Time aggregation visualization for embedded system program trace analysis

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

[1] INRIA

[2] Joseph Fourier University (UJF)
Grenoble

firstname.lastname@imag.fr

SONGS Meeting (June 2013)

Let’s play a game!

Embedded system multimedia applications

Embedded system growing complexity

- **Hardware**: multicore, heterogeneity
- **Software**: software stack, API, middlewares, multithread

Multimedia application characteristics

- Hardware and software **components**
- **Sequential** treatment on data
- **Cyclic** behavior, synchronization, QoS

QoS non respected!

- Video or sound perturbed!
- **Why?** overloaded buffer, IT conflict, drivers, HW, scheduling...

Embedded system multimedia applications

Embedded system growing complexity

- **Hardware**: multicore, heterogeneity
- **Software**: software stack, API, middlewares, multithread

Multimedia application characteristics

- Hardware and software **components**
- **Sequential** treatment on data
- **Cyclic** behavior, synchronization, QoS

QoS non respected!

- Video or sound perturbed!
- **Why?** overloaded buffer, IT conflict, drivers, HW, scheduling...

Embedded system multimedia applications

Embedded system growing complexity

- **Hardware**: multicore, heterogeneity
- **Software**: software stack, API, middlewares, multithread

Multimedia application characteristics

- Hardware and software **components**
- **Sequential** treatment on data
- **Cyclic** behavior, synchronization, QoS

QoS non respected!

- Video or sound perturbed!
- **Why?** overloaded buffer, IT conflict, drivers, HW, scheduling...

SoC-Trace project

Tracing-based analysis

- Get synchro **events**, process/function/IT **states**
- Low **intrusivity** but huge **data volume**
- How to relate trace information to application behavior?

SoC-Trace analysis framework

- Trace storage, data-model, trace/tools/results management
- Analysis flow : statistics, processing, data-mining, **visualization**

Project context

- Partners: INRIA, UJF, STMicroelectronics, ProbaYes, Magillem
- Granted by FUI

SoC-Trace project

Tracing-based analysis

- Get synchro **events**, process/function/IT **states**
- Low **intrusivity** but huge **data volume**
- How to relate trace information to application behavior?

SoC-Trace analysis framework

- Trace storage, data-model, trace/tools/results management
- Analysis flow : statistics, processing, data-mining, **visualization**

Project context

- Partners: INRIA, UJF, STMicroelectronics, ProbaYes, Magillem
- Granted by FUI

SoC-Trace project

Tracing-based analysis

- Get synchro **events**, process/function/IT **states**
- Low **intrusivity** but huge **data volume**
- How to relate trace information to application behavior?

SoC-Trace analysis framework

- Trace storage, data-model, trace/tools/results management
- Analysis flow : statistics, processing, data-mining, **visualization**

