ÉCOLE POLYTECHNIQUE MONTRÉAL

Trace Stream OLAP Analysis

Presented by: Naser Ezzati Jivan
Supervisor: Professor Michel Dagenais

DORSAL LAB

École Polytechnique, Montreal

Spring 2013

Trace Stream

- The importance of extending the existing tools and techniques to streaming mode, to capture today's needs.
  - Live performance monitoring
  - Live security monitoring
- Trace stream contains continuous high volumes (infinite) of events.
- Recording the whole trace or even a complete summary may not be possible.

Trace Stream

- Trace events usually represent low level behavior of the system.
  - Users might be interested in higher level analysis
- Trace events are multidimensional in nature.
  - A typical trace event represents interactions between different dimensions (i.e. system resources).
    - ”file read” event contains information of:
      - a file, running process, the current CPU , … .
    - processes or VMs use in the last 15 minutes, 30% more system resources than the last 24 hours average.
- Is multi-dimensional, multi-level OLAP Processing possible over trace stream?

OLAP?

- "**Online analytical processing,** or **OLAP** is an approach to answer multi-dimensional analytical (MDA) queries swiftly.
- OLAP tools enable users to analyze multi-dimensional data interactively from multiple perspectives." (wikipedia).
  - Usually used in Traditional DBMS.

- Is multi-dimensional, multi-level OLAP Processing possible over trace stream?

Challenges

- Large (infinite) number of events.
  - Referring back to the events is not possible.
- Keeping track of a long history is not possible.
  - Main memory or disk size limitations.
- Continuous Queries
  - answers are updated over time.
- Different levels and dimensions.
  - Different data aggregations

Solution: Cube Data Model

- Extending the state system to support multi-level multi-dimensional stream based queries.
  - How to model the cubes (dimensions and measures)?
  - How to materialize the cubes along the time axis?
  - How long we can keep history?
  - What types of queries are supported?

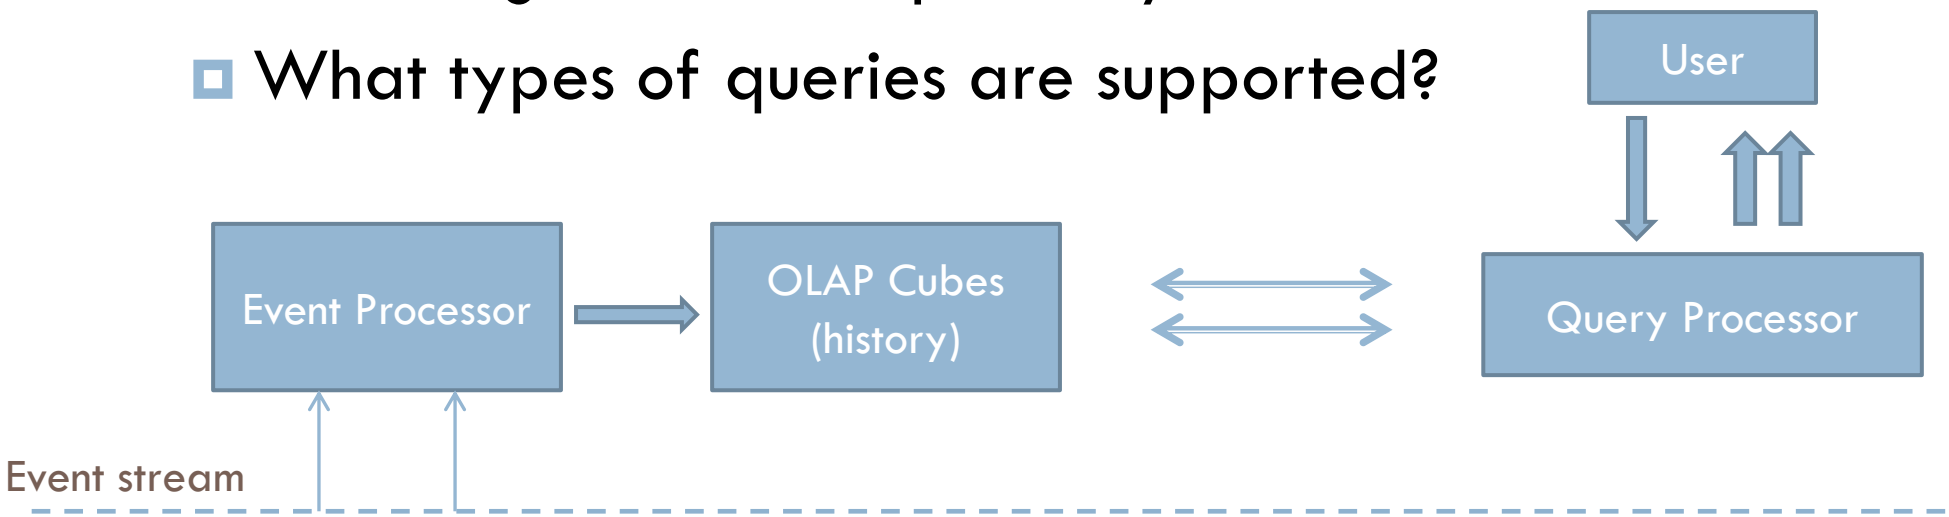

How to model the cubes?
(Dimensions and Measures)

- Dimension Schema

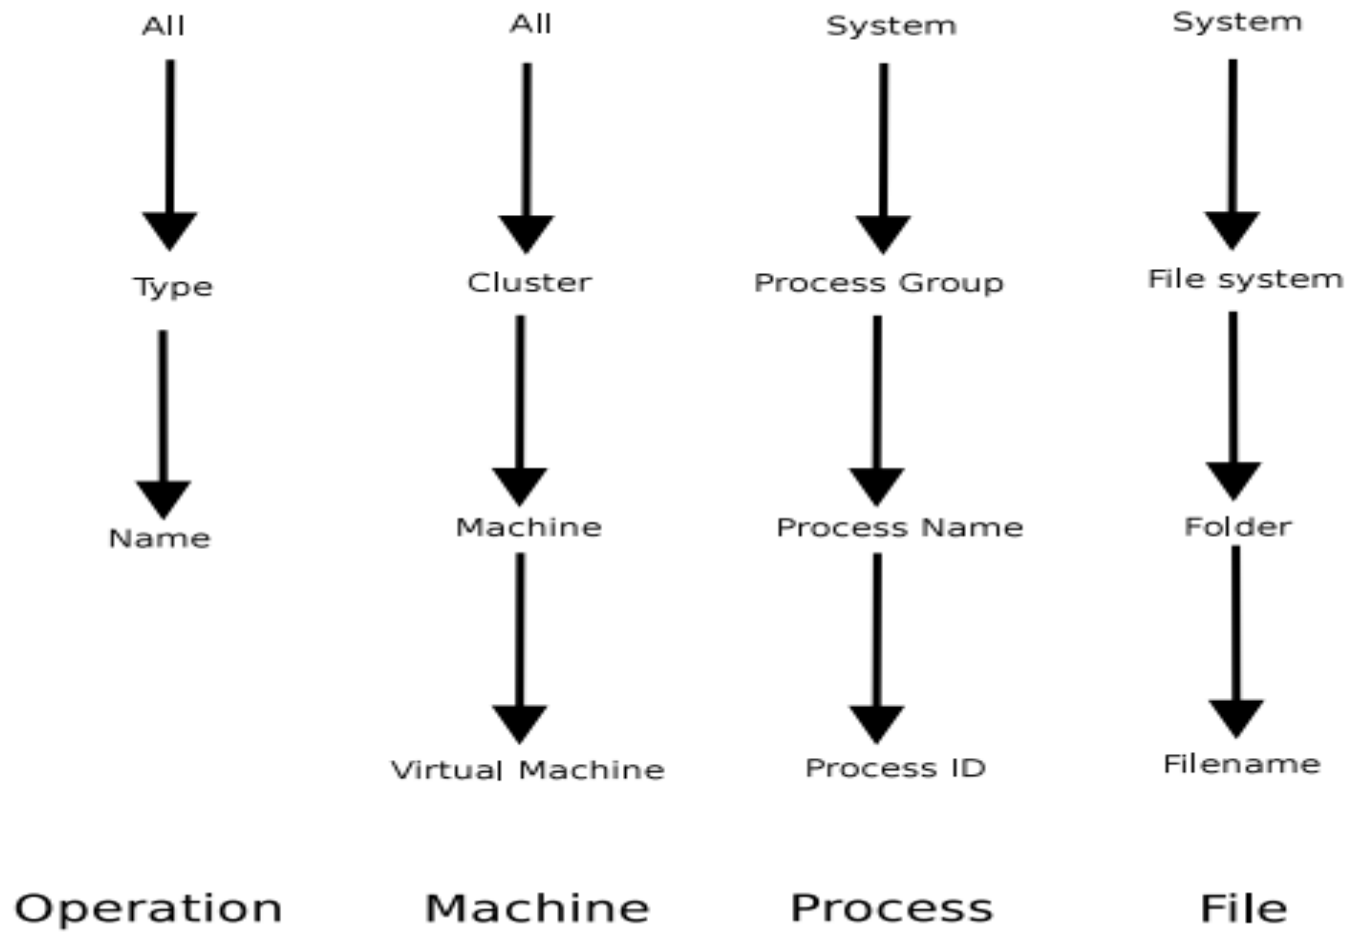

How to model dimensions?

- Dimension instance: a set of members from all levels.

