



Incremental View Maintenance over Array Data


Weijie Zhao¹, Florin Rusu^{1,2},
Bin Dong², Kesheng Wu², Peter Nugent²

University of California, Merced¹
Lawrence Berkeley National Laboratory²


May 16, 2017

- 1 Astronomical Data Products
- 2 Array Views
- 3 Array View Maintenance
- 4 Experimental Evaluation
- 5 Conclusions

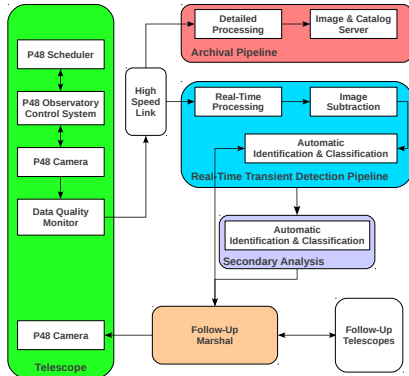
- Sloan Digital Sky Survey (SDSS), Palomar Transient Factory (PTF), etc.
- Transient identification: supernovae detection, exoplanet search




Array Construction



- Image $\xrightarrow{\text{Image processing}}$ Objects \rightarrow 3-D array [ra, dec, time]



- Image $\xrightarrow{\text{Image processing}}$ Objects \rightarrow 3-D array [ra, dec, time]
- Array is chunked and stored in a shared-nothing array database



- Image $\xrightarrow{\text{Image processing}}$ Objects \rightarrow 3-D array [ra, dec, time]
- Array is chunked and stored in a shared-nothing array database
- Array is sparse and skewed

- Gamma Ray Burst: extremely energetic explosions in distant galaxies that destroy everything around
 - For each detection, all its neighbor detections are within 12 hours with increasing magnitude
- M-Dwarf Flares: periodic flares with period larger than a month
 - No two neighbor detections within a month

- Neighborhood matching → Array similarity join (expensive)
- Astronomical data products → Materialized query result

- Neighborhood matching → Array similarity join (expensive)
- Astronomical data products → Materialized query result
- New detections every night → Base array updates
- Maintained data products → Incremental view maintenance
over arrays

Rare Supernova Discovery Ushers in New Era for Cosmology

Berkeley Lab astrophysicists develop novel method for finding gravitationally lensed Type 1a supernovae

News Release [LIR](#)

'Major breakthrough' as supernova four billion light years from Earth is captured from four different angles in a stunning world first

- Strange event occurred because light from the supernova bent through a galaxy
- Galaxies bend light through an effect that is called 'gravitational lensing'
- The galaxy magnified the supernova 50 times to give astronomers a unique view
- Supernovae are important to study as they reveal how the universe expands

By [HARRY PETTIT FOR MAILONLINE](#)

PUBLISHED: 14:00 EDT, 20 April 2017 | **UPDATED:** 14:01 EDT, 20 April 2017

Astronomers Get Rare View Of Type Ia Supernova Magnified 50 Times

21 April 2017, 9:01 am EDT By [Luan Chan](#) Tech Times

- Formal array view definition

- Formal array view definition
- Model incremental array view maintenance as an optimization formulation – including adaptive array and view reorganization

- Formal array view definition
- Model incremental array view maintenance as an optimization formulation – including adaptive array and view reorganization
- Propose an effective three-stage heuristics

- Formal array view definition
- Model incremental array view maintenance as an optimization formulation – including adaptive array and view reorganization
- Propose an effective three-stage heuristics
- Introduce an analytical cost model for query integration

- Formal array view definition
- Model incremental array view maintenance as an optimization formulation – including adaptive array and view reorganization
- Propose an effective three-stage heuristics
- Introduce an analytical cost model for query integration
- Experimentally compare against existing solutions on PTF catalog and LinkedGeoData

- Views over array similarity join (Zhao 2016)
 - Encompasses a larger variety of join predicates
 - Composable

