Ocelotl: Time Aggregation Visualization for Trace Overview

Damien Dosimont [1] Guillaume Huard [2] Jean-Marc Vincent [2]

[1] INRIA

[2] Grenoble University

firstname.lastname@imag.fr

LICIA 3rd Workshop (October 22nd 2013)

Context: SoC-Trace project

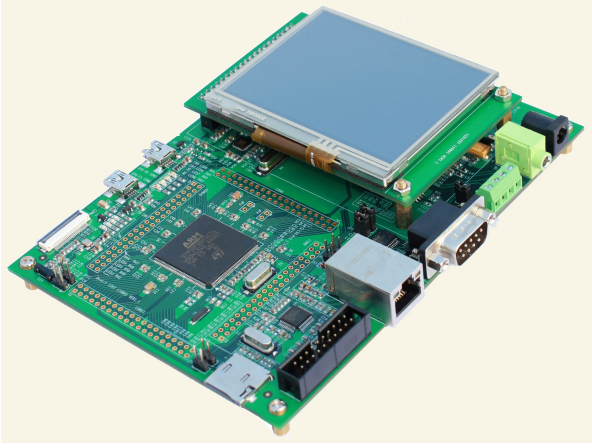

Embedded system trace analysis problematic

- **Hardware** and **software** complexity
- Trace **size** and **format** management
- Analysis technique **scalability**

Propositions

- **FrameSoC** infrastructure : storage, data-model, trace/tool/result management
- **Analysis flow** : statistics, data-mining, **visualization**...

Context: SoC-Trace project

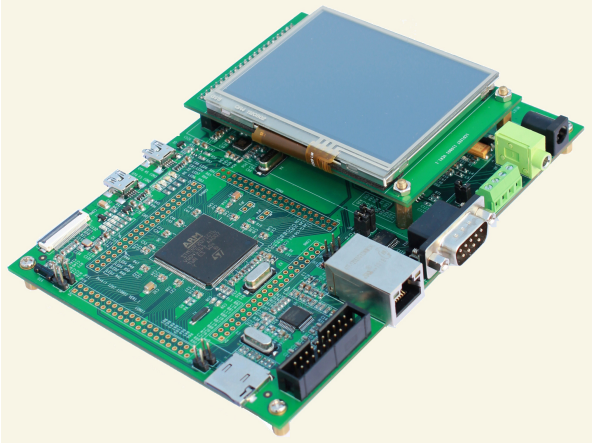

Embedded system trace analysis problematic

- **Hardware** and **software** complexity
- Trace **size** and **format** management
- Analysis technique **scalability**

Propositions

- **FrameSoC** infrastructure : storage, data-model, trace/tool/result management
- **Analysis flow** : statistics, data-mining, **visualization**...

Context: SoC-Trace project

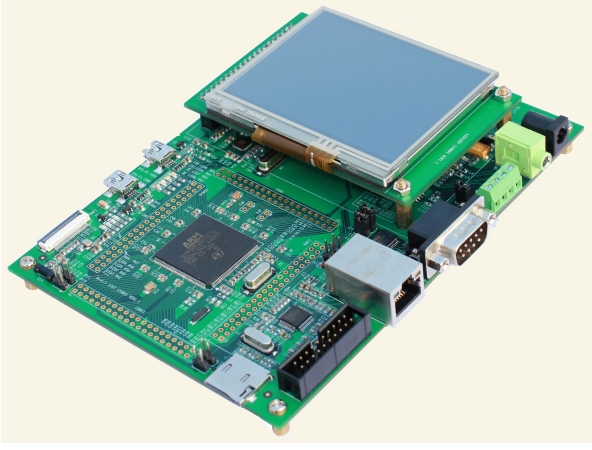

Embedded system trace analysis problematic

- **Hardware** and **software** complexity
- Trace **size** and **format** management
- Analysis technique **scalability**

Propositions

- **FrameSoC** infrastructure : storage, data-model, trace/tool/result management
- **Analysis flow** : statistics, data-mining, **visualization**...

Fil rouge: typical embedded system use case

Fil rouge: application tracing result

Only 20 second duration but...

- Almost **1500** different functions, **4** threads
- **One million** of events
- **100 Mo** trace (Pajé format)

For a 10 minute-long video

- Same number of functions, but...
- More than **30 millions** of events!
- **3 GB** trace!

We can easily obtain well bigger traces!!

Fil rouge: application tracing result

Only 20 second duration but...

- Almost **1500** different functions, **4** threads
- **One million** of events
- **100 Mo** trace (Pajé format)

For a 10 minute-long video

- Same number of functions, but...
- More than **30 millions** of events!
- **3 GB** trace!

We can easily obtain well bigger traces!!

Fil rouge: application tracing result

Only 20 second duration but...

- Almost **1500** different functions, **4** threads
- **One million** of events
- **100 Mo** trace (Pajé format)

For a 10 minute-long video

- Same number of functions, but...
- More than **30 millions** of events!
- **3 GB** trace!

We can easily obtain well bigger traces!!

How to represent this trace behavior over time?

Gantt Chart is the most common technique employed by analysts...

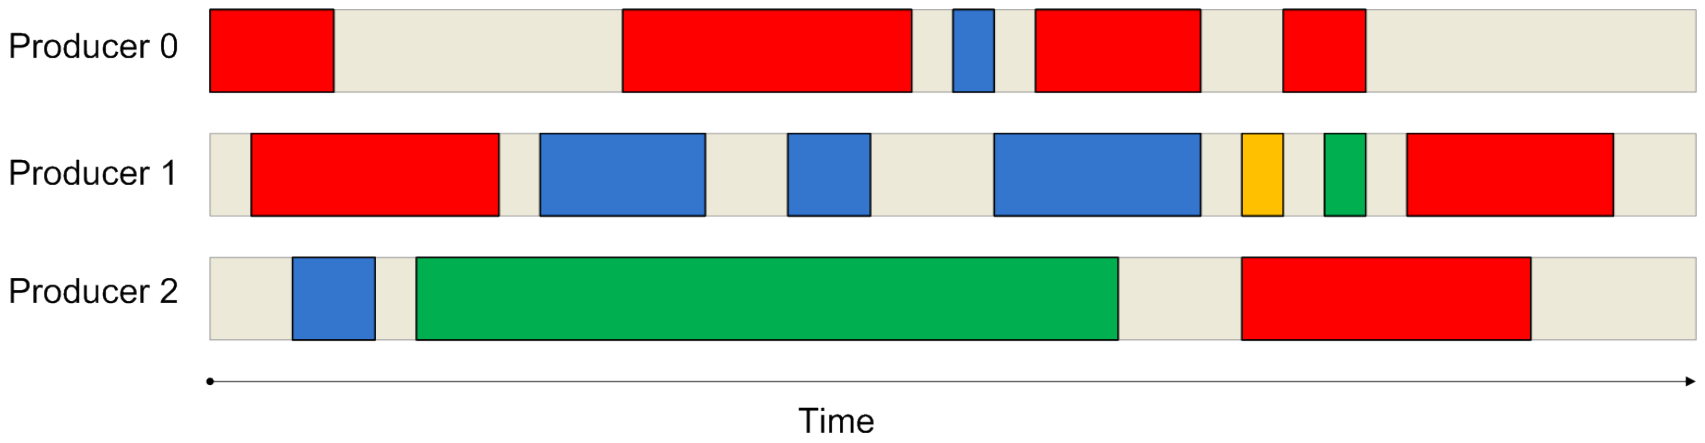

Figure 1: *Synthetic example of Gantt Chart*

How to represent this trace behavior over time?

Gantt Chart is the most common technique employed by analysts...

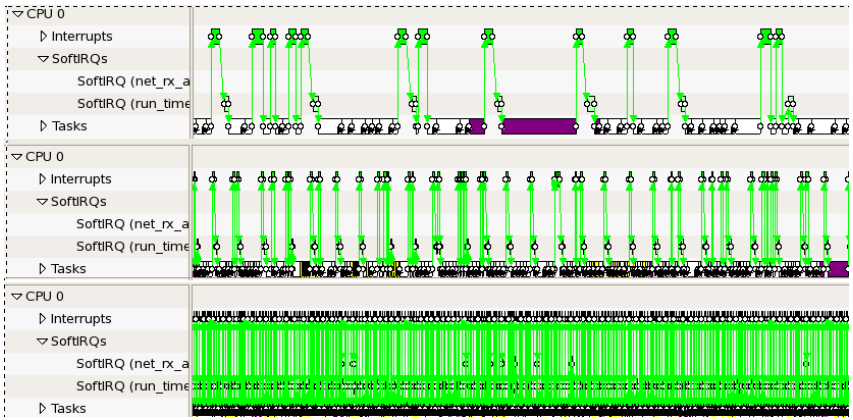

Figure 1: *KPTrace dezoom : example of time axis scalability issues*

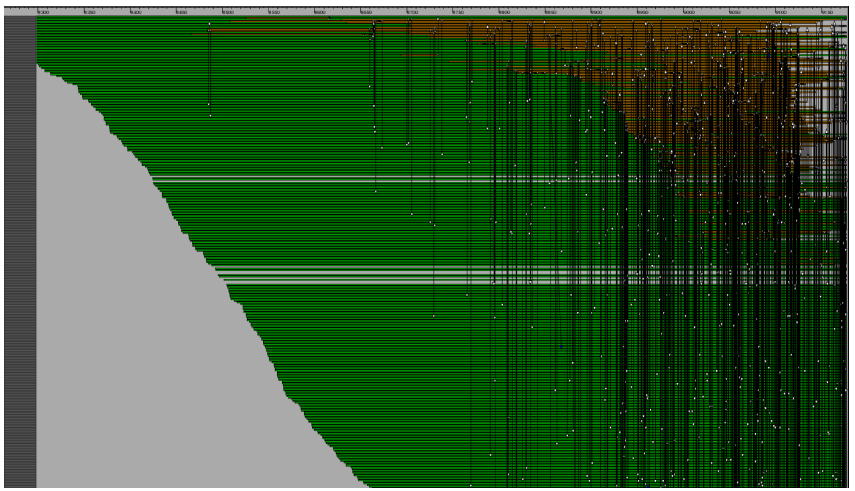

Figure 2: *Example of space limitations : Pajé trace with 700 producers*

... but it does not scale to voluminous traces

Our proposal: Ocelotl

Fit to Schneiderman's methodology...

- **Overview** first, zoom and filter, then details on demand

... build upon an algorithm proposed by Lamarche-Perrin

- Adapted to timestamped events using **time slicing**
- Extended to **multiple event sources**

Our proposal: Ocelotl

Fit to Schneiderman's methodology...

- **Overview** first, zoom and filter, then details on demand

... by providing a macroscopic description of the trace...

