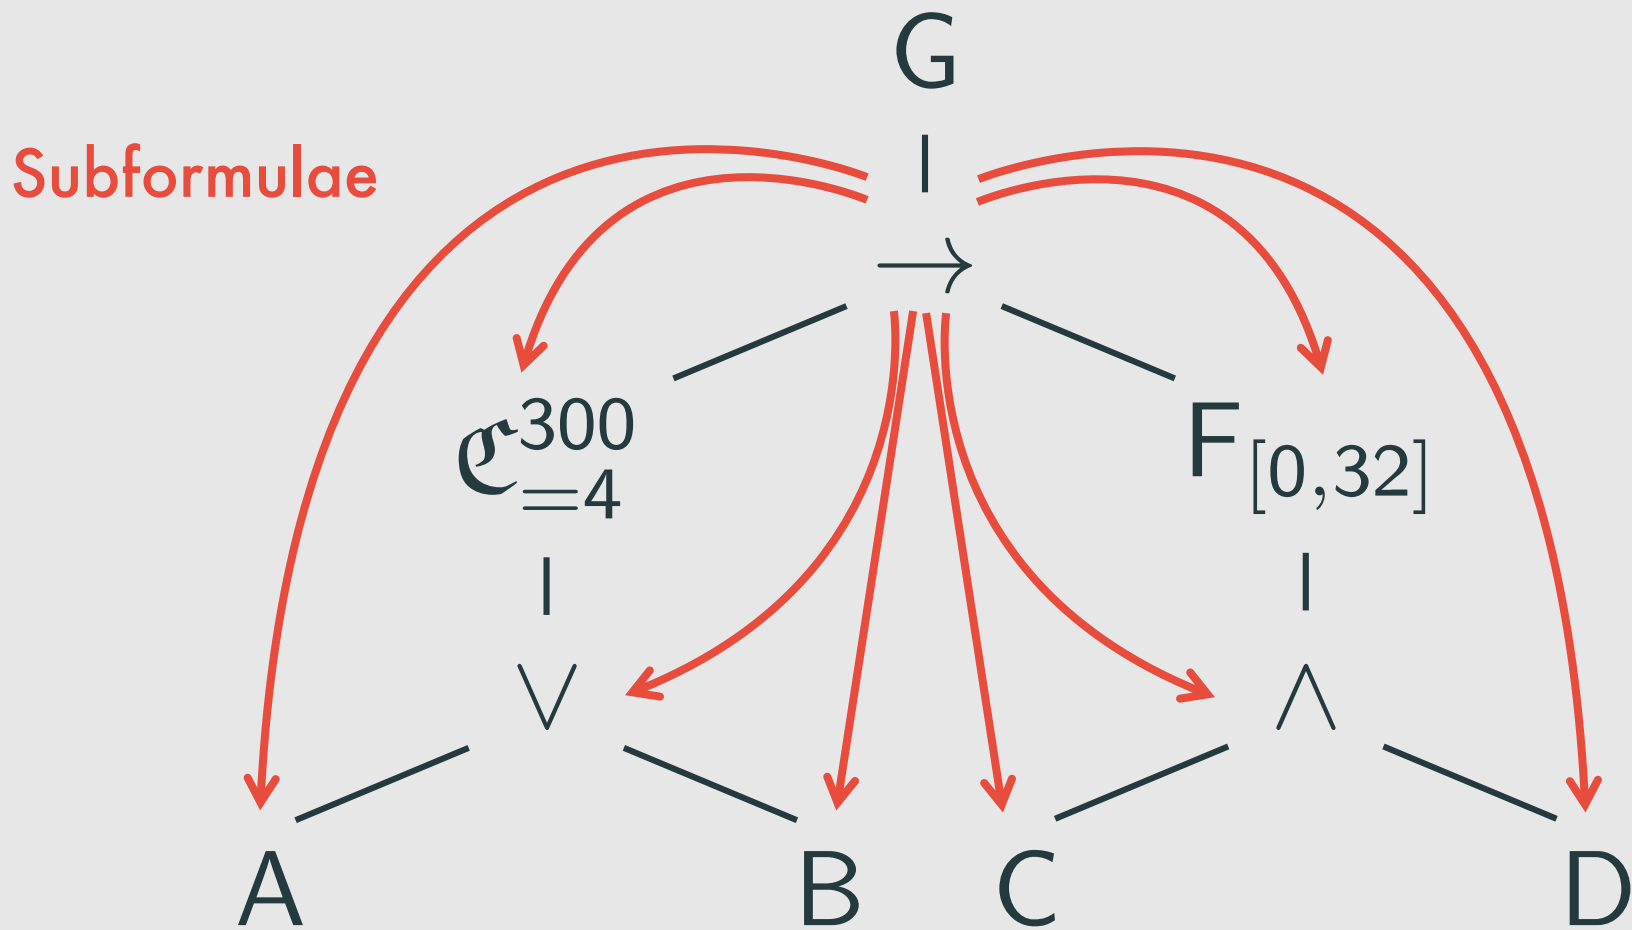


Terminology


Atomic propositions



Terminology




Terminology




Terminology

Direct superformula




Terminology



MapReduce

Sequential Juice Algorithm




Benefits




- A lot of fruit
- Different types of fruit