- Views over array similarity join (Zhao 2016)
 - Encompasses a larger variety of join predicates
 - Composable
- Batch updates – (He 2005, Katsis 2015)
 - Not only to reduce the per-transaction overhead
 - Share delta computation inside batch

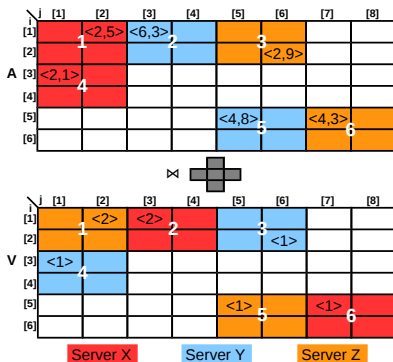
- Views over array similarity join (Zhao 2016)
 - Encompasses a larger variety of join predicates
 - Composable
- Batch updates – (He 2005, Katsis 2015)
 - Not only to reduce the per-transaction overhead
 - Share delta computation inside batch
- Update granularity – chunk (Brown 2010, Duggan 2014)

- Views over array similarity join (Zhao 2016)
 - Encompasses a larger variety of join predicates
 - Composable
- Batch updates – (He 2005, Katsis 2015)
 - Not only to reduce the per-transaction overhead
 - Share delta computation inside batch
- Update granularity – chunk (Brown 2010, Duggan 2014)
- Aggregates – standard SQL aggregate functions, e.g., SUM, COUNT, AVG (Yang 2001, Palpanas 2002)

- Views over array similarity join (Zhao 2016)
 - Encompasses a larger variety of join predicates
 - Composable
- Batch updates – (He 2005, Katsis 2015)
 - Not only to reduce the per-transaction overhead
 - Share delta computation inside batch
- Update granularity – chunk (Brown 2010, Duggan 2014)
- Aggregates – standard SQL aggregate functions, e.g., SUM, COUNT, AVG (Yang 2001, Palpanas 2002)
- Distributed processing (Luo 2003, Liu 2005)
 - Potential join pairs for each chunk are magnified

- Views over array similarity join (Zhao 2016)
 - Encompasses a larger variety of join predicates
 - Composable
- Batch updates – (He 2005, Katsis 2015)
 - Not only to reduce the per-transaction overhead
 - Share delta computation inside batch
- Update granularity – chunk (Brown 2010, Duggan 2014)
- Aggregates – standard SQL aggregate functions, e.g., SUM, COUNT, AVG (Yang 2001, Palpanas 2002)
- Distributed processing (Luo 2003, Liu 2005)
 - Potential join pairs for each chunk are magnified
- Recursive maintenance (Ahmad 2012, Nikolic 2016)

- 1 Astronomical Data Products
- 2 Array Views
- 3 Array View Maintenance
- 4 Experimental Evaluation
- 5 Conclusions



```

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
    
```

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

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



$$V \leftarrow \oplus_1 \left(\dots \oplus_k \left(\alpha_1 \bowtie_{\sigma_1, f_1}^{\mathcal{M}_1} \alpha_2 \bowtie_{\sigma_2, f_2}^{\mathcal{M}_2} \dots \bowtie_{\sigma_{n-1}, f_{n-1}}^{\mathcal{M}_{n-1}} \alpha_n \right) \dots \right)$$

