

DBToaster

Higher-Order Delta Processing for Dynamic, Frequently Fresh Views

Yanif Ahmad

Johns Hopkins

Oliver Kennedy

University at Buffalo

Christoph Koch

EPFL

Milos Nikolic

EPFL

Realtime Monitoring Programs...

...Monitor The State of the World

...React to Conditions in that State

Realtime Monitoring Programs are Everywhere

Traffic Sources

62 keywords

Right now

597

Realtime Analytics

DIRECT ORGANIC REFERRAL PAID CAMPAIGN



ALL TRAFFIC

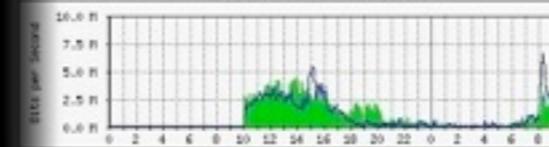
Metric Total: 597

Medium	Source
1. cpc	google
2. Direct	
3. Organic	google

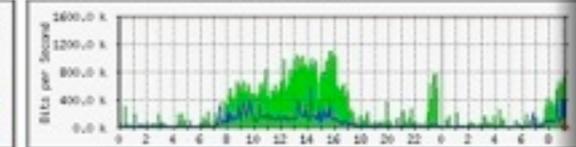
Pageviews



Corp to Engineering Ethernet - 172.21.0.17



Corp To Warehouse T1 - 172.20.0.41



Network Monitoring



Algorithmic Trading

4. Monitoring and

3. Process Enablers

- Ensure implementation and ongoing effectiveness

3. Process E

2. Policy

1. Level of Risk

Policy Monitoring

2. Policy Objectives

- Clear objectives and outcomes
- TBS/Dept well-defined responsibilities and accountabilities

1. Level of Risk

- Drives level of monitoring activities
- Resource effort

Clickstream Analysis

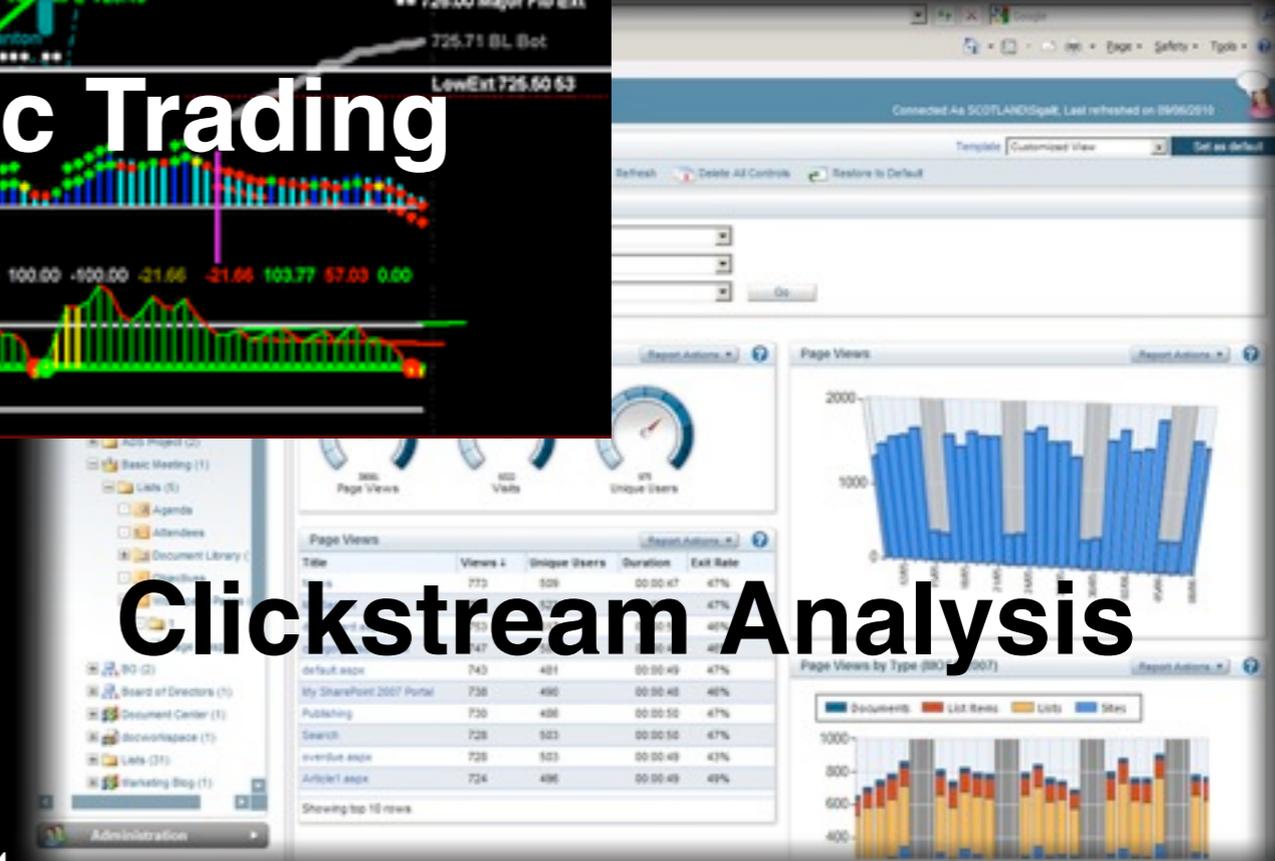
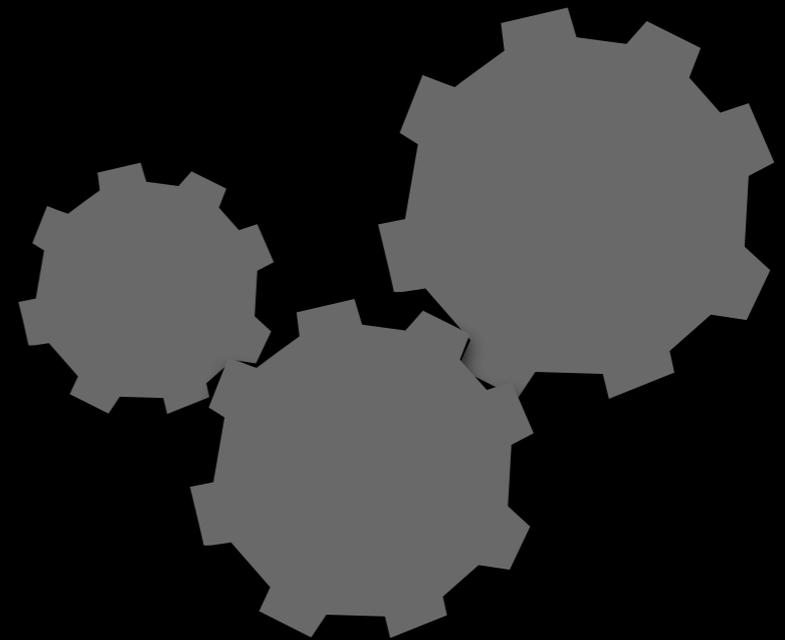
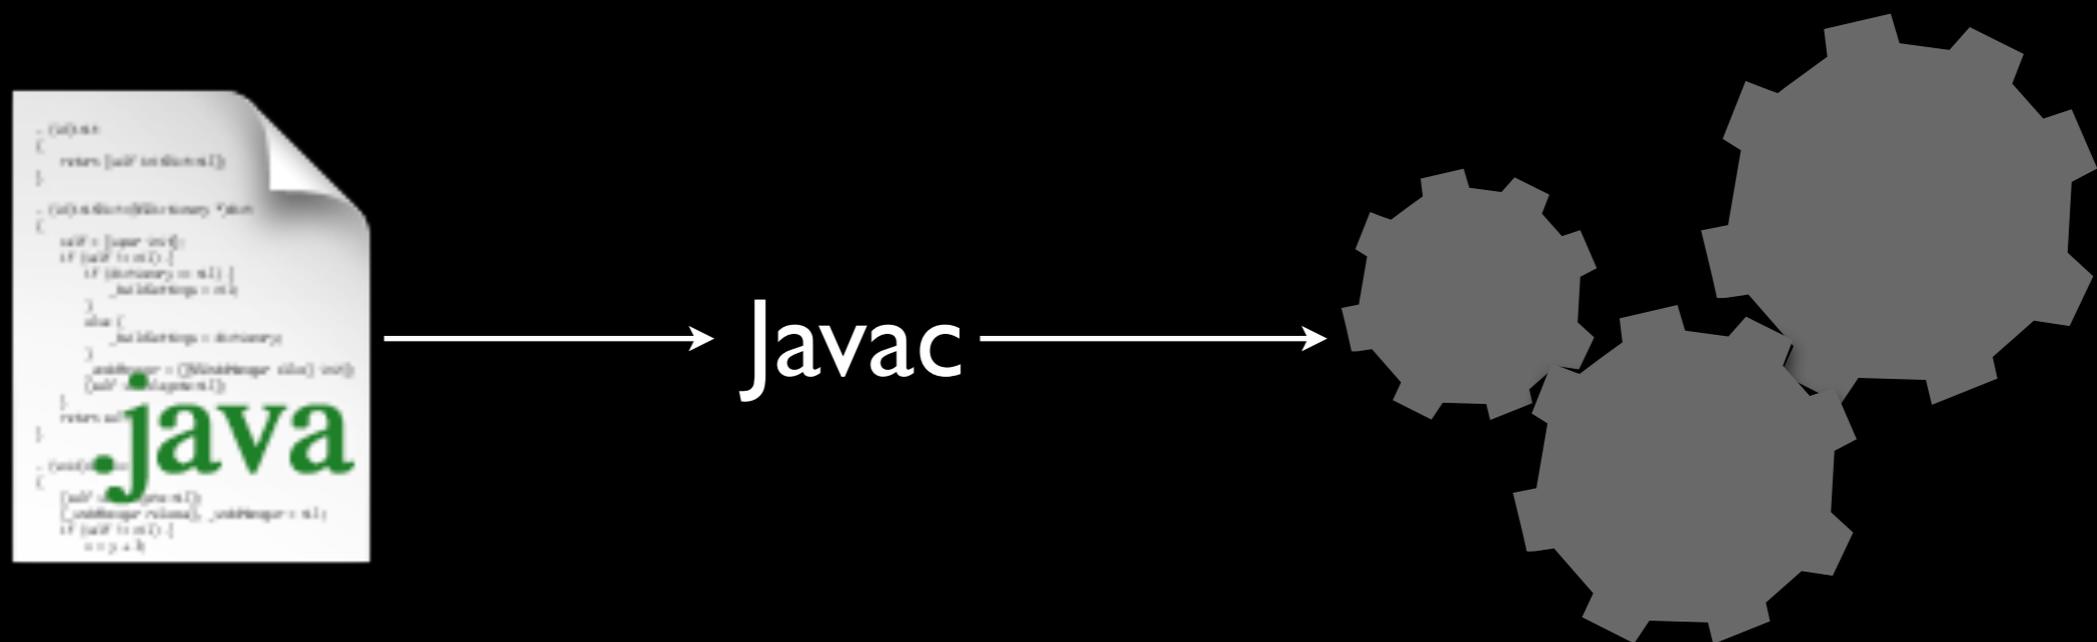


Figure 1

Monitoring Programs

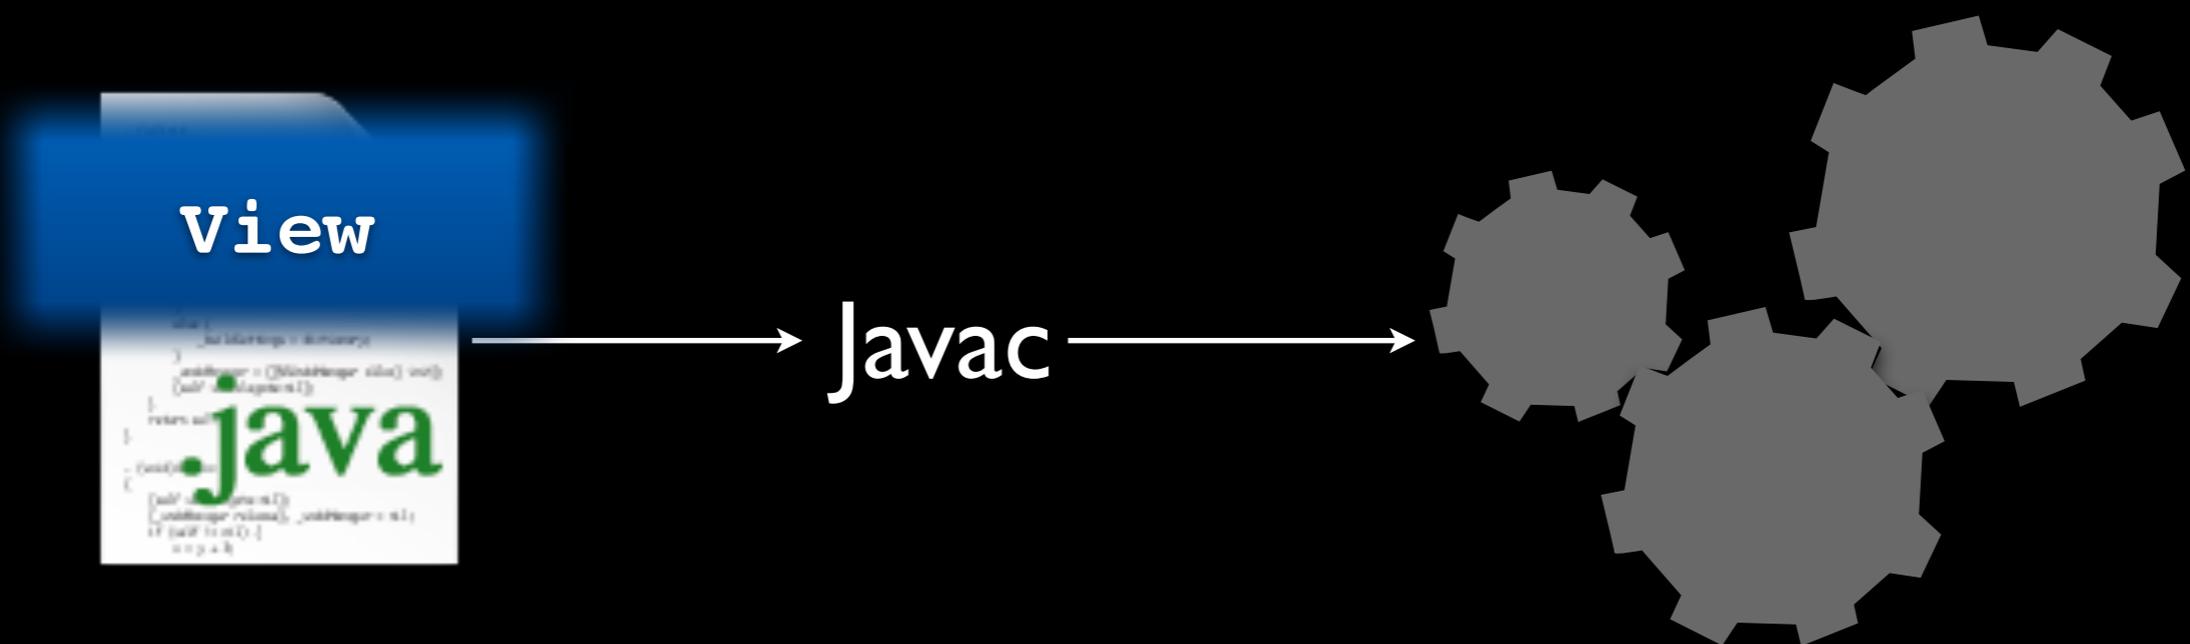


Monitoring Programs



Problem: People write monitoring programs by hand

Monitoring Programs



Problem: People write monitoring programs by hand

Monitoring Programs



View

Monitoring Programs



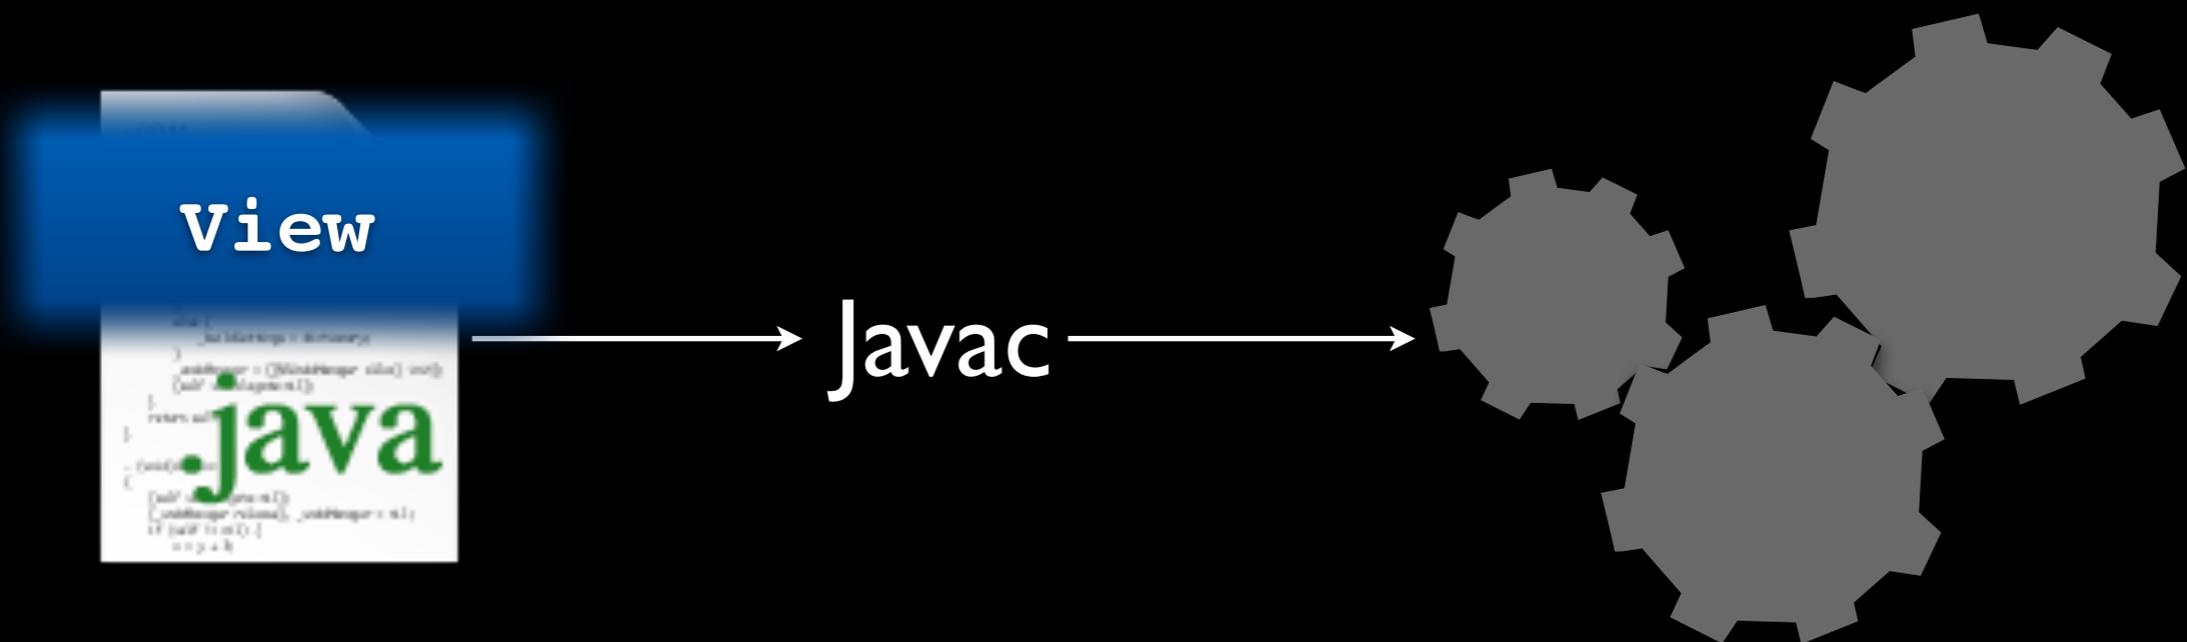
View

- An Aggregate Representation of the State of the World
- Maintained in Realtime as the State of the World Changes
- Needs to React to Changes In the World Quickly

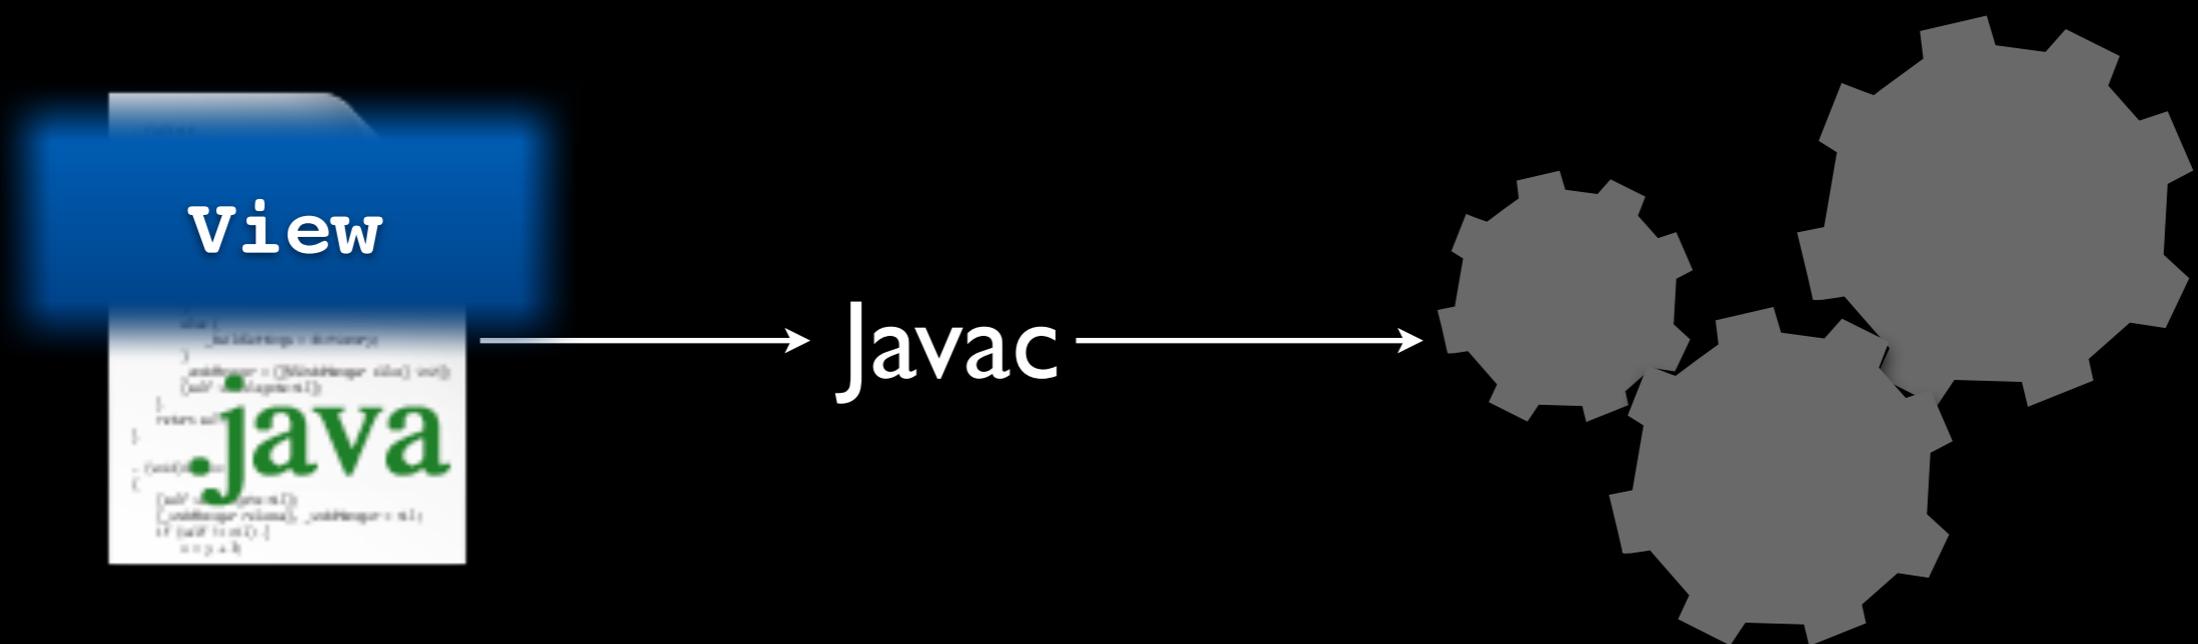
Not just Views

Frequently Fresh Views

Monitoring Programs

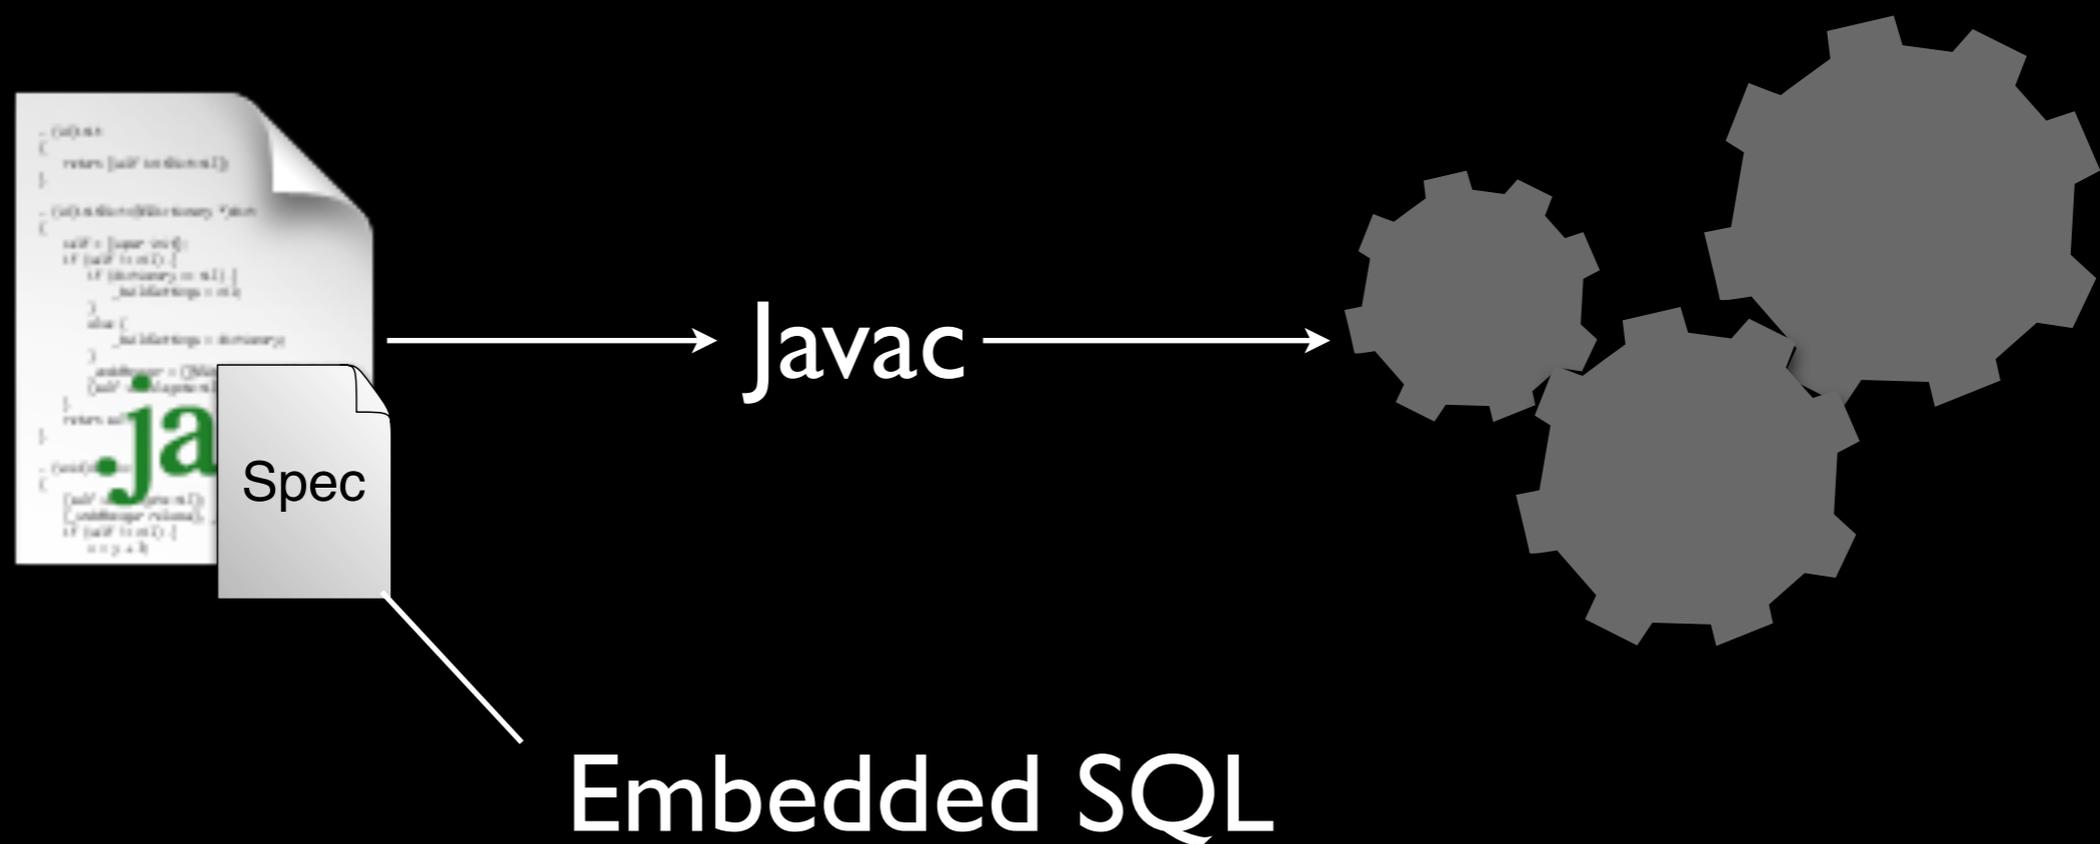


Monitoring Programs



(The Current State of the Art)

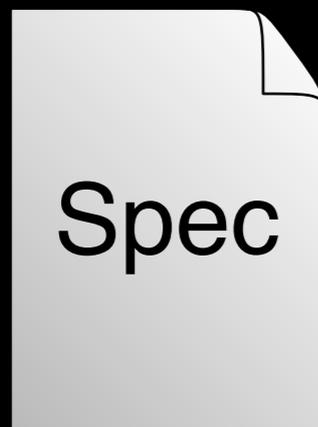
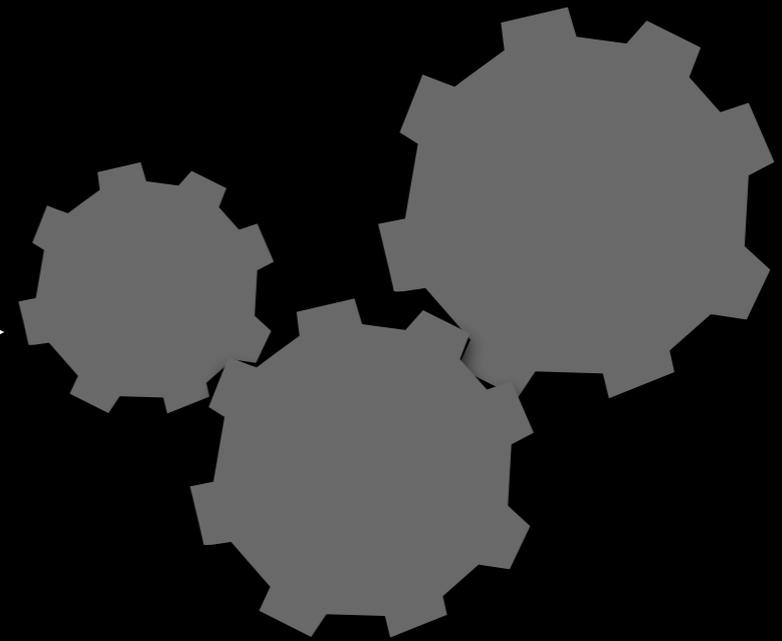
Monitoring Programs