Project context

- Partners: INRIA, UJF, STMicroelectronics, ProbaYes, Magillem
- Granted by FUI

FrameSoc: SoC-Trace infrastructure

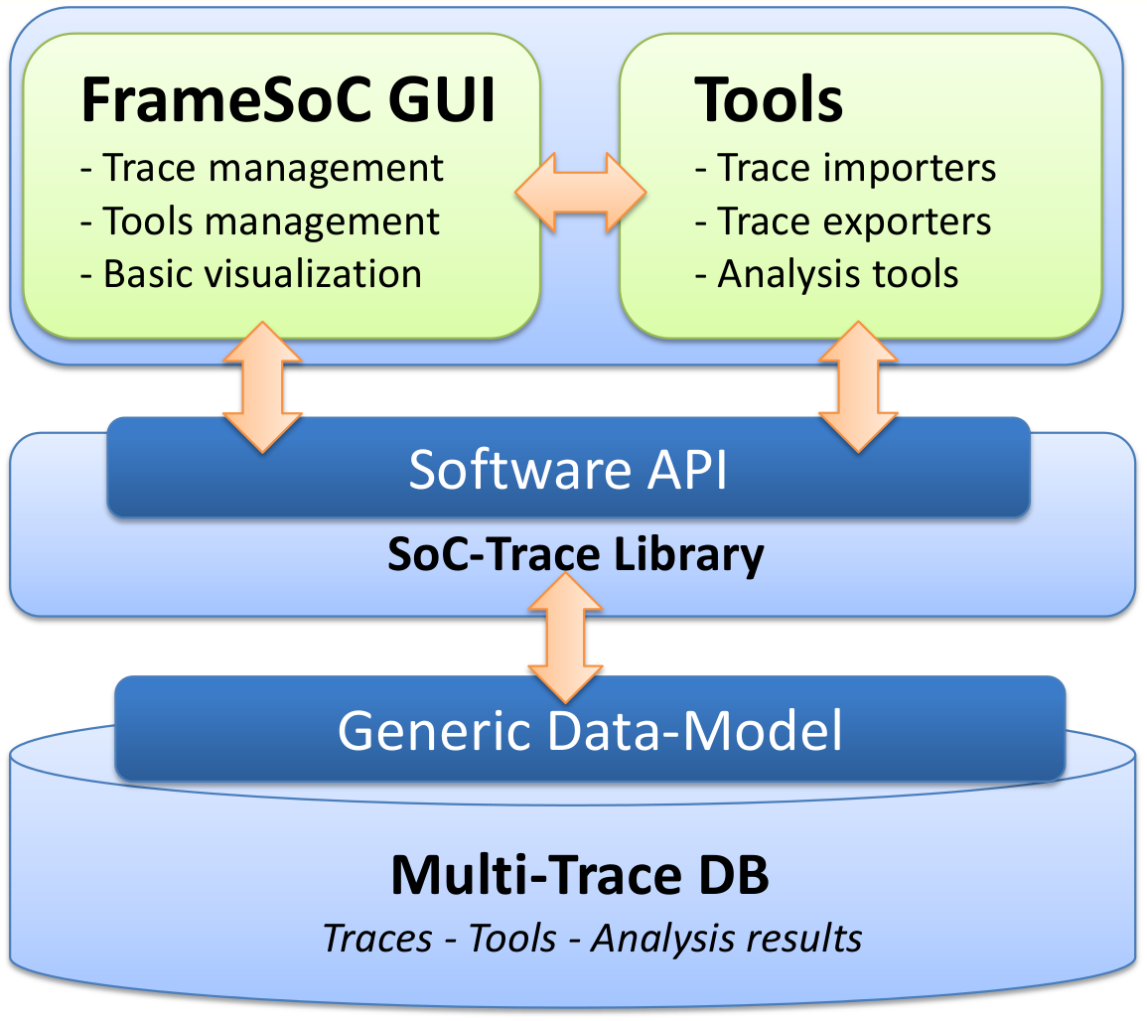

Figure 1 : *FrameSoC architecture and its features*

Table of Contents

Section

Visualization scalability issues

In traditional visualization techniques

- **Fidelity**: aliasing artifacts, proportions (zoom out)
- **Understanding**: loss of context (zoom in, scroll), information loss (aggregation)

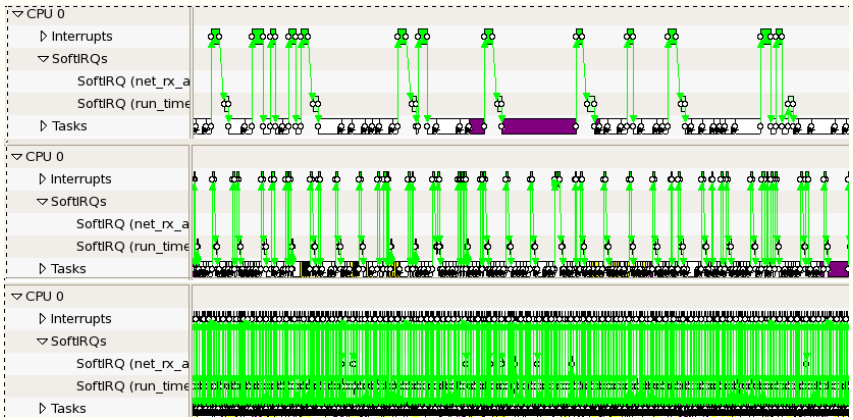

Figure 2 : *KPTrace dezoom : example of time axis scalability issues*

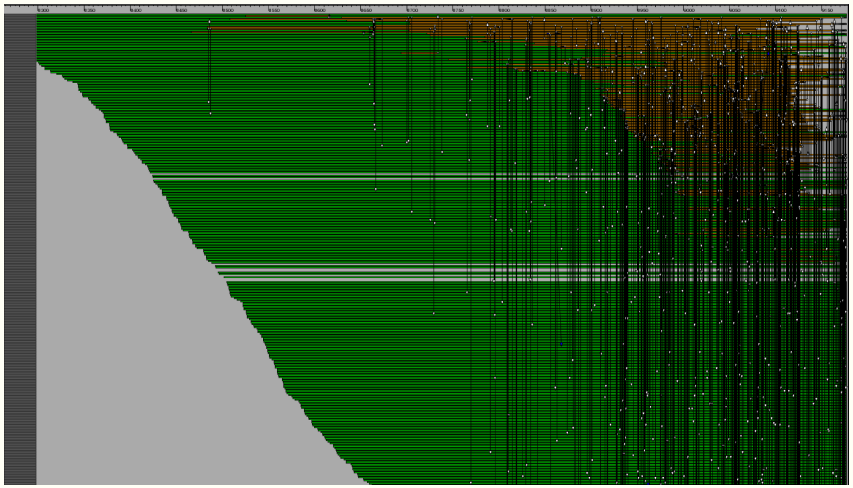

Figure 3 : *Example of space limitations : Pajé trace with 700 producers*

Viva: HPC system trace analysis

Hierarchical treemap and network topology representation

- **Time Aggregation**: time-slicing and animation
- **Space Aggregation**: hierarchy and/or best-cut partition

