



Performance Comparison for Scientific Computations on the Edge via Relative Performance

Speaker : **Aravind Sankaran**¹
Supervisor : **Paolo Bientinesi**²

3rd Workshop on Parallel AI and Systems for the Edge (PAISE)
May 21, 2021

¹ International Research Training Group for Modern Inverse Problem, RWTH Aachen University, Germany

² Department of Computing Science, Umeå Universitet, Sweden

Abstract Scientific Code

Scientific Code {

Math Task 1

Math Task 2

⋮

Math Task K

}

Setup:

We consider scientific codes consisting of sequence of math tasks.

Abstract Scientific Code

Scientific Code {

Math Task 1

Math Task 2

⋮

Math Task K

}

Matrix chain

ABCD + E

Setup:

We consider scientific codes consisting of sequence of math tasks.

Some examples of math tasks:

1. Linear algebra operations.

Abstract Scientific Code

Scientific Code {

Math Task 1

Math Task 2

⋮

Math Task K

}

Matrix chain

ABCD + E

Least Square problem

```
for i = 1, ..., n do:  
  update A  
   $Z = (A^T A + \lambda \cdot I)^{-1} A^T B$   
  loss =  $\|AZ - B\|^2$ 
```

Setup:

We consider scientific codes consisting of sequence of math tasks.

Some examples of math tasks:

1. Linear algebra operations.
2. Evaluation of some loss function in a loop.

Abstract Scientific Code

Scientific Code {

Math Task 1

Math Task 2

⋮

Math Task K

}

Matrix chain

ABCD + E

Least Square problem

```
for i = 1, ..., n do:  
  update A  
   $Z = (A^T A + \lambda \cdot I)^{-1} A^T B$   
  loss =  $\|AZ - B\|^2$ 
```

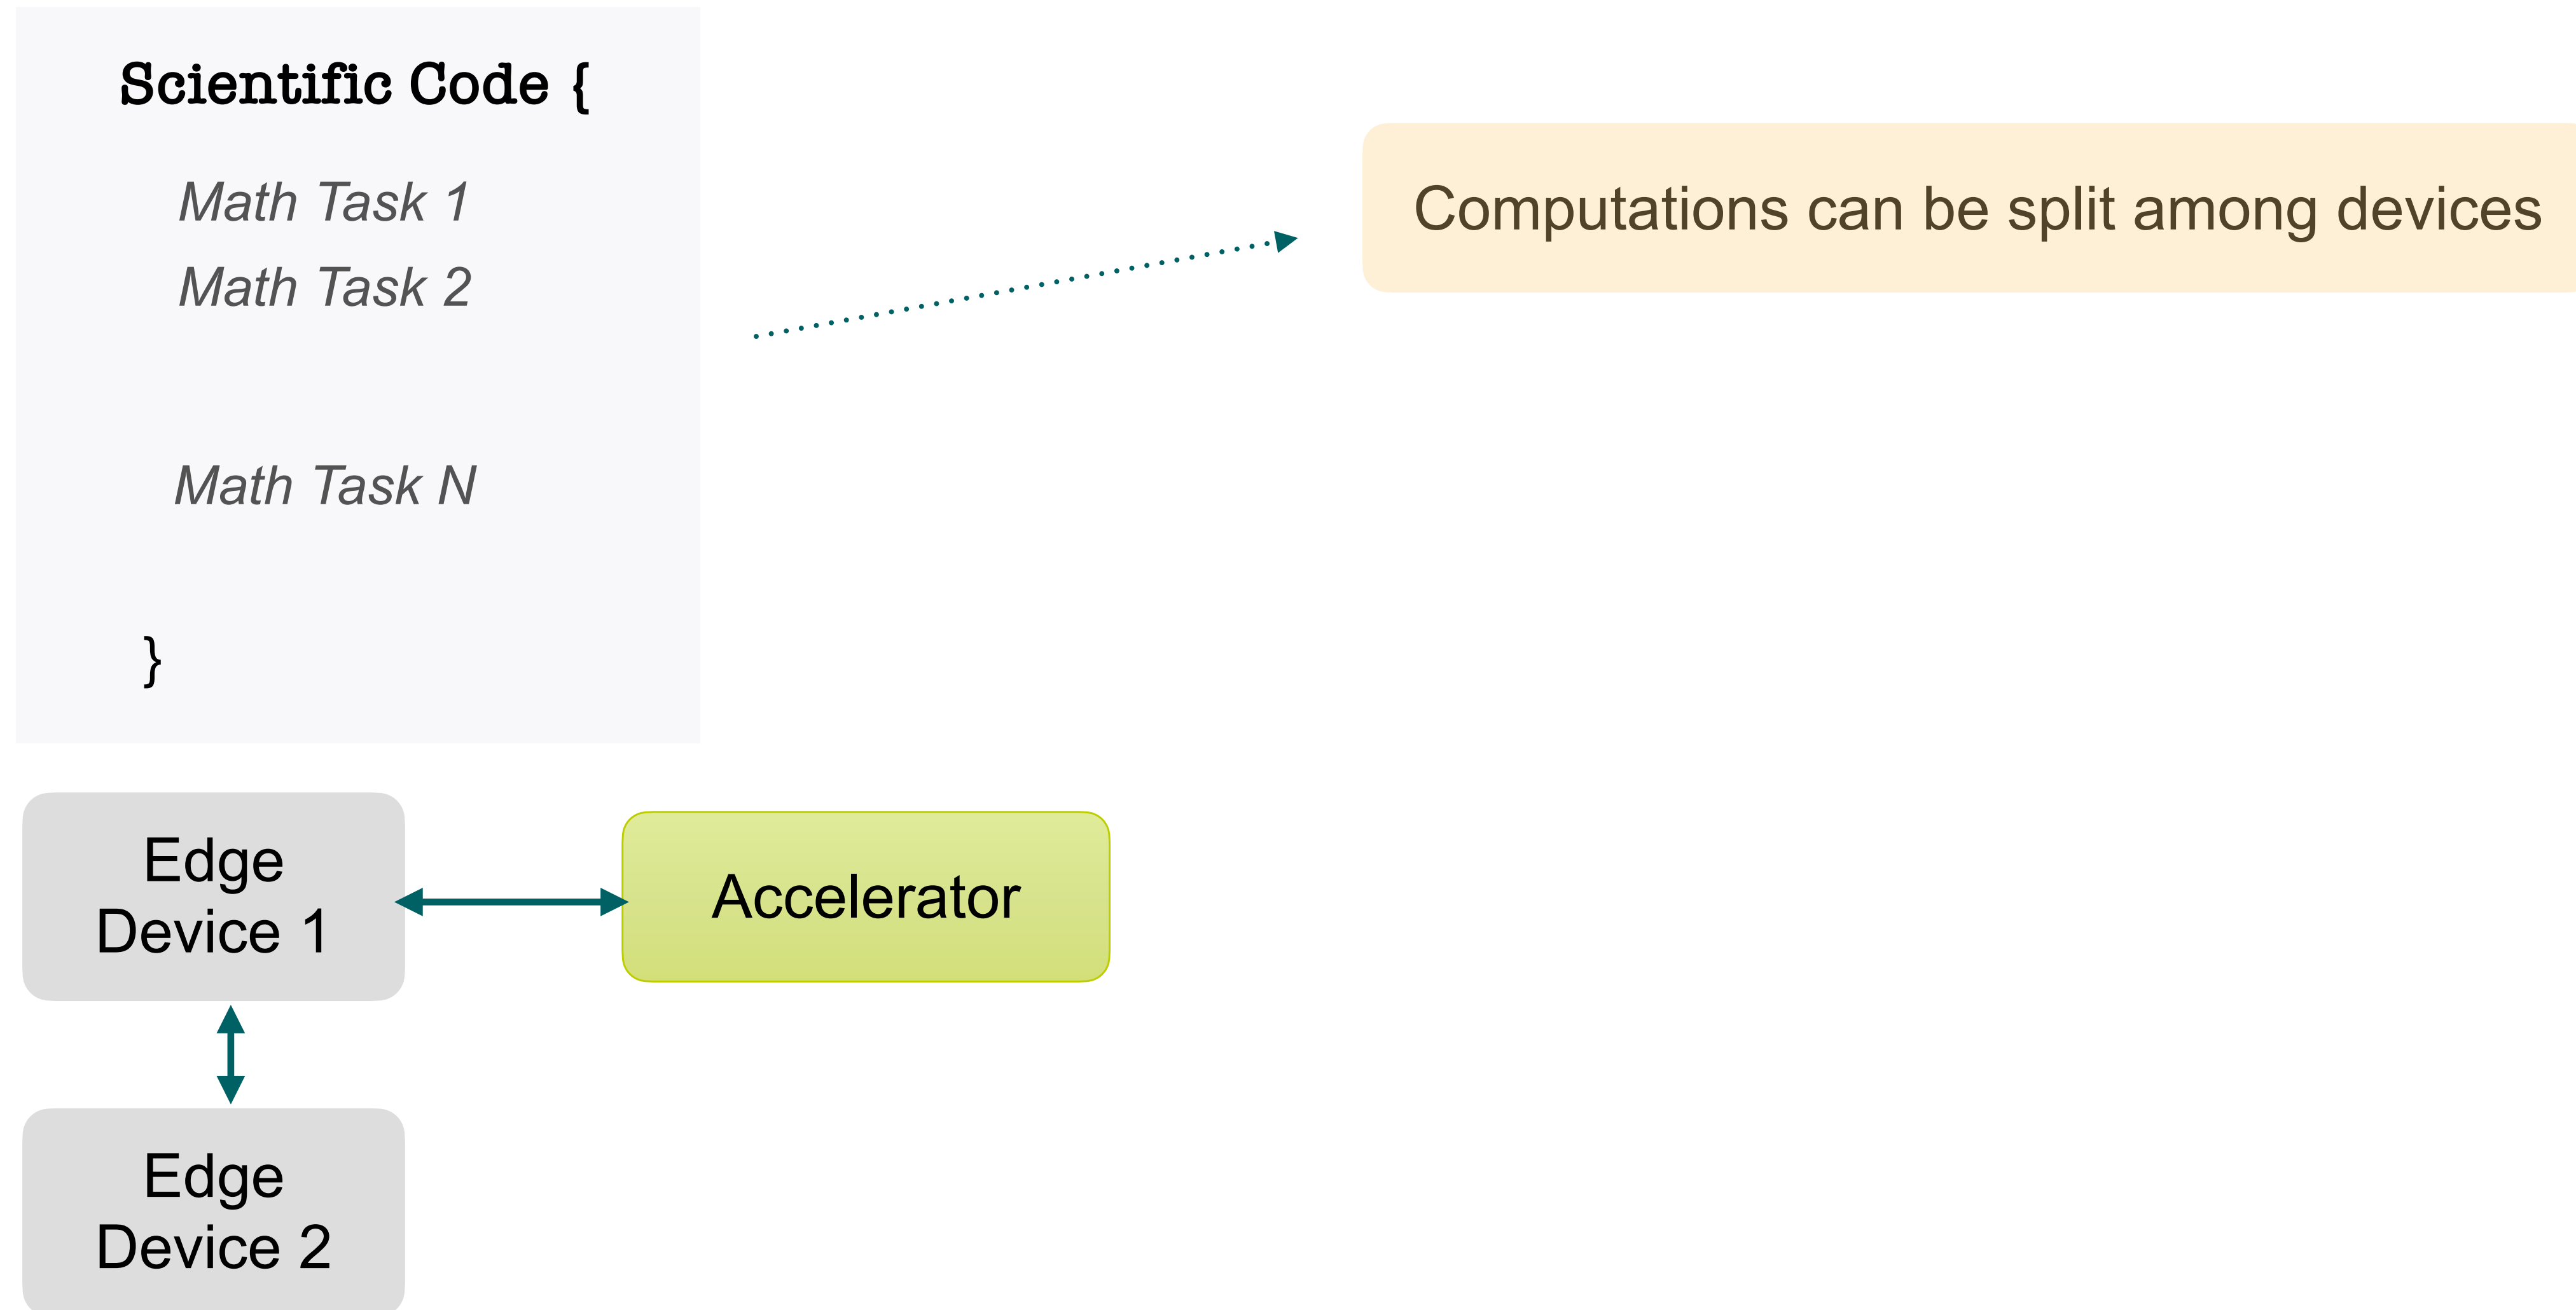
Setup:

We consider scientific codes consisting of sequence of math tasks.

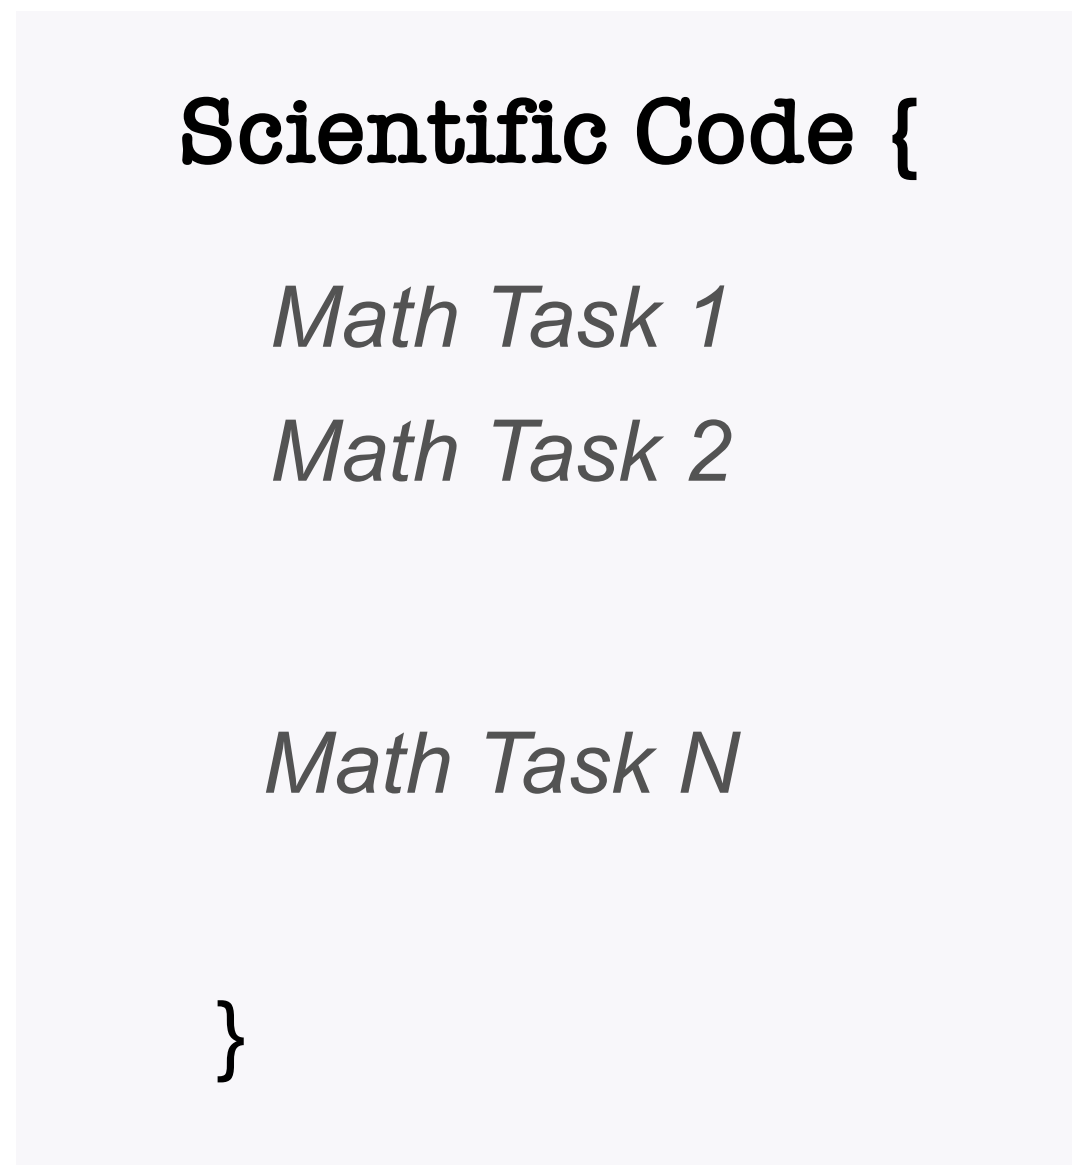
Some examples of math tasks:

1. Linear algebra operations.
2. Evaluation of some loss function in a loop.
3. Neural network.

Partitioning and distribution of Scientific Code among various devices



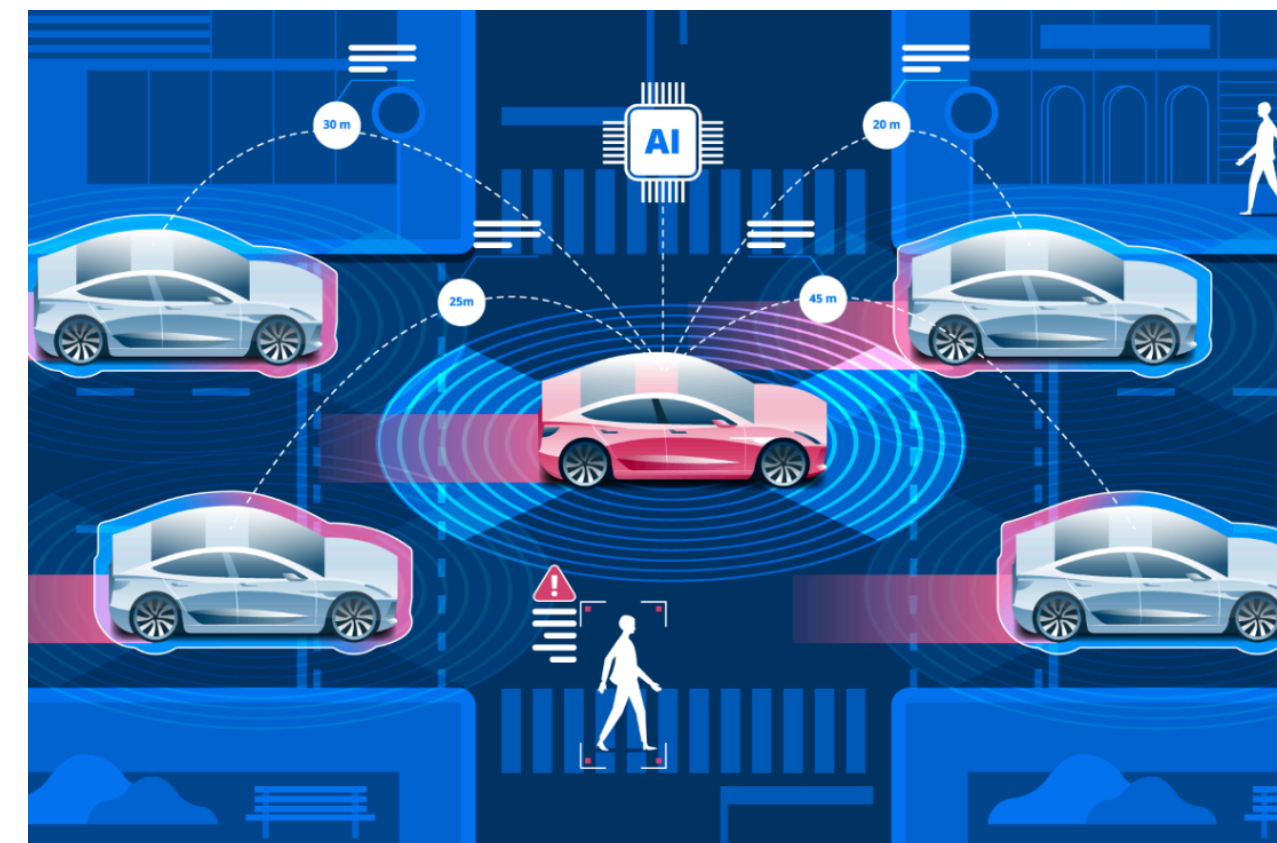
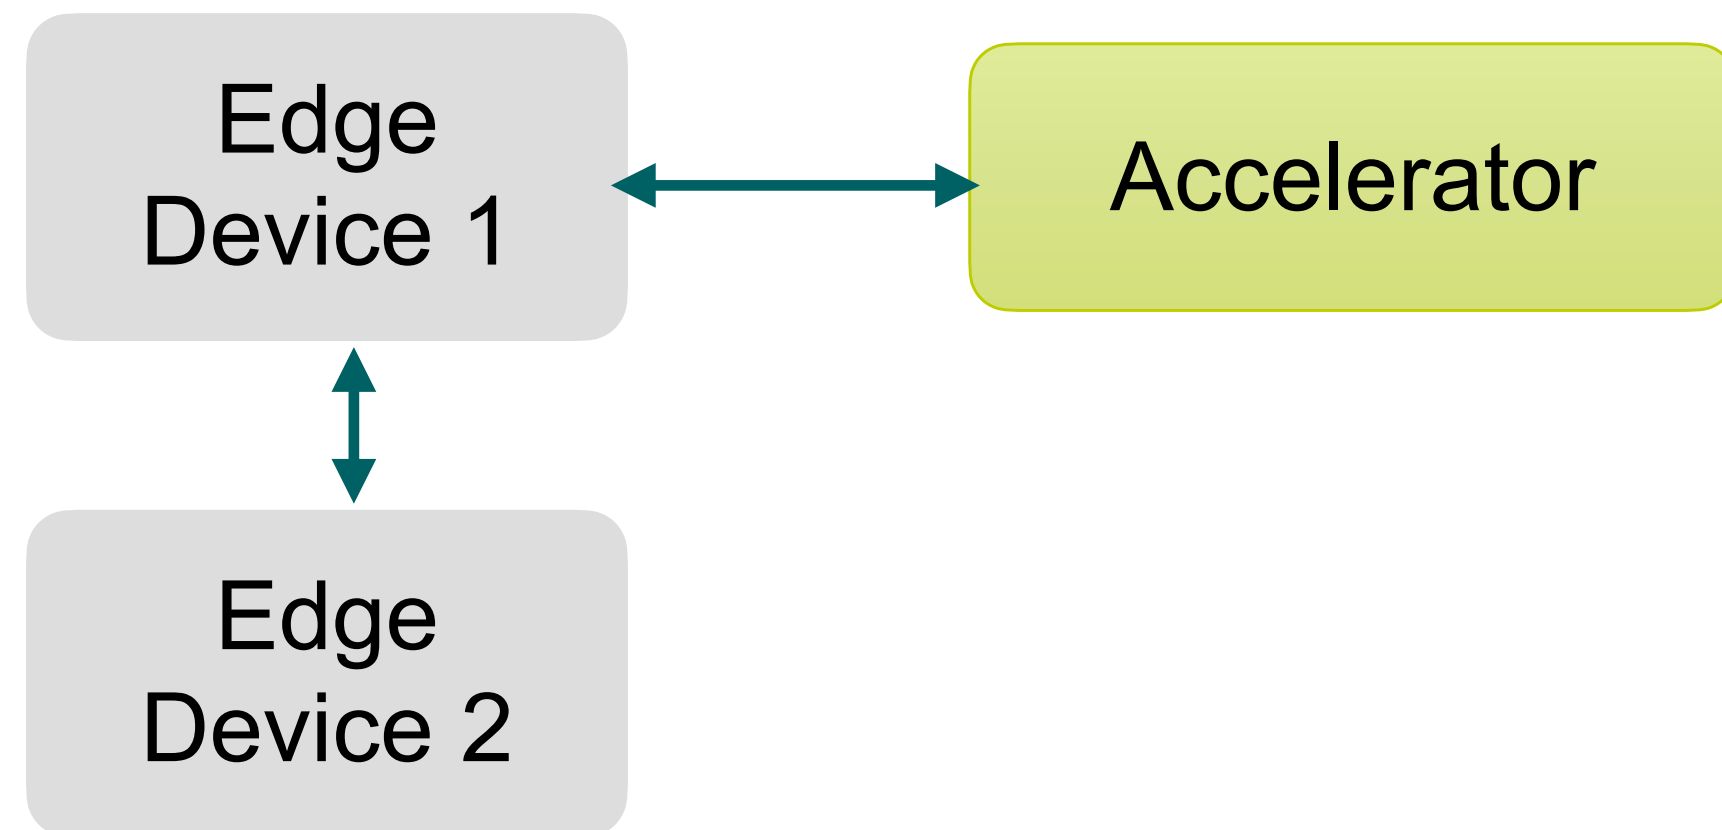
Partitioning and distribution of Scientific Code among various devices