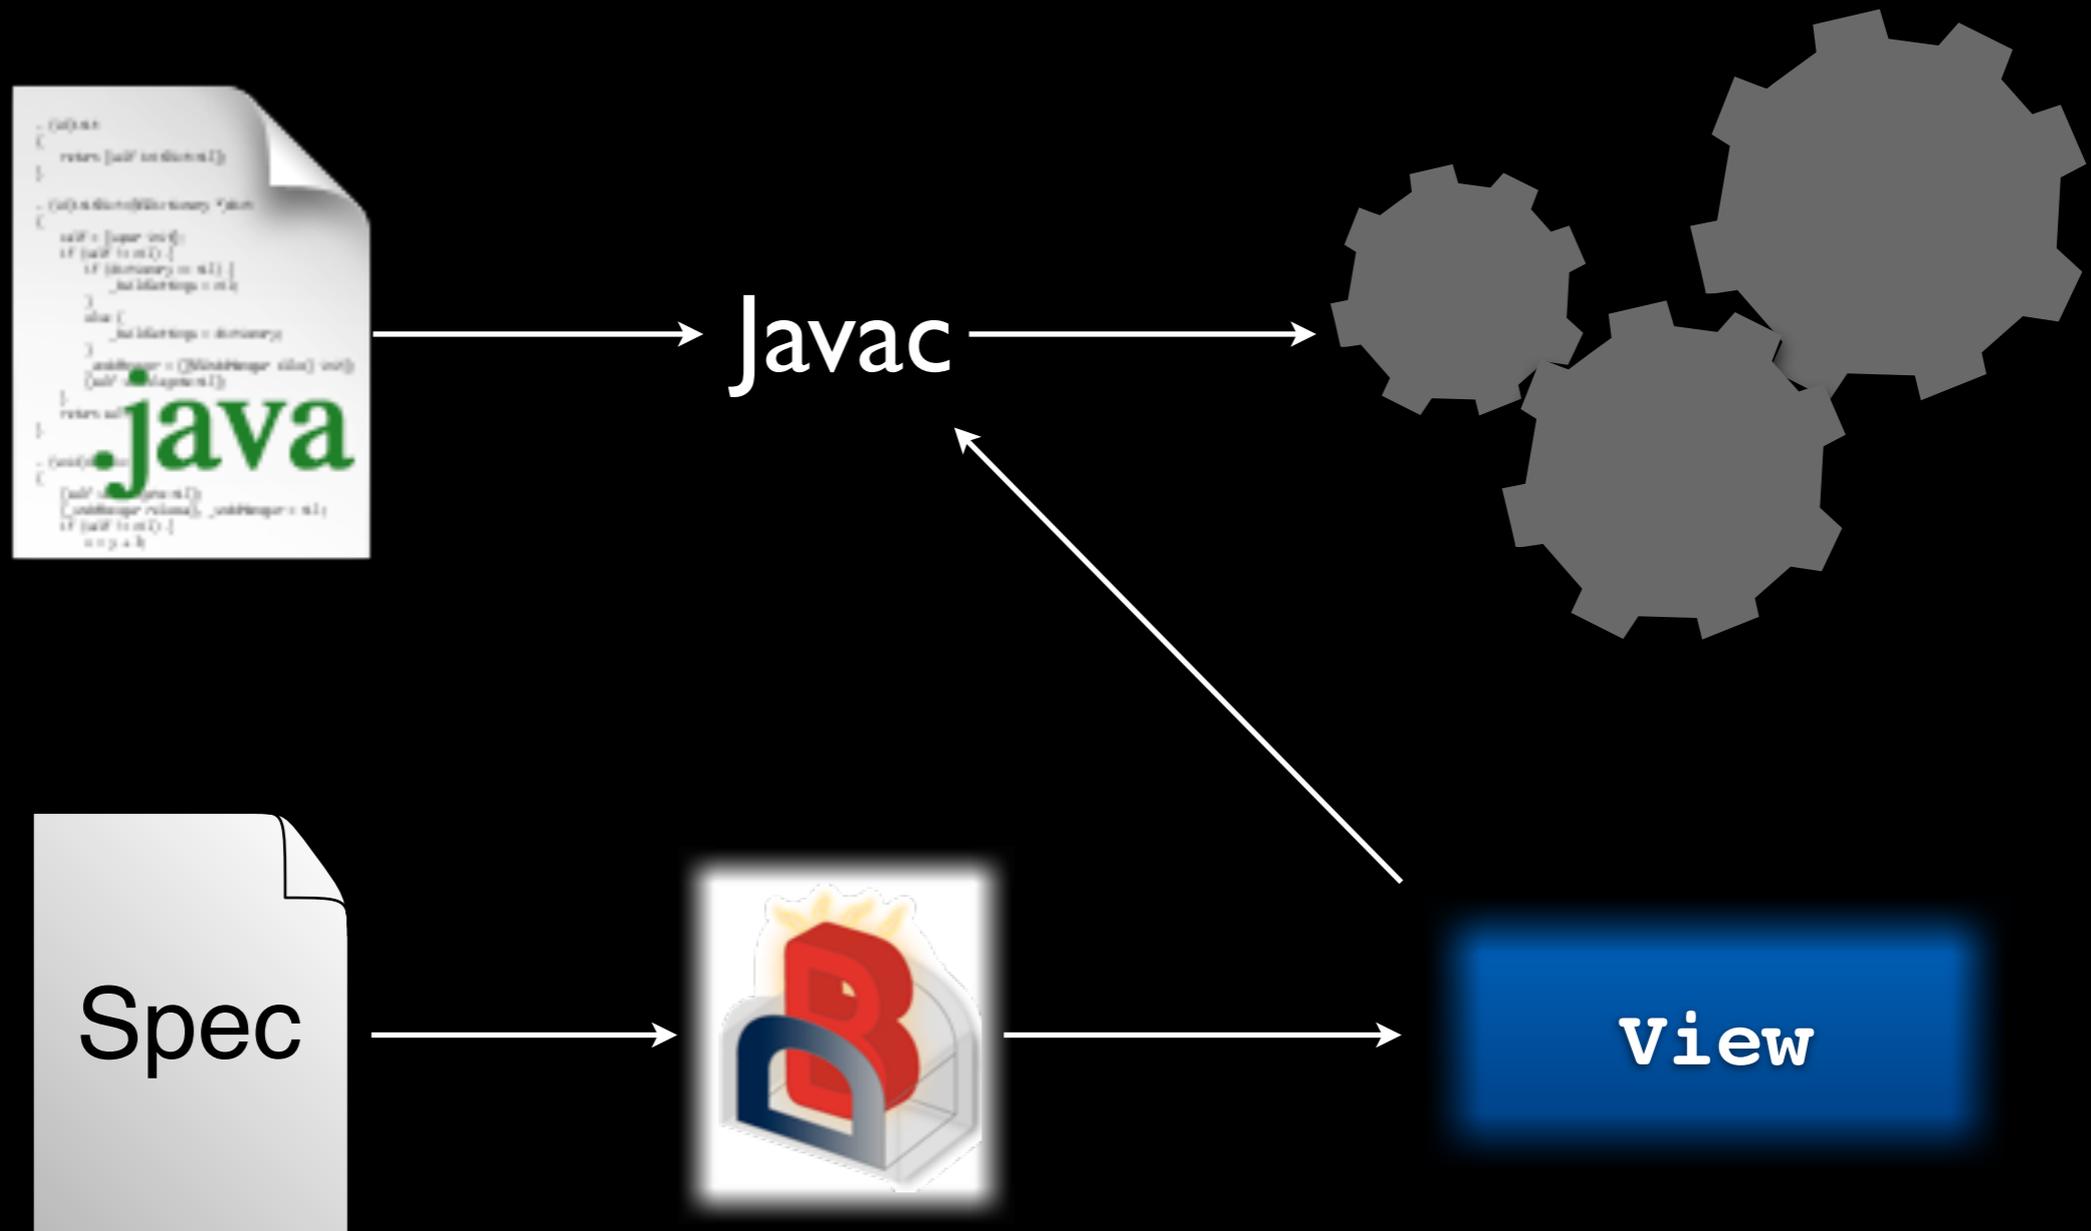
Monitoring Programs



Javac



Monitoring Programs



The DBToaster Compiler

The Viewlet Transform

The Viewlet Transform

Use Auxiliary Views to Speed Up View Maintenance

The Viewlet Transform

Use Auxiliary Views to Speed Up View Maintenance

The Delta of a Query Can Be Materialized!

The Viewlet Transform

```
SELECT SUM(R.A * S.C)
FROM   R, S
WHERE  R.B = S.B
```

A Simple 2-Way Join Aggregate

||

The Viewlet Transform

```
q[ ] := SELECT SUM(R.A * S.C)
        FROM   R, S
        WHERE  R.B = S.B
```

A Simple 2-Way Join Aggregate

||

The Viewlet Transform

ON $+R(\partial A, \partial B)$:

```
q[ ] += SELECT SUM(  $\partial A$  * S.C )  
        FROM S  
        WHERE  $\partial B = S.B$ 
```

Materialize and Incrementally Maintain The Query

The Viewlet Transform

ON $+R(\partial A, \partial B)$:

```
q[ ] += SELECT SUM(  $\partial A$  * S.C )  
        FROM S  
        WHERE  $\partial B = S.B$ 
```

The Delta of the Original Query

Materialize and Incrementally Maintain The Query

The Viewlet Transform

ON $+R(\partial A, \partial B)$:

$$q[] += \partial A * \left(\begin{array}{l} \text{SELECT SUM}(S.C) \\ \text{FROM S} \\ \text{WHERE } \partial B = S.B \end{array} \right)$$

Optimize

The Viewlet Transform

ON $+R(\partial A, \partial B)$:

$$q[] += \partial A * \left(\begin{array}{l} \text{SELECT SUM}(S.C) \\ \text{FROM } S \\ \text{WHERE } \partial B = S.B \end{array} \right)$$

Optimize

The Viewlet Transform

ON $+R(\partial A, \partial B)$:

$$q[] += \partial A * \left(\begin{array}{l} \text{SELECT } S.B, \text{SUM}(S.C) \\ \text{FROM } S \\ \text{GROUP BY } S.B \end{array} \right) [\partial B]$$

Optimize

The Viewlet Transform

ON $+R(\partial A, \partial B)$:

$q[] += \partial A * \left(\begin{array}{l} \text{SELECT } S.B, \text{SUM}(S.C) \\ \text{FROM } S \\ \text{GROUP BY } S.B \end{array} \right) [\partial B]$

Just an Ordinary Query



Optimize

The Viewlet Transform

ON $\Delta R(\Delta A, \Delta B)$:

$q[] \Delta = \Delta A * mR[\Delta B]$

```
mR[B] := SELECT S.B, SUM(S.C)
          FROM S
          GROUP BY S.B
```

Extract and Materialize The Delta View

The Viewlet Transform

ON $+R(\partial A, \partial B)$:

$q[] += \partial A * mR[\partial B]$

↙ A Hash Map (indexed by S.B)

```
mR[B] := SELECT S.B, SUM(S.C)
        FROM S
        GROUP BY S.B
```

Extract and Materialize The Delta View

The Viewlet Transform

ON +R(∂A , ∂B):

$q[] += \partial A * mR[\partial B]$

ON +S(∂B , ∂C):

$mR[B] += \text{SELECT } \partial B, \text{SUM}(\partial C)$

Incrementally Maintain The Delta View

The Viewlet Transform

ON +R(∂A , ∂B):

$$q[] += \partial A * mR[\partial B]$$

ON +S(∂B , ∂C):

$$mR[\partial B] += \partial C$$

Optimize

The Viewlet Transform

ON +R(∂A , ∂B):

$$q[] += \partial A * mR[\partial B]$$

$$mS[\partial B] += \partial A$$

ON +S(∂B , ∂C):

$$mR[\partial B] += \partial C$$

$$q[] += \partial C * mS[\partial B]$$

Repeat for the Other Deltas of the Query

The Viewlet Transform

The Viewlet Transform

- Take the Deltas
 - Optimize and Materialize Them
 - Take the Deltas
 - Optimize and Materialize Them
 - ...

The Viewlet Transform

- Take the Deltas
 - Optimize and Materialize Them
 - Take the Deltas
 - Optimize and Materialize Them
 - ...

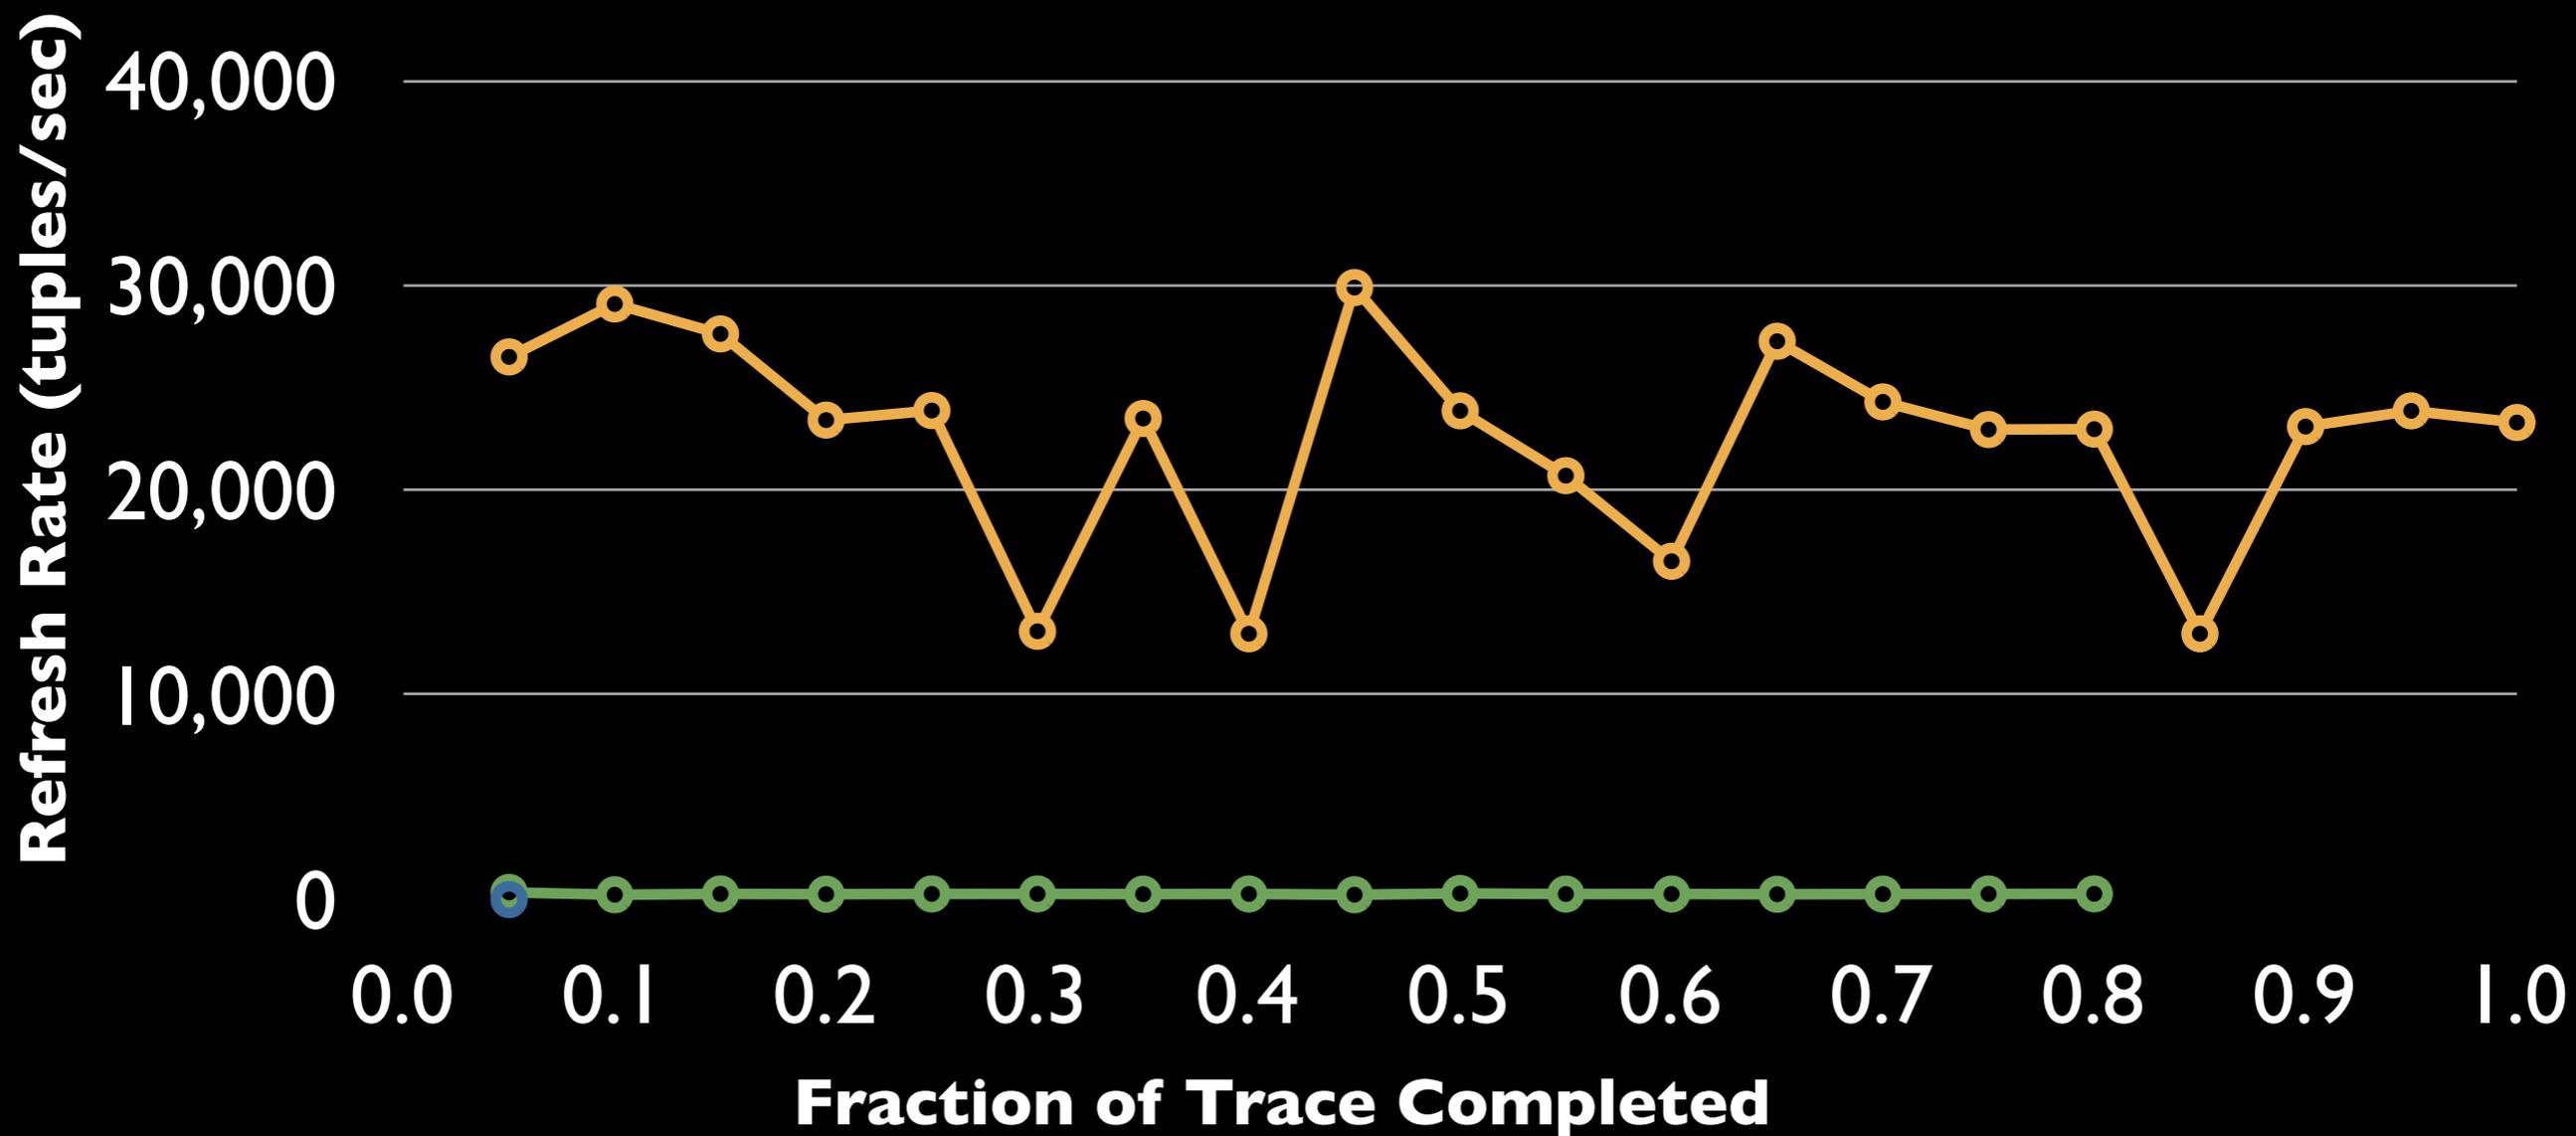
Performance

- TPC-H Workload
 - Simulated Realtime Data Warehouse
 - Update Stream Derived from TPC-H Gen
- Financial Benchmark
 - 24 hr Trace for an Actively Traded Stock.

TPCH: Q3

Refresh Rate

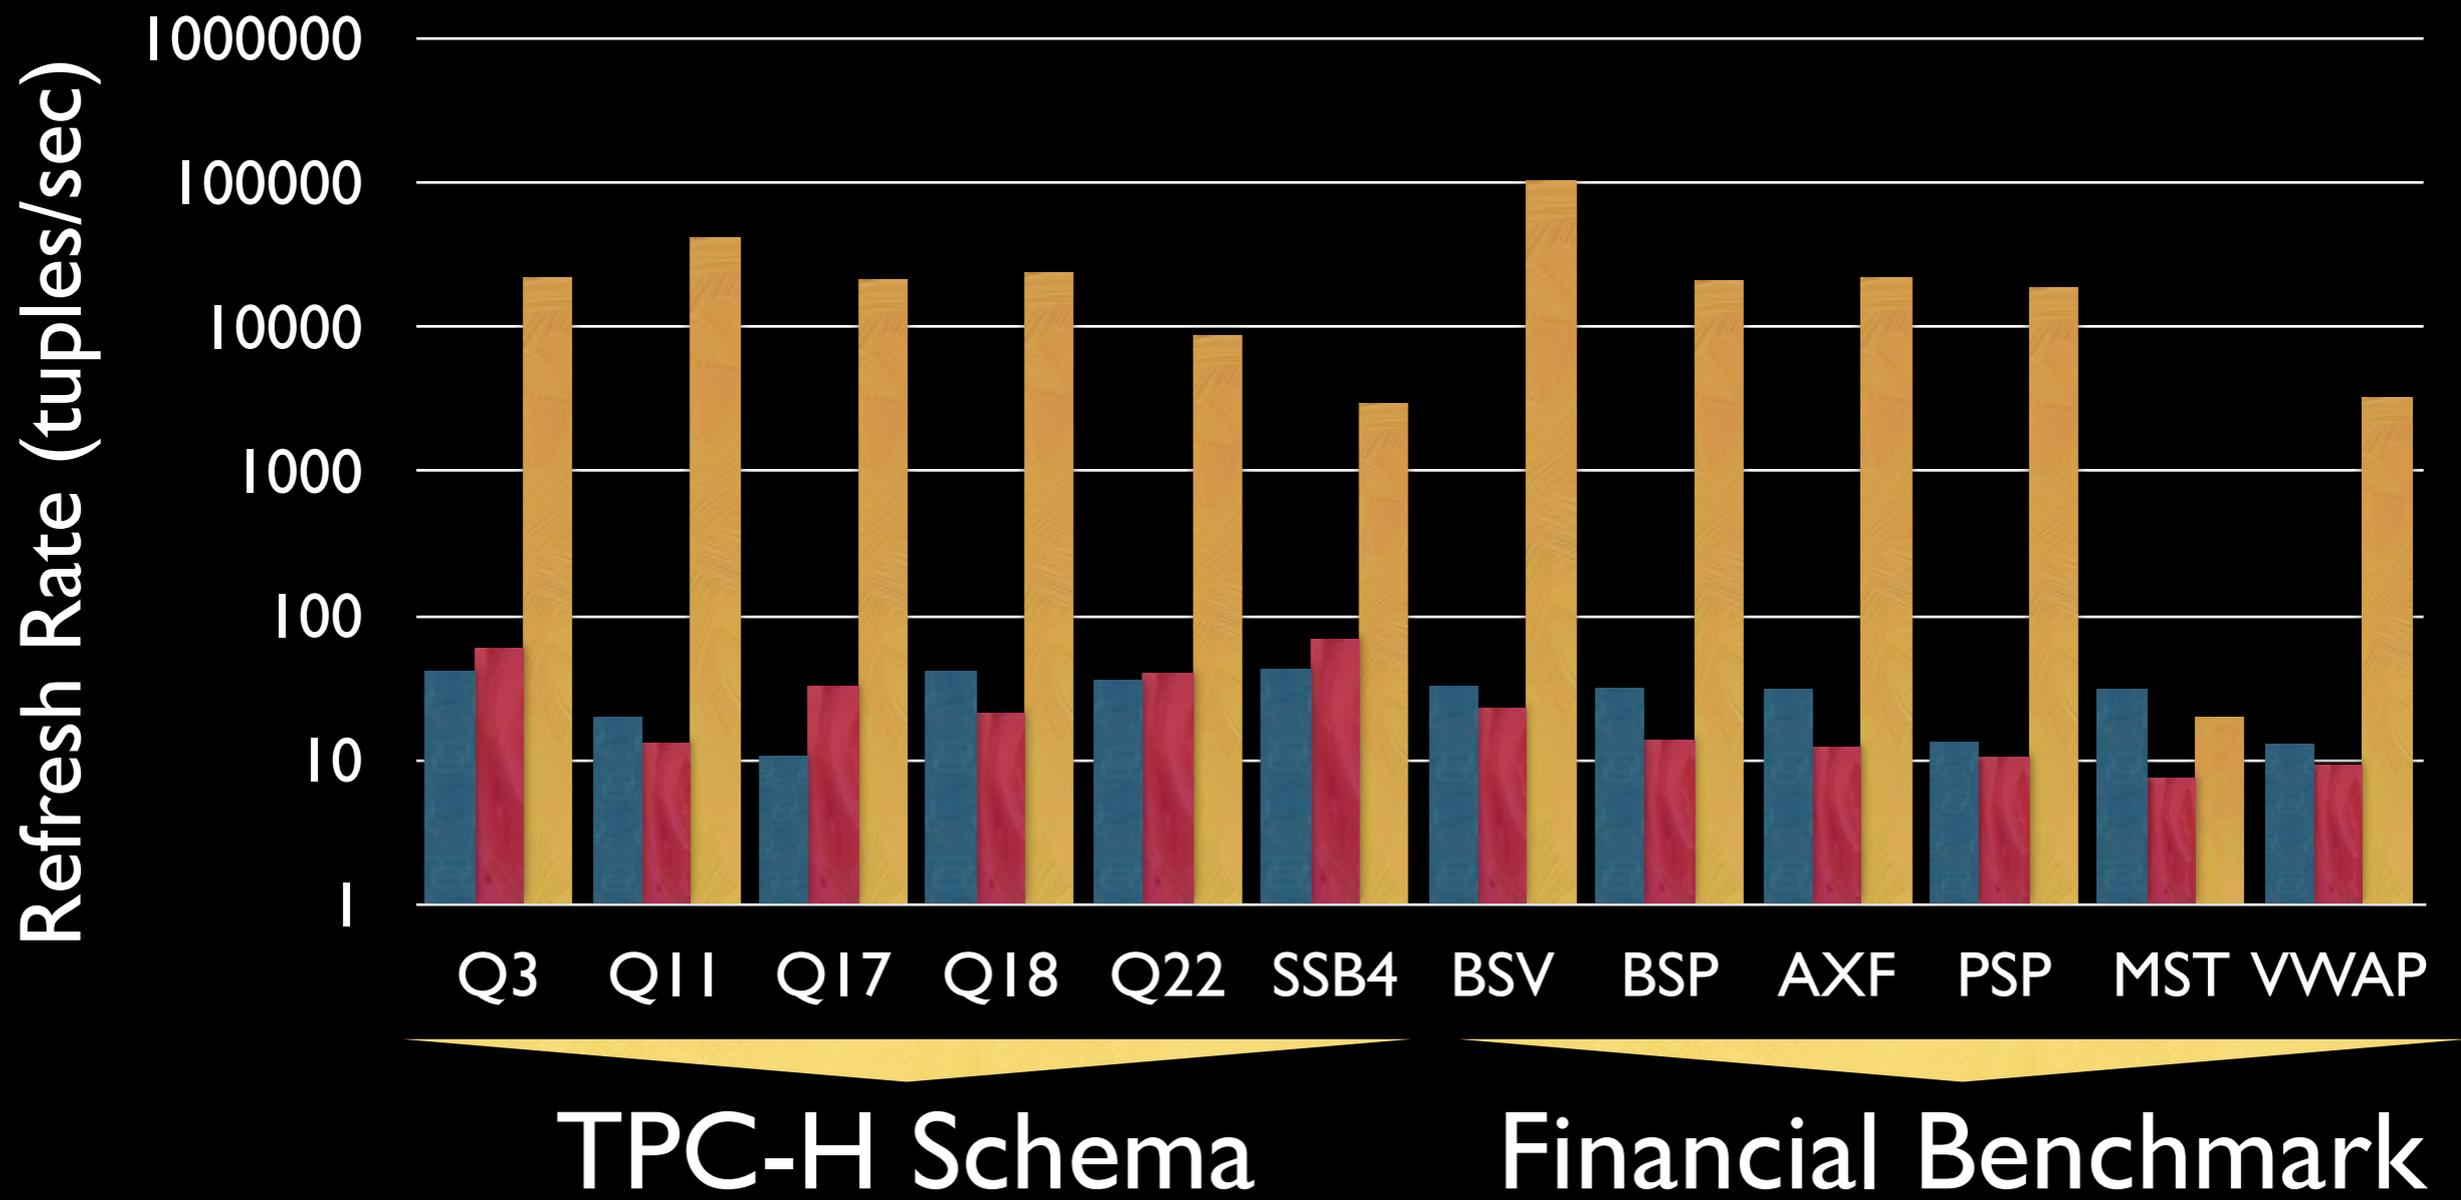
(3-Way Join)