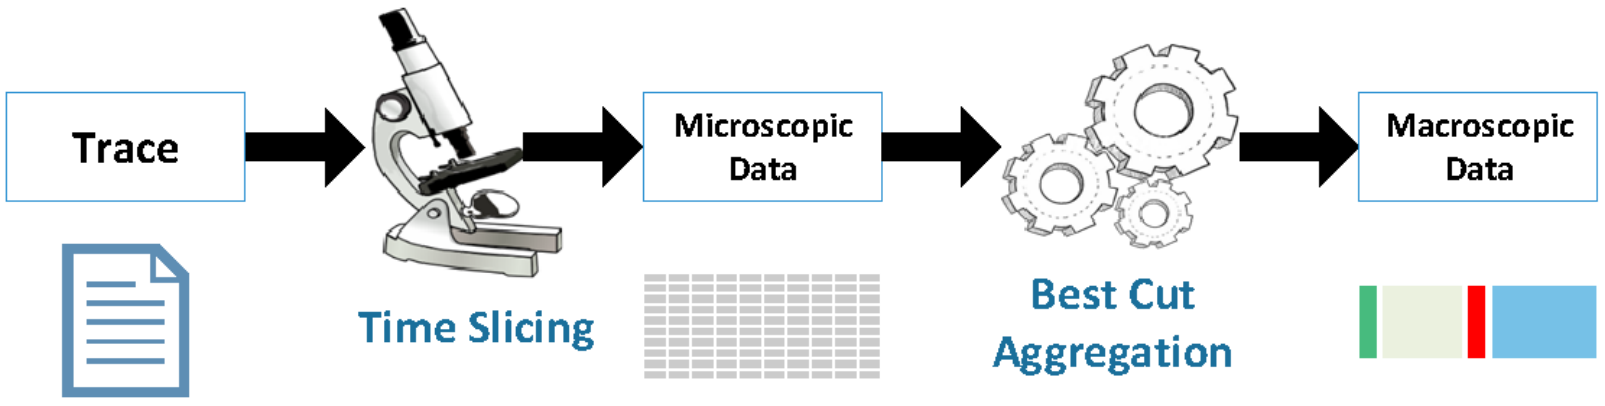

... build upon an algorithm proposed by Lamarche-Perrin

- Adapted to timestamped events using **time slicing**
- Extended to **multiple event sources**

Our proposal: Ocelotl

Fit to Schneiderman's methodology...

- **Overview** first, zoom and filter, then details on demand

... by providing a macroscopic description of the trace...

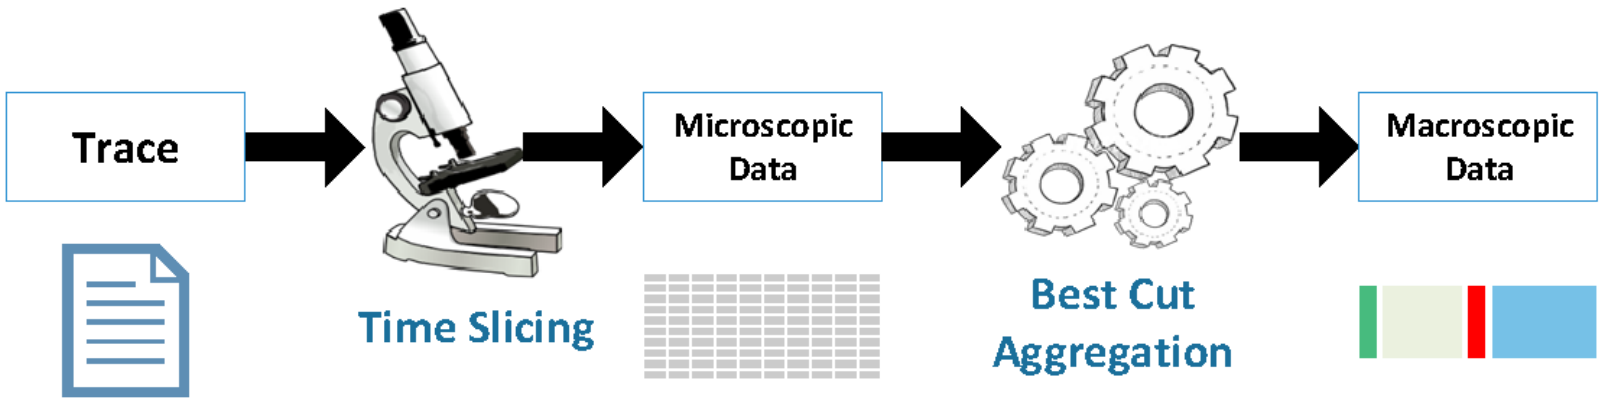

... build upon an algorithm proposed by Lamarche-Perrin

- Adapted to timestamped events using **time slicing**
- Extended to **multiple event sources**

User controlled level of details

Parametrized aggregation enables interaction

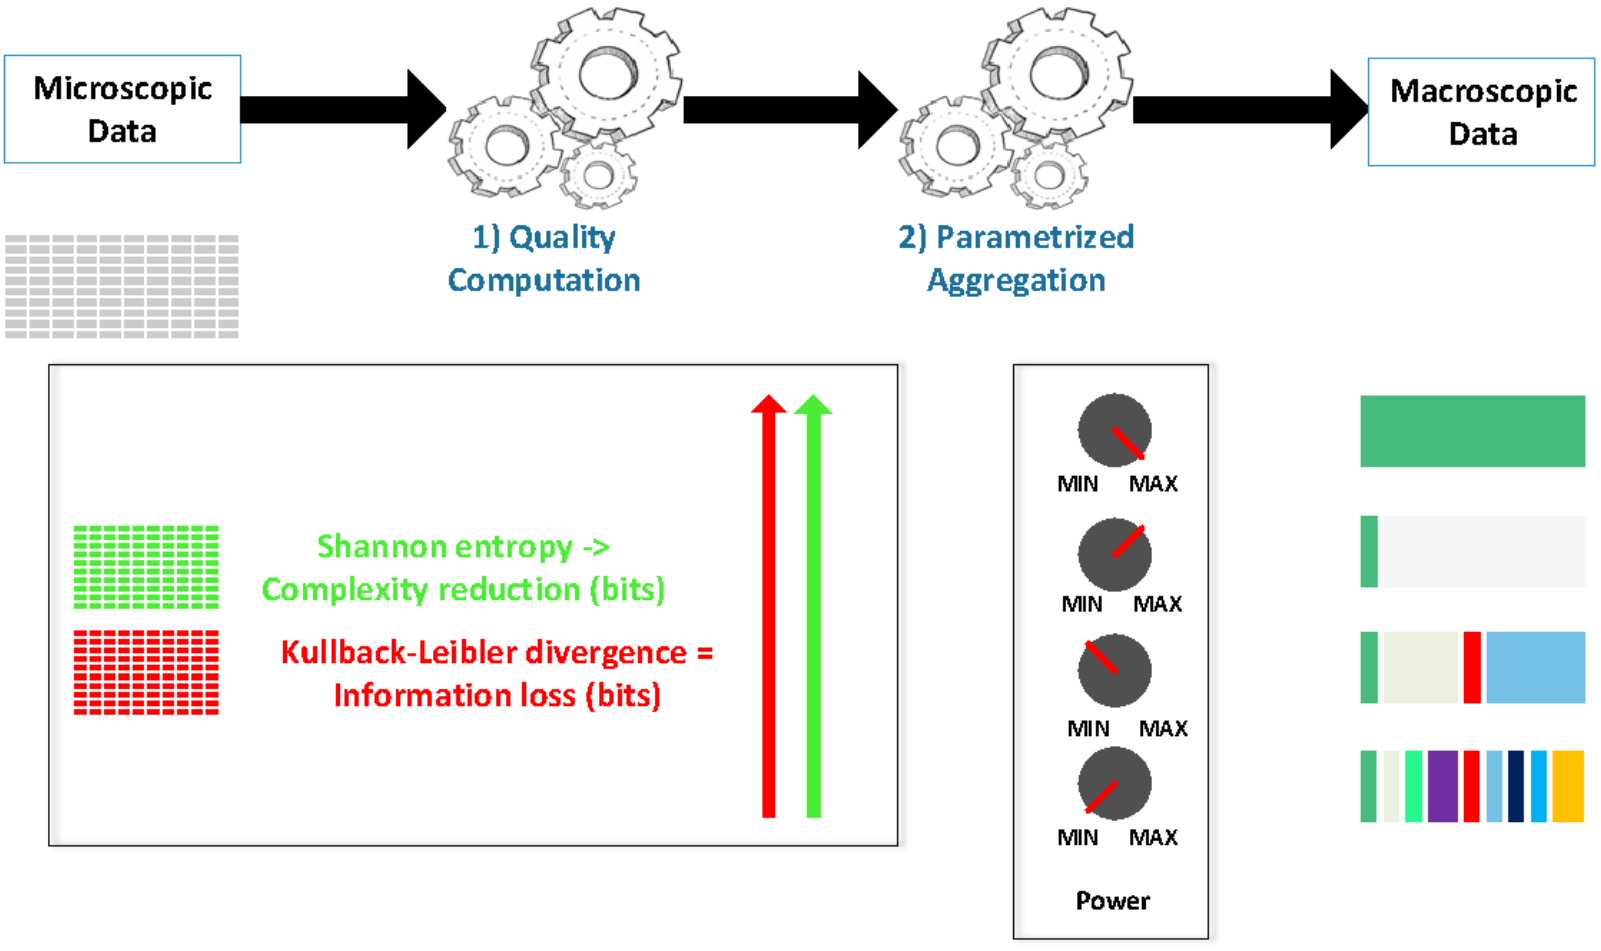

Analysis with Ocelotl (Settings)

Analysis with Ocelotl (Overview, Qualities)

Analysis with Ocelotl (Zoom and details)

Results

Information based time aggregation

- ... describes behavior by **highlighting phases** and **perturbations**
- **Interaction** helps to focus on these points
- **Performance**:
  - 20s to visualize 1M event trace
  - 2h to visualize 30M event trace (bounded by database query time)

But...

- Lack of **space dimension** representation

Results

Information based time aggregation

- ... describes behavior by **highlighting phases** and **perturbations**
- **Interaction** helps to focus on these points
- **Performance**:
  - 20s to visualize 1M event trace
  - 2h to visualize 30M event trace (bounded by database query time)

But...

- Lack of **space dimension** representation

Current Focus

New features

- **Spatio-temporal** aggregation

Use-cases

- **HPC/Distributed** system relevant **use-cases**

Current Focus

New features

- **Spatio-temporal** aggregation

Use-cases

- **HPC/Distributed** system relevant **use-cases**

Links

My website

http://moais.imag.fr/membres/damien.dosimont/

Tools and libraries are available on my github

http://github.com/dosimont

Merci pour votre attention!

