




Exploiting Incrementality with DBToaster










Monitoring Programs




Network Monitoring

Server Status

Task Allocations

Task Properties



Servers Per Task $>$ Task QOS?


Move Task to New Servers

Computational Advertising

Site Information

Available Ads

User Clicks





Good Ad Offers

Which Ad To Show?

Monitoring Programs

State Updates



Actions


Read

Views

On-Change


Reactions

Internal
State




Actions

Monitoring Programs




- **Existing Tools**
- DBToaster
- Cumulus


Stream Processors




Stream Processors



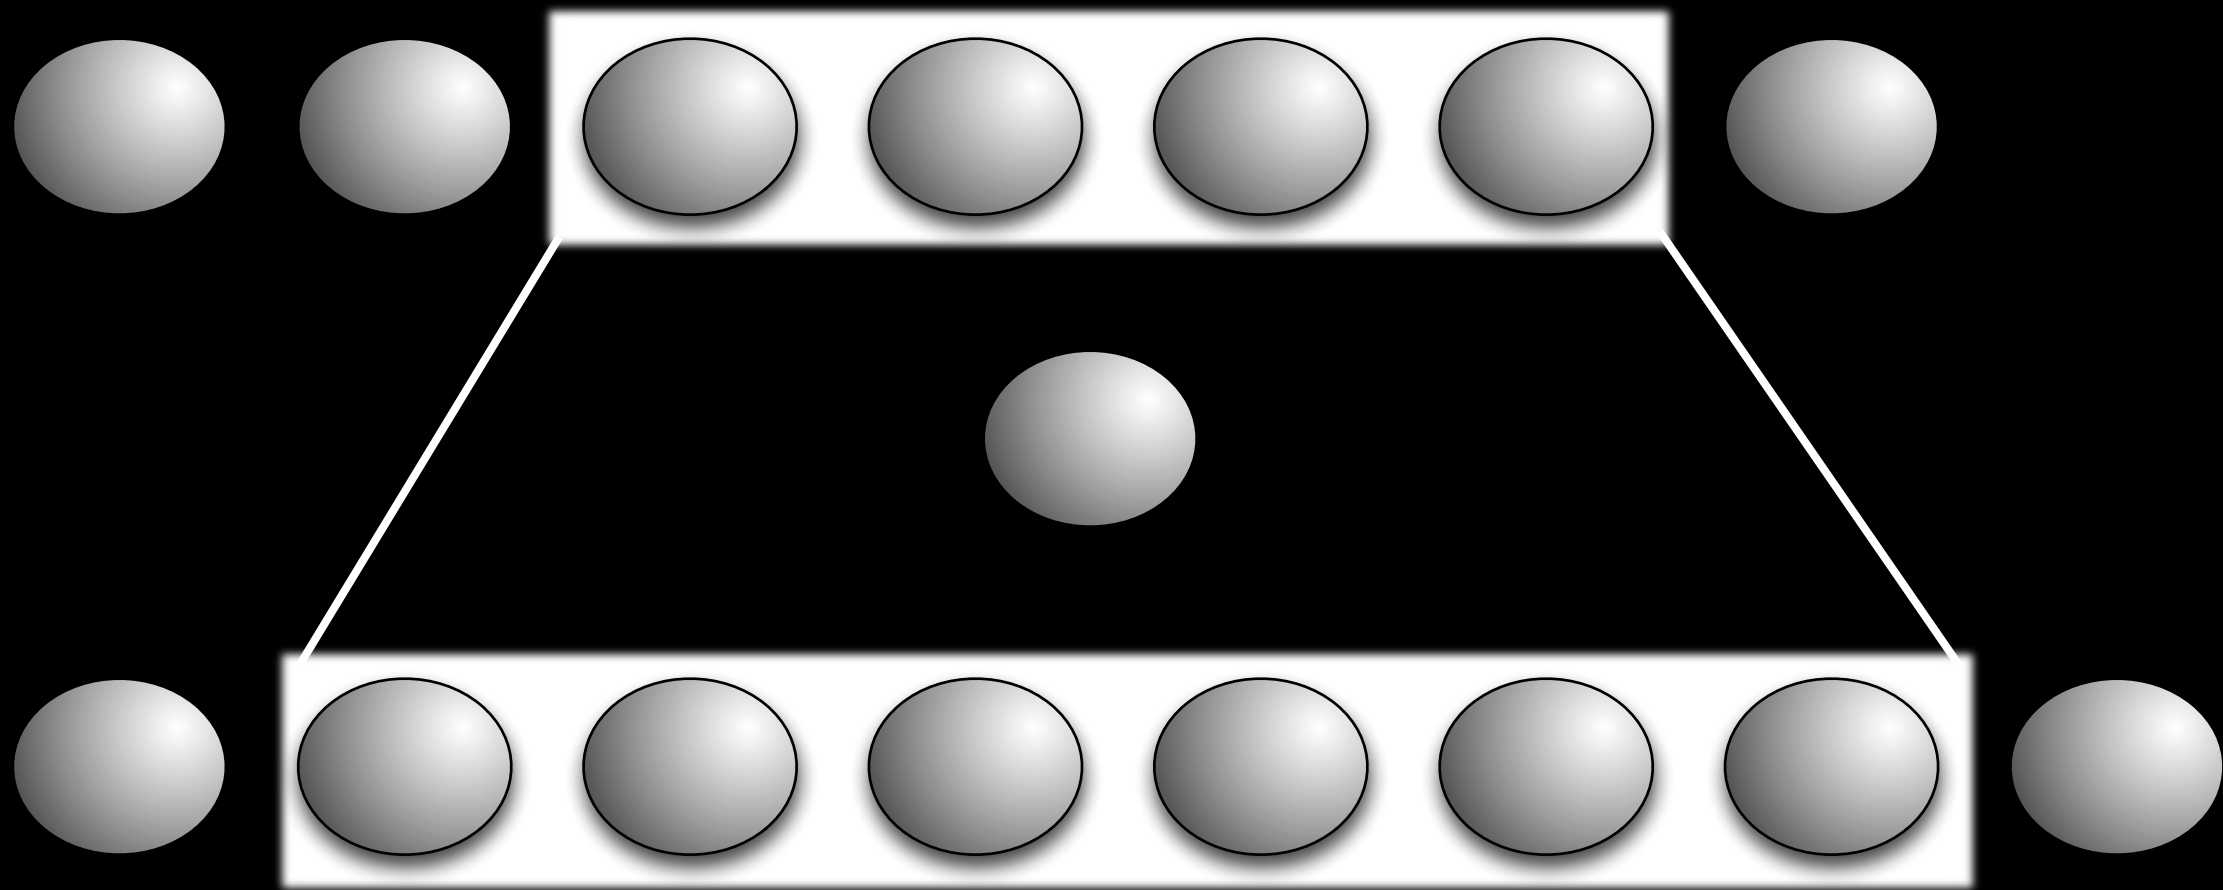
Stream Processors




Stream Processors




Stream Processors



Stream Processors



Stream Processors



Stream Processors

No Persistent State

Stream Processors

but also dynamic
No Persistent State
^

Incremental View Maintenance

QUERY

Incremental View Maintenance

ON CHANGE :