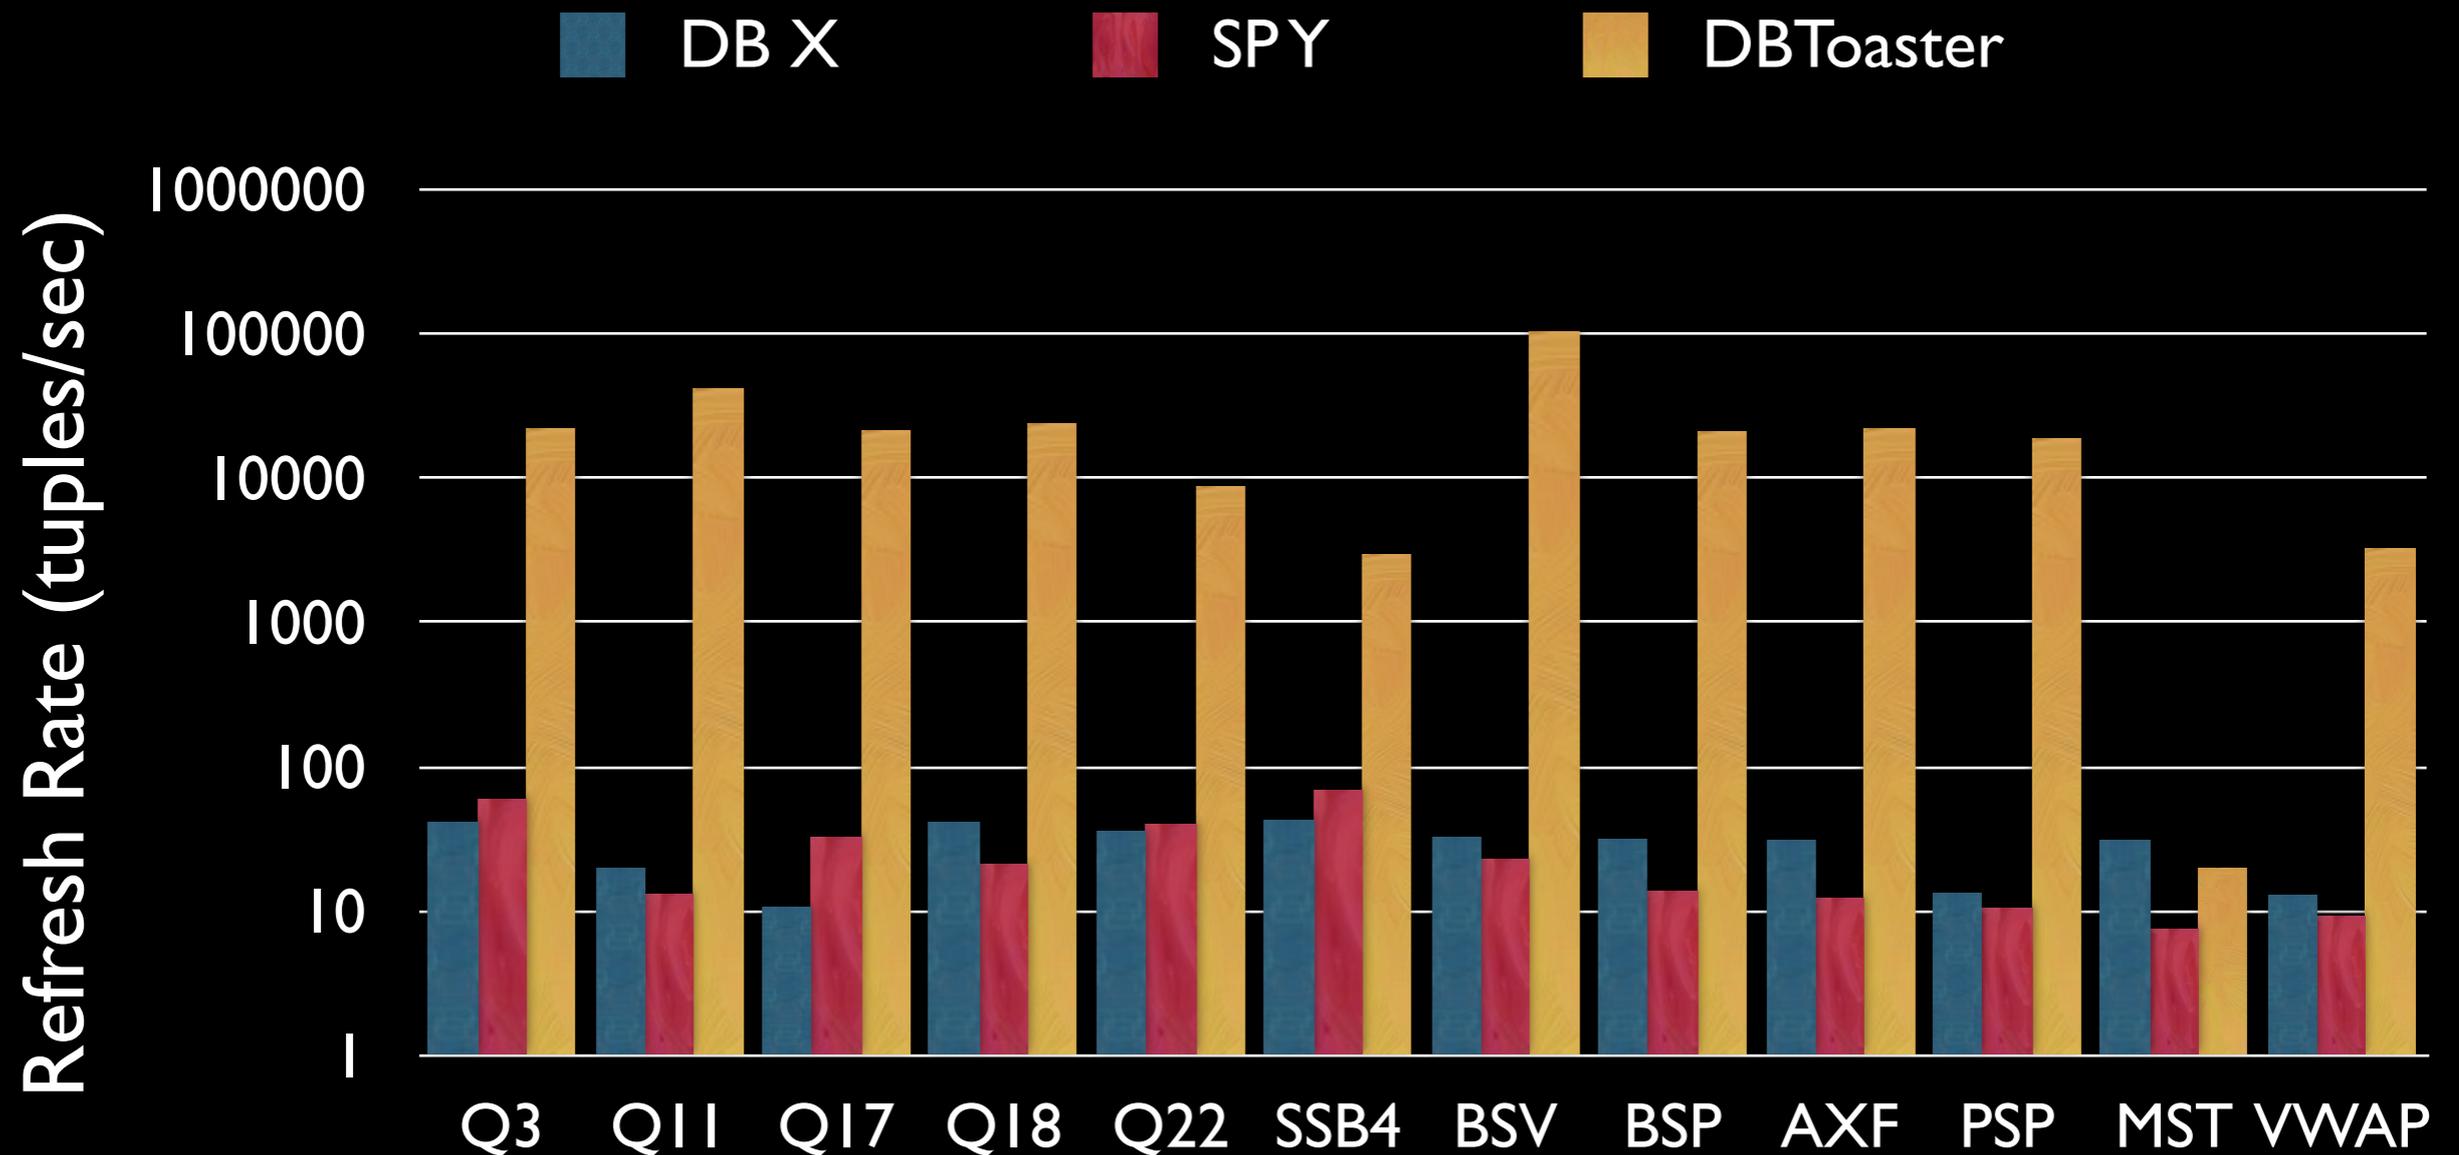
Naive Re-evaluation Traditional IVM DBToaster

DBToaster vs Commercial Engines

DB X SPY DBToaster



DBToaster vs Commercial Engines



DBToaster is consistently 3 OOM better!

Limitations of Commercial Systems

- OLTP IVM is not designed for aggregating Low-Latency/Single-Tuple Updates.
- OLTP IVM doesn't support our full query workload.
- Stream Processors are not designed for rapidly changing long-lived data.

Limitations of Commercial Systems

- OLTP IVM is not designed for aggregating Low-Latency/Single-Tuple Updates.
- OLTP IVM doesn't support our full query workload.
- Stream Processors are not designed for rapidly changing long-lived data.

DBToaster opens entirely new application domains!

Conclusions

- The Viewlet Transform generates auxiliary views that make incremental maintenance fast.
- Materializing only part of an auxiliary view can sometimes be faster.
- DBToaster is commonly 3 OoM faster than Commercial Systems.

Download Now: <http://www.dbtoaster.org>

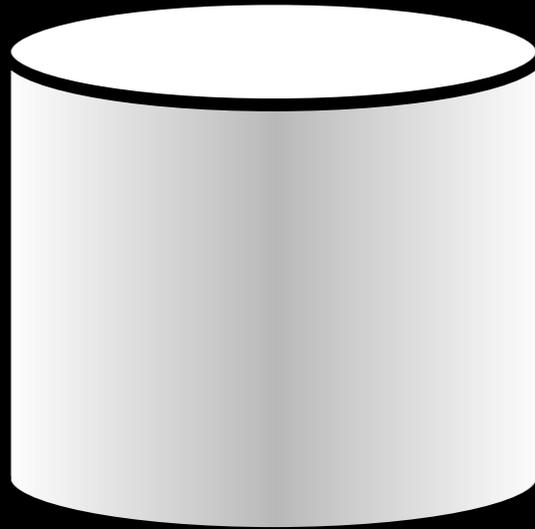
Cumulus: Hybrid Consistency Aggregation Queries (in the cloud)

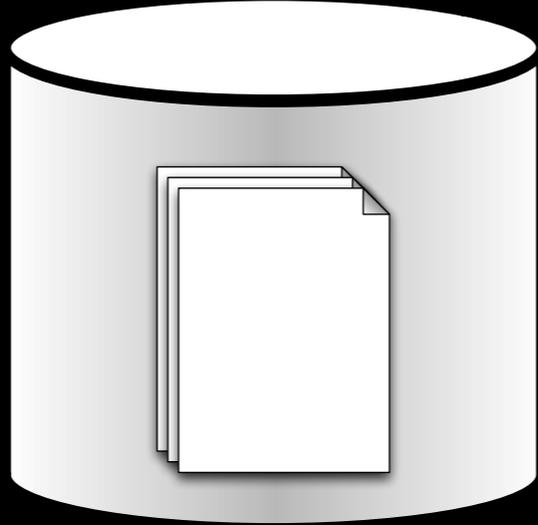


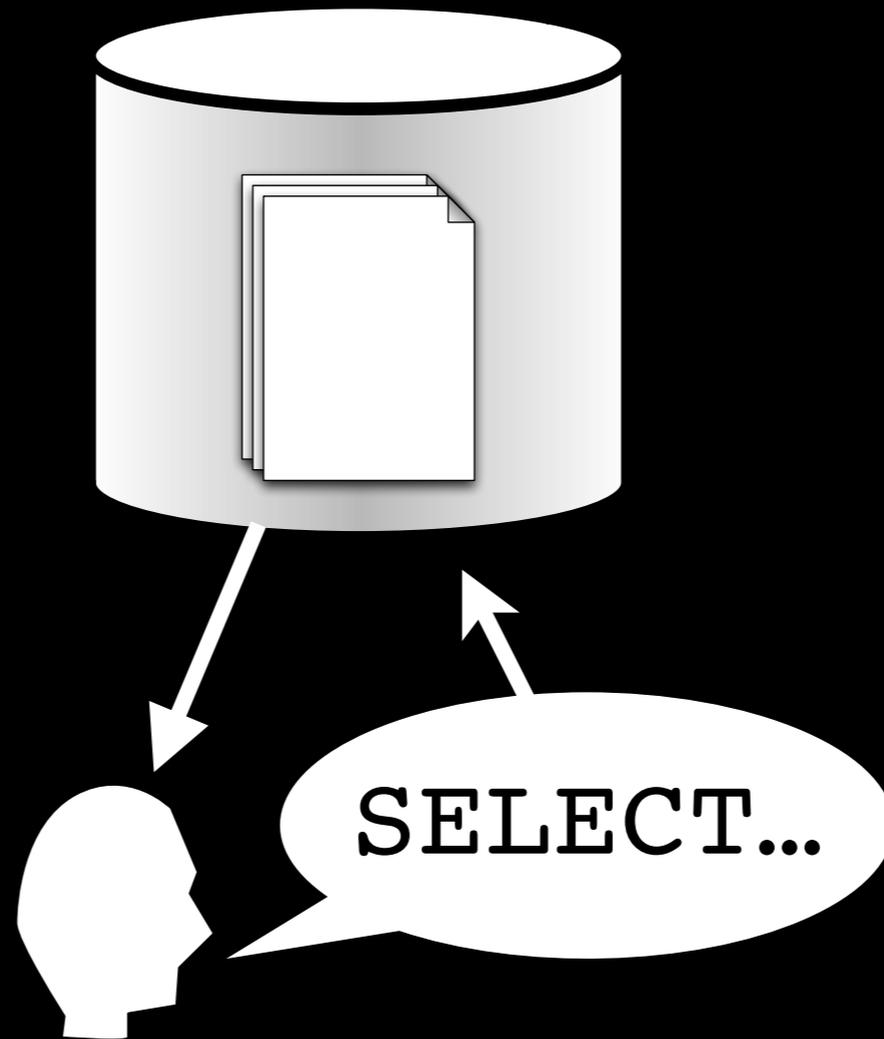
Oliver Kennedy

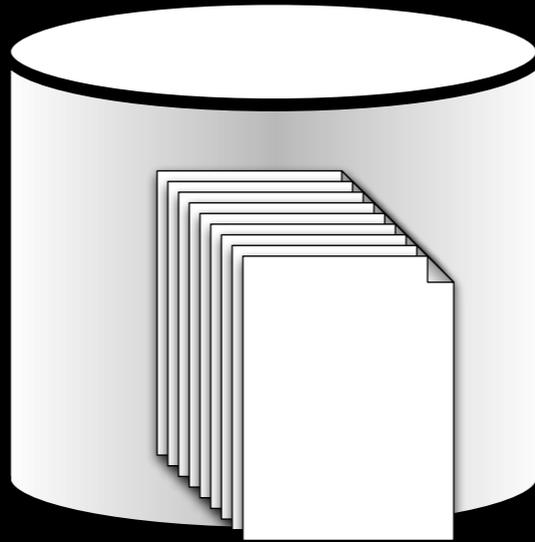


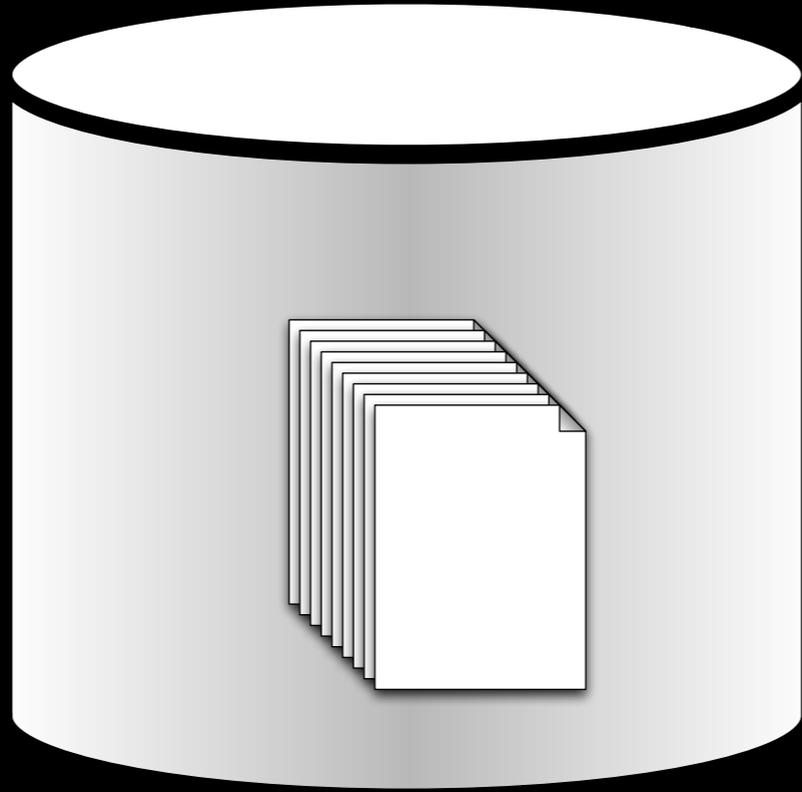
(joint work with JHU's Yanif Ahmad and Yotam Barnoy)

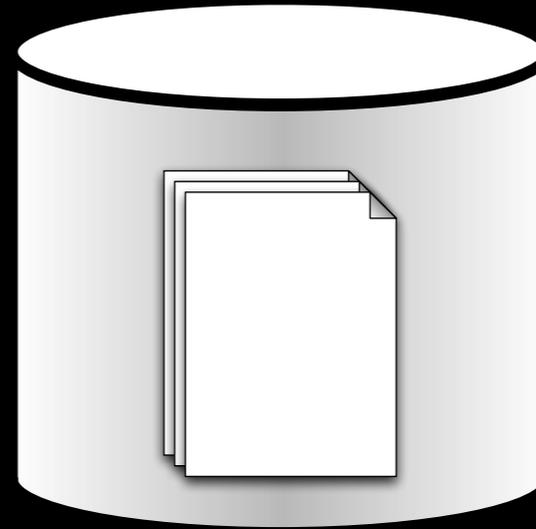
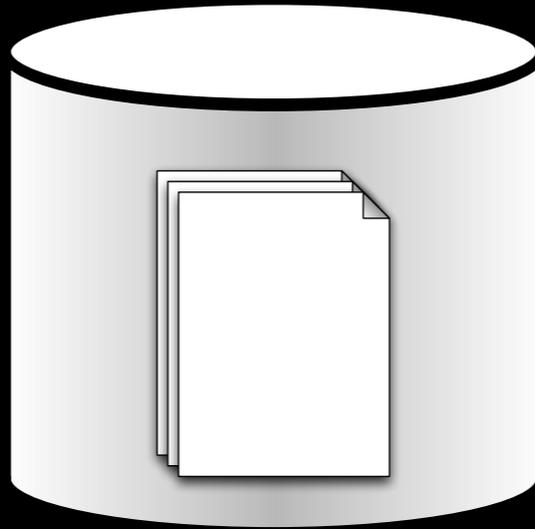
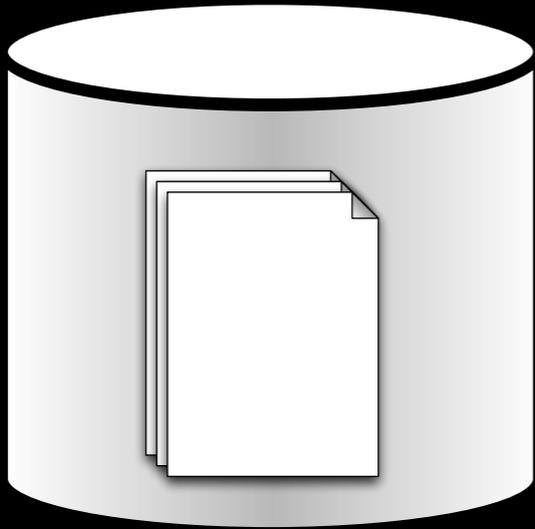


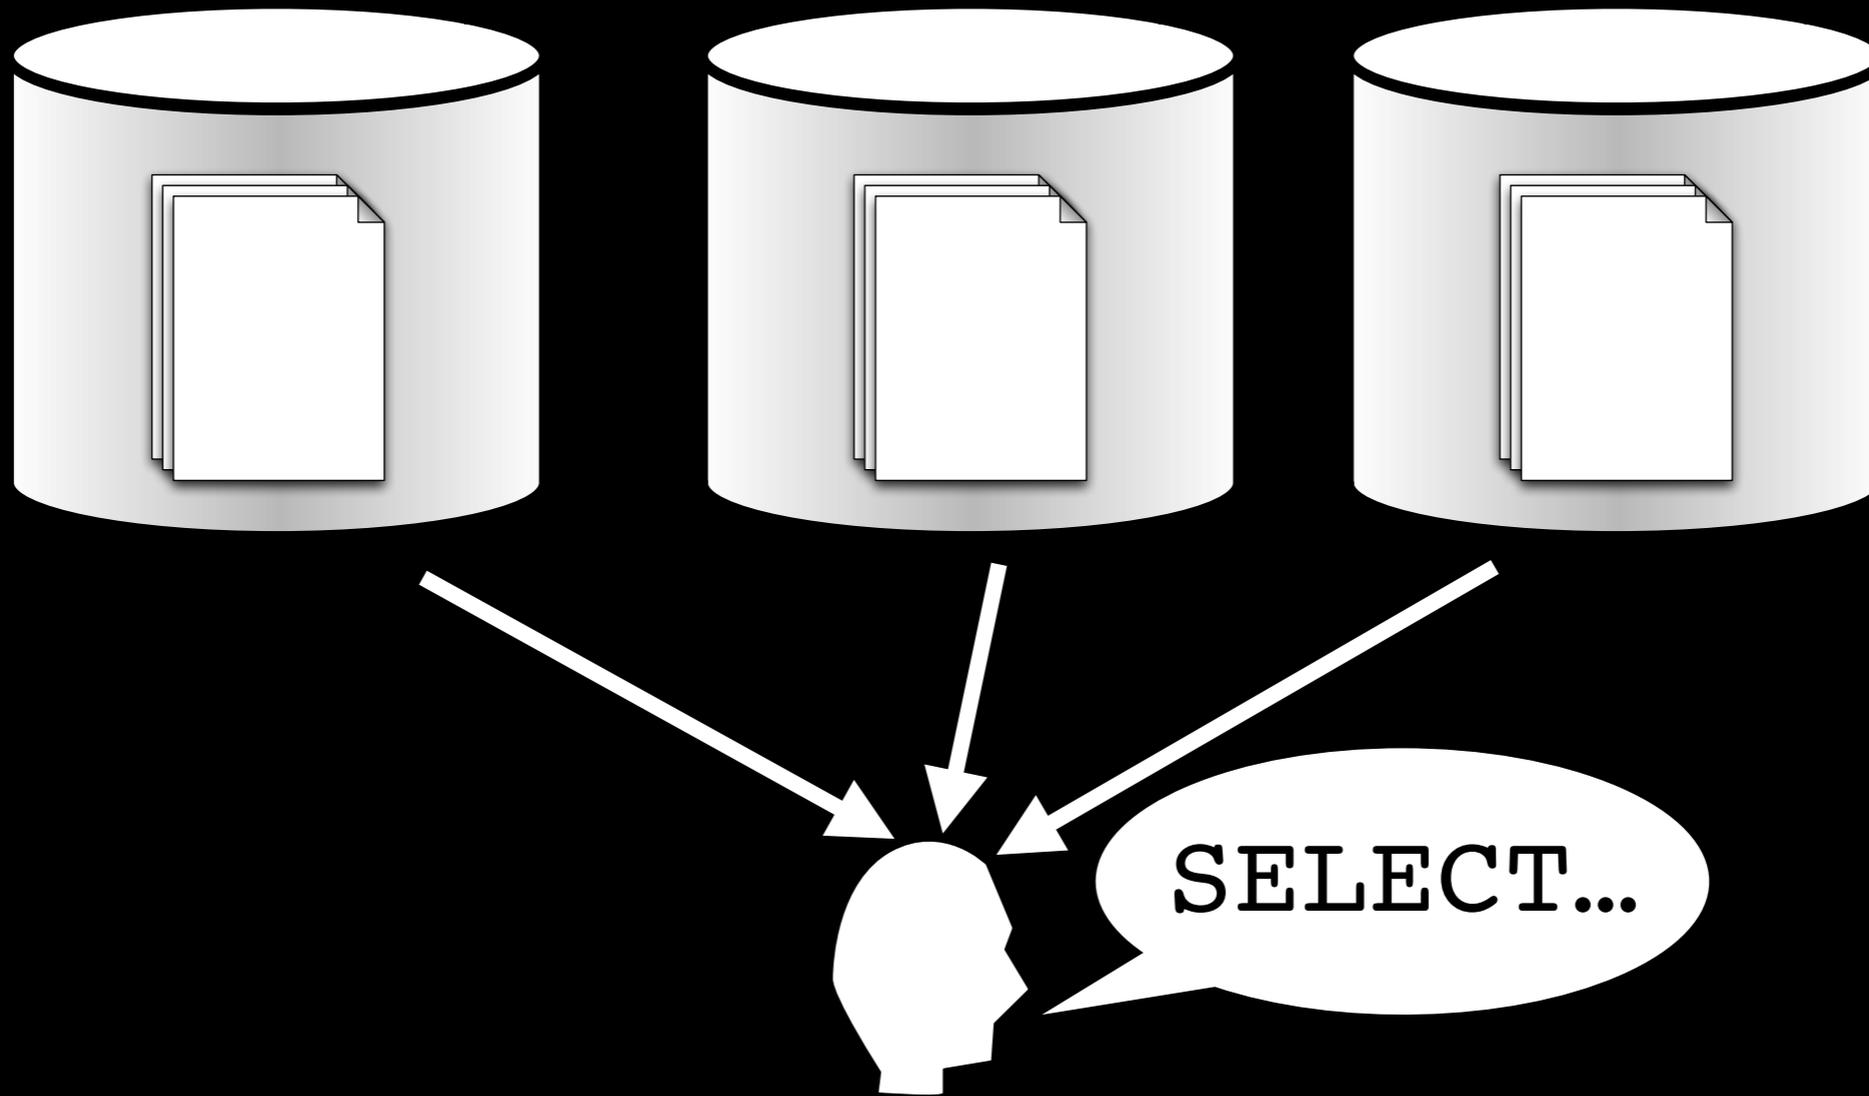


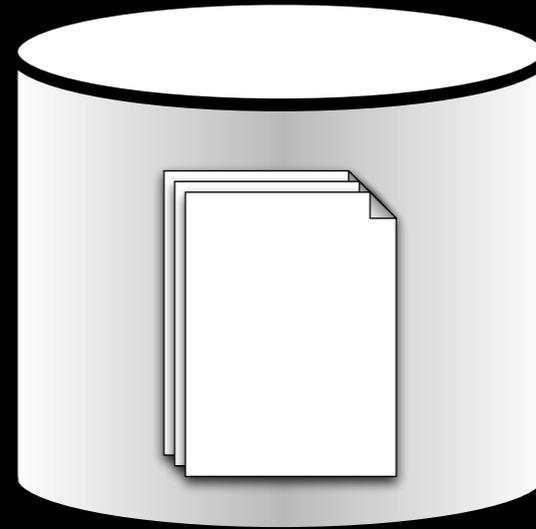
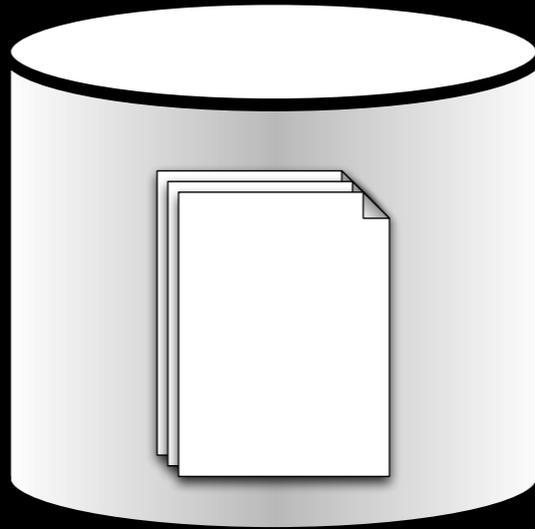
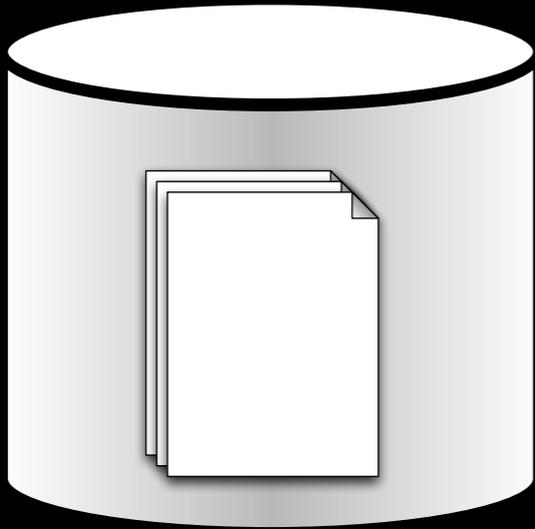


