




Raw Array Data Model

j \ i	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]
[1]	<X,1>		<Y,1>	<Z,1>		<Z,3>	<Y,2>	
[2]		<Y,1>	<X,1>		<Z,2>	<X,2>	<Y,2>	
[3]	<X,1>		<Z,1>		<X,2>			<Z,3>
[4]		<Z,1>						
[5]					<Y,2>	<Z,2>		
[6]				<Z,2>				<X,2>

Raw Array Chunking



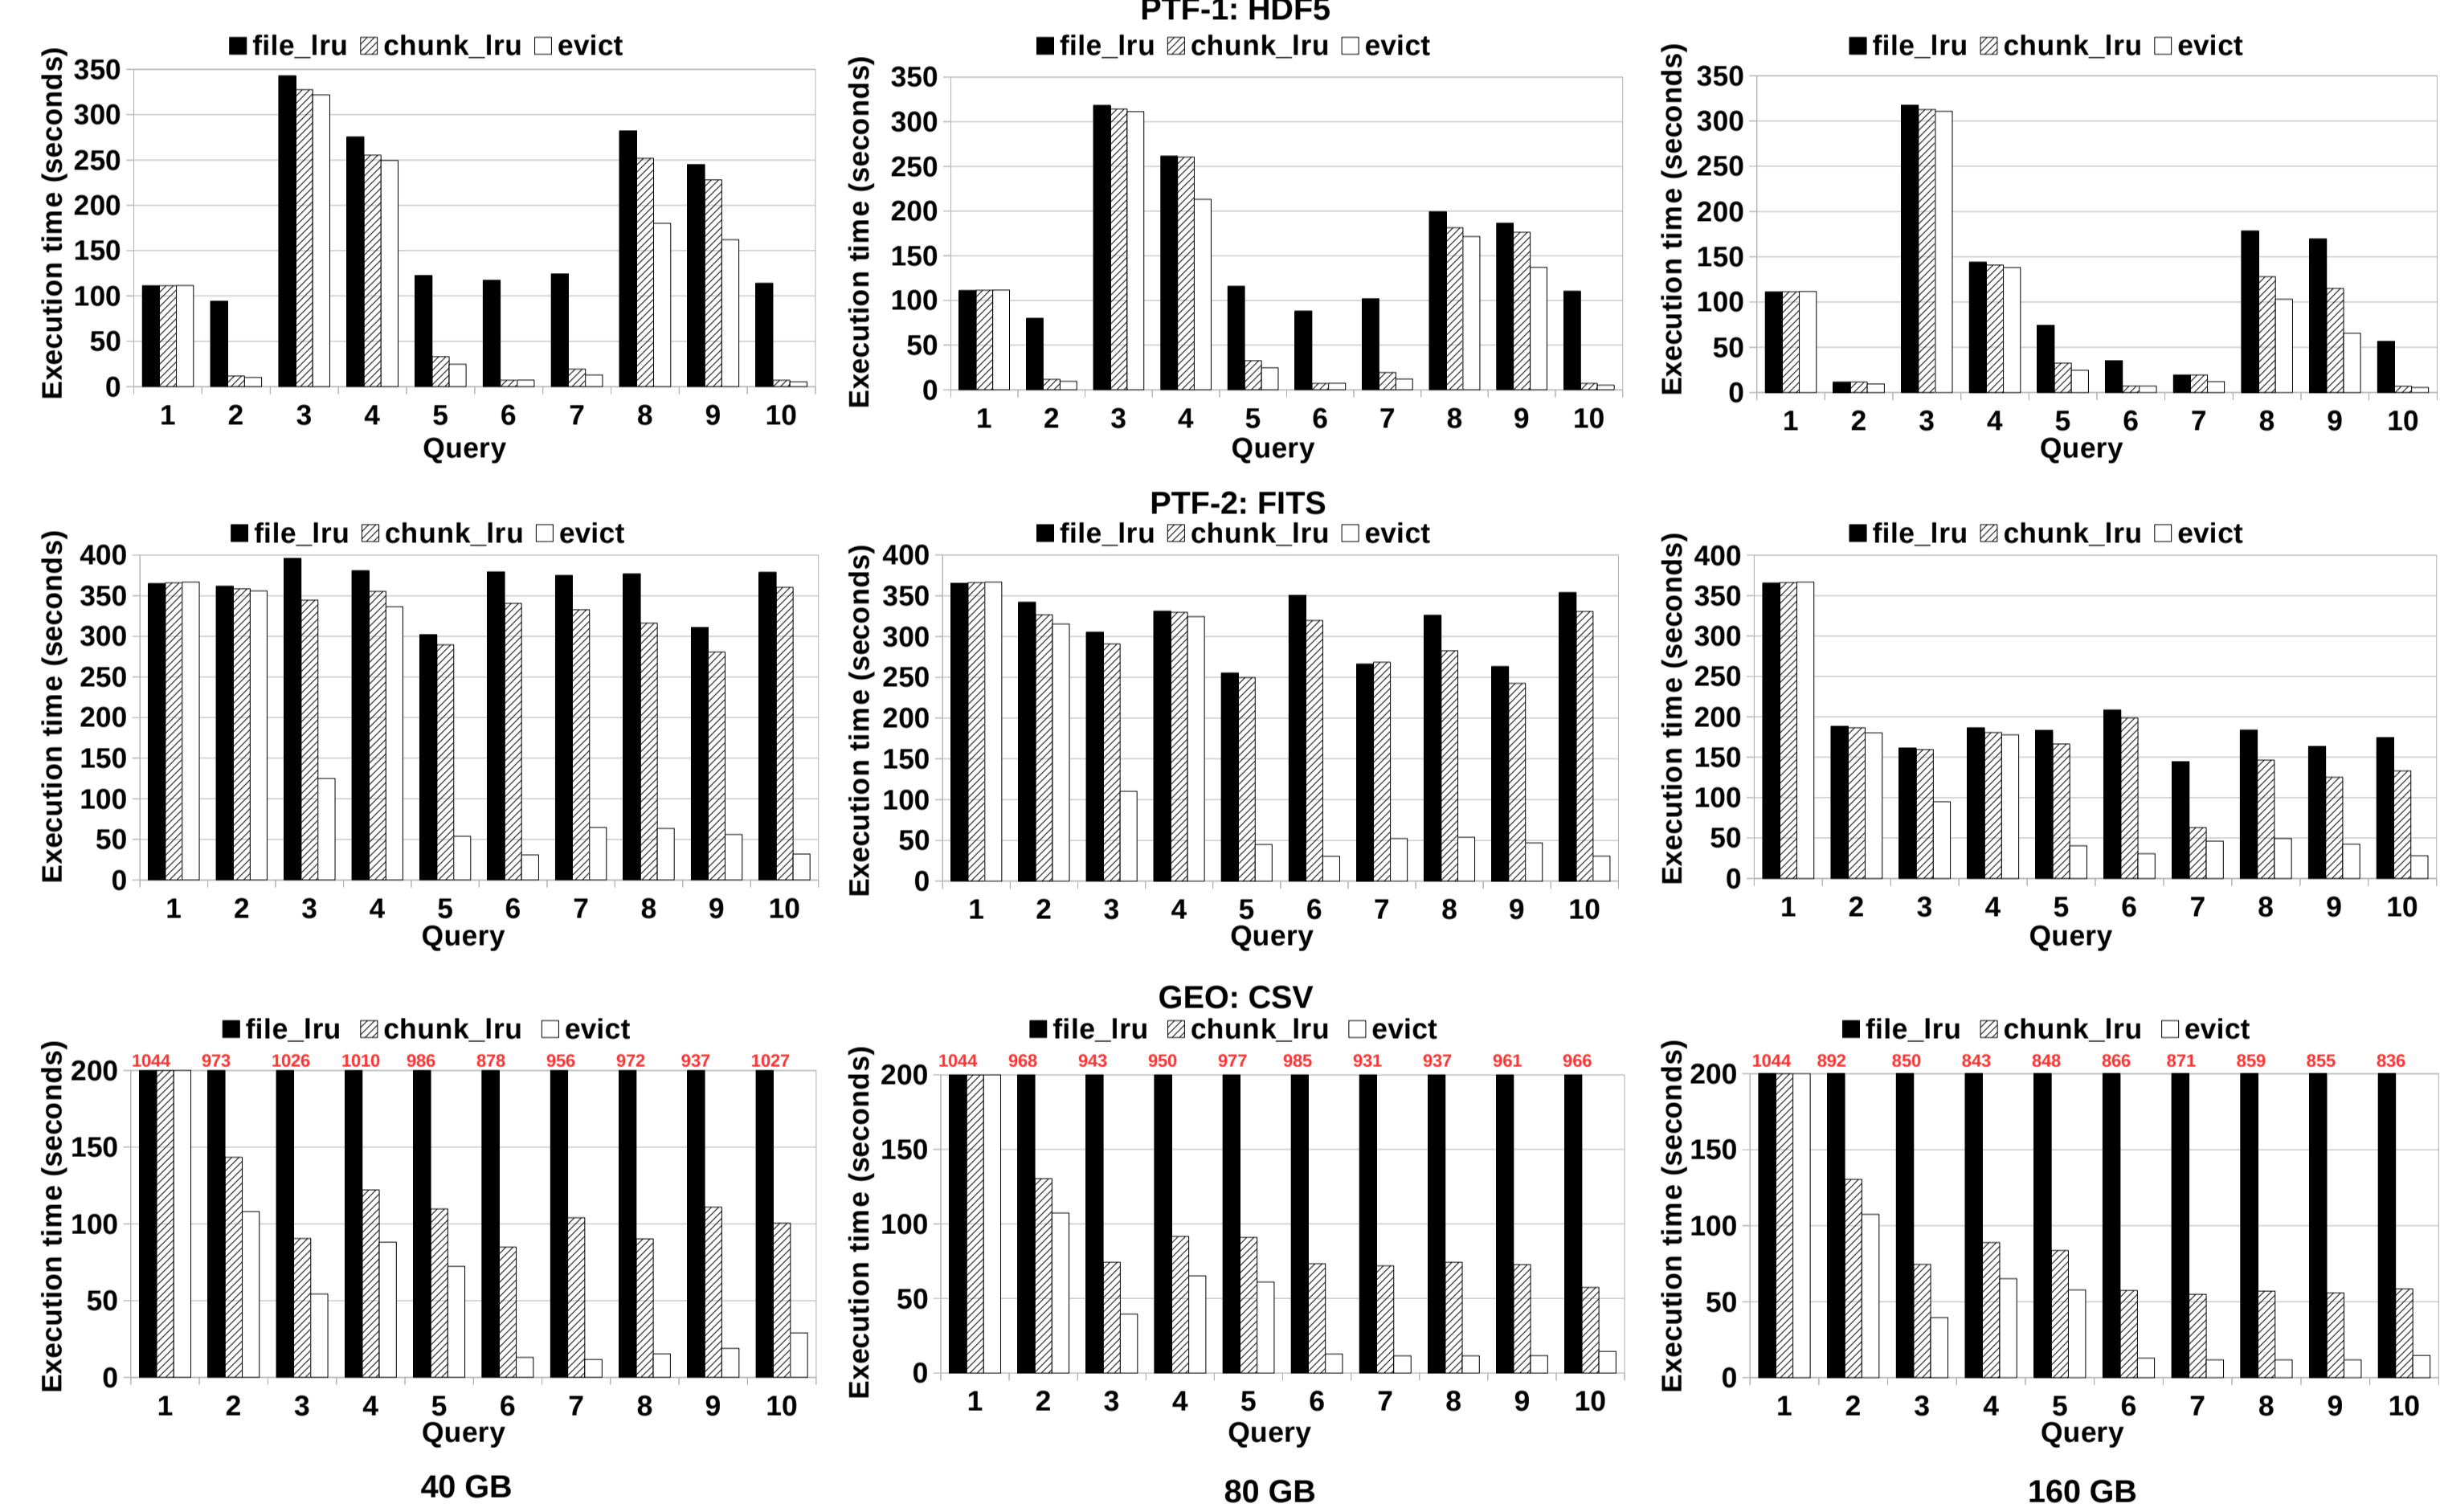
Experiments

PTF catalog: PTF [time=1, 153064; ra=1, 100000; dec=1, 50000]
1 billion objects, 343 GB in CSV, 262 GB in HDF5, and 221 GB in FITS


LinkedGeoData: GEO [long=1, 100000; lat=1, 50000]
30 billion objects, 1.6 TB in CSV.

3 query patterns: real workload, shifting ranges and alternative queries.


Query execution time:



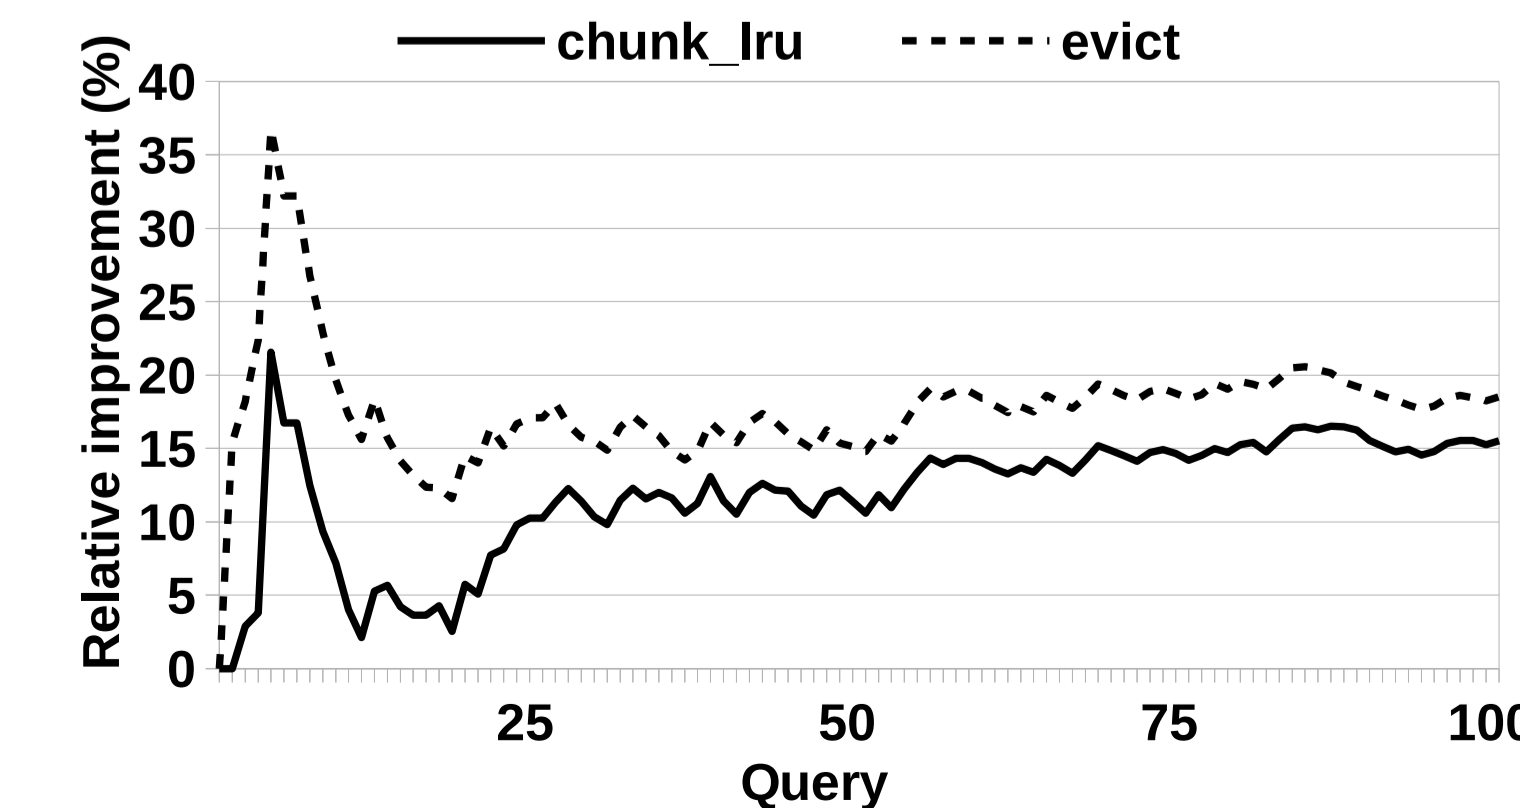
Similarity join execution time:



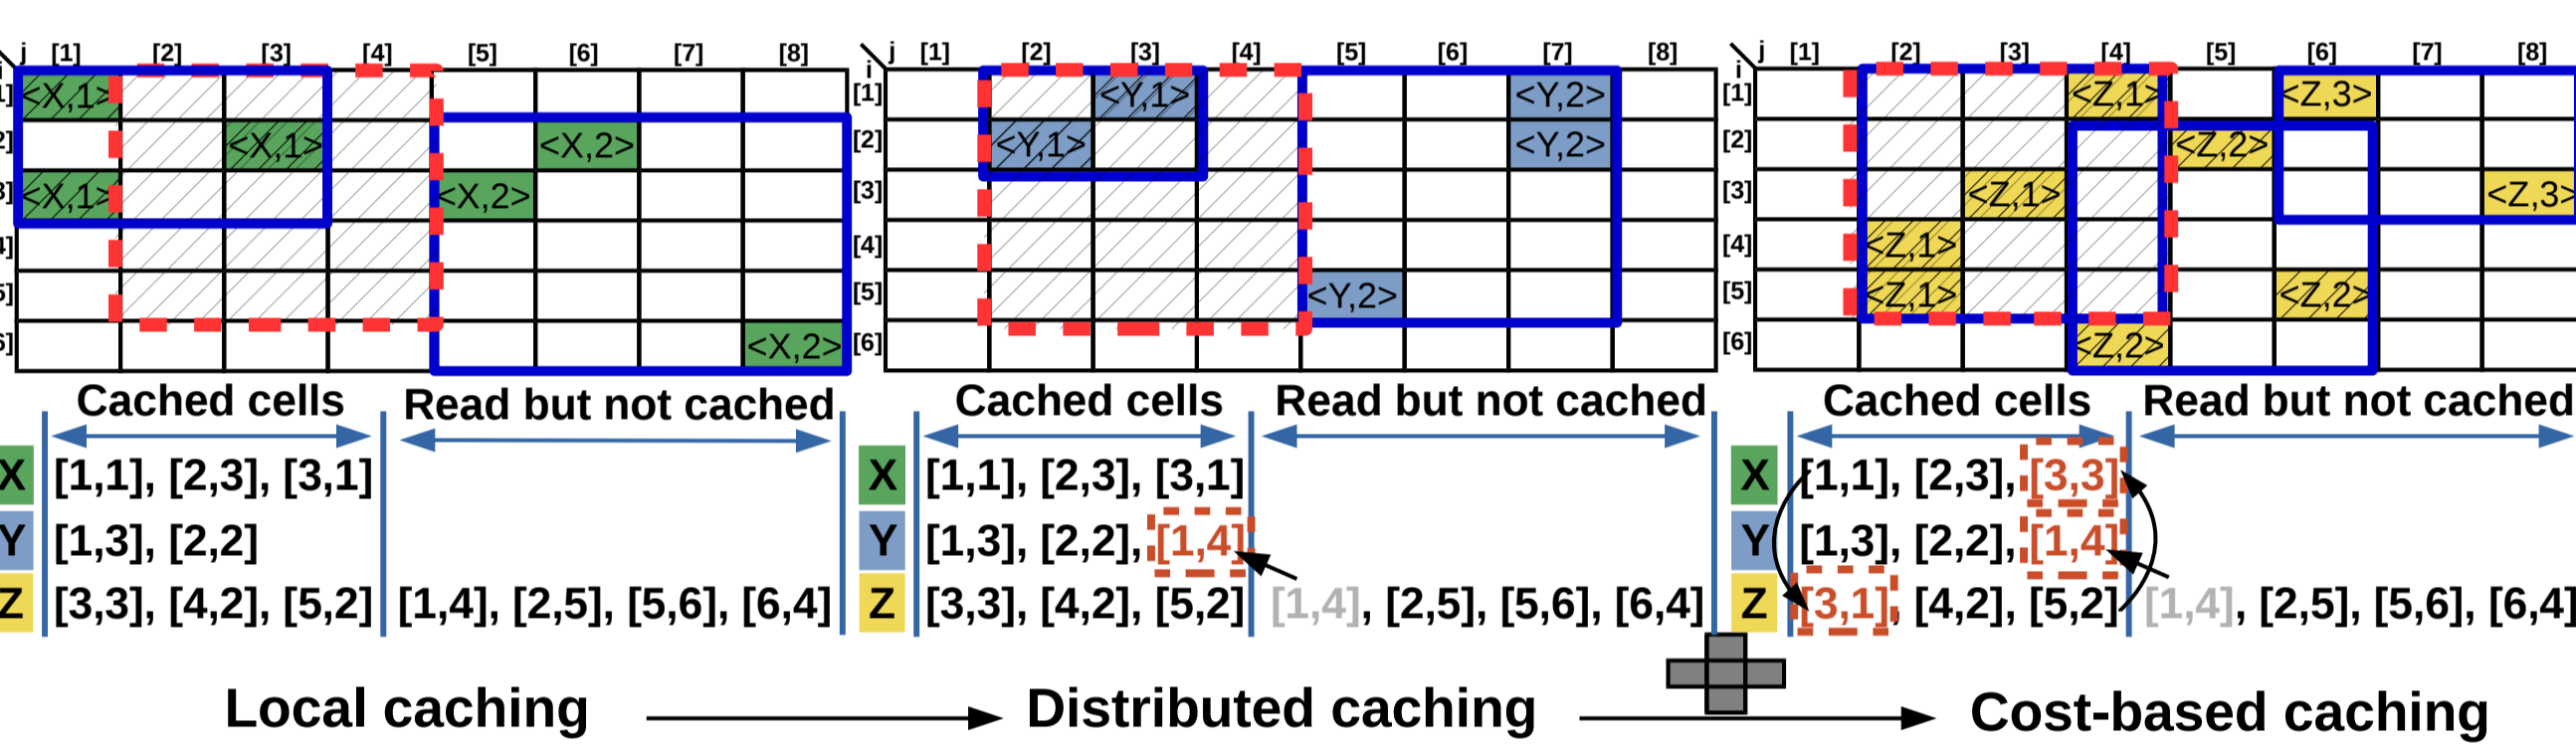
Optimization time



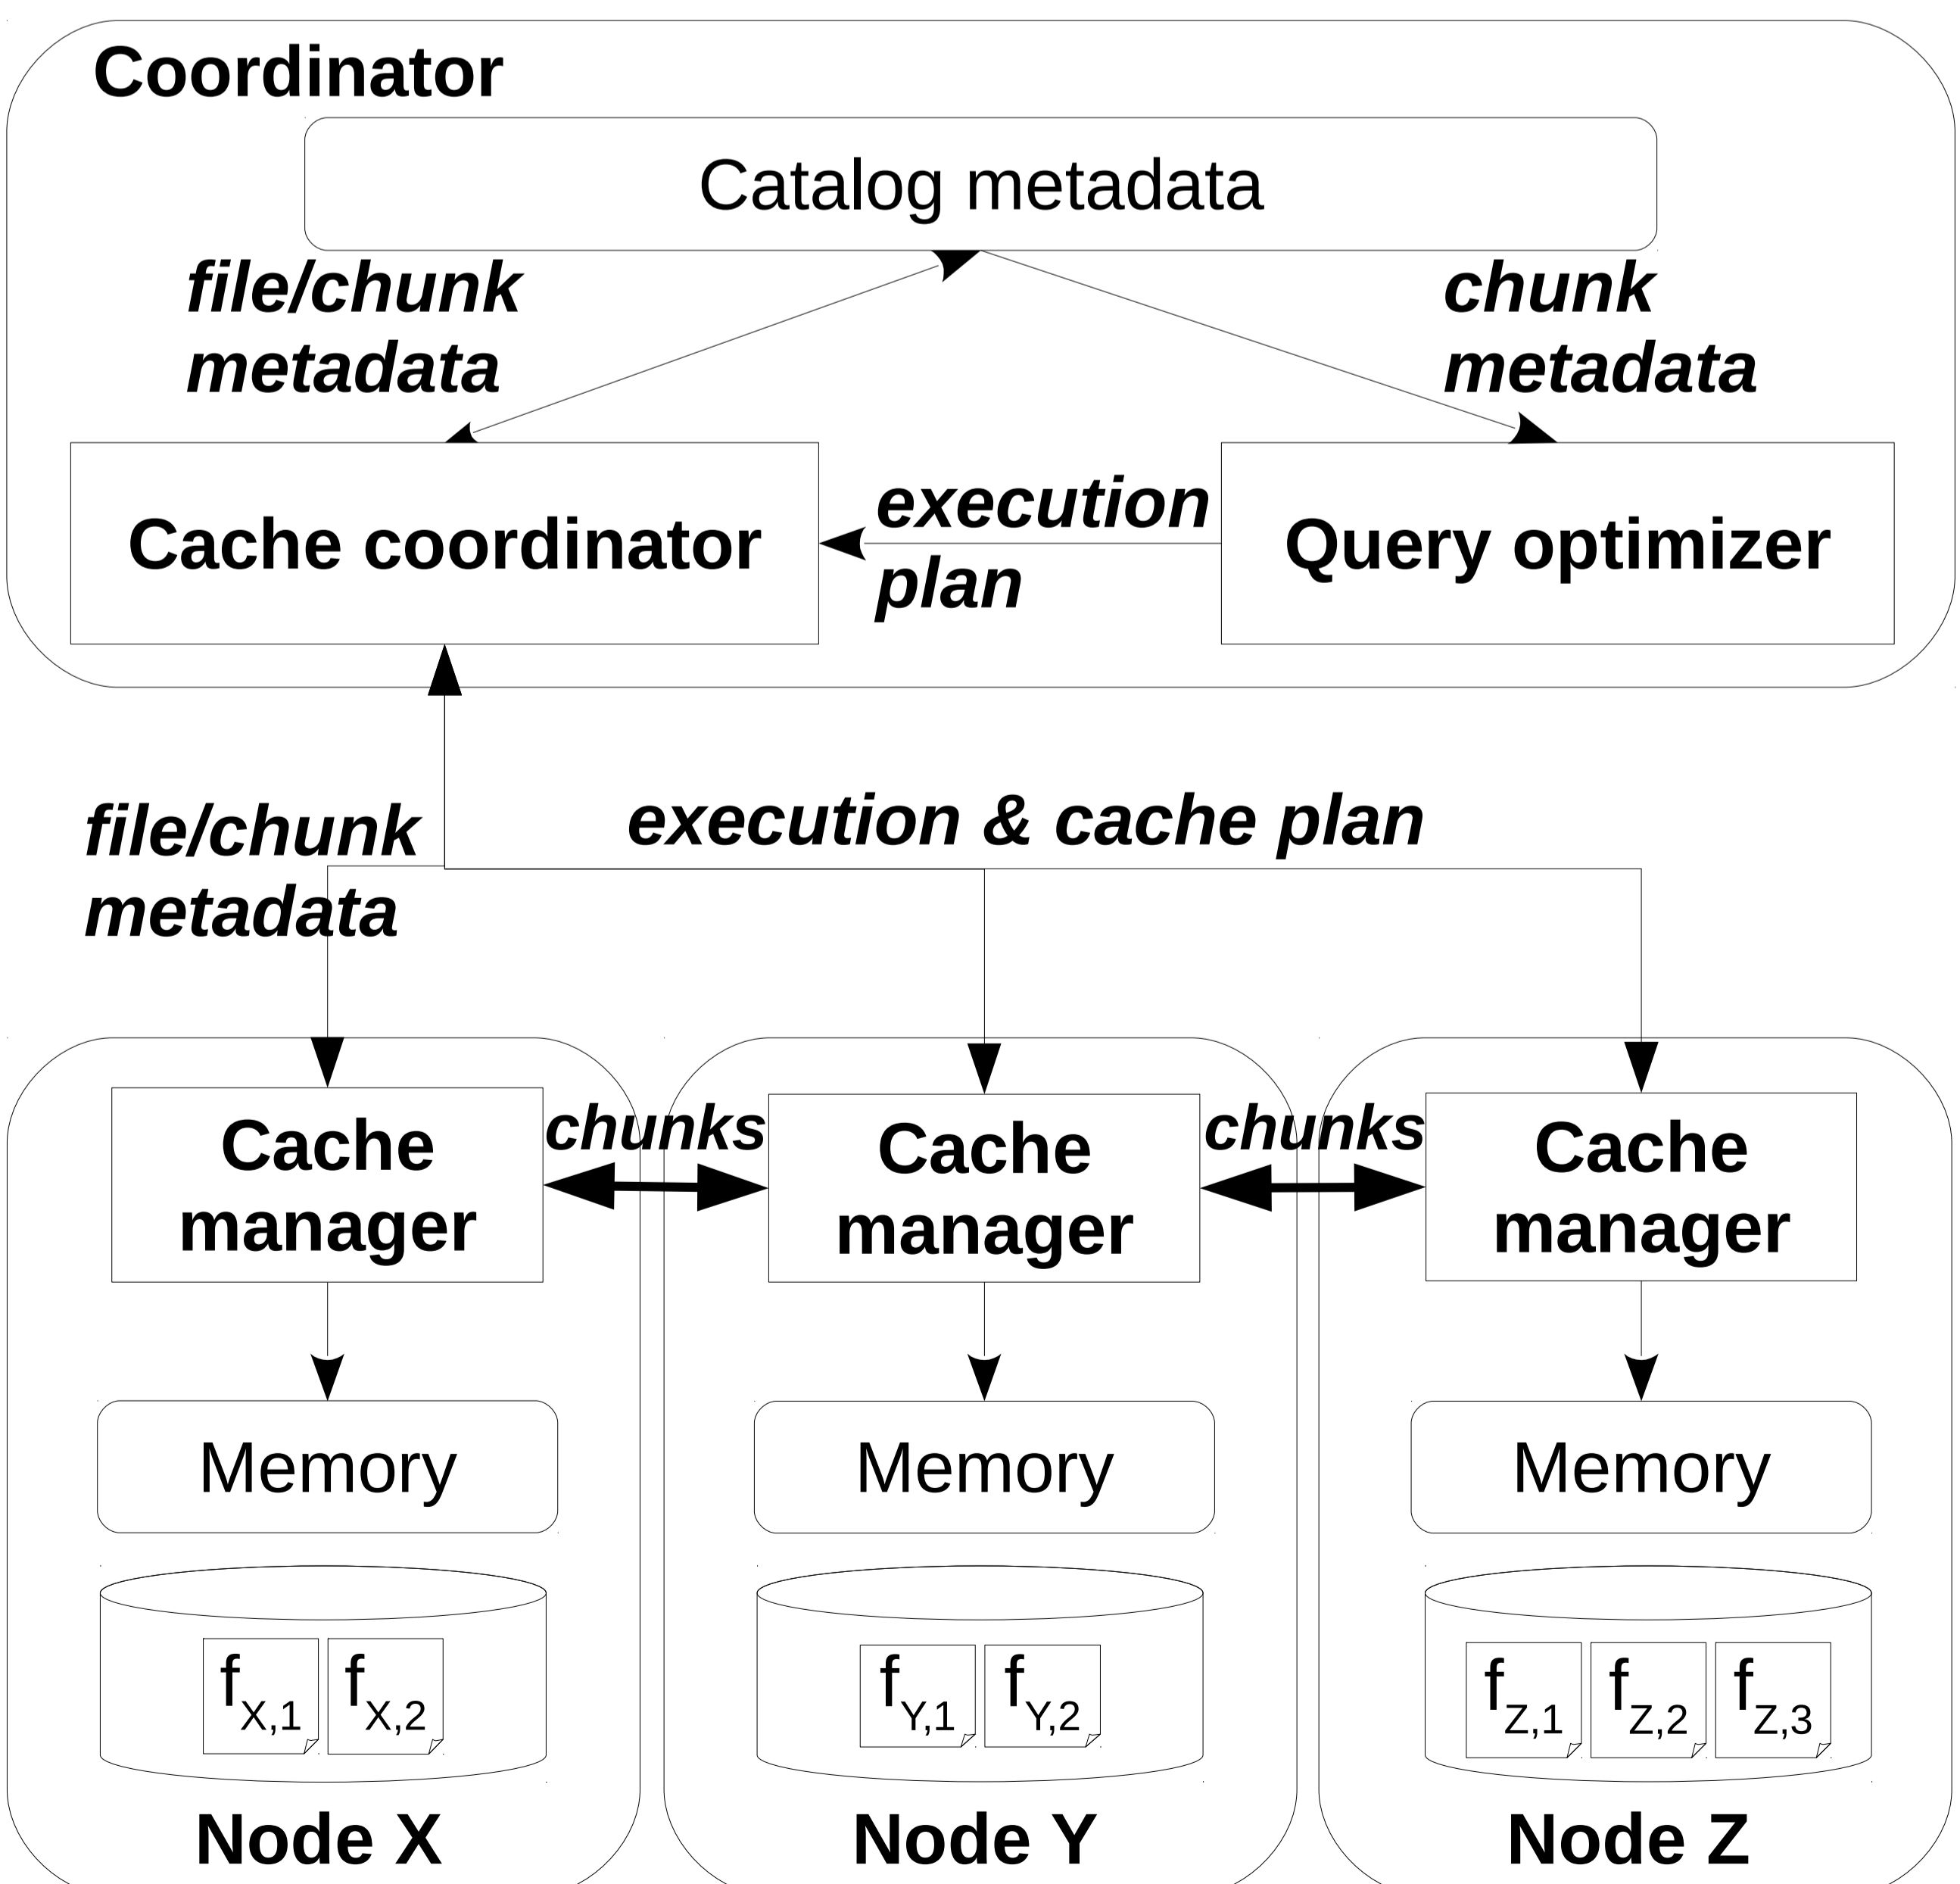
evict v.s chunk lru improvement



Raw Array Distributed Caching



Distributed Caching Architecture



When to split?

- Chunk is sufficiently large.
- Query subarray does not contain any cells.

How to split?

- According to the query subarray boundaries?

Algorithm 1 Chunk Split