- Create different juice

Abstract cutting



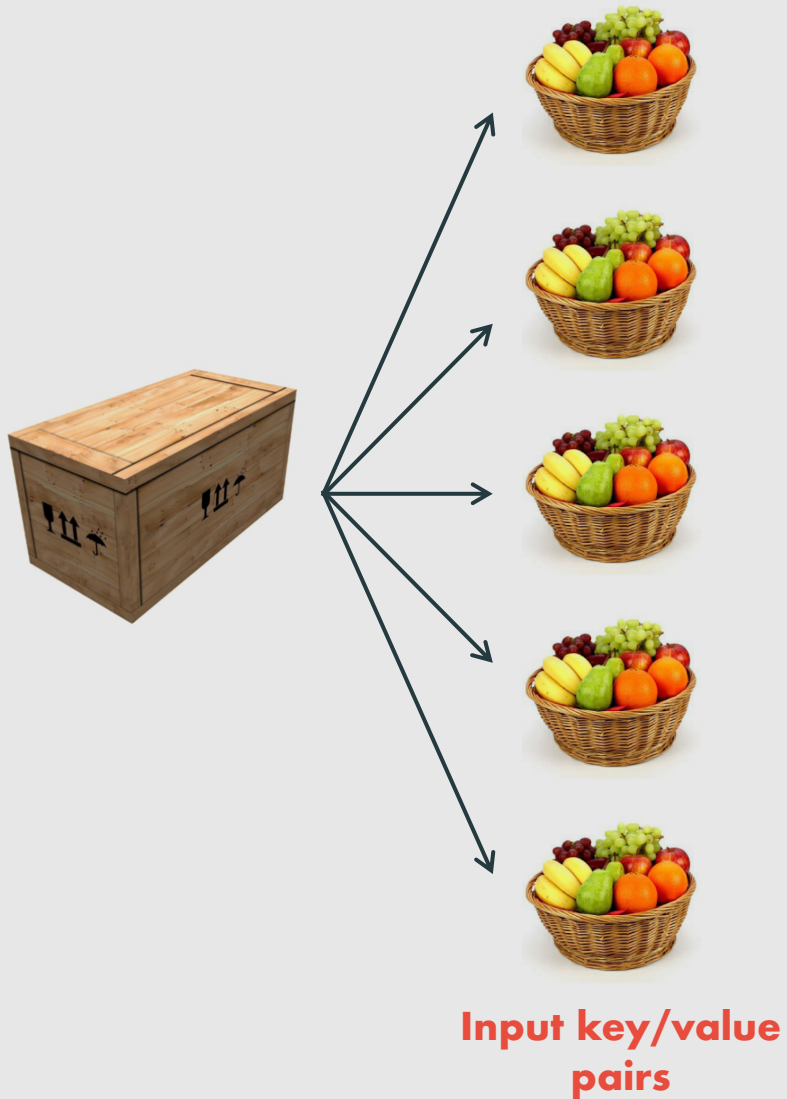
Abstract blending



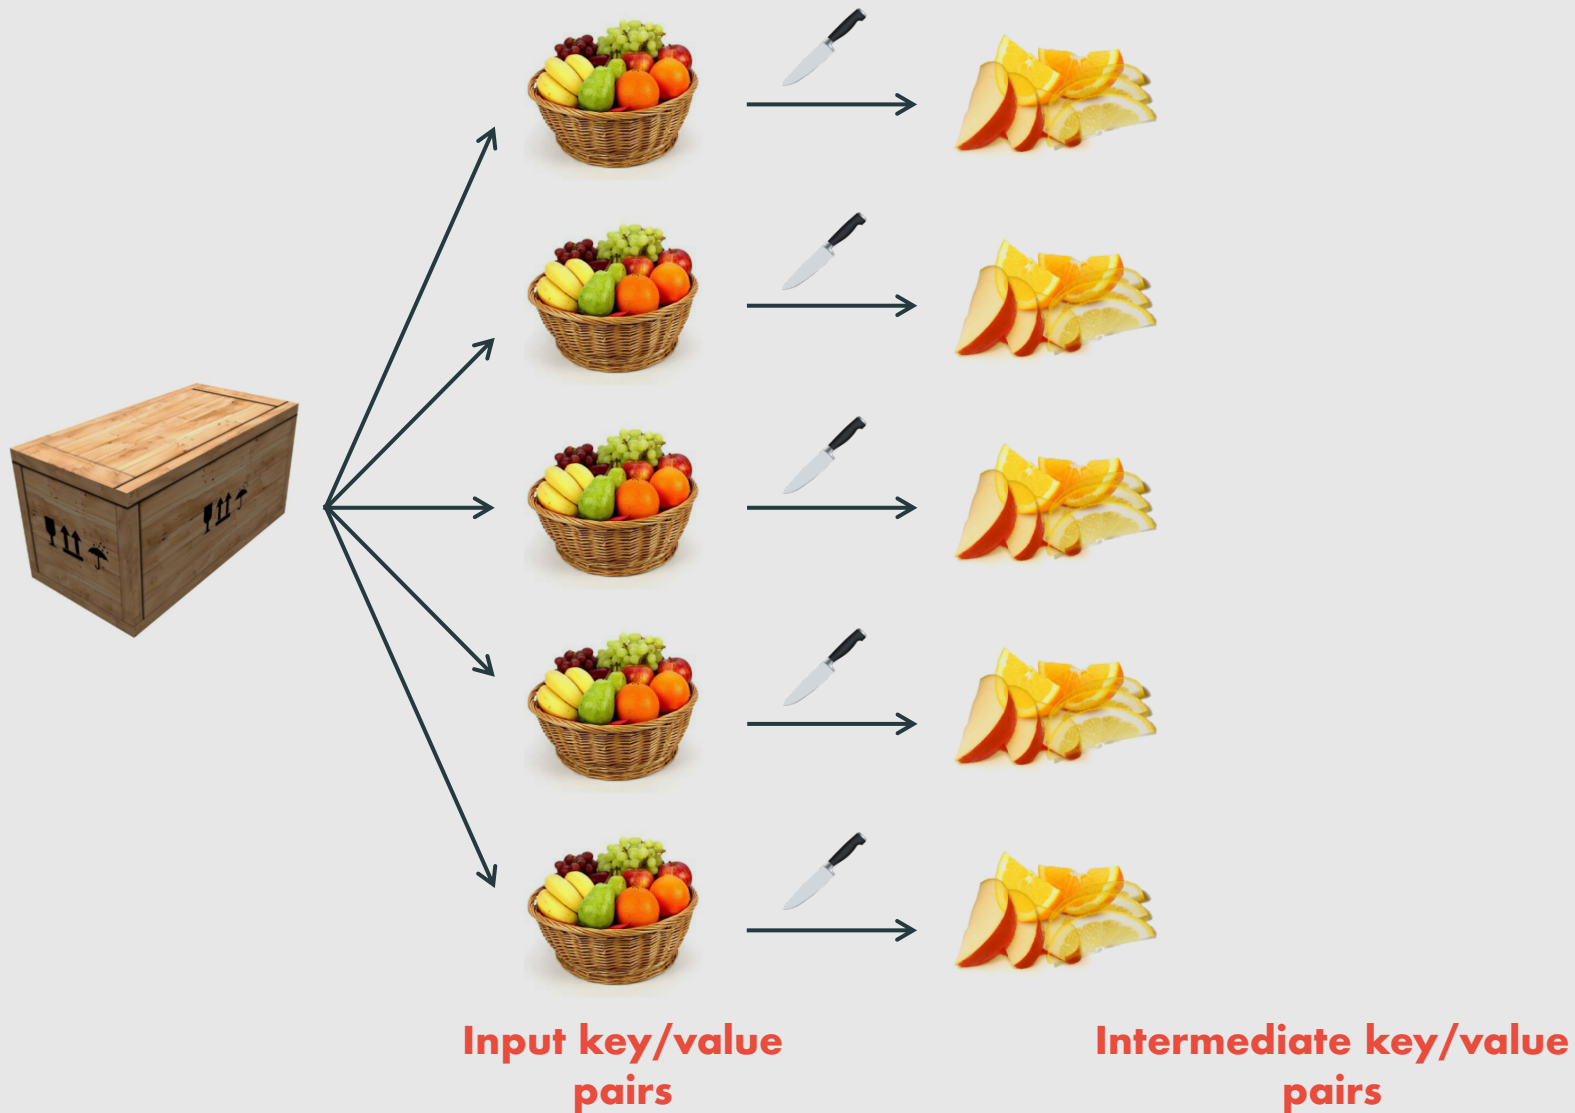
Parallel Juice Algorithm



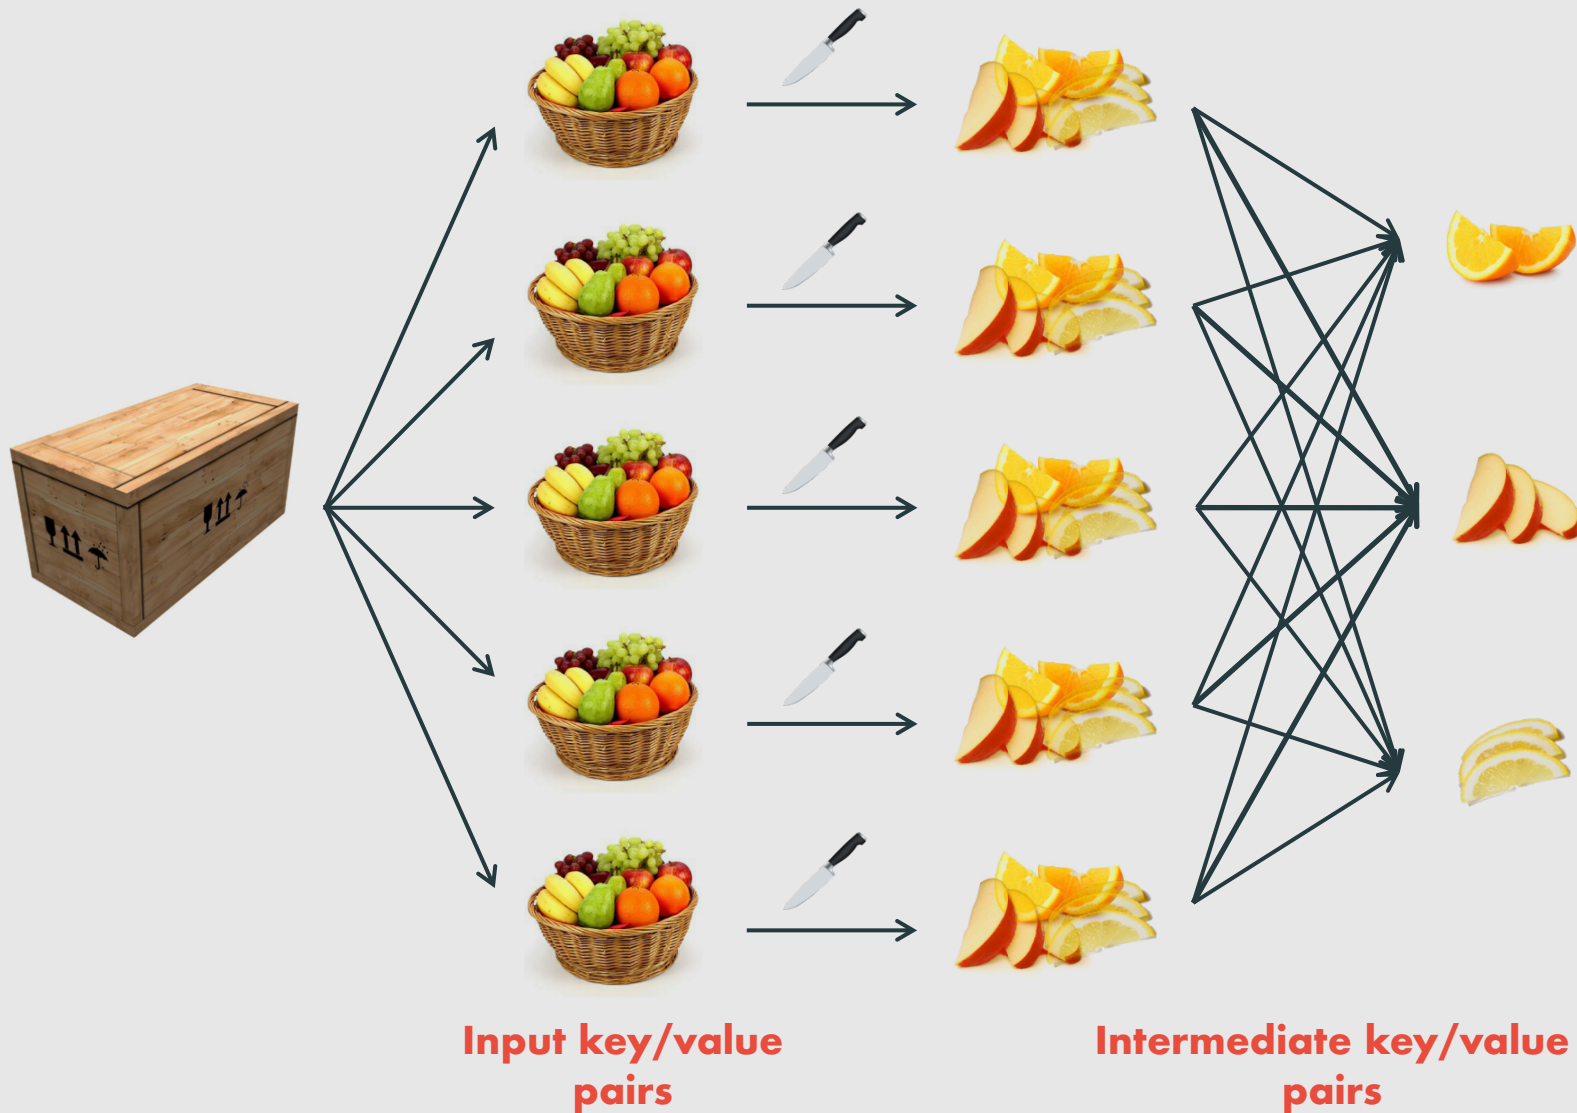
Parallel Juice Algorithm



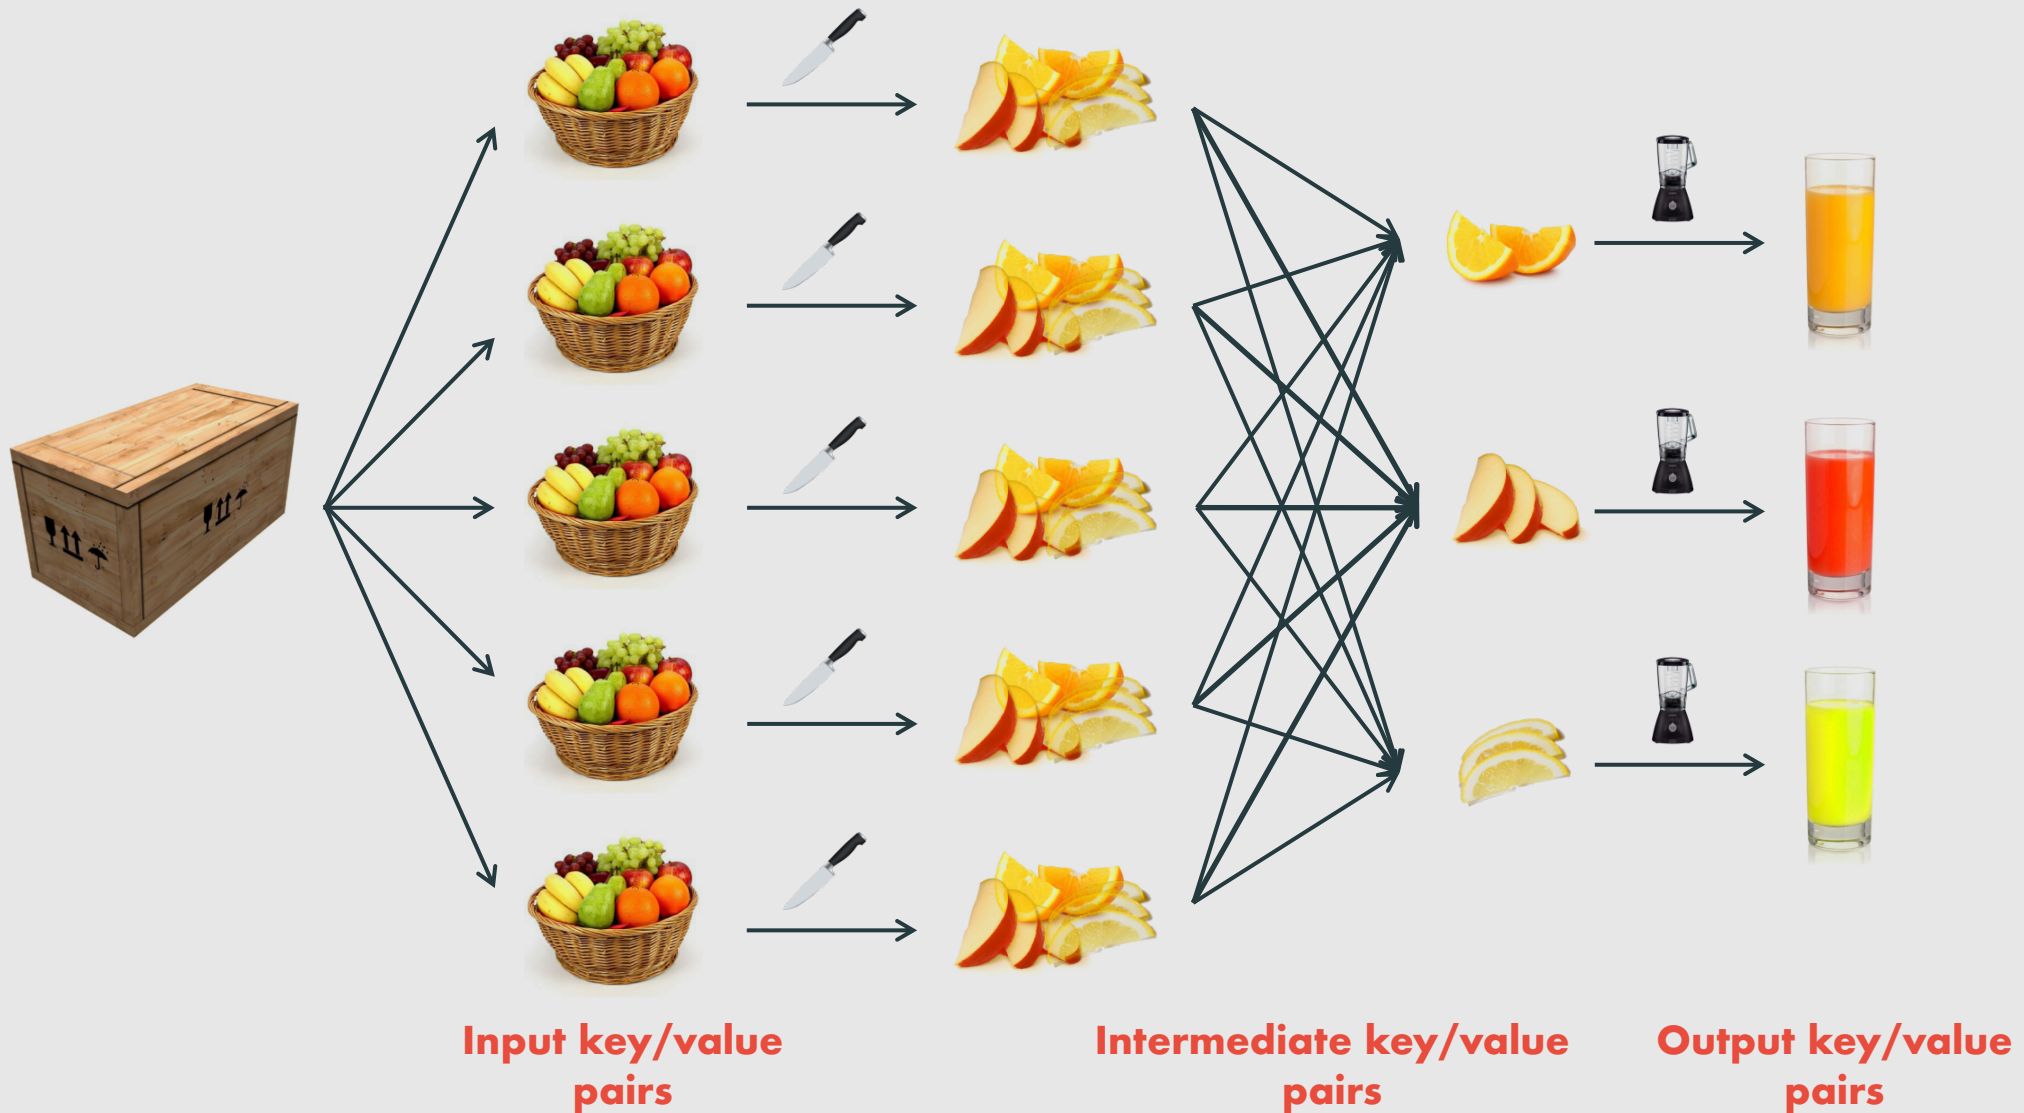