Computations can be split among devices

Applications:

- 1) Connected Autonomous Vehicles.



Source: <https://www.smartcitiesworld.net/>

Partitioning and distribution of Scientific Code among various devices

Scientific Code {

Math Task 1

Math Task 2

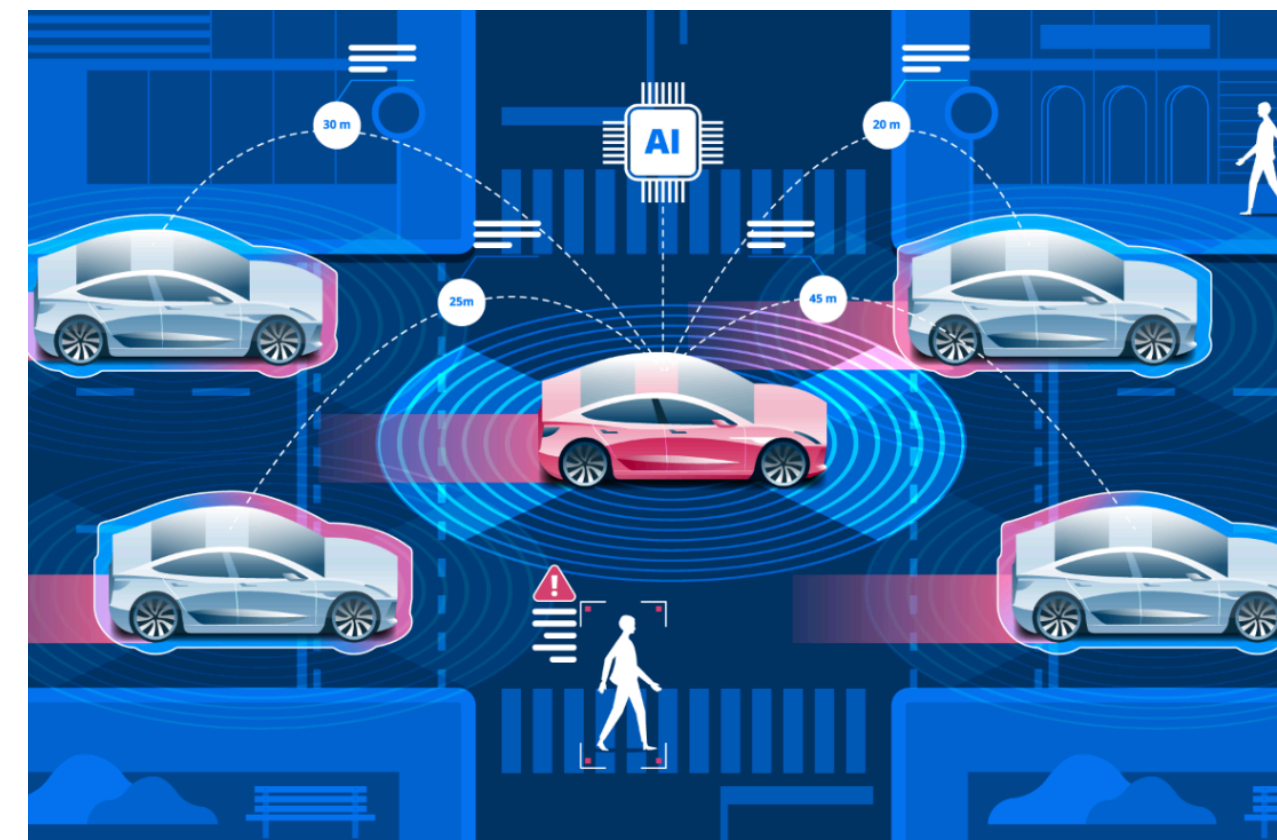
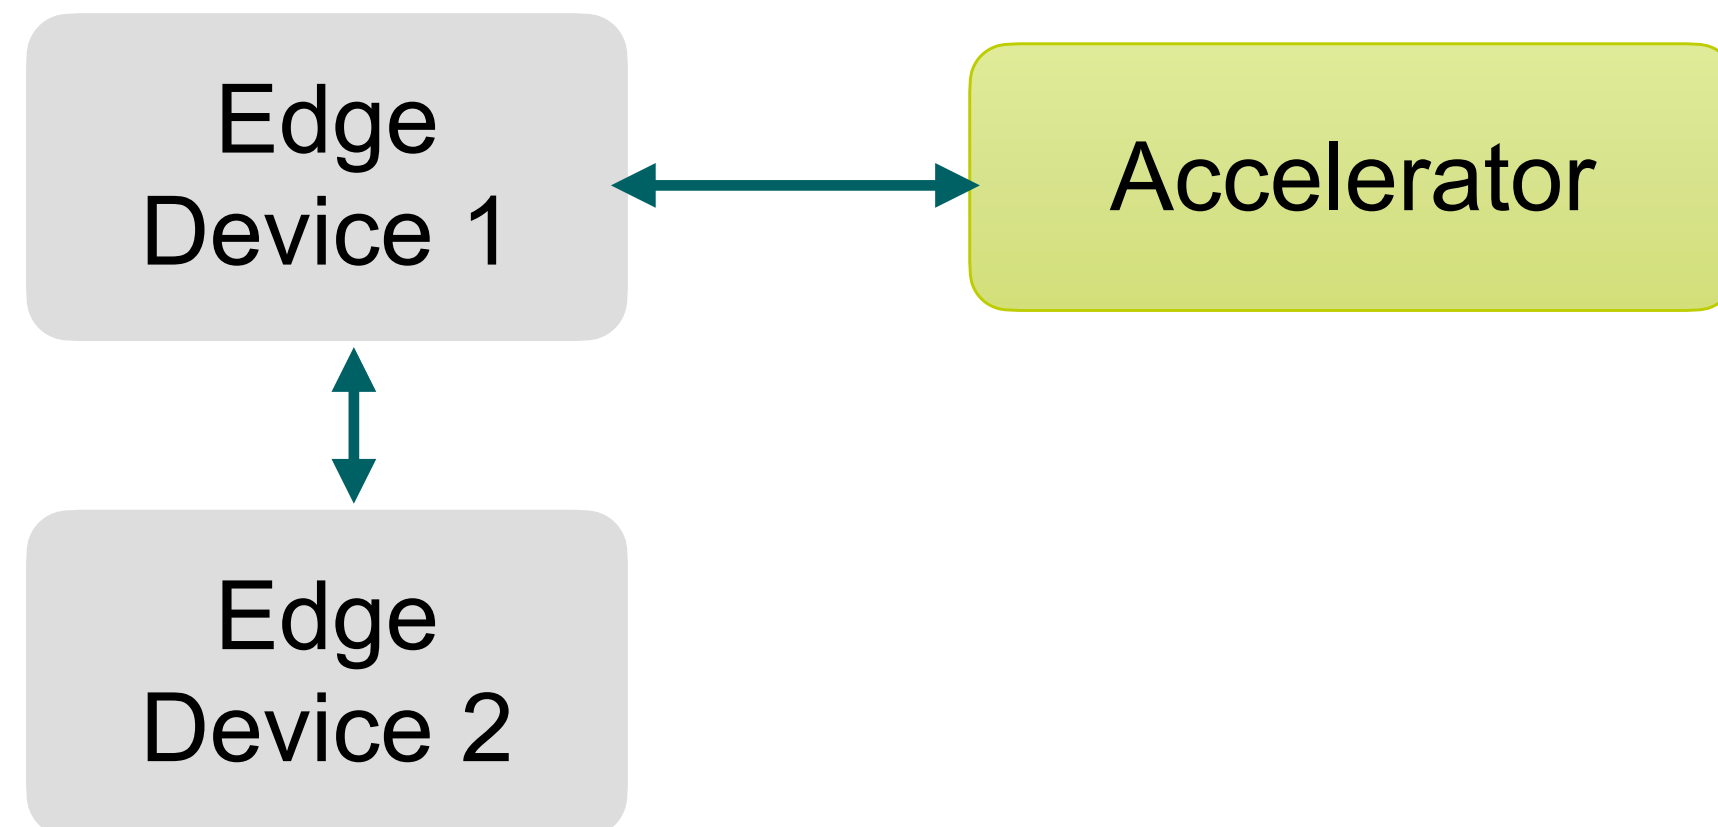
Math Task N

}

Computations can be split among devices

Applications:

- 1) Connected Autonomous Vehicles.
- 2) Digital Twins.