Implementation

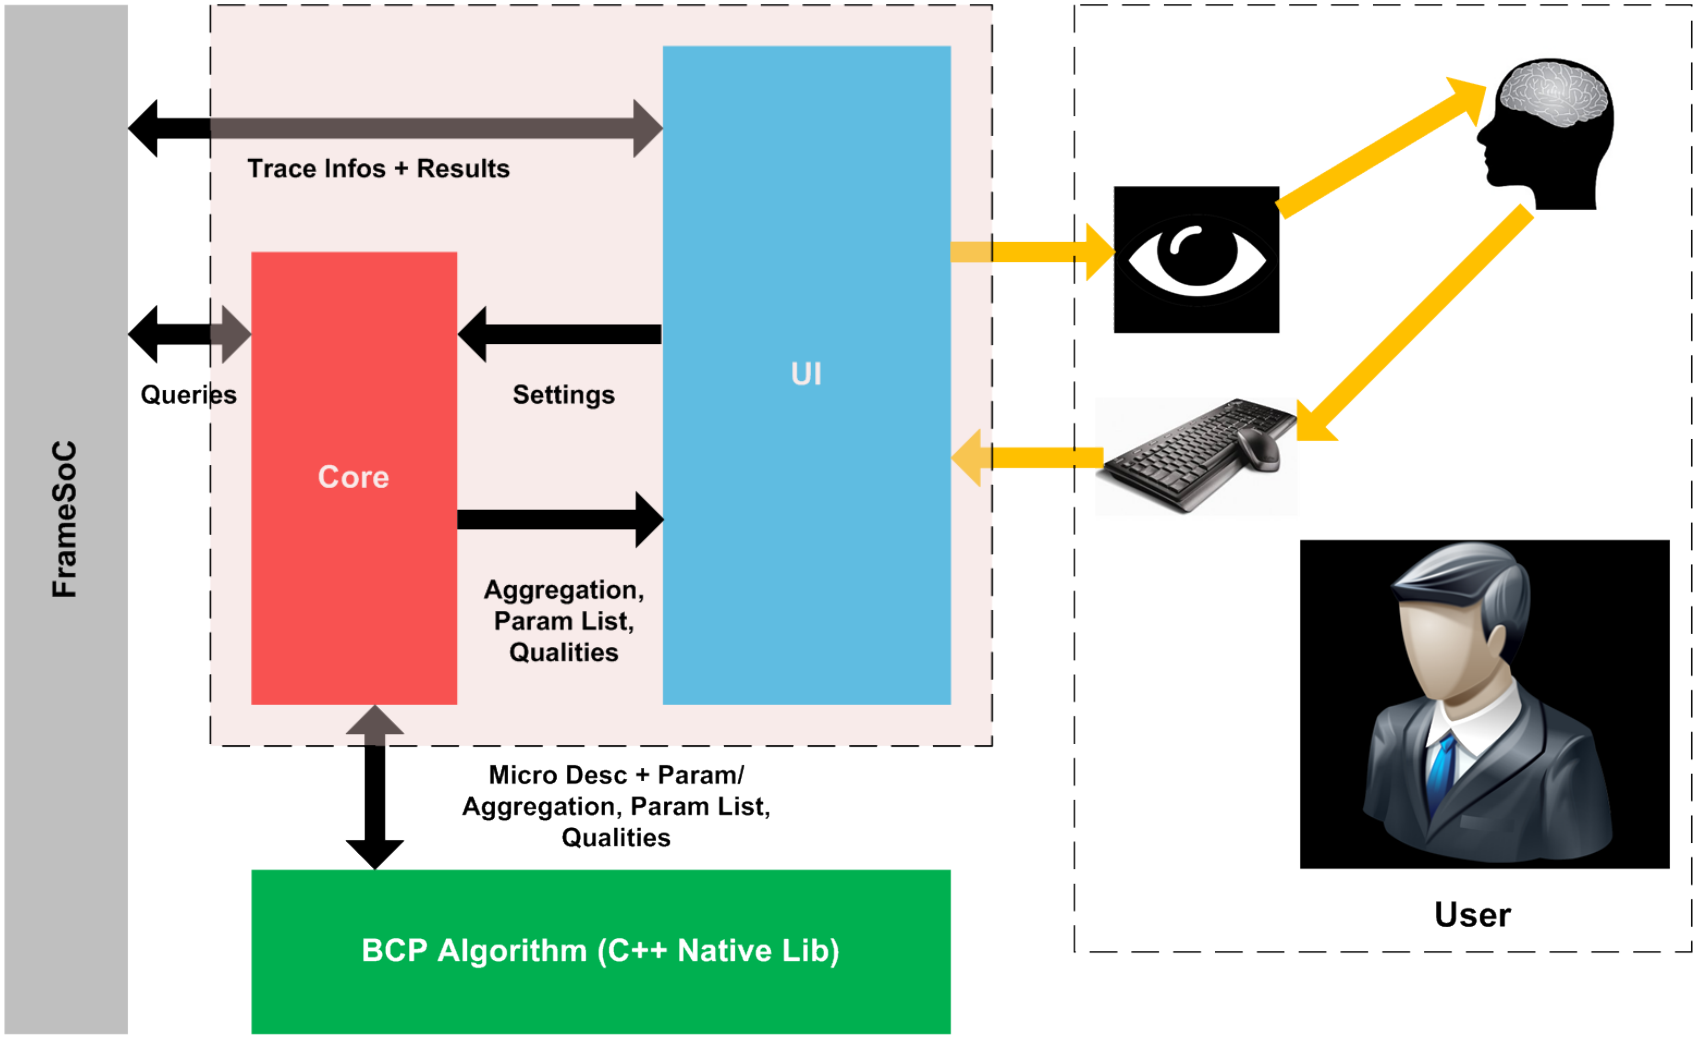

Interface Overview

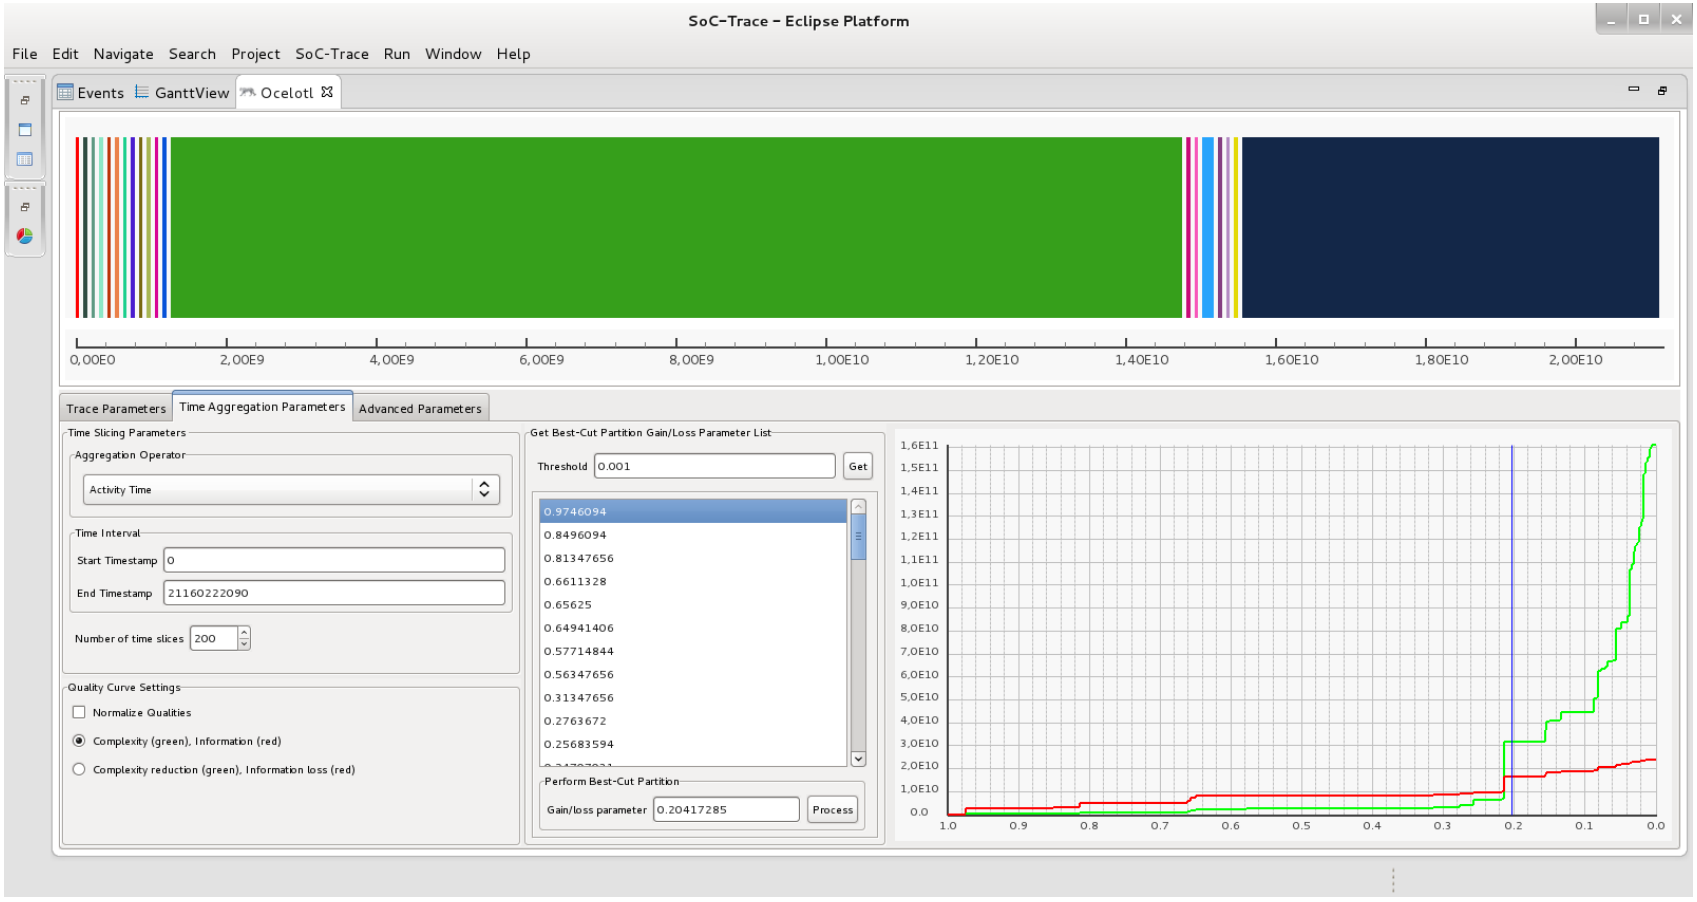

Interface Overview

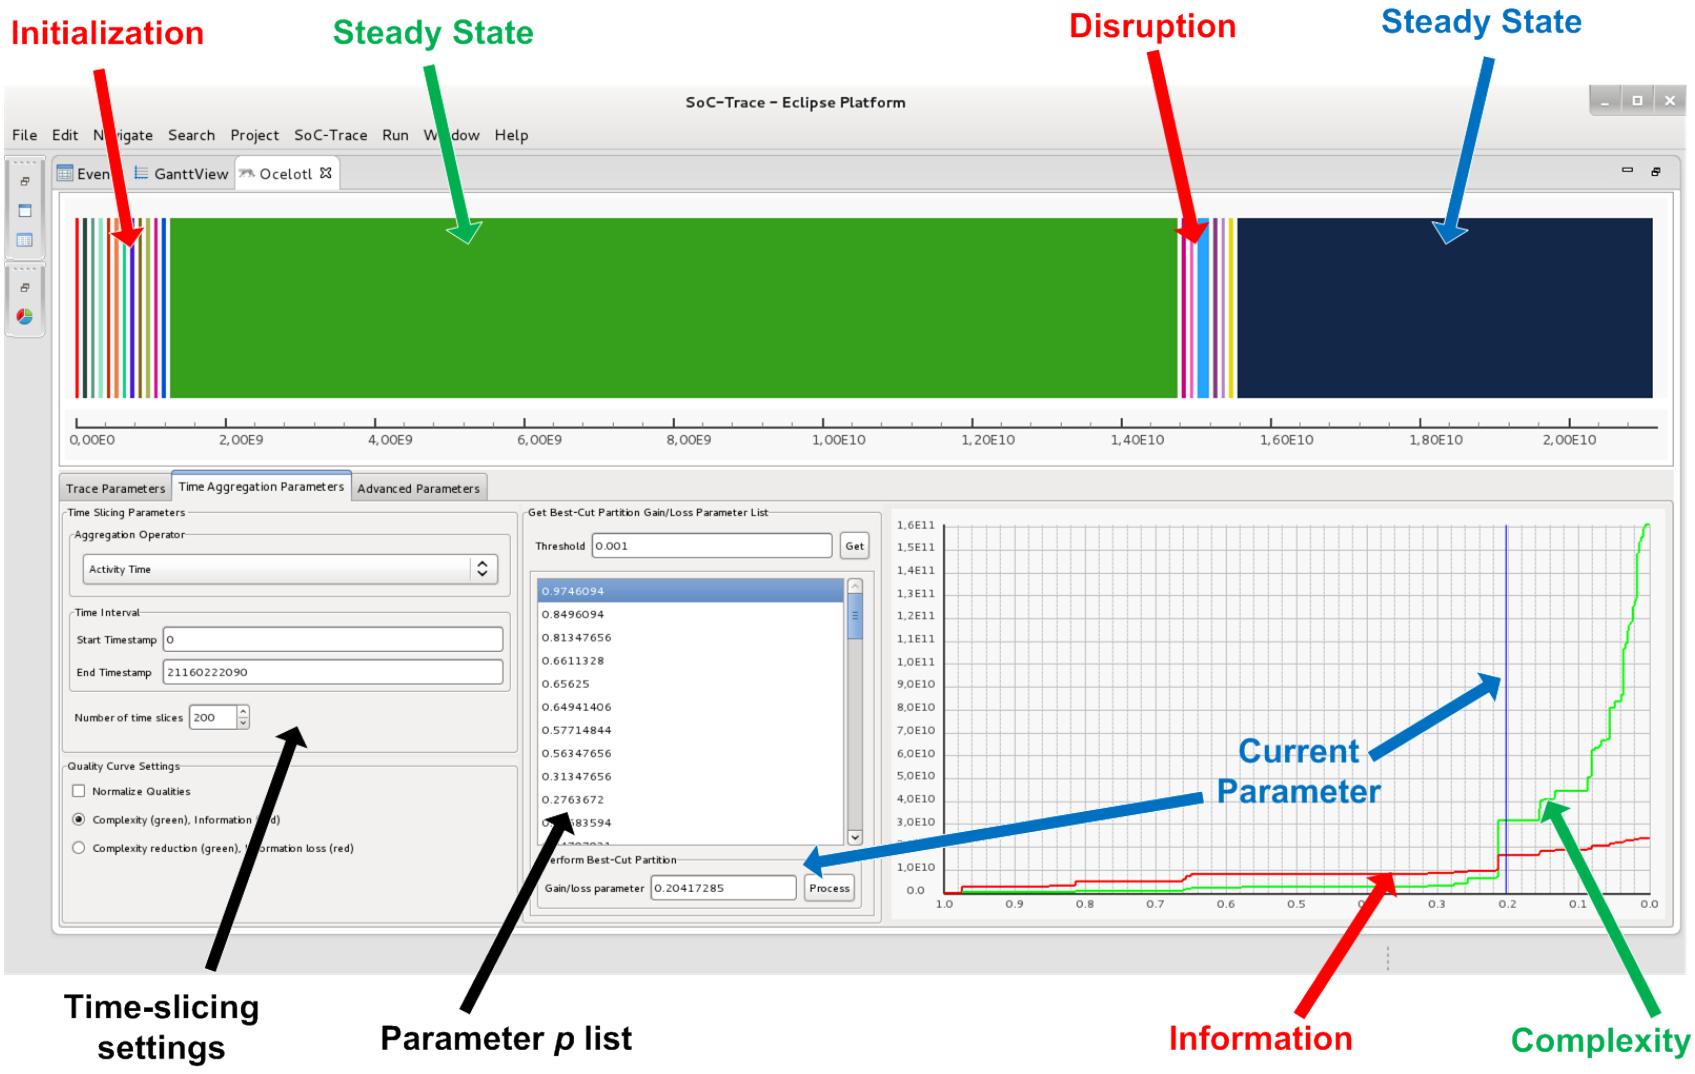

Lamarche-Perrin Works: Multi-Agent Systems

How to Build a Meaningful Macroscopic Description?