Parallel Juice Algorithm



Parallel Juice Algorithm

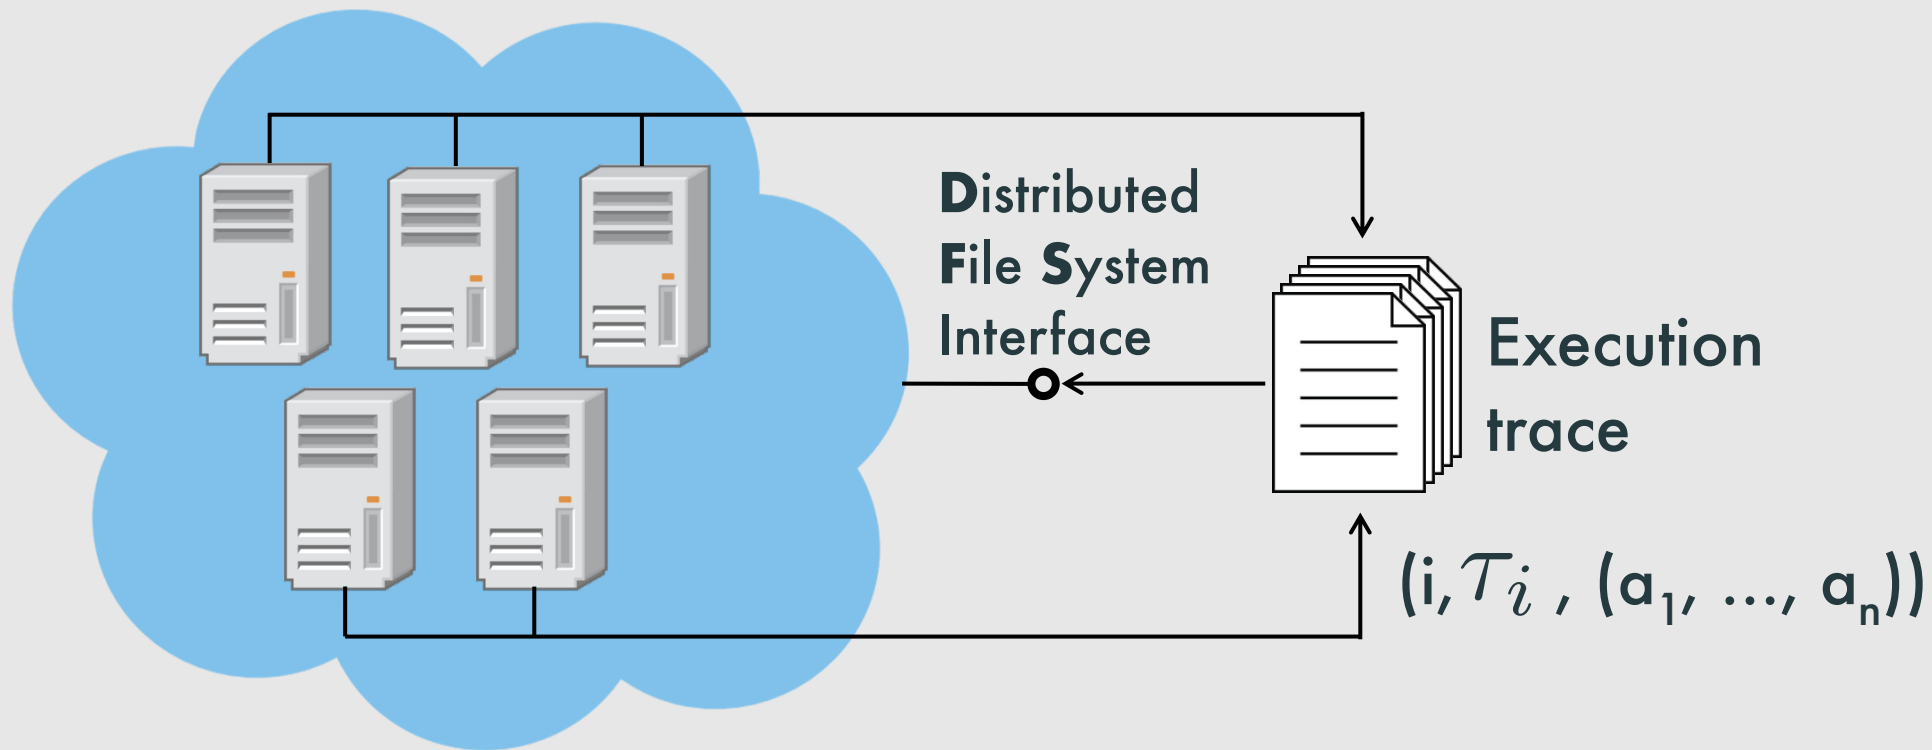


Parallel Juice Algorithm




Algorithm


Assumptions on Traces



Formula to Evaluate

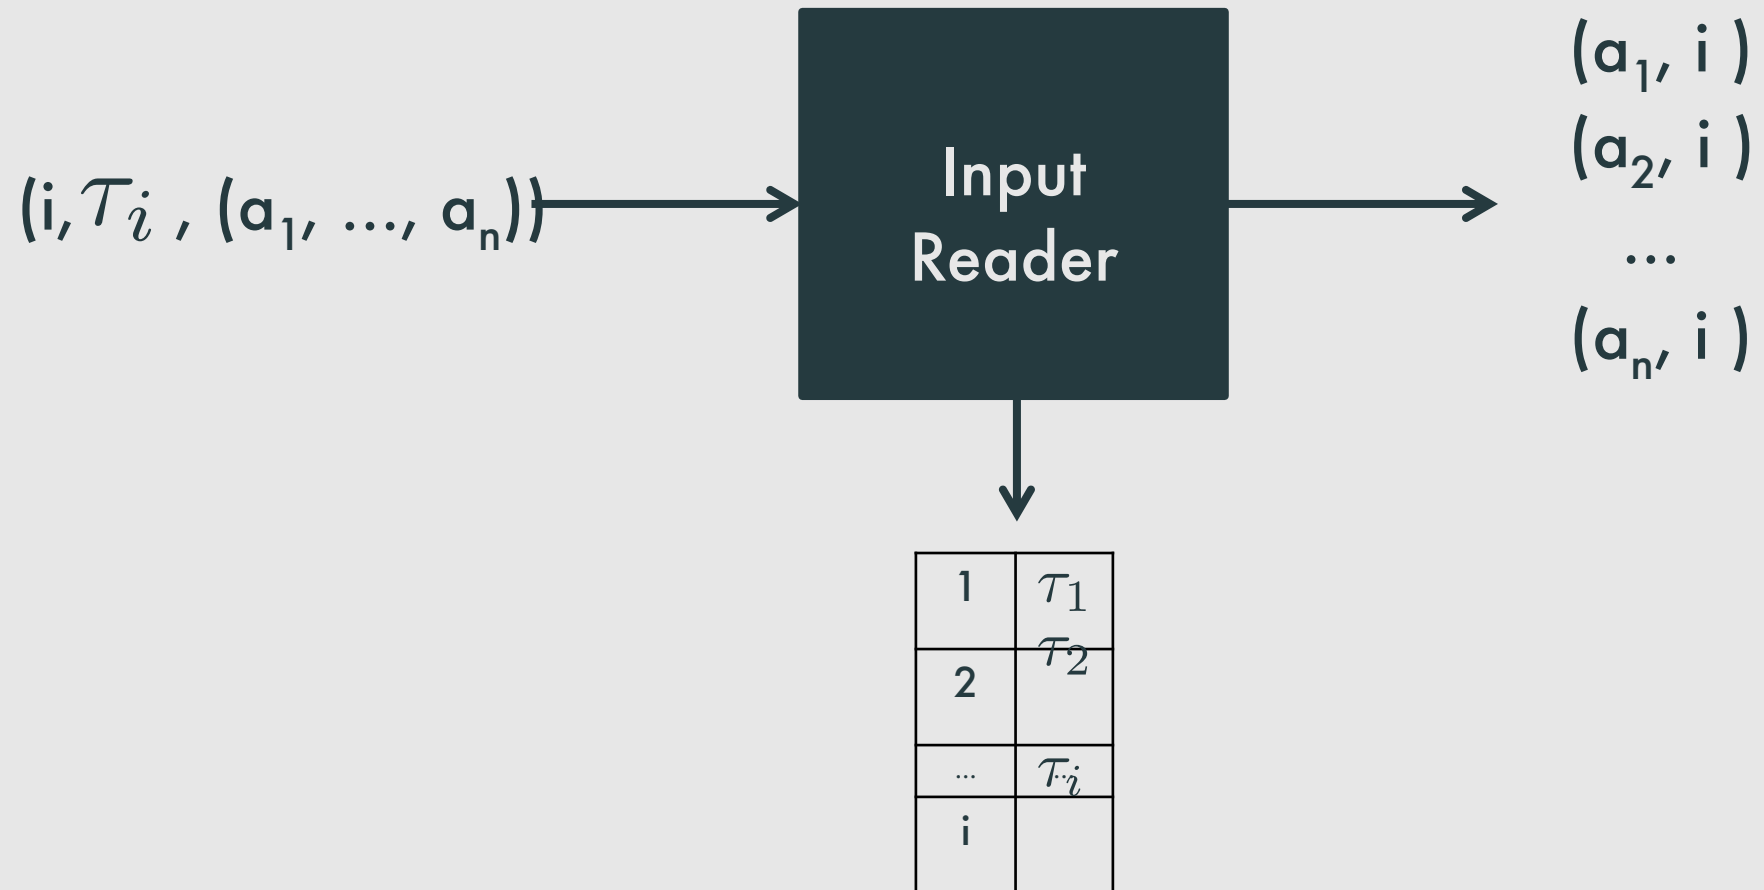


Formula to Evaluate





Algorithm Iteration 1