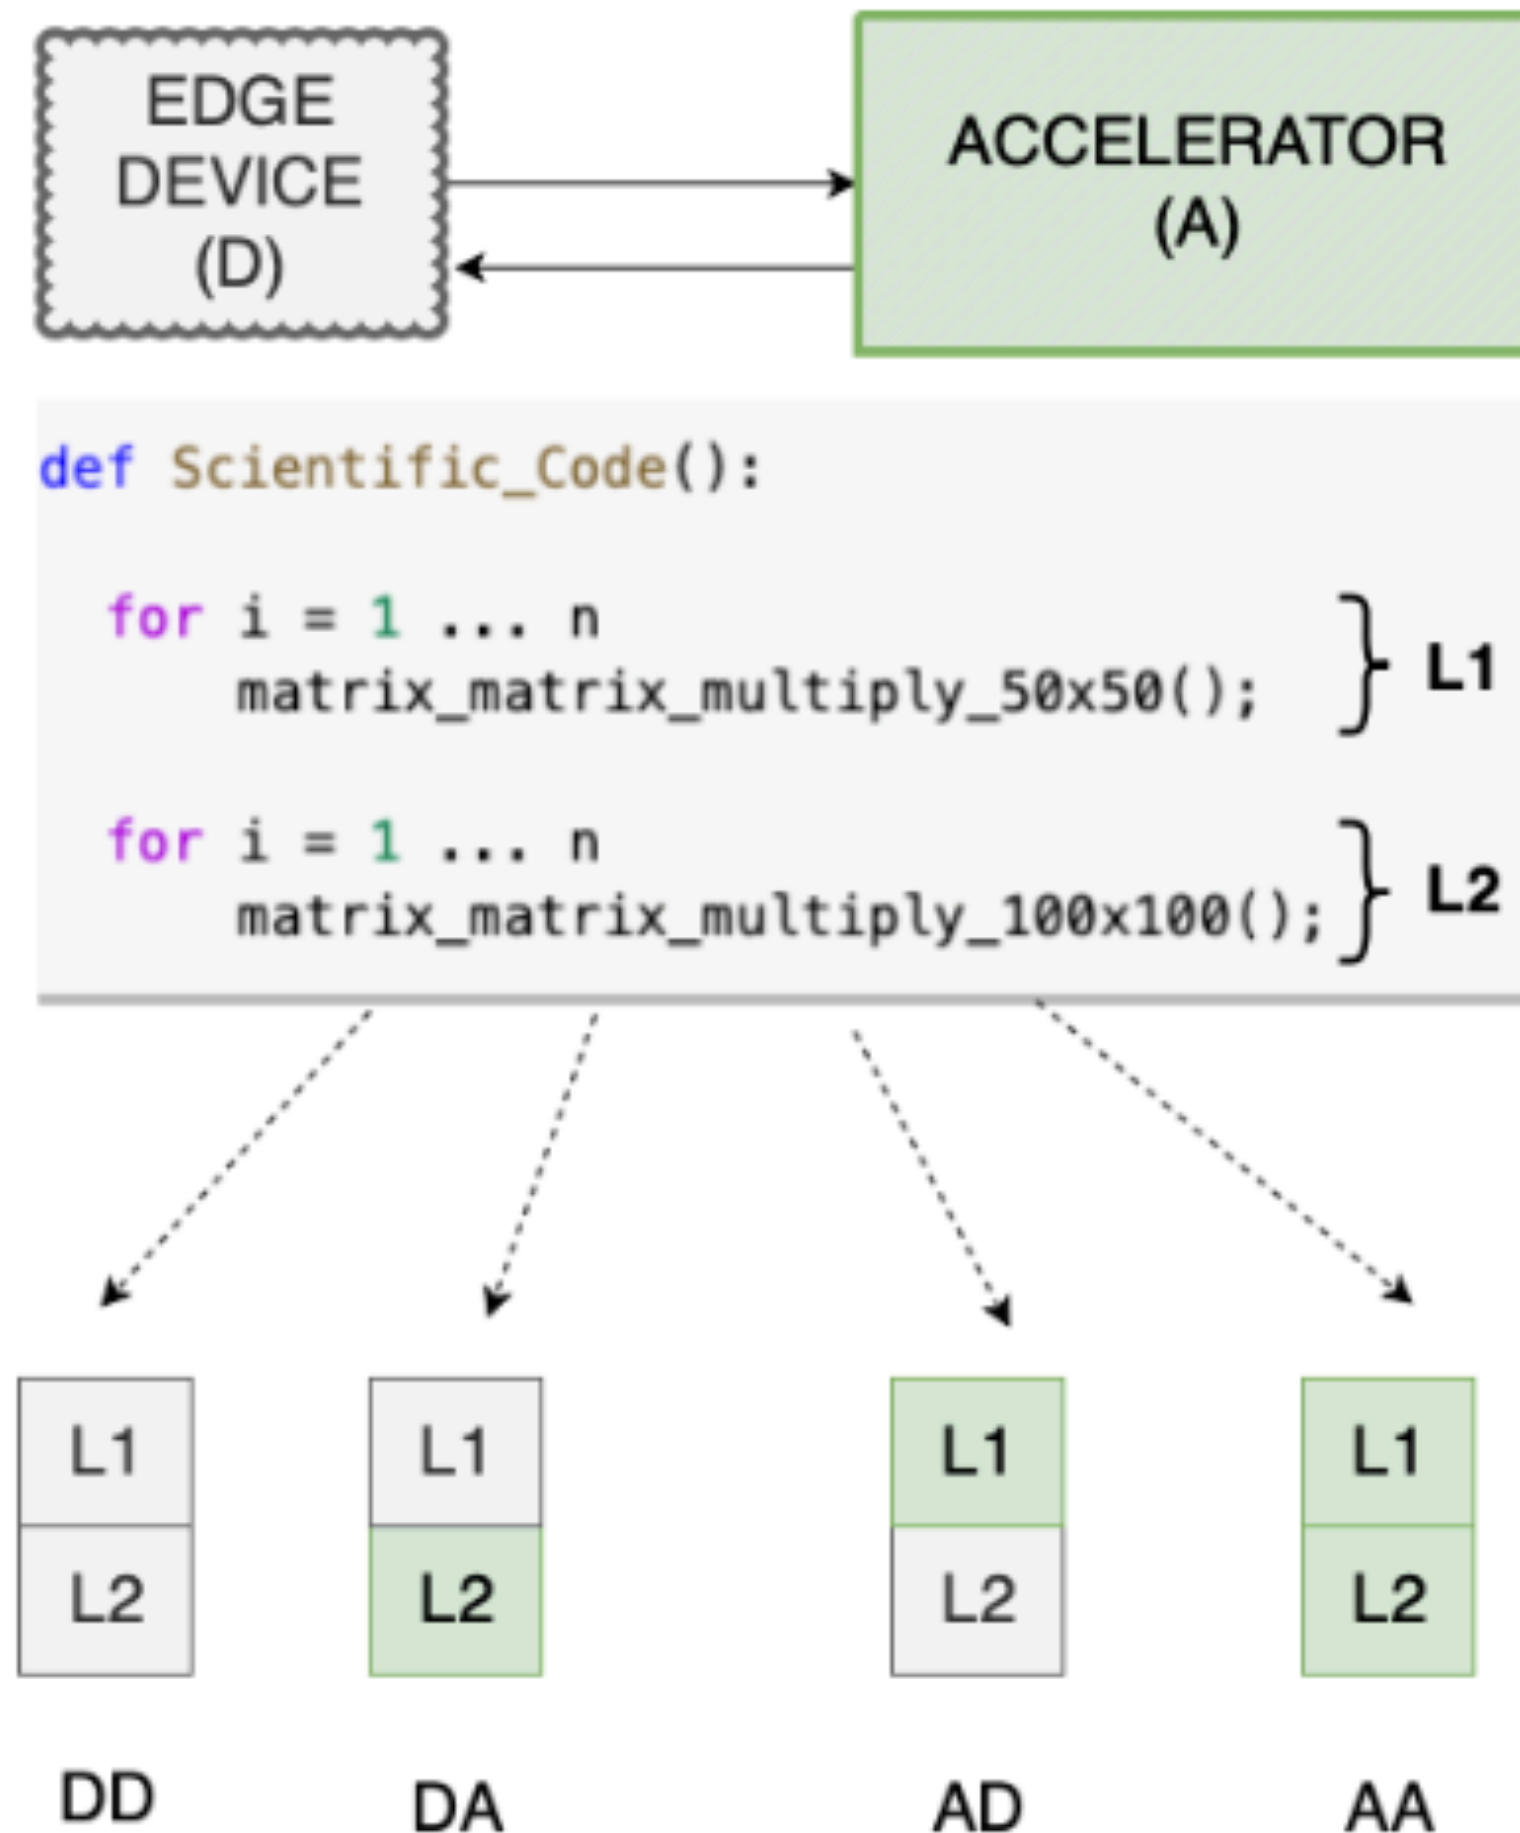


Source: <https://www.smartcitiesworld.net/>



Source: <https://metrology.news>

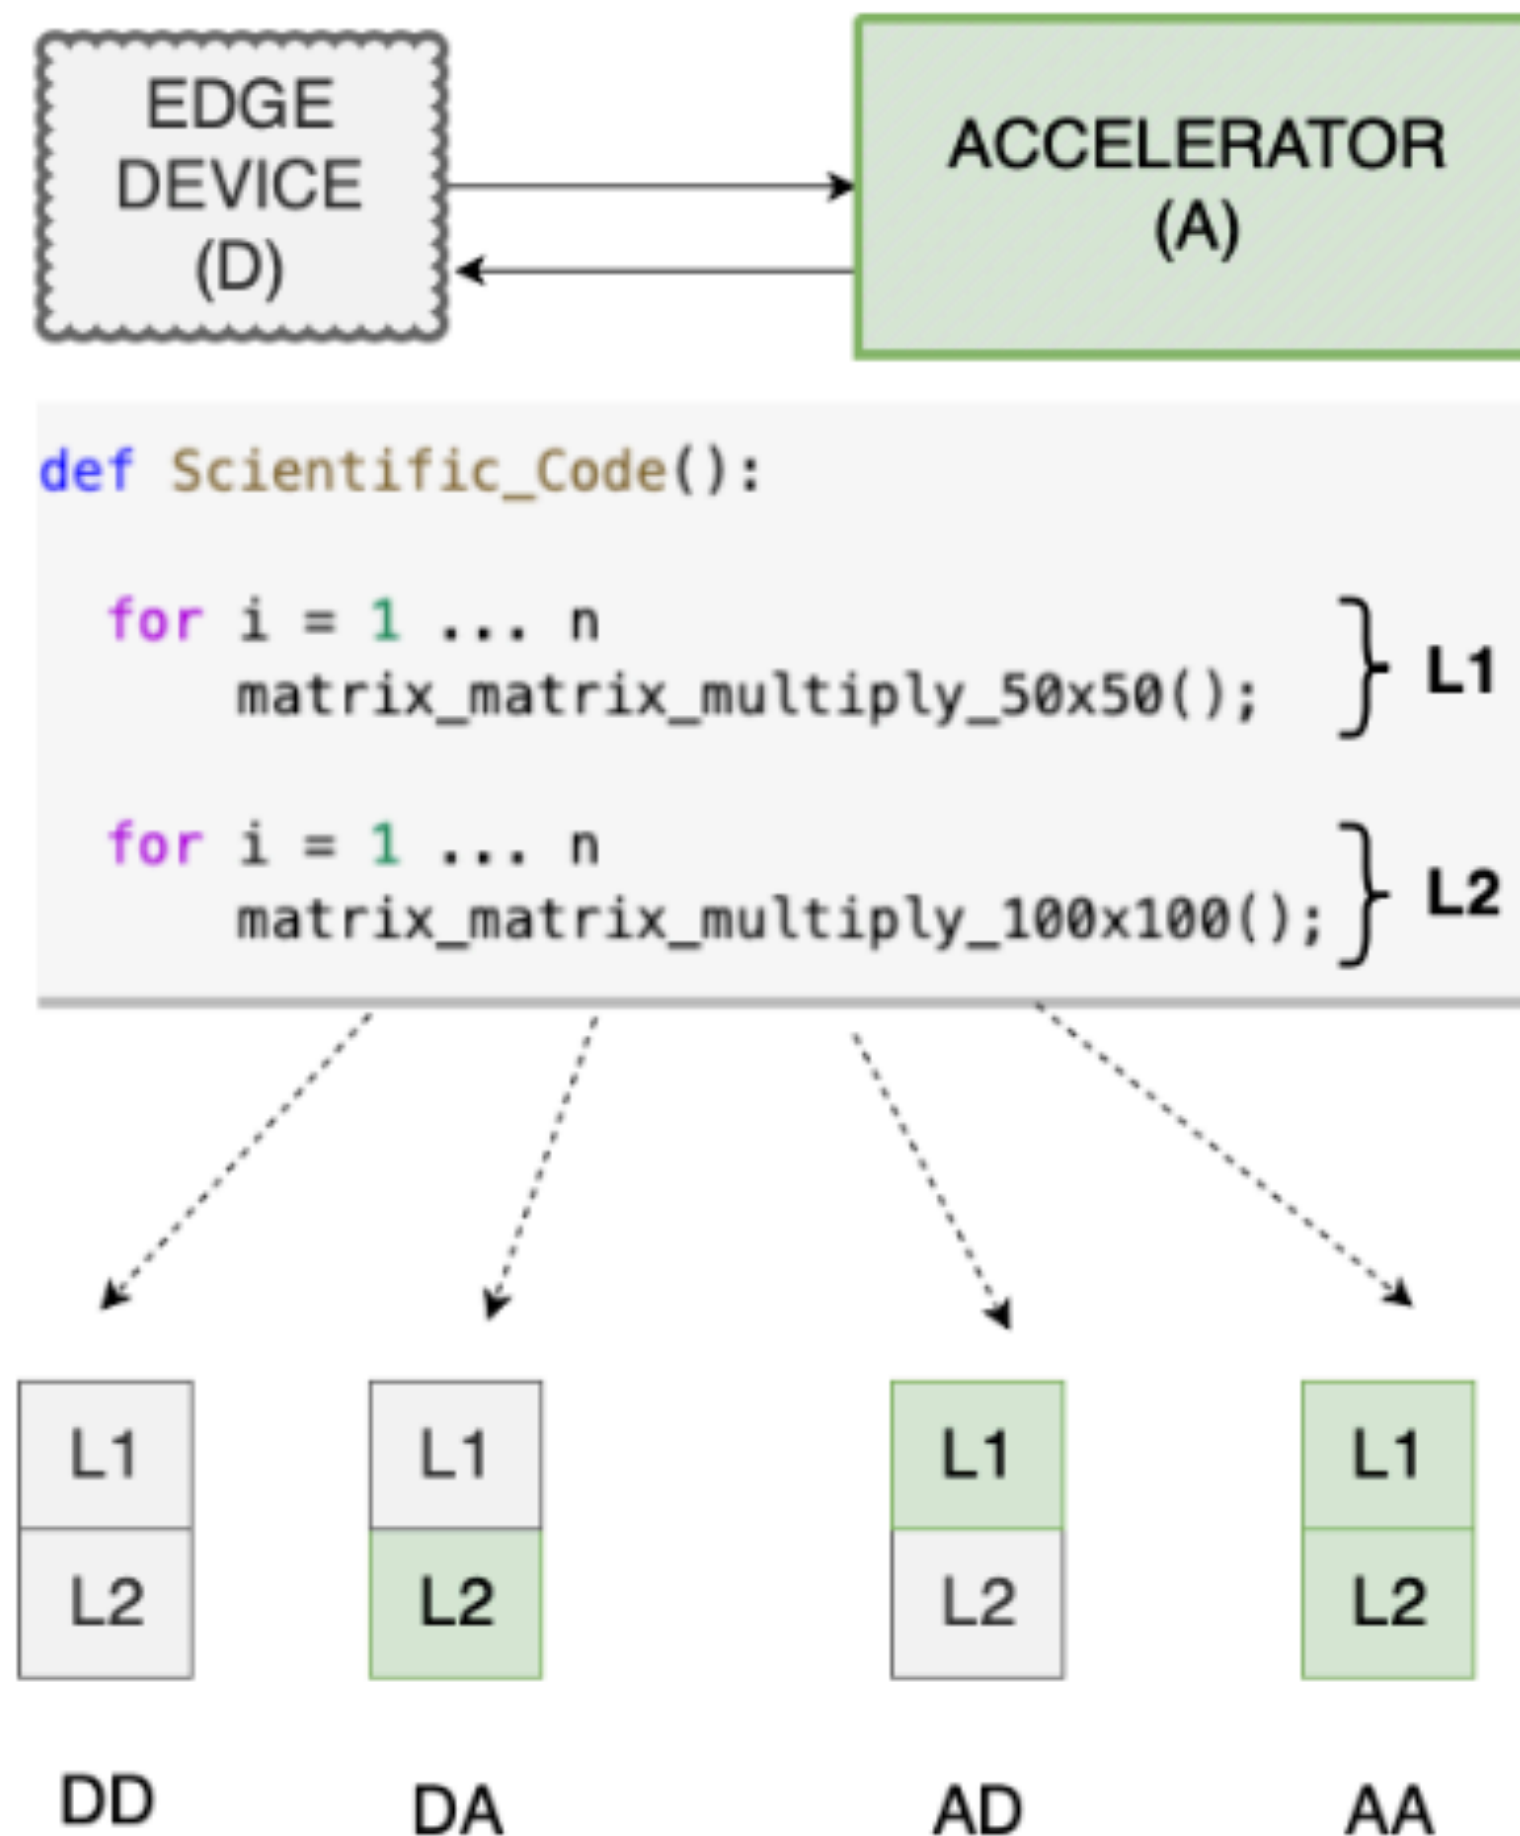
Capturing Performance variations



An example showing four possible ways of splitting a scientific code between CPU and GPU:

To simplify our explanation, we assume L1 and L2 **cannot** be executed concurrently.

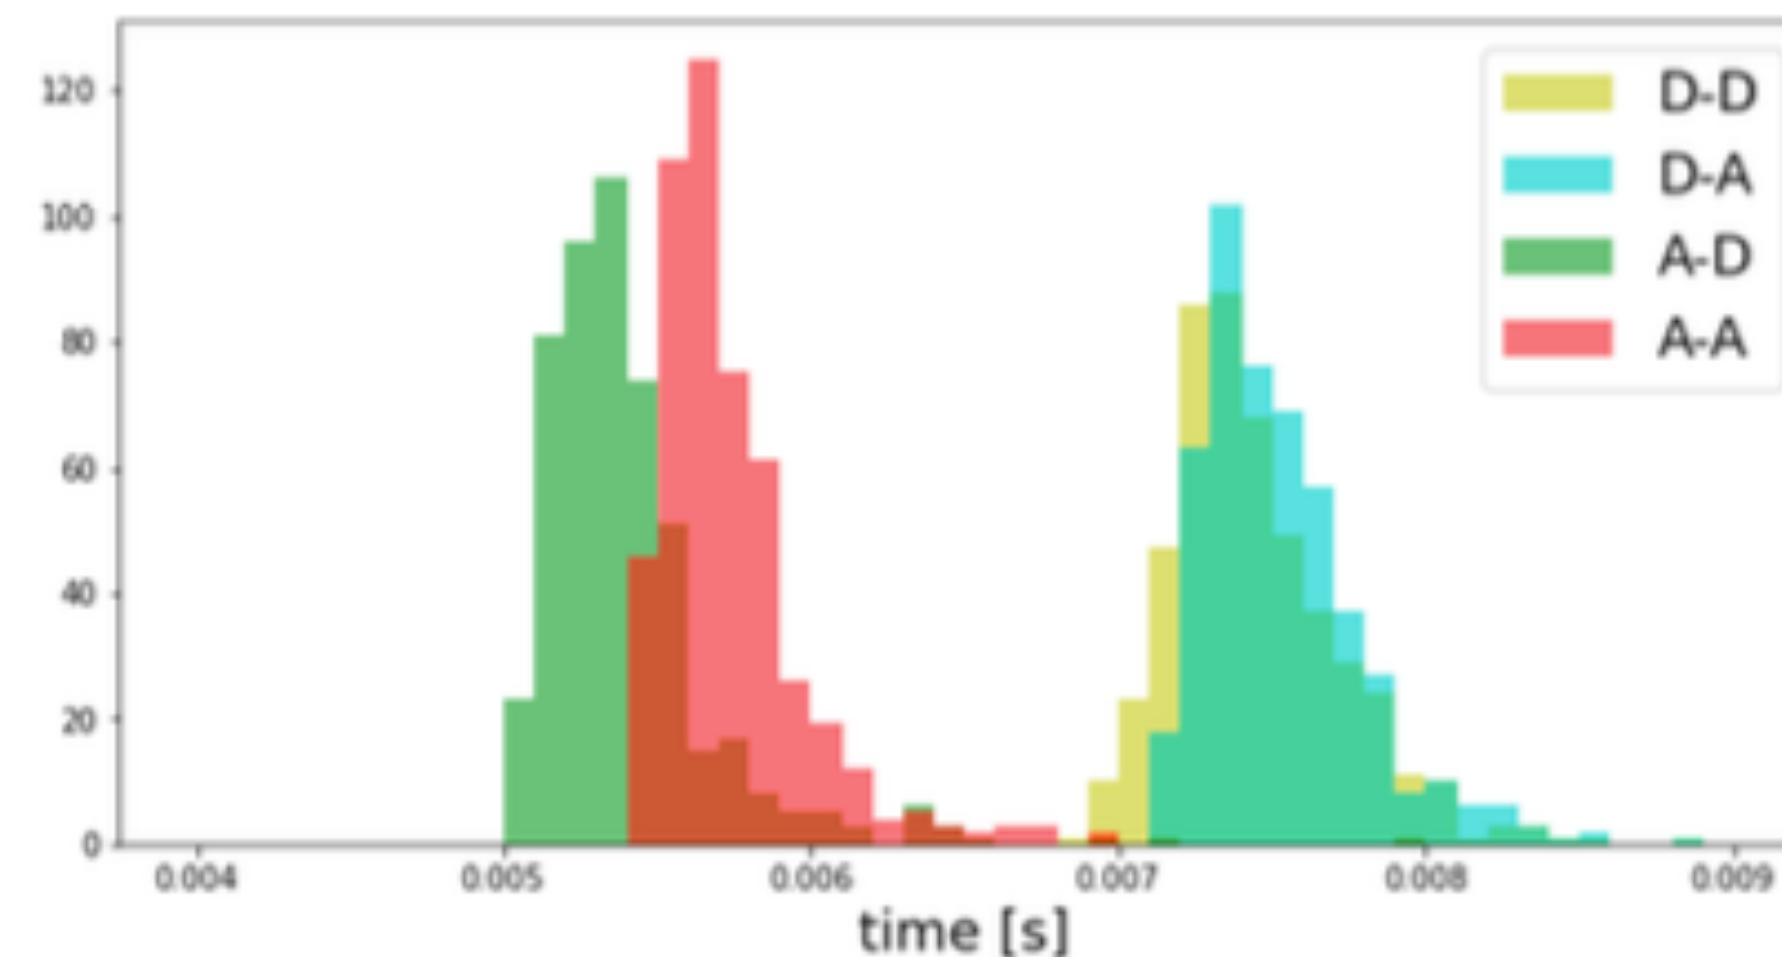
Capturing Performance variations



An example showing four possible ways of splitting a scientific code between CPU and GPU:

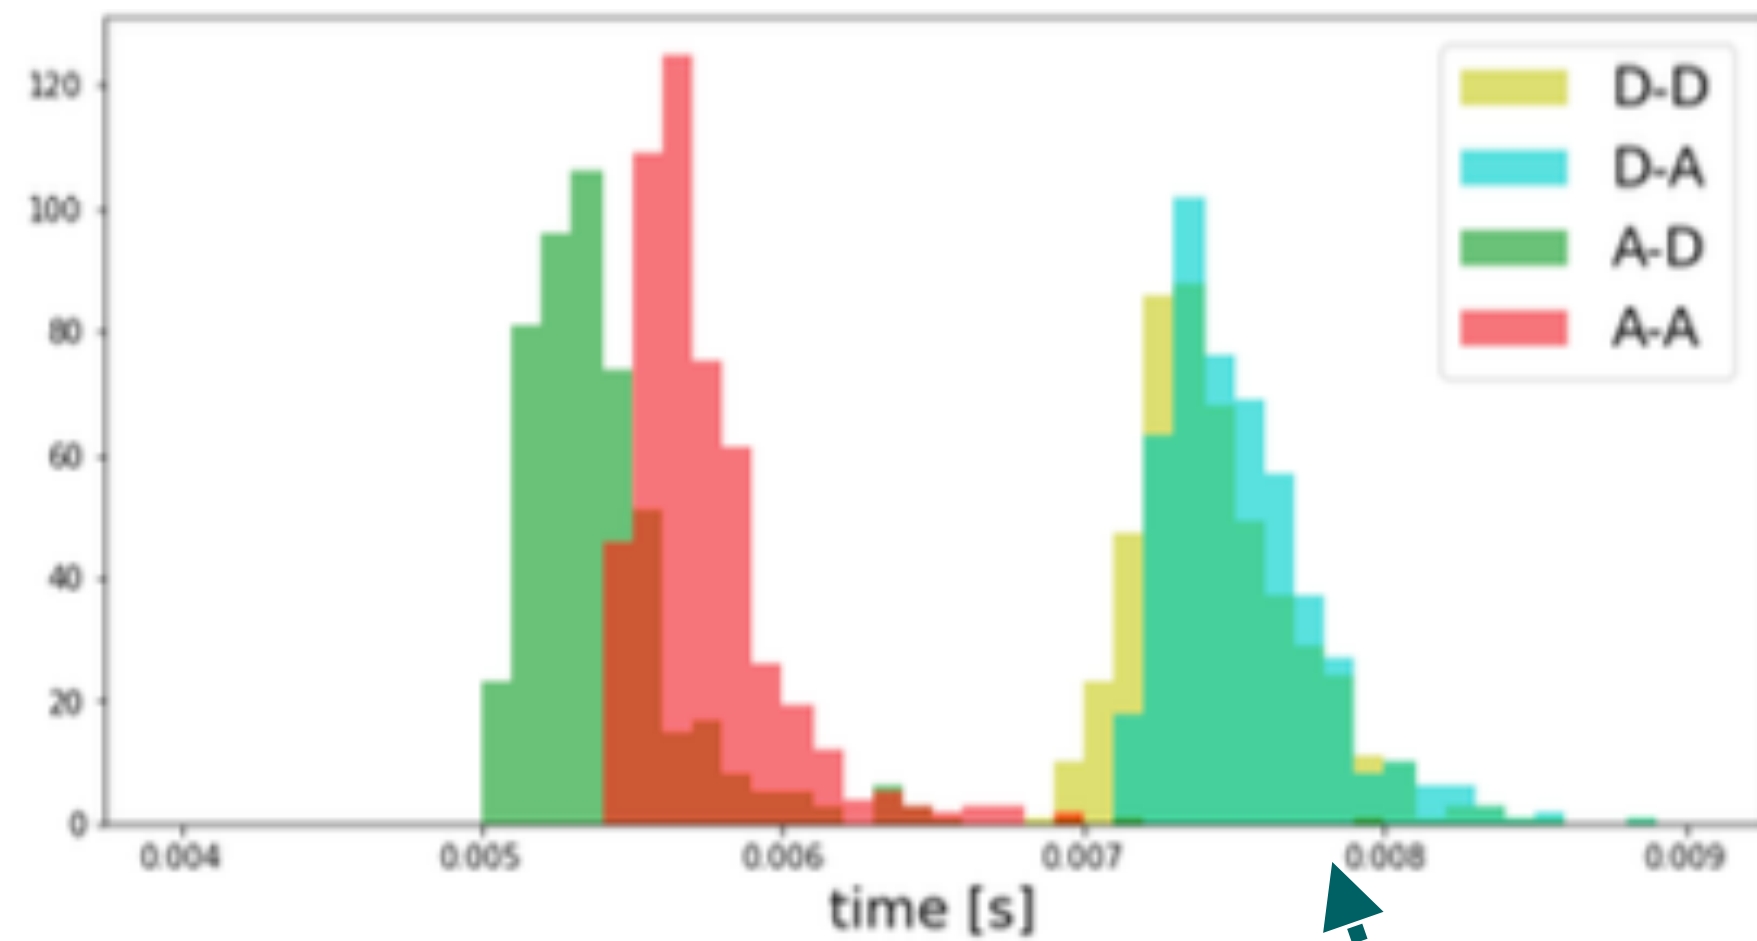
We assume L1 and L2 **cannot** be executed concurrently.

Histogram of execution times indicates differences in performance.