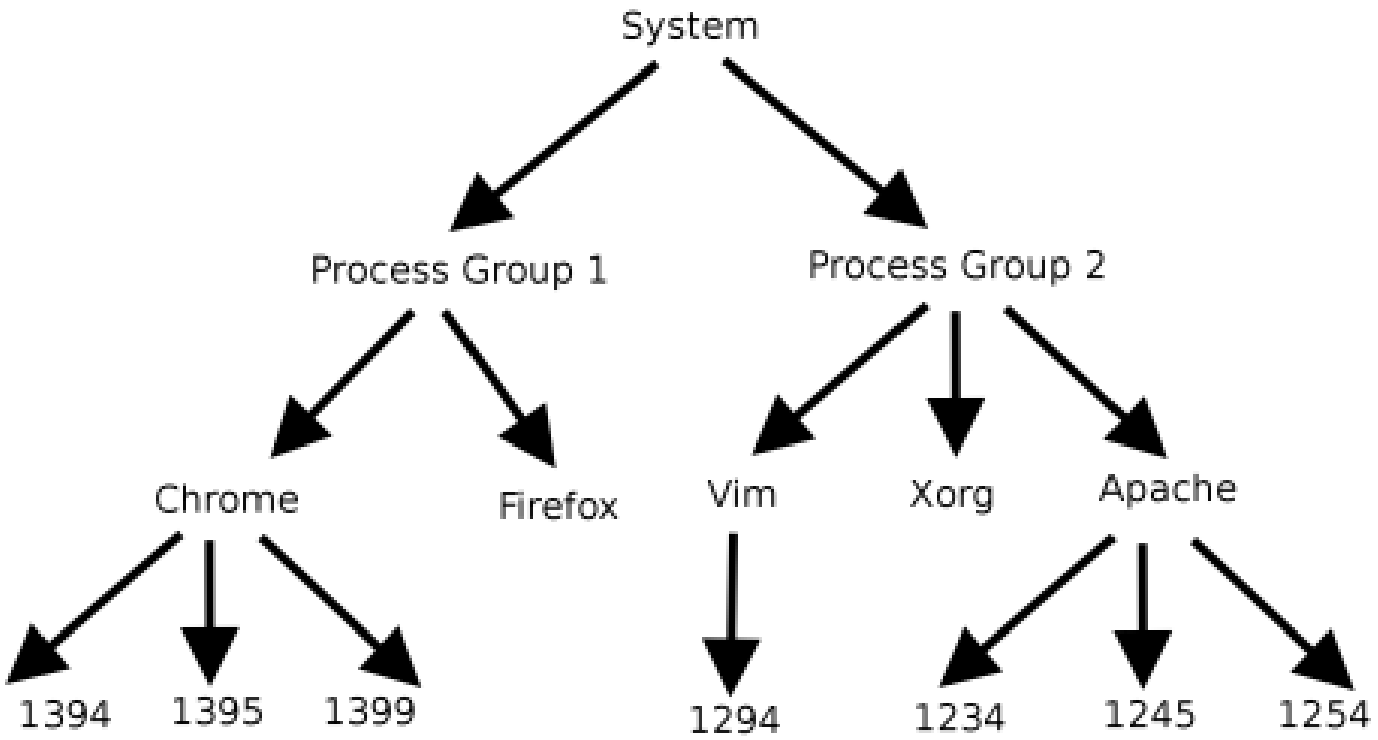

Measures

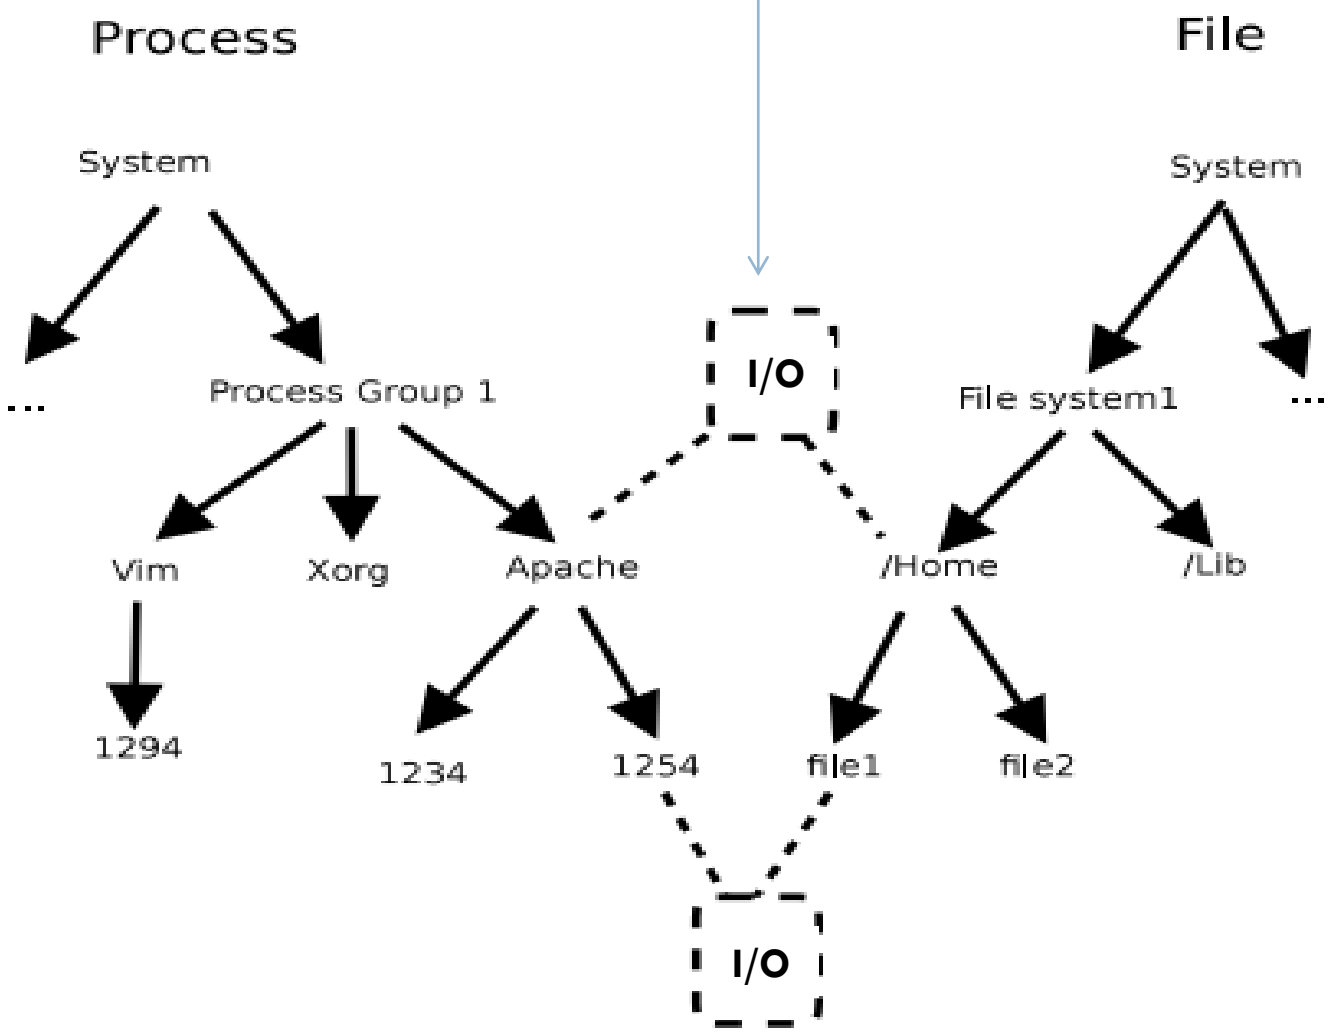
Process
File
System
System
...
Process Group 1
I/O
File system1
...
Vim
Xorg
Apache
/Home
/Lib
1294
1234
1254
file1
file2
I/O

Measures at Different Levels

- Selection is based on:
  - User/App requirements
  - Available resources
    - Storage,
    - time ,
    - …
  - Static / Dynamic

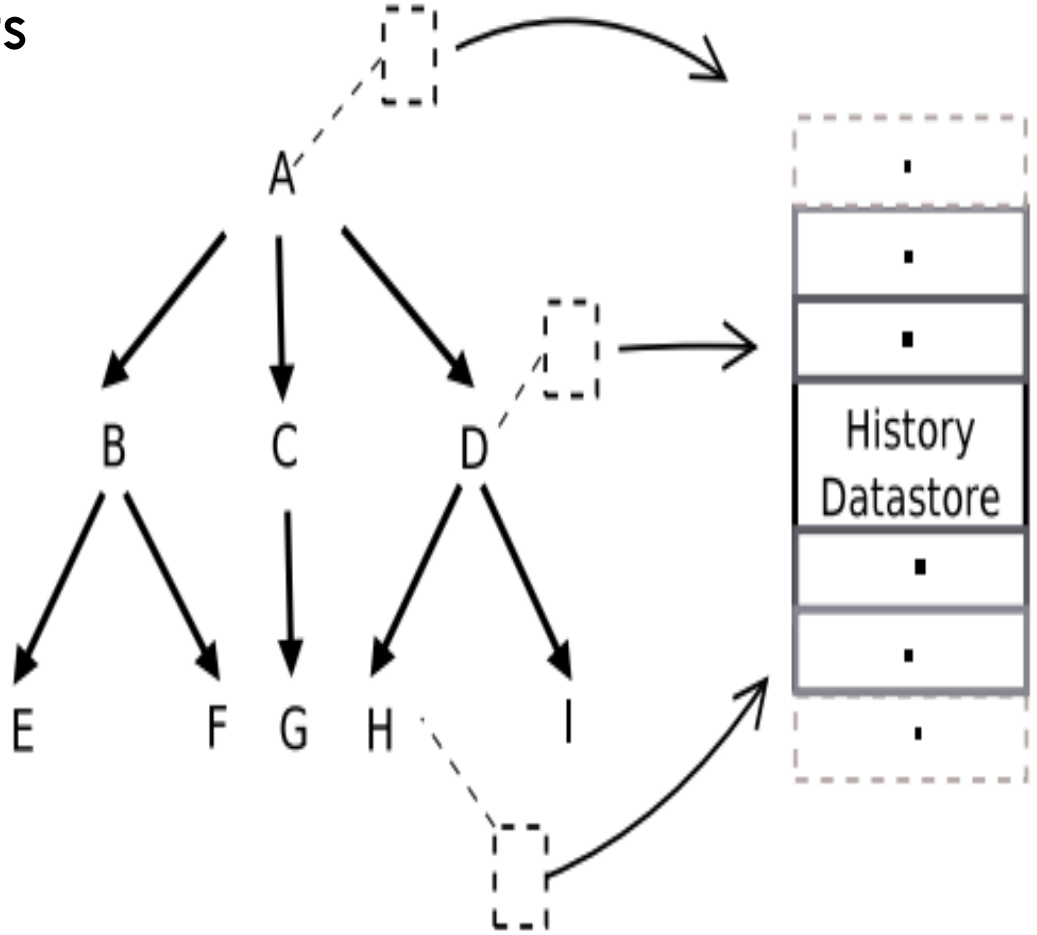

Cube Materialization

- **Minimal** materialization
  - Only the finer level nodes are materialized.

