Grid-based Data Stream Processing in e-Science

Richard Kuntschke[1], Tobias Scholl[1], Sebastian Huber[1],
Alfons Kemper[1], Angelika Reiser[1],
Hans-Martin Adorf[2], Gerard Lemson[3], and Wolfgang Voges[3]

[1]Lehrstuhl Informatik III:
Datenbanksysteme
Fakultät für Informatik
Technische Universität München

[2]Max-Planck-Institut
für Astrophysik

[3]Max-Planck-Institut
für extraterrestrische
Physik

Important Challenges in e-Science

- In general:
  - Large and exponentially growing amounts of data
  - Distributed data archives
  - No unique identifiers
  - Uncertainty
- In astrophysics:

  Spectral Energy Distributions (SEDs)
  - Used to classify celestial objects (active galactic nuclei, brown dwarfs, neutron stars, ...)
  - Generation requires spatial (astrometric) matching

Spatial (Astrometric) Matching

Current solutions …

- … load all data into main memory
  - Uses a lot of memory
  - Infeasible if memory size is insufficient
- … process all data at once and deliver the complete result at the end
  - Inefficient
  - No results until all processing has completed

Our Contributions

- StarGlobe
  - Grid-based P2P Data Stream Management System implemented on top of Globus
  - In-network processing
    - Early filtering
    - Parallelization
    - Pipelining
    - Load-balancing
  - Mobile user-defined operators
- Astrophysical Example Workflow
  - Astrometric matching
  - Performance evaluation