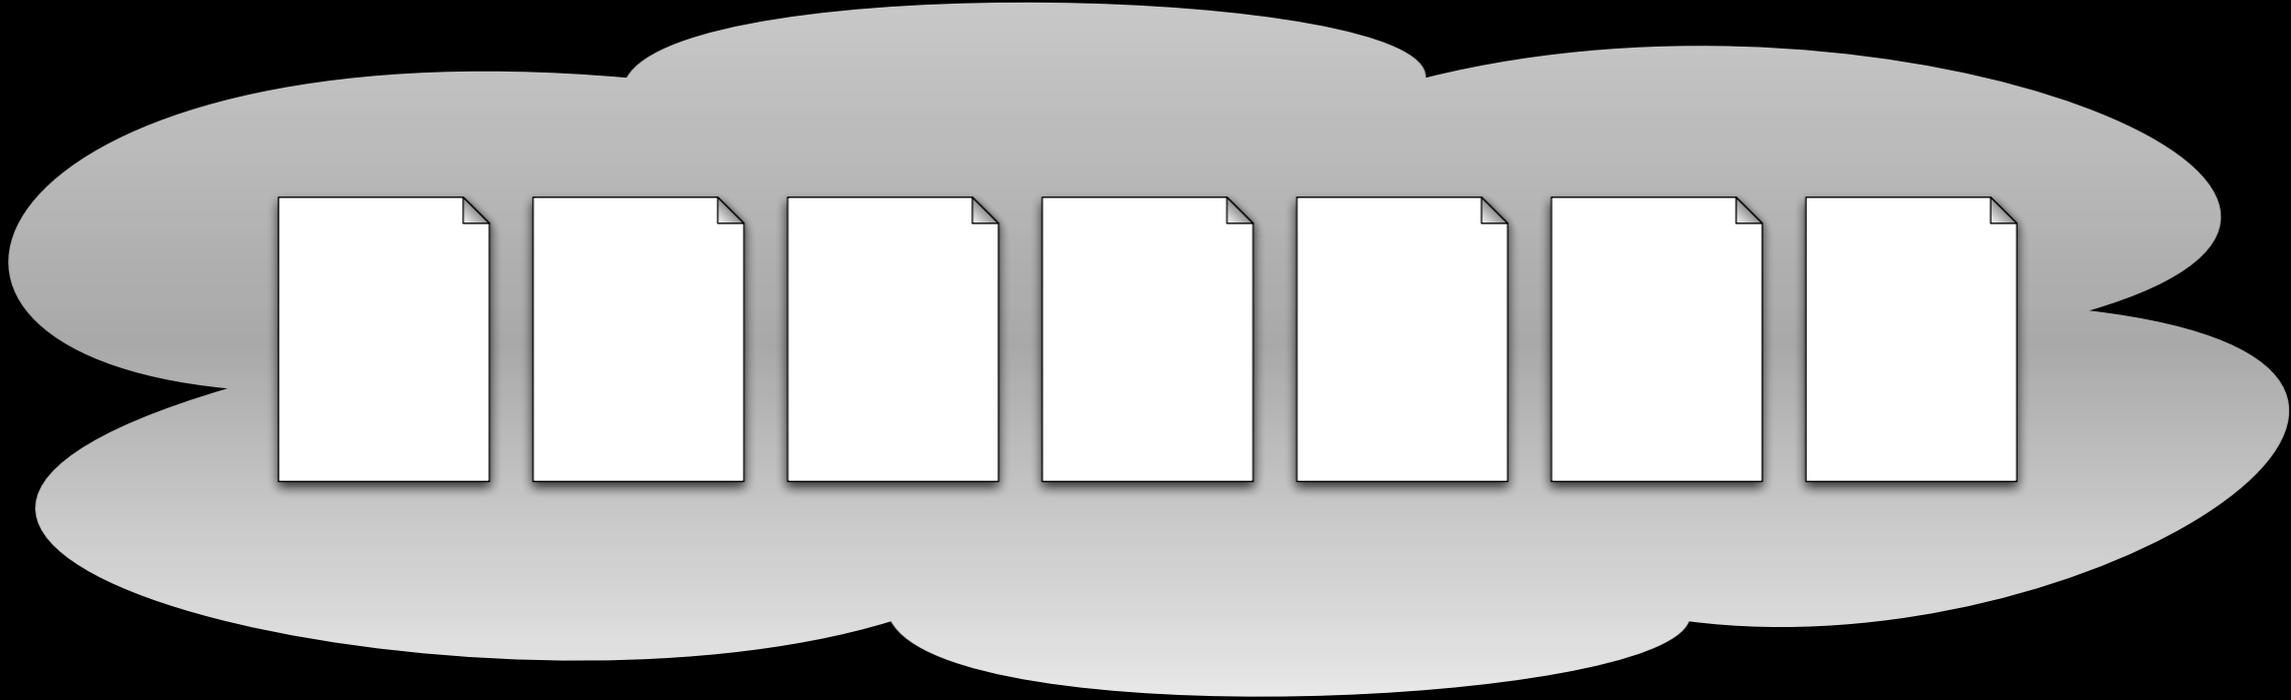


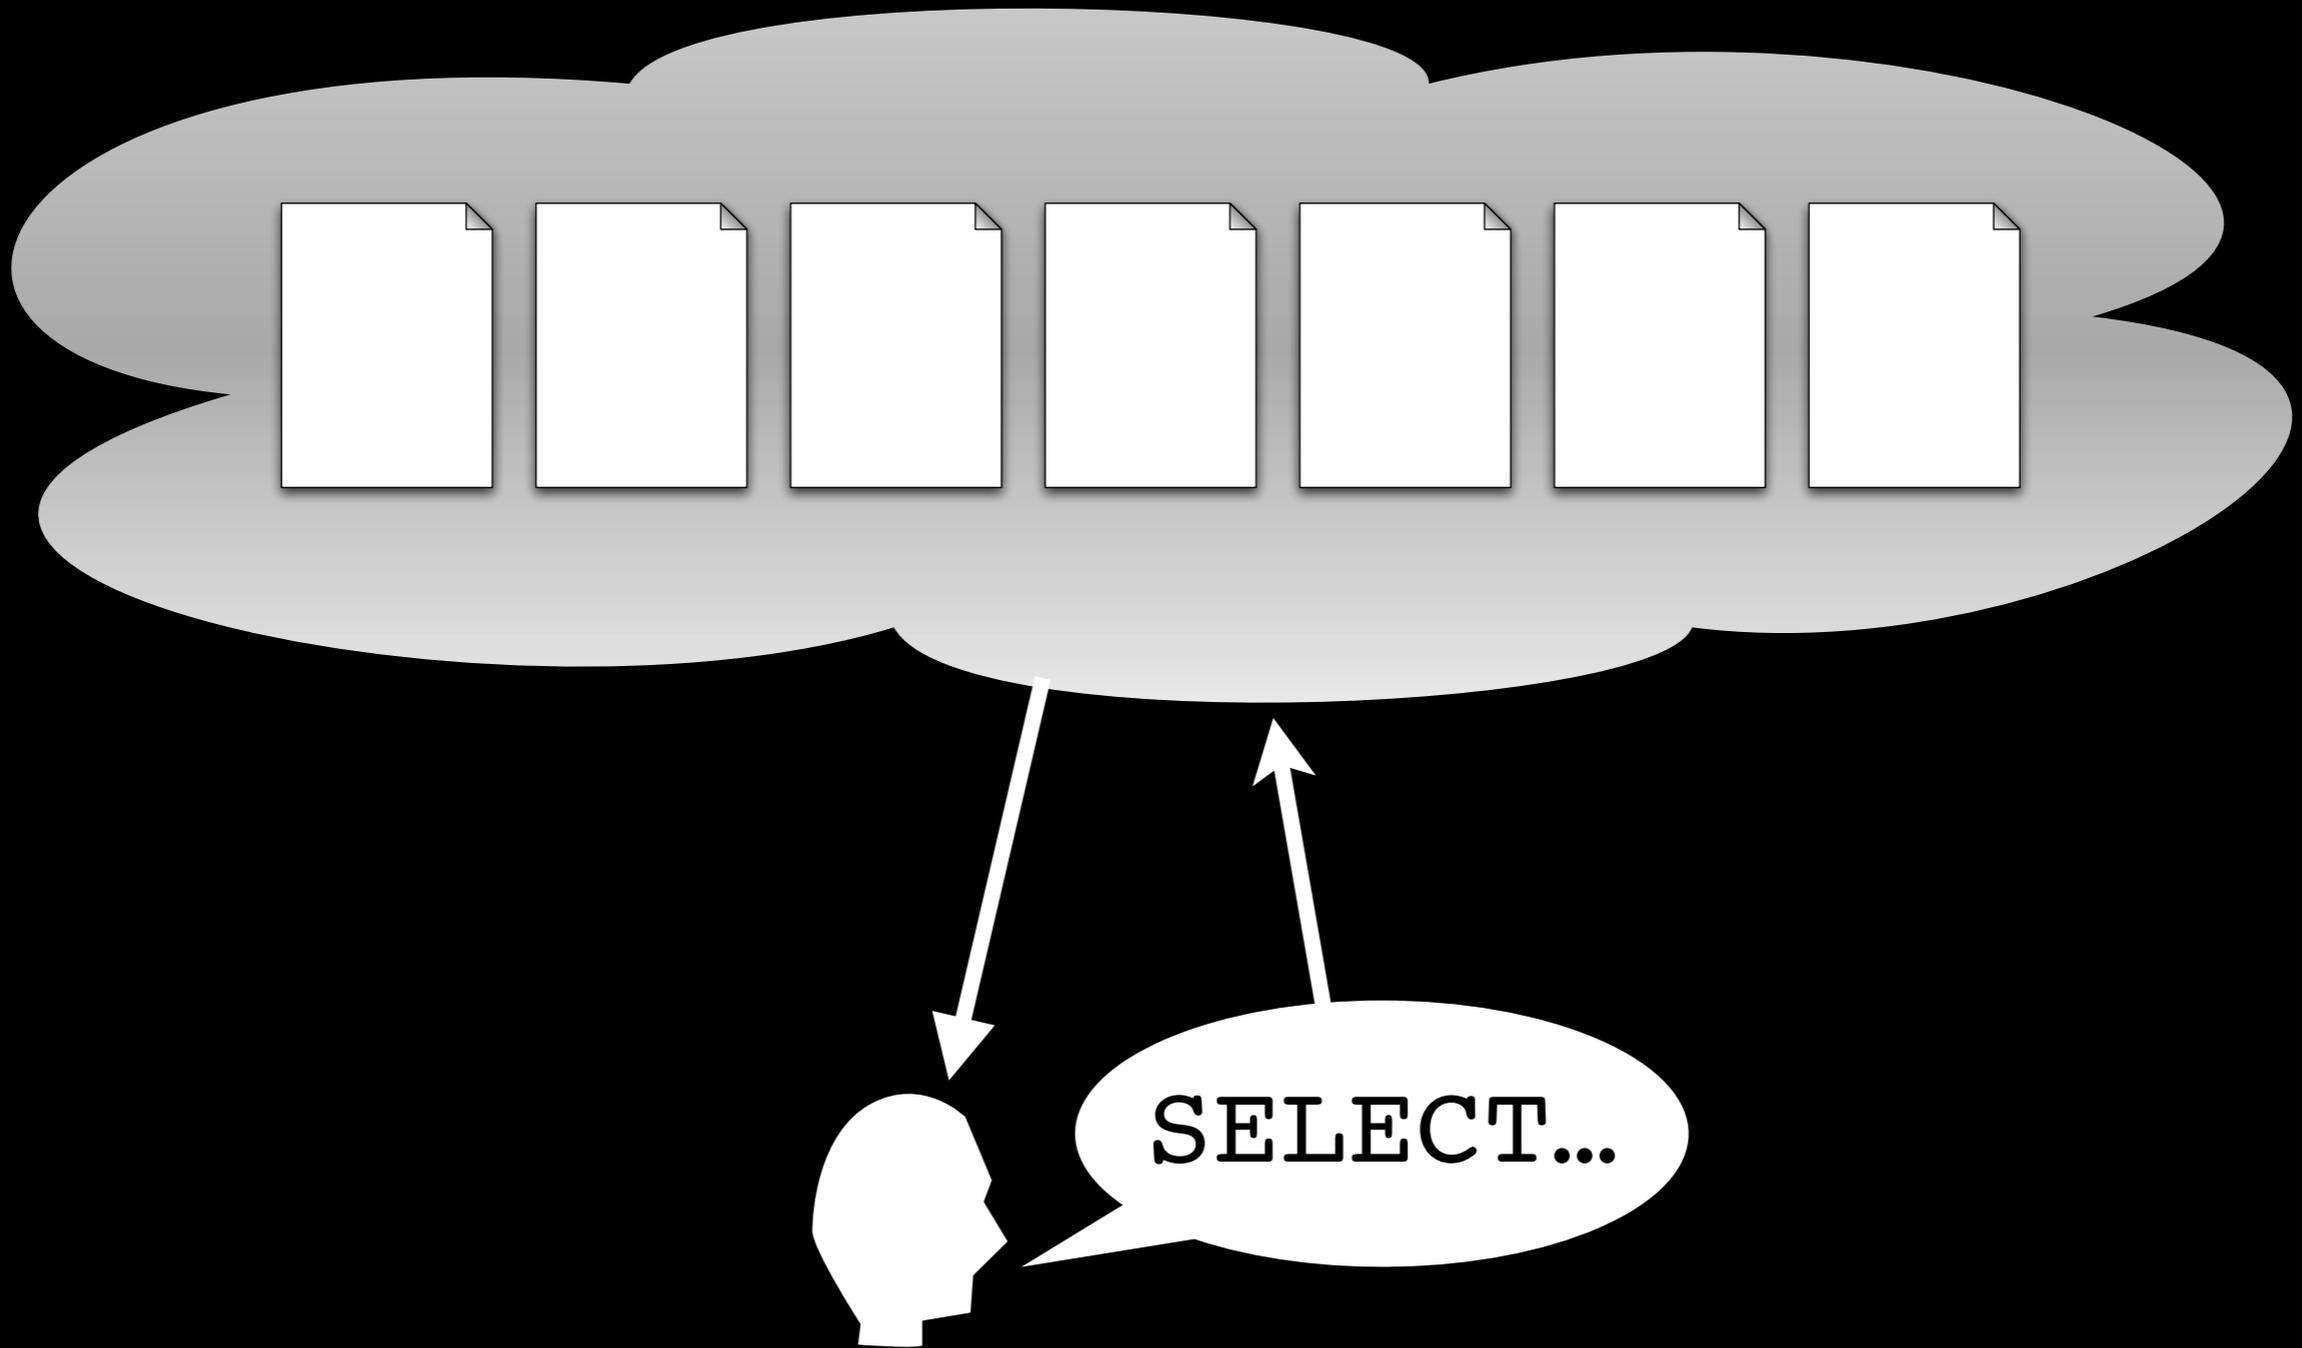


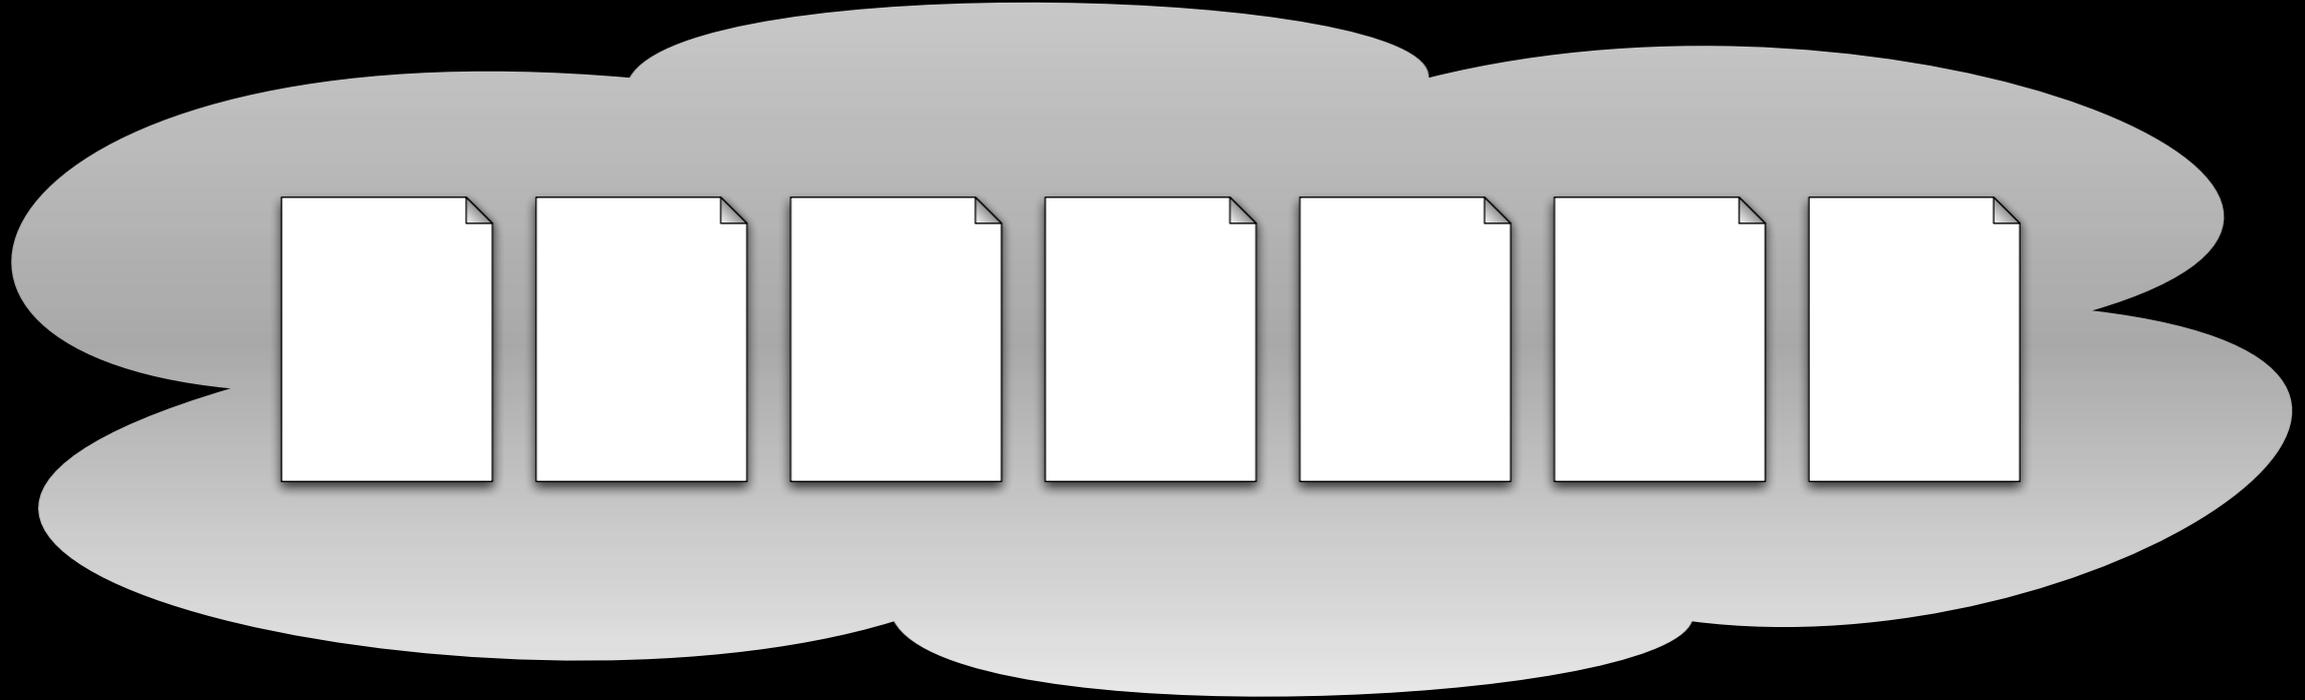


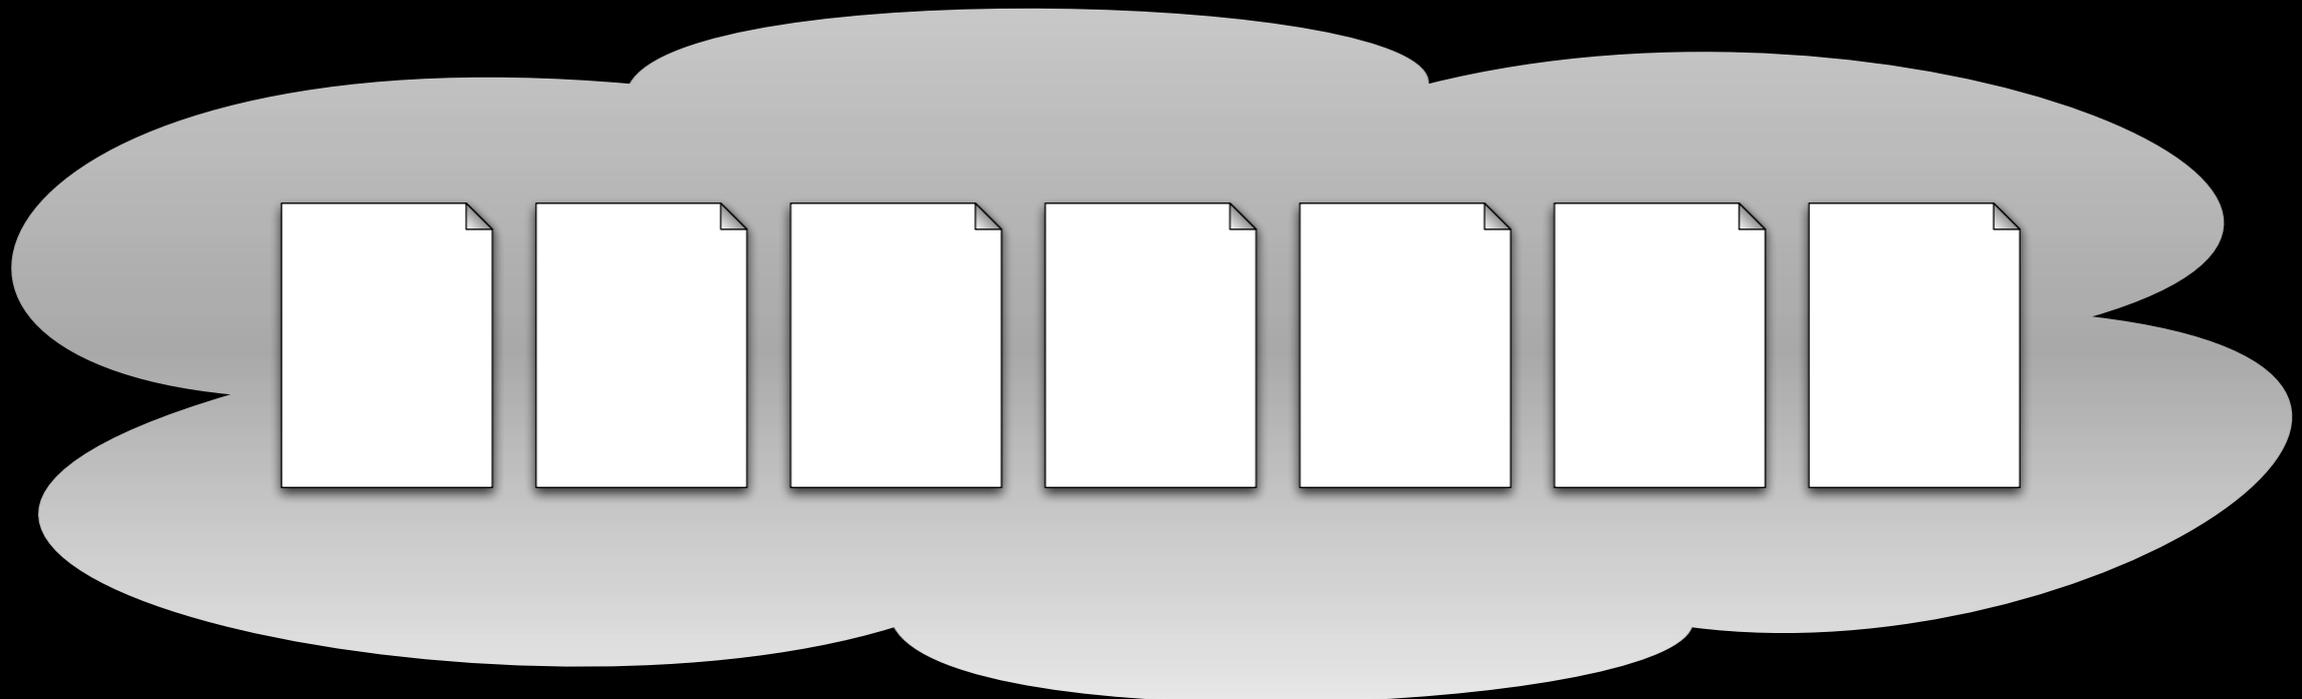






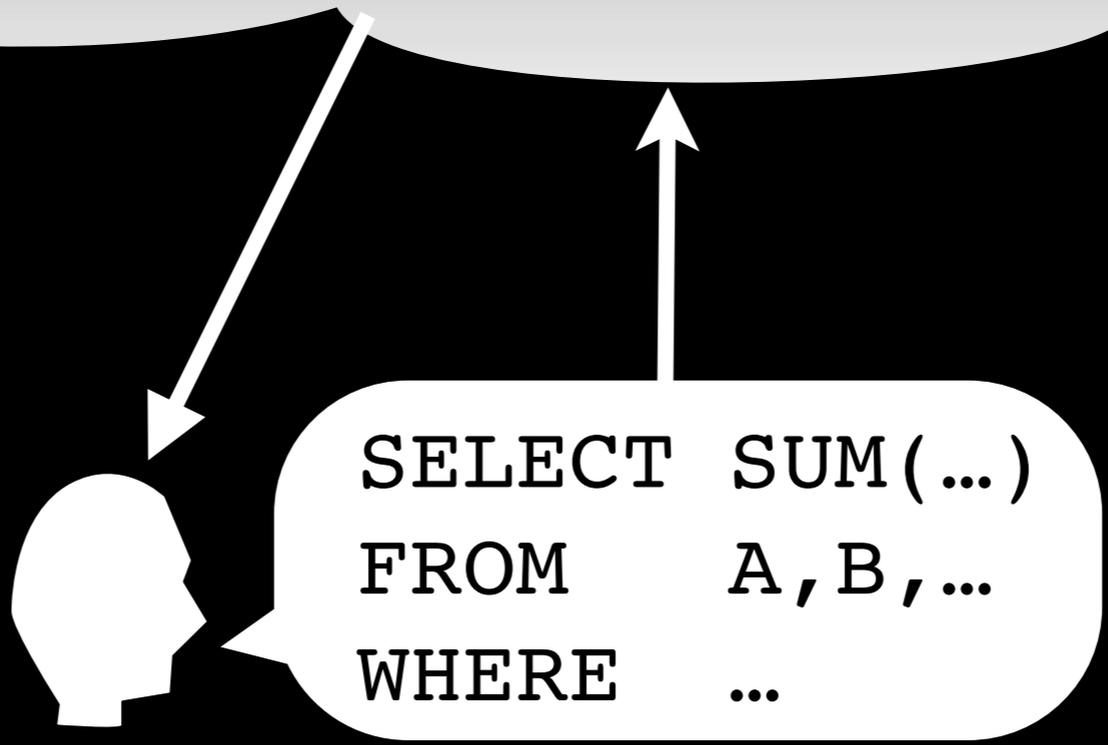
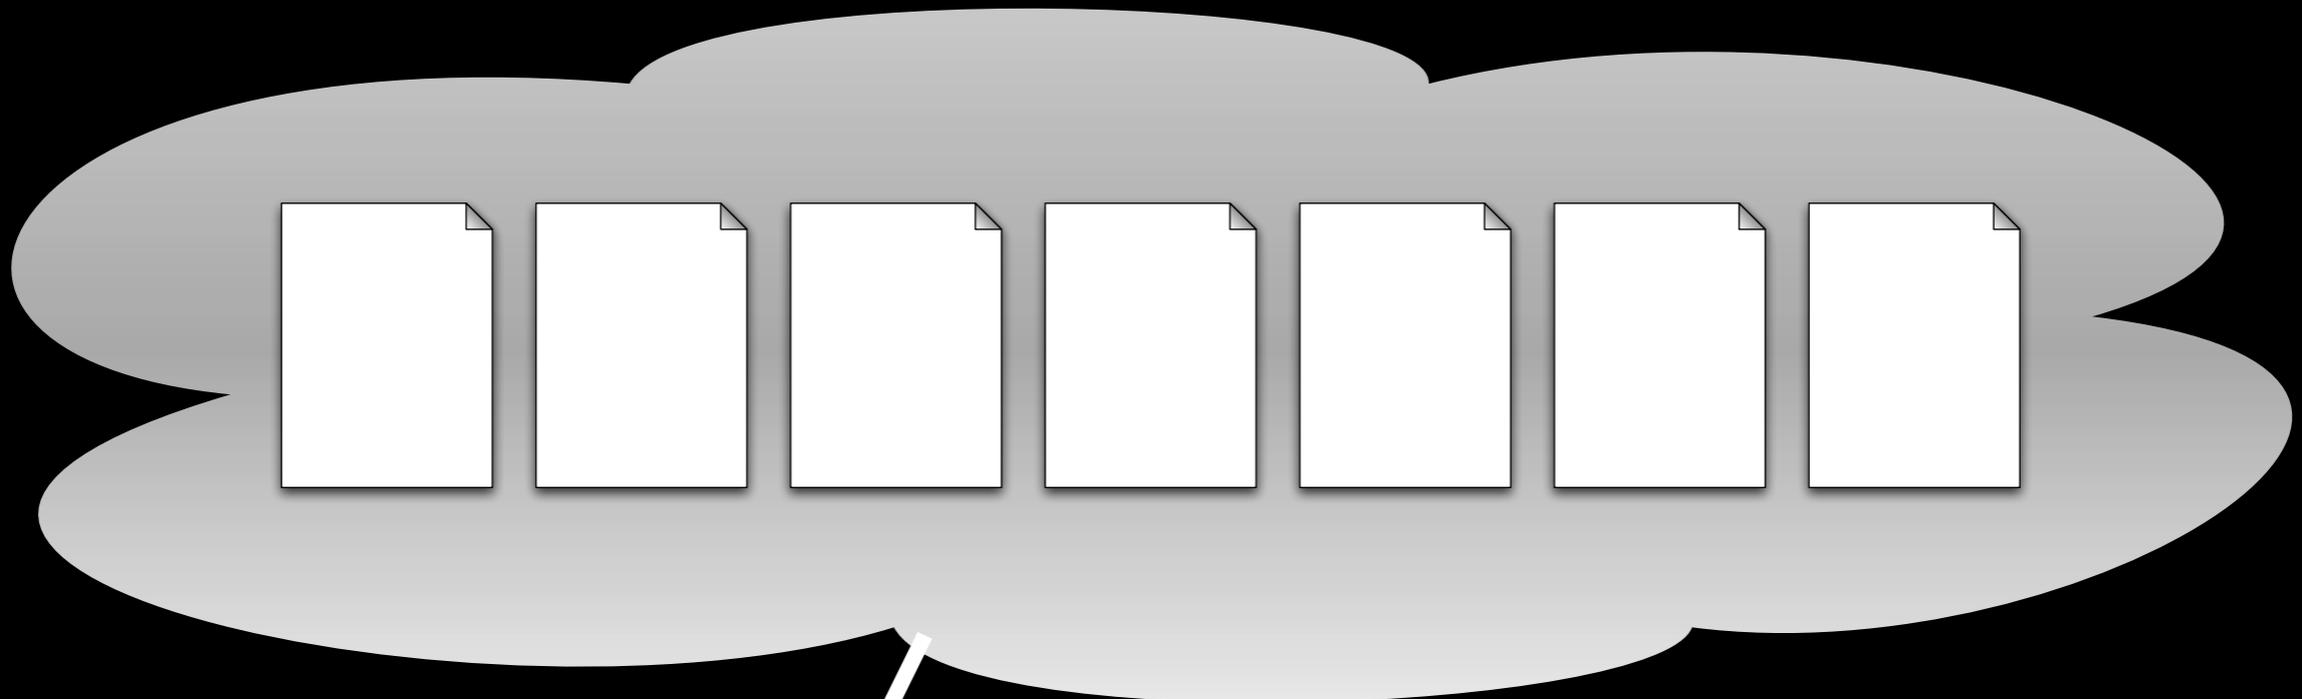




