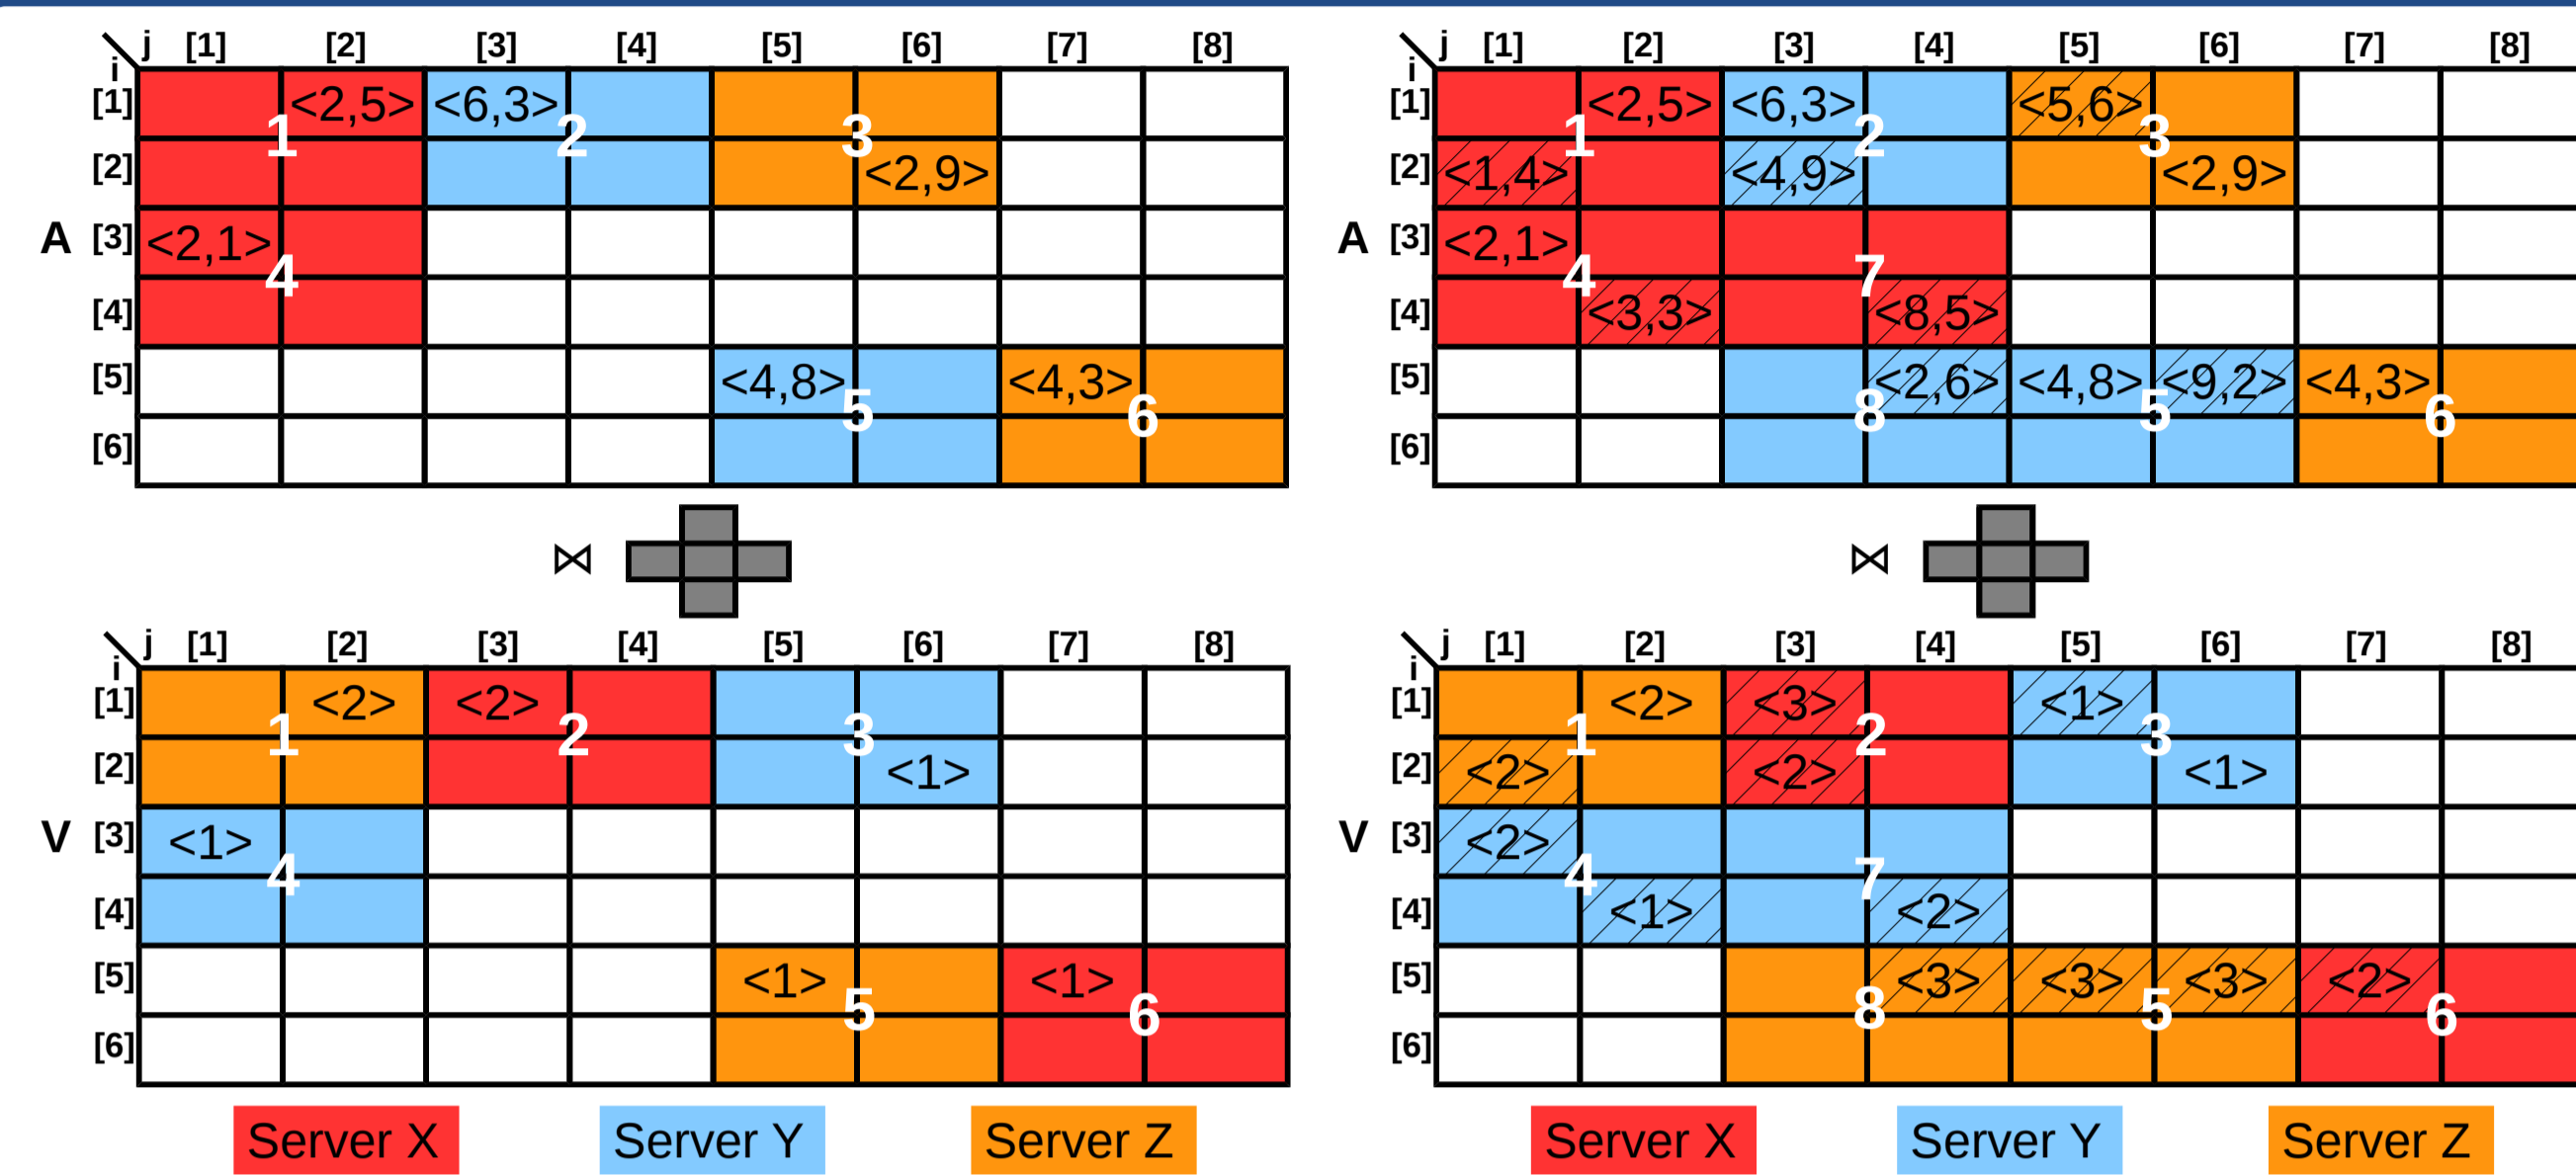




Shared-Nothing Architecture



Array View

n multi-dimensional arrays: $\alpha_1, \dots, \alpha_n$

k unary array operators $\oplus_1, \dots, \oplus_k$

We define an array view:

$$V \leftarrow \oplus_1 \left(\dots \oplus_k \left(\alpha_1 \bowtie_{\sigma_1, f_1}^{M_1} \alpha_2 \bowtie_{\sigma_2, f_2}^{M_2} \dots \bowtie_{\sigma_{n-1}, f_{n-1}}^{M_{n-1}} \alpha_n \right) \dots \right)$$

For example:

```
CREATE ARRAY VIEW V AS
SELECT COUNT(*) AS cnt FROM A A1 SIMILARITY JOIN A A2
ON (A1.i=A2.i) AND (A1.j=A2.j) WITH SHAPE L1(1)
GROUP BY A1.i, A1.j
```

View Maintenance

Batch update: a sequence of batch updates.

For each update batch ΔD :

$$M(D + \Delta D) = M(D) + \Delta Q(D, \Delta D)$$

- Assign ΔD to node according to chunking strategy
- $Q(D, \Delta D)$ – differential view computation
- $M(D) + \Delta Q(D, \Delta D)$ – view aggregation

- Excessive communication
- Load imbalance

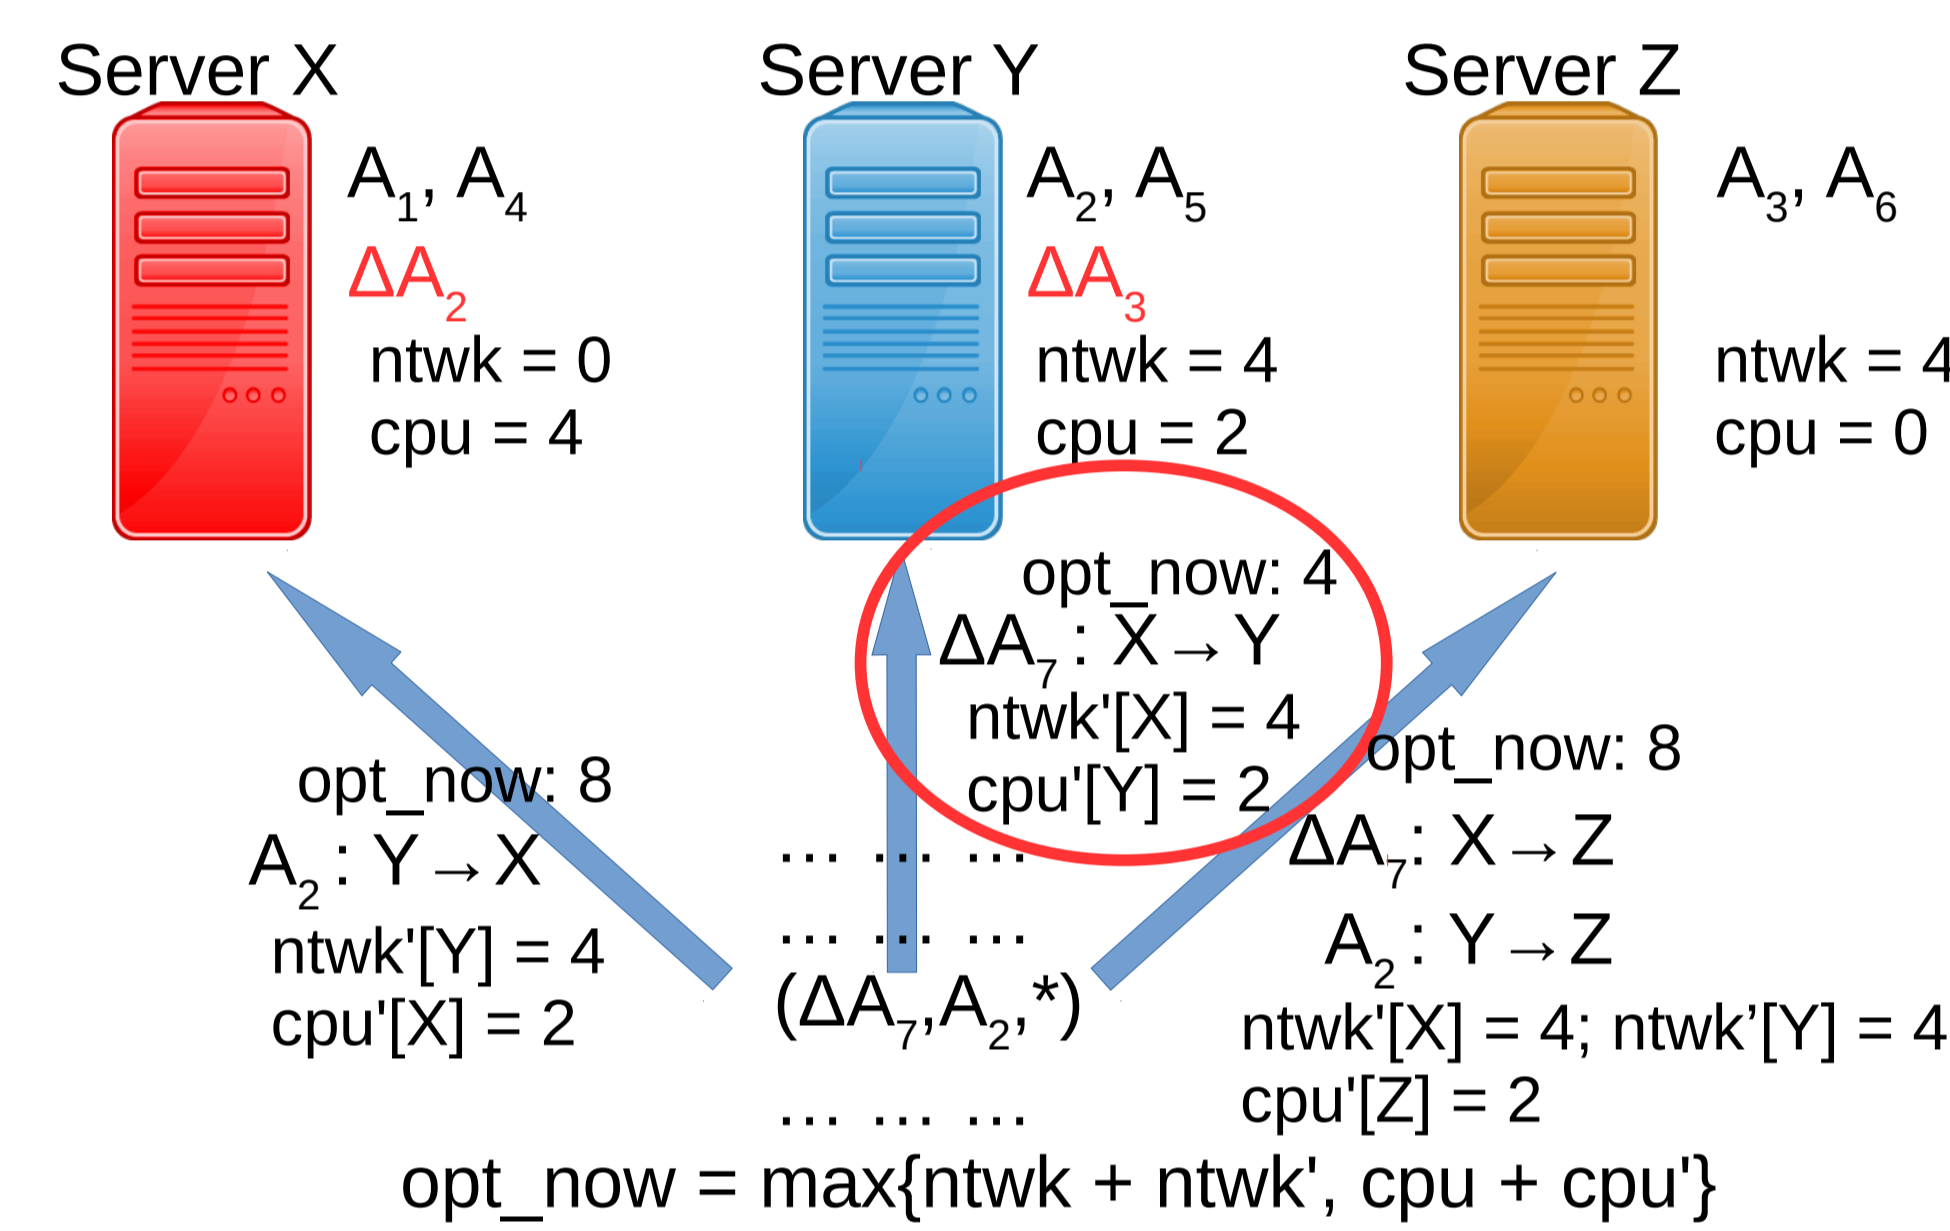
Optimal View Maintenance

- Array chunk (view chunk) reassignment** – ensures that we do not get stuck with an unfavorable static chunking strategy.
- Incoming chunks are not first assigned to a node based on a pre-determined chunking strategy.
- Piggyback** on the chunk replication incurred by view maintenance when computing the reassignment.
- A window of past batch updates is considered to avoid frequent unstable reassignments.