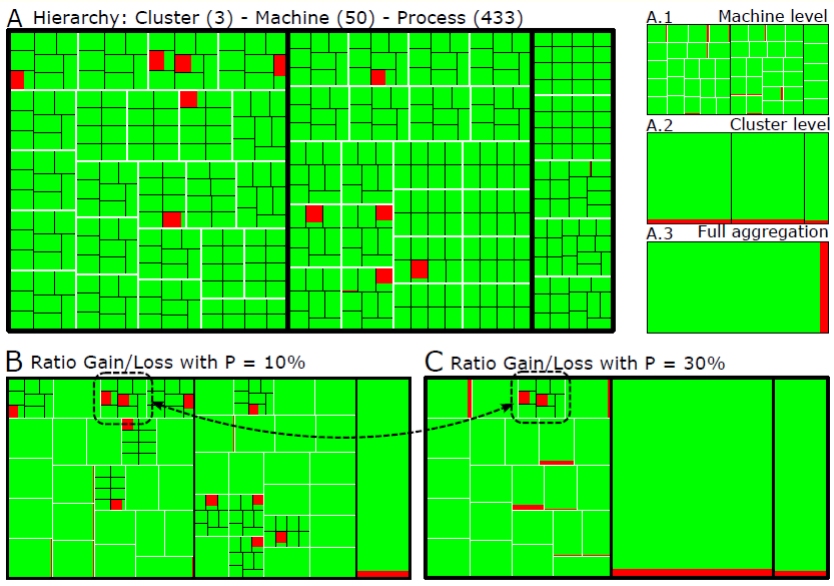

Figure 4 : *Treemap, showing the same application with different hierarchical aggregation settings*

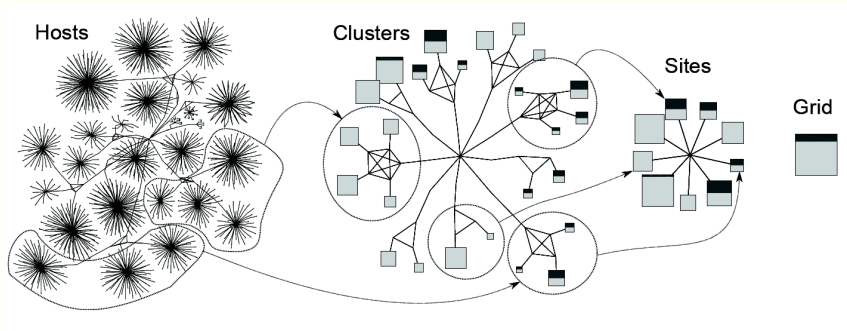

Figure 5 : *Network topology, with progressive hierarchical aggregation (left to right)*

Embedded system analysis needs?

Visual representation that enables to...

- ... show **behavior** evolution over **time**
- ... spot **disruptions**, phases
- ... relate time behavior to **space** dimension

Scalability issues solving

- Avoid **artifacts** due to aliasing
- Keep context
- Get control on information **loss**
- Keep reasonable **performance**

Embedded system analysis needs?

Visual representation that enables to...

- ... show **behavior** evolution over **time**
- ... spot **disruptions**, phases
- ... relate time behavior to **space** dimension

Scalability issues solving

- Avoid **artifacts** due to aliasing
- Keep context
- Get control on information **loss**
- Keep reasonable **performance**

Section

Time Aggregation Visualization : Presentation

Principle

- Trace is divided in **time slices**
- **Variable parameter** enables to aggregate **consecutive slices**
- **Aggregates** are related to phases, disruptions

Theoretical aspects

- Trace time-slicing (Schnorr)
- **Best-Cut partition** algorithm (Lamarche-Perrin)

Implementation

- C++ library (best partition algorithm)
- FrameSoC module/Java (GUI, database queries, time-slicing)

Time Aggregation Visualization : Presentation

Principle

- Trace is divided in **time slices**
- **Variable parameter** enables to aggregate **consecutive slices**
- **Aggregates** are related to phases, disruptions

Theoretical aspects

- Trace time-slicing (Schnorr)
- **Best-Cut partition** algorithm (Lamarche-Perrin)

Implementation

- C++ library (best partition algorithm)
- FrameSoC module/Java (GUI, database queries, time-slicing)