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
```




- 1 Astronomical Data Products
- 2 Array Views
- 3 Array View Maintenance**
- 4 Experimental Evaluation
- 5 Conclusions

Baseline Array View Maintenance



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

Baseline Array View Maintenance



- For each update batch ΔD :
 - 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 Array View Maintenance

- Array chunk (view chunk) reassignment – not stuck with an unfavorable static chunking strategy

Optimal Array View Maintenance

- Array chunk (view chunk) reassignment – not stuck with an unfavorable static chunking strategy
- Incoming chunks are not first assigned to node based on pre-determined chunking strategy

Optimal Array View Maintenance





- Array chunk (view chunk) reassignment – not stuck with an unfavorable static chunking strategy
- Incoming chunks are not first assigned to node based on pre-determined chunking strategy
- Piggyback on the chunk replication incurred by view maintenance when computing the reassignment

Optimal Array View Maintenance

- Array chunk (view chunk) reassignment – not stuck with an unfavorable static chunking strategy
- Incoming chunks are not first assigned to node based on 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 unstable reassignments

Optimal Array View Maintenance

- Array chunk (view chunk) reassignment – not stuck with an unfavorable static chunking strategy
- Incoming chunks are not first assigned to node based on 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 unstable reassignments



$$\min \left\{ \lambda \cdot \max \left\{ \max_k \left\{ \sum_{i,j}^{k \neq j} x_{ikj} \cdot B_i T_{ntwk} + \sum_{p,q,v,j}^{k \neq j, (p,q,v) \in U_0} z_{pqk} \cdot y_{vj} \cdot B_{pq} T_{ntwk} \right\}, \right. \right. \\ \left. \left. \max_k \left\{ \sum_{p,q} z_{pqk} \cdot B_{pq} T_{cpu} \right\} \right\} \right. \\ \left. + (1 - \lambda) \cdot \max \left\{ \max_k \left\{ \sum_{i,j}^{k \neq j} x'_{ikj} \cdot B_i T_{ntwk} + \sum_{p,q,v,j,l}^{k \neq j, (p,q,v) \in U_l} z'_{pqk} \cdot y_{vj} \cdot W_l B_{pq} T_{ntwk} \right\}, \right. \right. \\ \left. \left. \max_k \left\{ \sum_{p,q,l}^{(p,q) \in U_l} z'_{pqk} \cdot W_l B_{pq} T_{cpu} \right\} \right\} \right\}$$

Network time to transfer chunks

Network time to transfer partial join results

CPU time to compute aggregations

Complicated formulation \rightarrow decoupling

Complicated formulation \rightarrow decoupling

- Differential view computation $\rightarrow x_{ikj}, z_{pqk}$
- View chunk reassignment $\rightarrow y_{vj}$
- Array chunk reassignment $\rightarrow y_{aj}$

Differential View Computation

Server X



A_1, A_4

ΔA_2

ntwk = 0

cpu = 4

Server Y



A_2, A_5

ΔA_3

ntwk = 4

cpu = 2

Server Z



A_3, A_6

ntwk = 4

cpu = 0

opt_now: 8
 $A_2: Y \rightarrow X$
 ntwk'[Y] = 4
 cpu'[X] = 2


opt_now: 4
 $\Delta A_7: X \rightarrow Y$
 ntwk'[X] = 4
 cpu'[Y] = 2

... ..

 $(\Delta A_7, A_2, *)$

opt_now: 8
 $\Delta A_7: X \rightarrow Z$
 $A_2: Y \rightarrow Z$
 ntwk'[X] = 4; ntwk'[Y] = 4
 cpu'[Z] = 2


$$\text{opt_now} = \max\{\text{ntwk} + \text{ntwk}', \text{cpu} + \text{cpu}'\}$$



View Chunk Reassignment


server	<i>ntwk</i>	<i>cpu</i>	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	<i>ntwk'</i>	<i>cpu'</i>	<i>opt_now</i>