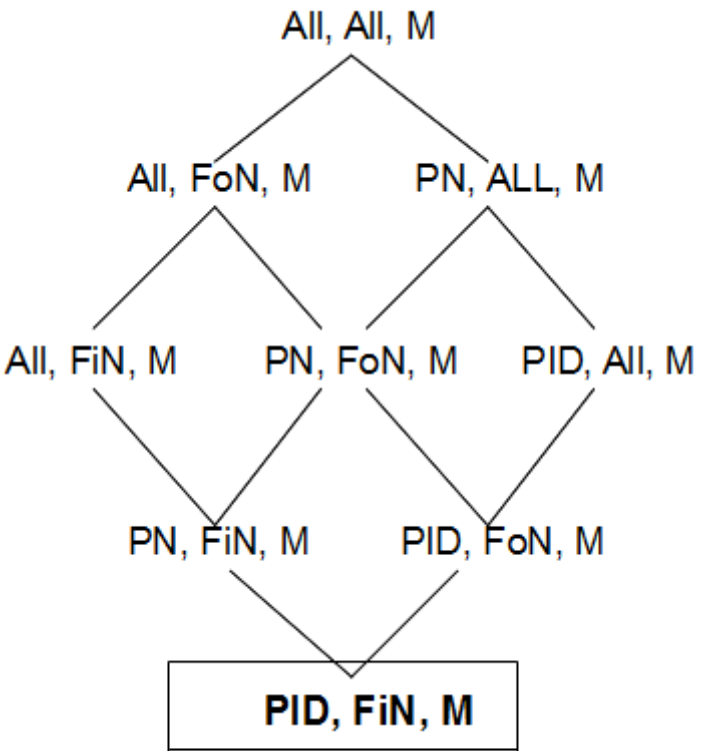

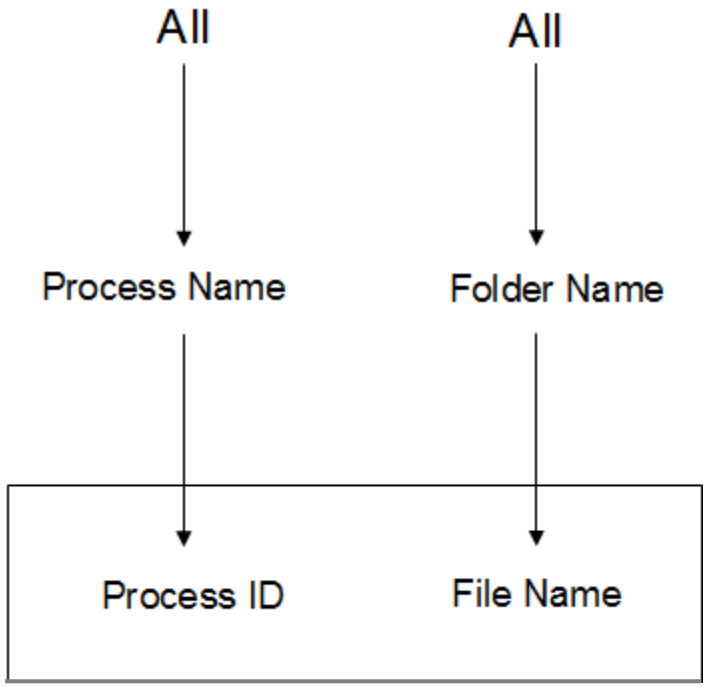

Cube Materialization

- **Partial** materialization
  - The finer level and some high level nodes are materialized.

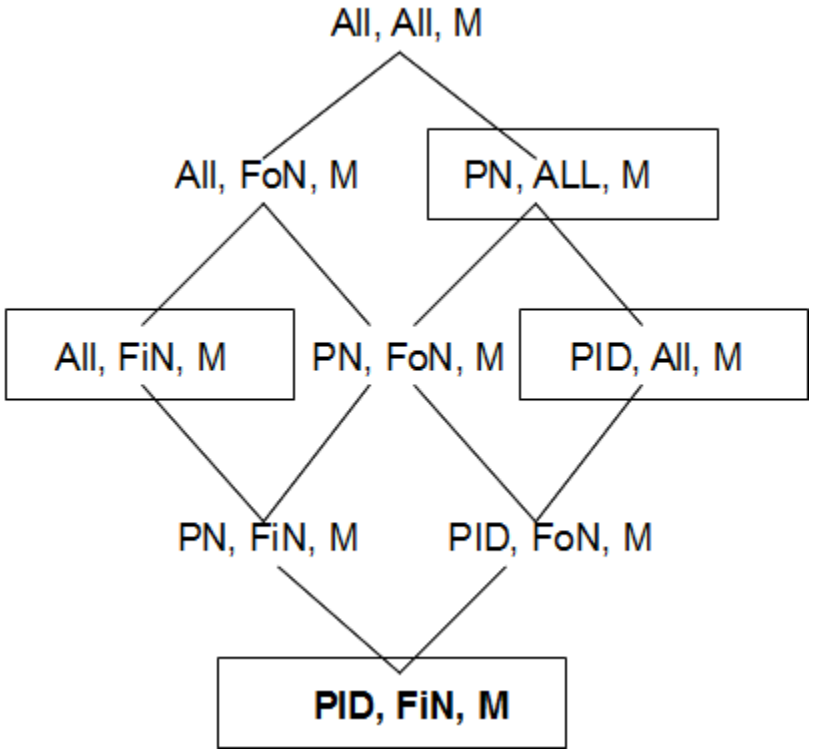

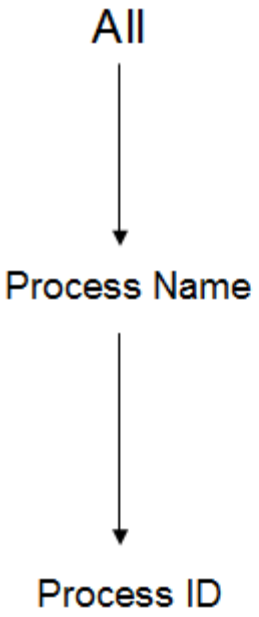

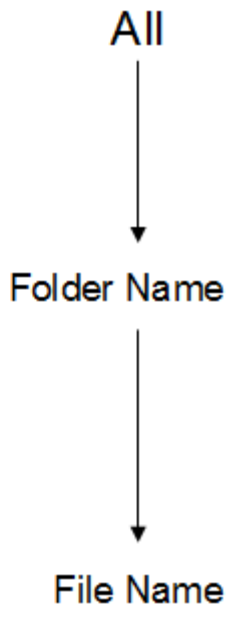

Cube Materialization

- **Maximal** materialization
  - All levels are materialized.