Clustering into performance classes

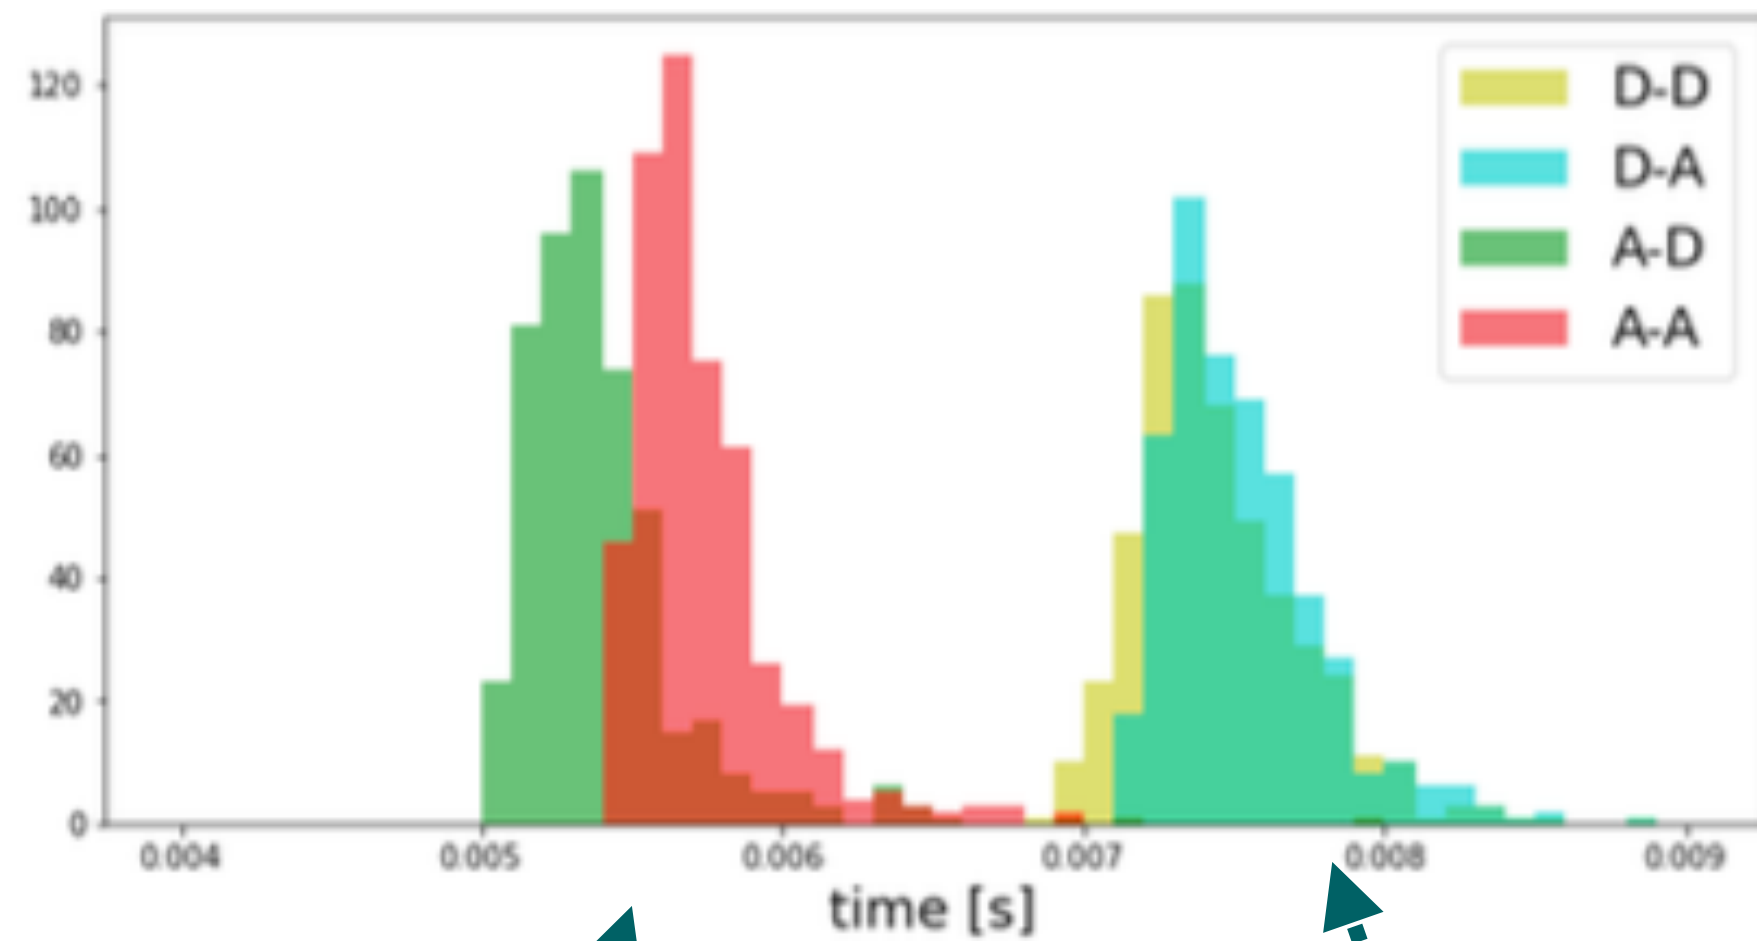
Histogram of execution times



Same cluster

Clustering into performance classes

Histogram of execution times

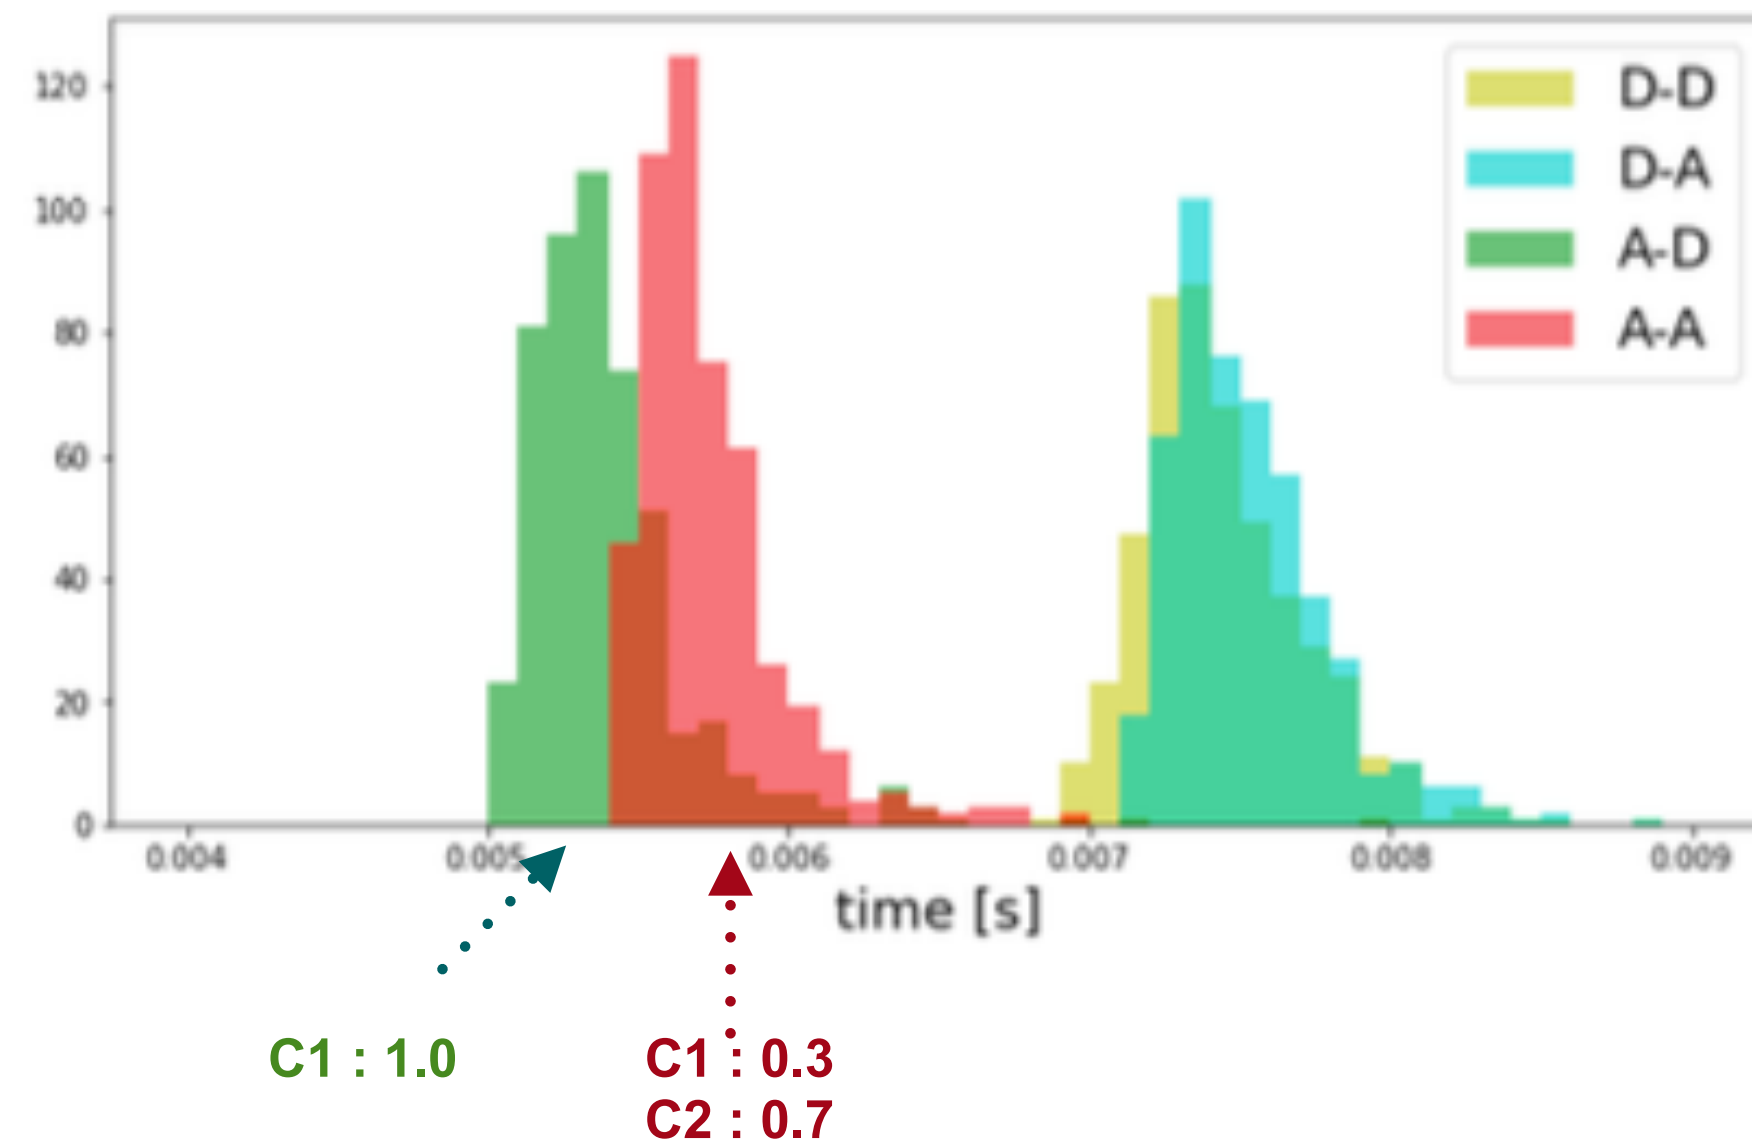


Same cluster?!

Same cluster

Clustering into performance classes

Histogram of execution times



Soft clustering

An implementation can be assigned to more than one cluster.

The relative score indicates the confidence of assignment to a particular cluster.

Cluster	Algorithm	Relative Score
C1	AD	1.0
	AA	0.3
C2	AA	0.7
	DD	0.3
	DA	0.3
C3	DD	0.7
	DA	0.6
C4	DA	0.1

Clustering into performance classes

Consider another example with three math tasks:

Scientific Code {

Math Task 1

Math Task 2

Math Task 3

}

Cluster	Algorithm	Relative Score
C_1	algDDA	1.0
	algDAA	0.6
C_2	algDDD	1.0
	algDAA	0.4
C_3	algADA	1.0
	algADD	1.0
	algDAD	0.7
C_4	algAAA	1.0
	algDAD	0.3
C_5	algAAD	1.0

Clustering into performance classes

Consider the following example with three math tasks:

Scientific Code {

Math Task 1

Math Task 2

Math Task 3

}

Cluster	Algorithm	Relative Score
C_1	algDDA	1.0
	algDAA	0.6
C_2	algDDD	1.0
	algDAA	0.4
C_3	algADA	1.0
	algADD	1.0
	algDAD	0.7
C_4	algAAA	1.0
	algDAD	0.3
C_5	algAAD	1.0

Applications:

- 1) Energy optimization.

Clustering into performance classes

Consider the following example with three math tasks:

Scientific Code {

Math Task 1

Math Task 2

Math Task 3

}

Cluster	Algorithm	Relative Score
C_1	algDDA	1.0
	algDAA	0.6
C_2	algDDD	1.0
	algDAA	0.4
C_3	algADA	1.0
	algADD	1.0
	algDAD	0.7
C_4	algAAA	1.0
	algDAD	0.3
C_5	algAAD	1.0

Applications:

- 1) Energy optimization.
- 2) Chart fault tolerant policies.

Clustering into performance classes

Consider the following example with three math tasks:

Scientific Code {

Math Task 1

Math Task 2

Math Task 3

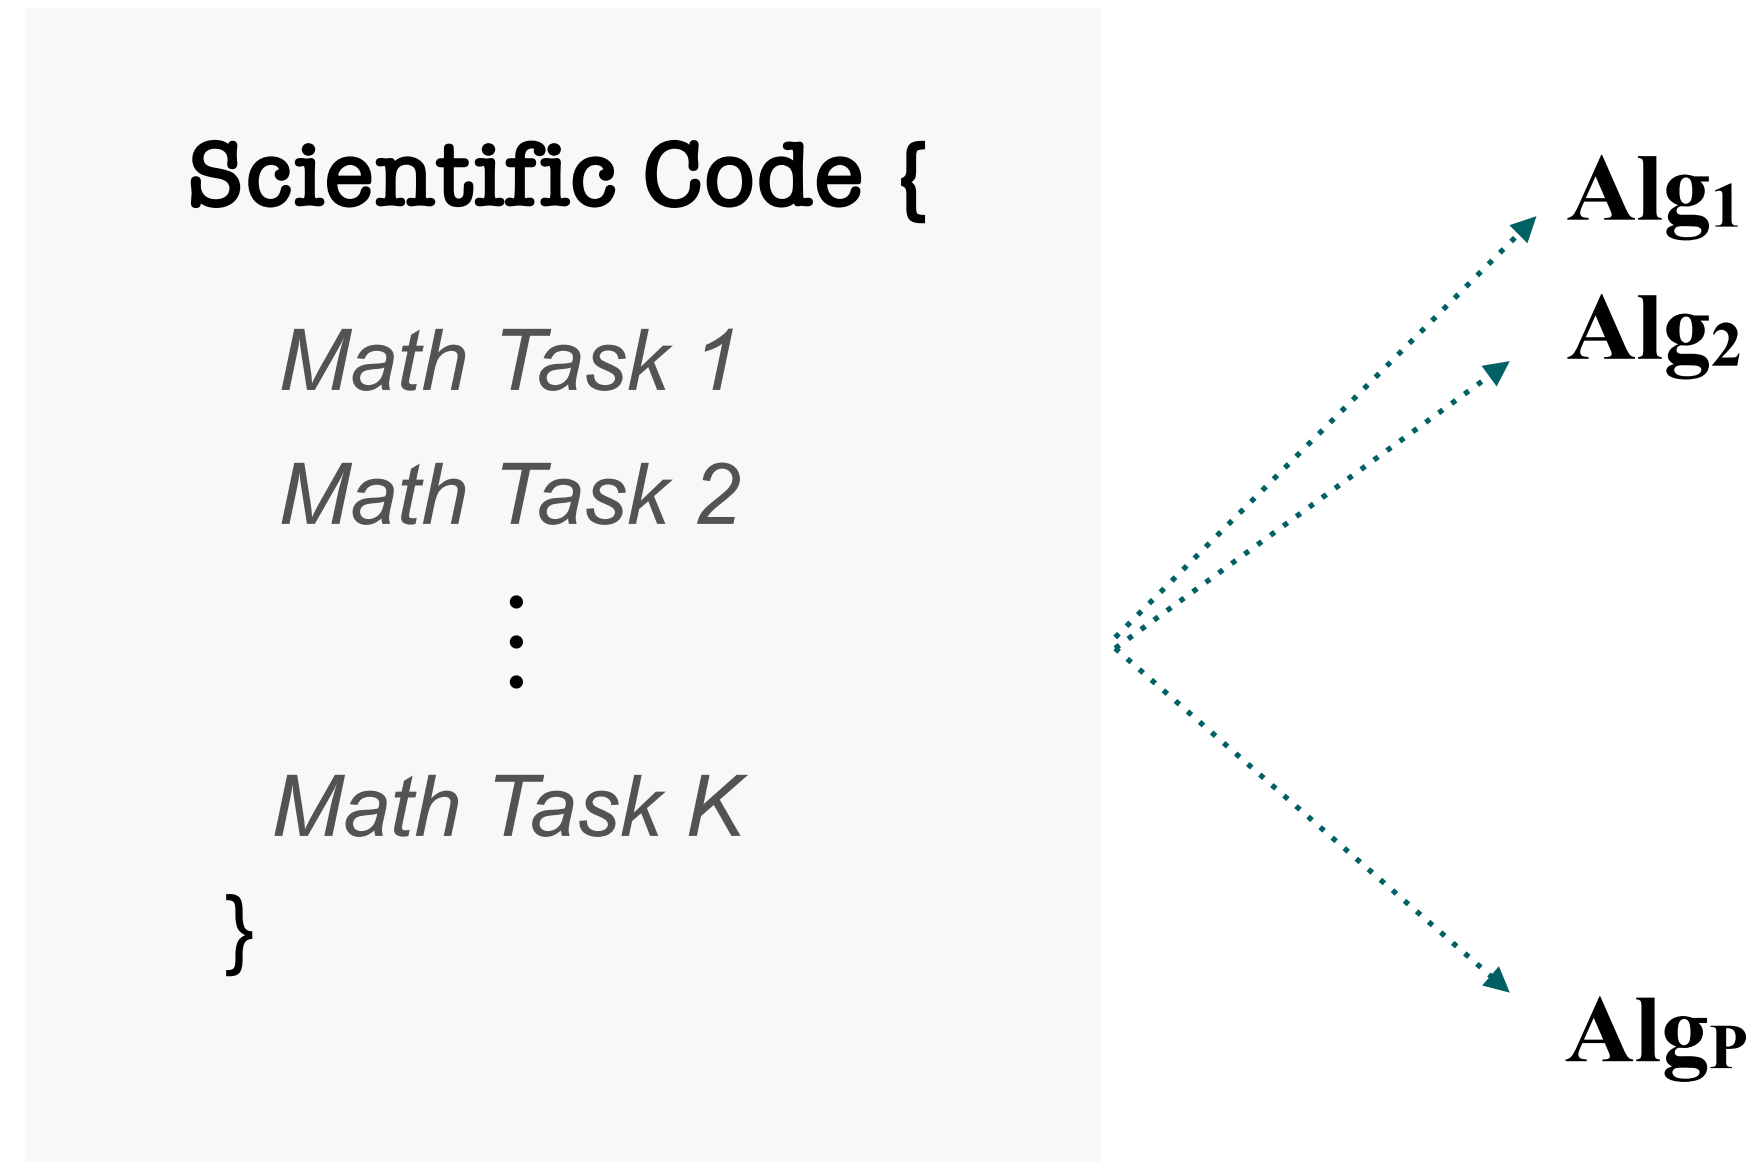
}

Cluster	Algorithm	Relative Score
C_1	algDDA	1.0
	algDAA	0.6
C_2	algDDD	1.0
	algDAA	0.4
C_3	algADA	1.0
	algADD	1.0
	algDAD	0.7
C_4	algAAA	1.0
	algDAD	0.3
C_5	algAAD	1.0

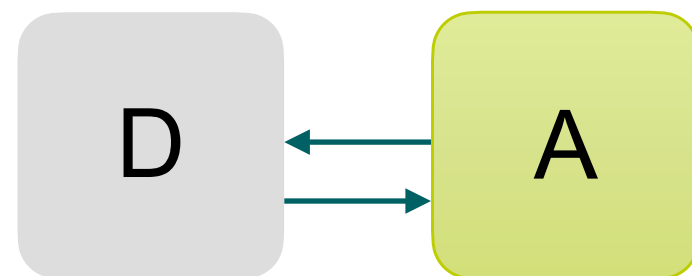
Applications:

- 1) Energy optimization.
- 2) Chart fault tolerant policies.
- 3) Derive performance models for automatic clustering.