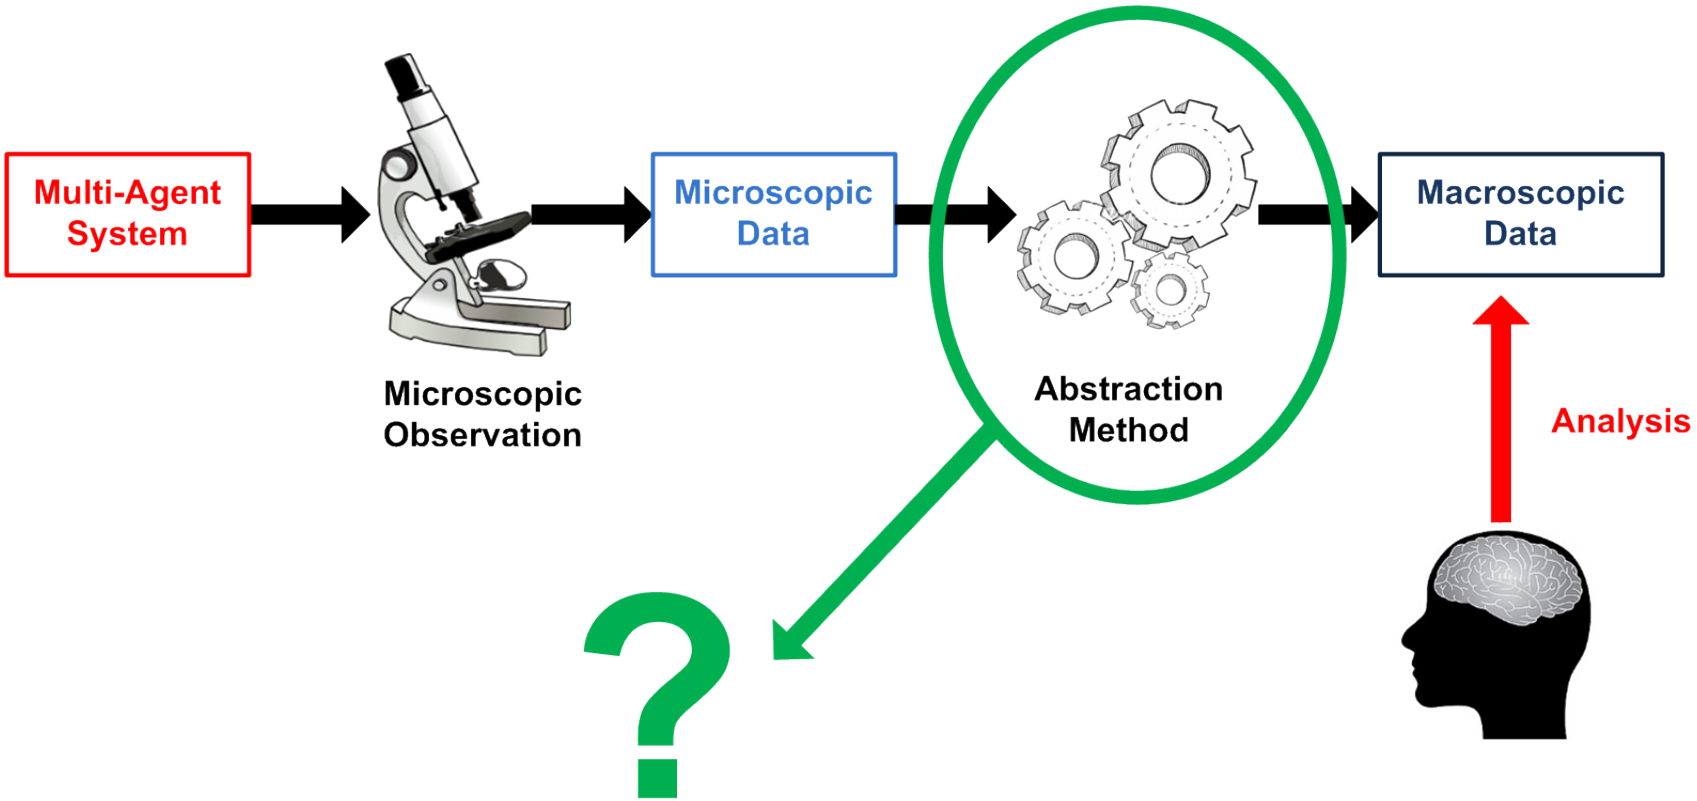

Example: Geomedia Project

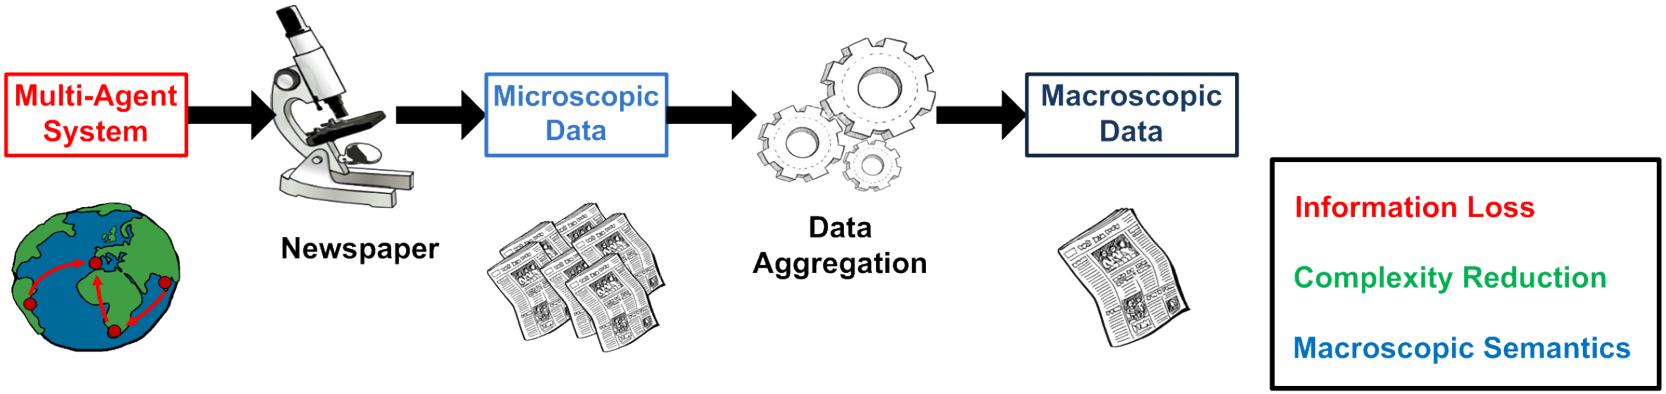

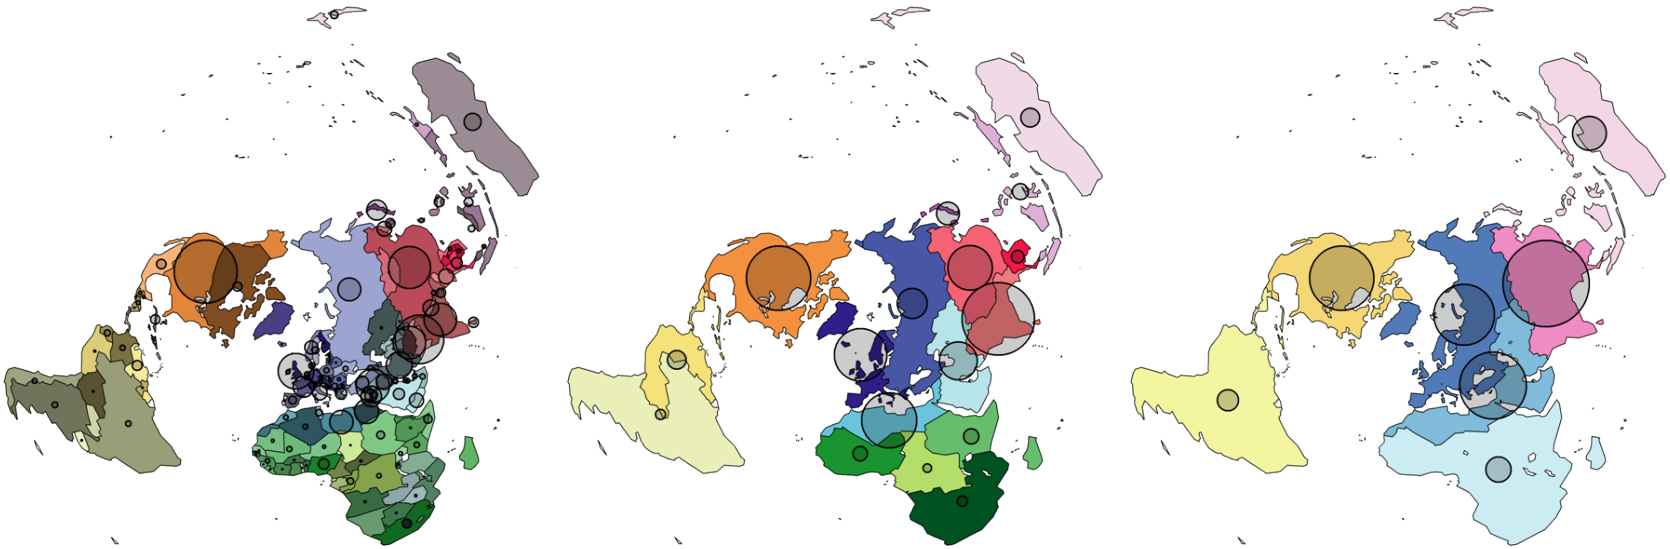

Resolution: Max | Resolution: Intermediary | Resolution: Min

Example: Viva

Represent Hierarchical Structure according to Value Heterogeneity

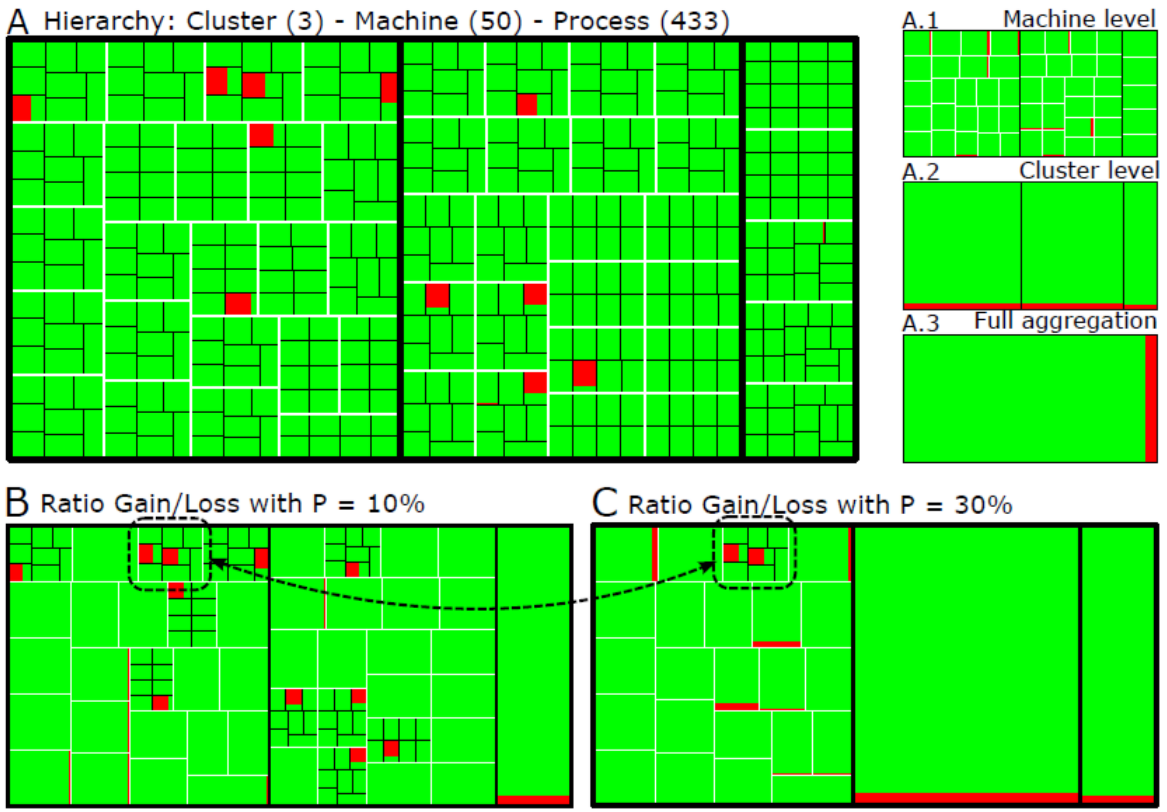

Information Loss