Input Reader



Input Reader





Input Reader



1	0
---	---

(C,1)


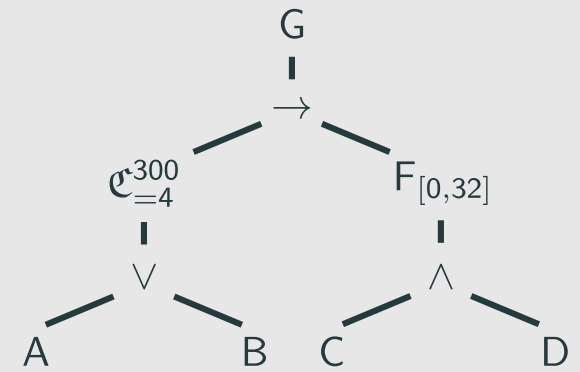
Input Reader



1	0
2	50

(C,1) (A,2)
(B,2)



Input Reader



1	0
2	50
3	70

(C,1) (A,2) (A,3)
(B,2)



Input Reader



1	0
2	50
3	70
4	160

(C,1) (A,2) (A,3) (B,4)
 (B,2) (C,4)

Input Reader




1	0
2	50
3	70
4	160
5	320
6	340
7	351
8	640

(C,1)
(A,2)
(A,3)
(B,4)
(B,5)
(A,6)
(C,7)
(D,8)


(B,2)
(C,4)
(D,7)

Mapper




ϕ is direct subformula of ψ

Mapper




(C,1)
(A,2)
(B,2)
(A,3)
(B,4)
(C,4)
(B,5)
(A,6)
(C,7)
(D,7)
(D,8)