Time Aggregation Visualization : Presentation

Principle

- Trace is divided in **time slices**
- **Variable parameter** enables to aggregate **consecutive slices**
- **Aggregates** are related to phases, disruptions

Theoretical aspects

- Trace time-slicing (Schnorr)
- **Best-Cut partition** algorithm (Lamarche-Perrin)

Implementation

- C++ library (best partition algorithm)
- FrameSoC module/Java (GUI, database queries, time-slicing)

Time-Slicing: example of a synthetic trace

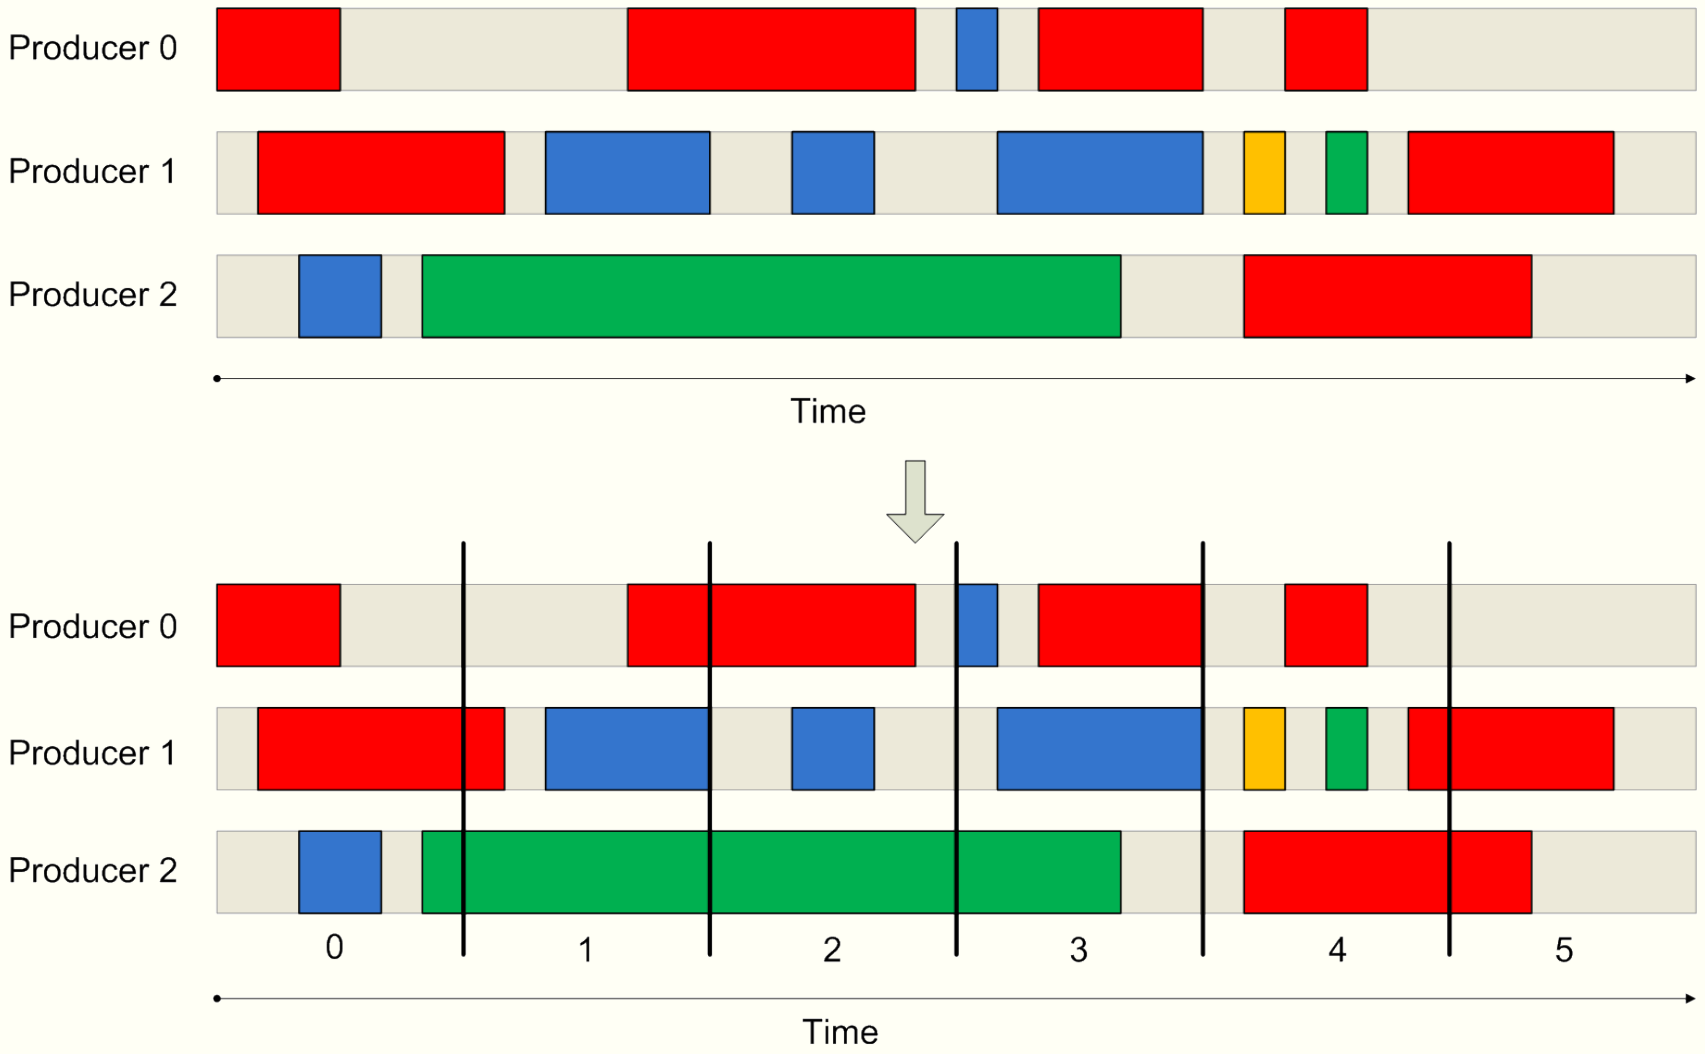

Time-Slicing: activity time matrix generation

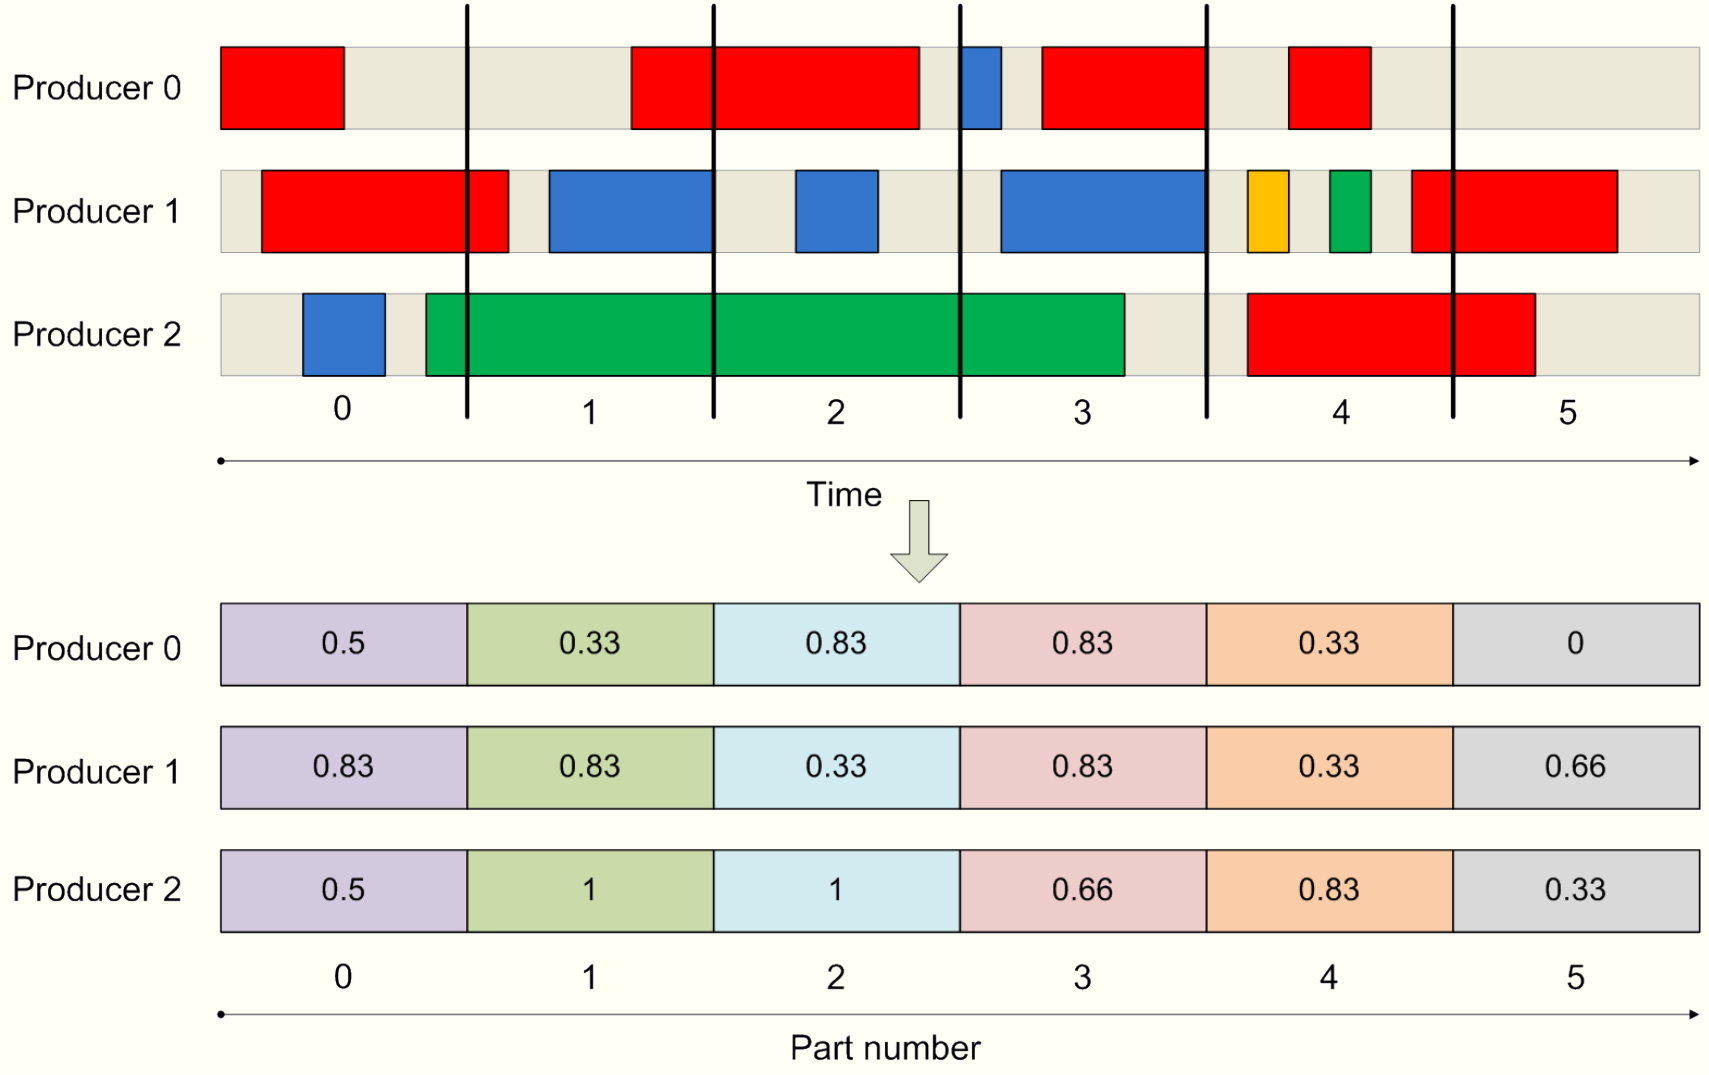

Best-Cut Partition algorithm: qualities

Qualities: Gain and loss