The StarGlobe Architecture

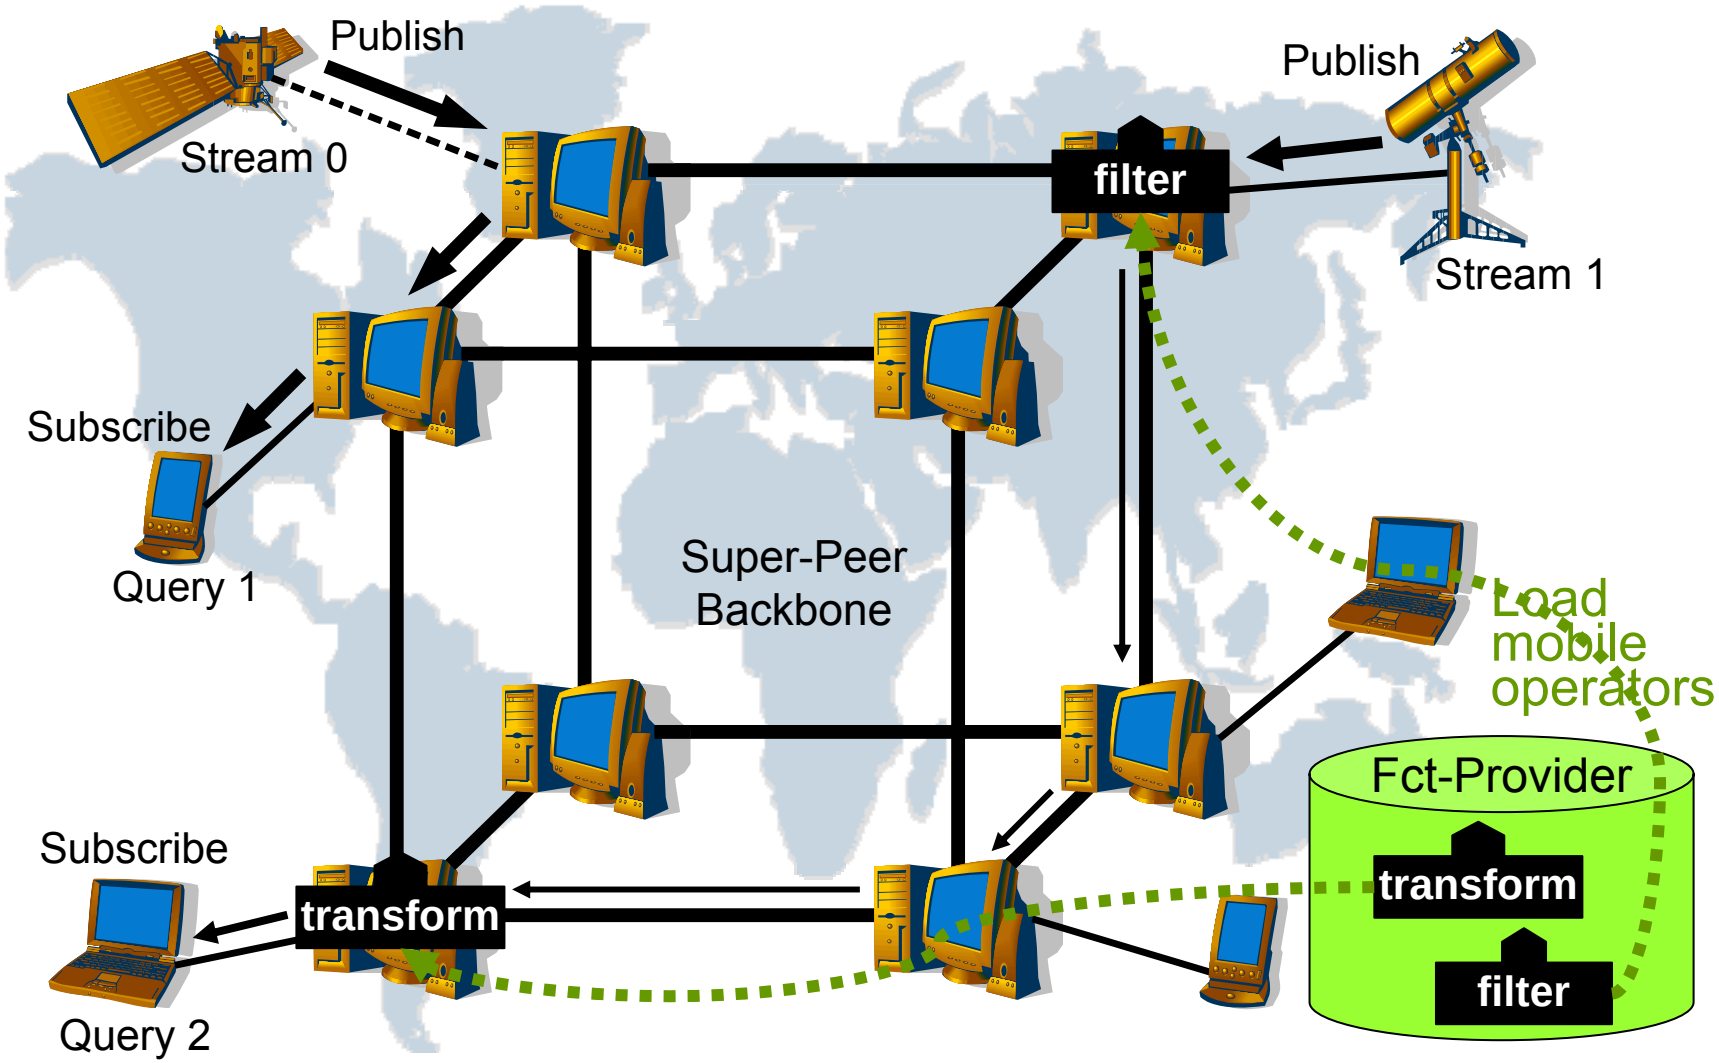

Traditional Approach: Bring Data to Code

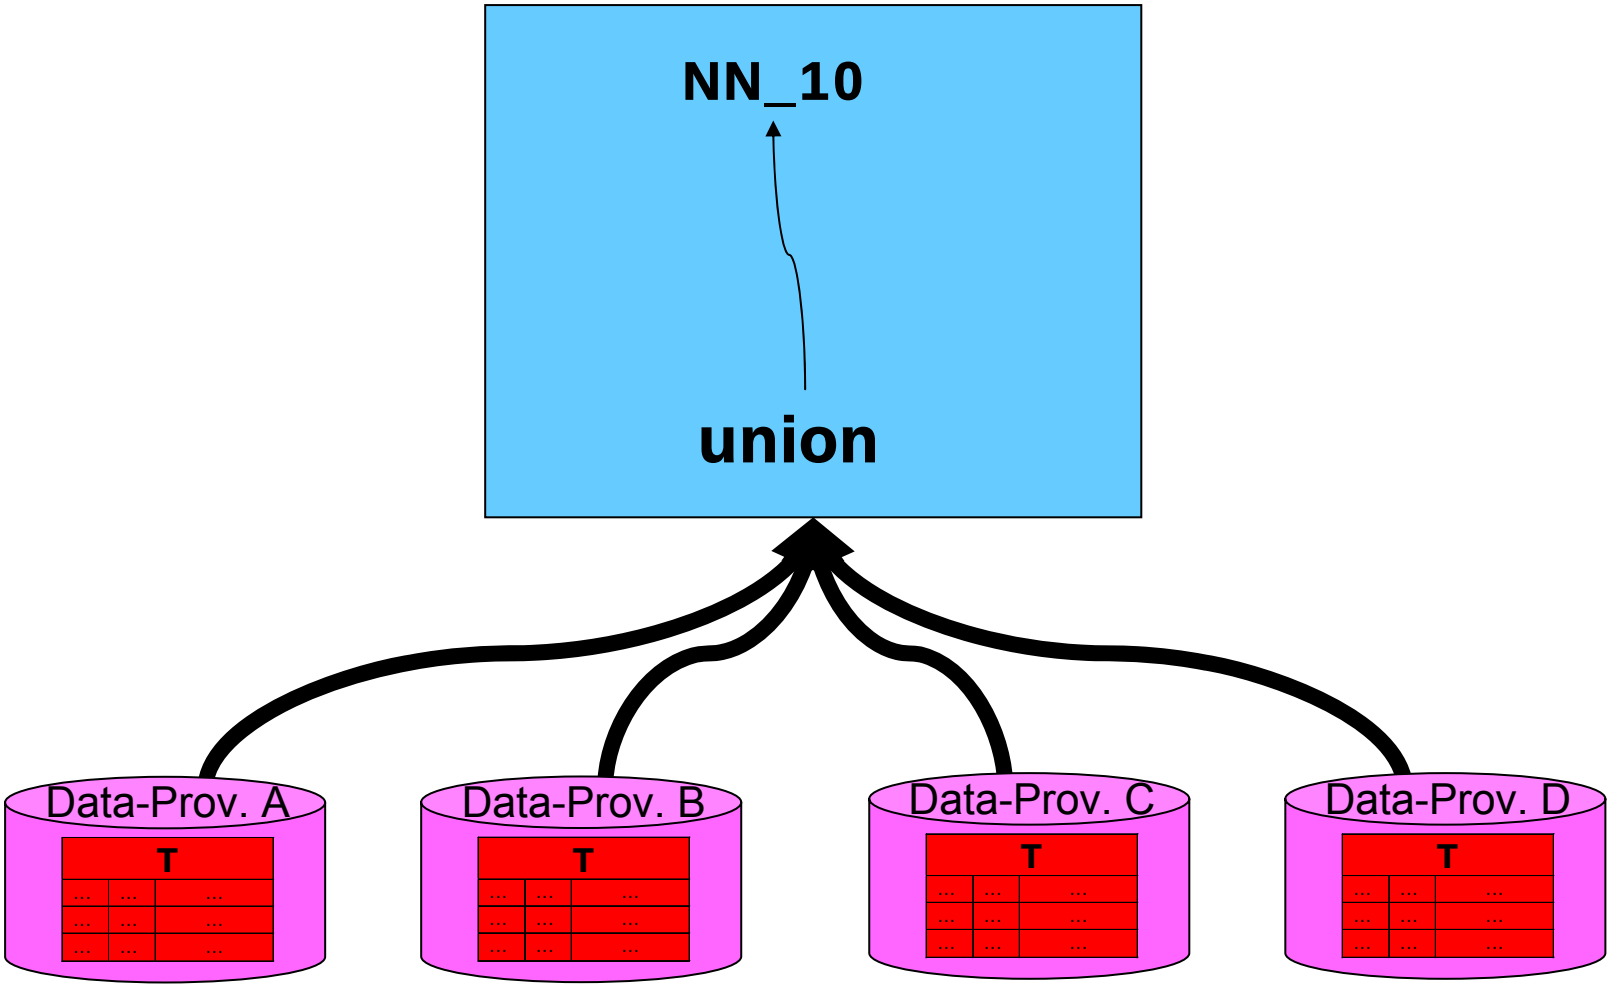

New Approach: Bring Code to Data

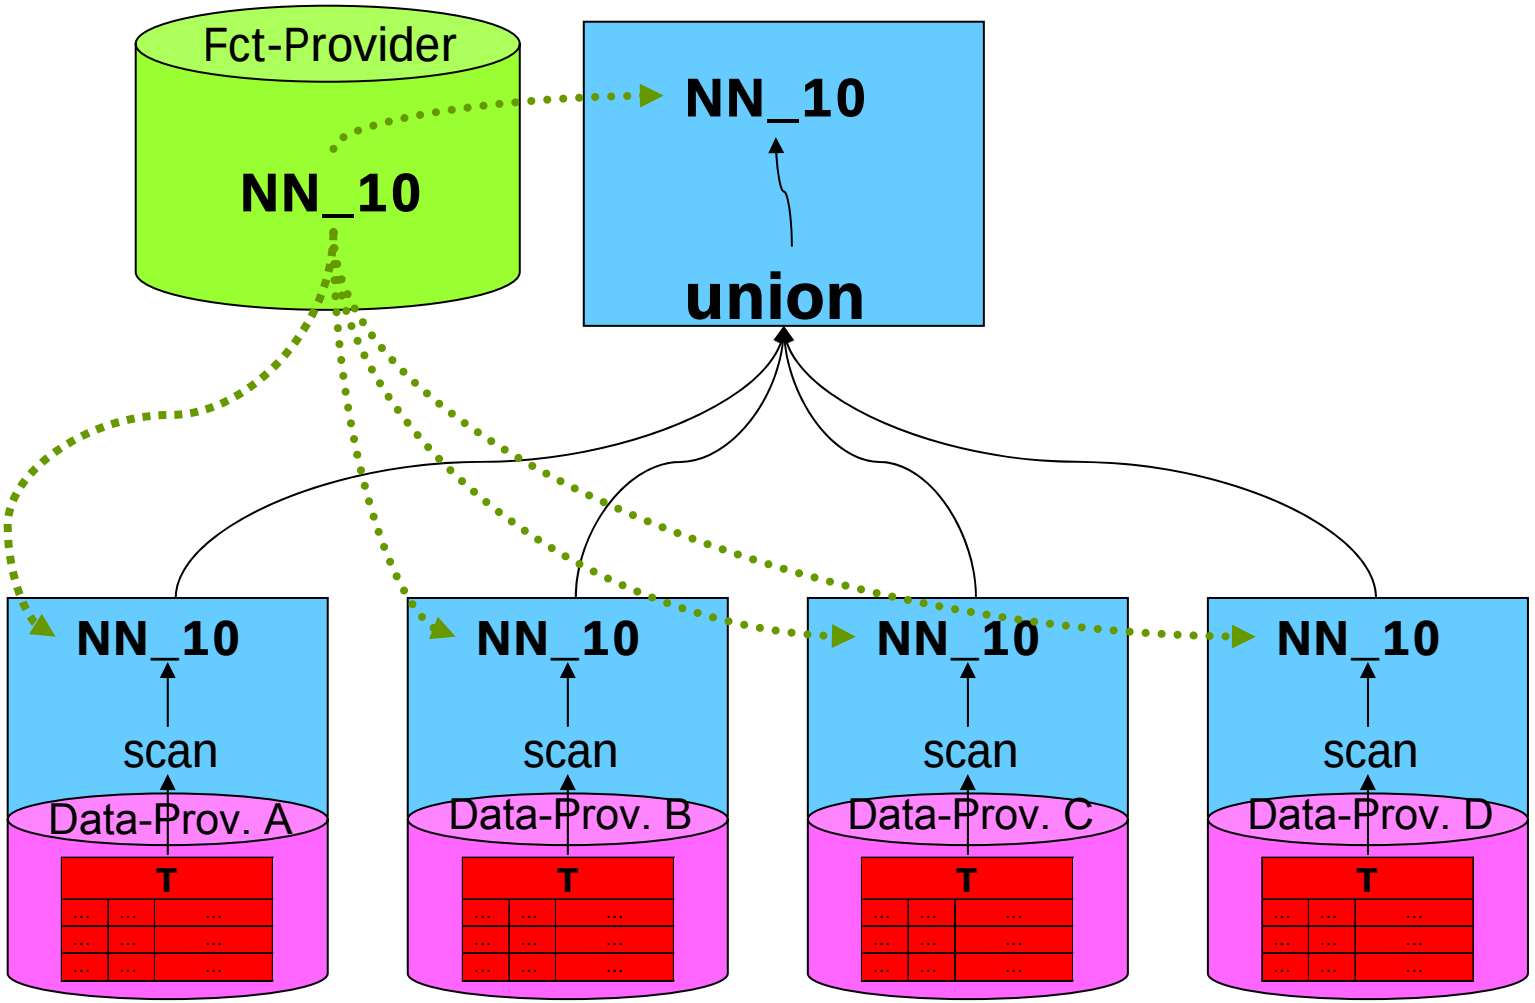

Mobile User-Defined Operators

- Load user-defined operators from function provider servers in the network
- Common interface for integrating external operators
- Push-based iterator
- Flexibility

StreamIterator Interface

- open(Config, StreamWriter)
  - Configuration parameters
  - Writer for result stream
- next(StreamIteratorEvent)
  - Next element in input stream
  - Writing output to result stream using StreamWriter.write()
- close()