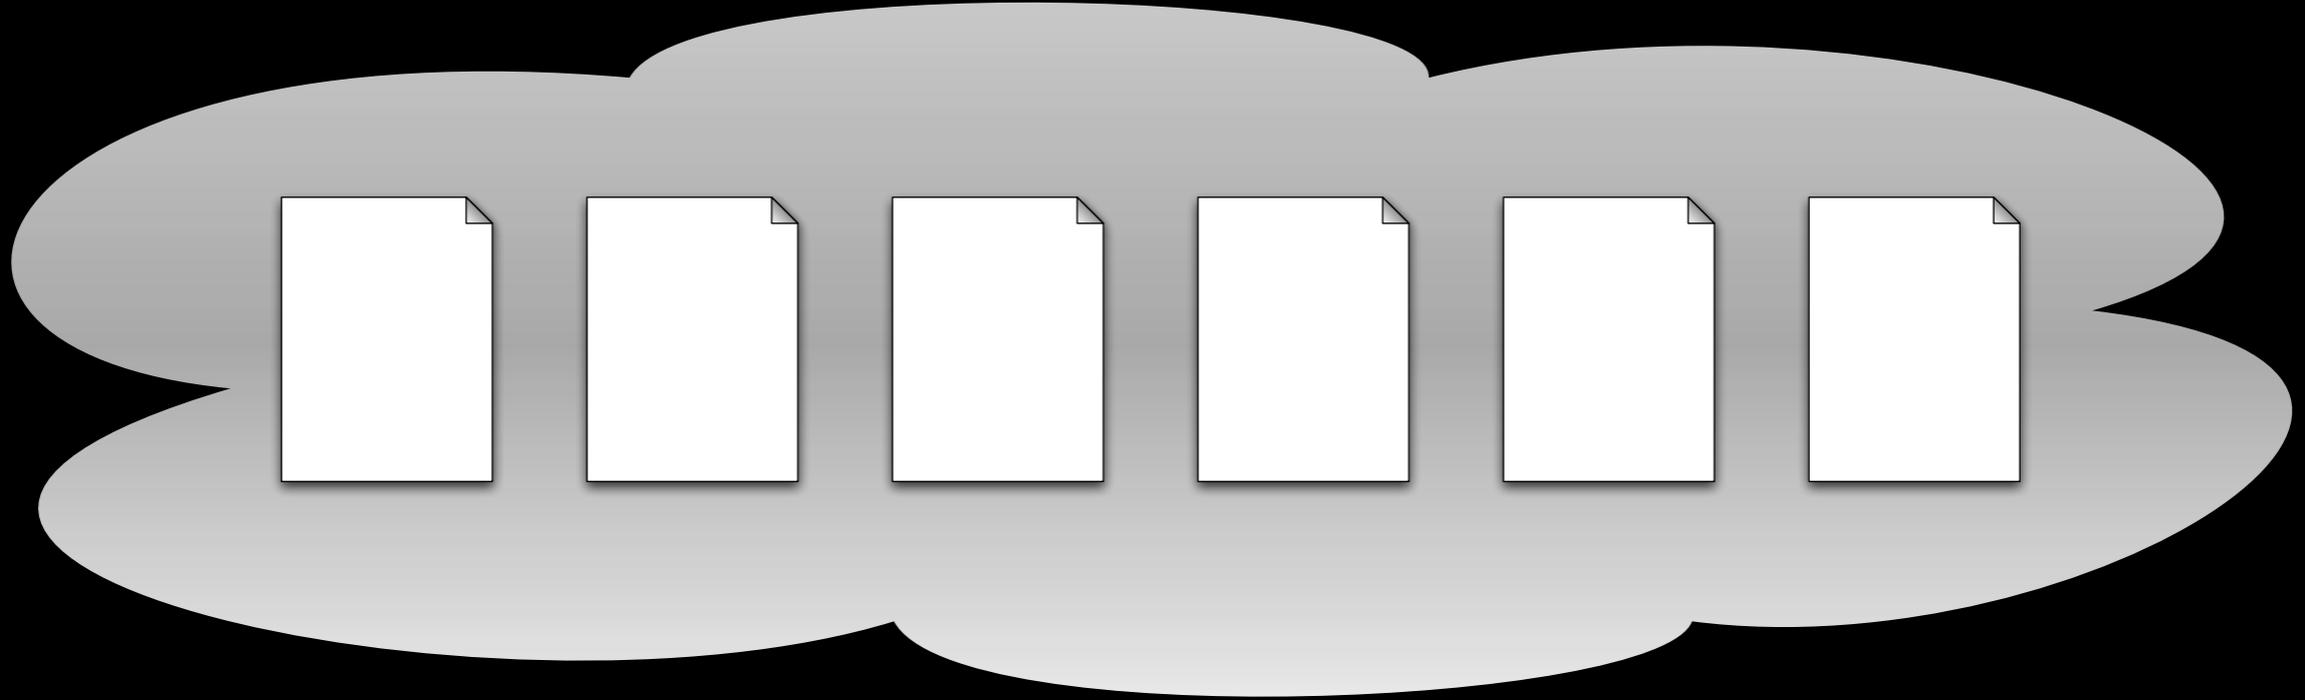


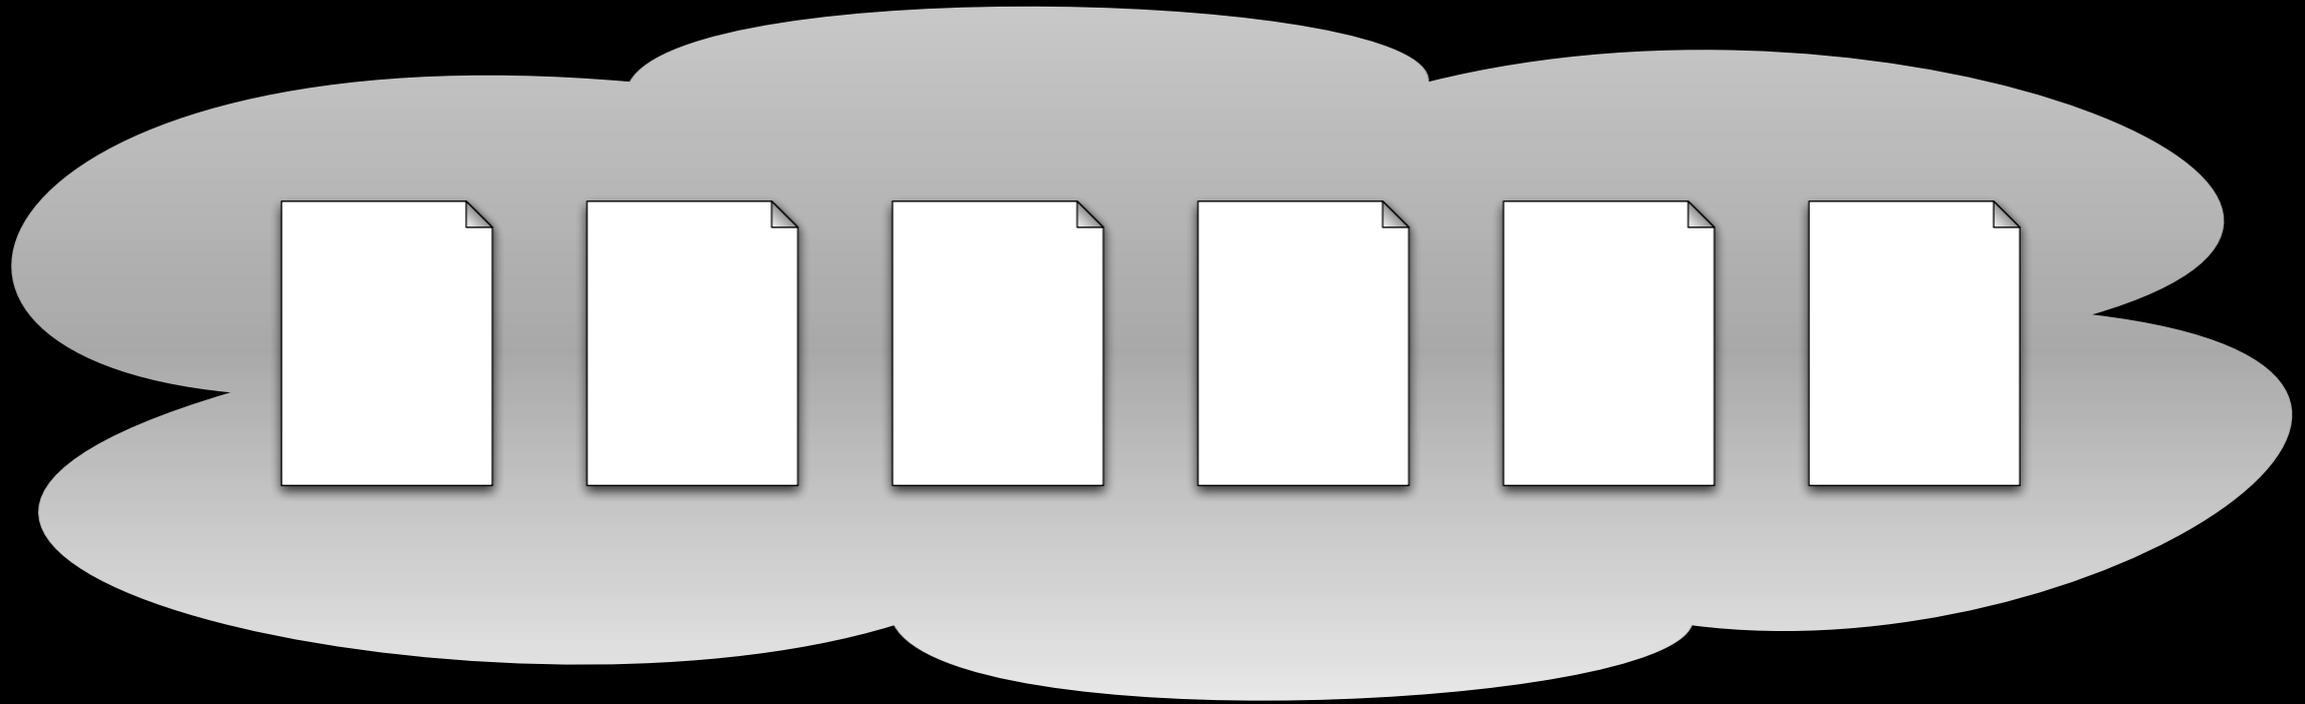
SELECT SUM (...)
FROM A, B, ...
WHERE ...



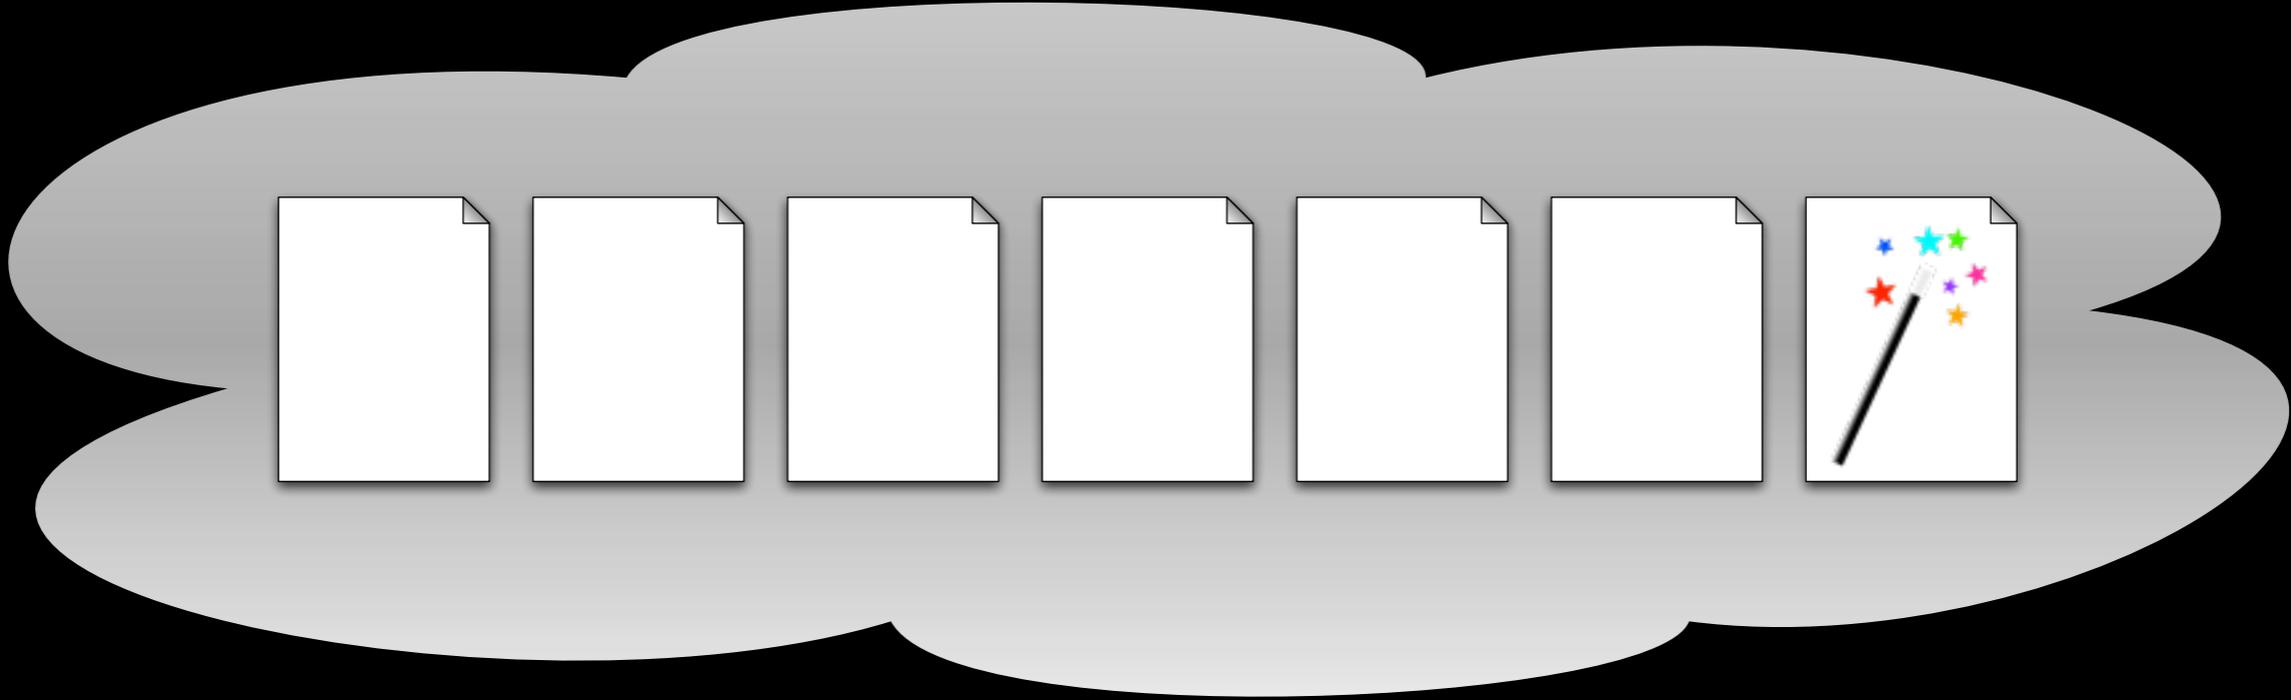


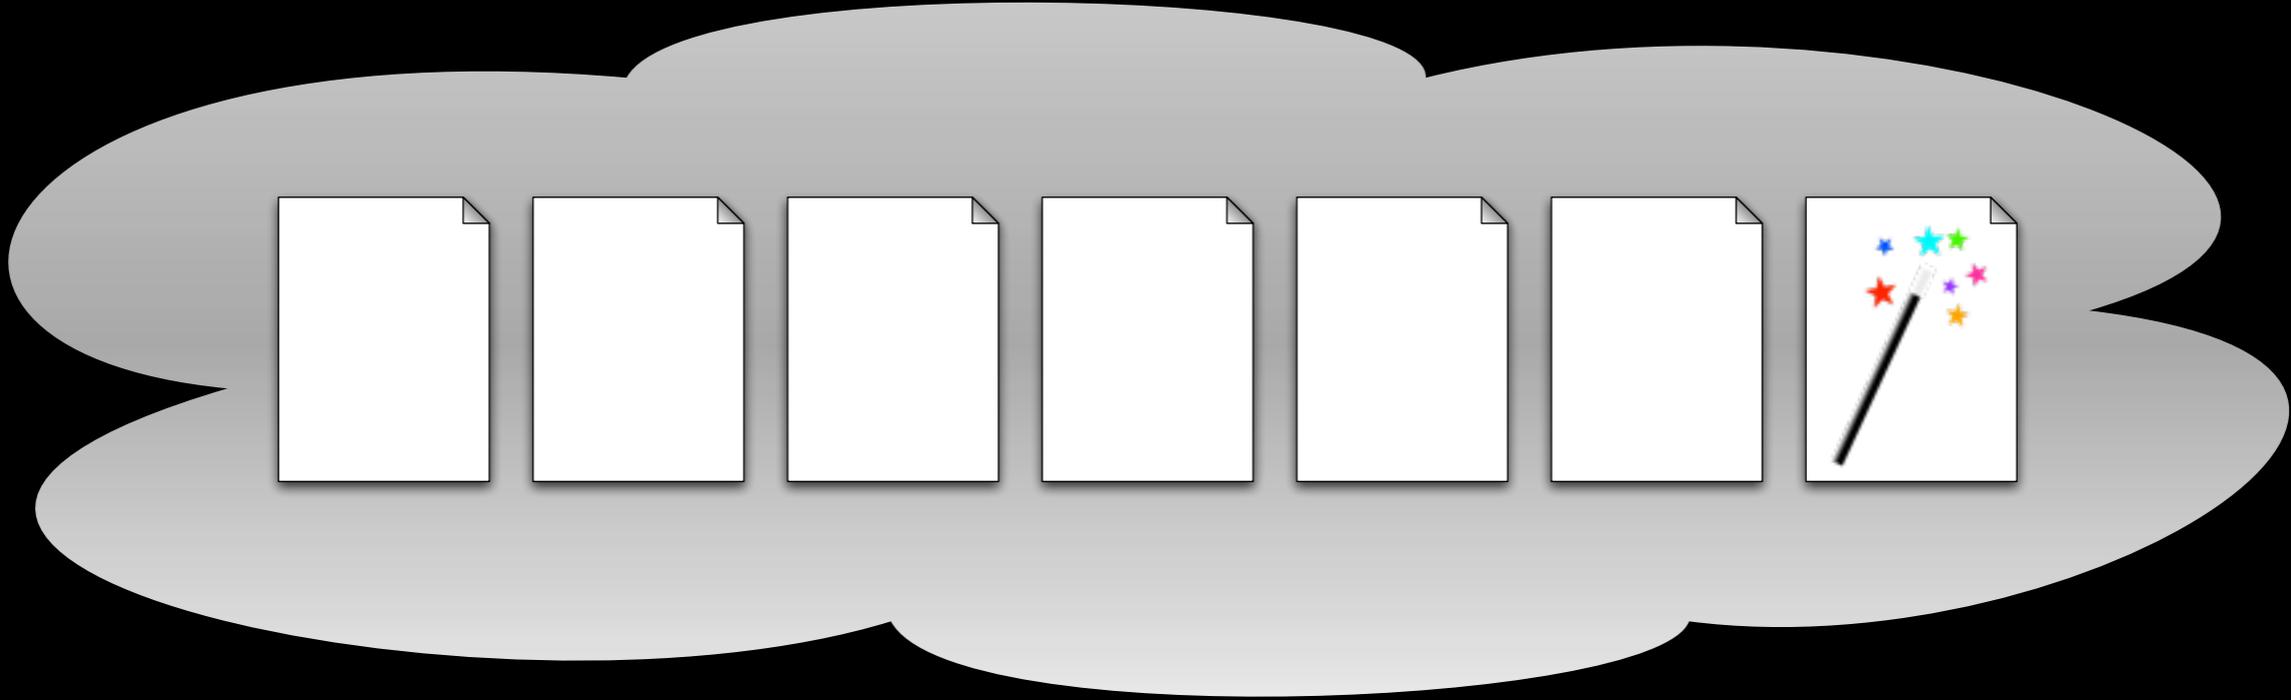
```
SELECT SUM(...)  
FROM A, B, ...  
WHERE ...
```

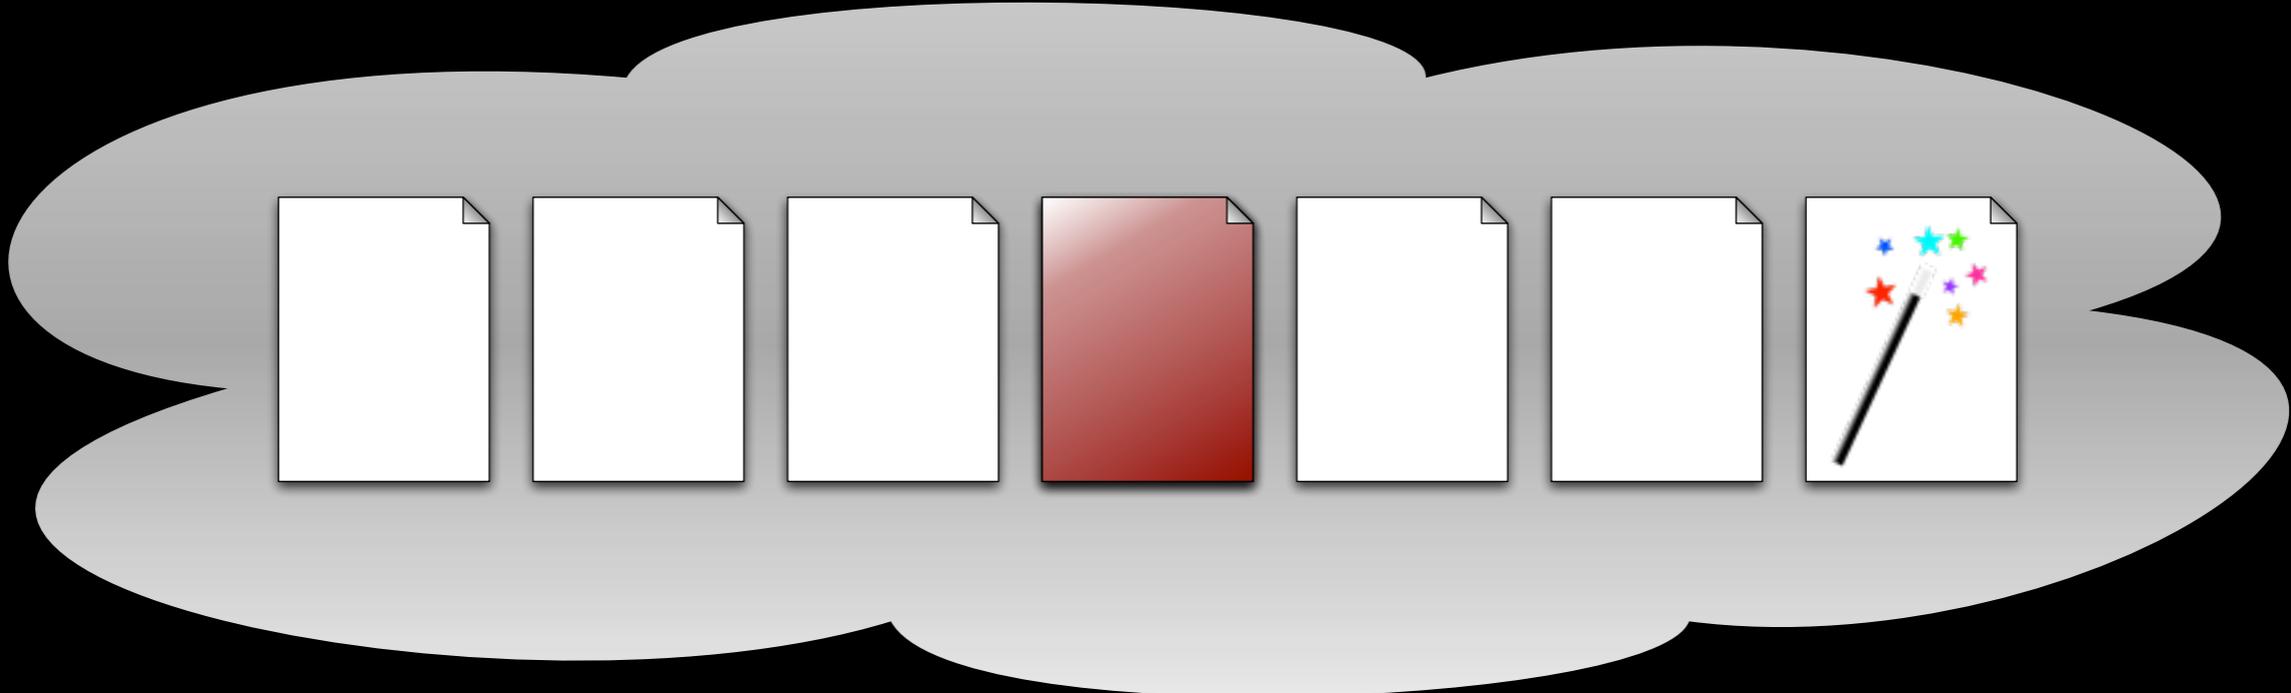




```
SELECT SUM(...)  
FROM A, B, ...  
WHERE ...
```







“Magic Maps”



Map Maintenance Messages

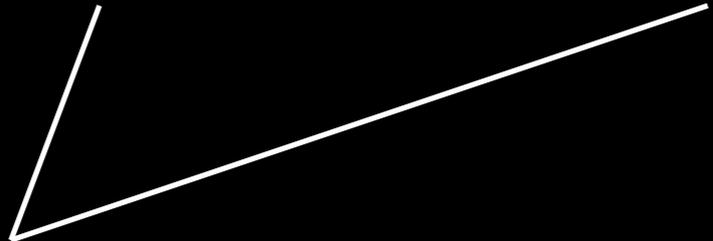
System Design

Hybrid Consistency

$\text{Map}_1 [V_1 , V_2 , \dots]$

$\text{Map}_1[V_1, V_2, \dots] \text{ += } f(\text{Map}_2, \text{Map}_3, \dots, V_1, V_2, \dots)$

Insert Delete



```
ON Event(V1, V2, ...) DO {  
  Map1[V1, V2, ...] += f(Map2, Map3, ..., V1, V2, ...)  
}
```

```
ON Event (V1, V2, ...) DO {  
  Map1[V1, V2, ...] += f (Map2, Map3, ..., V1, V2, ...)  
  Map2[V1, V2, ...] += f (Map1, Map3, ..., V1, V2, ...)  
}
```

```
SELECT SUM(R.A * T.D)
FROM R, S, T
WHERE R.B = S.B
      AND S.C = T.C
```

```
ON Insert_R(VA,VB) DO {  
  
    SELECT SUM(R.A * T.D)  
    FROM   R, S, T  
    WHERE  R.B = S.B  
          AND S.C = T.C  
  
}
```

```
ON Insert_R(VA,VB) DO {  
  
    Result[] +=  
  
        SELECT SUM(R.A * T.D)  
        FROM   R, S, T  
        WHERE  R.B = S.B  
              AND S.C = T.C  
  
}
```

```
ON Insert_R(VA,VB) DO {  
  
    Result[] +=  
  
        SELECT VA * SUM(T.D)  
        FROM      S, T  
        WHERE     VB = S.B  
                AND S.C = T.C  
  
}
```

```
ON Insert_R(VA, VB) DO {
```

```
    Result[] +=
```

```
        VA * Map2[VB]
```

```
}
```

```
ON Insert_R(A,B) DO {  
  Result[] += A * Map2[B]  
  Map1[C] += A * Map3[B,C]  
  Map5[B] += A }
```

```
ON Insert_T(C,D) DO {  
  Result[] += Map1[C] * D  
  Map2[B] += D * Map3[B,C]  
  Map4[C] += D }
```

```
ON Insert_S(B,C) DO {  
  Result[] += Map5[B] * Map4[C]  
  Map1[C] += Map5[B]  
  Map2[B] += Map4[C]  
  Map3[B,C] += 1 }
```

