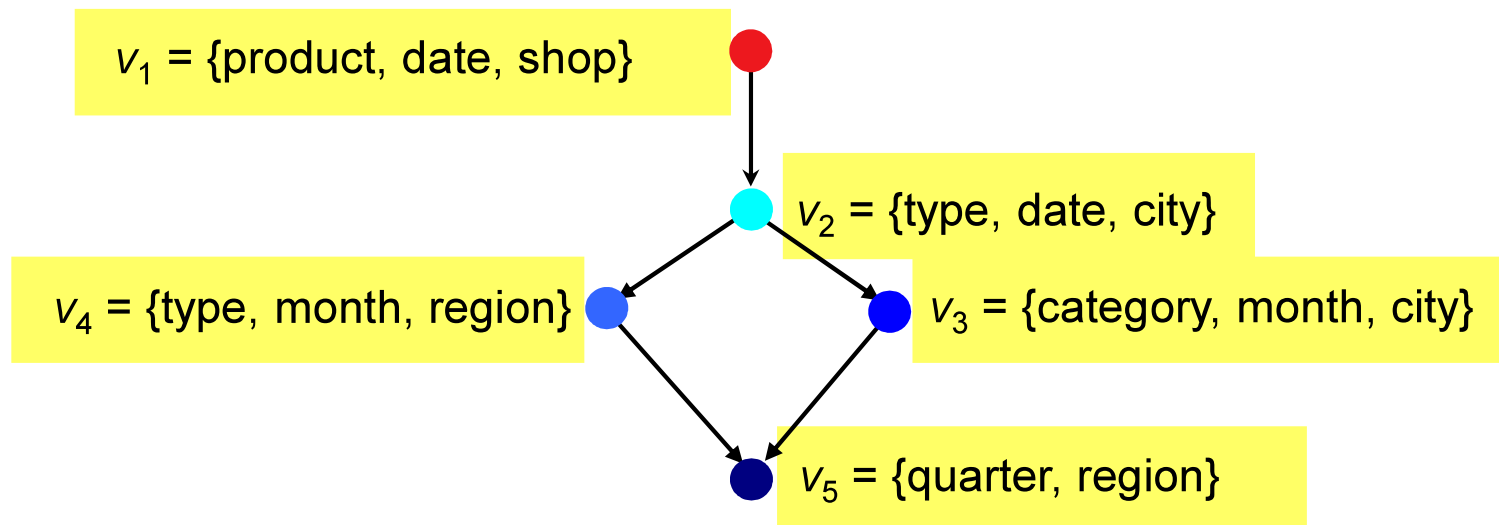


# *Materialized views*

Elena Baralis  
Politecnico di Torino

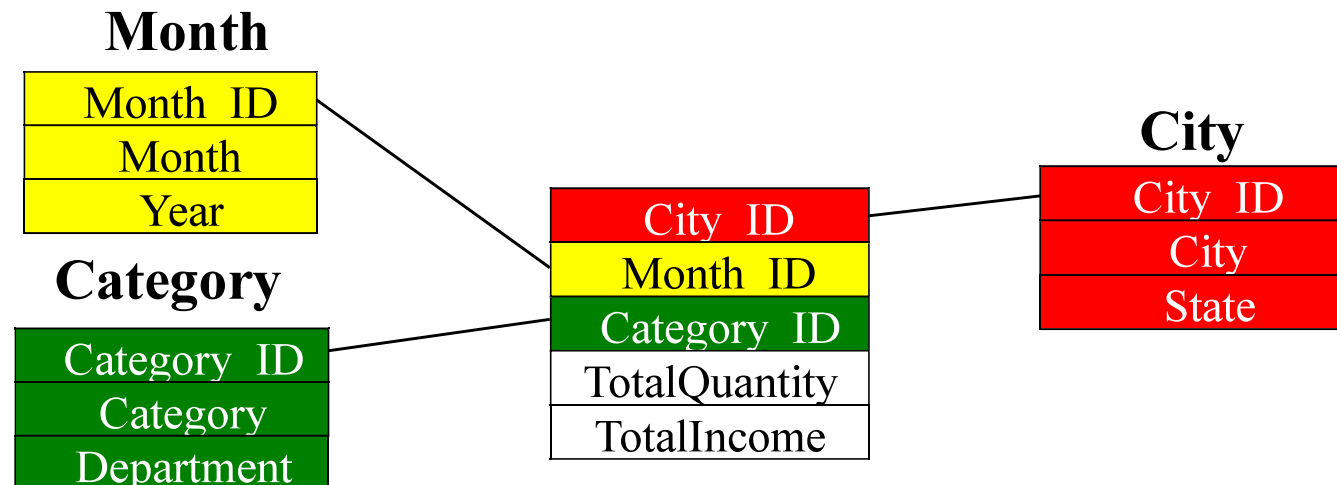
# Materialized views

- Precomputed summaries for the fact table
  - explicitly stored in the data warehouse
  - provide a performance increase for aggregate queries






# Materialized views

- Defined by SQL statements
- Example: definition of  $v_3$ 
  - Starting from base tables or views with higher granularity
    - group by City, Category, Month
  - Aggregation (SUM) on Quantity, Income measures
  - Reduction of detail in dimensions



# Materialized views

- Materialized views may be exploited for answering several different queries
  - not for all aggregation operators




**Multidimensional lattice**



# Materialized view selection

- Huge number of allowed aggregations
  - most attribute combinations are eligible
- Selection of the “best” materialized view set
- Cost function minimization
  - query execution cost
  - view maintenance (update) cost
- Constraints
  - available space
  - time window for update
  - response time
  - data freshness


# Materialized view selection




**Multidimensional lattice**

 +  = *candidate views*,  
 possibly useful to  
 increase workload  
 query performance



# Materialized view selection



Space and time minimization




# Materialized view selection





# Materialized view selection



All constraints