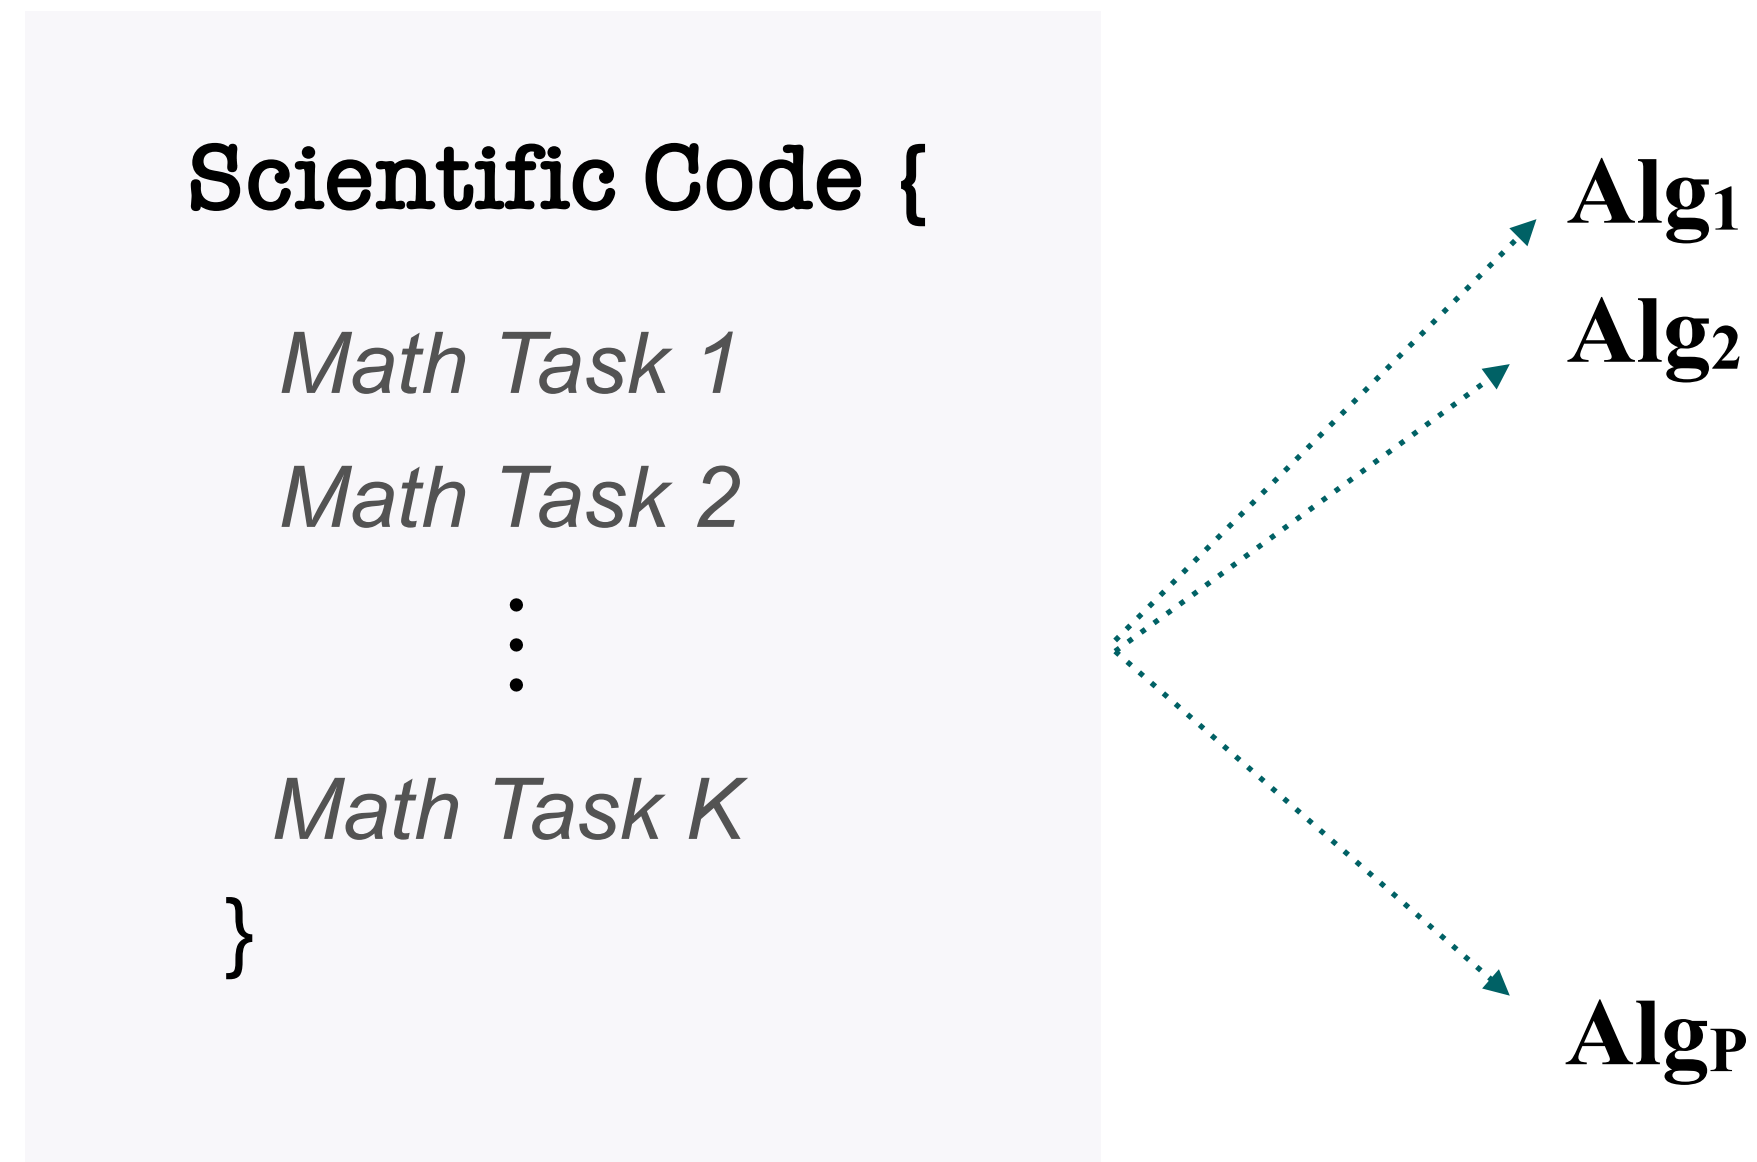
Clustering Methodology



Let Alg₁, Alg₂ ... Alg_p be different implementations of the **Scientific Code**



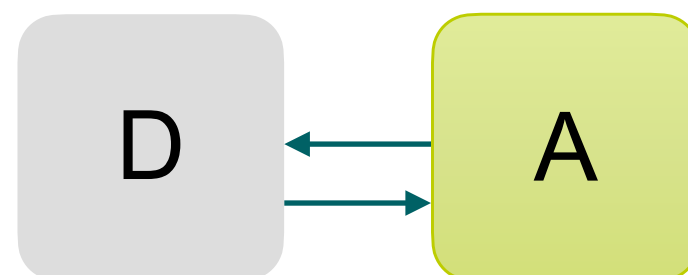
Clustering Methodology



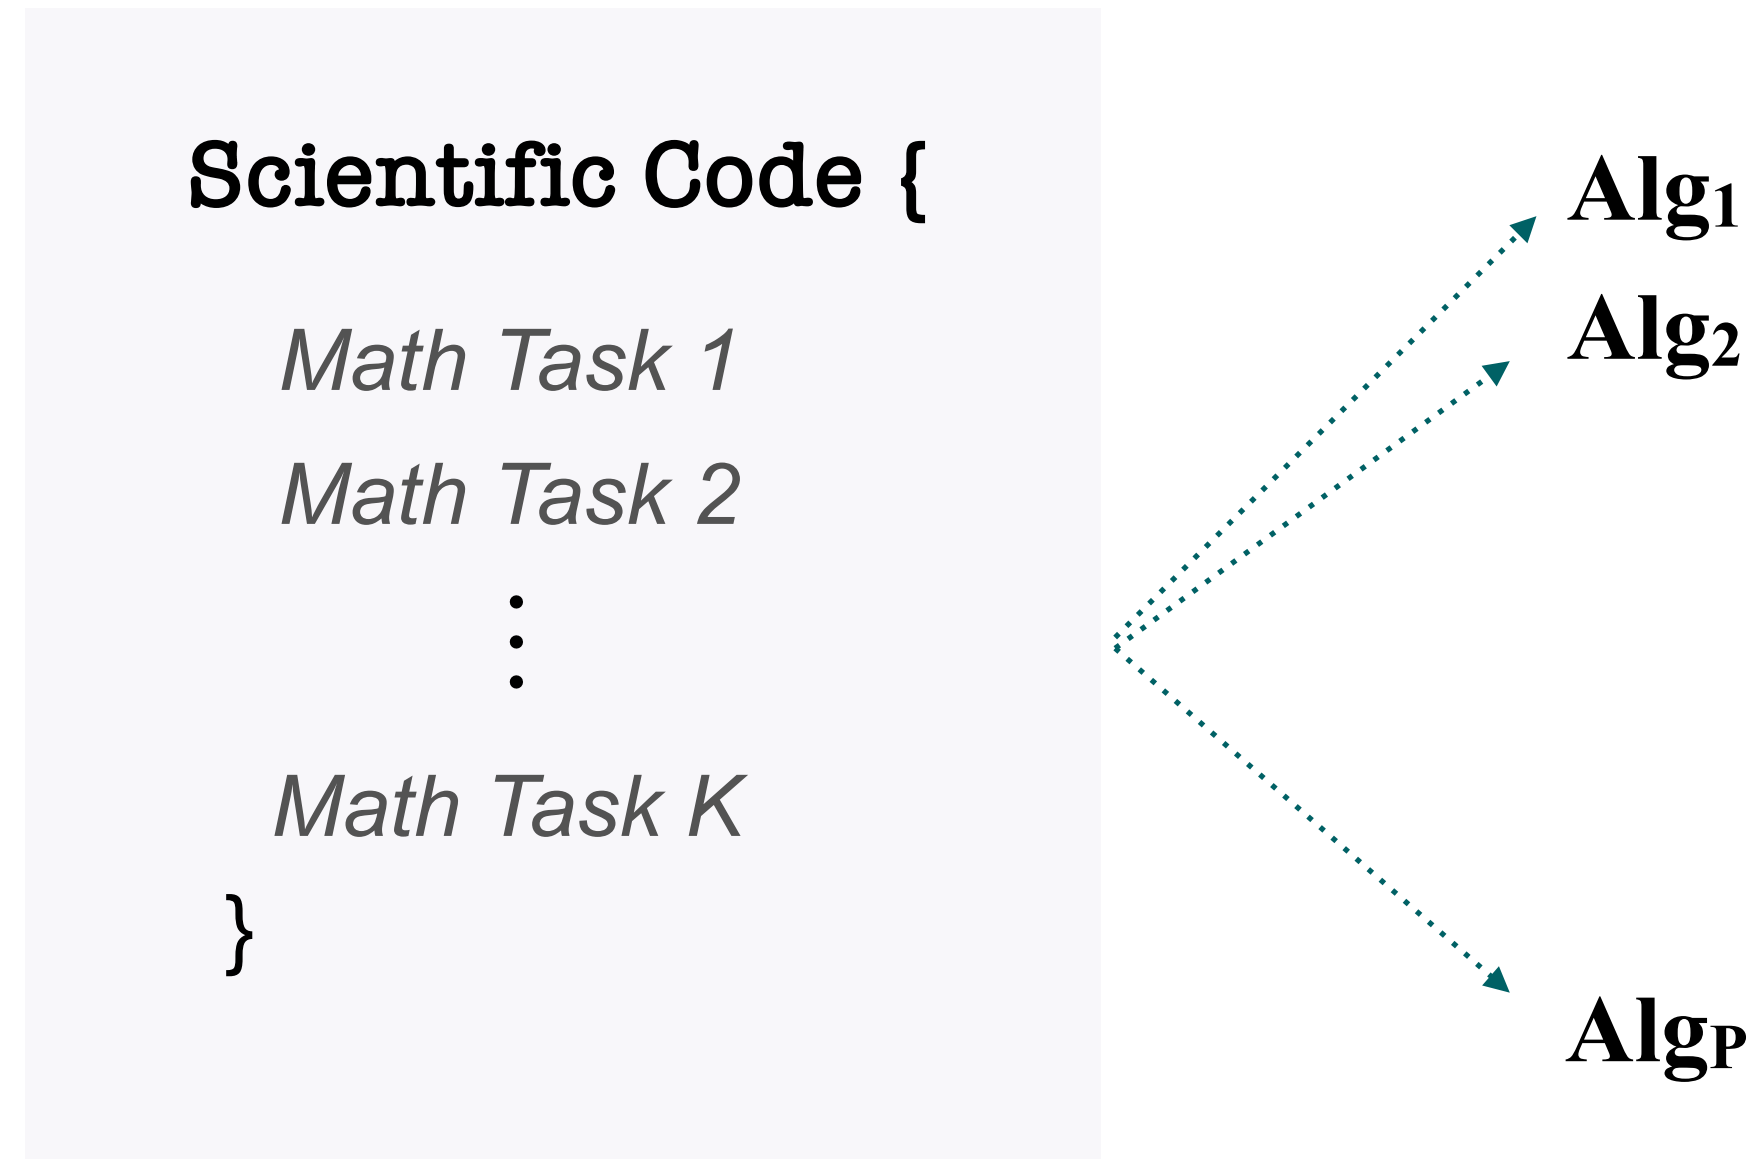
Let Alg₁ , Alg₂ ... Alg_p be different implementations of the **Scientific Code**

Steps:

- 1) Measure each algorithm N times



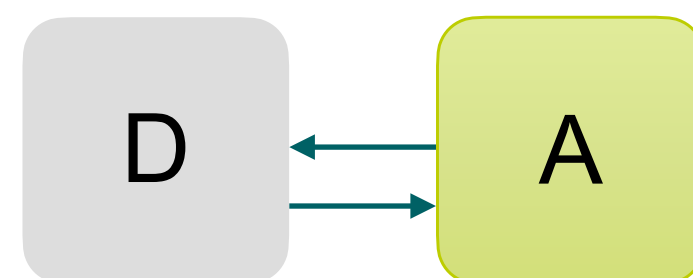
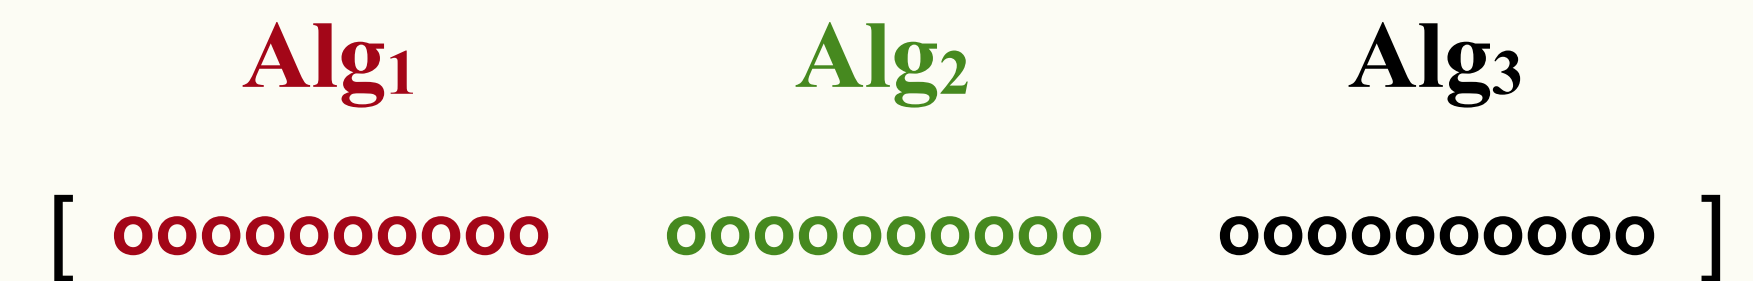
Clustering Methodology



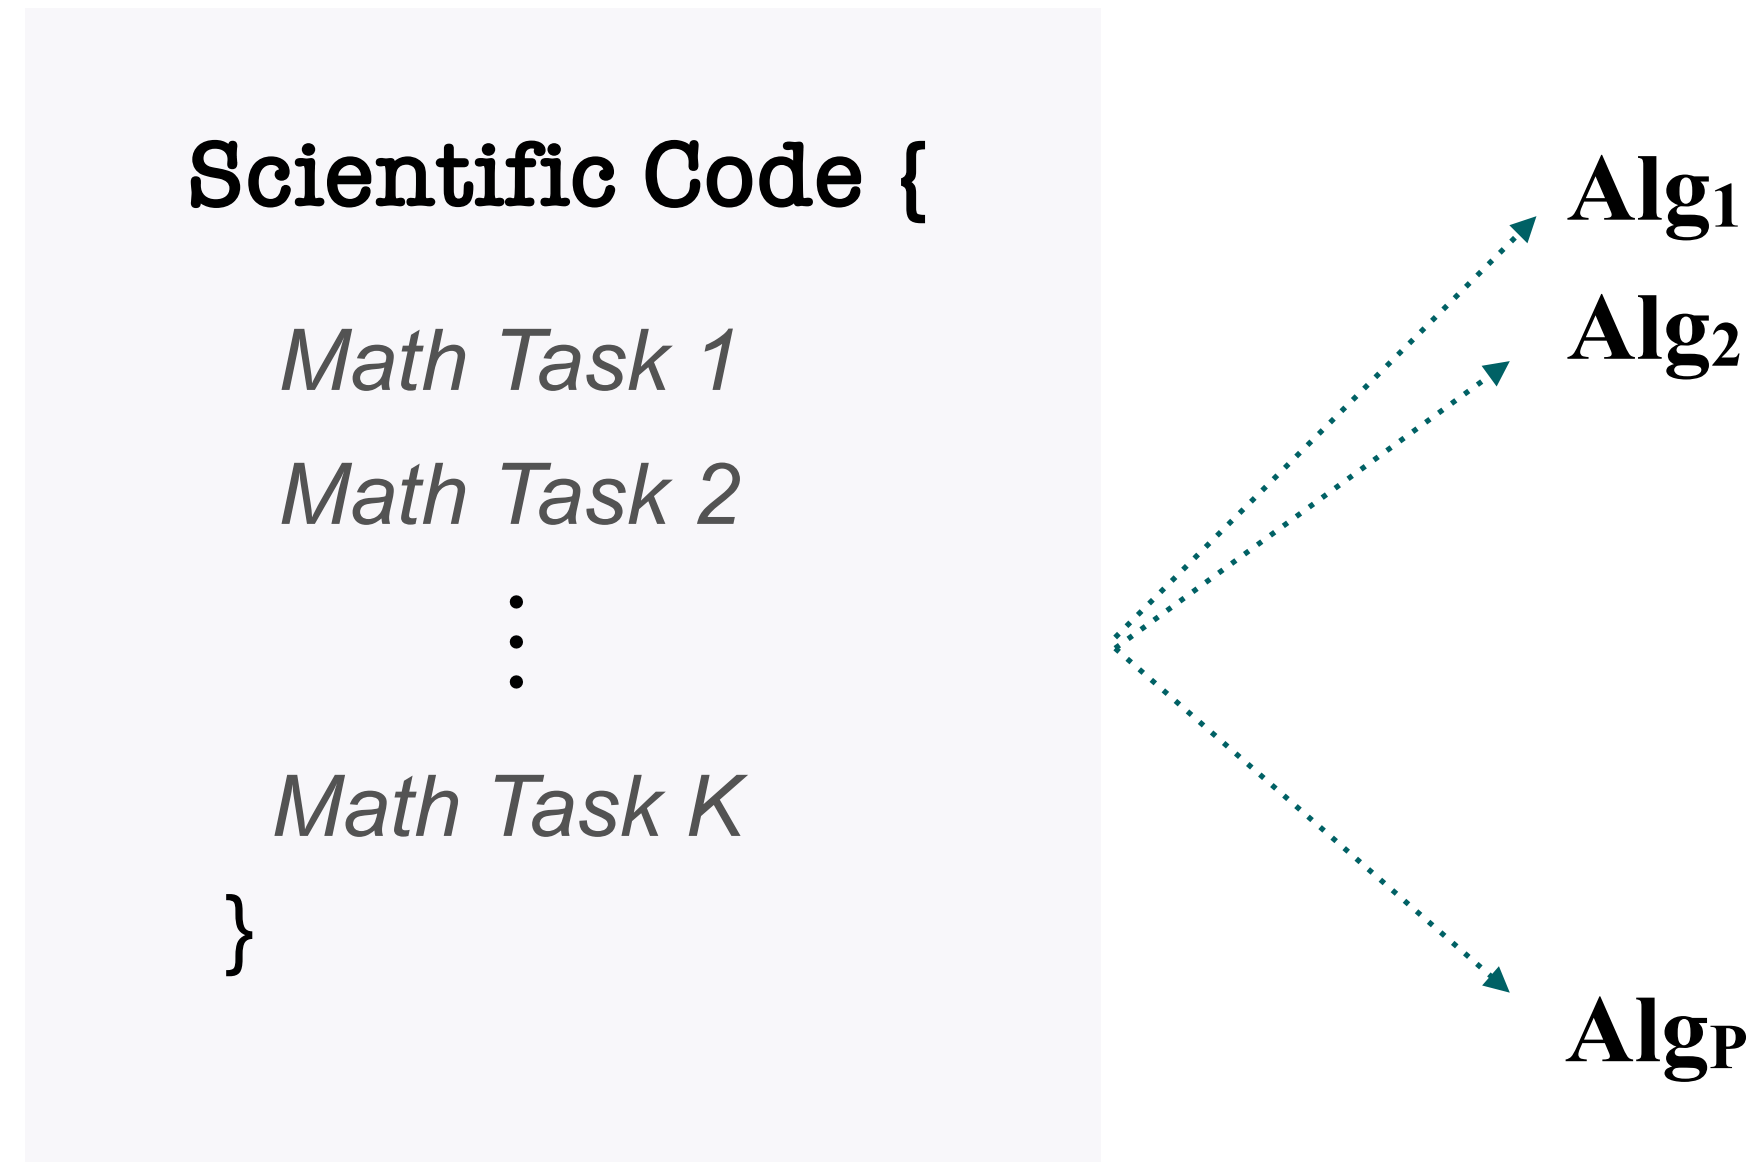
Steps:

- 1) Measure each algorithm N times

Ensure measurements are invariant system noise.