QUERY ::= $R \bowtie S \bowtie T$

Incremental View Maintenance

ON CHANGE :


QUERY $\text{+=} \Delta (R \bowtie S \bowtie T)$

↑
Simpler

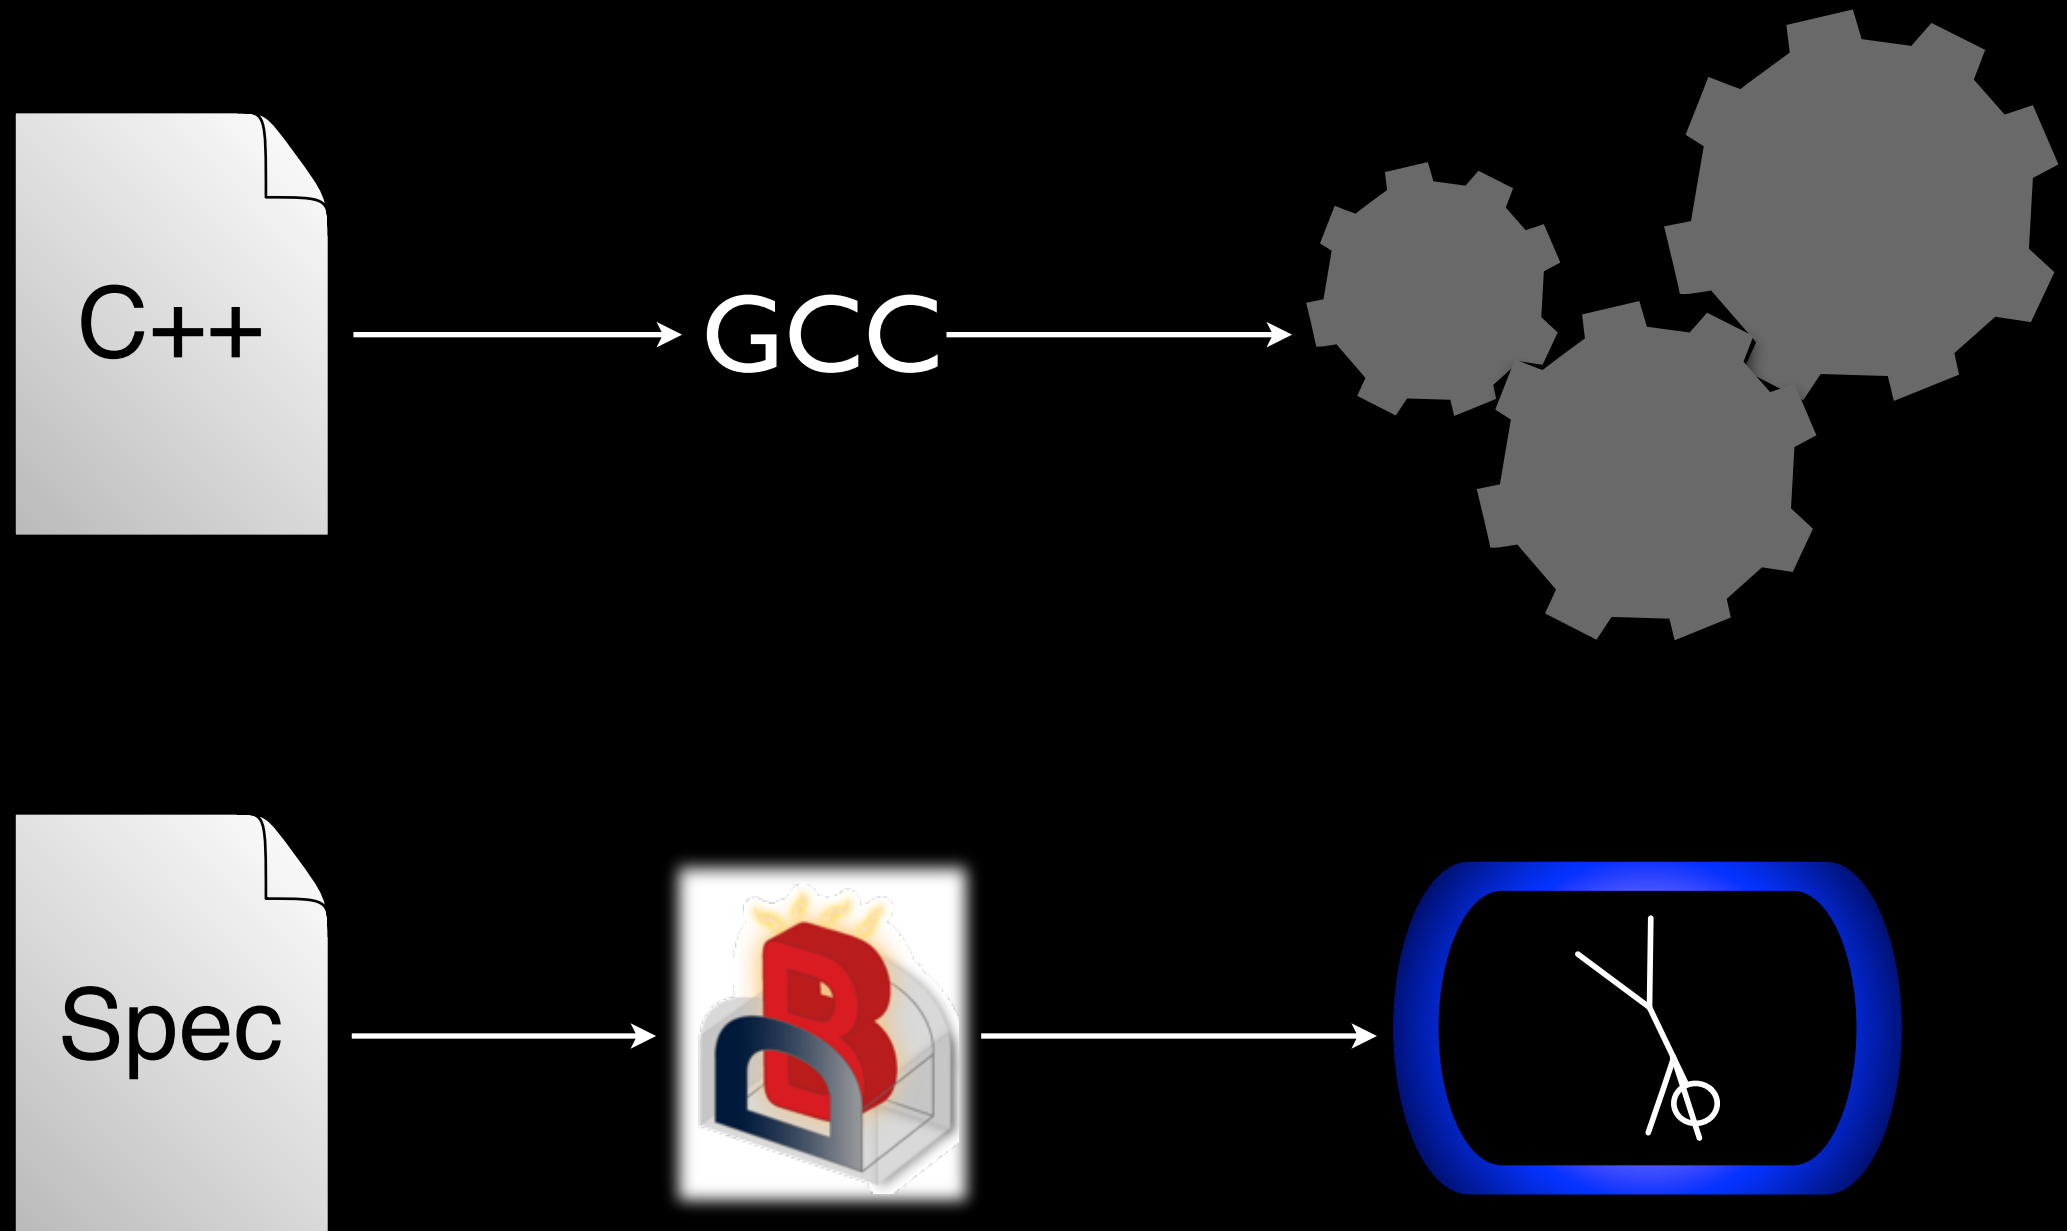
But still slow

- Existing Tools
- DBToaster
- Cumulus

DBToaster



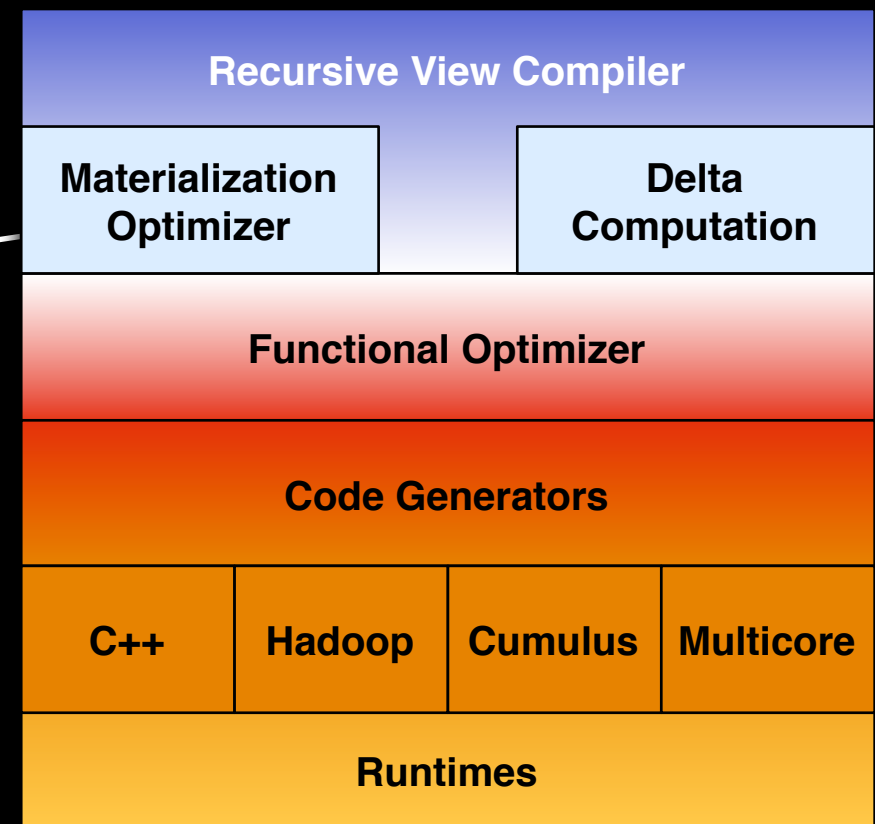
DBToaster





Toaster

- Exploit Incrementality ← Deltas
- Pick the Right Data Model
 - ... the right representation
 - ... understand the platform
- Borrow (Liberally) from Other Fields
 - ... use set-at-a-time optimizations (PL)
 - ... generate machine code (Compilers)
 - ... and others



Recursive Delta Compilation

ON ΔR :

$$\mathbf{QUERY} \ += \Delta_{\Delta R} (R \bowtie S \bowtie T)$$

Recursive Delta Compilation

ON ΔR :

$$\mathbf{QUERY} \ += \Delta_{\Delta R} (R \bowtie S \bowtie T)$$

$\Delta_{\Delta R} (R \bowtie S \bowtie T)$ is ^{Usually} \wedge **Simpler** than $R \bowtie S \bowtie T$

Δ is **Closed**

$\Delta_{\Delta R} (R \bowtie S \bowtie T)$ ^{Usually} \wedge has **Finite Support** for ΔR

(Koch, PODS '10)

Recursive Delta Compilation

QUERY ::=

SUM (R.A * T.D)
of

R (A, B) ⋈_B S (B, C) ⋈_C T (C, D)