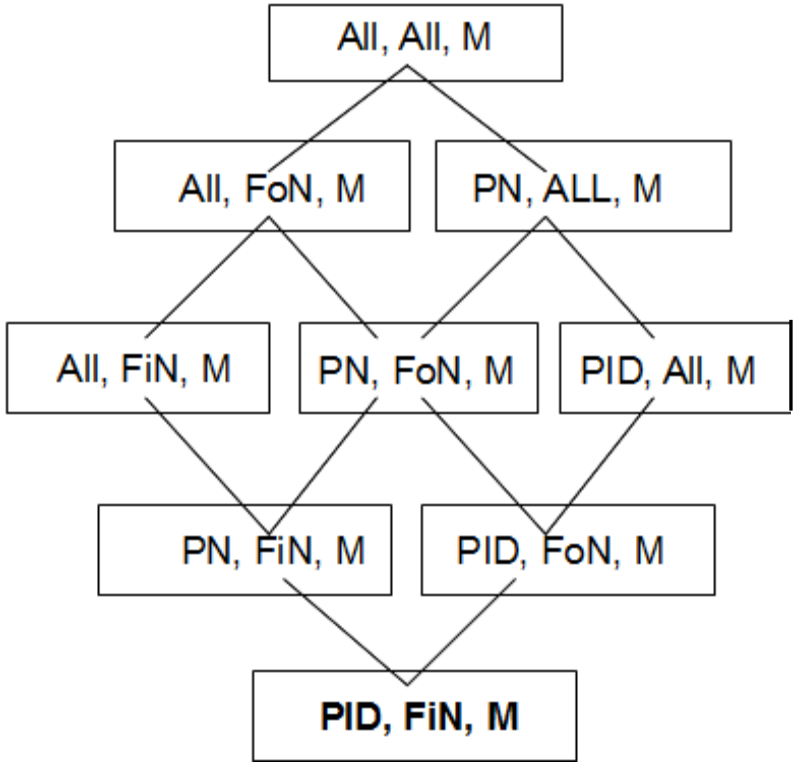

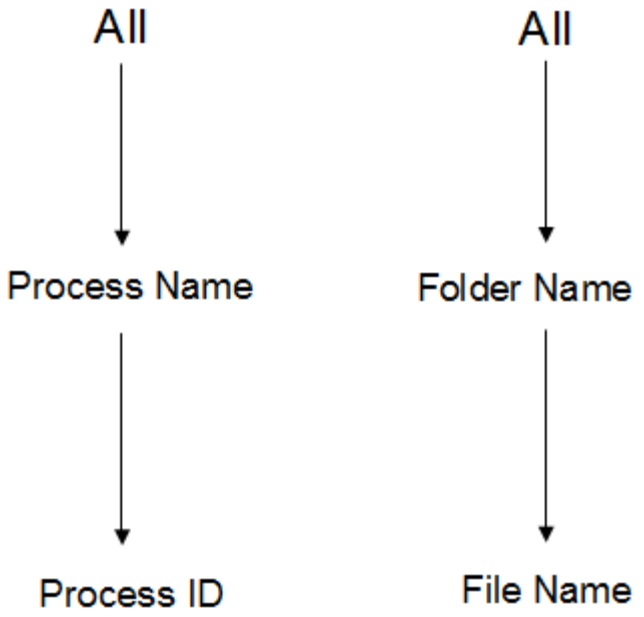

Cube Materialization & Time Dimension

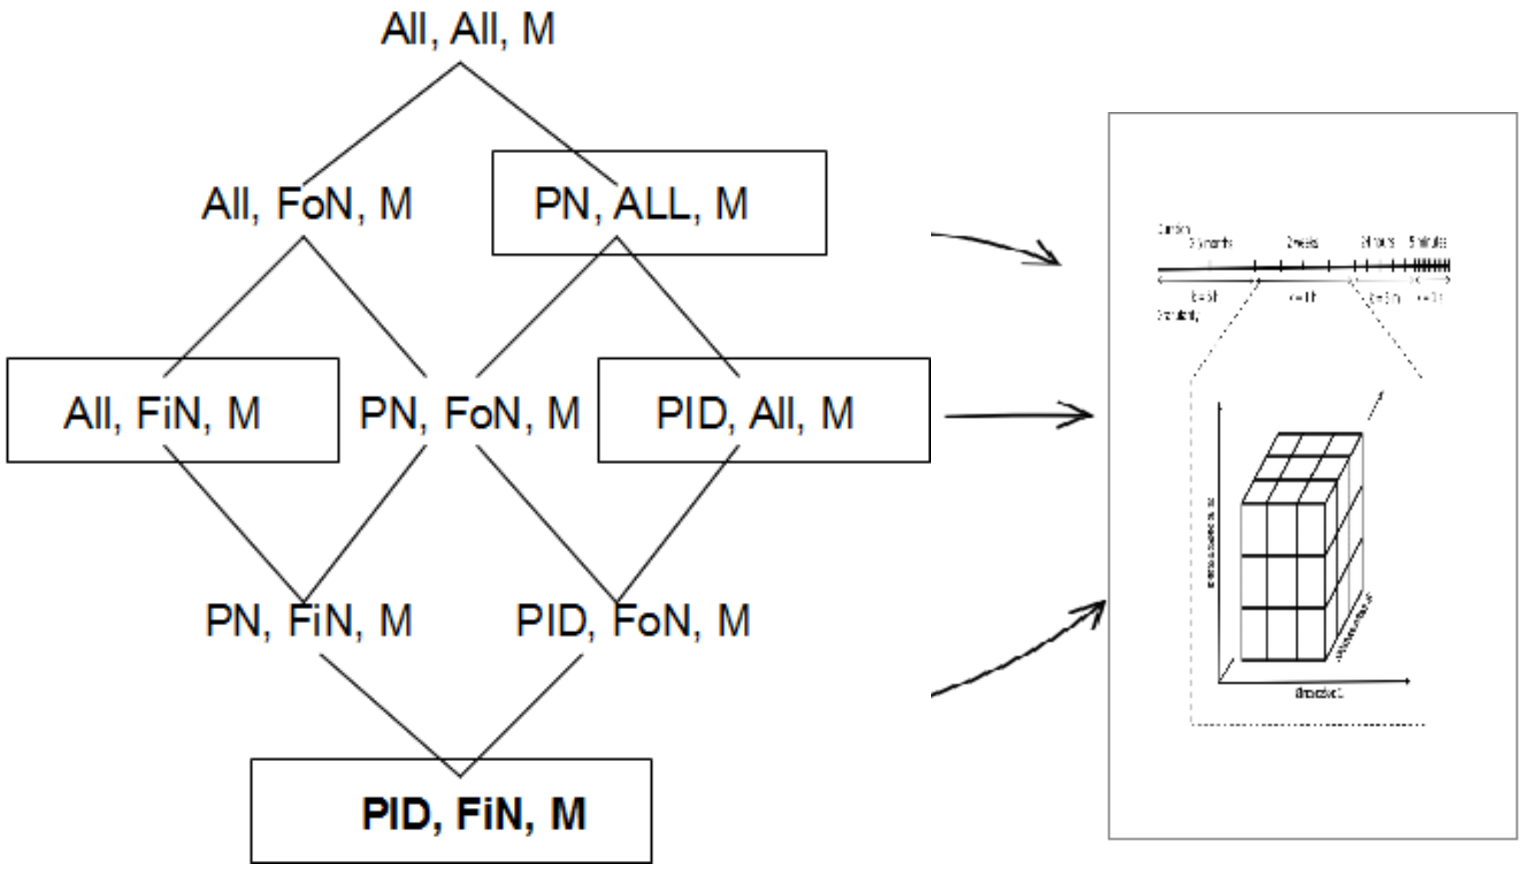

How long we can keep history?

- Memory or disk size limit.
- Solutions:
  - Cut off after a while
    - Decent data might be needed
  - **A New Time Frame !**

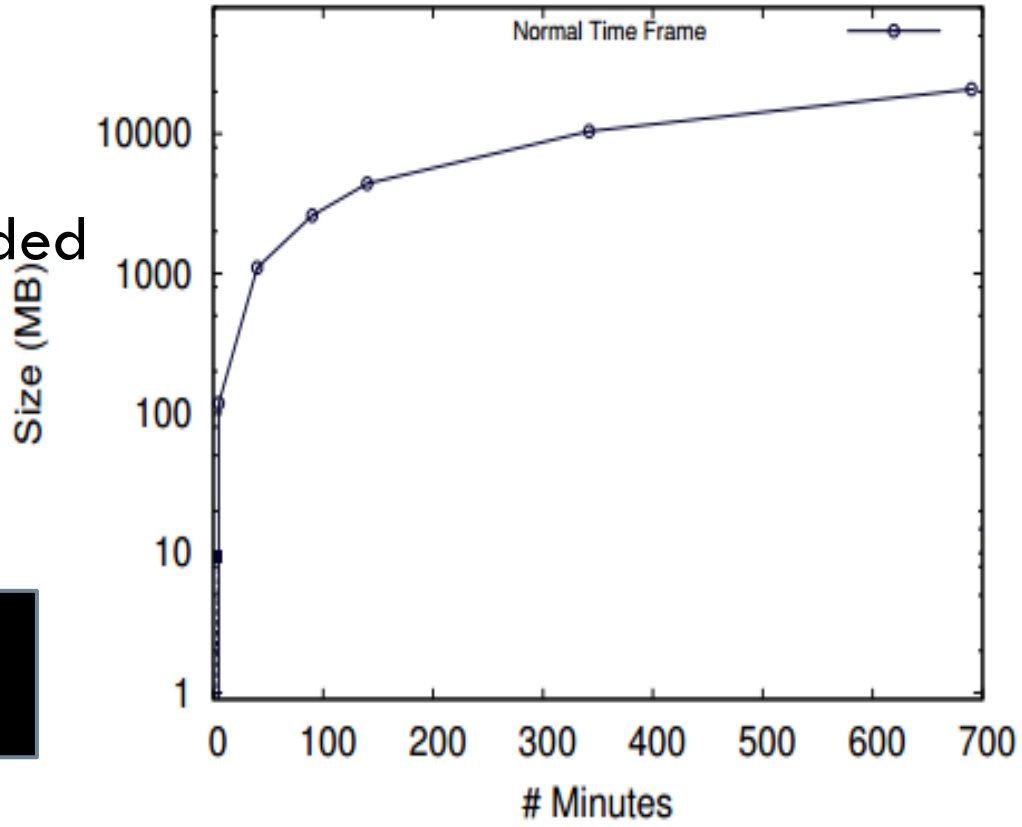

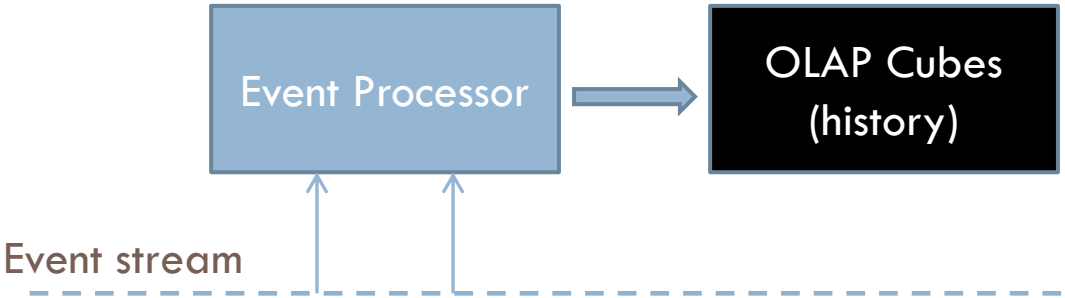

Comparison

Normal Time Frame

T

Granularity

Tilted Time Frame

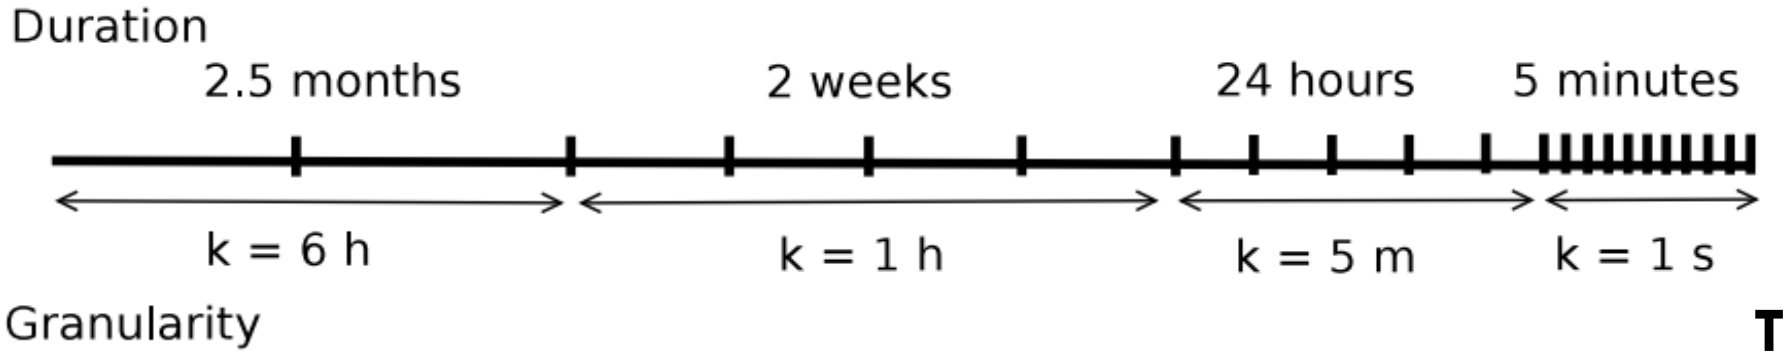

Tilted Time Frame

- Different time granularities for recent and decent information.
  - Coarser granularity for the too old history, but a finer granularity for the most recent history.
  - Compress the data over the time dimension.