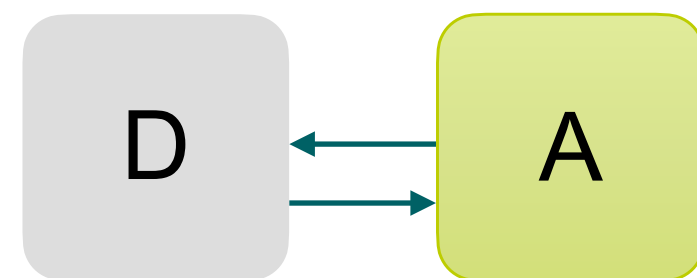
Clustering Methodology



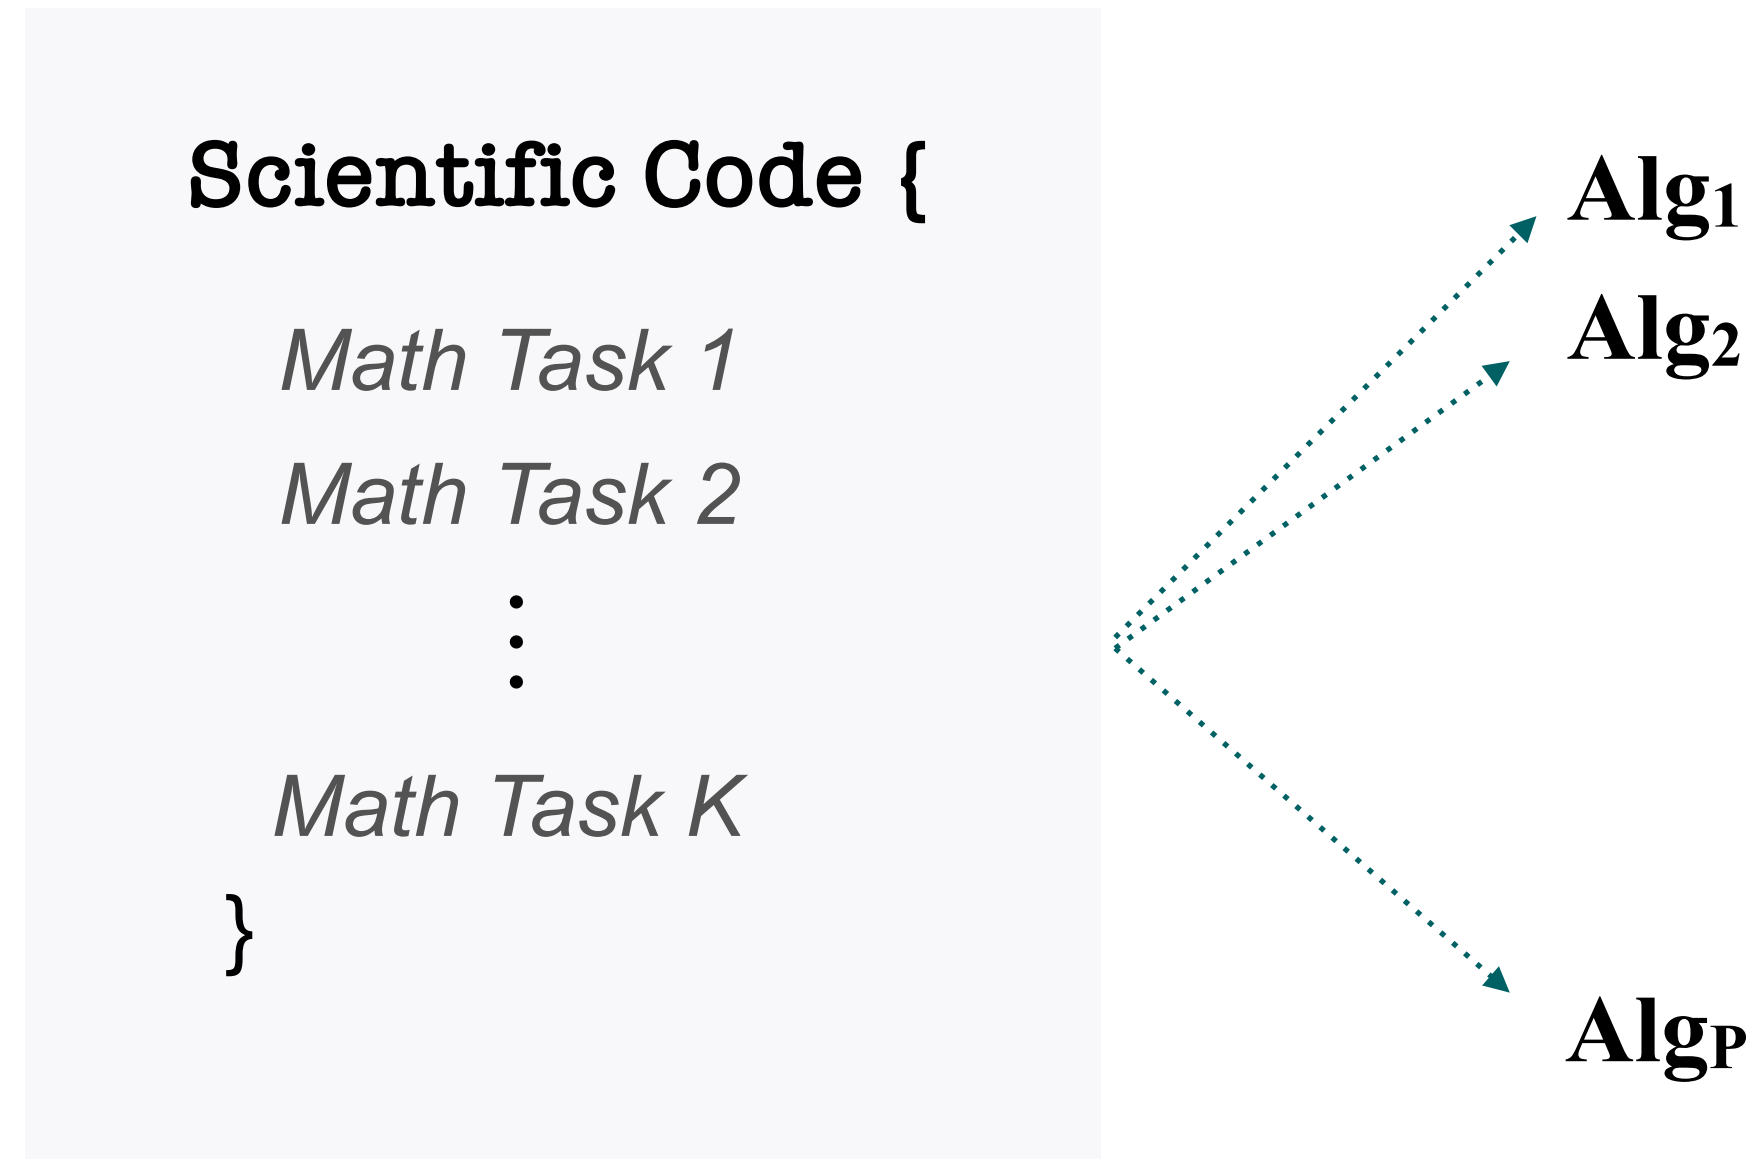
Steps:

- 1) Measure each algorithm N times

Ensure measurements are invariant system noise.

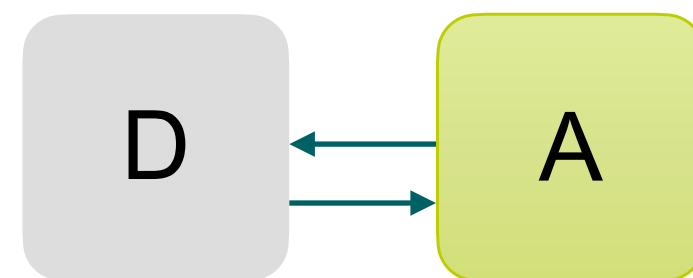


Clustering Methodology



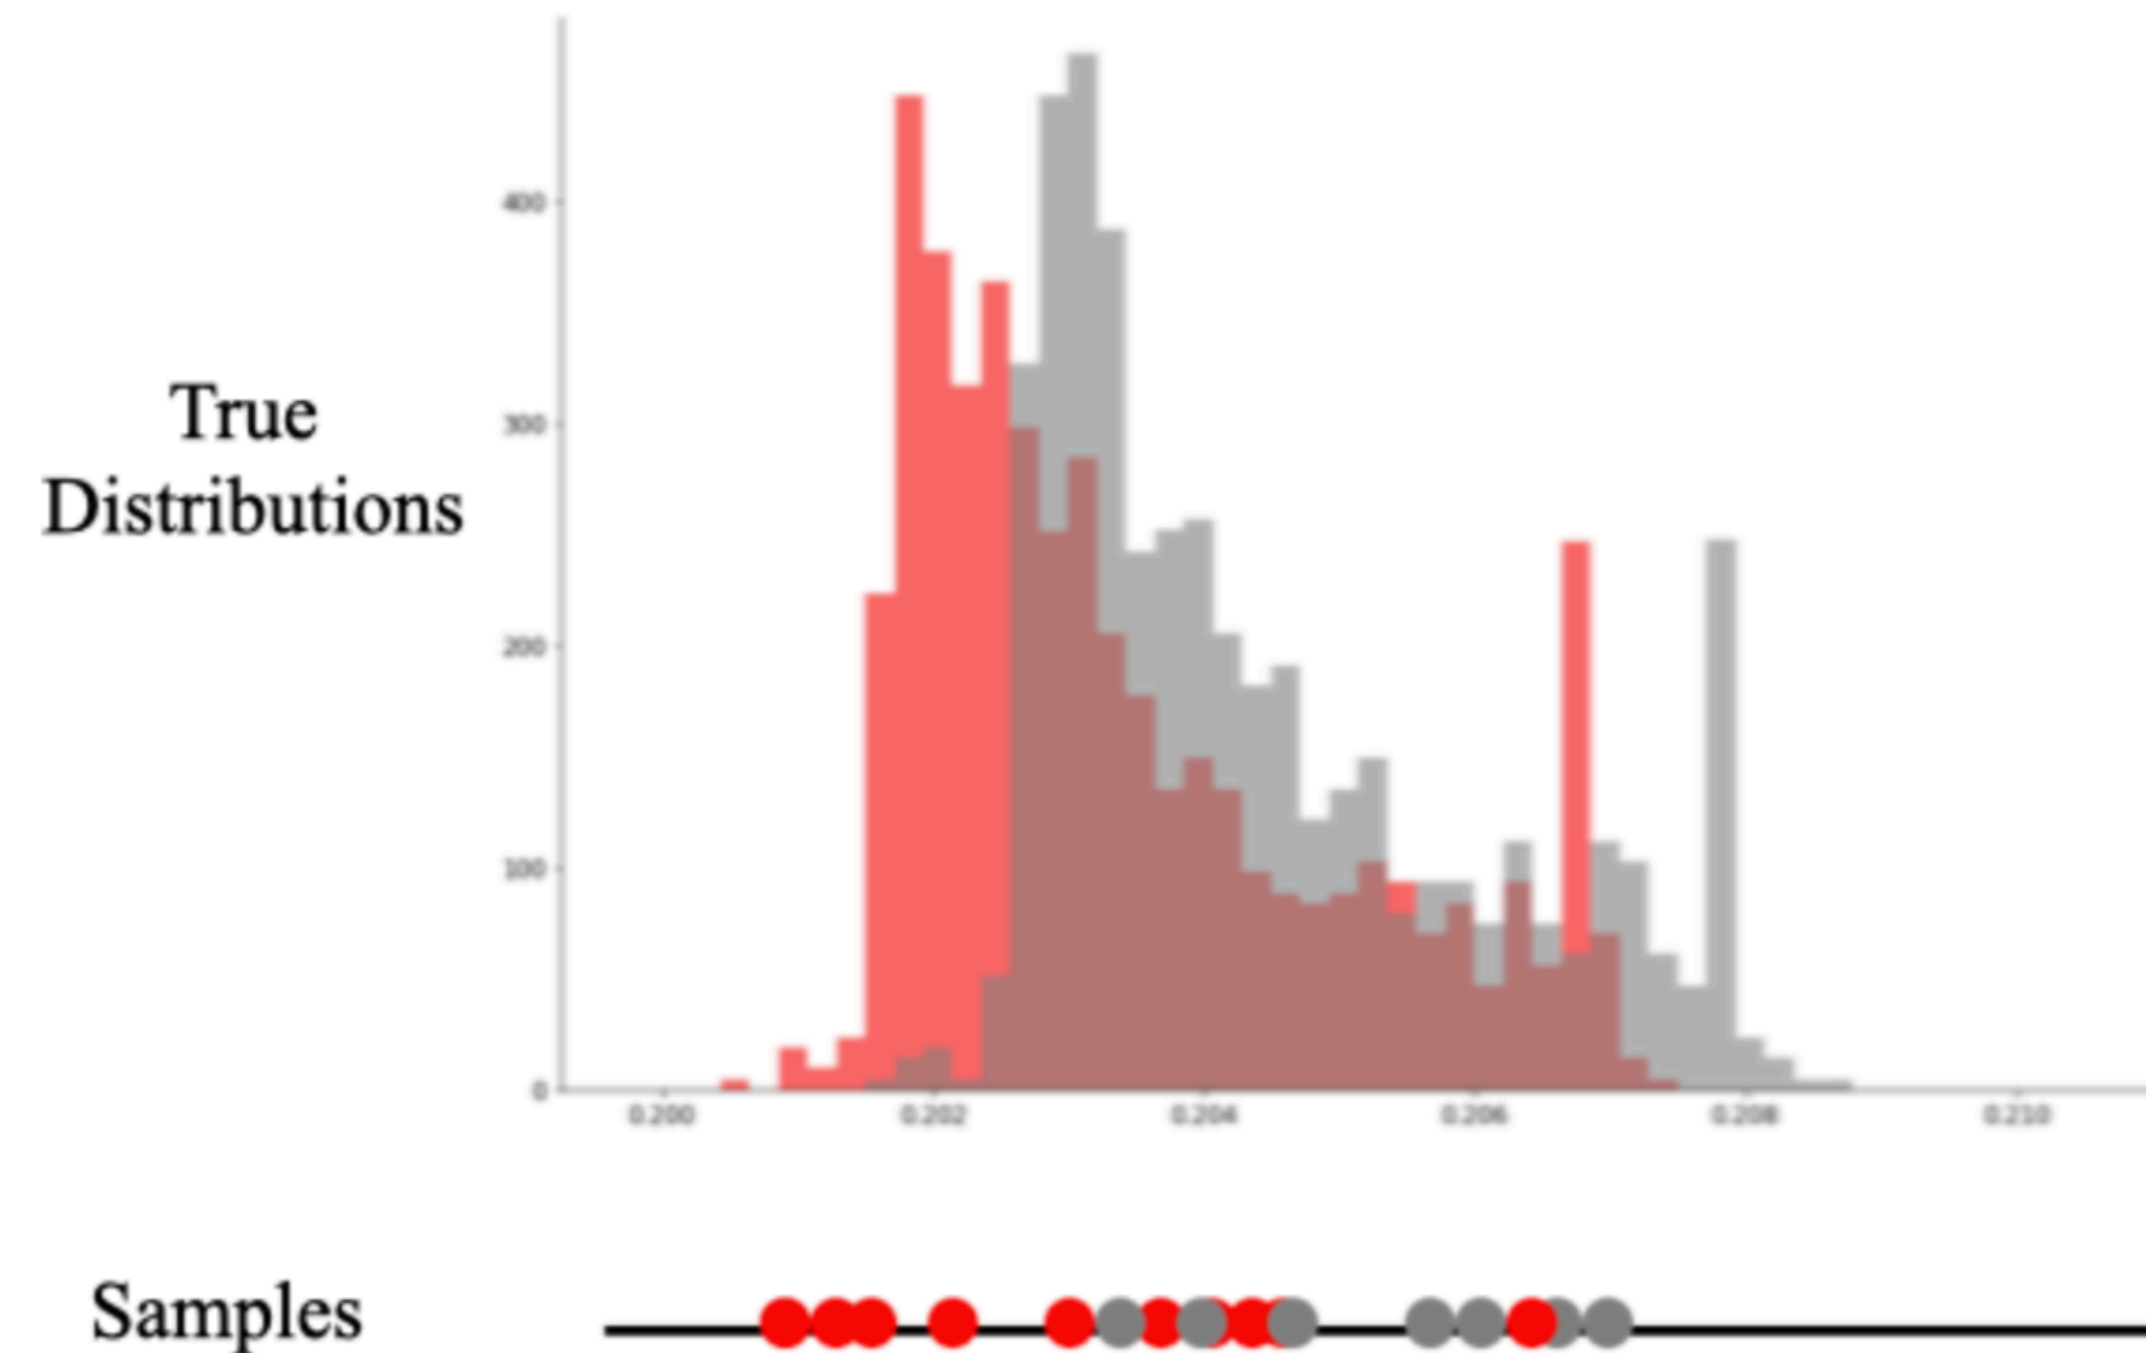
Steps:

- 1) Measure each algorithm N times.
- 2) Compare adjacent pairs of algorithms and sort them.



Clustering Methodology

Comparing two algorithms:



True distributions of the algorithms are not available.

We have just a snapshot of the true distribution.