Analytical Cost Model

$$\min \left\{ \lambda \cdot \max \left\{ \max_{k, diff \in U_0} \{ ntwk_{chunk} + ntwk_{diff} \}, \max_k \{ cpu_{aggr} \} \right\} \right. \\ \left. + (1 - \lambda) \cdot \max \left\{ \max_{k, diff \in U} \{ ntwk_{chunk} + ntwk_{diff} \}, \max_k \{ cpu_{aggr} \} \right\} \right\}$$

Differential View Computation




View Chunk Reassignment

server	ntwk	cpu	join result	server	
X	32	36	$J_1: \Delta A_1 \bowtie A_1$	X	
Y	36	30	$J_2: \Delta A_4 \bowtie A_1$	X	
Z	30	35	$J_3: \Delta A_2 \bowtie A_1$	Y	
$V_1 \rightarrow$	transfers		$ntwk'$	cpu'	opt_now
X	J_3		$Y = 4$	$X = 6$	42
Y	J_1, J_2		$X = 8$	$Y = 6$	40
Z	J_1, J_2, J_3	$X = 8, Y = 4$	$Z = 6$		41

Array Chunk Reassignment

score	server	views	cpu_thr
A_2, V_1	X	V_2, V_6	4
A_1, V_1	Y	V_1, V_4, V_7	3
A_1, V_2	Z	V_3, V_5, V_8	1
A_2, V_3		A_1, A_2, A_3, \dots	
A_3, V_3	size	1 1 1	...
...	replica	X, Z Y, Z Z, Y	