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

parametrized Information Criteria

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

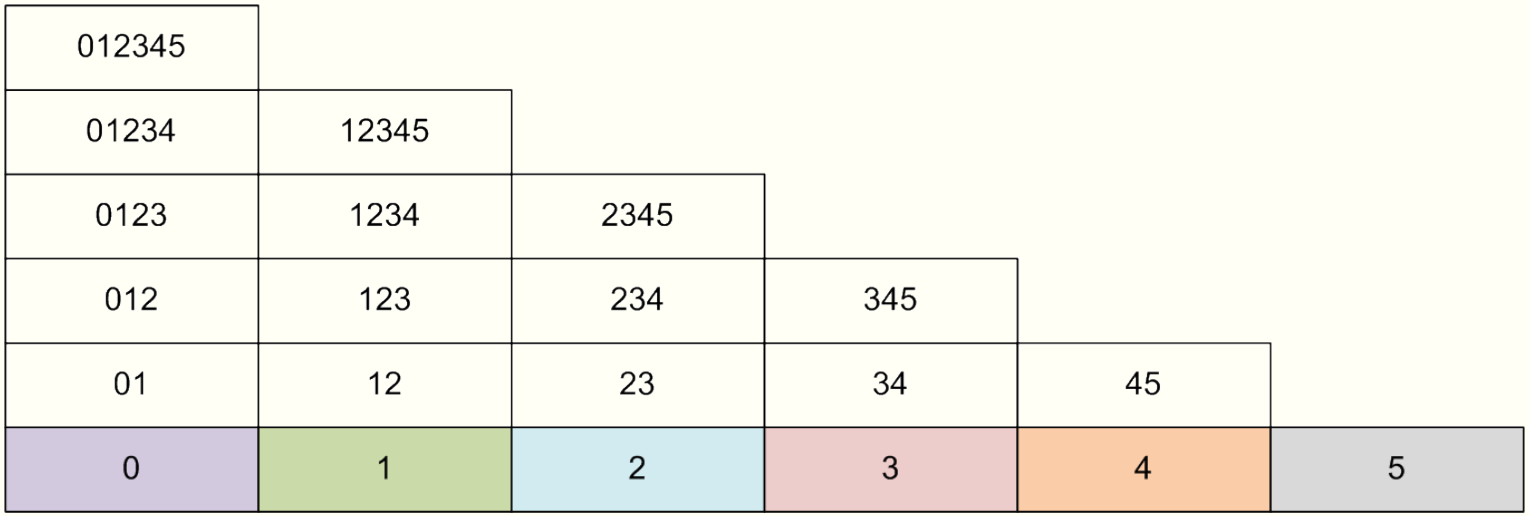

Best-Cut Partition algorithm: parts aggregation

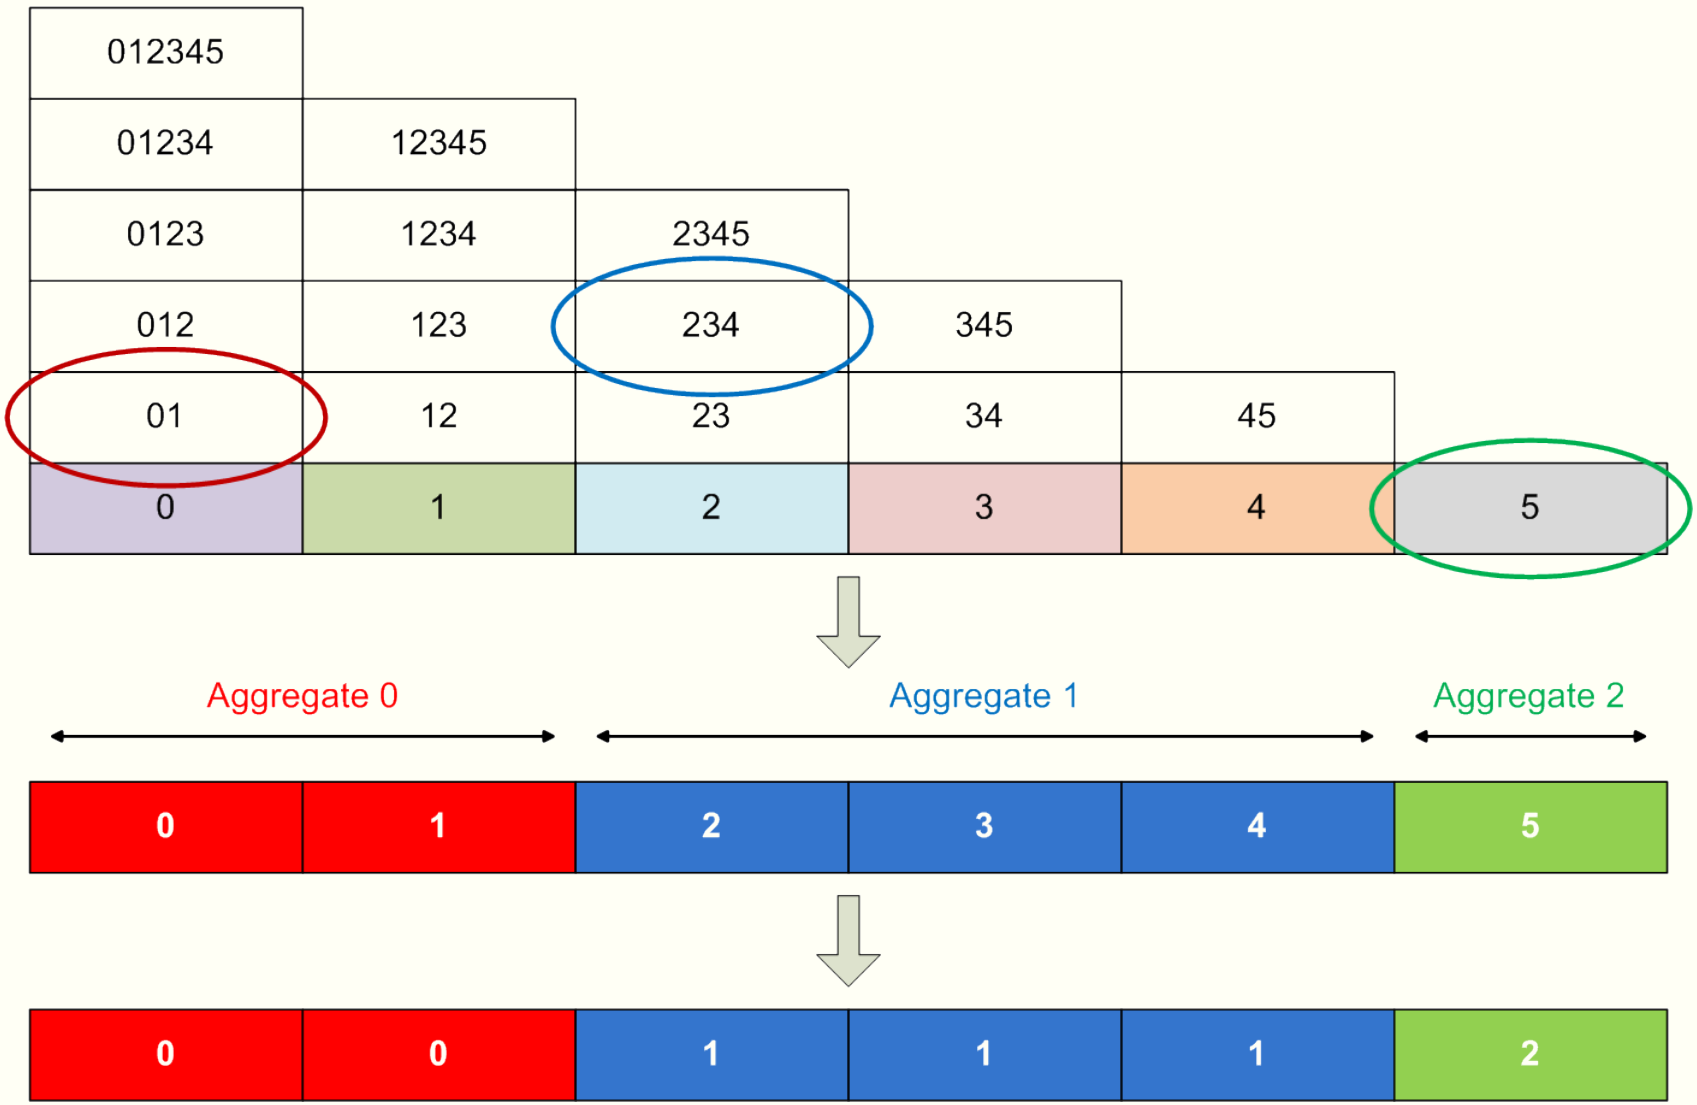

Tracing video execution: summary

Context

- GStreamer application playing a video, traced with `GST_DEBUG`
- Perturbation by *stress* program
- Trace converted into Pajé trace format
- Pajé trace imported to FrameSoC Data-Model

Use Case	Behavior	Duration	Trace Size	E.P. Number	Event Number
0 (ref)	Normal	20s	159 MB	1500	944303
1	Perturbation (@ 10s)	21s	167 MB	1500	990995