Clustering Methodology

Comparing two algorithms by bootstrapping Distribution Statistics :

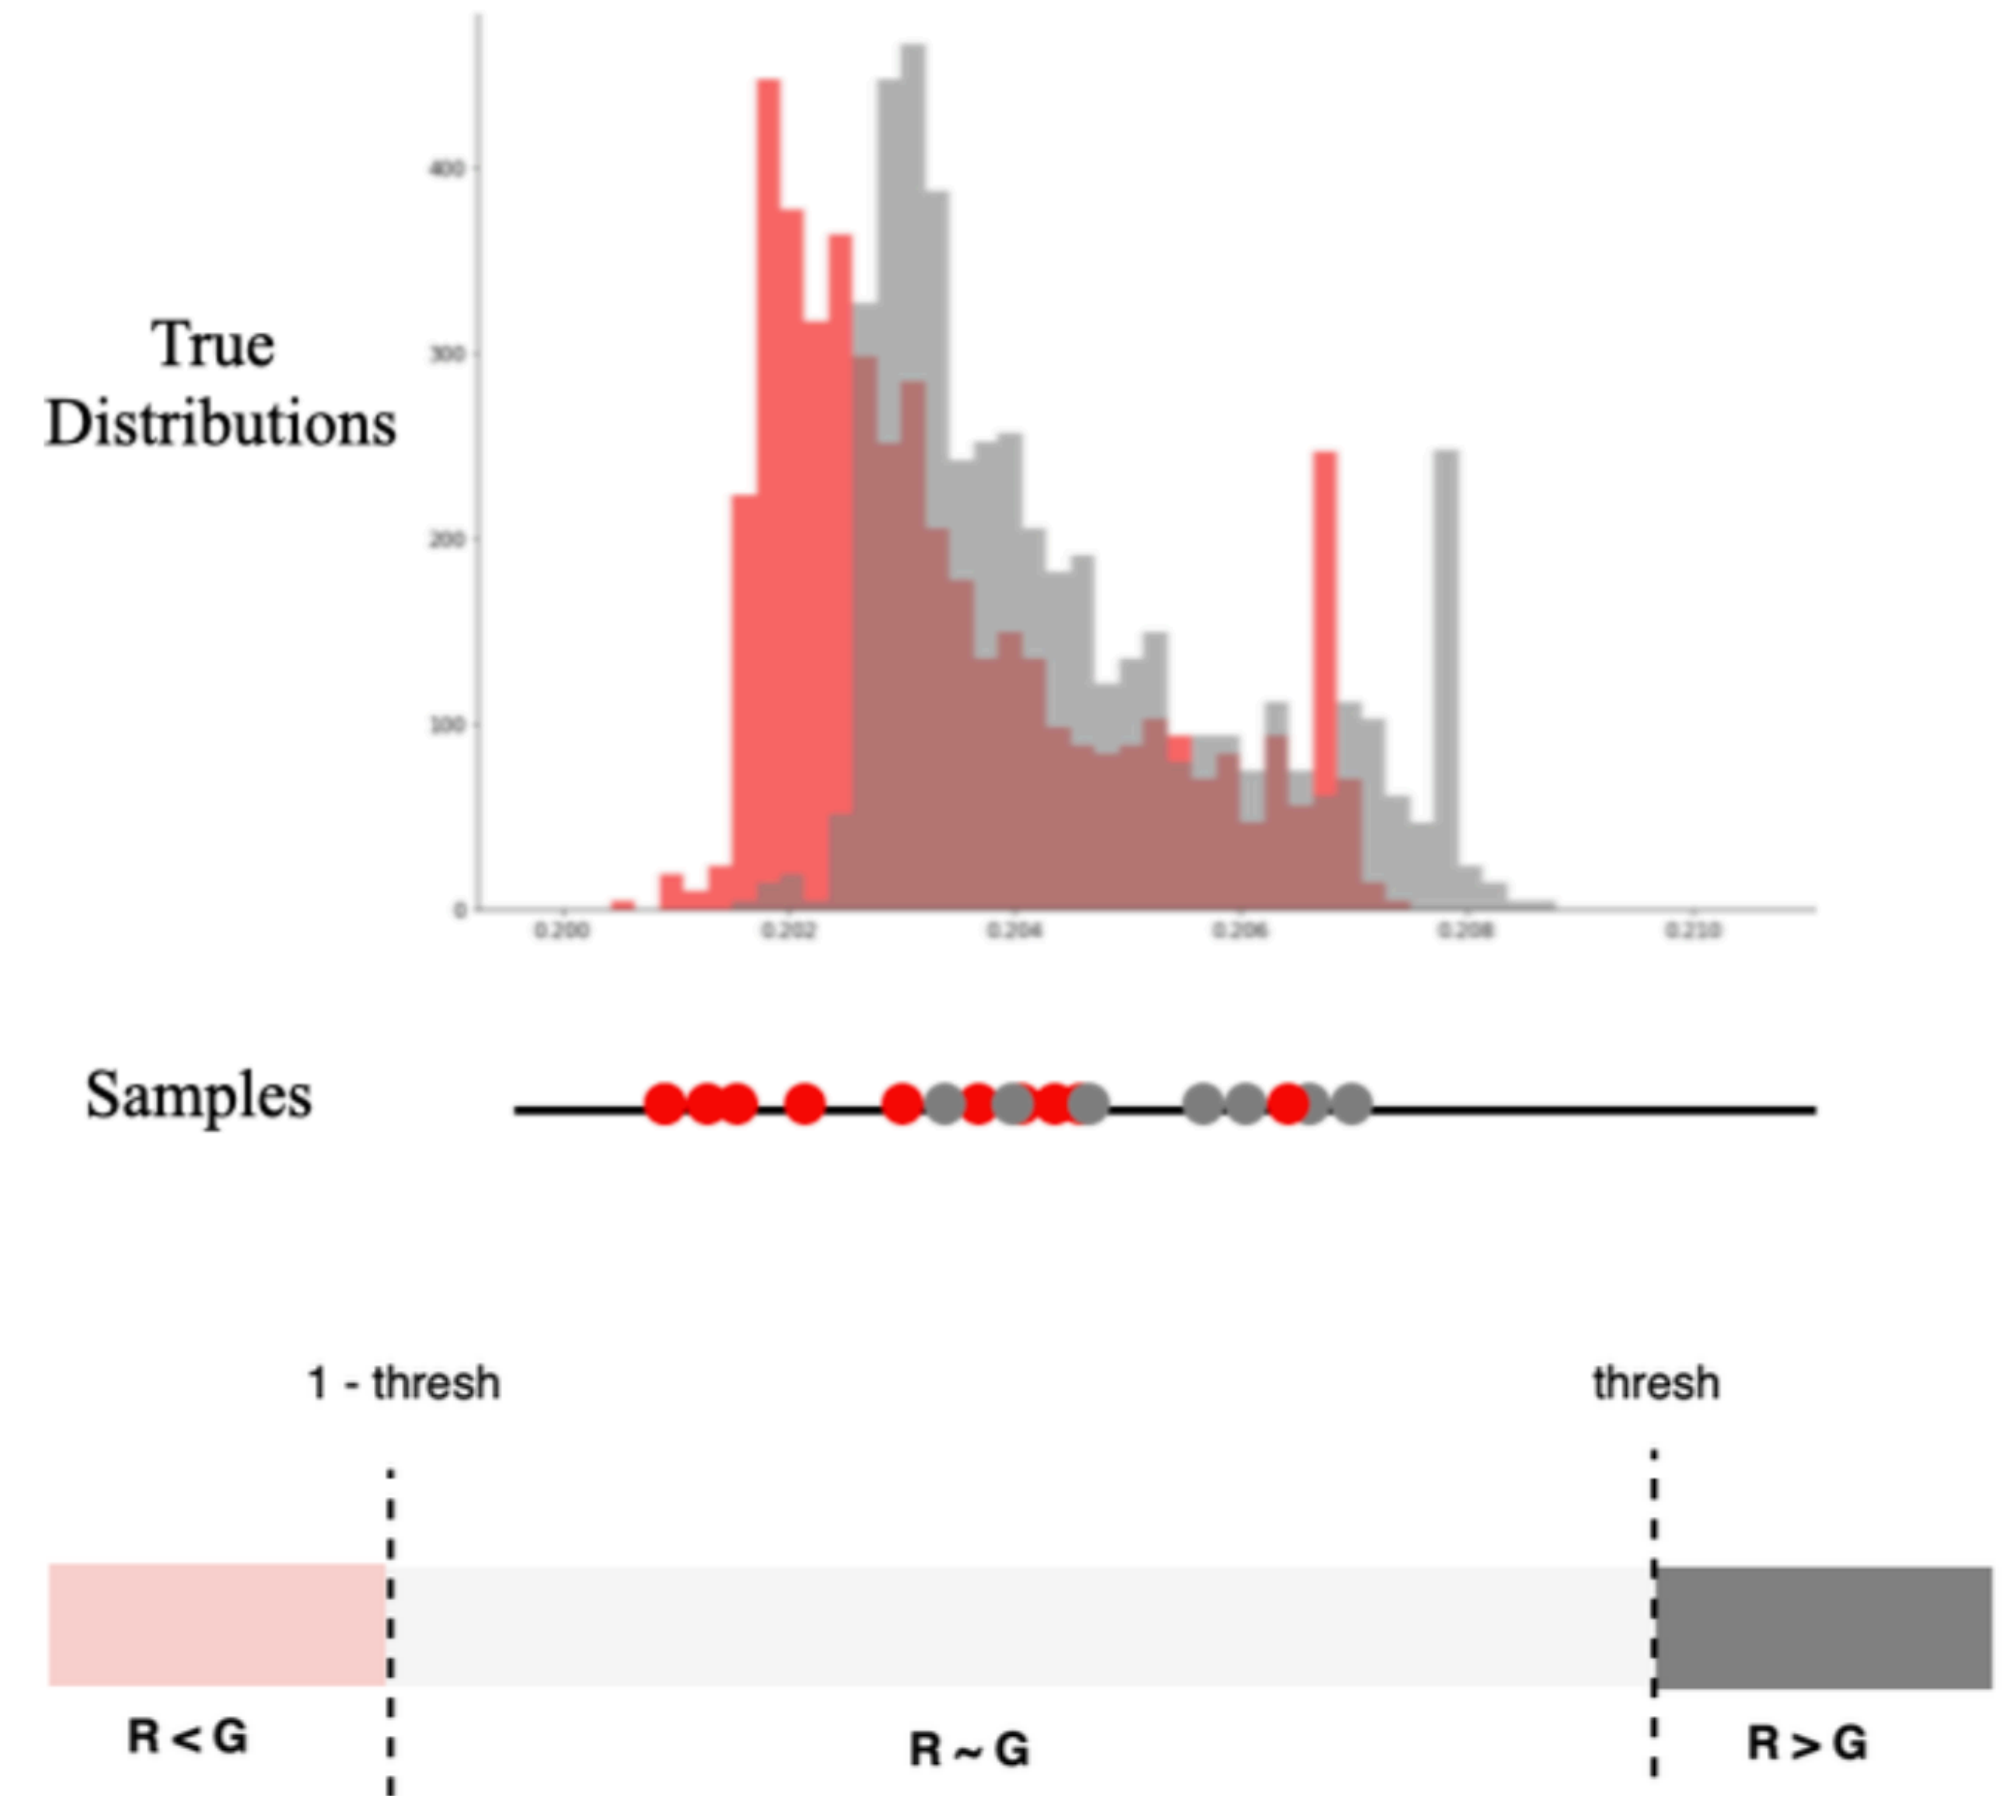
$$\mathbf{R} = \{tR_1, tR_2, \dots, tR_n\}$$

$$\mathbf{G} = \{tG_1, tG_2, \dots, tG_n\}$$

$$\text{Compare}(\mathbf{R}, \mathbf{G}) \in \{<, >, \sim\}$$

$$\text{Sample } R^s \subset \mathbf{R}, G^s \subset \mathbf{G}$$

$$\text{Compare}(\mathbf{R}, \mathbf{G}) = \begin{cases} \mathbf{R} > \mathbf{G} & \text{if } \frac{\#(\text{stat}(R^s) > \text{stat}(G^s))}{M} > \text{thresh} \\ \mathbf{R} < \mathbf{G} & \text{if } \frac{\#(\text{stat}(R^s) > \text{stat}(G^s))}{M} < 1 - \text{thresh} \\ \mathbf{R} \sim \mathbf{G} & \text{otherwise} \end{cases}$$



Sorting the algorithms using the three-way comparison:

- 1) Initialise algorithms with consecutive ranks.
- 2) Compare adjacent pairs of algorithms and swap their ranks if an algorithm occurring later in the sequence performs better than the one occurring earlier.
- 3) Merge ranks if the two algorithms are performance equivalent.

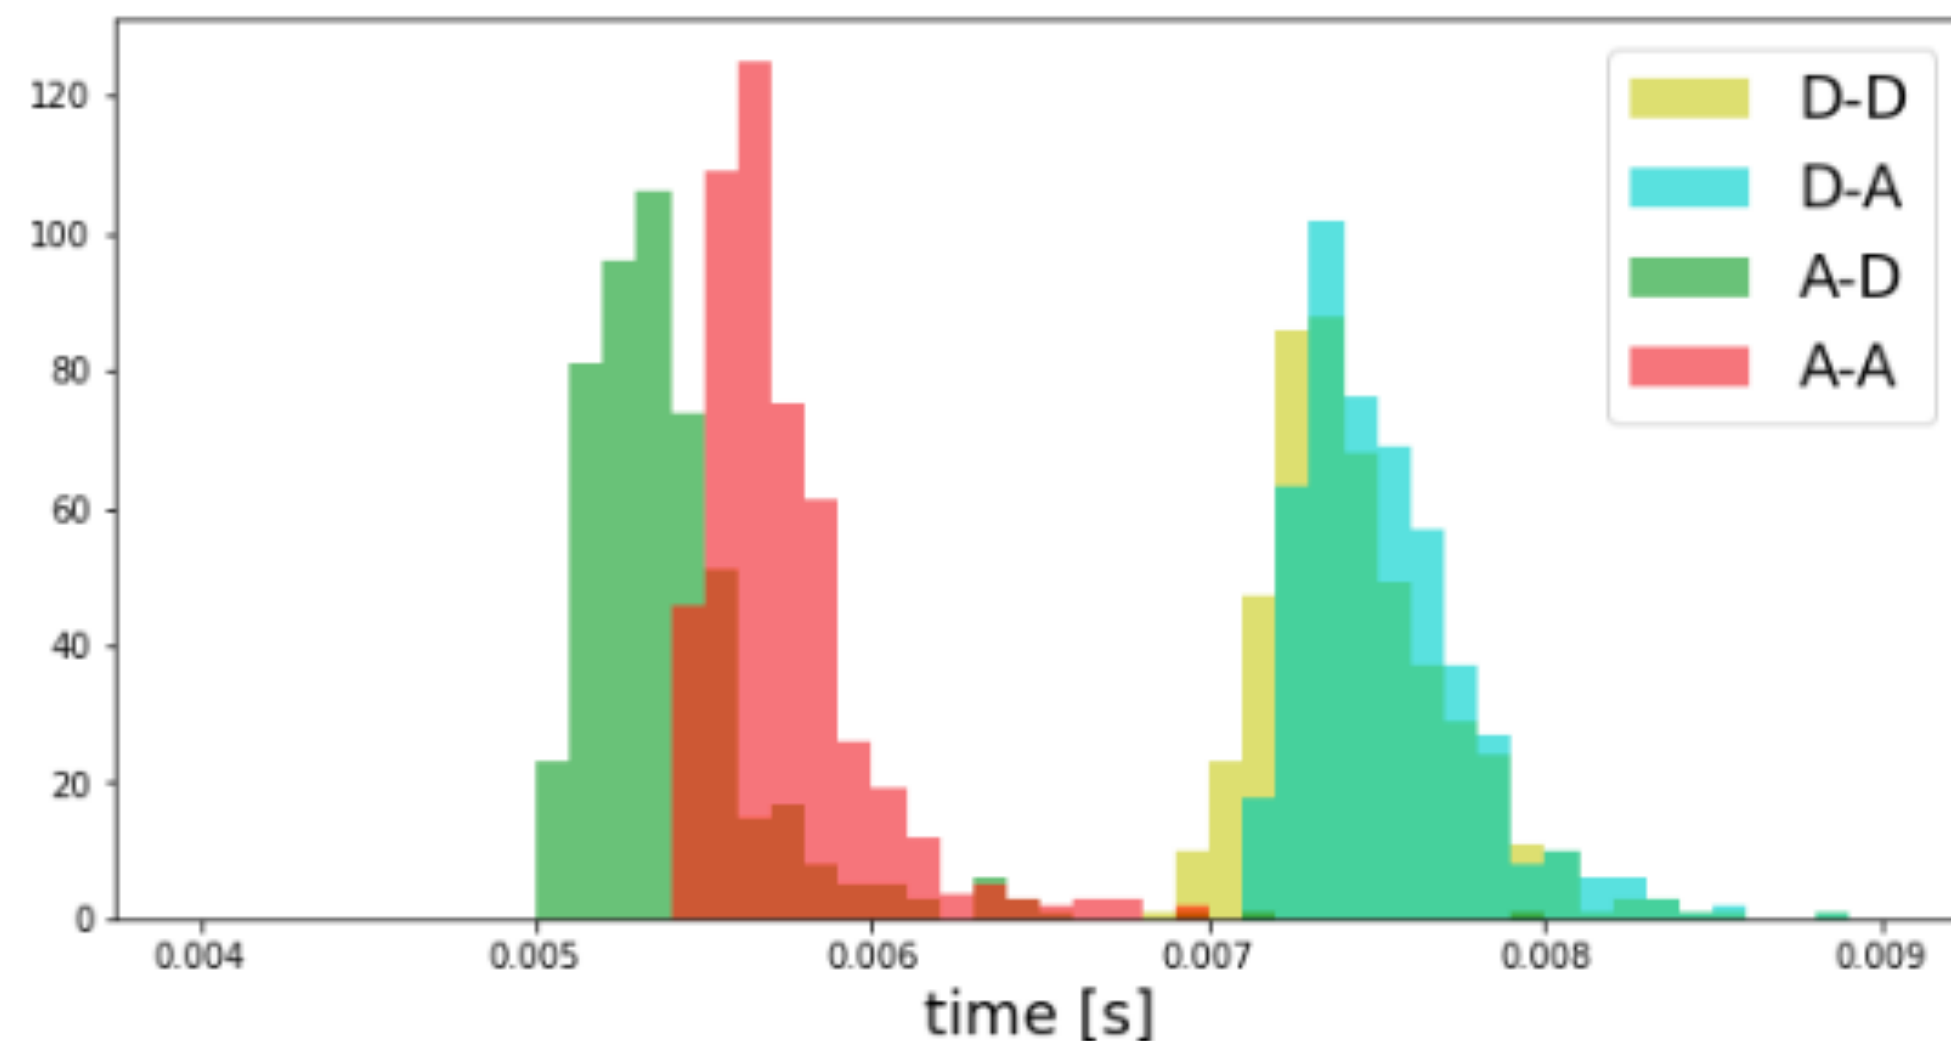
Algorithm	DD	AA	DA	AD
Rank	1	2	3	4
	"DD" < "AA"			
Algorithm	AA	DD	DA	AD
Rank	1	2	3	4
		"DD" ~ "DA"		
Algorithm	AA	DD	DA	AD
Rank	1	2	2	3
			"DA" < "AD"	
Algorithm	AA	DD	AD	DA
Rank	1	2	2	2
		"DD" < "AD"		
Algorithm	AA	AD	DD	DA
Rank	1	2	3	3
	"AA" < "AD"			
Algorithm	AD	AA	DD	DA
Final Rank	1	2	3	3

Clustering Methodology

Sorting the algorithms using the three-way comparison:

- 1) Initialise algorithms with consecutive ranks.
- 2) Compare adjacent pairs of algorithms and sort them.
- 3) Merge ranks if the two algorithms are performance equivalent.