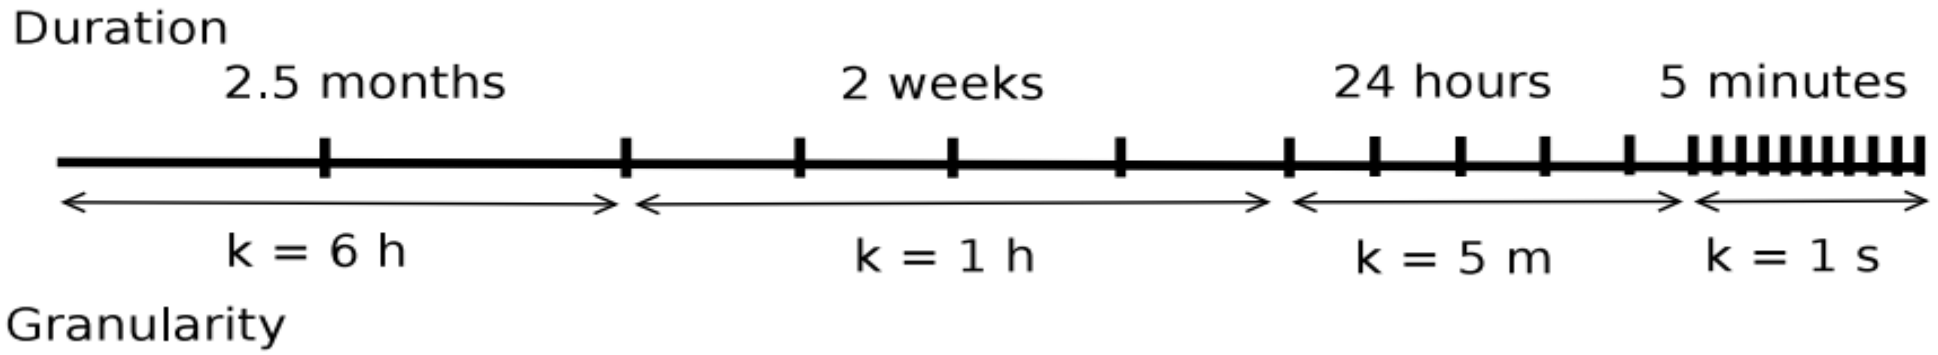

Comparison

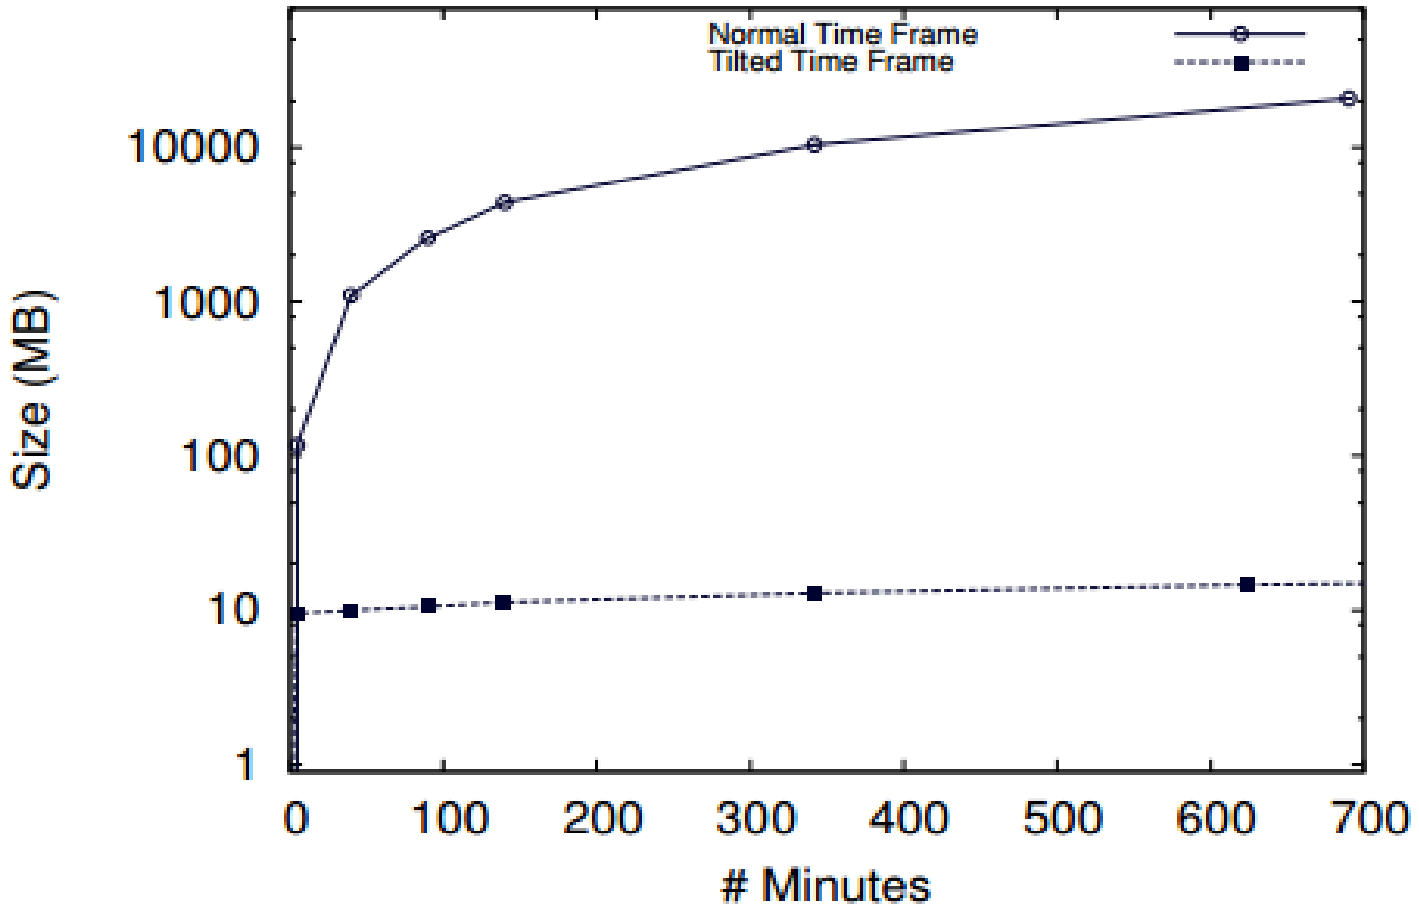
Normal Time Frame
Tilted Time Frame
Size (MB)
10000
1000
100
10
1
0
100
200
300
400
500
600
700
Minutes