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	8	X	V_2, V_6	4			
A_1, V_1	6	Y	V_1, V_4, V_7	3			
A_1, V_2	4	Z	V_3, V_5, V_8	1			
A_2, V_3	4		A_1	A_2	A_3	...	
A_3, V_3	2	size	1	1	1	...	
...	...	replica	X, Z	Y, Z	Z, Y		



- 1 Astronomical Data Products
- 2 Array Views
- 3 Array View Maintenance
- 4 Experimental Evaluation**
- 5 Conclusions

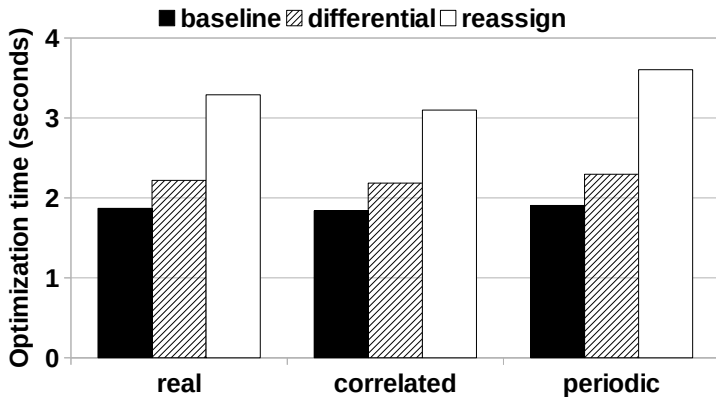
PTF [time=1,153064;ra=1,100000;dec=1,50000]

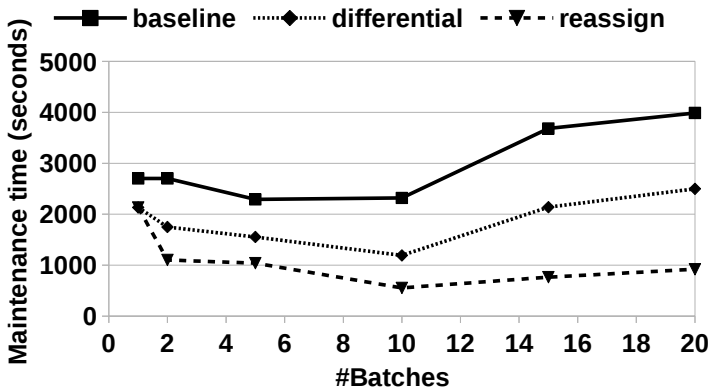
1 billion time-stamped objects, 343 GB



- Real: real updates from PTF pipeline
- Correlated: updates focusing on the same place
- Periodic: updates from several regions periodically – 1, 2, 3, 3, 2, 1, ...

Average Optimization Time per Update Batch





- 1 Astronomical Data Products
- 2 Array Views
- 3 Array View Maintenance
- 4 Experimental Evaluation
- 5 Conclusions**

- Introduce materialized array views as a database construct for derived data products in science

- Introduce materialized array views as a database construct for derived data products in science
- Model incremental array view maintenance with batch updates as an MIP optimization and give a three-stage heuristic that finds effective update plans

- Introduce materialized array views as a database construct for derived data products in science
- Model incremental array view maintenance with batch updates as an MIP optimization and give a three-stage heuristic that finds effective update plans
- Heuristic repartitions the array and the view continuously based on a window of past updates as a side-effect of view maintenance

- Introduce materialized array views as a database construct for derived data products in science
- Model incremental array view maintenance with batch updates as an MIP optimization and give a three-stage heuristic that finds effective update plans
- Heuristic repartitions the array and the view continuously based on a window of past updates as a side-effect of view maintenance
- Design an analytical cost model for integrating materialized array views in queries

- Introduce materialized array views as a database construct for derived data products in science
- Model incremental array view maintenance with batch updates as an MIP optimization and give a three-stage heuristic that finds effective update plans
- Heuristic repartitions the array and the view continuously based on a window of past updates as a side-effect of view maintenance
- Design an analytical cost model for integrating materialized array views in queries
- Experimental results confirm the effectiveness of the heuristics and the quality of the maintenance plan

- Introduce materialized array views as a database construct for derived data products in science
- Model incremental array view maintenance with batch updates as an MIP optimization and give a three-stage heuristic that finds effective update plans
- Heuristic repartitions the array and the view continuously based on a window of past updates as a side-effect of view maintenance
- Design an analytical cost model for integrating materialized array views in queries
- Experimental results confirm the effectiveness of the heuristics and the quality of the maintenance plan
- Incrementally maintain the production PTF “association table” and help find rare supernovae at Berkeley Lab and CalTech

Thank you!
Questions?