(\wedge , C,1)
(\vee , A,2)
(\vee , B,2)
(\vee , A,3)
(\vee , B,4)
(\wedge , C,4)
(\vee , B,5)
(\vee , A,6)
(\wedge , C,7)
(\wedge , D,7)
(\wedge , D,8)

Shuffle & Sort

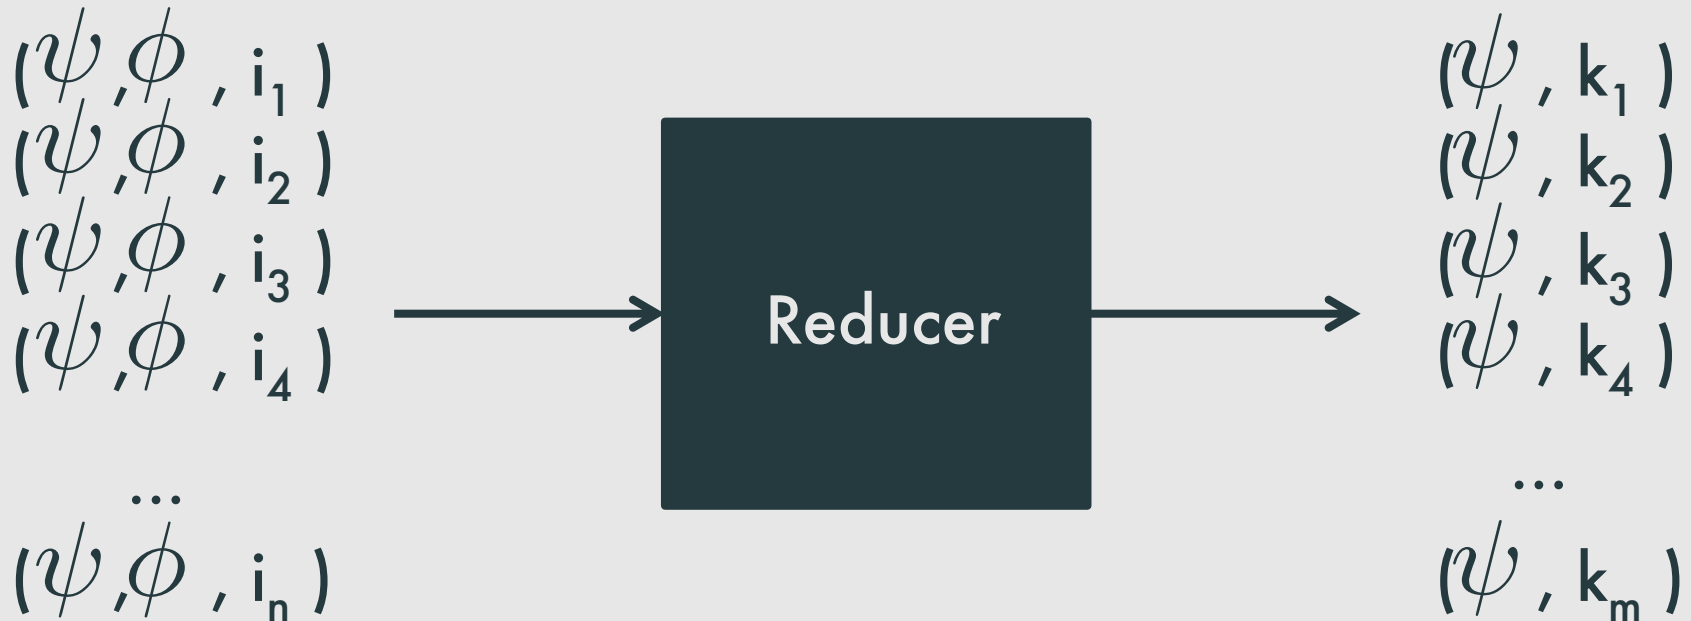
(^ , C,1)
(v , A,2)
(v , B,2)
(v , A,3)
(v , B,4)
(^ , C,4)
(v , B,5)
(v , A,6)
(^ , C,7)
(^ , D,7)
(^ , D,8)



(v , A,2)
(v , B,2)
(v , A,3)
(v , B,4)
(v , B,5)
(v , A,6)
(v , B,7)

(^ , C,1)
(^ , C,4)
(^ , C,7)
(^ , D,7)
(^ , D,8)

Reducer



Formula ψ holds at positions (k_1, k_2, \dots, k_m)


Reducer

(√, A, 2)
(√, B, 2)
(√, A, 3)
(√, B, 4)
(√, B, 5)
(√, A, 6)




(√, 2)
(√, 3)
(√, 4)
(√, 5)
(√, 6)

Reducer



Algorithm Iteration 2

Mapper



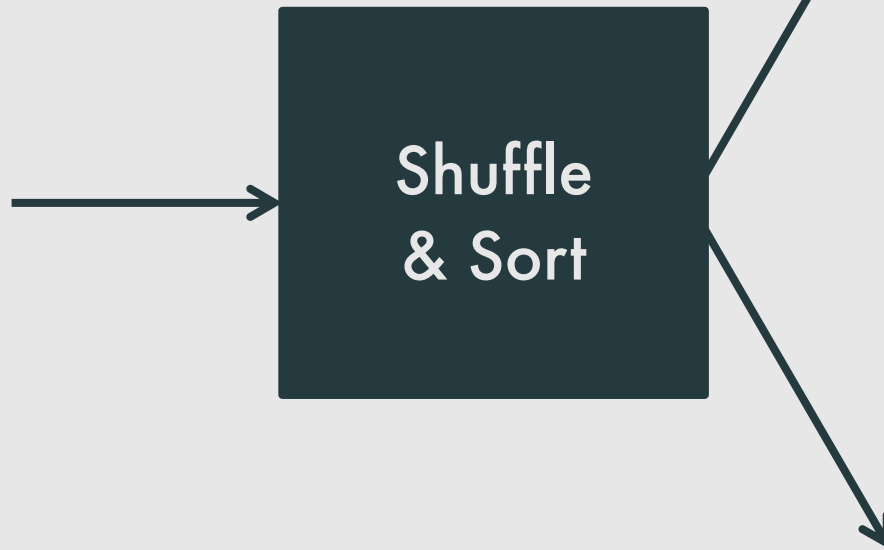
- (V,2)
- (V,3)
- (V,4)
- (V,5)
- (V,6)
- (^,7)



- ($e_{=4}^{300}$, V,2)
- ($e_{=4}^{300}$, V,3)
- ($e_{=4}^{300}$, V,4)
- ($e_{=4}^{300}$, V,5)
- ($e_{=4}^{300}$, V,6)
- ($F_{[0,32]}$, ^,7)

Shuffle & Sort


$(e_{=4}^{300}, \vee, 2)$
 $(e_{=4}^{300}, \vee, 3)$
 $(e_{=4}^{300}, \vee, 4)$
 $(e_{=4}^{300}, \vee, 5)$
 $(e_{=4}^{300}, \vee, 6)$
 $(F_{[0,32]}, \wedge, 7)$



$(e_{=4}^{300}, \vee, 2)$
 $(e_{=4}^{300}, \vee, 3)$
 $(e_{=4}^{300}, \vee, 4)$
 $(e_{=4}^{300}, \vee, 5)$
 $(e_{=4}^{300}, \vee, 6)$