Cube Materialization & Time Dimension

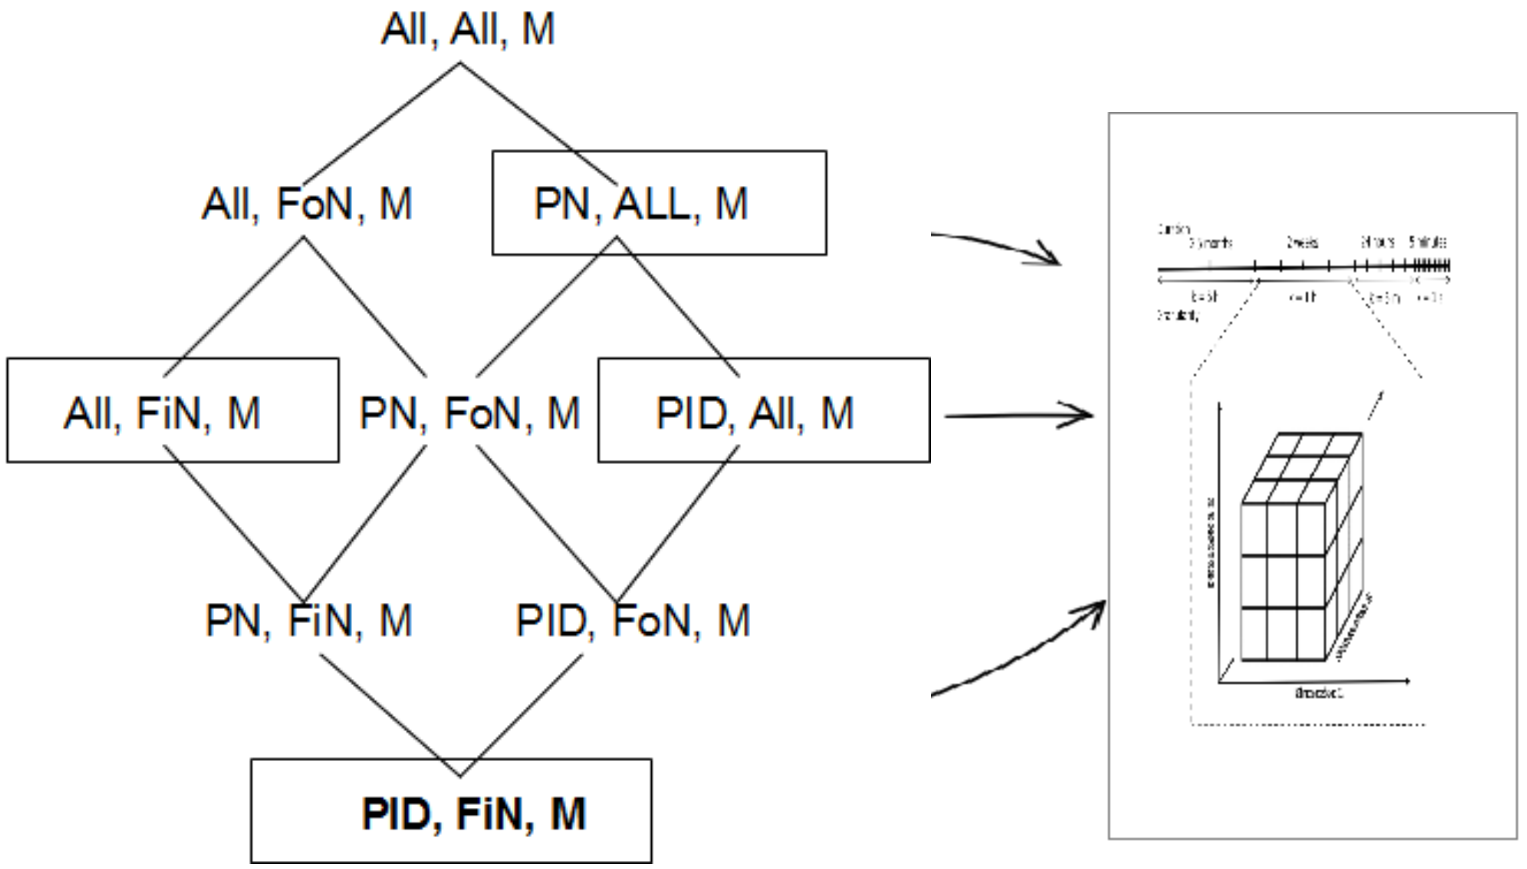

Several History Cubes

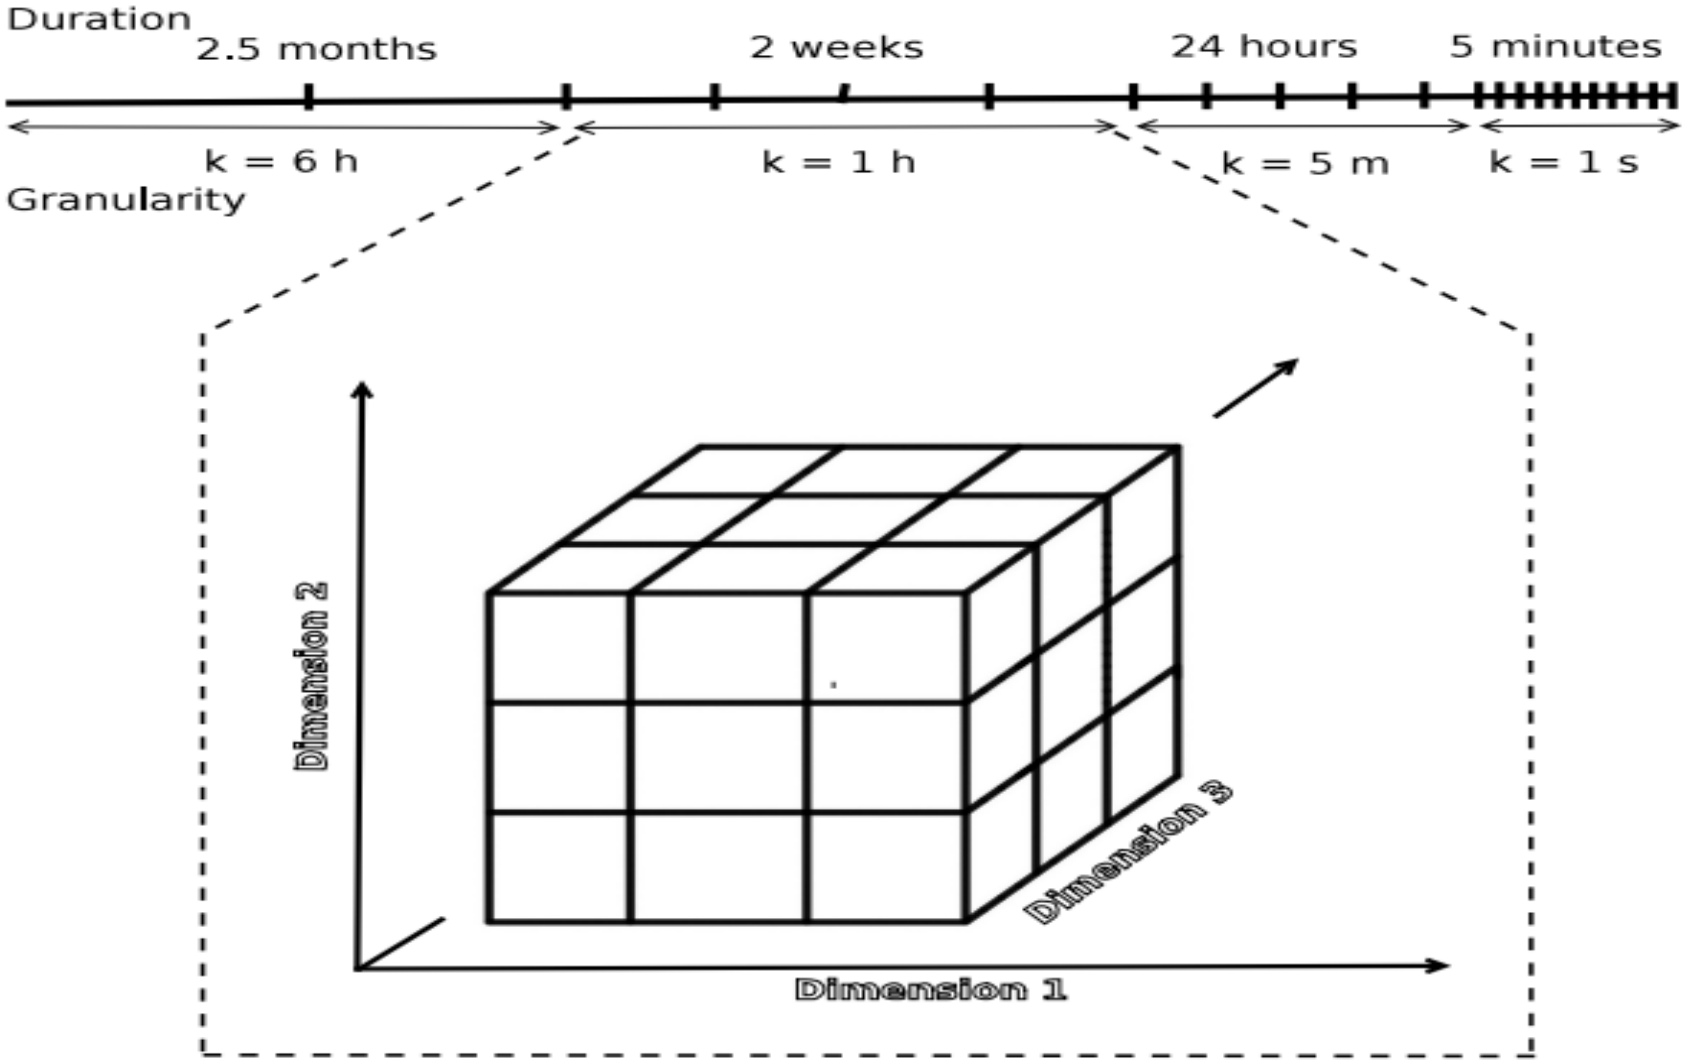

History Updates

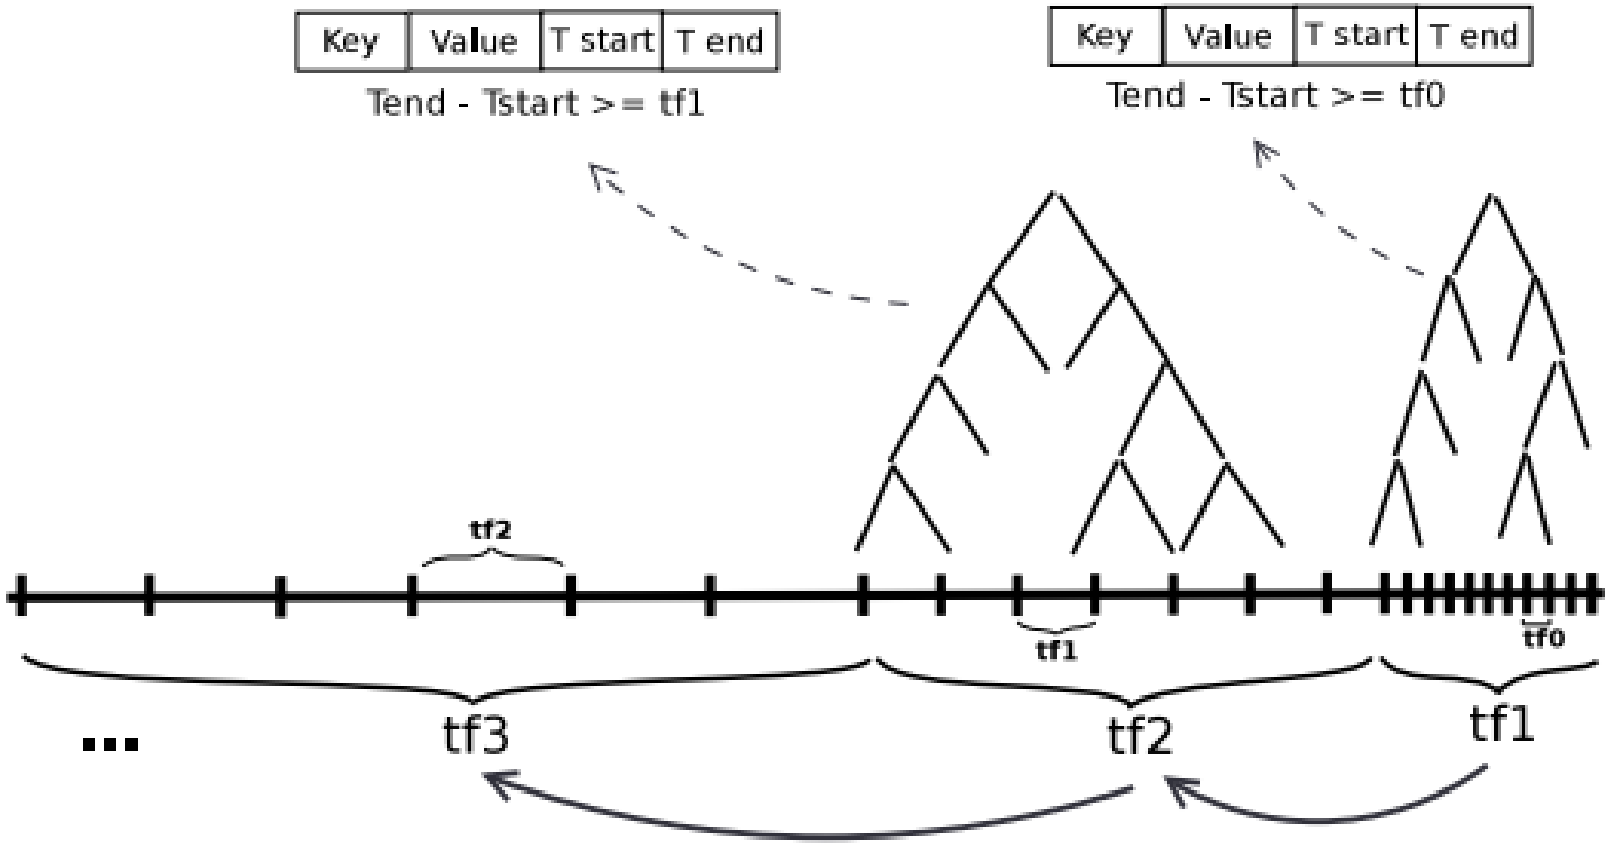

Query Types

Query Types (1)

- Range Queries (multi-dimensional)

Algorithm 1 Stabbing query.

Require: an interval tree v and a query point t.

```
1: if root node r contains the point t then
2:     add r to the result list L.
3: end if
4: if there is any children for node r then
5:     for any children of v like c(v) do
6:         call the algorithm for c(v) , t.
7:     end for
8: end if
9: return list L as result;
```

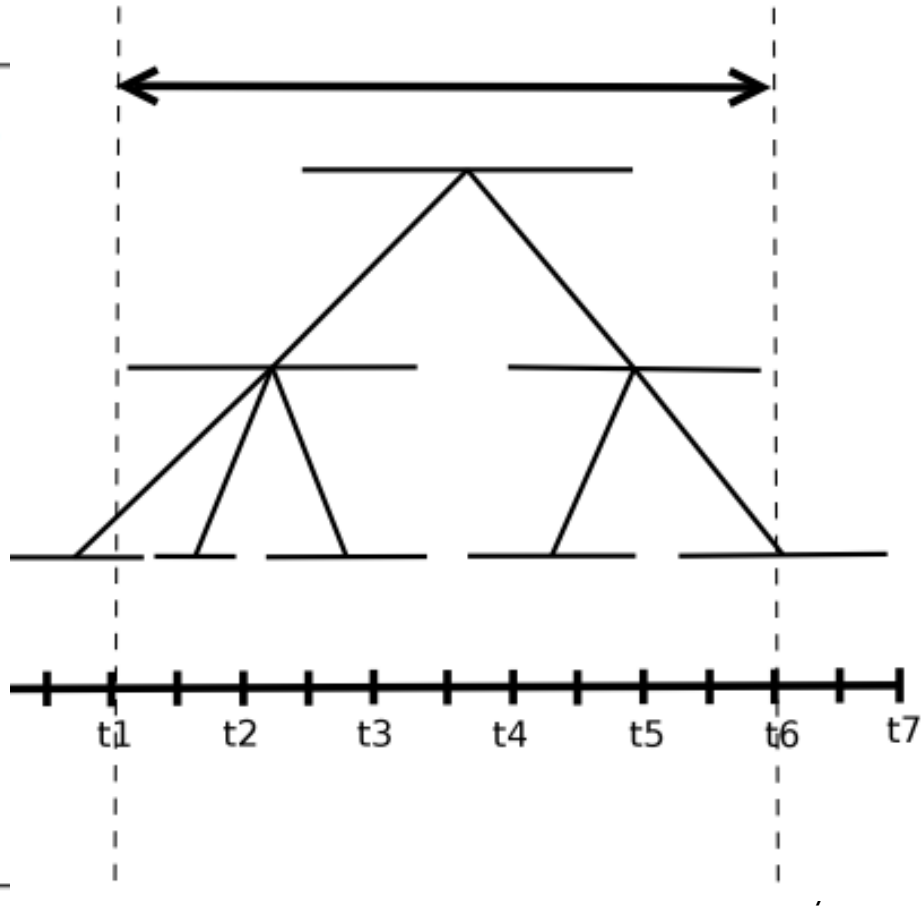

Query Types (2)

- Top-K Queries
  - Find virtual machines or processes which use more system resources than others
  - Values exceed a pre-defined threshold.
  - Algorithm:
    - Range query + Sort

Query Types (3)

- Sliding Window Queries (**Continuous queries**):
  - Statistics values for the last k time units, for the fixed or moving values of k.
- **Fixed**: reporting the statistics after each k time units.
  - Summarize data using a set of buckets. Can be used to draw histograms.
  - Example: network throughput each 1 second or each 1 minute of execution

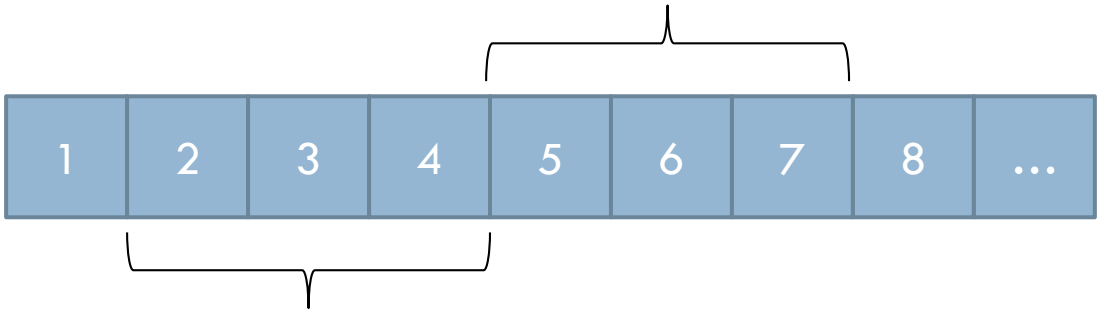

Query Types (3)

- Sliding Window Queries (**Continuous queries**):
  - Statistics values for the last k time units, for the fixed or moving values of k.
  - **Moving**: recent k time units.
    - Example: CPU usage for the last 3 minutes,
      - that is updated at each second!