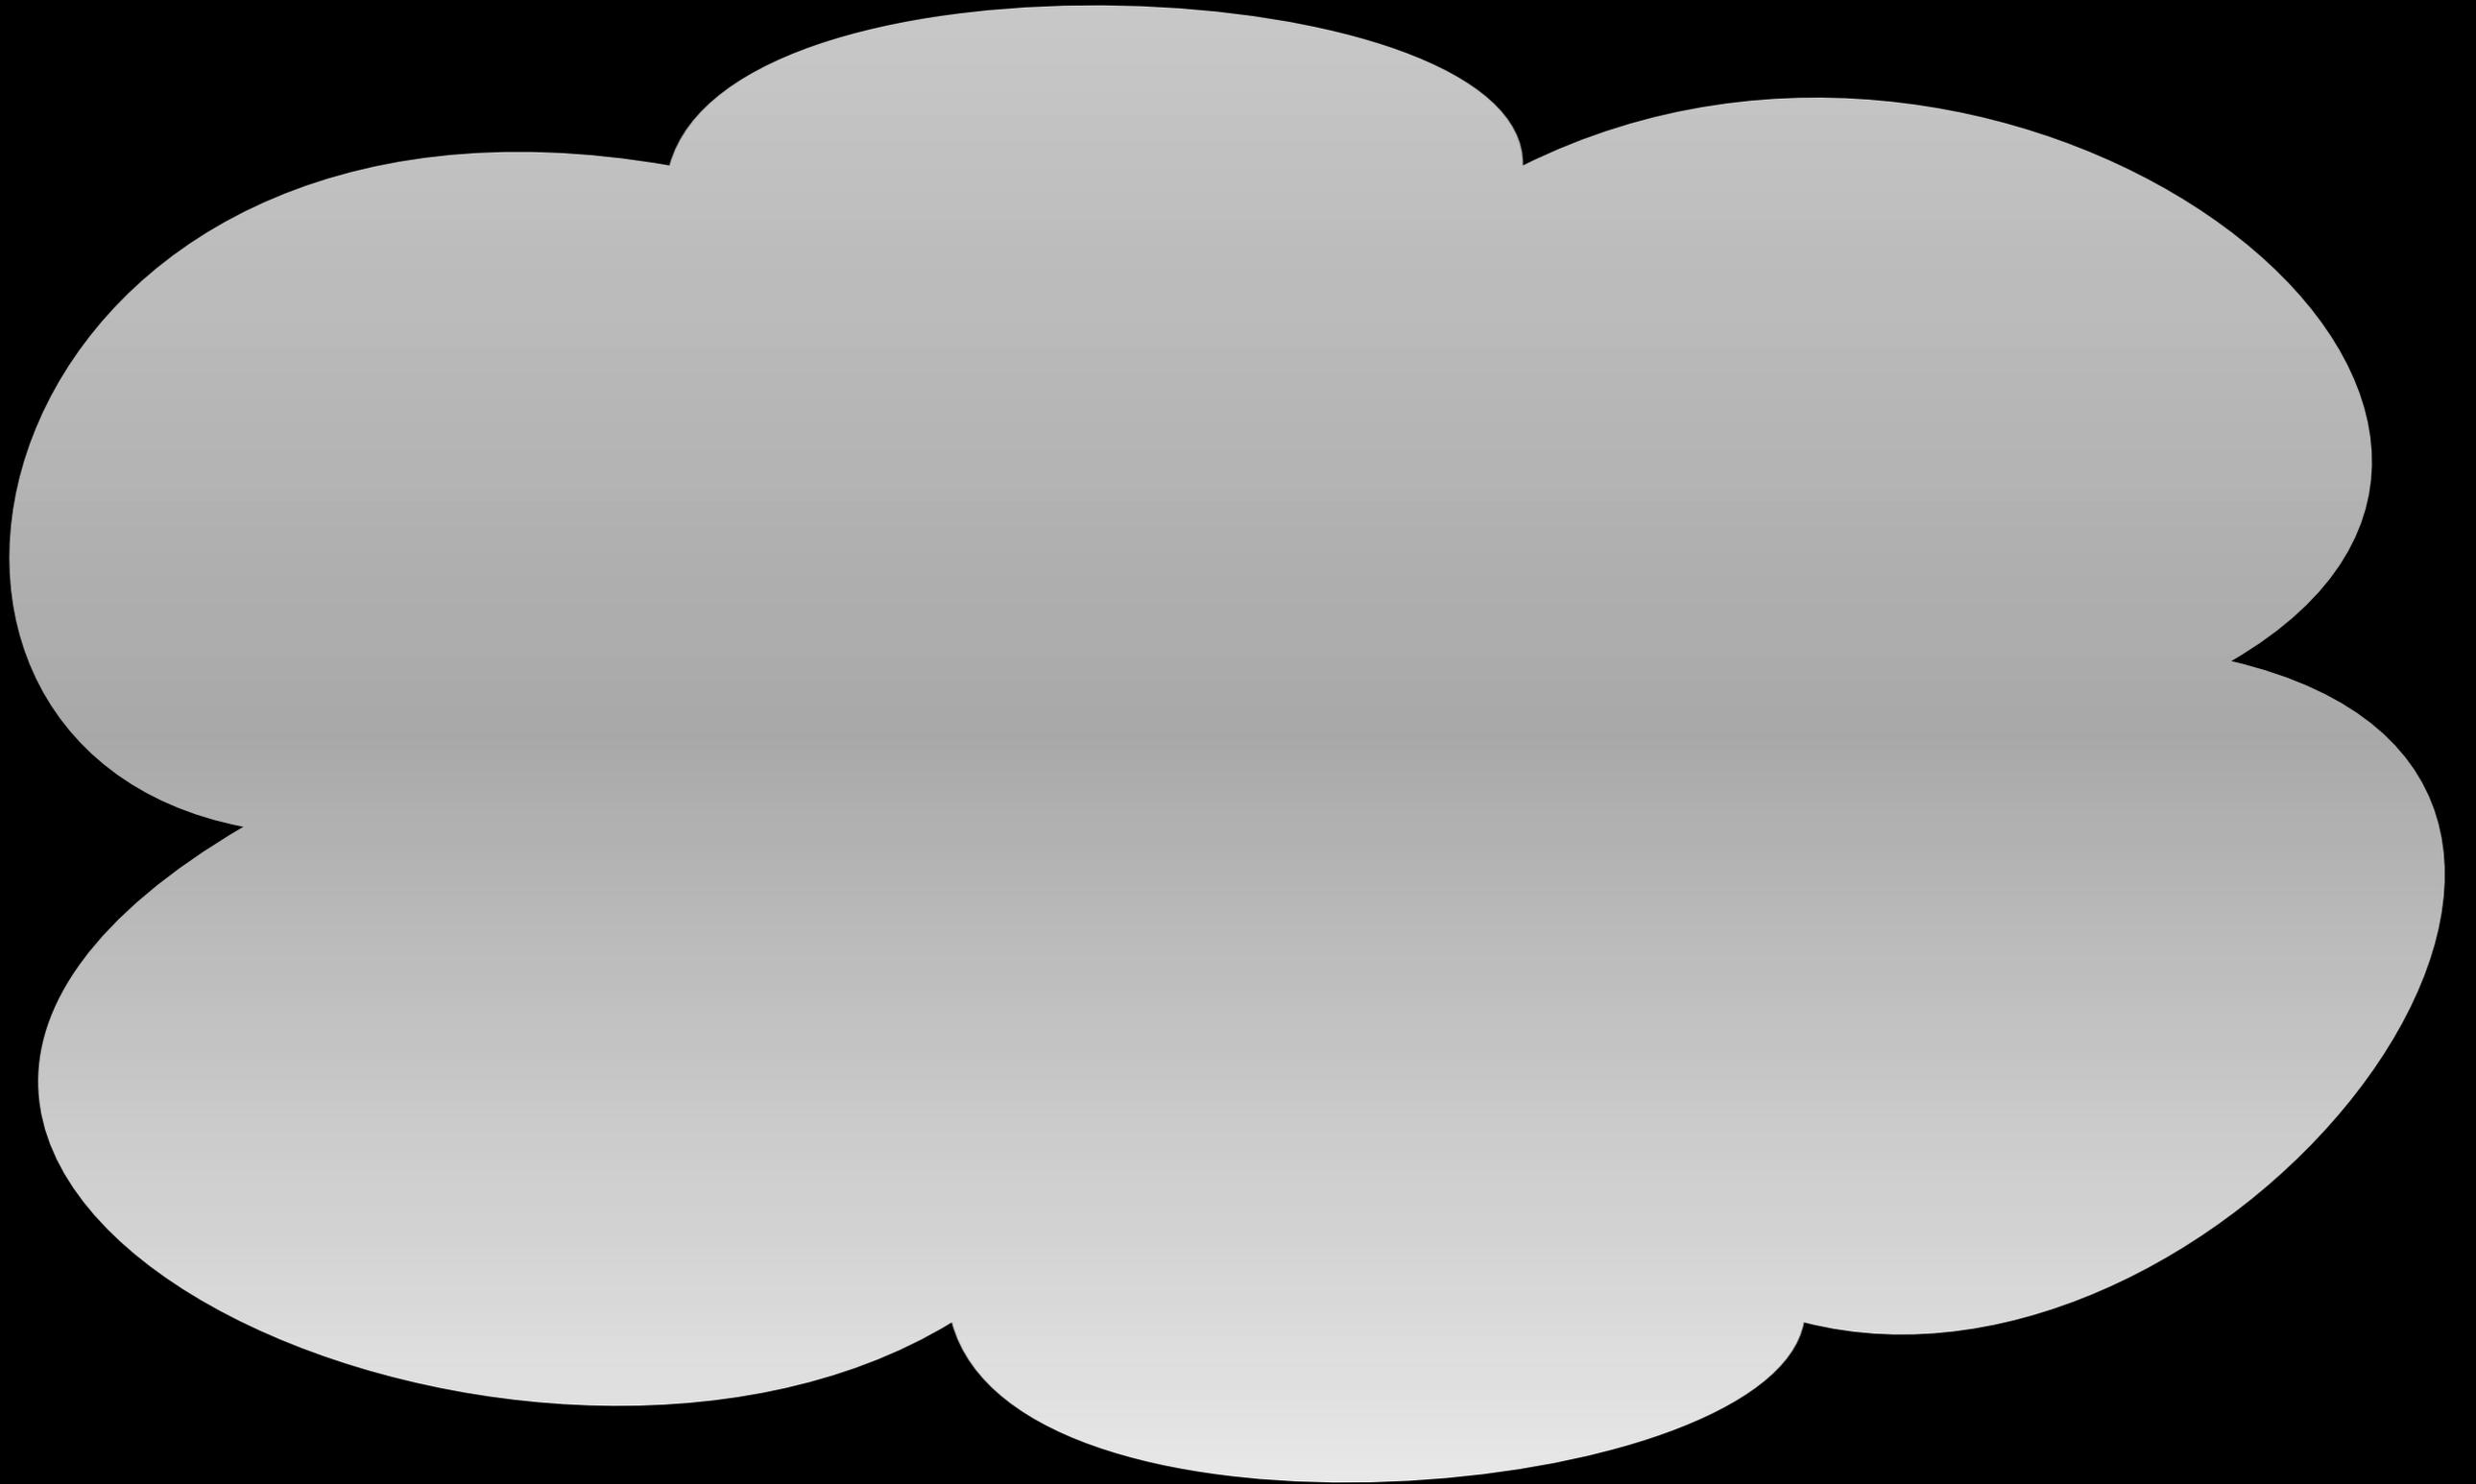
“Magic Maps”

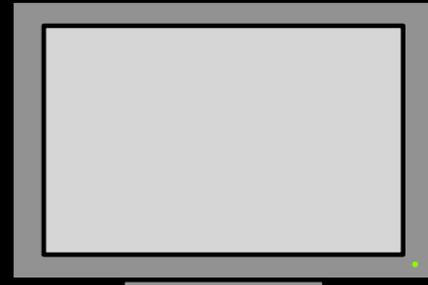
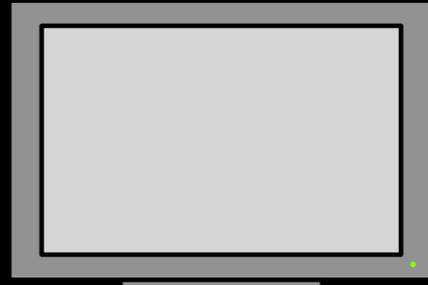
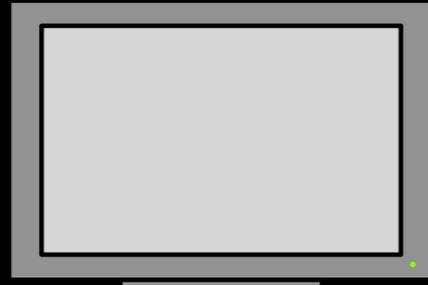
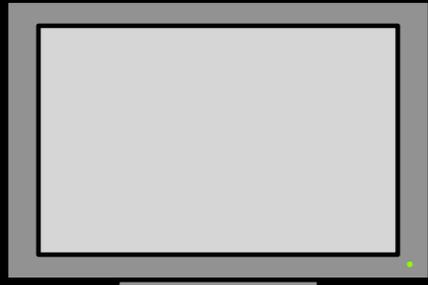
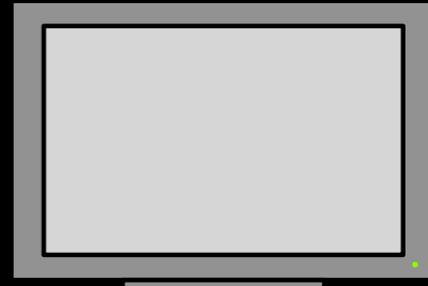
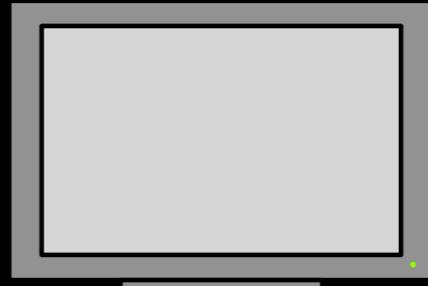
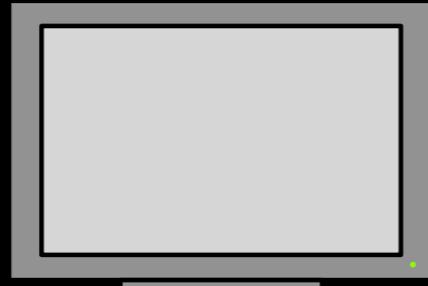
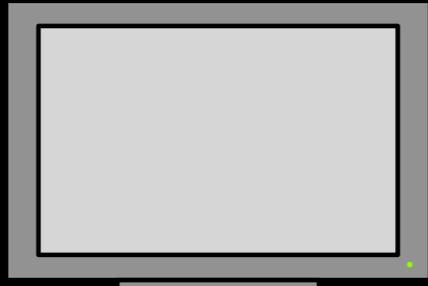
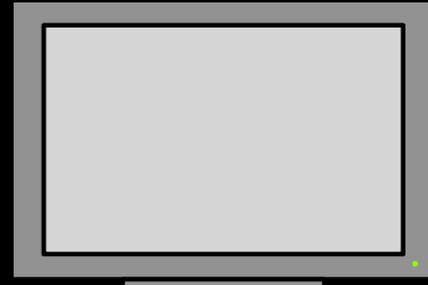
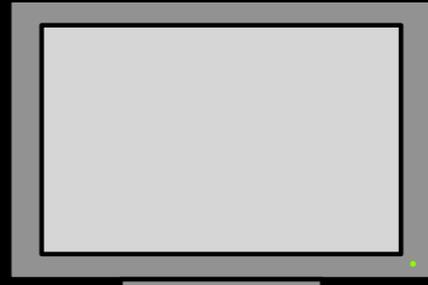
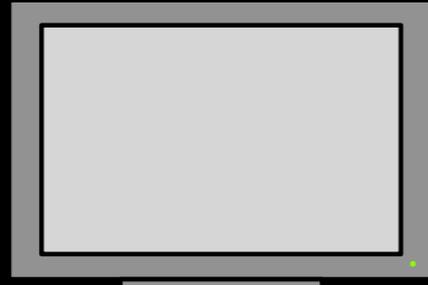
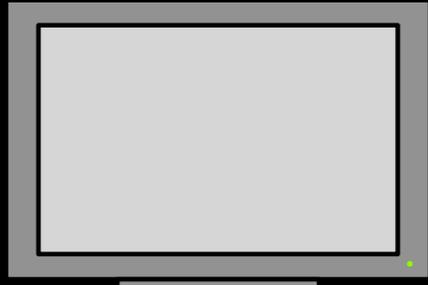
Map Maintenance Messages



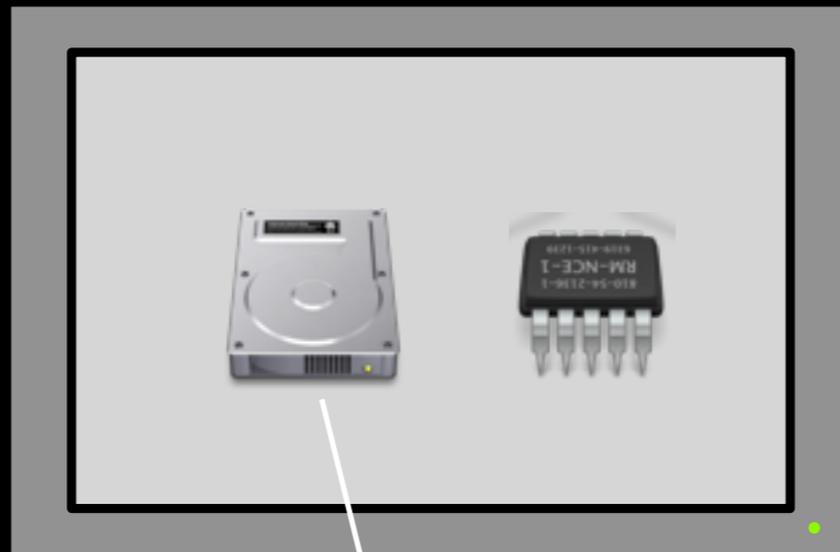
System Design

Hybrid Consistency









Map₁

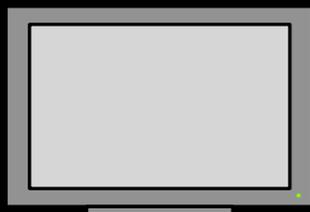
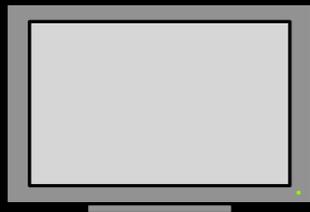
Map₂

Map₃

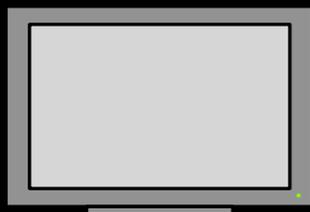
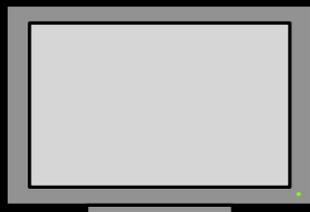
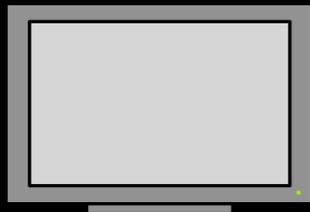
...

$\text{Map}_1[] += \text{Map}_2[] * \text{Map}_3[]$

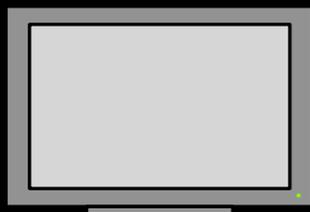
$Map_1 [] += Map_2 [] * Map_3 []$



$\text{Map}_1 [] += \text{Map}_2 [] * \text{Map}_3 []$

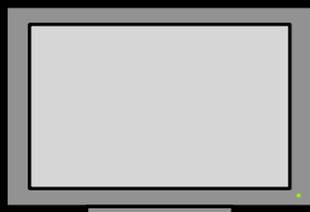
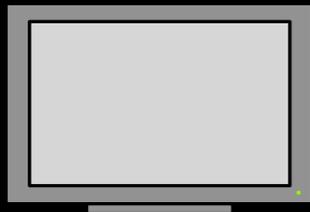


$$\text{Map}_1 [] += \text{Map}_2 [] * \text{Map}_3 []$$



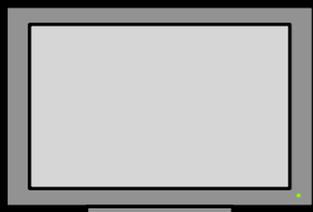
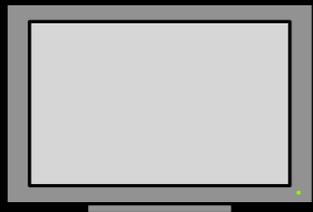
Map ₁	
Map ₂	
Map ₃	

$$\text{Map}_1[] += \text{Map}_2[] * \text{Map}_3[]$$

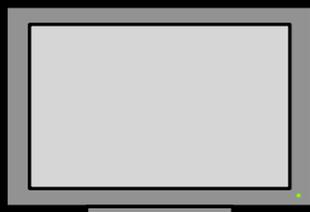
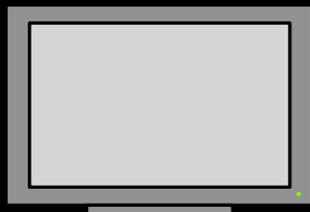
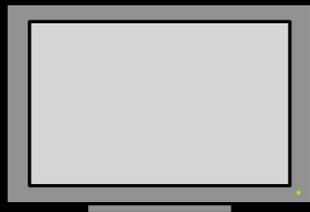


Map_1	$\text{Map}_2 * \text{Map}_3$
Map_2	
Map_3	

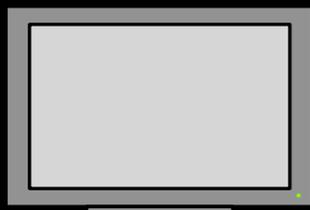
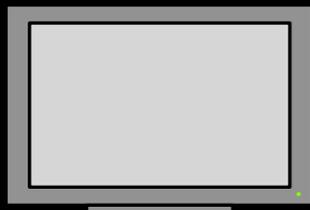
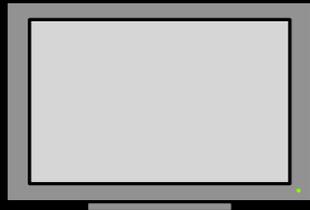
$$\text{Map}_1[] += \text{Map}_2[] * \text{Map}_3[]$$



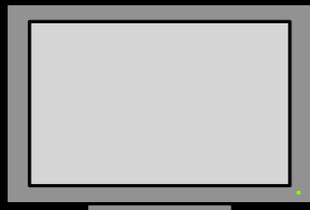
$\text{Map}_1 [] += \text{Map}_2 [] * \text{Map}_3 []$



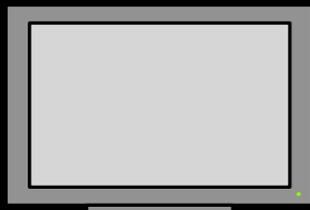
$\text{Map}_1 [] += \text{Map}_2 [] * \text{Map}_3 []$



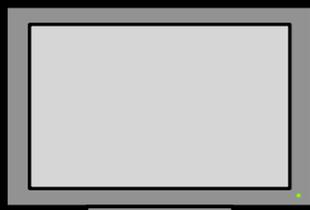
$$\text{Map}_1 [] += \text{Map}_2 [] * \text{Map}_3 []$$



$\text{Map}_2, \text{Map}_3$

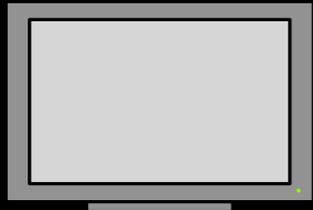


Map_1



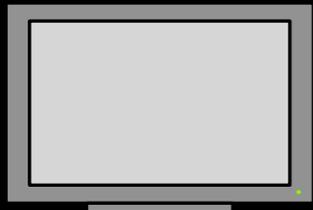
Map_1

$$\text{Map}_1[] += \text{Map}_2[] * \text{Map}_3[]$$

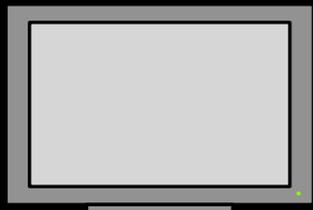


$\text{Map}_2, \text{Map}_3$

$\text{Map}_2 * \text{Map}_3$

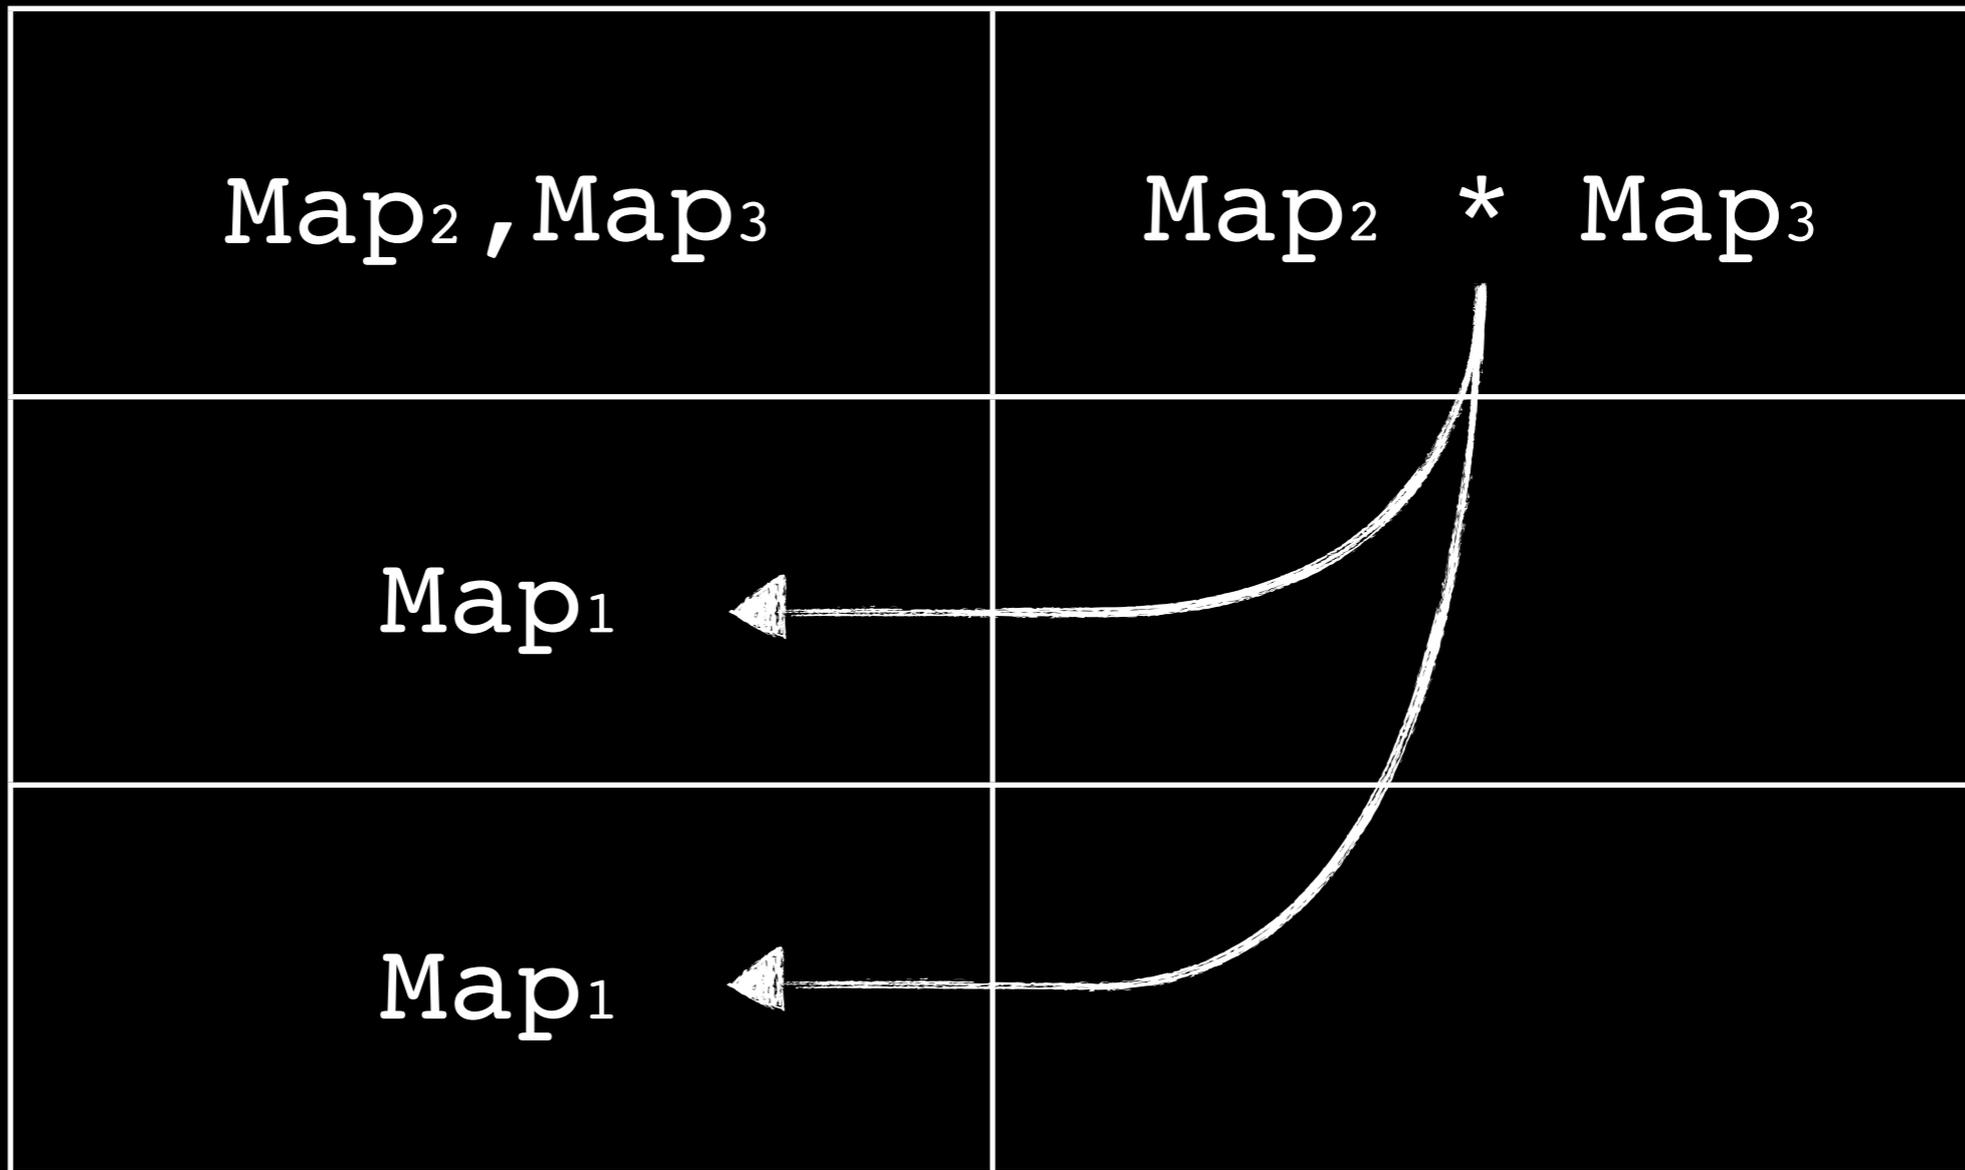
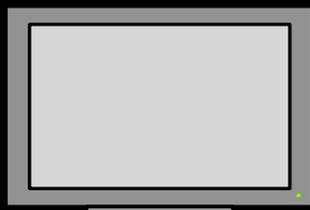
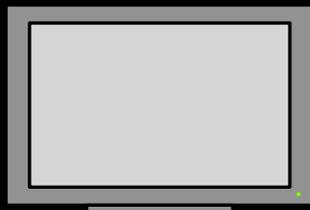
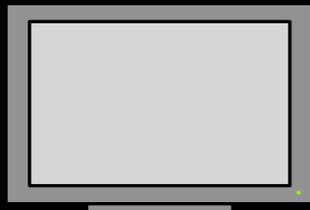


Map_1



Map_1

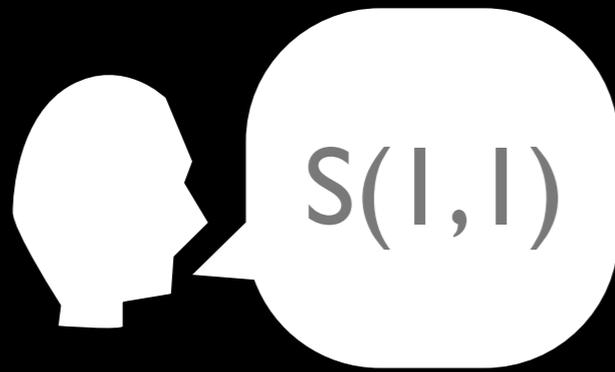
$$\text{Map}_1[] += \text{Map}_2[] * \text{Map}_3[]$$



```
ON Insert_R(A,B) DO {  
  Result[] += A * Map2[B]  
  Map1[C] += A * Map3[B,C]  
  Map5[B] += A }
```

```
ON Insert_T(C,D) DO {  
  Result[] += Map1[C] * D  
  Map2[B] += D * Map3[B,C]  
  Map4[C] += D }
```

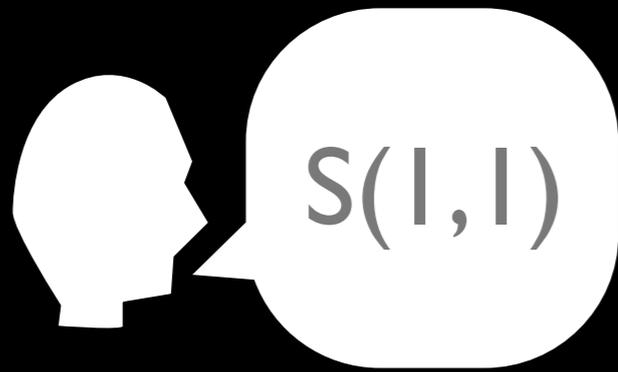
```
ON Insert_S(B,C) DO {  
  Result[] += Map5[B] * Map4[C]  
  Map1[C] += Map5[B]  
  Map2[B] += Map4[C]  
  Map3[B,C] += 1 }
```



```
ON Insert_R(A,B) DO {  
  Result[] += A * Map2[B]  
  Map1[C]   += A * Map3[B,C]  
  Map5[B]   += A }
```

```
ON Insert_T(C,D) DO {  
  Result[] += Map1[C] * D  
  Map2[B]   += D * Map3[B,C]  
  Map4[C]   += D }
```

```
ON Insert_S(B,C) DO {  
  Result[] += Map5[B] * Map4[C]  
  Map1[C]   += Map5[B]  
  Map2[B]   += Map4[C]  
  Map3[B,C] += 1 }
```