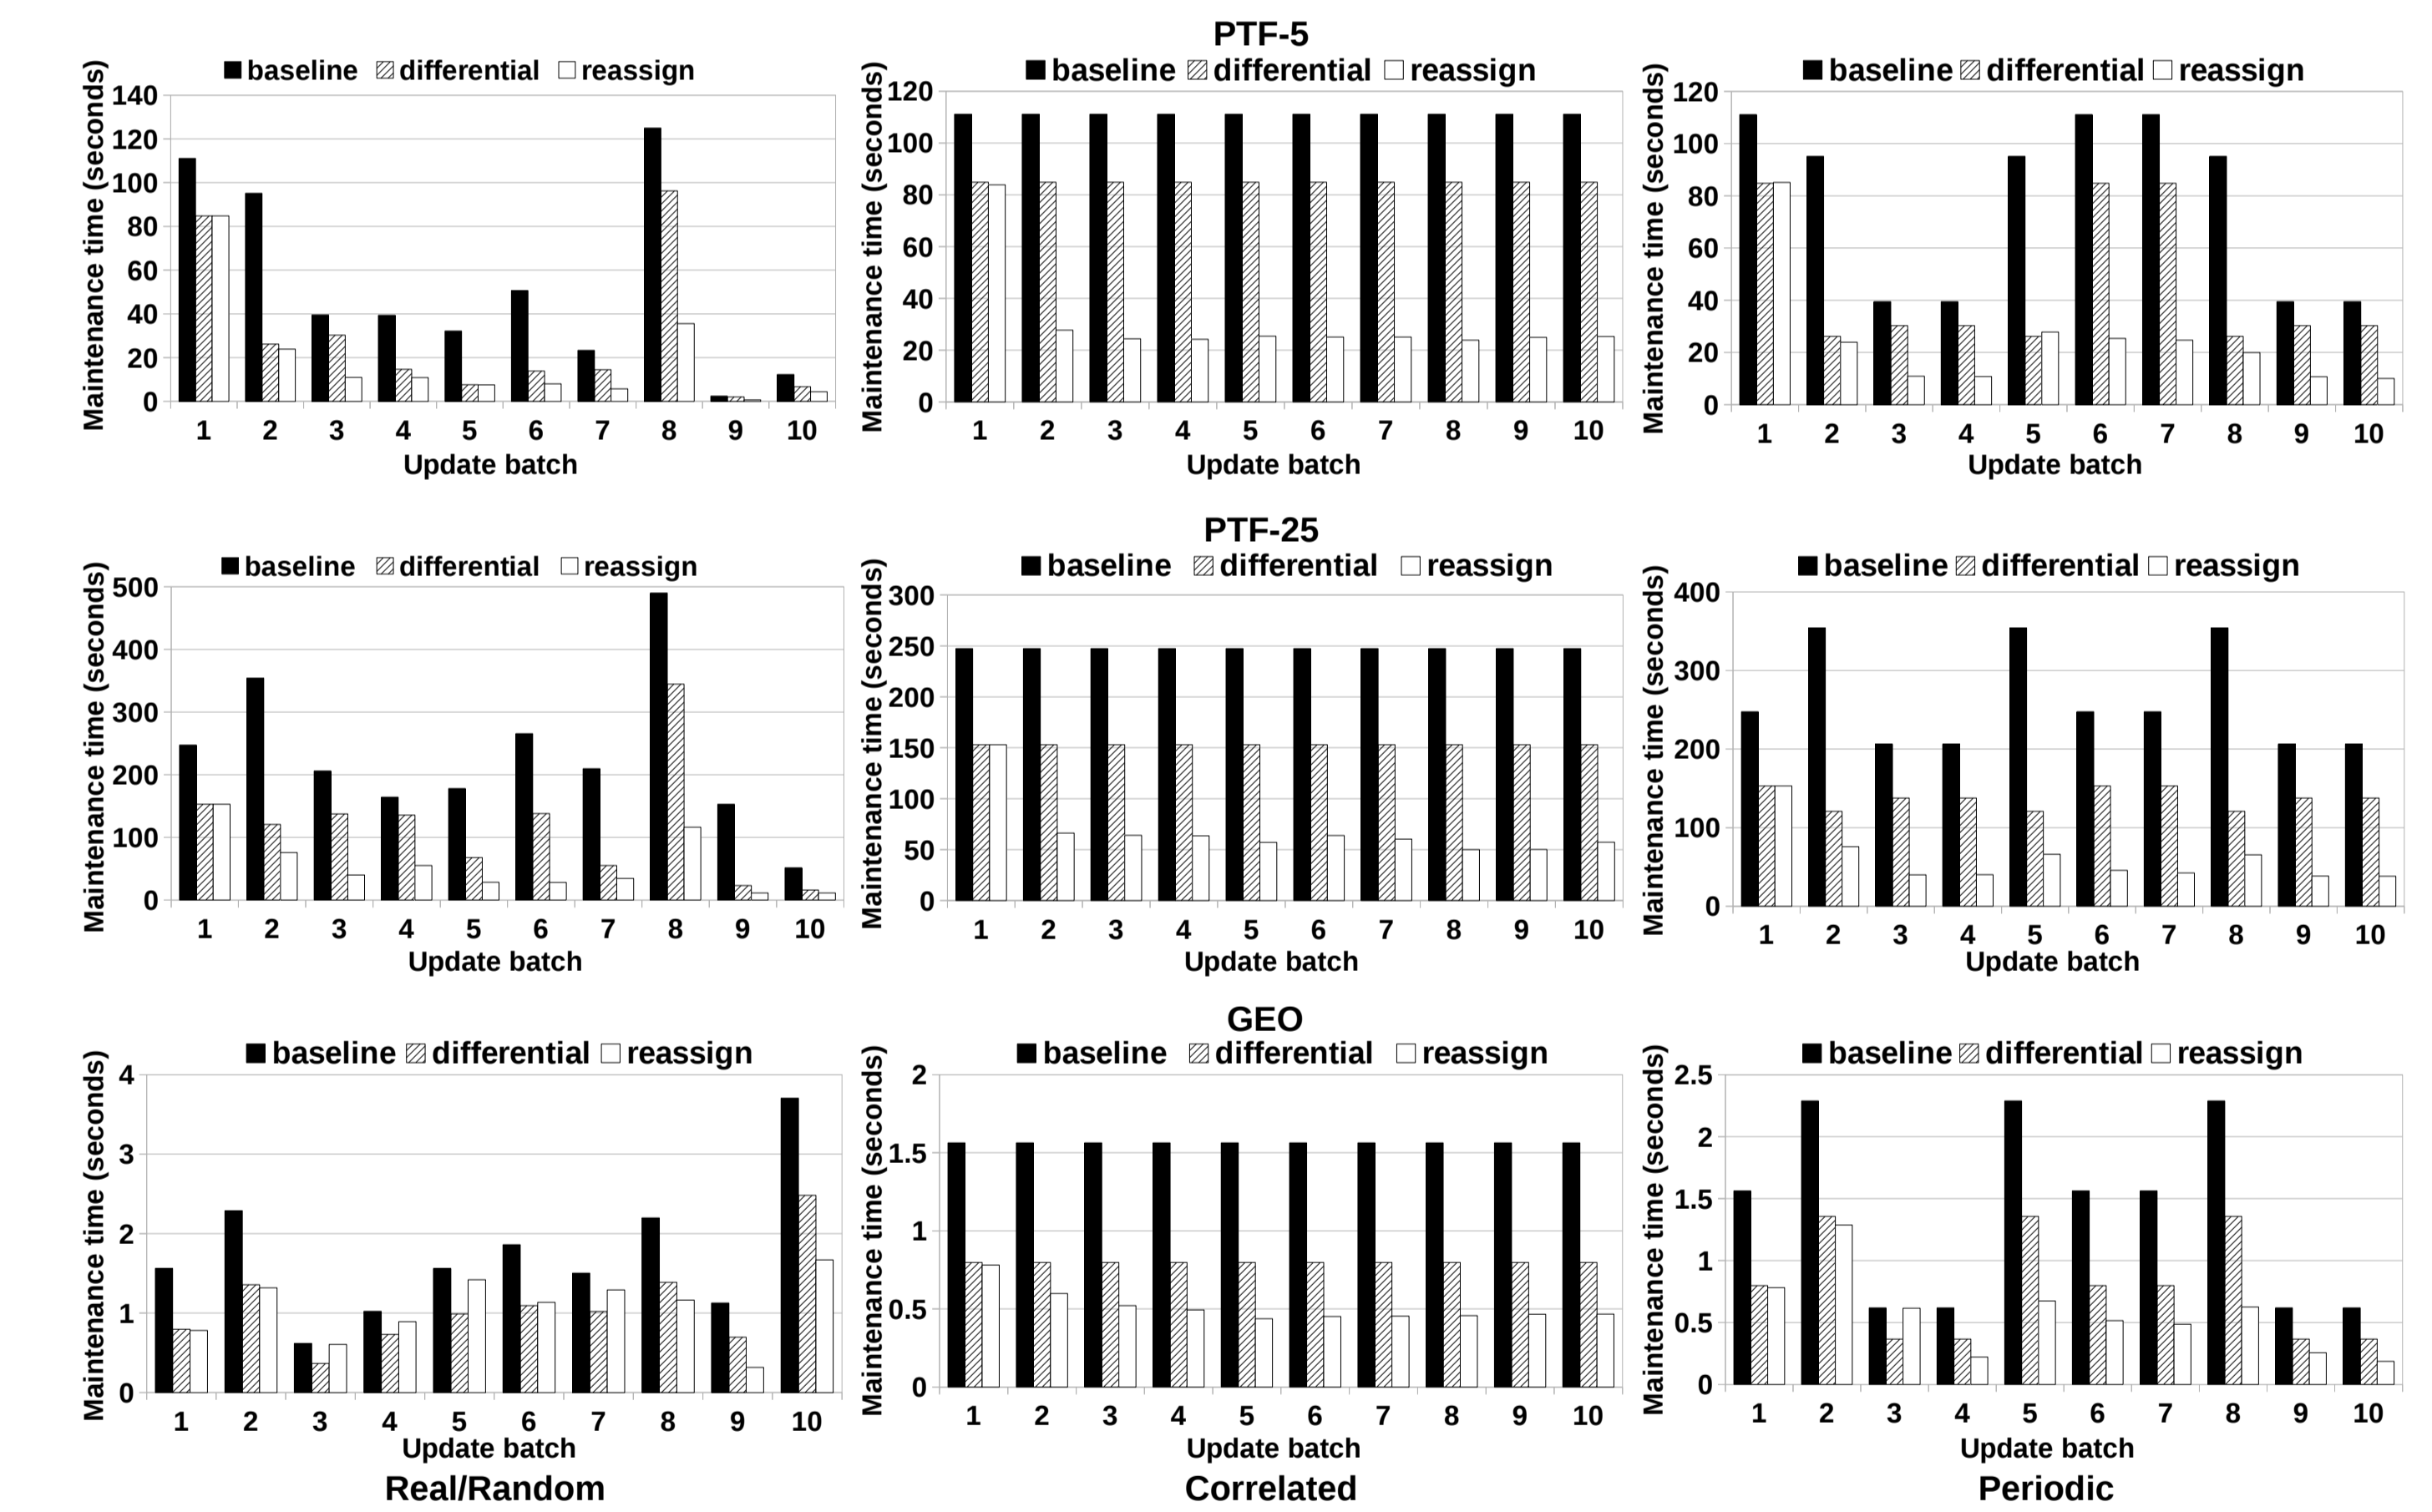
Query Integration




Experiments

PTF catalog: PTF [time=1, 153064; ra=1, 100000; dec=1, 50000] 1 billion objects, 343 GB
LinkedGeoData: GEO [long=1, 100000; lat=1, 50000] 30 million objects, 1 GB