Recursive Delta Compilation

ON $+R(\alpha, \beta) :$

QUERY $+= \alpha * \text{SUM}(T.D)$
of

$S(\beta, C) \bowtie_C T(C, D)$

Recursive Delta Compilation

ON $+R(\alpha, \beta) :$

QUERY $+= \alpha * \text{SUM}_{\beta}(T.D)$
of

$S(\beta, C) \bowtie_C T(C, D)$

$m_1[\beta] := \text{SUM}(T.D)$
of

$S(\beta, C) \bowtie_C T(C, D)$

Recursive Delta Compilation

ON $+R(\alpha, \beta) :$

QUERY $+= \alpha * m_1[\beta]$

ON $+S(\beta', \gamma) :$

$m_1[\beta'] += \text{SUM}(T.D)$
of

$T(\varnothing, D)$

Recursive Delta Compilation

ON $+R(\alpha, \beta) :$

QUERY $+= \alpha * m_1[\beta]$

ON $+S(\beta', \gamma) :$

$m_1[\beta'] += m_2[\gamma]$

$m_2[\gamma] := \text{SUM}(T.D) \text{ of } T(\gamma, D)$

Recursive Delta Compilation

ON $+R(\alpha, \beta) :$

QUERY $+= \alpha * m_1[\beta]$

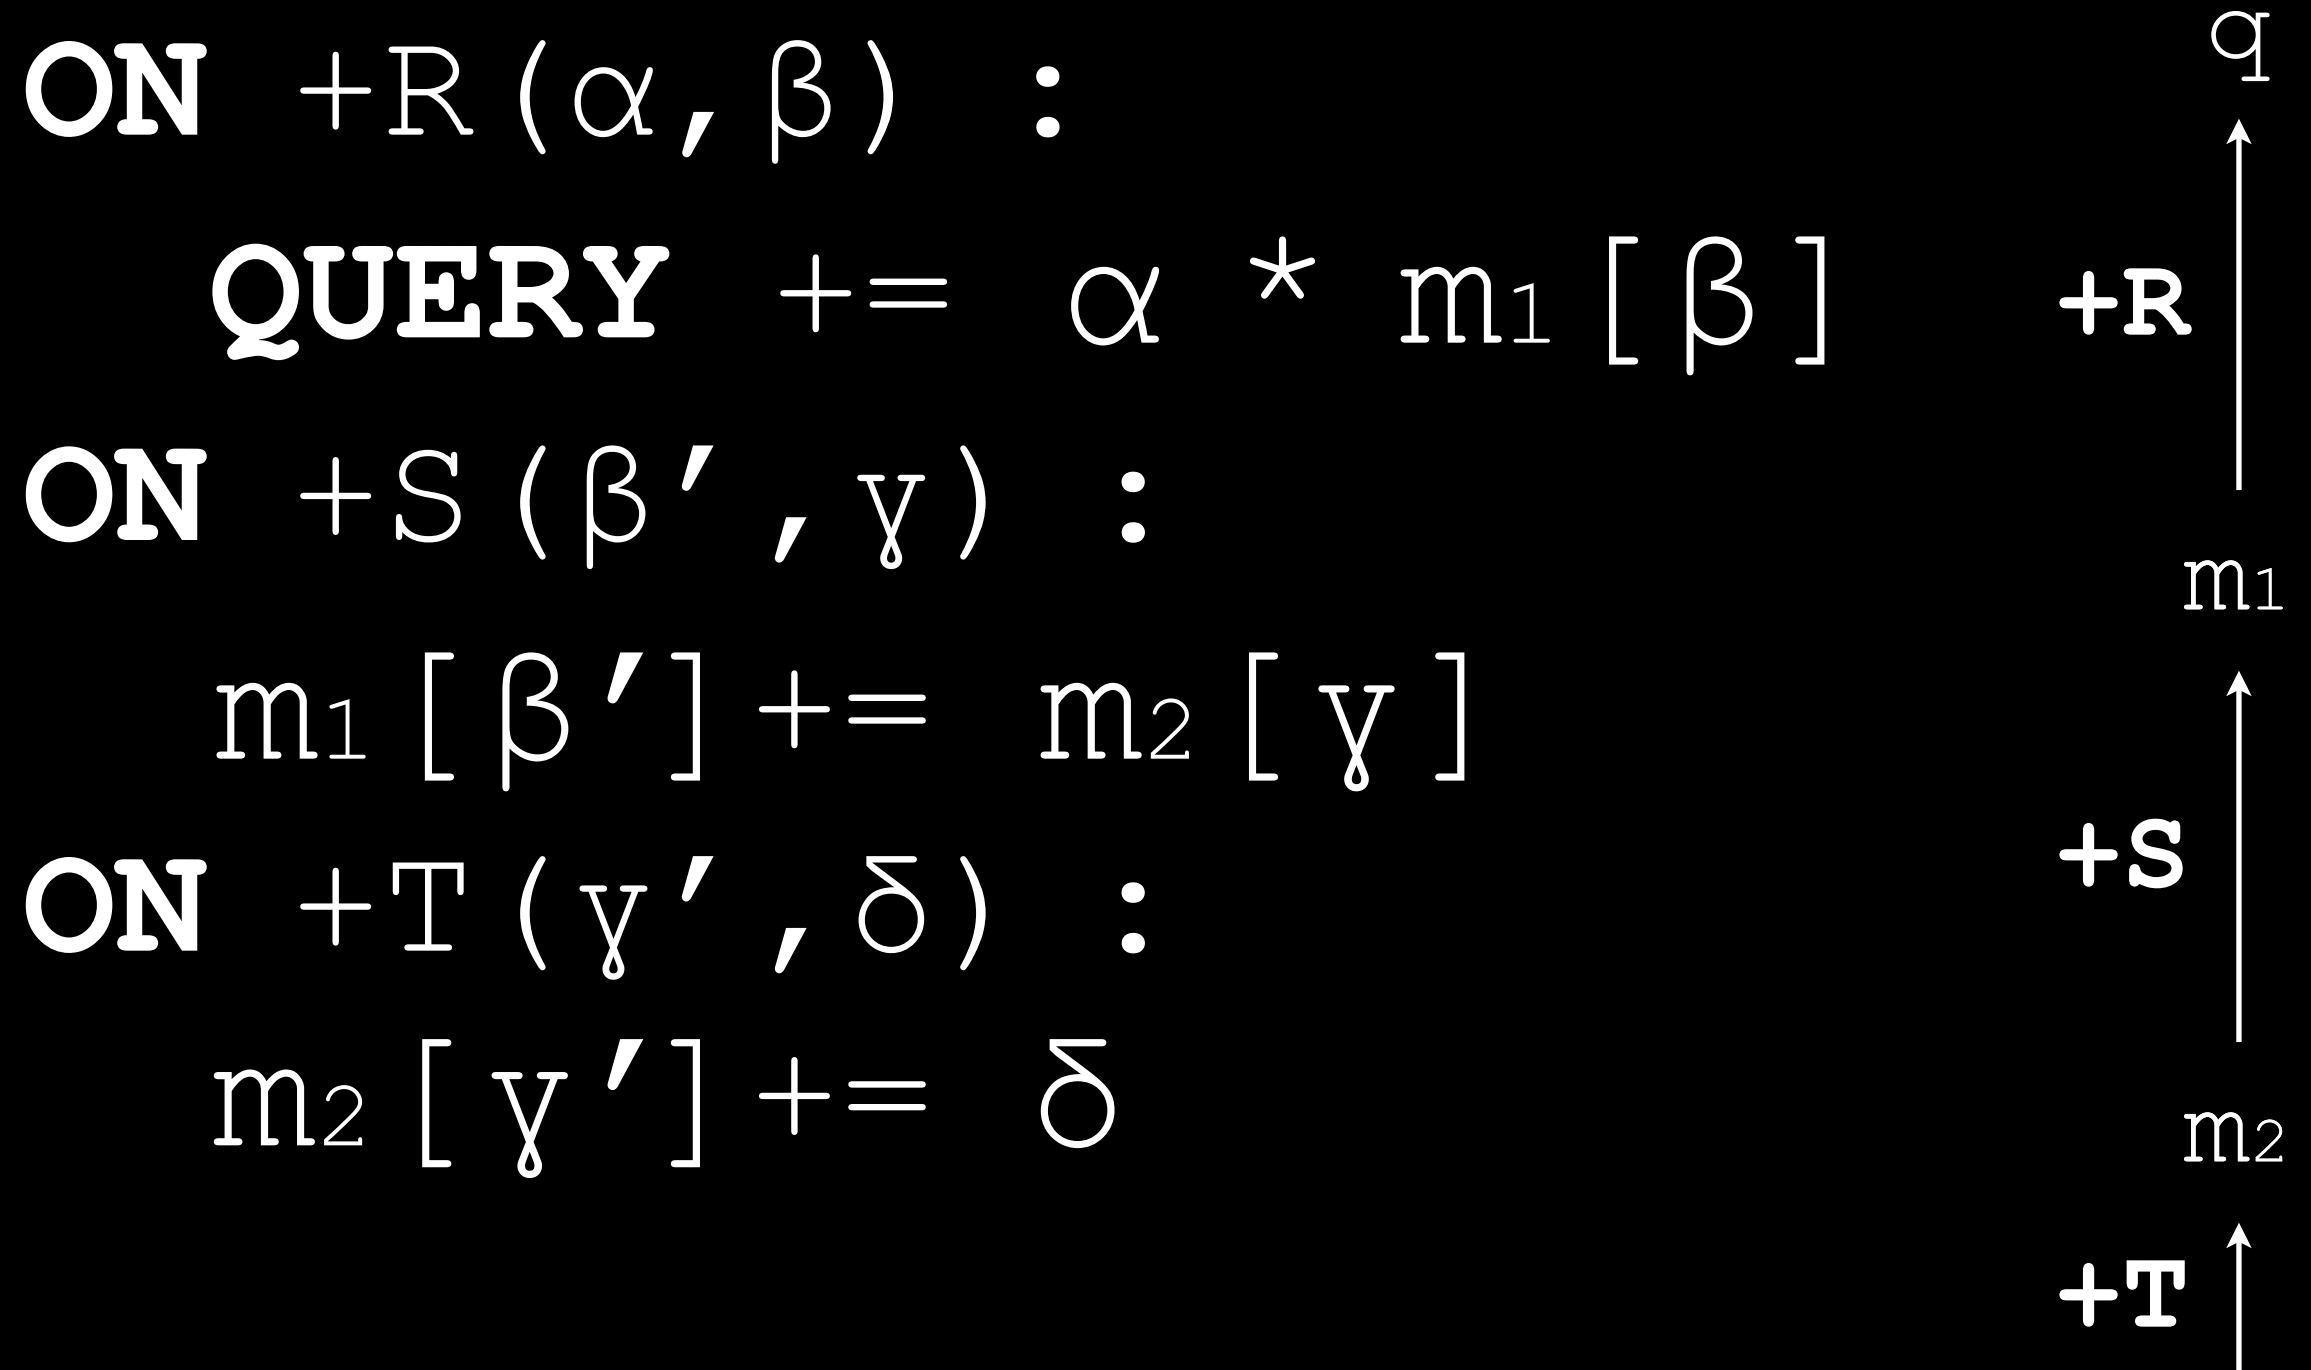
ON $+S(\beta', \gamma) :$

$m_1[\beta'] += m_2[\gamma]$

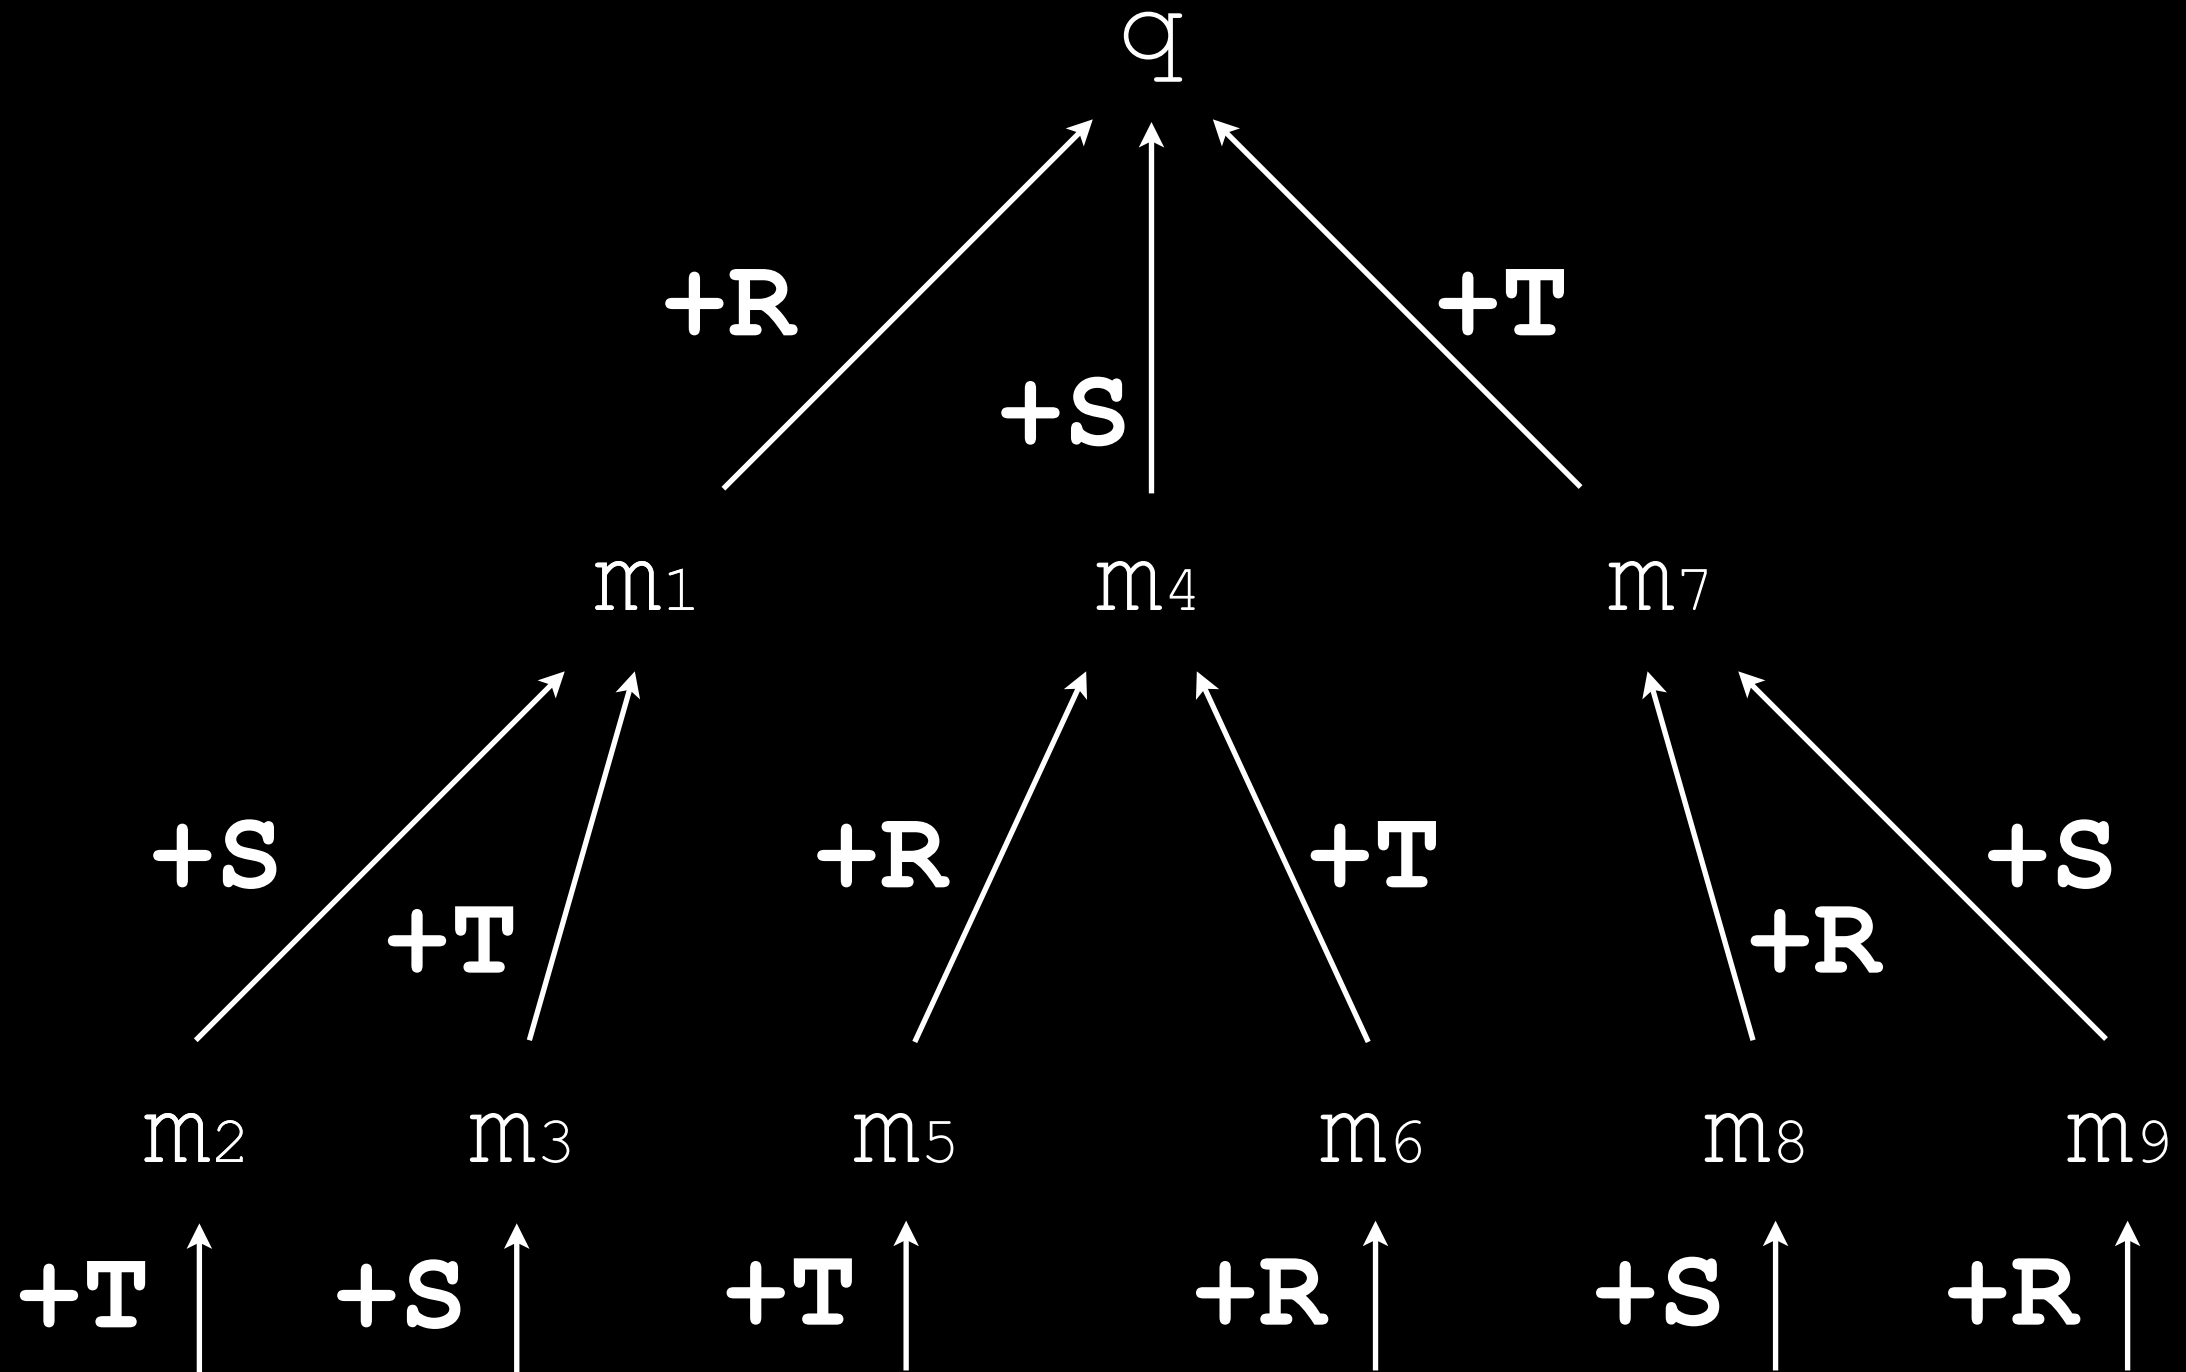
ON $+T(\gamma', \delta) :$

$m_2[\gamma'] += \text{SUM}(\delta)$

Recursive Delta Compilation



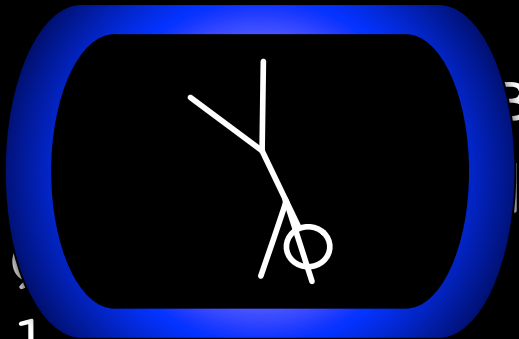
View Hierarchy



Maintenance Program

```
ON +R[ A, B ]:  
  QUERY[ ] += ( A * QUERY_dR[ B ] )  
  QUERY_dT[ C ] += FORALL C:( A * QUERY_dR_dT[ B, C ] )  
  QUERY_dS[ B ] += A
```


```
ON +S[ B, C ]:  
  QUERY[ ] += ( B * QUERY_dR_dS[ C ] )  
  QUERY_dT[ C ] += ( C * QUERY_dR_dS[ C ] )  
  QUERY_dR[ B ] += ( B * QUERY_dR_dS[ C ] )  
  QUERY_dR_dT[ B, C ] += 1.
```



```
ON +T[ C, D ]:  
  QUERY[ ] += ( QUERY_dT[ C ] * D )  
  QUERY_dR[ B ] += FORALL B:( D * QUERY_dR_dT[ B; C ] )  