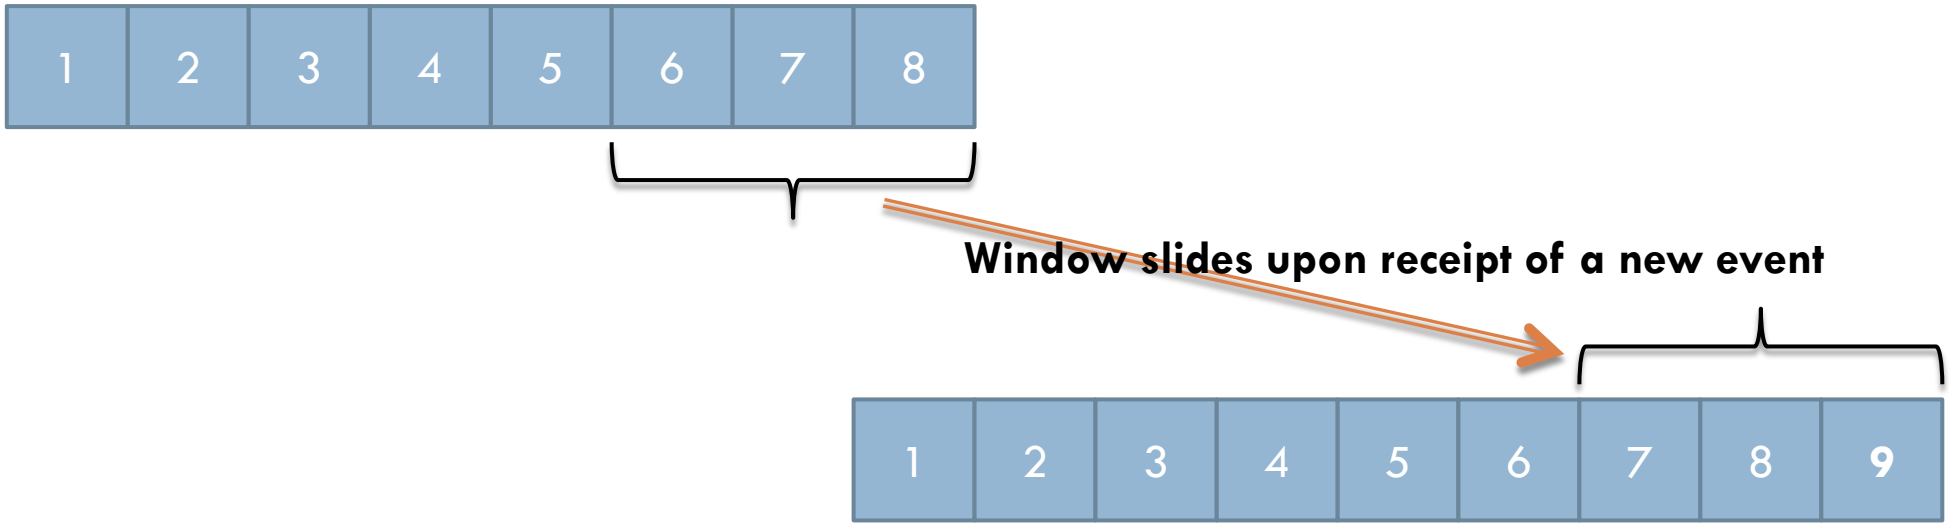

Query Types (4)

- Multilevel Queries
  - **Group by, drill down or roll up operations over any dimensions.**
  - With minimum materialization strategy:
    - Any other level statistics (higher levels) are computed on the fly using the leaf nodes.
  - With partial materialization strategy:
    - Using directly for the nodes that are stored in the database + on the fly for the other values
  - With full materialization strategy
    - All queries are answered directly using the stored values.

Query Types (4)

- Memory usage

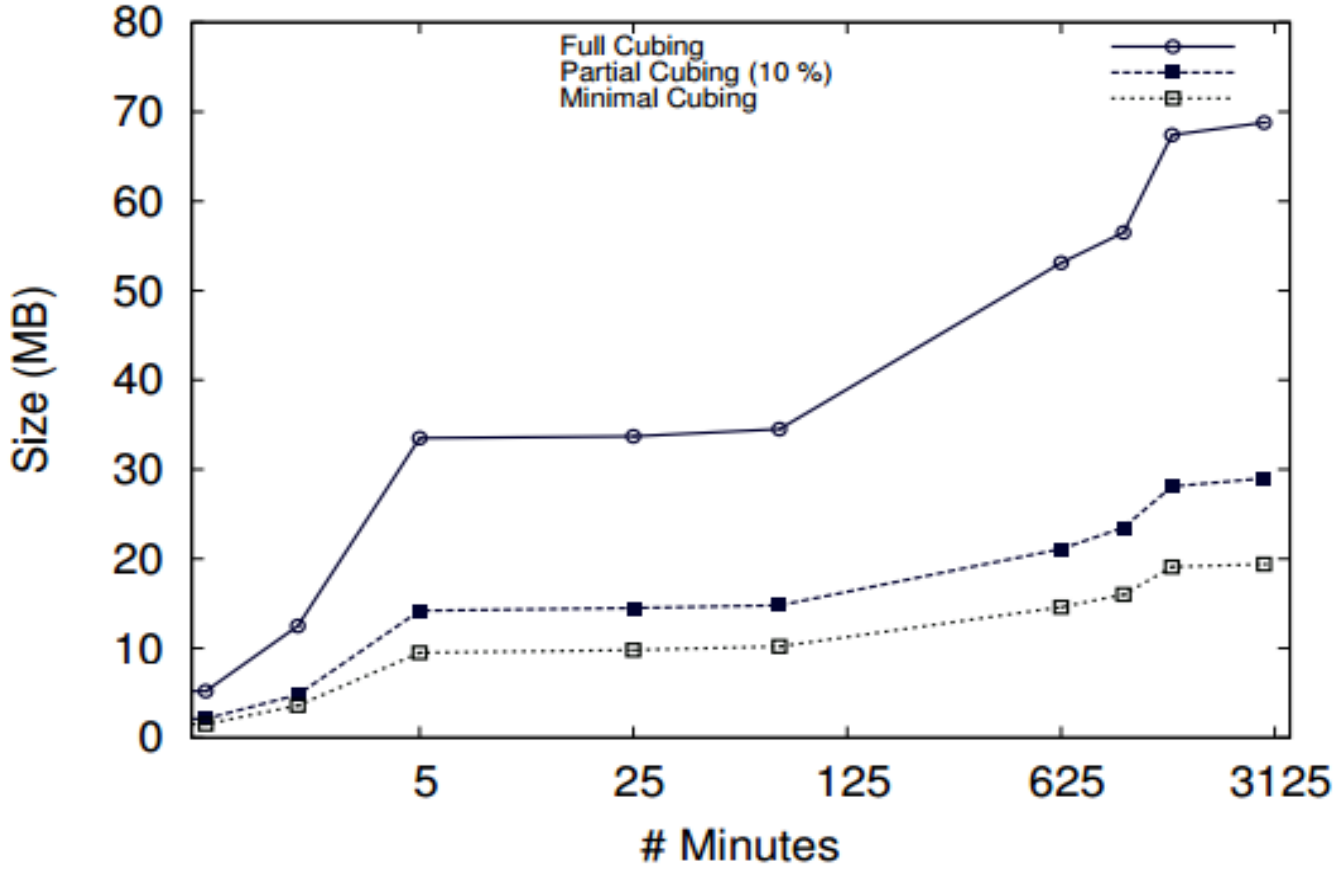

Memory usages for different cube materialization strategies.

Query Types (4)

- Response time

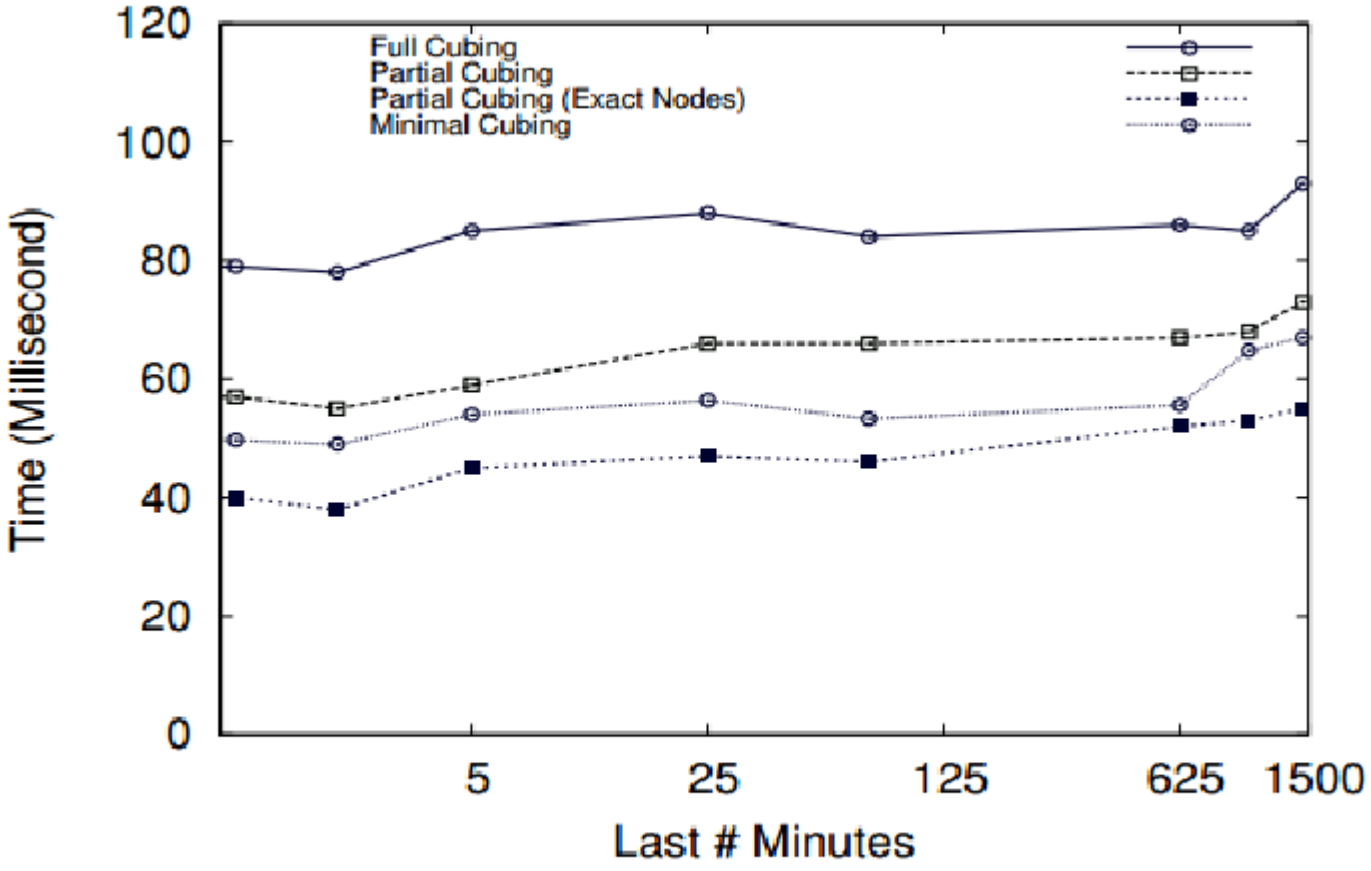

Response time for roll-up queries.