$(F_{[0,32]}, \wedge, 7)$

Reducer



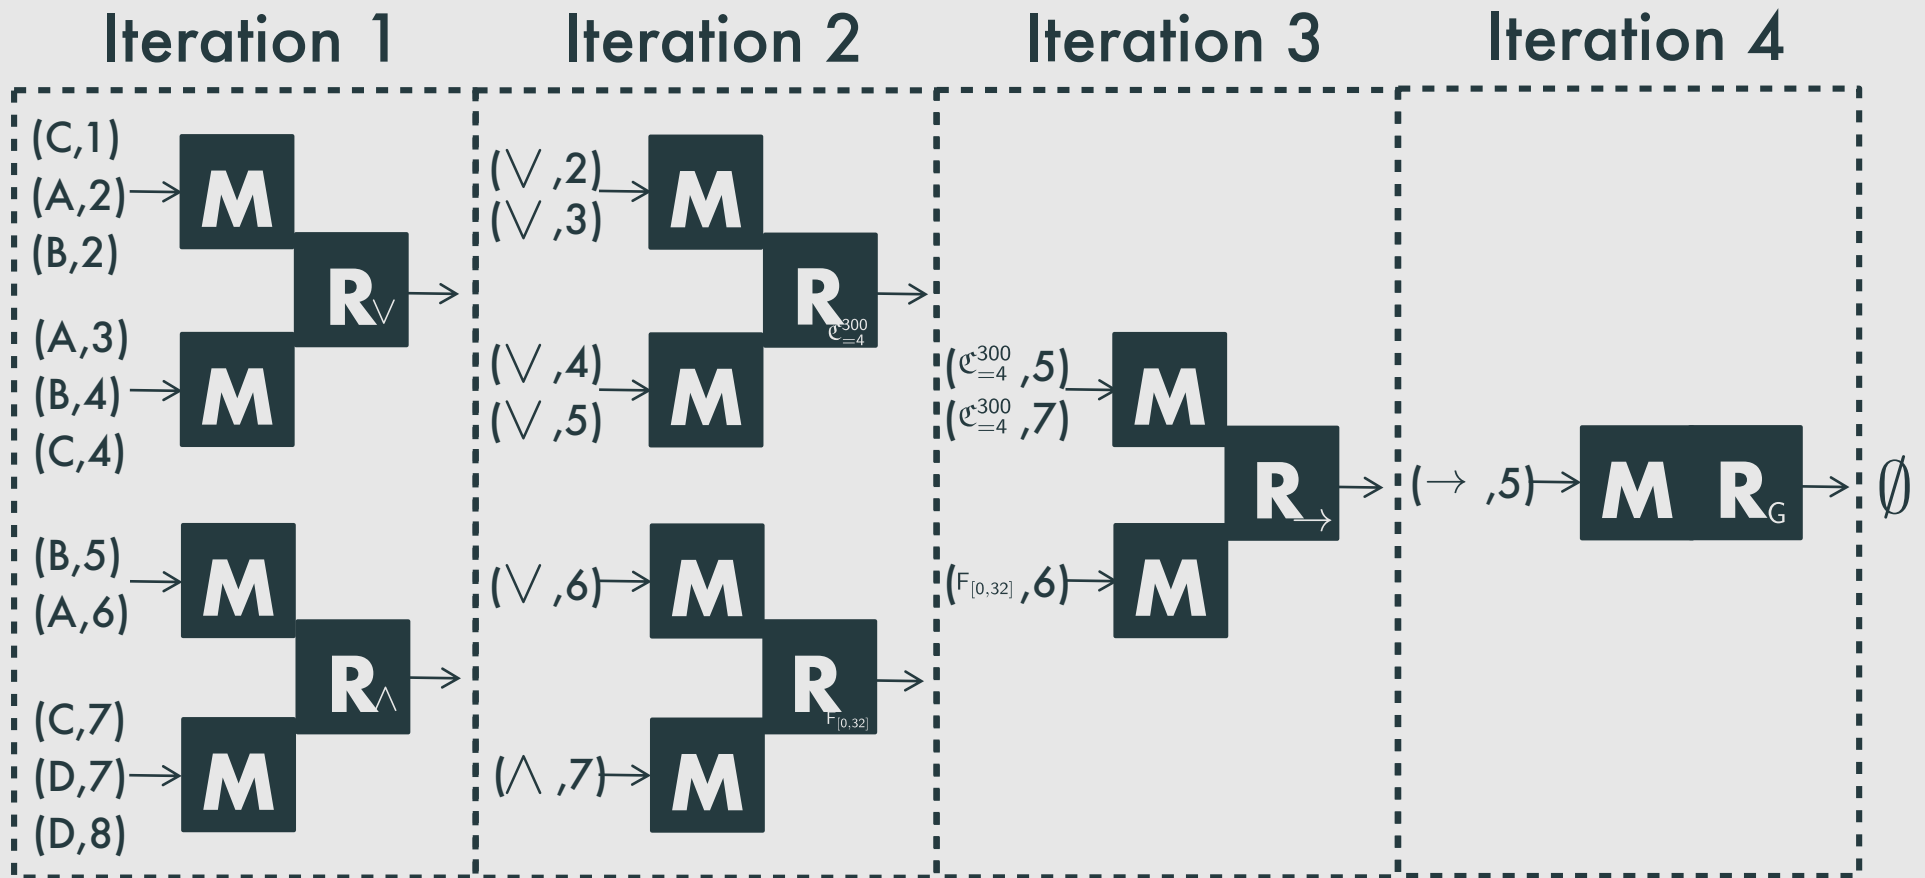
1	0
2	50
3	70
4	160
5	320
6	340
7	351
8	640