View maintenance time:



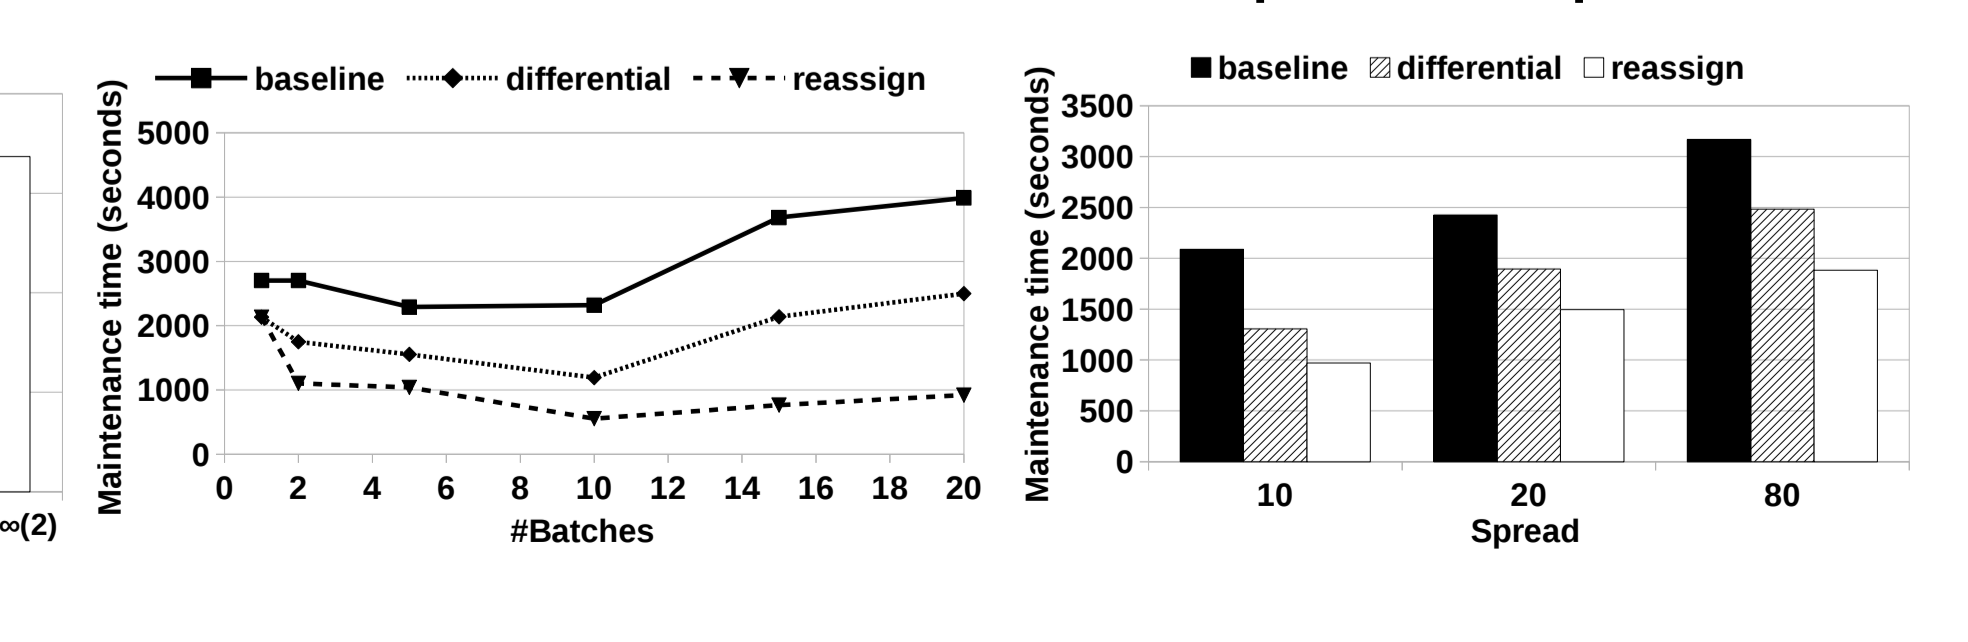
Average optimization time per update batch:



Query Integration



#Batches



Update Spread