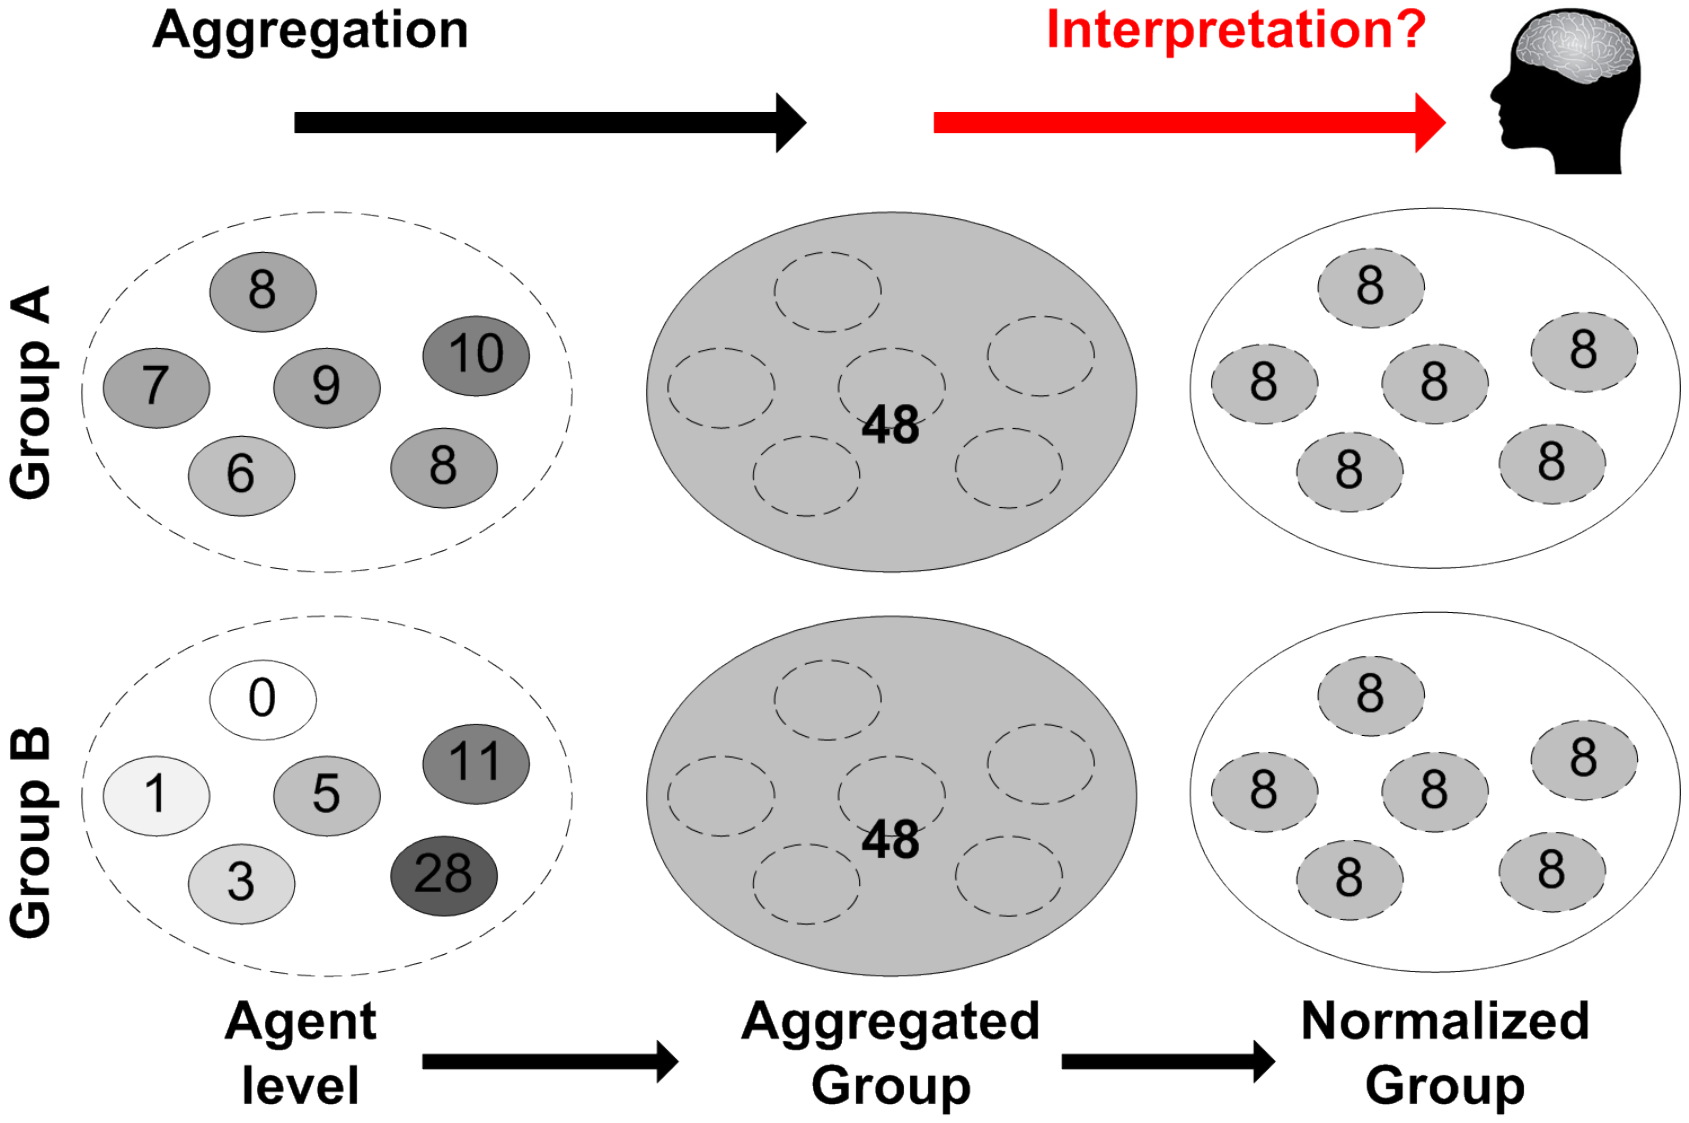

Information Loss

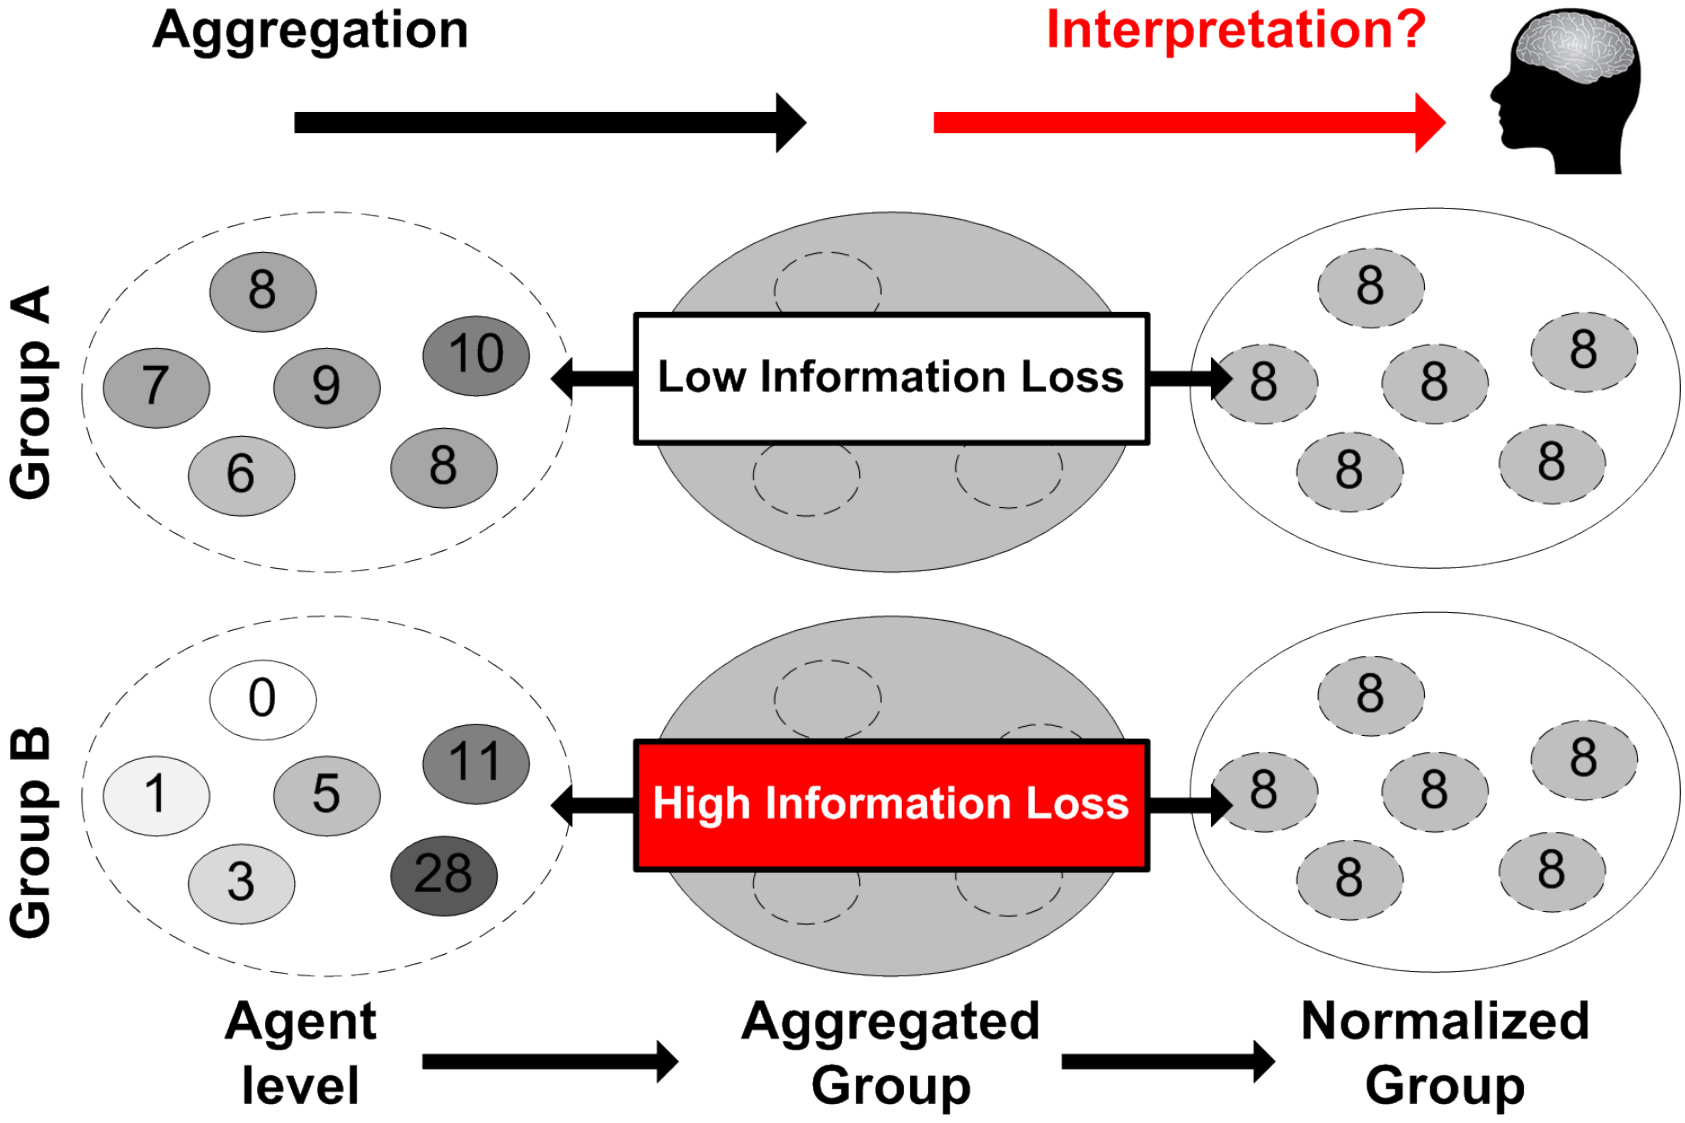

Information Loss Measure

Kullback-Leibler Divergence

$$\text{loss}(A||e) = \sum_{e \in A} v(e) \times \log_2 \left(\frac{v(e)}{v(A)} \right) \text{ in bits/x}$$

- Quantity of information than one **loses** by using an **aggregated description** instead of the **microscopic description**