Input: Chunk α with bounding box $B\alpha$ that intersects query subarray Q ;
Minimum number of cells threshold $MinC$

Output: Chunks β and γ after splitting α

- 1: **if**(cells in $\alpha < MinC$) **and** (\exists cell in $\alpha \in Q$) **then return**
- 2: $min_vol = +\infty$
- 3: **for each** boundary $b \in Q$ that intersects with $B\alpha$ **do**
- 4: $(\beta_b, \gamma_b) \leftarrow$ split cells in α into two sets by boundary b
- 5: **if** $vol(\beta_b) + vol(\gamma_b) < min_vol$ **then**
- 6: $min_vol \leftarrow vol(\beta_b) + vol(\gamma_b)$
- 7: $\beta \leftarrow$ bounding_box(β_b), $\gamma \leftarrow$ bounding_box(γ_b)
- 8: **end if**
- 9: **end for**

Cost-Based Caching

Cache Eviction

We must scan a file entirely even if only one accessed chunk isn't cached.
→ We aim to cache all the queried chunks in a file.

$$cost_{evict}(Q_i, f_i, \{C_j\}) = w_{Q_i} \cdot \frac{size(f_i)}{\sum size(uncached C_j)}$$

Cache Placement

We piggyback on the replication induced by query execution.

$$cost_{placement}(C_i, n, P', W) = \sum_{Q \in W} w_Q \cdot |C_j \in P'_n \wedge (C_i, C_j) \in Q|$$

C_i : current chunk to place, n : candidate node – typically chosen from where a replica of C_i exists

P' : location of already placed chunks, W : query workload.