Communication between StreamProcessor and StreamIterator

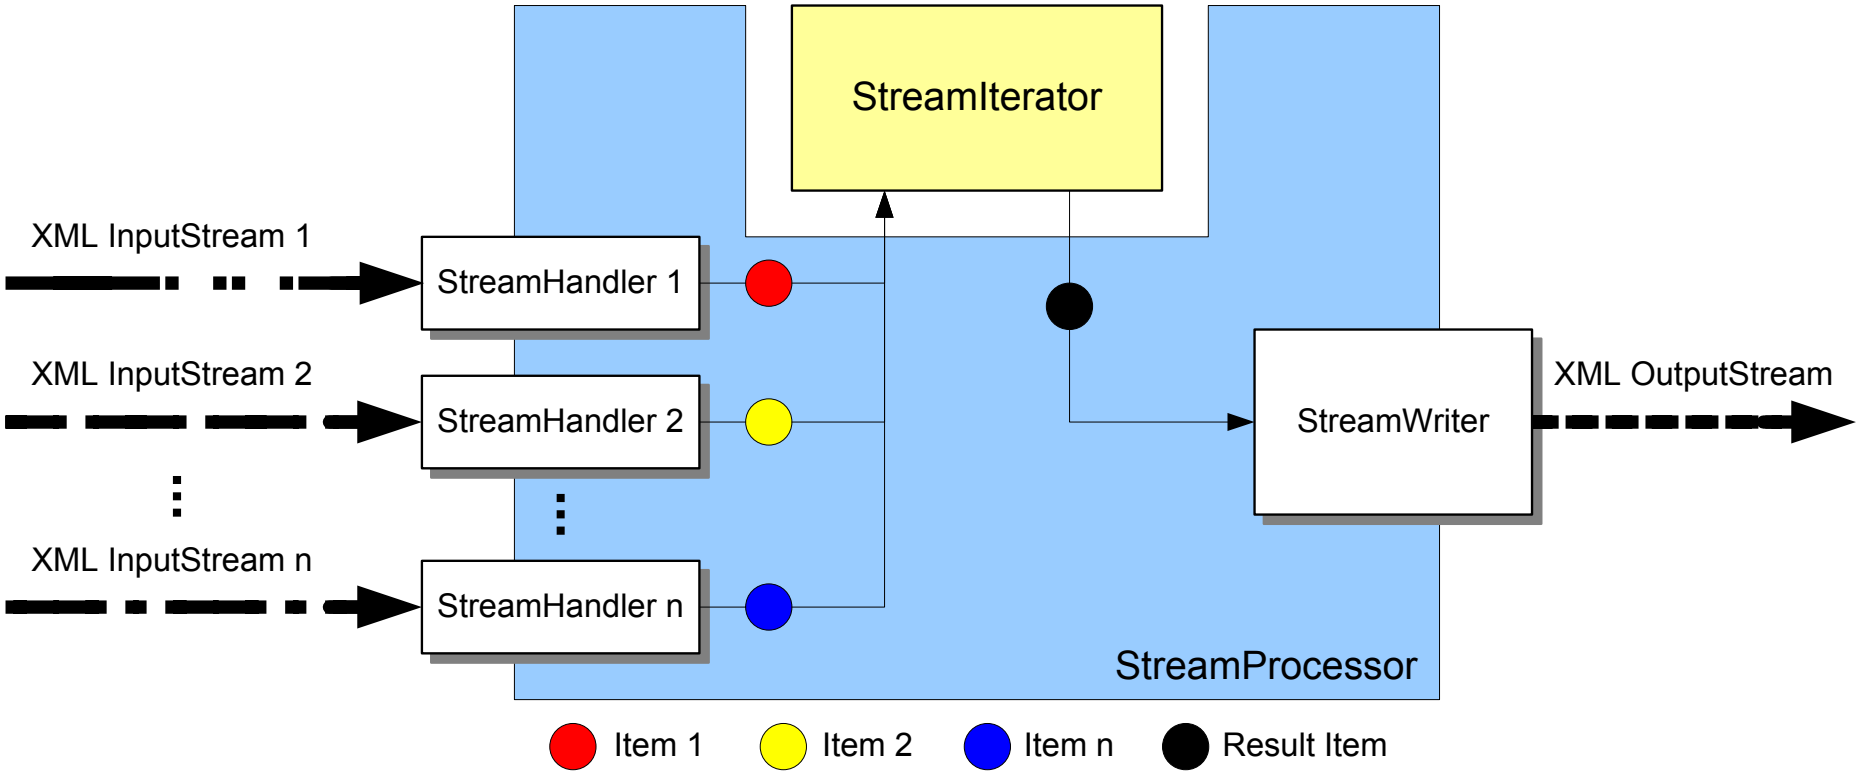

Astrophysical Example Workflow

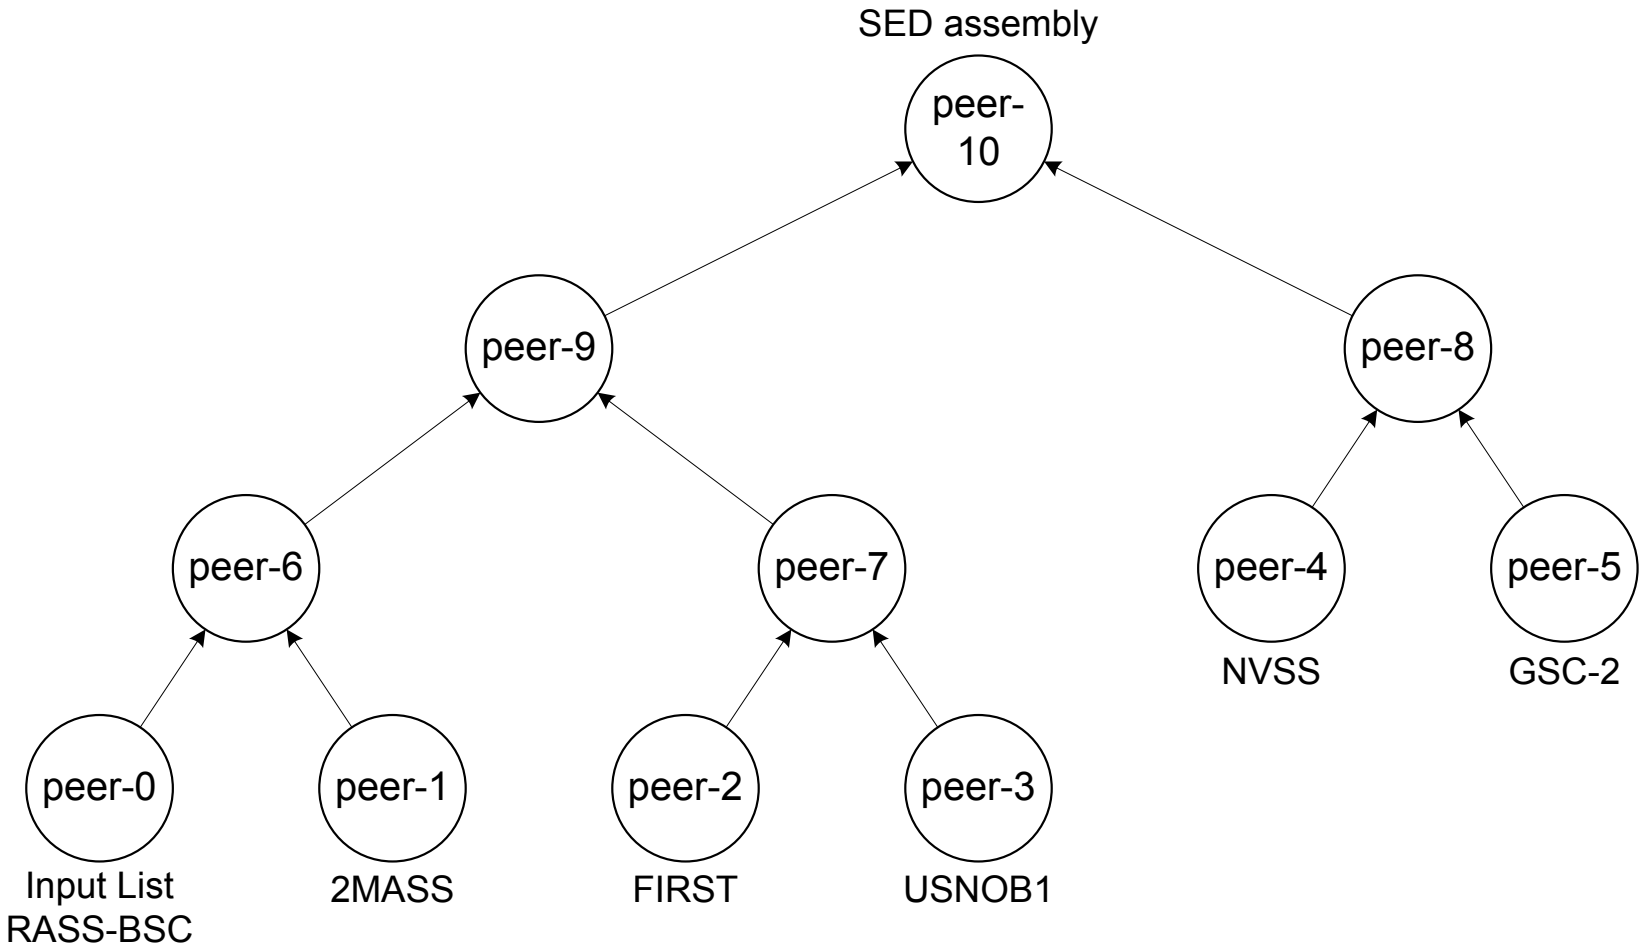

Distributed Query Evaluation Plan

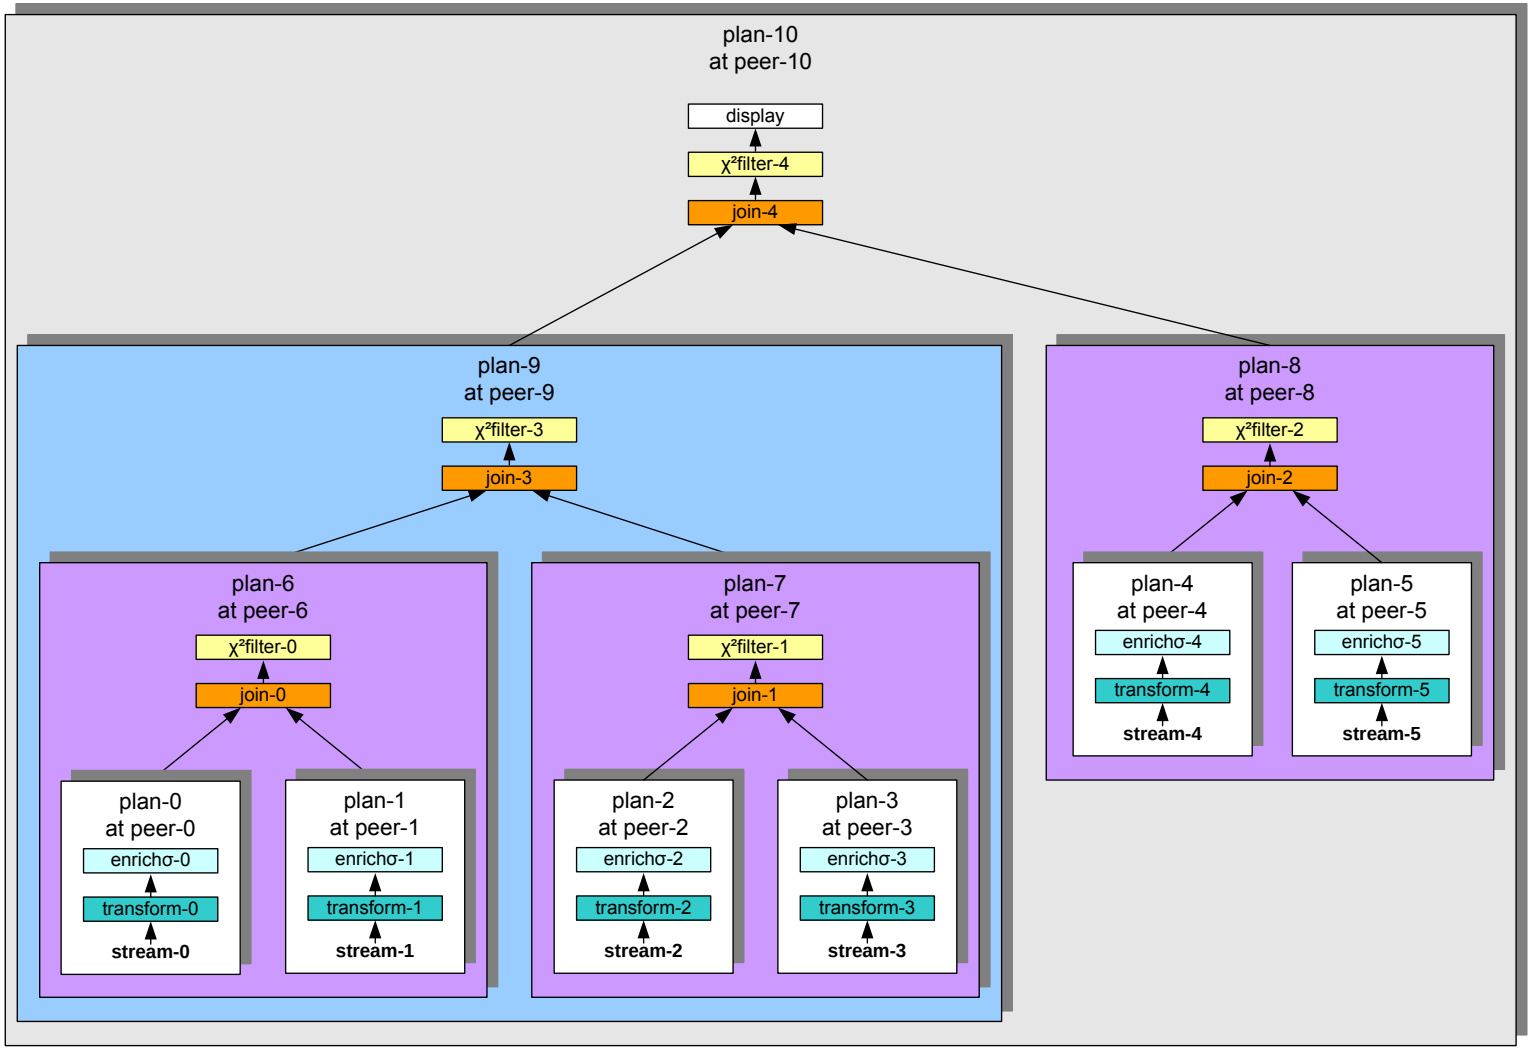

Distributed Query Evaluation Plan

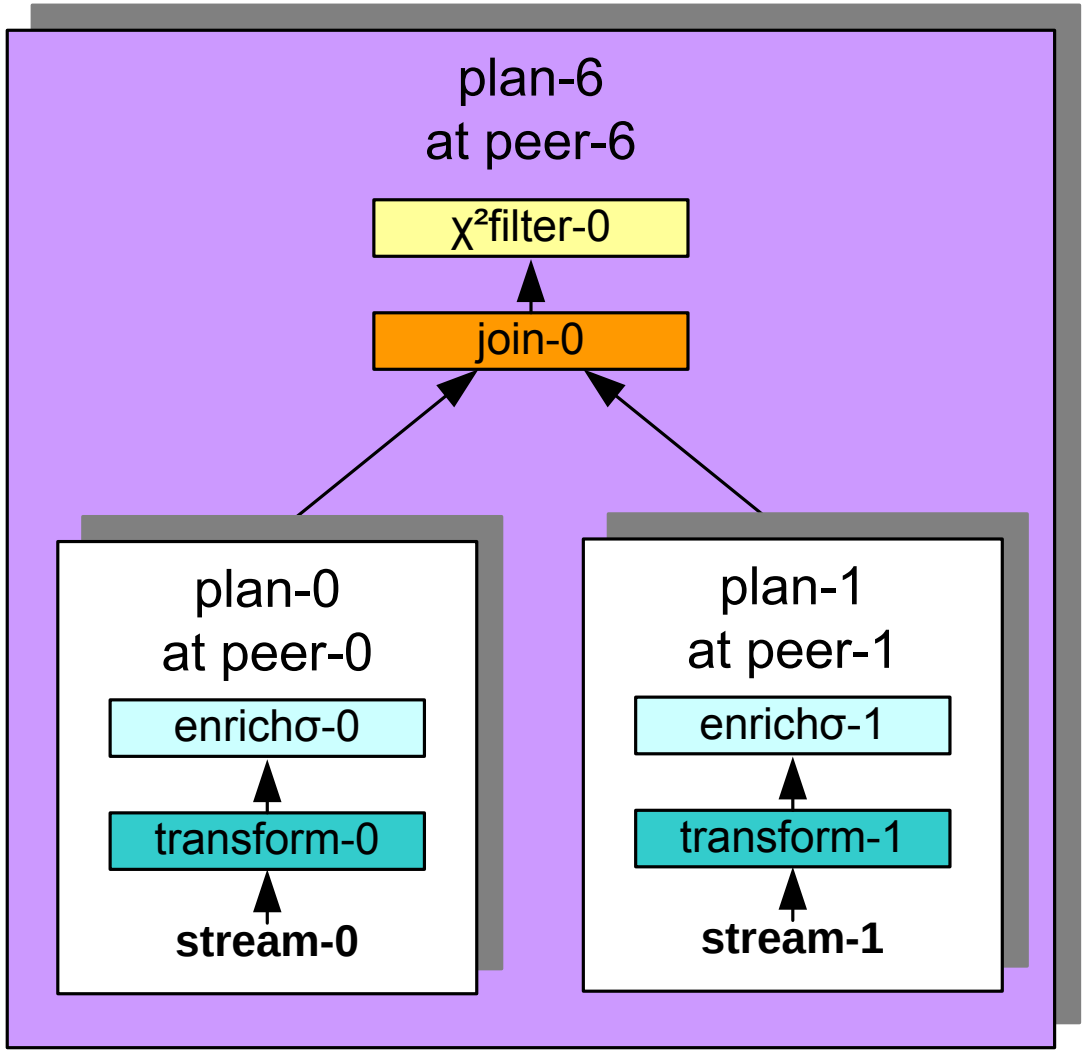

Distributed Query Evaluation Plan