2	Perturbations (@ 7s, 14s)	26s	192 MB	1500	1140449
3	Light Perturbation (@ 15s)	21s	166 MB	1500	985003

Analysis with FrameSoC module

Section

Conclusion

Embedded Systems Trace Visualization Issues

- Space and time axis scalability problems
- Fidelity, reliability :-S
- Loss of information, coherence

Time Aggregation Visualization

- Able to represent application **behavior over time**
- Solves some **time scalability** issues

But...

- Lack of **space dimension** representation/aggregation
- Lack of **interaction** (details-on-demand)

Conclusion

Embedded Systems Trace Visualization Issues

- Space and time axis scalability problems
- Fidelity, reliability :-S
- Loss of information, coherence

Time Aggregation Visualization

- Able to represent application **behavior over time**
- Solves some **time scalability** issues

But...

- Lack of **space dimension** representation/aggregation
- Lack of **interaction** (details-on-demand)

Conclusion

Embedded Systems Trace Visualization Issues

- Space and time axis scalability problems
- Fidelity, reliability :-S
- Loss of information, coherence

Time Aggregation Visualization

- Able to represent application **behavior over time**
- Solves some **time scalability** issues

But...

- Lack of **space dimension** representation/aggregation
- Lack of **interaction** (details-on-demand)

Future Works

New features

- Discontinue parts **similarity**
- **Hierarchical** aggregation
- Aggregation **metrics**
- Visualization/parts representation improvement
- User **interaction**

Outside embedded system domain

- Does this representation fit to parallel system analysis needs?

Future Works

New features

- Discontinue parts **similarity**
- **Hierarchical** aggregation
- Aggregation **metrics**
- Visualization/parts representation improvement
- User **interaction**

Outside embedded system domain

- Does this representation fit to parallel system analysis needs?

Merci de votre attention!

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