Conclusion

- The “Cube Data Model” and corresnponding algorithms are presented for performing OLAP analysis over trace streams.
  - As an extension to the **STATE SYSTEM.**
  - Tilted time frame instead of normal time frame.
  - Different Query types:
    - Multi-dimensional, multi-level .
    - Fixed and moving sliding windows queries.
    - Top-k queries.
    - Range queries.

n.ezzati@polymtl.ca

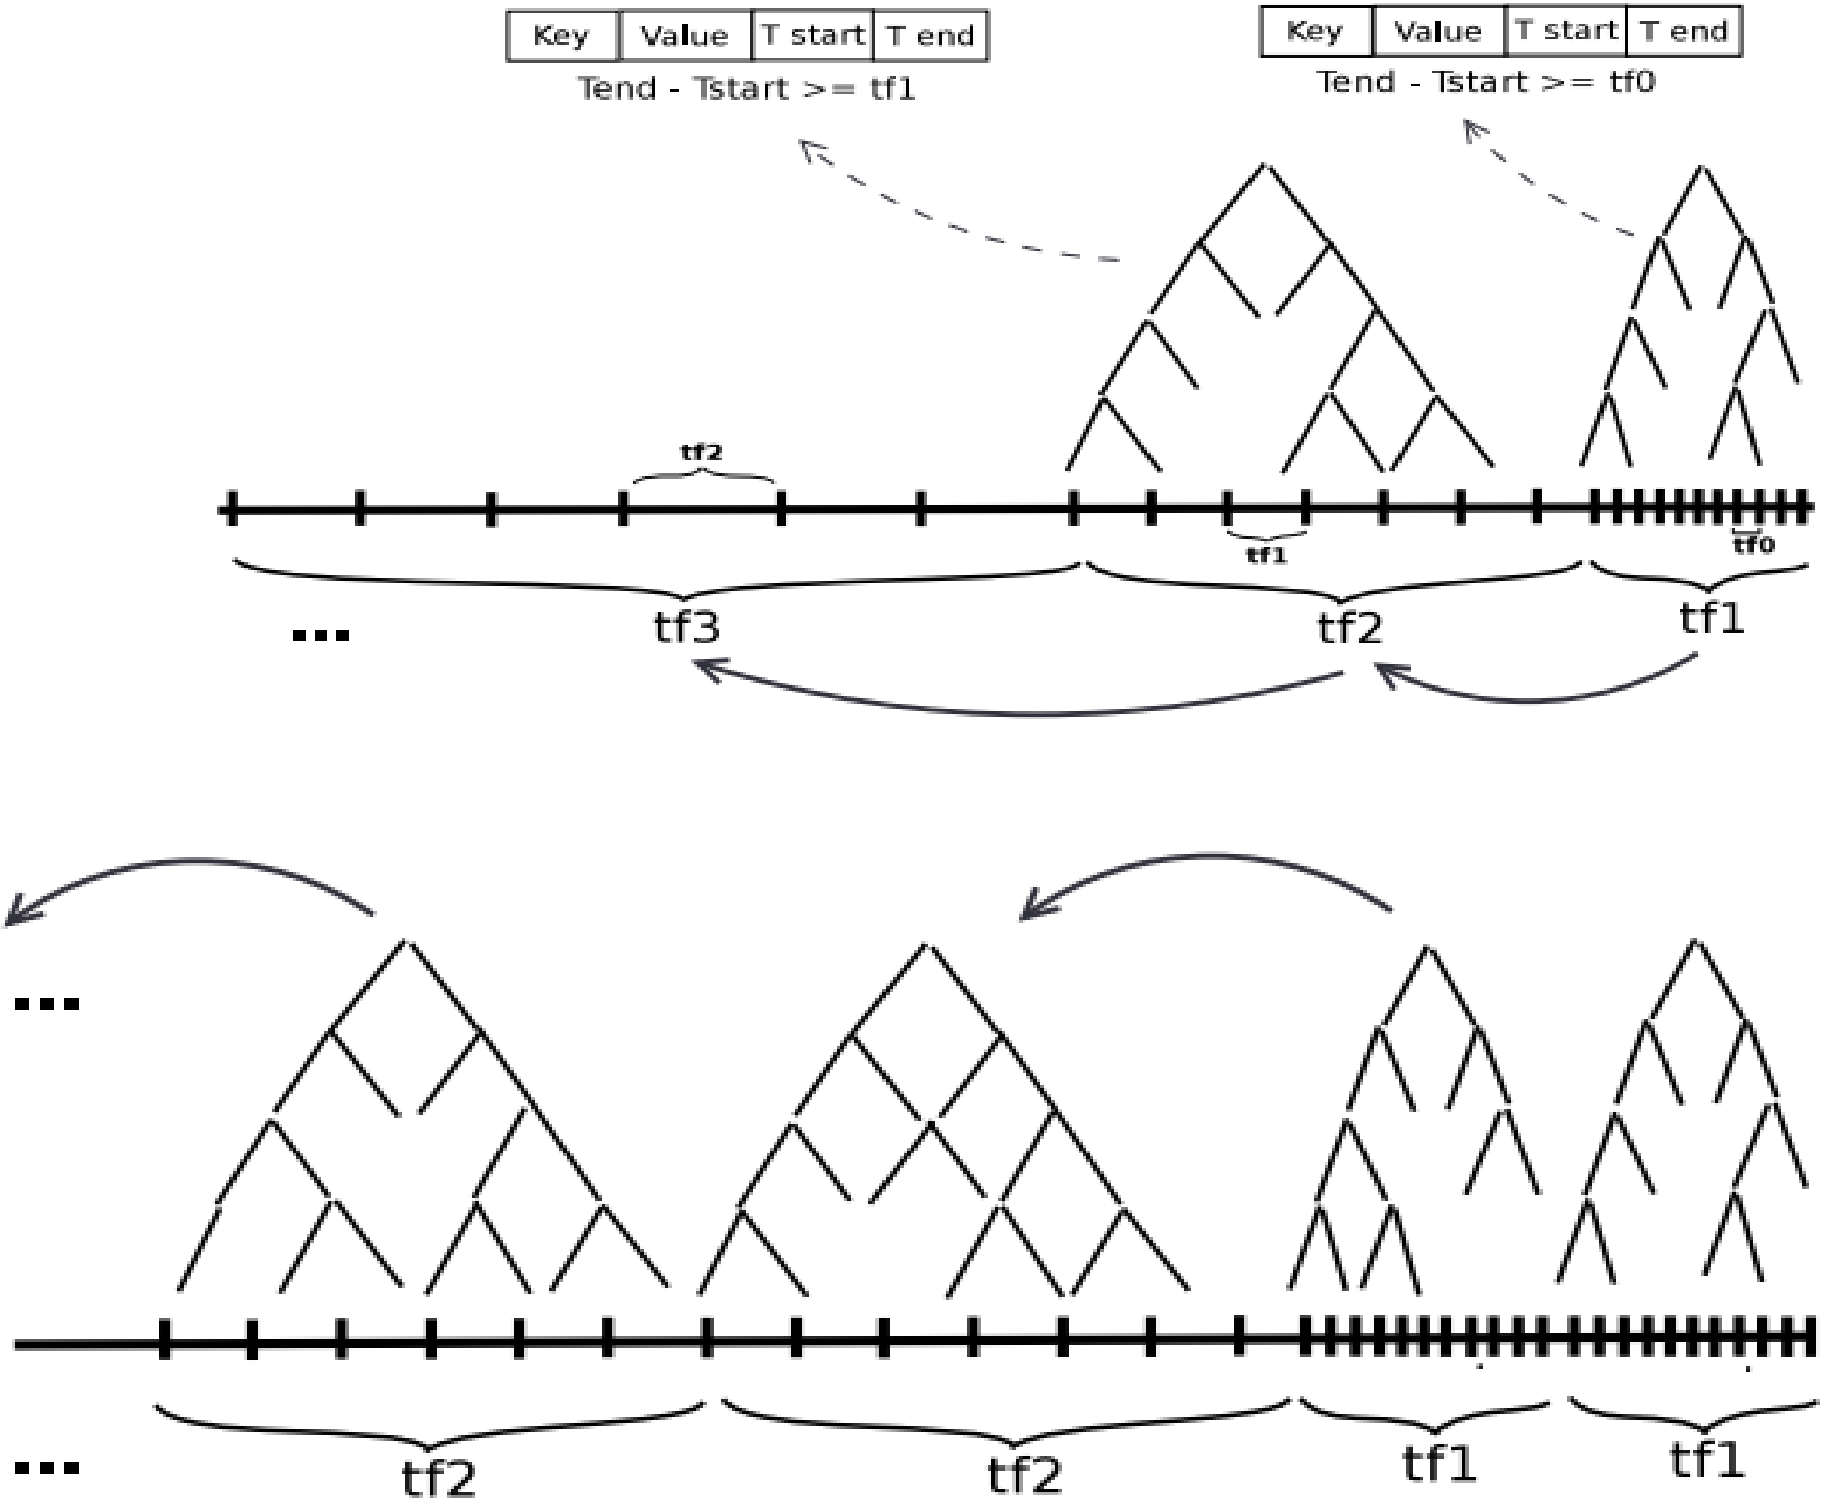
Key
Value
T start
T end
Tend - Tstart >= tf1
Key
Value
T start
T end
Tend - Tstart >= tf0
tf2
tf1
tf0
...
tf3
tf2
tf1
...
...
tf2
tf2
tf1
tf1