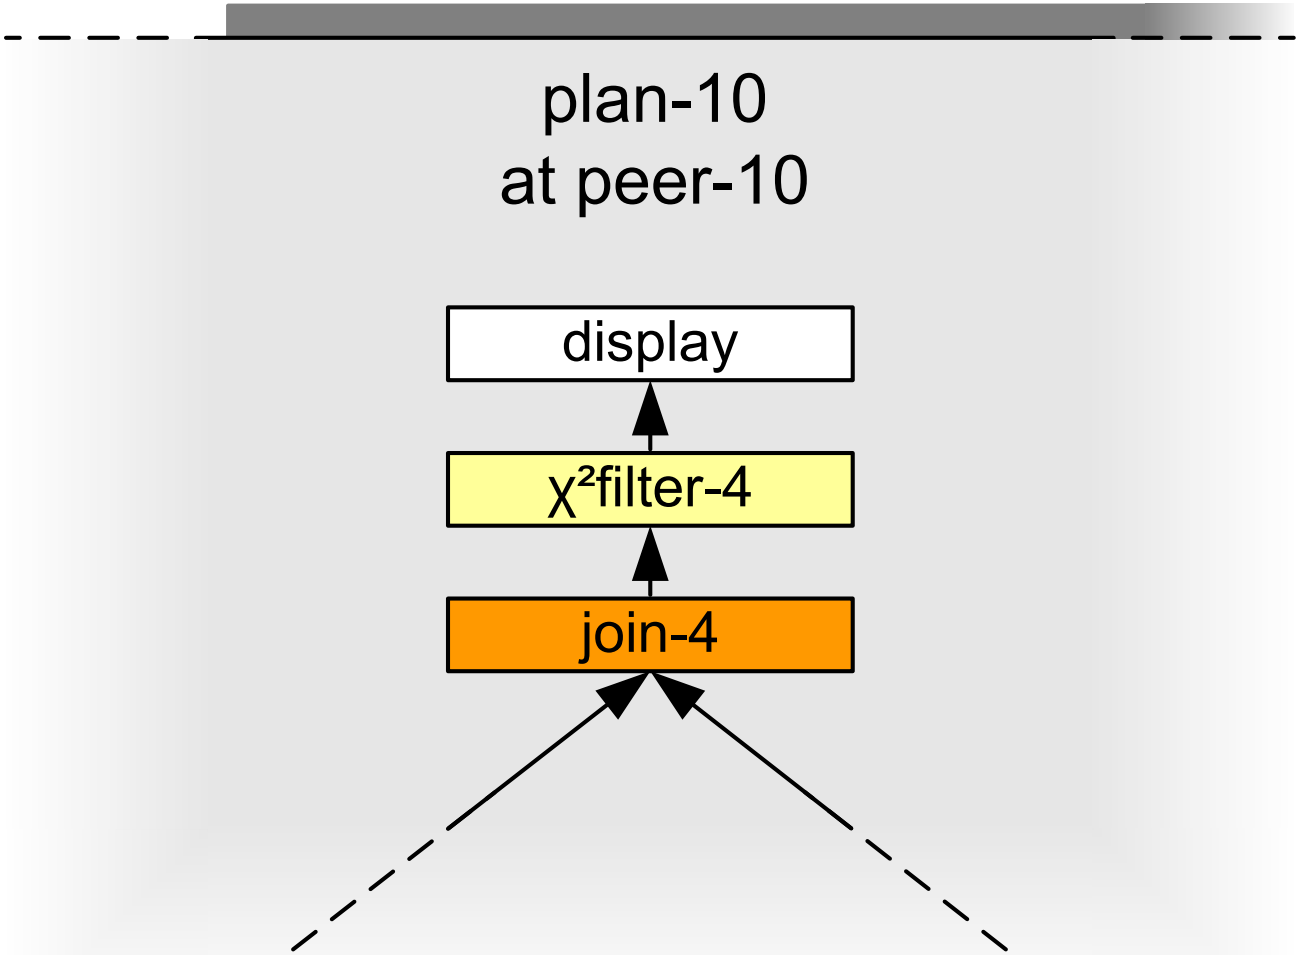

Evaluation of Early Filtering

	WITH EARLY FILTERING		WITHOUT EARLY FILTERING	
	Stream size	# Match candidates	Stream size	# Match candidates
After join-0	808 KB	611	808 KB	611
After join-1	1,874 KB	1,138	1,874 KB	1,138
After join-2	1,387 KB	826	1,387 KB	826
After join-3	6,355 KB	2,522	46,525 KB	15,489
After join-4	14,356 KB	3,815	1,838,648 KB	364,299
After filter-4	1,364 KB	318	1,364 KB	318
Duration h:m:s	00:02:58		02:46:00	

	# MATCH CANDIDATES		FILTER RATIO
	Before filtering	After filtering	
At join-0	611	289	47.3%
At join-1	1,138	452	39.7%
At join-2	826	458	55.4%
At join-3	2,522	400	15.9%
At join-4	3,815	318	8.3%

Conclusion

- Synergies between research in computer science and other scientific disciplines, e.g., astrophysics
- StarGlobe
  - Handling large data volumes efficiently
    - Early filtering, parallelization, pipelining
  - Returning first results early on
    - Pipelining
  - Flexible support of domain-specific application logic
    - Mobile user-defined operators
- Results also applicable to other domains