Complexity Reduction

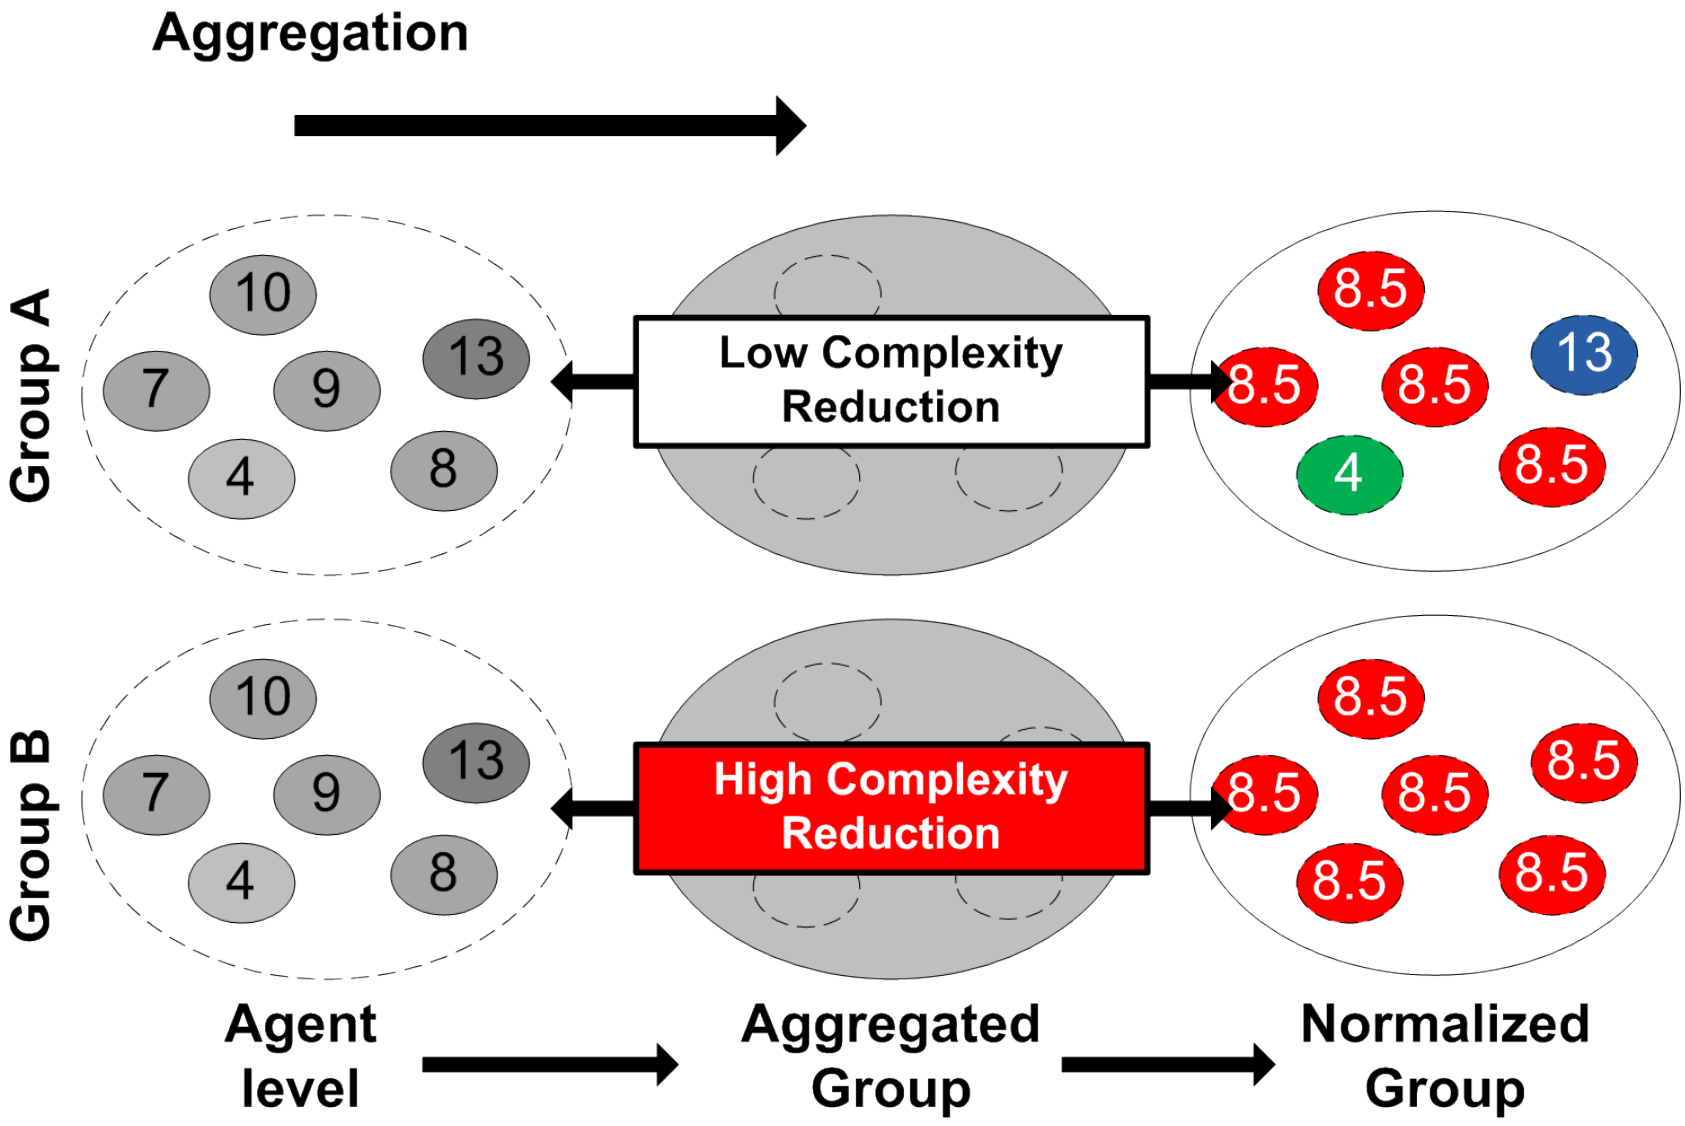

Complexity Reduction Measure

Shannon Entropy

$H(v) = \sum (v(i) \times \log_2 v(i))$ in bits/x

Entropy Reduction

$gain(A||e) = H(A) - H(e)$ in bits/x

- Quantity of information than one **saves** by encoding the **aggregated description** instead of the **microscopic description**