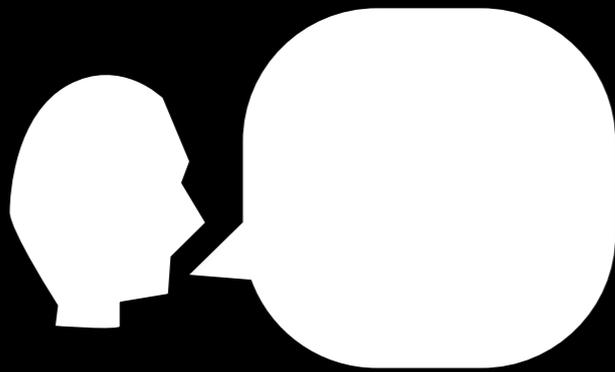


```
ON Insert_R(A,B) DO {  
  Result[] += A * Map2[B]  
  Map1[C]   += A * Map3[B,C]  
  Map5[B]   += A }
```

```
ON Insert_T(C,D) DO {  
  Result[] += Map1[C] * D  
  Map2[B]   += D * Map3[B,C]  
  Map4[C]   += D }
```

```
ON Insert_S(B,C) DO {  
  Result[] += Map5[B] * Map4[C]  
  Map1[C]   += Map5[B]  
  Map2[B]   += Map4[C]  
  Map3[B,C] += 1 }
```

Map1	
Map2	
Map3	
Map4	
Map5	
Result	

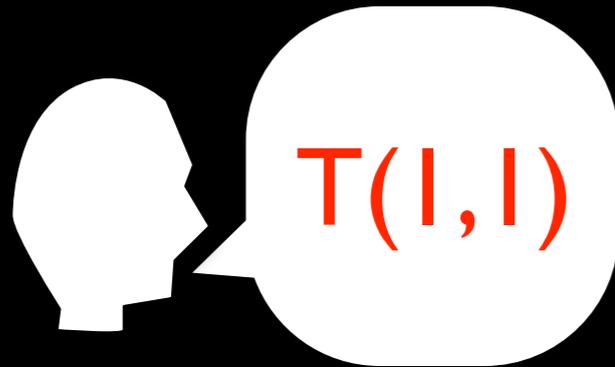
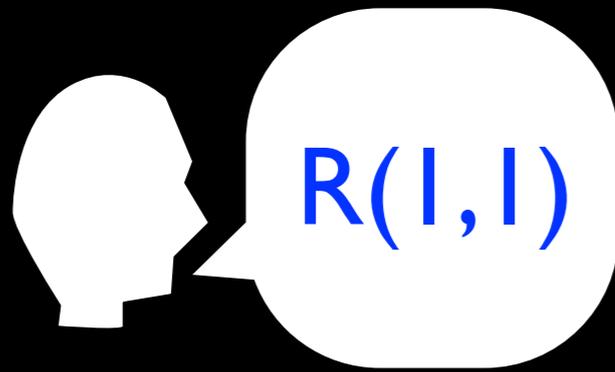


```
ON Insert_R(A,B) DO {  
  Result[] += A * Map2[B]  
  Map1[C]   += A * Map3[B,C]  
  Map5[B]   += A }
```

```
ON Insert_T(C,D) DO {  
  Result[] += Map1[C] * D  
  Map2[B]   += D * Map3[B,C]  
  Map4[C]   += D }
```

```
ON Insert_S(B,C) DO {  
  Result[] += Map5[B] * Map4[C]  
  Map1[C]   += Map5[B]  
  Map2[B]   += Map4[C]  
  Map3[B,C] += 1 }
```

Map1	
Map2	
Map3	$\langle 1, 1 \rangle \rightarrow 1$
Map4	
Map5	
Result	

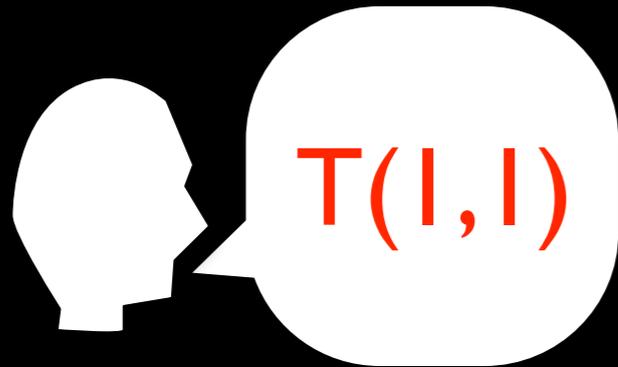
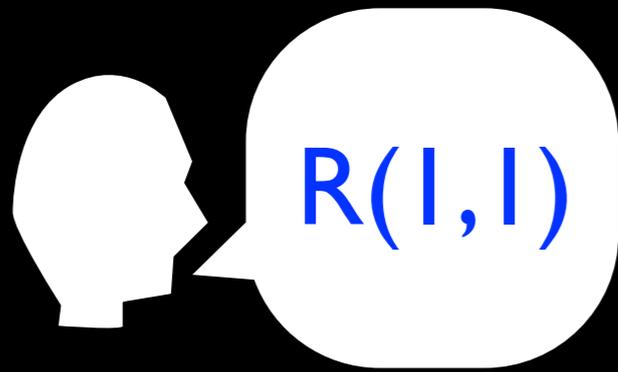


```
ON Insert_R(A,B) DO {  
  Result[] += A * Map2[B]  
  Map1[C]   += A * Map3[B,C]  
  Map5[B]   += A }  
}
```

```
ON Insert_T(C,D) DO {  
  Result[] += Map1[C] * D  
  Map2[B]   += D * Map3[B,C]  
  Map4[C]   += D }  
}
```

```
ON Insert_S(B,C) DO {  
  Result[] += Map5[B] * Map4[C]  
  Map1[C]   += Map5[B]  
  Map2[B]   += Map4[C]  
  Map3[B,C] += 1 }  
}
```

Map1	
Map2	
Map3	<1,1>→1
Map4	
Map5	
Result	



Map1
Map3
Map4

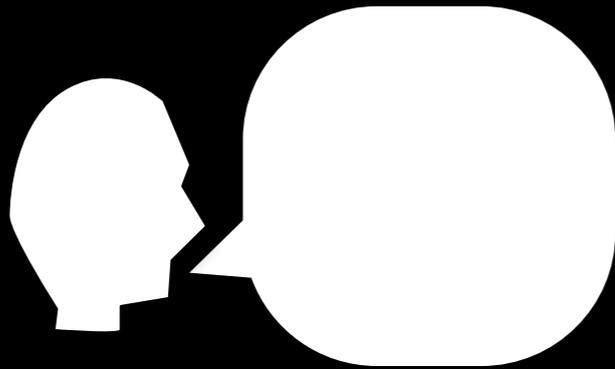
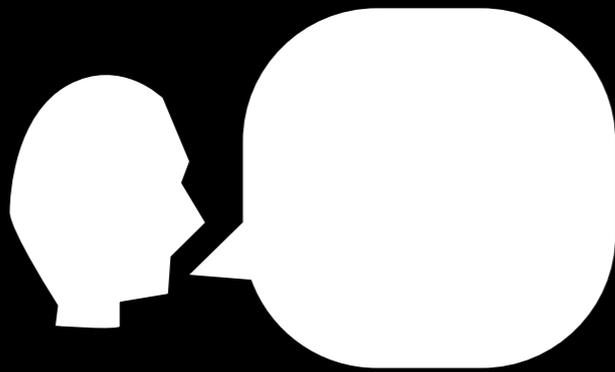
Map2
Map3
Map5

```
ON Insert_R(A,B) DO {  
  Result[] += A * Map2[B]  
  Map1[C]   += A * Map3[B,C]  
  Map5[B]   += A }  
}
```

```
ON Insert_T(C,D) DO {  
  Result[] += Map1[C] * D  
  Map2[B]  += D * Map3[B,C]  
  Map4[C]  += D }  
}
```

```
ON Insert_S(B,C) DO {  
  Result[] += Map5[B] * Map4[C]  
  Map1[C]  += Map5[B]  
  Map2[B]  += Map4[C]  
  Map3[B,C] += 1 }  
}
```

Map1	
Map2	
Map3	<1,1>→1
Map4	
Map5	
Result	



Map1
Map3
Map4

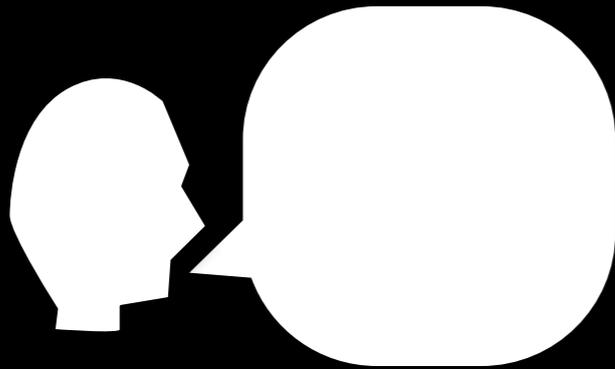
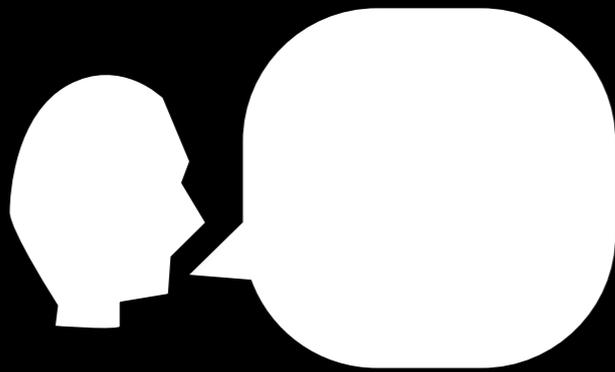
Map2
Map3
Map5

```
ON Insert_R(A,B) DO {
  Result[] += A * Map2[B]
  Map1[C] += A * Map3[B,C]
  Map5[B] += A }
```

```
ON Insert_T(C,D) DO {
  Result[] += Map1[C] * D
  Map2[B] += D * Map3[B,C]
  Map4[C] += D }
```

```
ON Insert_S(B,C) DO {
  Result[] += Map5[B] * Map4[C]
  Map1[C] += Map5[B]
  Map2[B] += Map4[C]
  Map3[B,C] += 1 }
```

Map1	
Map2	
Map3	<1, 1>→1
Map4	
Map5	
Result	

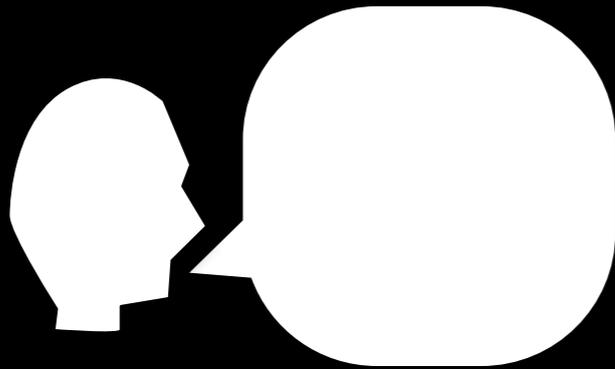
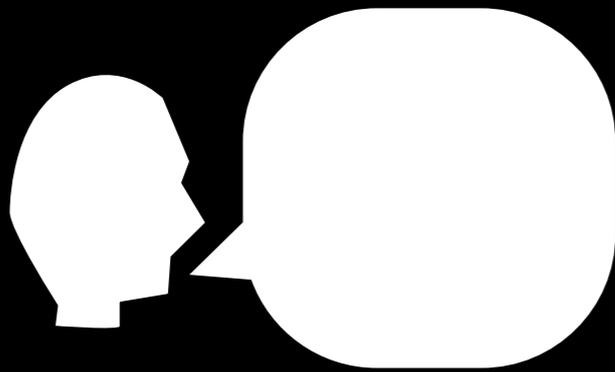


Map1
Map3
Map4

Map2
Map3
Map5

```
ON Insert_R(A,B) DO {  
  Result[] += A * Map2[B]  
  Map1[C] += A * Map3[B,C]  
  Map5[B] += A }  
  
ON Insert_T(C,D) DO {  
  Result[] += Map1[C] * D  
  Map2[B] += D * Map3[B,C]  
  Map4[C] += D }  
  
ON Insert_S(B,C) DO {  
  Result[] += Map5[B] * Map4[C]  
  Map1[C] += Map5[B]  
  Map2[B] += Map4[C]  
  Map3[B,C] += 1 }
```

Map1	
Map2	
Map3	<1,1>→1
Map4	<1>→1
Map5	<1>→1
Result	



Map1
Map3
Map4

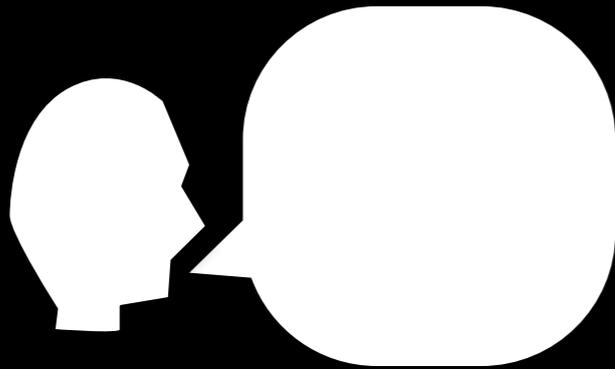
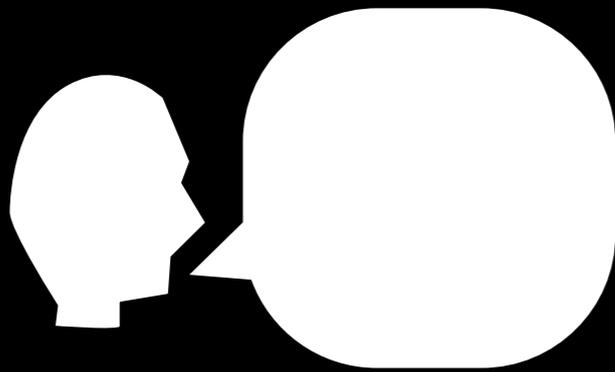
Map2
Map3
Map5

```
ON Insert_R(A,B) DO {  
  Result[] += A * Map2[B]  
  Map1[C] += A * Map3[B,C]  
  Map5[B] += A }
```

```
ON Insert_T(C,D) DO {  
  Result[] += Map1[C] * D  
  Map2[B] += D * Map3[B,C]  
  Map4[C] += D }
```

```
ON Insert_S(B,C) DO {  
  Result[] += Map5[B] * Map4[C]  
  Map1[C] += Map5[B]  
  Map2[B] += Map4[C]  
  Map3[B,C] += 1 }
```

Map1	<1>→1
Map2	<1>→1
Map3	<1, 1>→1
Map4	<1>→1
Map5	<1>→1
Result	



```

ON Insert_R(A,B) DO {
  Result[] += A * Map2[B]
  Map1[C] += A * Map3[B,C]
  Map5[B] += A }

```

```

ON Insert_T(C,D) DO {
  Result[] += Map1[C] * D
  Map2[B] += D * Map3[B,C]
  Map4[C] += D }

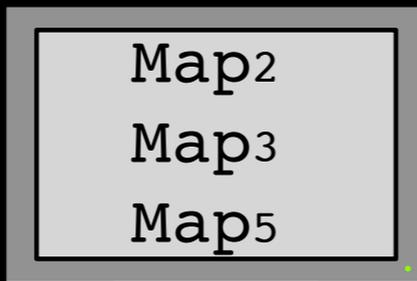
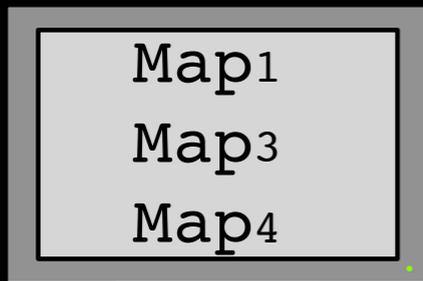
```

```

ON Insert_S(B,C) DO {
  Result[] += Map5[B] * Map4[C]
  Map1[C] += Map5[B]
  Map2[B] += Map4[C]
  Map3[B,C] += 1 }

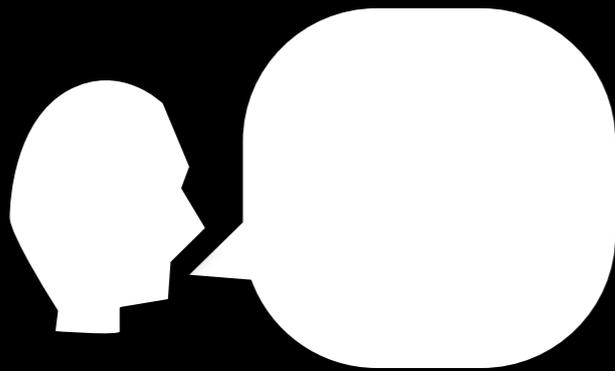
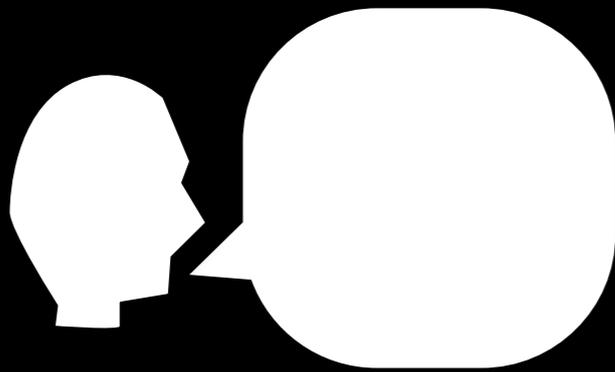
```

Map2[1] += 1



Map1[1] += 1

Map1	<1>→1
Map2	<1>→1
Map3	<1,1>→1
Map4	<1>→1
Map5	<1>→1
Result	



```

ON Insert R(A,B) DO {
  Result[] += A * Map2[B]
  Map1[C] += A * Map3[B,C]
  Map5[B] += A }

```

```

ON Insert T(C,D) DO {
  Result[] += Map1[C] * D
  Map2[B] += D * Map3[B,C]
  Map4[C] += D }

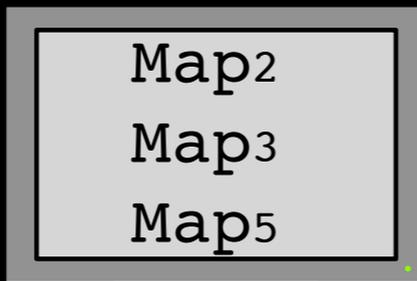
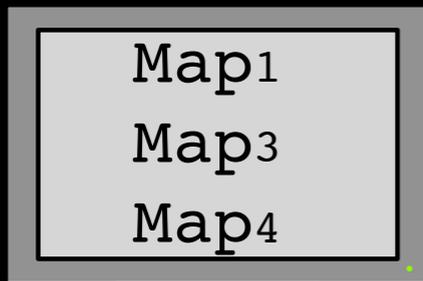
```

```

ON Insert_S(B,C) DO {
  Result[] += Map5[B] * Map4[C]
  Map1[C] += Map5[B]
  Map2[B] += Map4[C]
  Map3[B,C] += 1 }

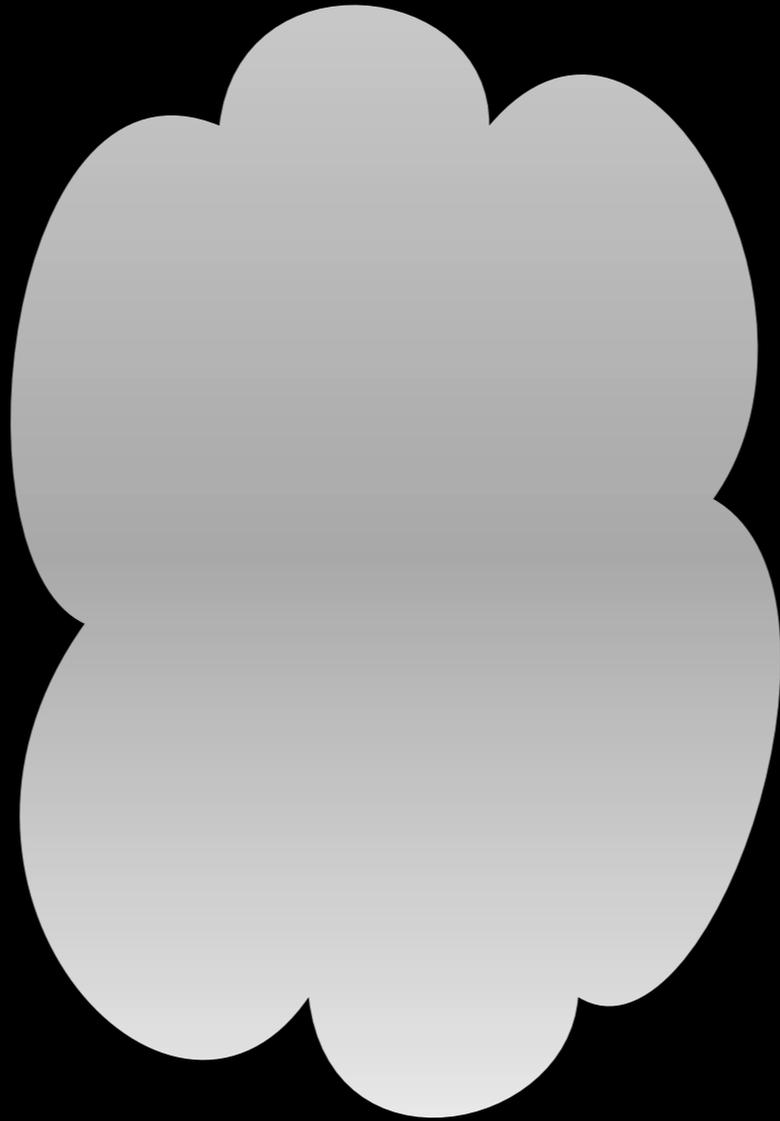
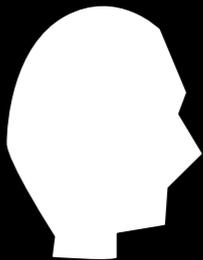
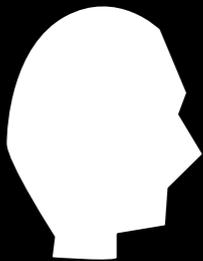
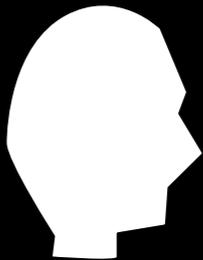
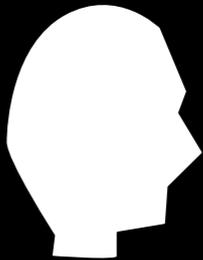
```