Reducer



1	0
2	50
3	70
4	160
5	320
6	340
7	351
8	640

Iterations

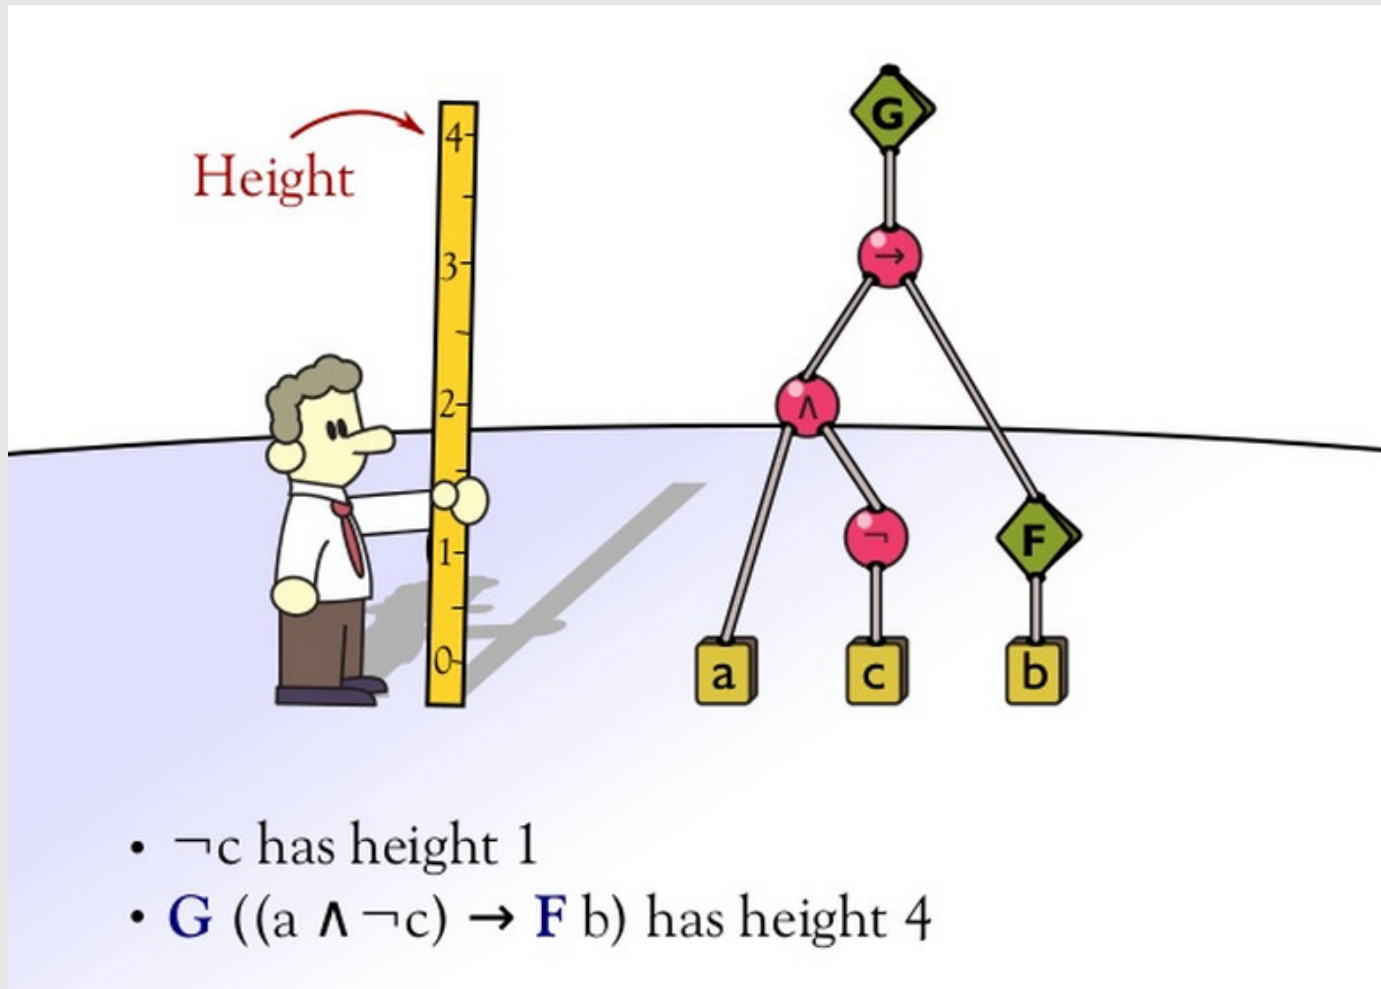


Remarks

- .Exploits structure of the formula to obtain parallelization
- .Handles traces that cannot fit into the operating memory
- .Checks complex formal specifications that include timing information and aggregate metrics

Evaluation

Related Work



Benjamin Barre, Mathieu Klein, Maxime Soucy-Boivin, Pierre-Antoine Ollivier, Sylvain Hallé:
MapReduce for Parallel Trace Validation of LTL Properties

Implementation

- .Algorithm implemented in Java + Hadoop framework
- .Executed on Windows Azure cloud-based infrastructure
- .10 VMs with 1 CPU core and 1.75GB of RAM
- .100 Mappers and 9 Reducers

Evaluation Setting

.Randomly generated traces

~ 350,000 time instants

~ 100 proposition per time instant

~ time span 578.7 days on average

Evaluation Setting

.Scalability

Formulae:

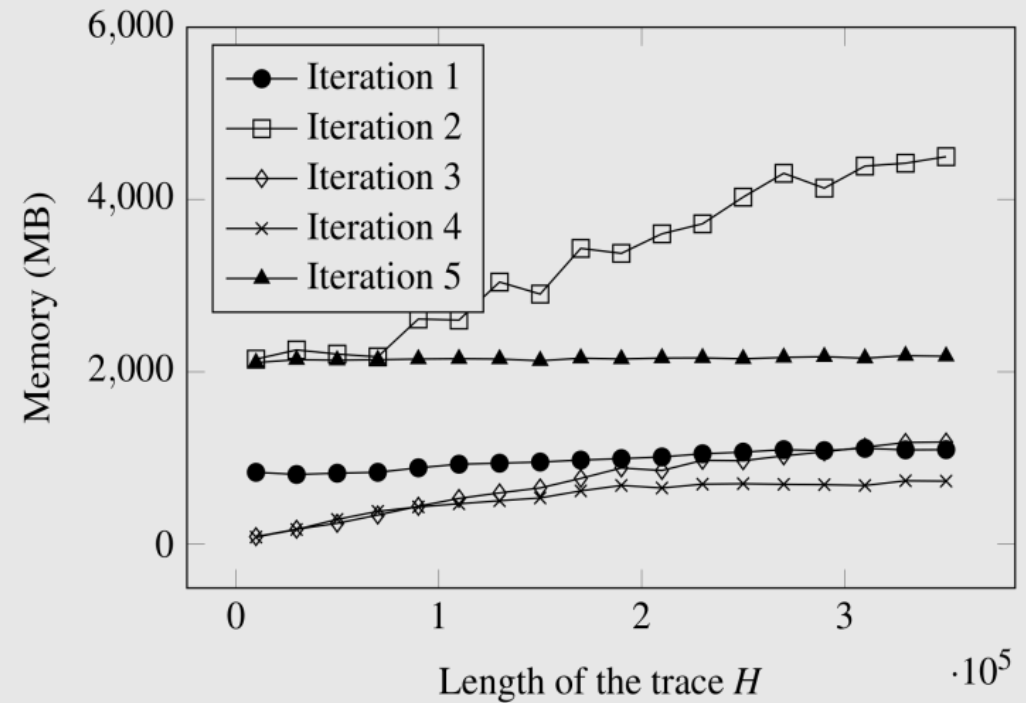

$$\mathcal{E}_{<10}^{50000}(a_0)$$

$$\mathcal{D}_{<10}^{50000}(a_1, a_2)$$

$$(a_0 \wedge (a_1 \wedge a_2))U_{(50,200)}((a_1 \wedge a_2) \vee a_1)$$

$$\exists j \in \{0 \dots 9\} \forall i \in \{0 \dots 8\} : G_{(50,500)}(a_{i,j} \rightarrow X_{(50,500)}(a_{i+1,j}))$$

Scalability



Formula: $\exists j \in \{0 \dots 9\} \forall i \in \{0 \dots 8\} : G_{(50,500)}(a_{i,j} \rightarrow X_{(50,500)}(a_{i+1,j}))$

Evaluation Setting

.Comparison with LTL algorithm*

Formulae:

$$G_{(50,500)}(\neg a_0)$$

$$G_{(50,500)}(a_0 \rightarrow X_{(50,500)}(a_1))$$

$$\forall i \in \{0 \dots 8\} : G_{(50,500)}(a_i \rightarrow X_{(50,500)}(a_{i+1}))$$

$$\exists j \in \{0 \dots 9\} \forall i \in \{0 \dots 8\} : G_{(50,500)}(a_{i,j} \rightarrow X_{(50,500)}(a_{i+1,j}))$$

*Benjamin Barre, Mathieu Klein, Maxime Soucy-Boivin, Pierre-Antoine Ollivier, Sylvain Hallé:
MapReduce for Parallel Trace Validation of LTL Properties

Comparison with LTL Algorithm*

	Property 1		Property 2		Property 3		Property 4	
	LTL	SOLOIST	LTL	SOLOIST	LTL	SOLOIST	LTL	SOLOIST
Number of tuples	55K	16K	120K	24K	599K	215K	4.9M	1.7M
Time per event (μ s)	19	1.2	21	1.9	14	3.7	30	7.2

* Benjamin Barre, Mathieu Klein, Maxime Soucy-Boivin, Pierre-Antoine Ollivier, Sylvain Hallé:
MapReduce for Parallel Trace Validation of LTL Properties

Wrapping up

Overview

Future Work

- .Distributed timestamp management
- .Splitting the tuples received by reducers

Trace checking of Metric Temporal Logic with Aggregating Modalities using MapReduce

Srdan Krstić

with

Domenico Bianculli and Carlo Ghezzi