  QUERY_dR_dS[ C ] += D
```

Maintenance Program

(DBToaster; CIDR '11)



But...

$\Delta_{\Delta R} (R \bowtie S \bowtie T)$ is ^{Usually} **Simpler** than $R \bowtie S \bowtie T$
Nested Subqueries

$\Delta_{\Delta R} (R \bowtie S \bowtie T)$ ^{Usually} has **Finite Support** for ΔR
Non-Equi-Joins

Nested Subqueries

QUERY ::=

COUNT () **of** R (A, B)

where A = (**SUM** (C) **of** S (C))

Nested Subqueries

Step 1:

SUM (C) **of** [S (C)]



Step 2:

COUNT () **of** R (A, B)

where A = [result of step 1]

Nested Subqueries

Step 2:

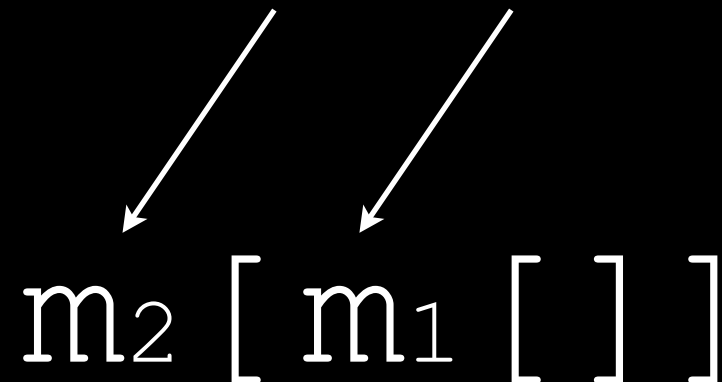
~~count~~ **COUNT** () **of** ~~step 1~~ **R** (A, B)

where A = [result of step 1]

$m_2 [A] := \text{COUNT} () \text{ of } R (A, B)$

Partial Materialization

Materialize the query in parts



Perform computations at maintenance-time

Non-Equality Predicates

QUERY ::=

COUNT () **of** $R(A) \times S(B, C)$
where $A < B$

Non-Equality Predicates

ON $+R(\alpha)$:

QUERY += $\text{SUM}(\text{m}_1([B]) \text{ of } S(B, C))$

where $\alpha < B$

Partial Materialization

$\text{m}_1[B] := \text{COUNT}() \text{ of } S(B, C)$

group by B

Materialization Optimizer


- Nested Subqueries
- Non-equality predicates
- Memory Constraints
- High Maintenance Cost
- Specialized Datastructures

WWAP

```
SELECT sum(b1.price
          * b1.volume)
FROM cumulativeTime
WHERE 0.25 *
      (SELECT sum(b3.volume)
       FROM Views Refreshing)
      >
      (SELECT sum(b2.volume)
       FROM bids b2
       WHERE b2.price >
            b1.price);
```