Compromise Finding between Information Loss and Complexity Reduction

Parametrized Information Criterion

$$\mathrm{pIC}(\mathcal{A}) = p \times \mathrm{gain}(\mathcal{A}) - (1-p) \times \mathrm{loss}(\mathcal{A})$$

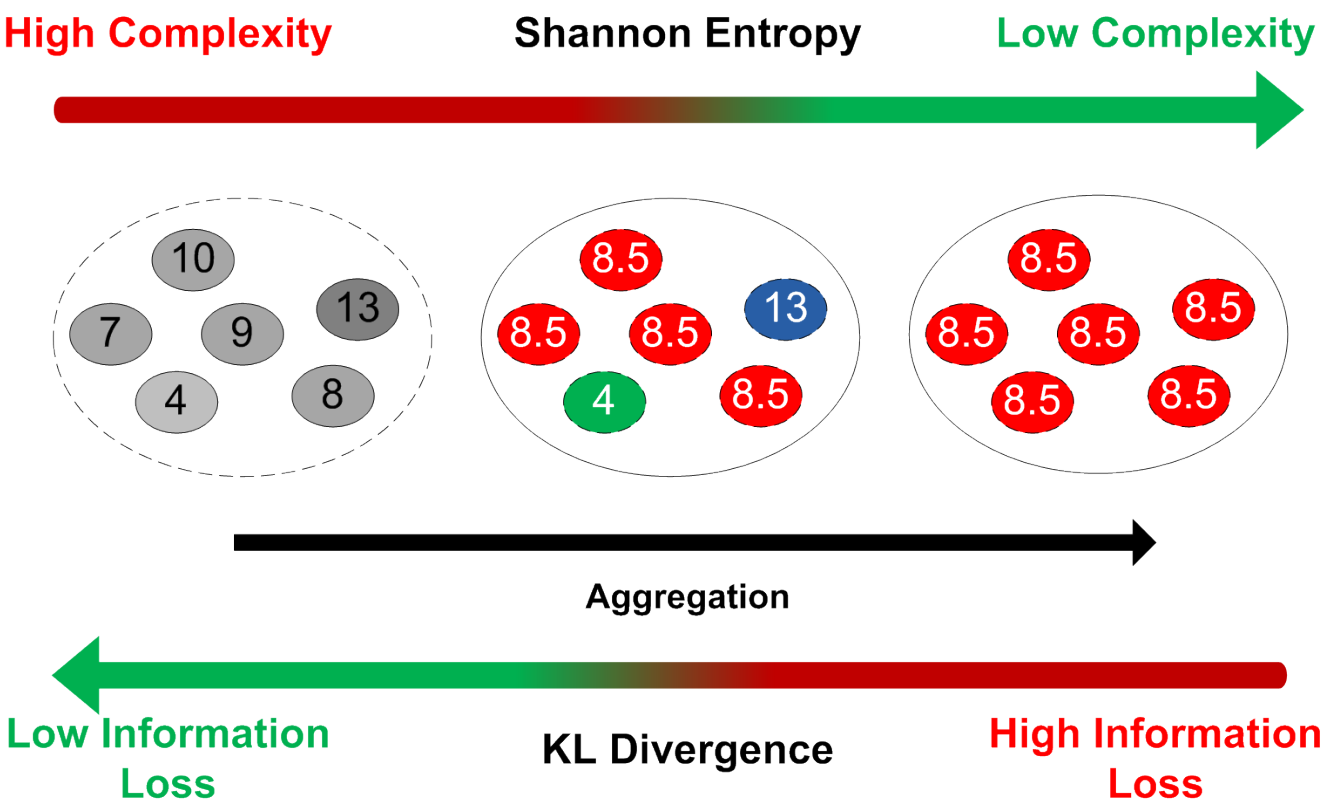

Temporal Aggregation

Temporal Aggregation principle

- Same principle but only consecutive data can be aggregated

Ex: Tunisia citation

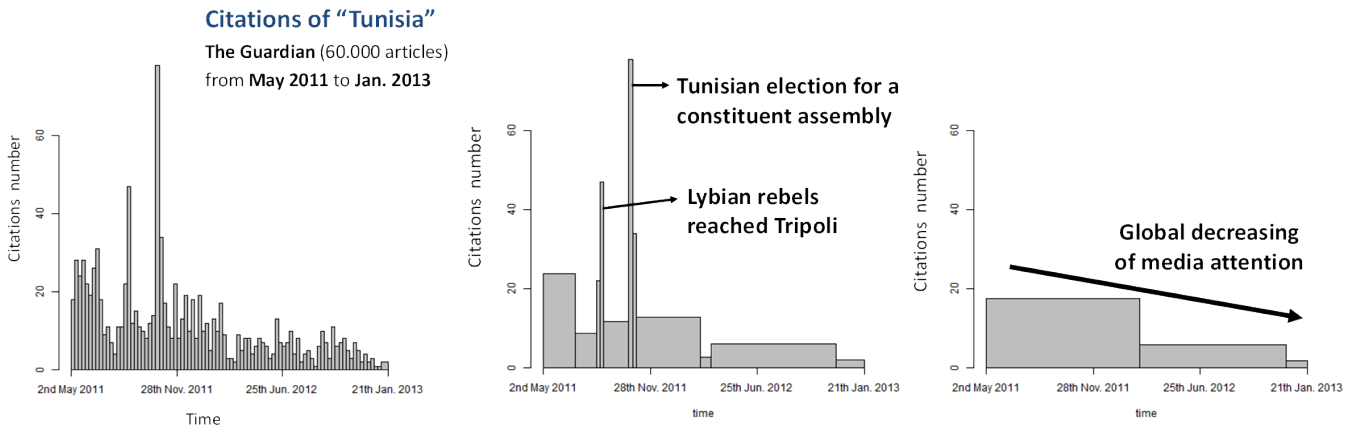

Need of a microscopic level description

Microscopic Level: Time-Slicing

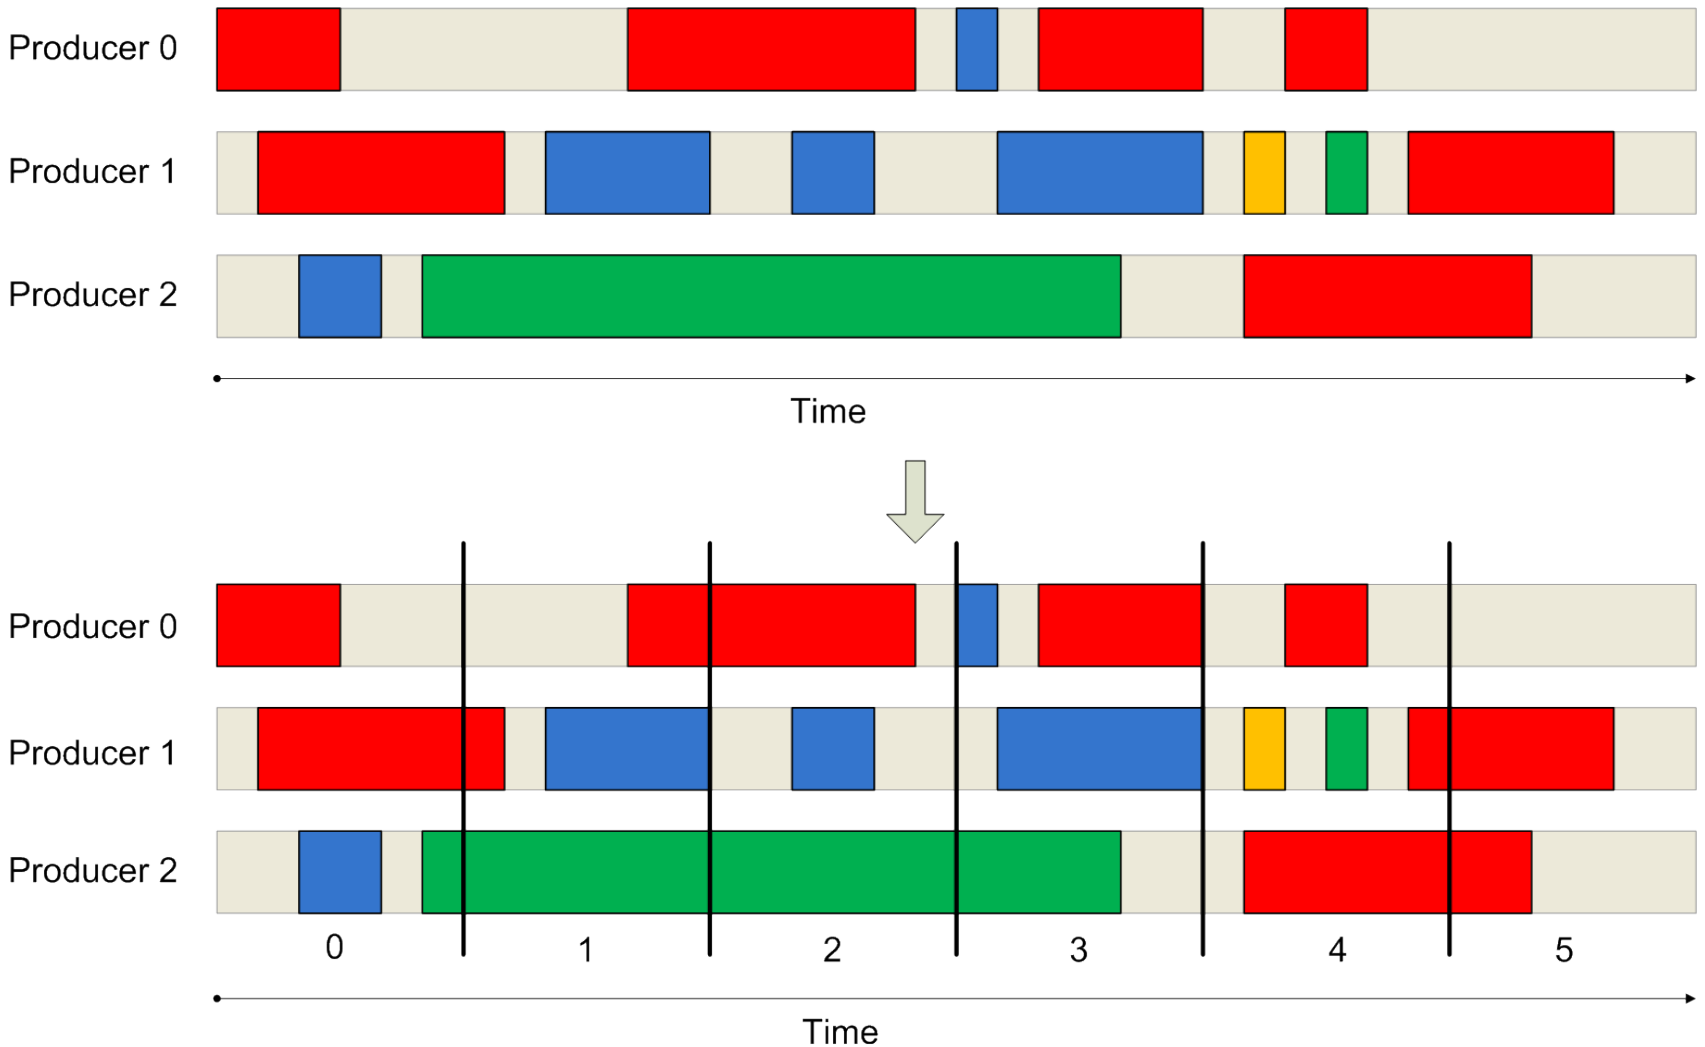

Microscopic Level: Producer Activity Time Matrix

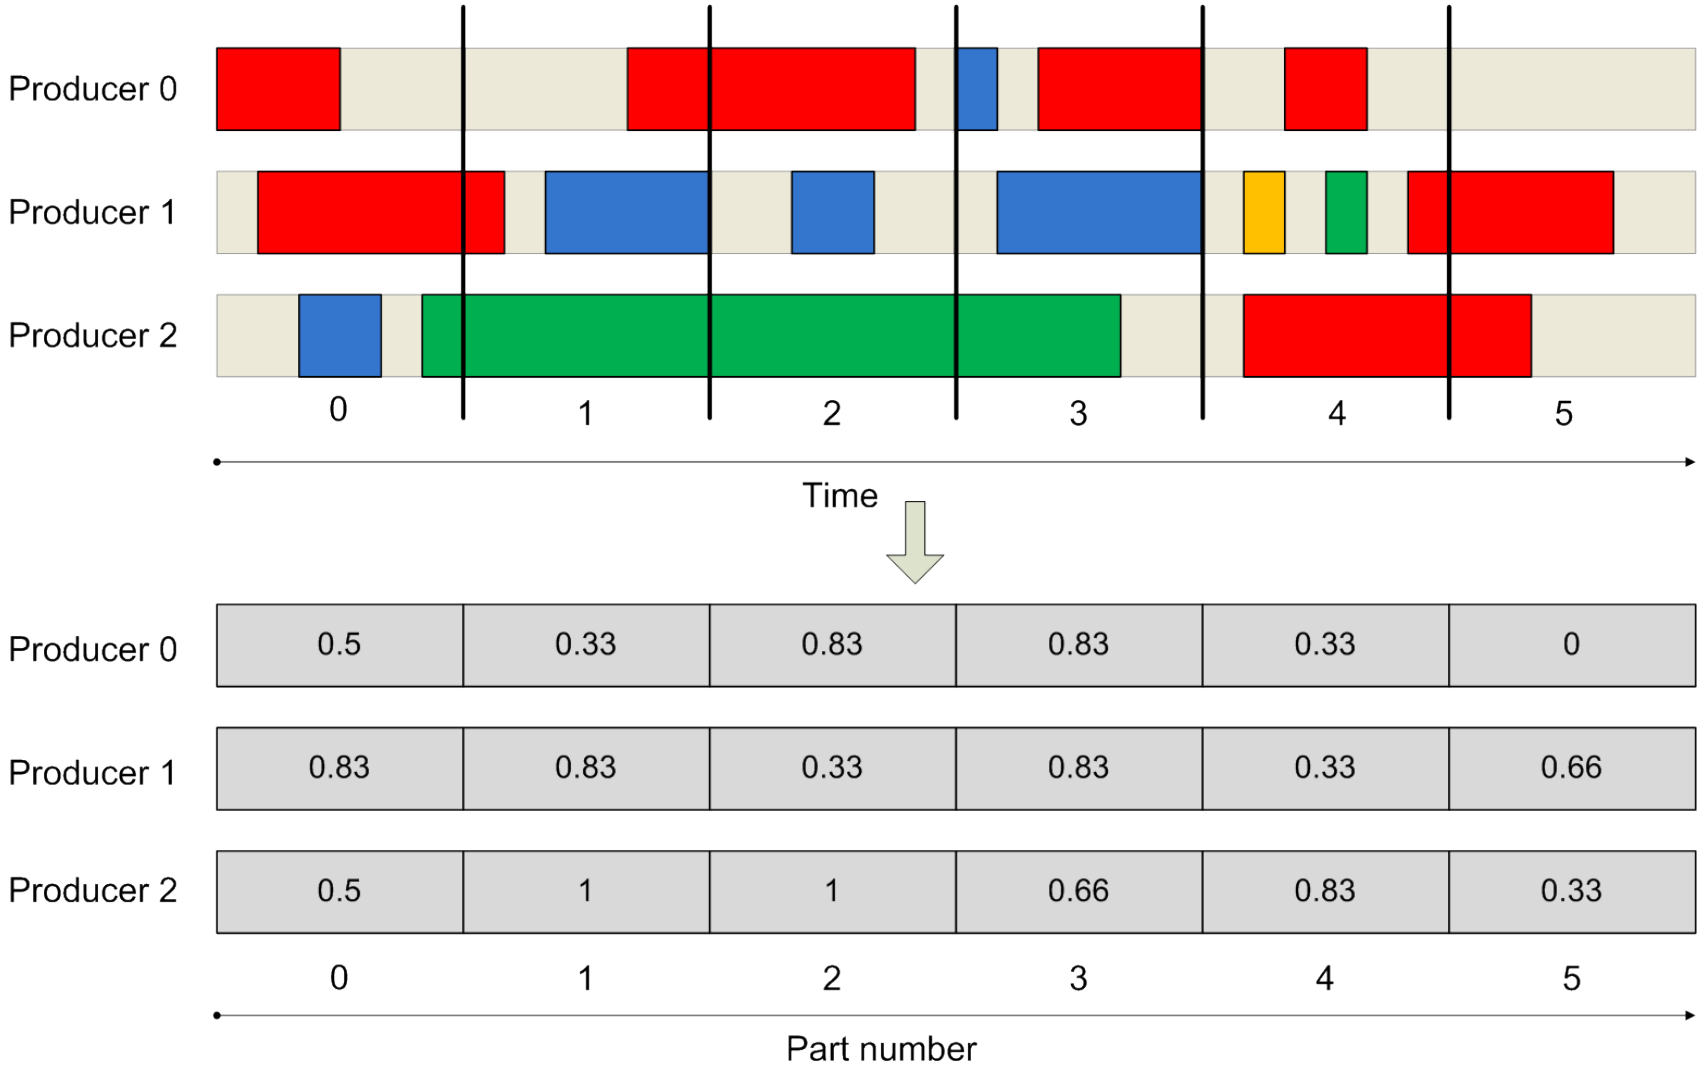

Microscopic Level: State Activity Time Cubic Matrix

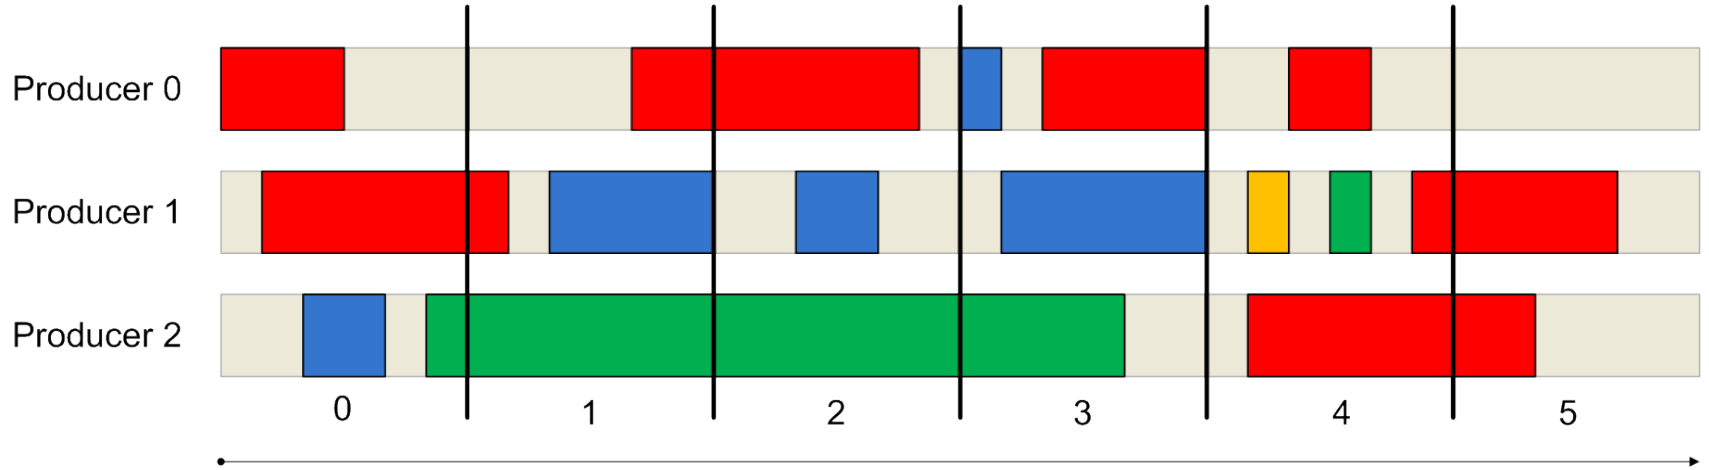

Producer 0	0.5	0.33	0.83	0.66	0.33	0
	0	0	0	0.16	0	0
	0	0	0	0	0	0
	0	0	0	0	0	0
Producer 1	0.83	0.16	0	0	0.16	0.66
	0	0	0.33	0.83	0	0
	0	0	0	0	0.16	0
	0	0	0	0	0.16	0
Producer 2	0	0	0	0	0.83	0.33
	0.33	0	0	0	0.33	0
	0.16	1	1	0.66	0.33	0
	0	0	0	0	0	0
	0	1	2	3	4	5

Part number

Quality Computation

Gain and loss formulas: originally for scalars

012345					
01234	12345				
0123	1234	2345			
012	123	234	345		
01	12	23	34	45	
0	1	2	3	4	5

Adaptation for time-sliced description

- Vector (ex: activity time per process)
  $\text{quality}(A) = \sum_{i \in n} \text{quality}(A[i])$
- Matrix (ex: activity time per state type)
  $\text{quality}(A) = \sum_{i \in n} (\sum_{j \in m} \text{quality}(A[i][j]))$

Best-Cut Partition for a given *p*