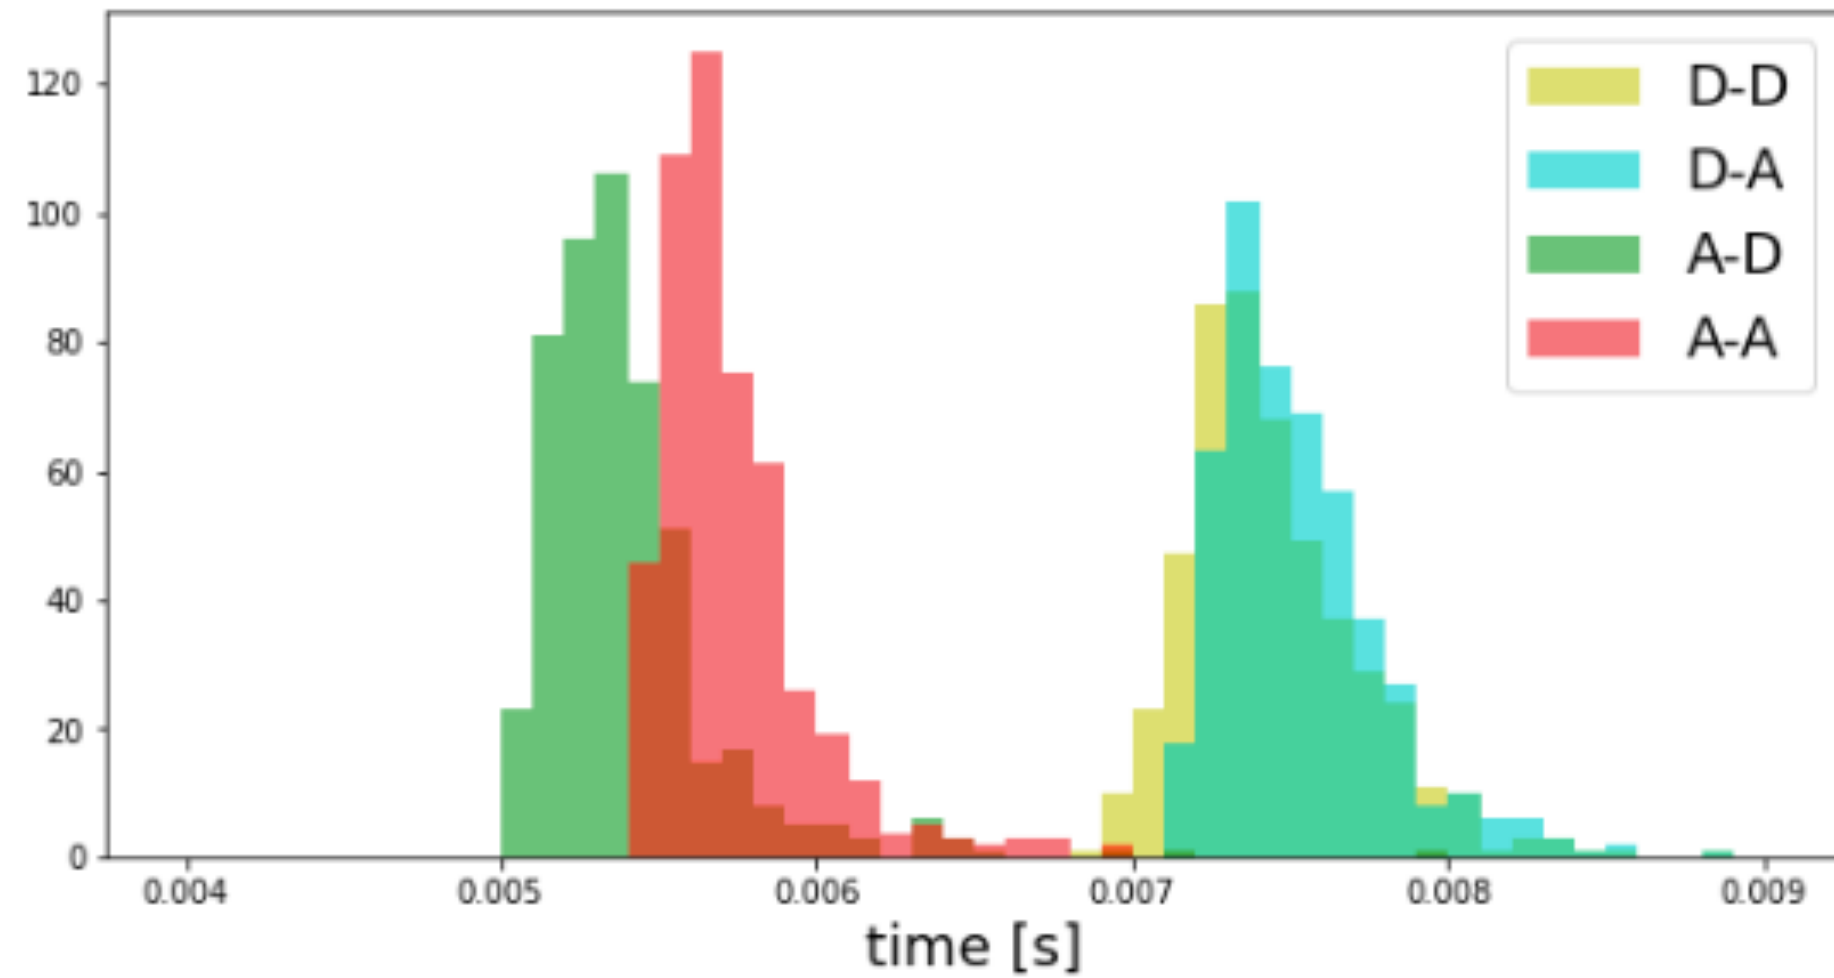
Example : {DD, DA, AD, AA}



Algorithm	DD	AA	DA	AD
Rank	1	2	3	4
	"DD" < "AA"			
Algorithm	AA	DD	DA	AD
Rank	1	2	3	4
	"DD" ~ "DA"			
Algorithm	AA	DD	DA	AD
Rank	1	2	2	3
			"DA" < "AD"	
Algorithm	AA	DD	AD	DA
Rank	1	2	2	2
			"DD" < "AD"	
Algorithm	AA	AD	DD	DA
Rank	1	2	3	3
	"AA" < "AD"			
Algorithm	AD	AA	DD	DA
Final Rank	1	2	3	3

Clustering Methodology

Example : {DD, DA, AD, AA}



Steps:

- 1) Measure each algorithm N times.
- 2) Compare adjacent pairs of algorithms and sort them.
- 3) Repeat step (2) K times to compute relative scores.

Algorithm	DD	AA	DA	AD
Rank	1	2	3	4
"DD" < "AA"				
Algorithm	AA	DD	DA	AD
Rank	1	2	3	4
"DD" ~ "DA"				
Algorithm	AA	DD	DA	AD
Rank	1	2	2	3
"DA" < "AD"				
Algorithm	AA	DD	AD	DA
Rank	1	2	2	2
"DD" < "AD"				
Algorithm	AA	AD	DD	DA
Rank	1	2	3	3
"AA" < "AD"				
Algorithm	AD	AA	DD	DA
Final Rank	1	2	3	3

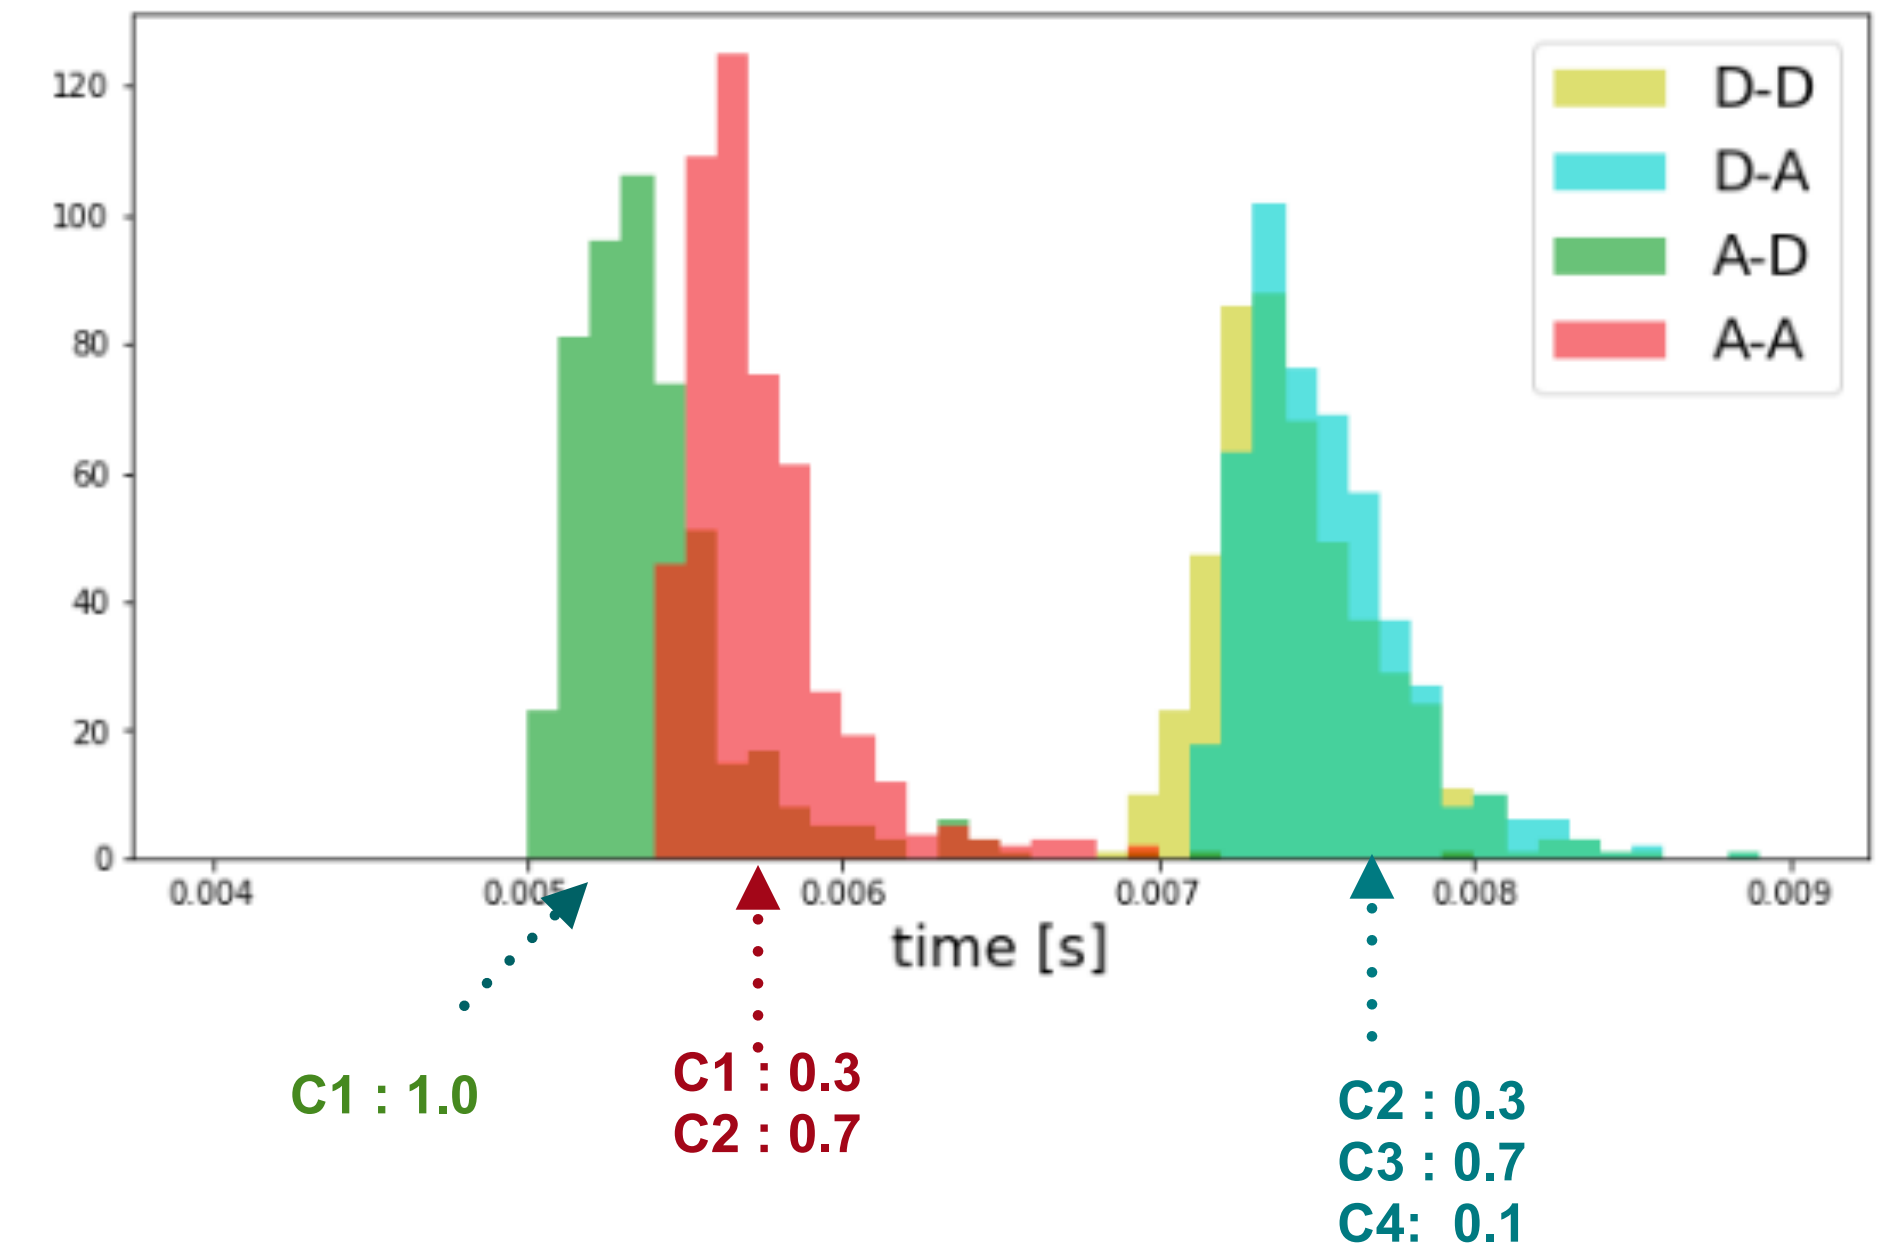
Clustering Methodology

Steps:

- 1) Measure each algorithm N times.
- 2) Compare adjacent pairs of algorithms and sort them.
- 3) Repeat step (2) K times to compute relative scores.

Cluster	Algorithm	Relative Score
C1	AD	1.0
	AA	0.3
C2	AA	0.7
	DD	0.3
	DA	0.3
C3	DD	0.7
	DA	0.6
C4	DA	0.1

Example : {DD, DA, AD, AA}



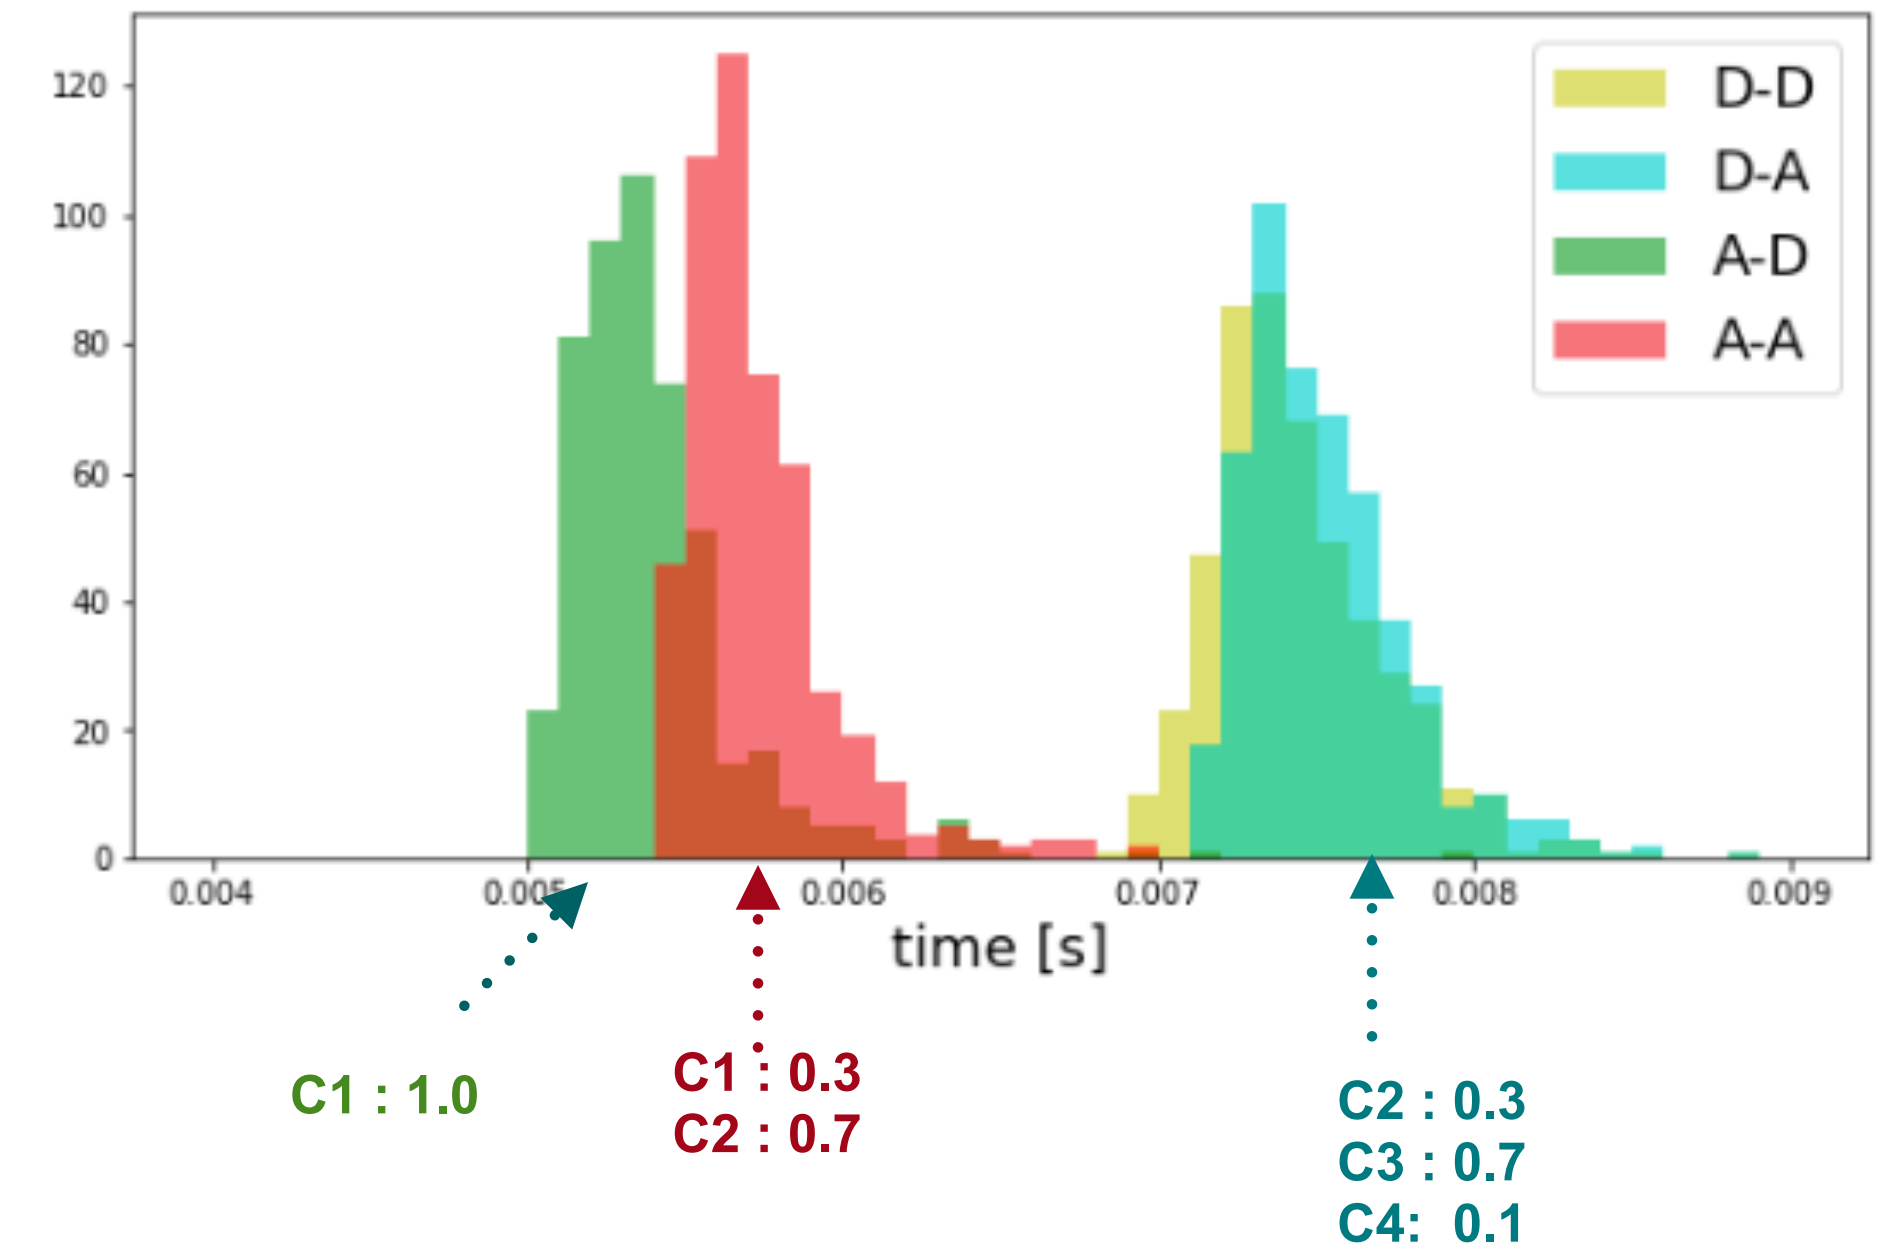
Clustering Methodology

Steps:

- 1) Measure each algorithm N times.
- 2) Compare adjacent pairs of algorithms and sort them.
- 3) Repeat step (2) K times to compute relative scores.

Cluster	Algorithm	Relative Score
C1	AD	1.0
	AA	0.3
C2	AA	0.7
	DD	0.3
	DA	0.3
C3	DD	0.7
	DA	0.6
C4	DA	0.1

Example : {DD, DA, AD, AA}



Each cluster is NOT a SET, but a sequence.
The order of algorithms within a cluster is important!

Summary

- We consider scientific codes that can have many different ways of implementation.
- Cluster the implementations into performance classes.
- Our clustering methodology creates a discrimination among the candidate implementations.
- The information of discrimination can be used as a basis to train neural network based approaches.

Summary

- We consider scientific codes that can have many different ways of implementation.
- Cluster the implementations into performance classes.
- Our clustering methodology creates a discrimination among the candidate implementations.
- The information of discrimination can be used as a basis to train neural network based approaches.

Acknowledgement

Financial support from the
Deutsche Forschungsgemeinschaft (DFG)
through grant IRTG-2379 is
gratefully acknowledged