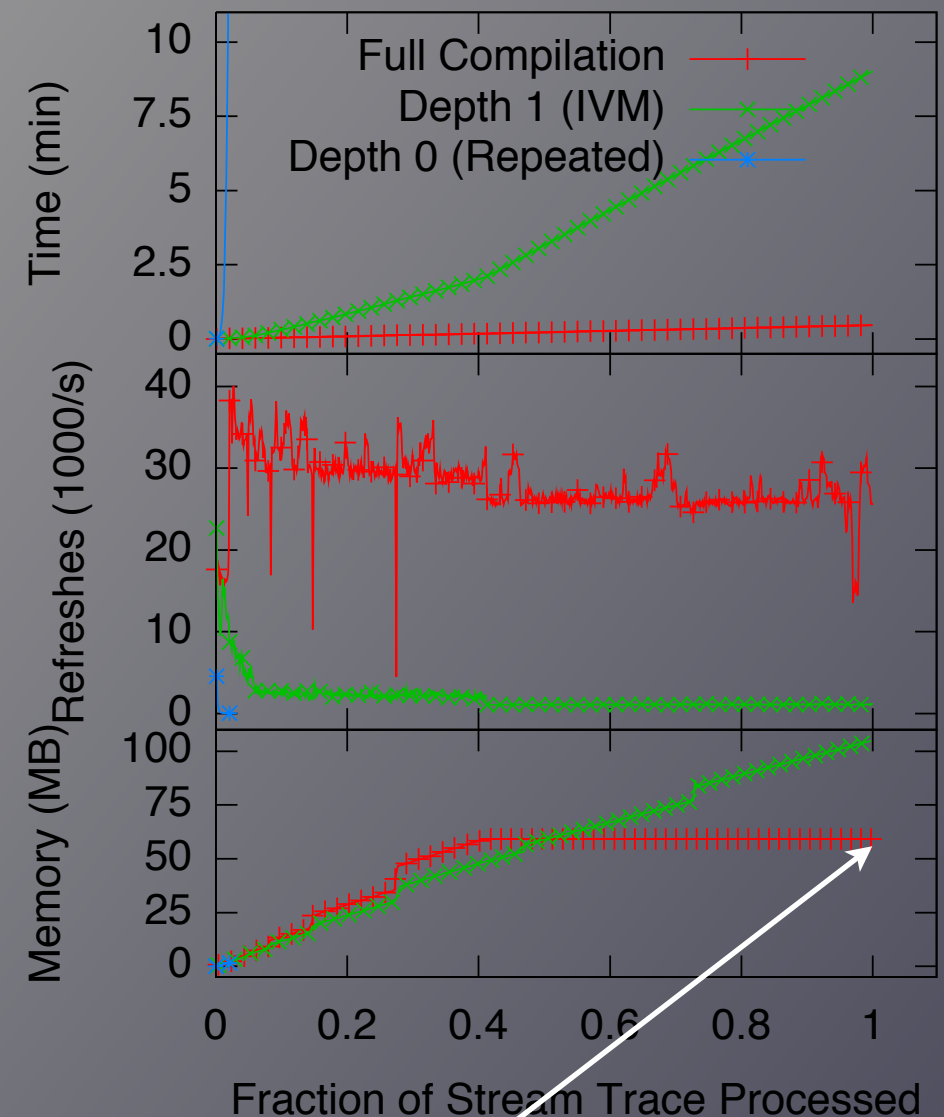
IVM & Naive

DBToaster



TPC-H Q3

```
SELECT ORDERS.orderkey,  
       ORDERS.orderdate,  
       ORDERS.shippriority,  
       SUM(extendedprice  
           * (1 - discount))  
FROM CUSTOMER, ORDERS, LINEITEM  
WHERE CUSTOMER.mktsegment  
       = 'BUILDING'  
AND ORDERS.custkey  
     = CUSTOMER.custkey  
AND LINEITEM.orderkey  
     = ORDERS.orderkey  
AND ORDERS.orderdate  
     < DATE('1995-03-15')  
AND LINEITEM.SHIPDATE  
     > DATE('1995-03-15')  
GROUP BY ORDERS.orderkey,  
         ORDERS.orderdate,  
         ORDERS.shippriority;
```



Half the Memory Usage

- Existing Tools
- DBToaster
- Cumulus

Maintenance Programs

Data-Parallel Computations

```
ON +R[ A, B ]:  
  QUERY[ ] += ( A * QUERY_dR[ B ] )  
  QUERY_dT[ C ] += FORALL C:( A * QUERY_dR_dT[ B, C ] )  
  QUERY_dS[ B ] += A
```

```
ON +S[ B, C ]:  
  QUERY[ ] += ( QUERY_dS[ B ] * QUERY_dR_dS[ C ] )  
  QUERY_dT[ C ] += QUERY_dS[ B ]  
  QUERY_dR[ B ] += QUERY_dR_dS[ C ]  
  QUERY_dR_dT[ B, C ] += 1.
```

```
ON +T[ C, D ]:  
  QUERY[ ] += ( QUERY_dT[ C ] * D )  
  QUERY_dR[ B ] += FORALL B:( D * QUERY_dR_dT[ B; C ] )  
  QUERY_dR_dS[ C ] += D
```

Key/Value Style Datastructures

Execution Model

ON Event(param 1, param2, ...)

Statement 1

Statement 2

Statement 3


...

Statement Execution

Read

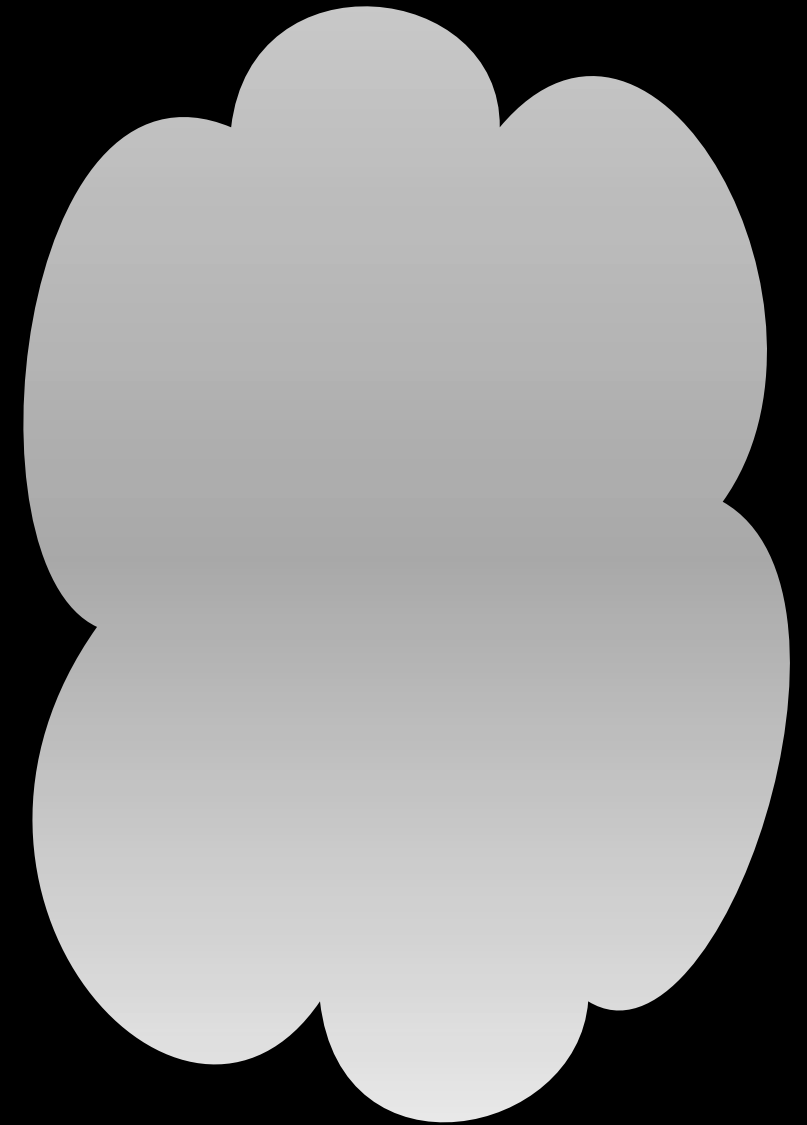




Compute

Write







[Old Version]




[New Version]











Epoch: 0








$S := \langle S, 0 \rangle$







Epoch: 0



Epoch: 0



Epoch: 0











$R: \langle R, 1 \rangle$

$T: \langle T, 0 \rangle$

2

Epoch: 0







Epoch: 0


R: <0,0,1>

T: <0,1,0>



$$M_2[1,3] += 2$$
$$\langle 0, 2, 4 \rangle$$




$$R$$
$$\langle 0, 1, 3 \rangle$$



R
<0,1,3>




$M_2 [\dots] +=$
 $M_3 [\dots] * M_4 [\dots]$






$\Sigma \delta$



$M_3 [\dots]$
 $M_4 [\dots]$





$M_3 [\dots] * M_4 [\dots]$





$M_2 [\dots] +=$...
<0,1,3>








$$M_2[1,3] += 2$$
$$\langle 0, 2, 4 \rangle$$







$$M_2[1,3] += 2$$
$$\langle 0, 2, 4 \rangle$$



$$\delta$$
$$\langle 0, 2, 5 \rangle$$



Open Challenges

- Data Placement
 - Migration
- Batch Processing
- Live Program Management