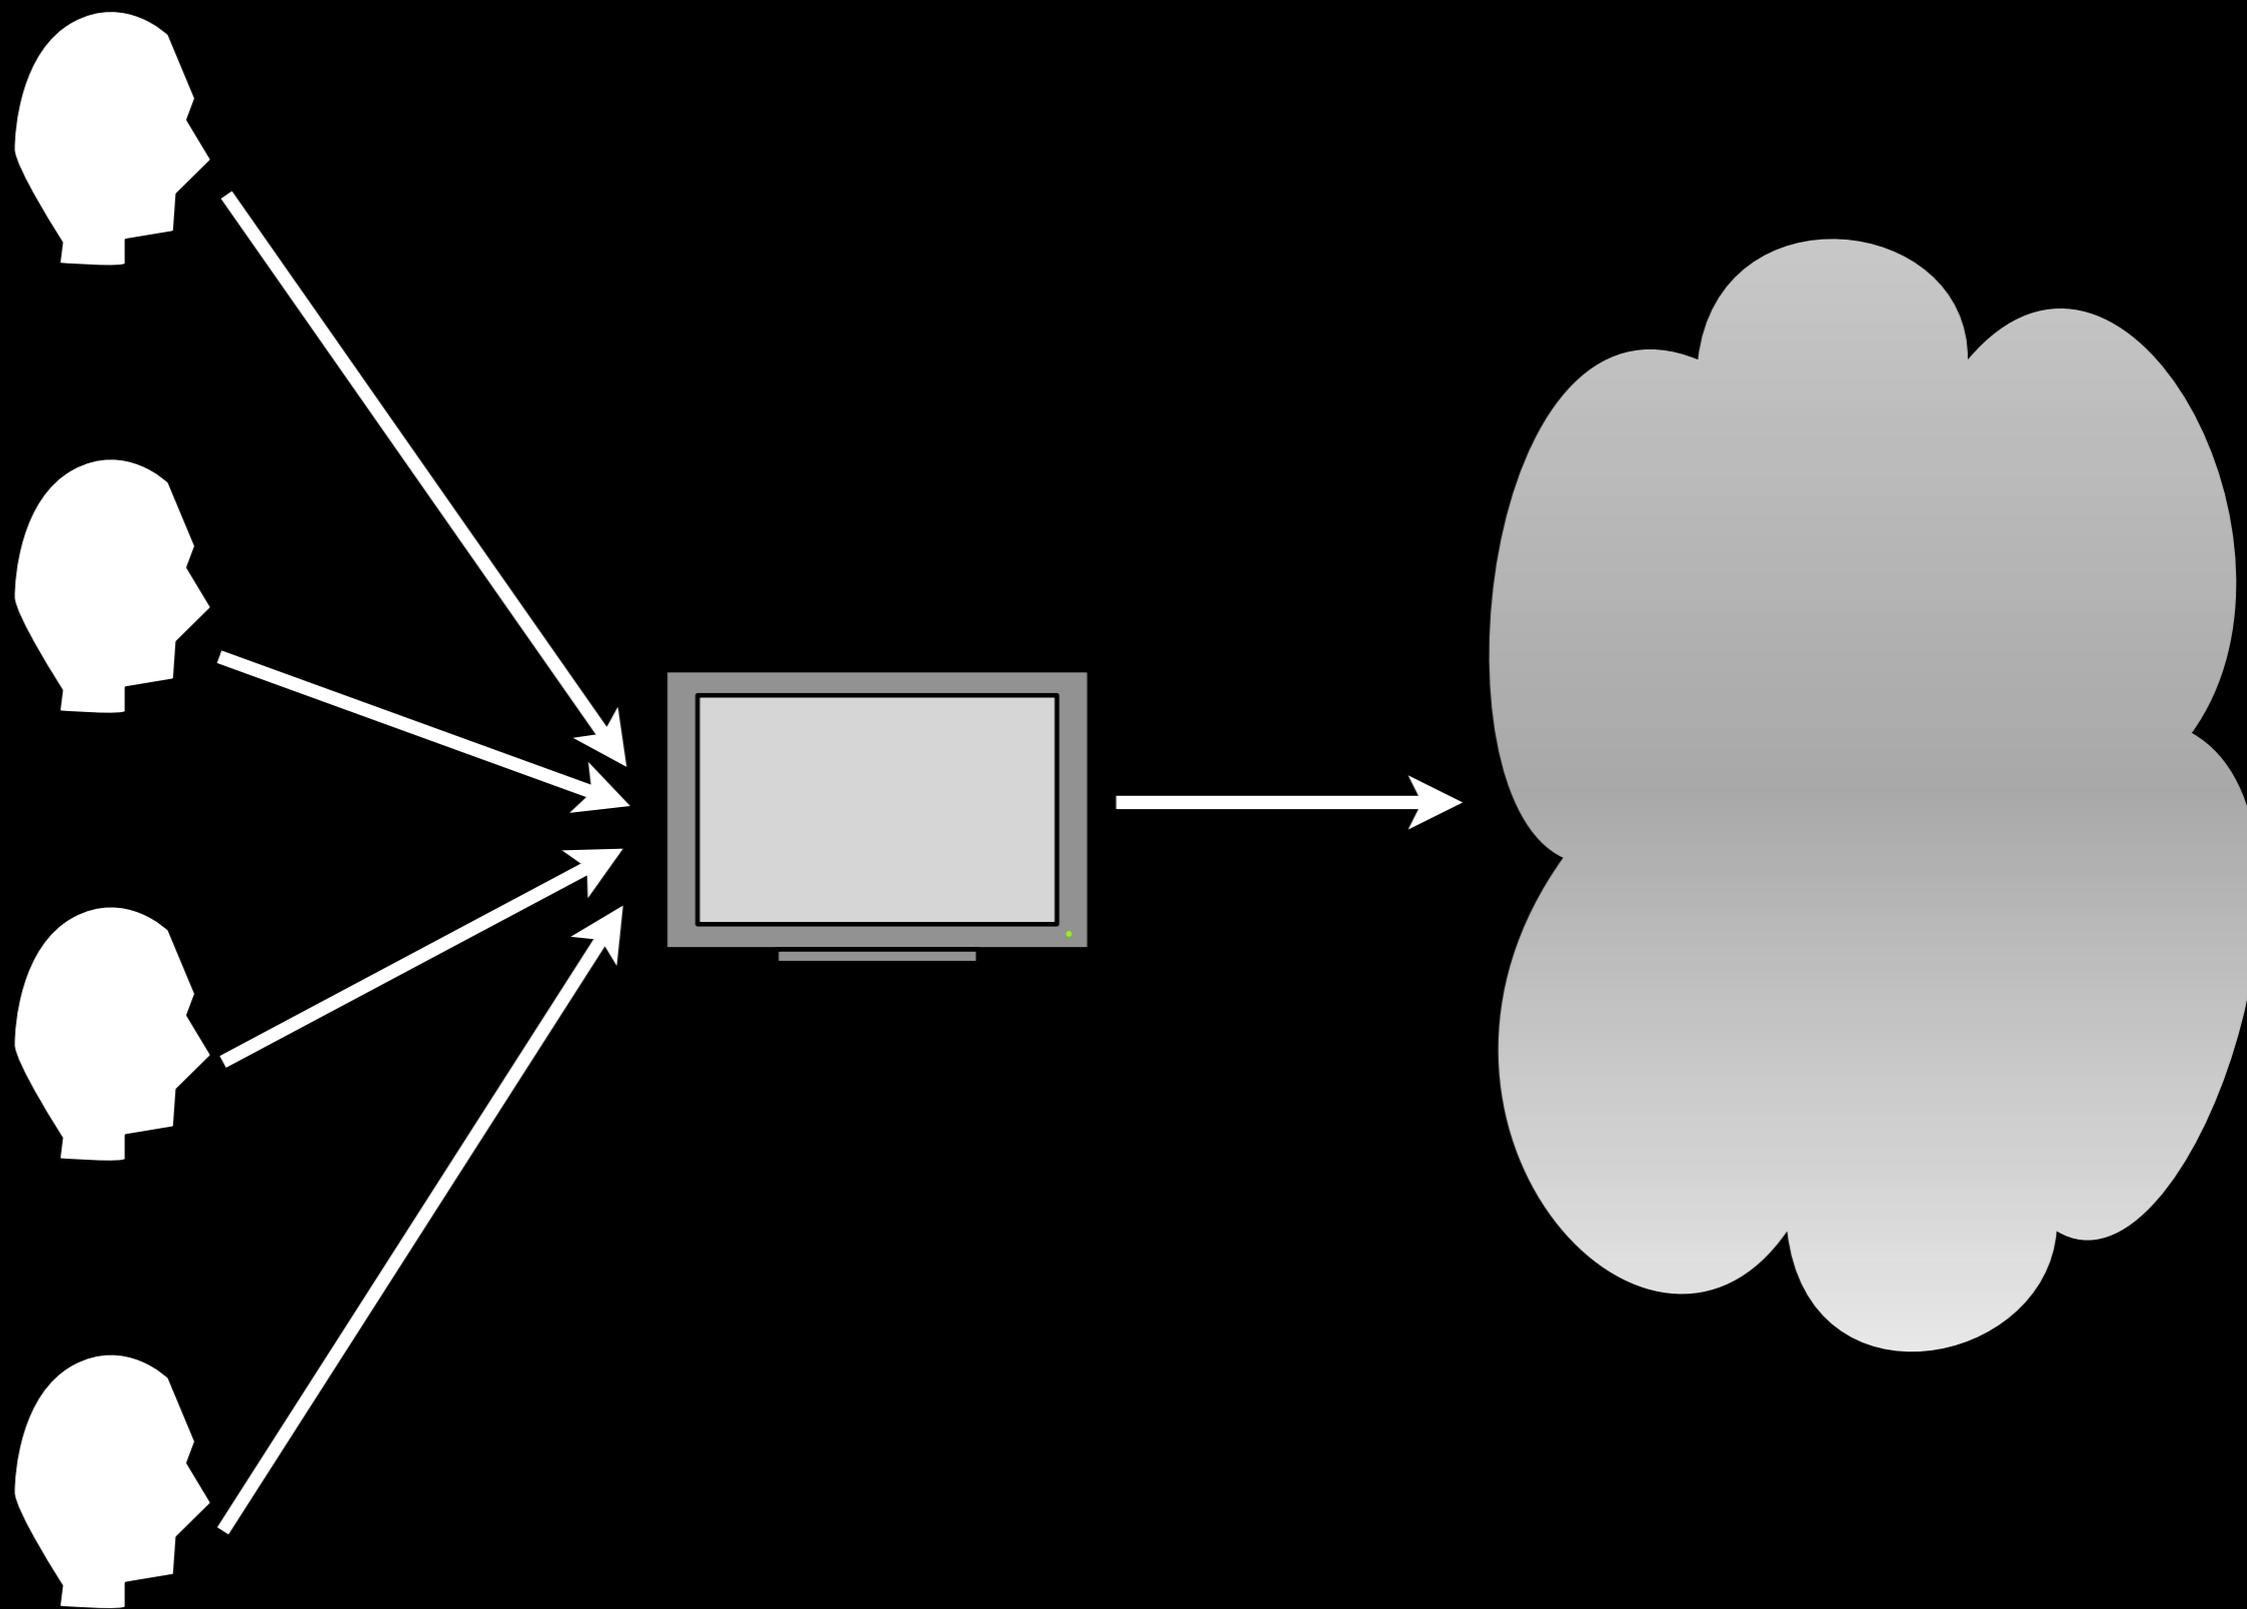
Map2[1] += 1

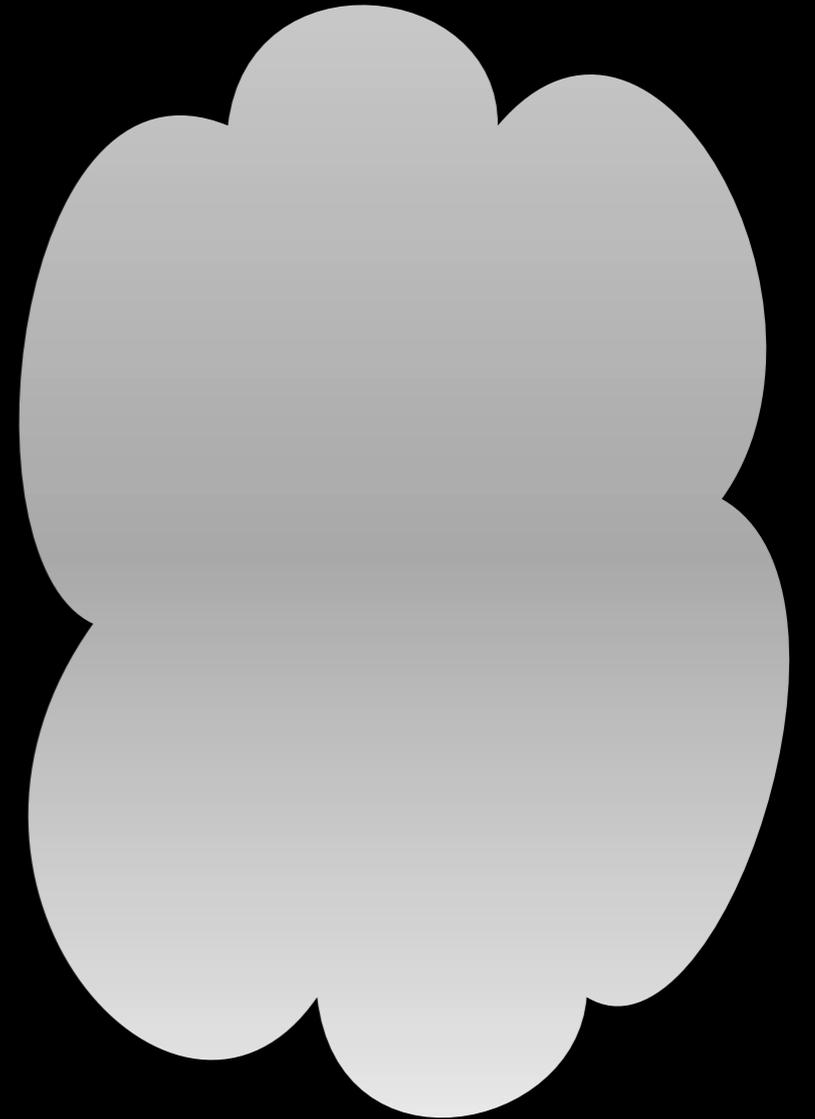
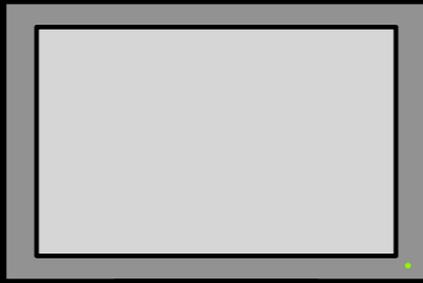
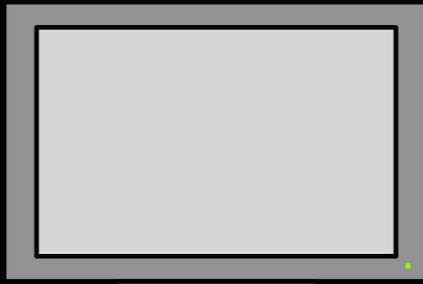
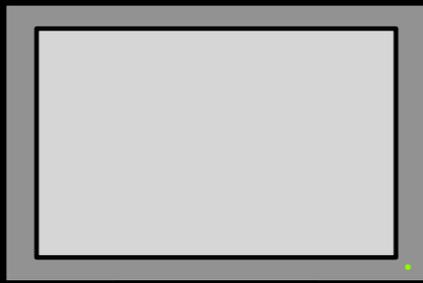
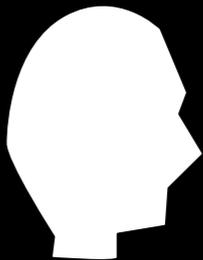
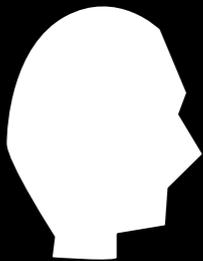
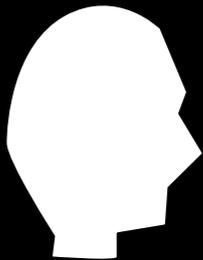


Map1[1] += 1

Map1	<1>→1
Map2	<1>→1
Map3	<1,1>→1
Map4	<1>→1
Map5	<1>→1
Result	??







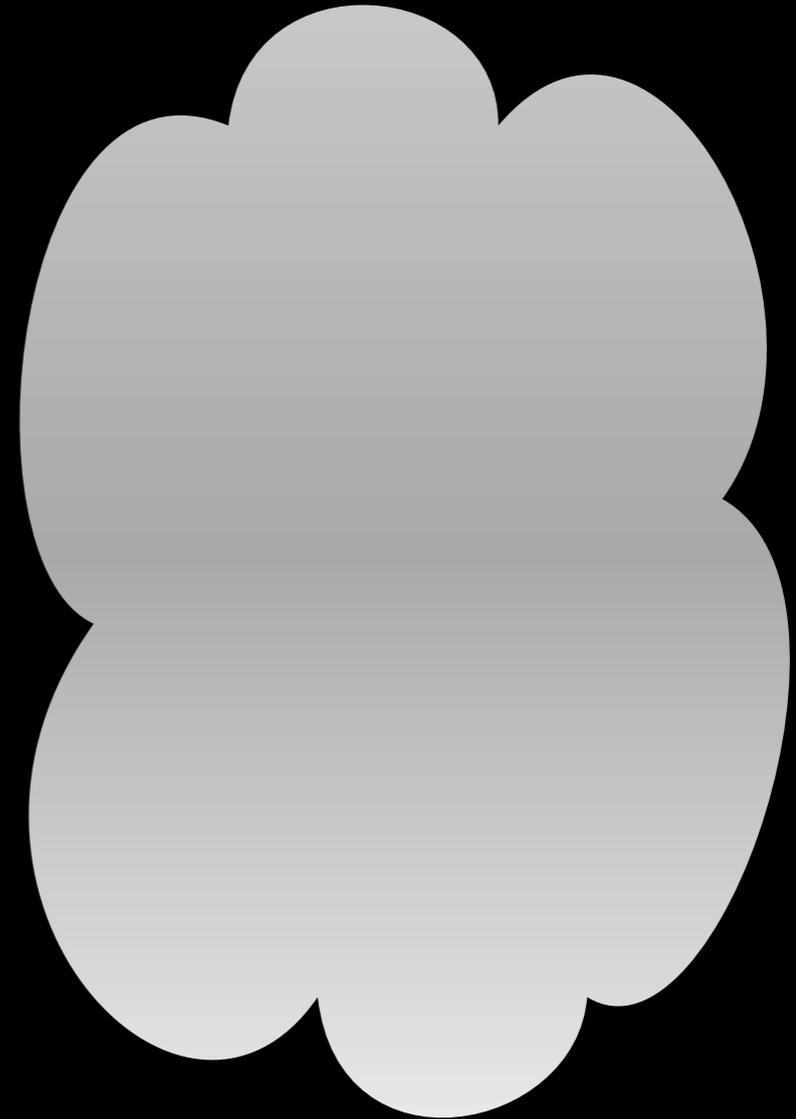
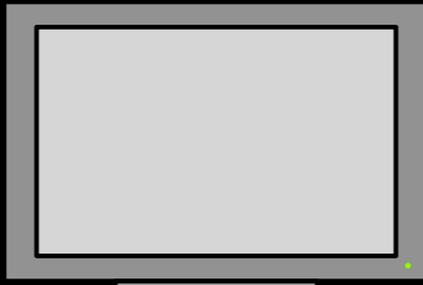
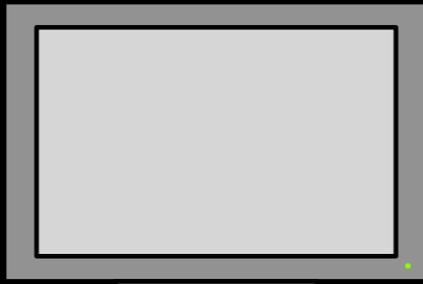
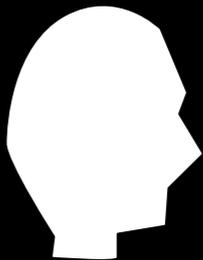
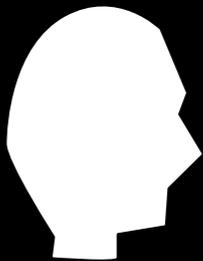
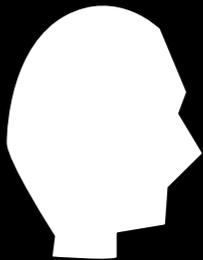
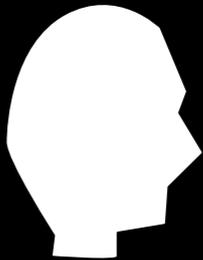
“Magic Maps”

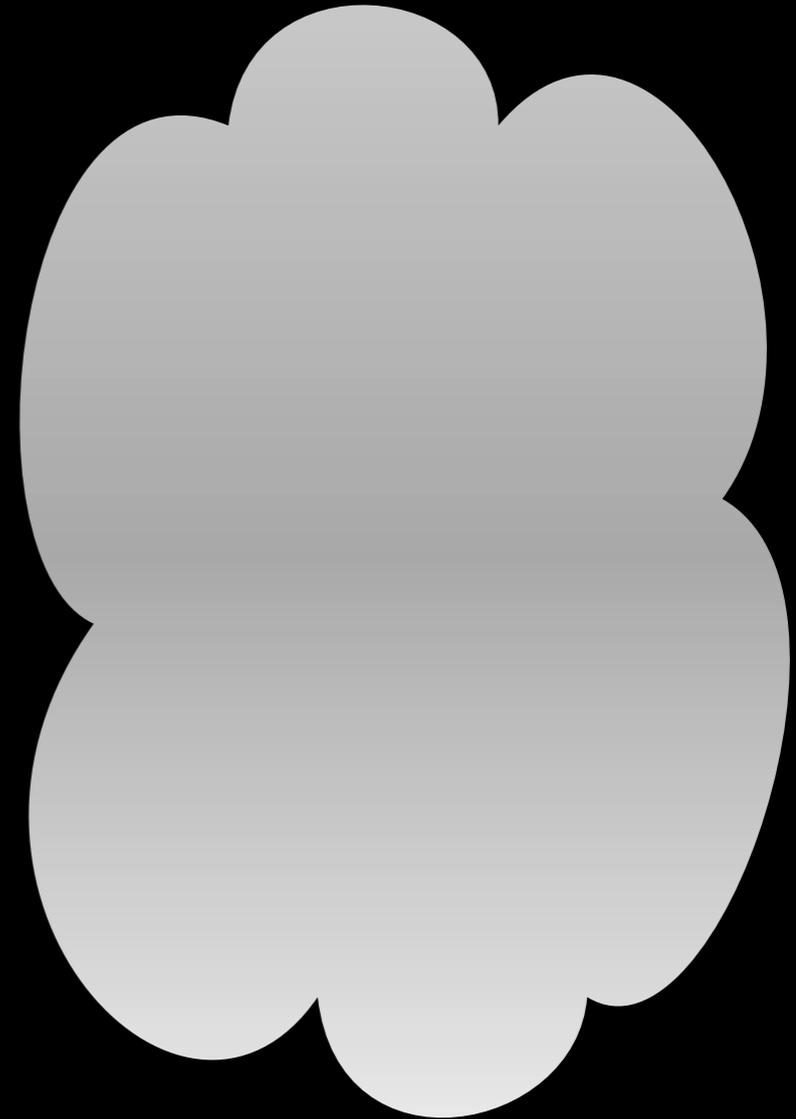
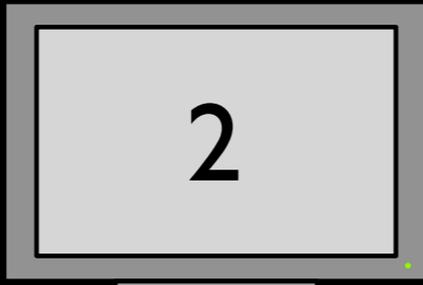
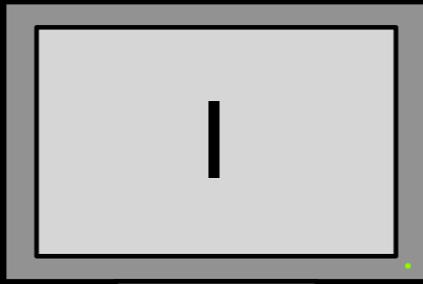
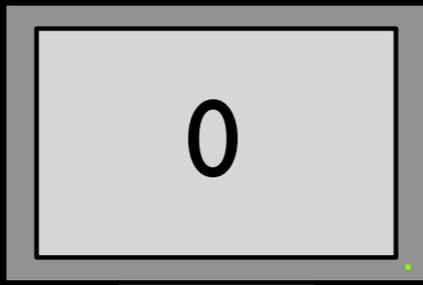
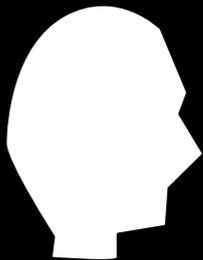
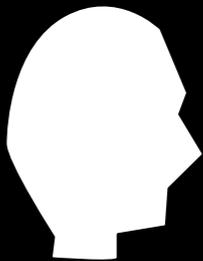
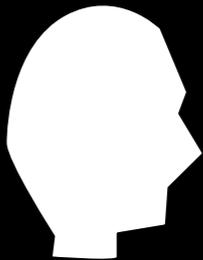
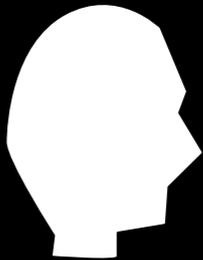
Map Maintenance Messages

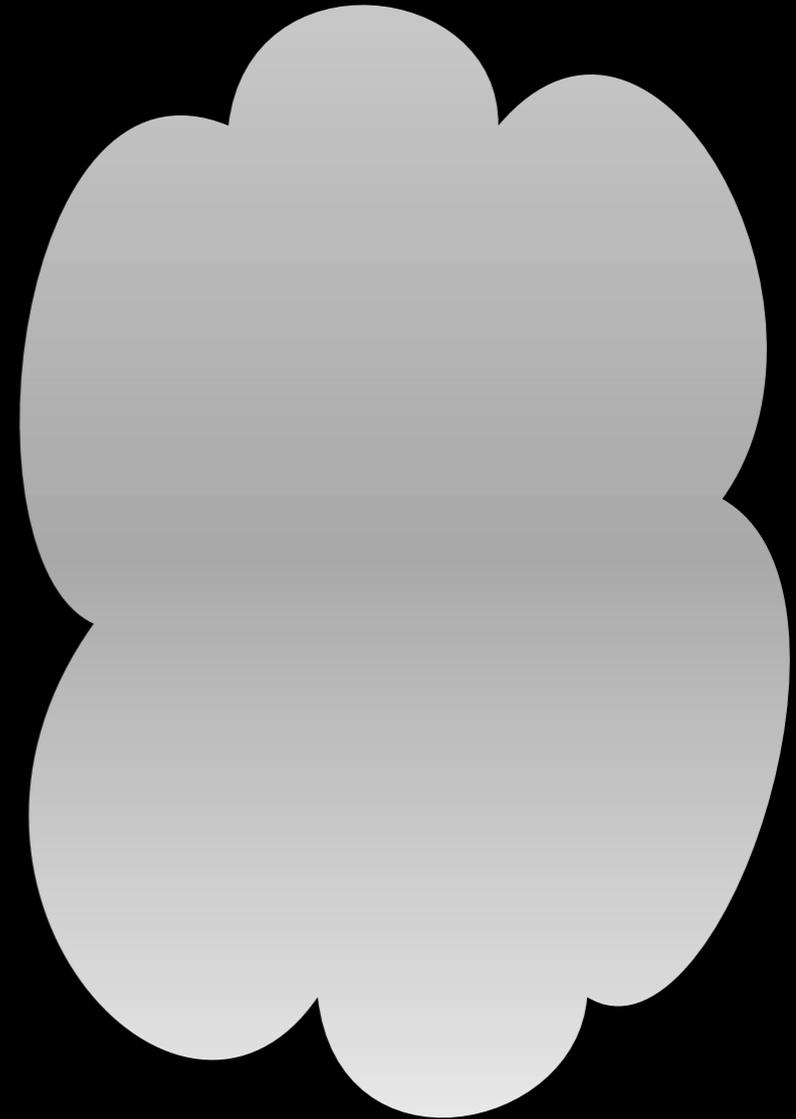
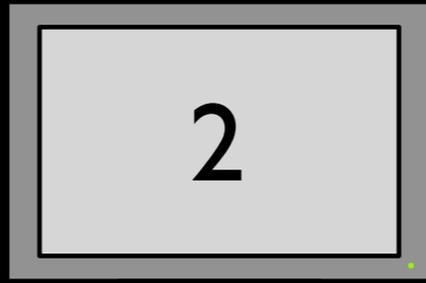
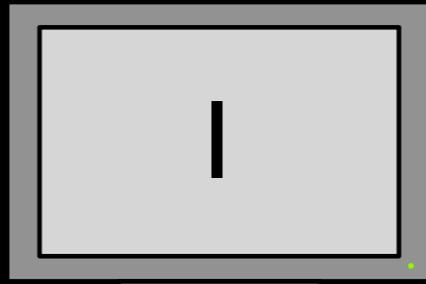
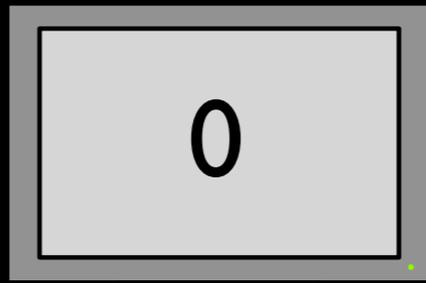
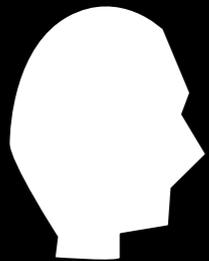
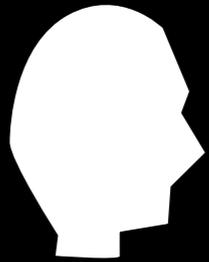
System Design



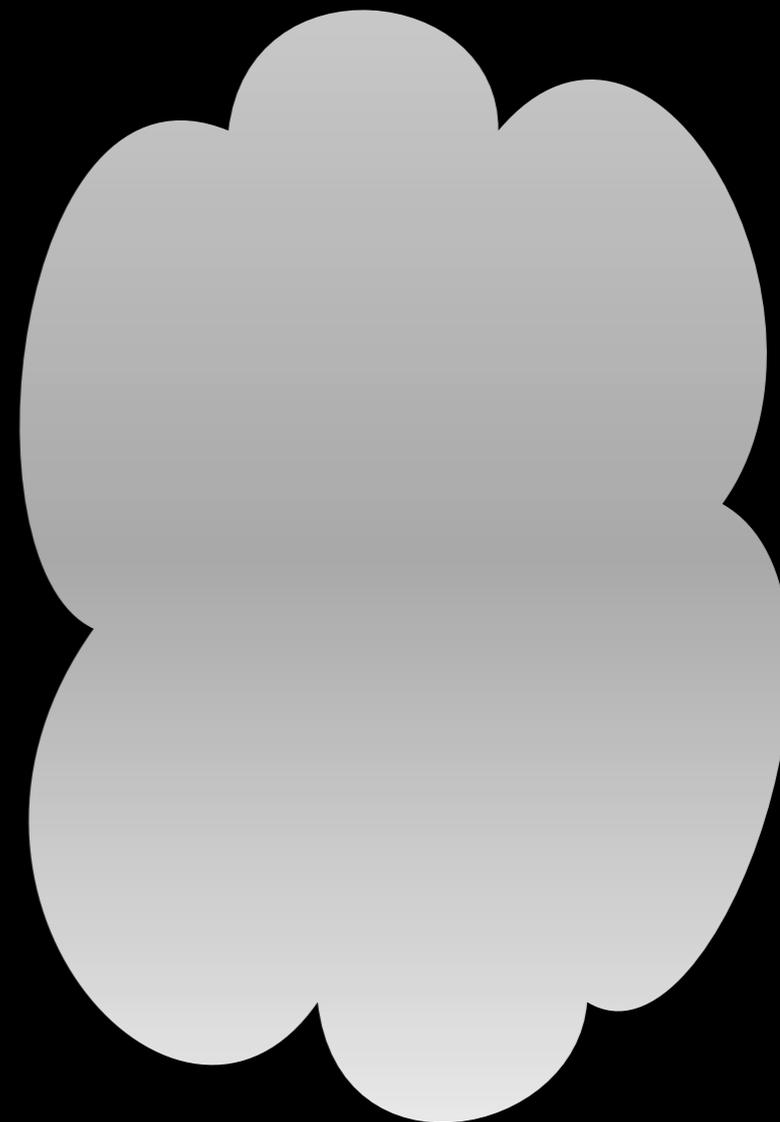
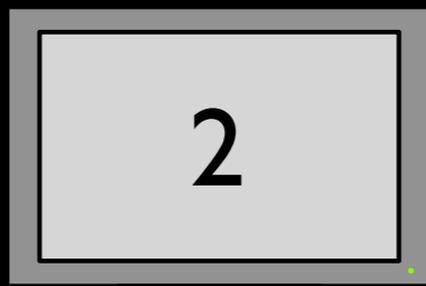
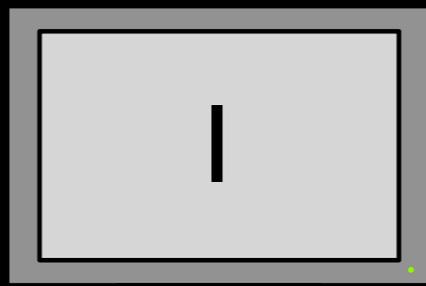
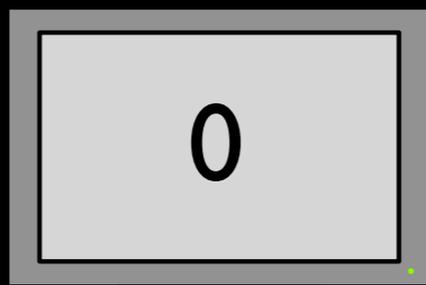
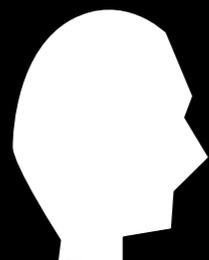
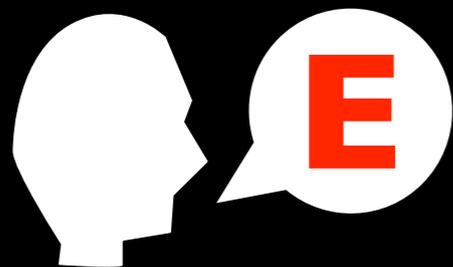
Hybrid Consistency



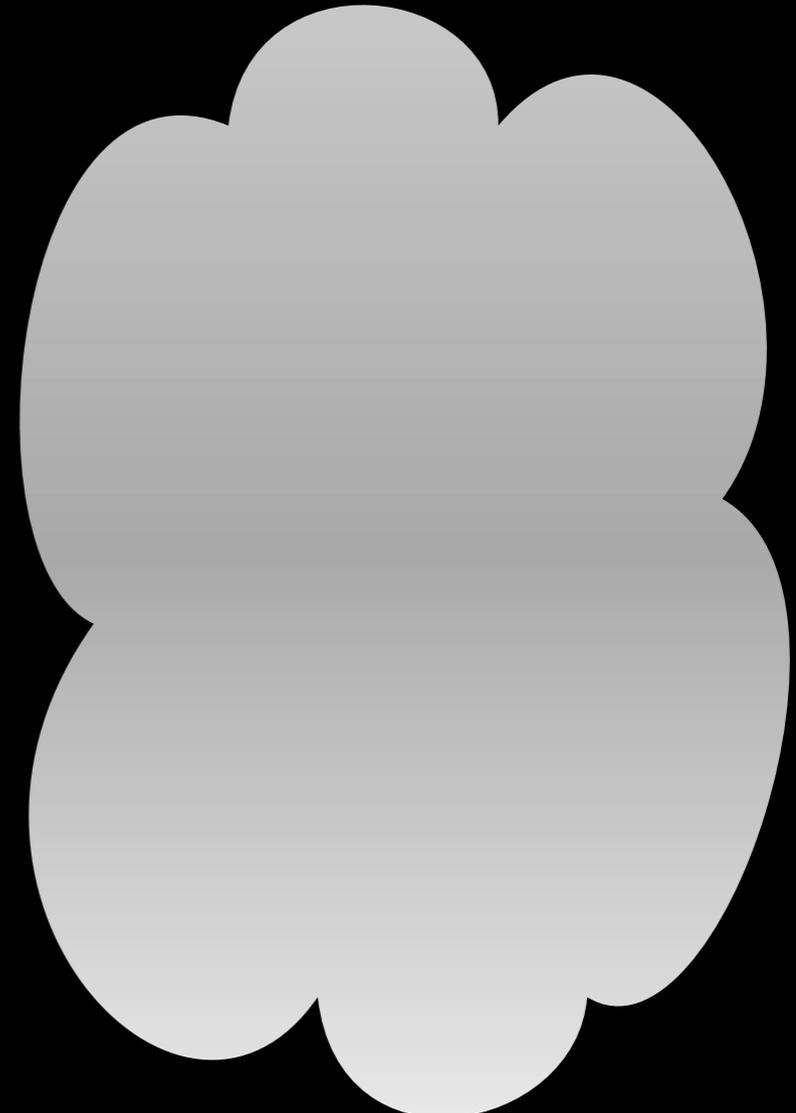
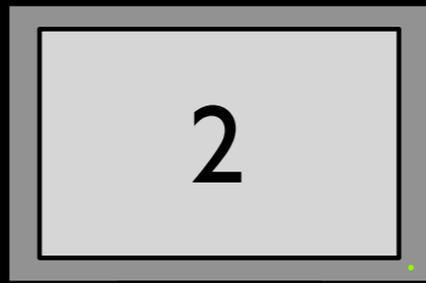
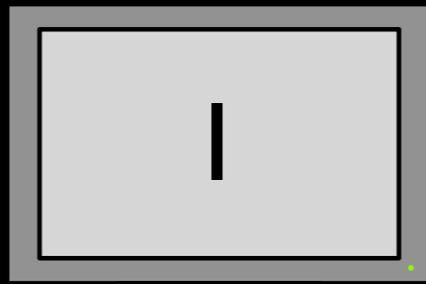
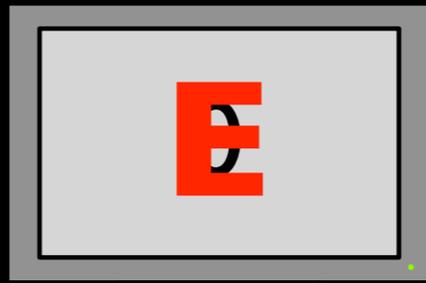
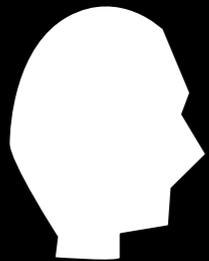
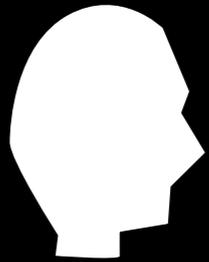




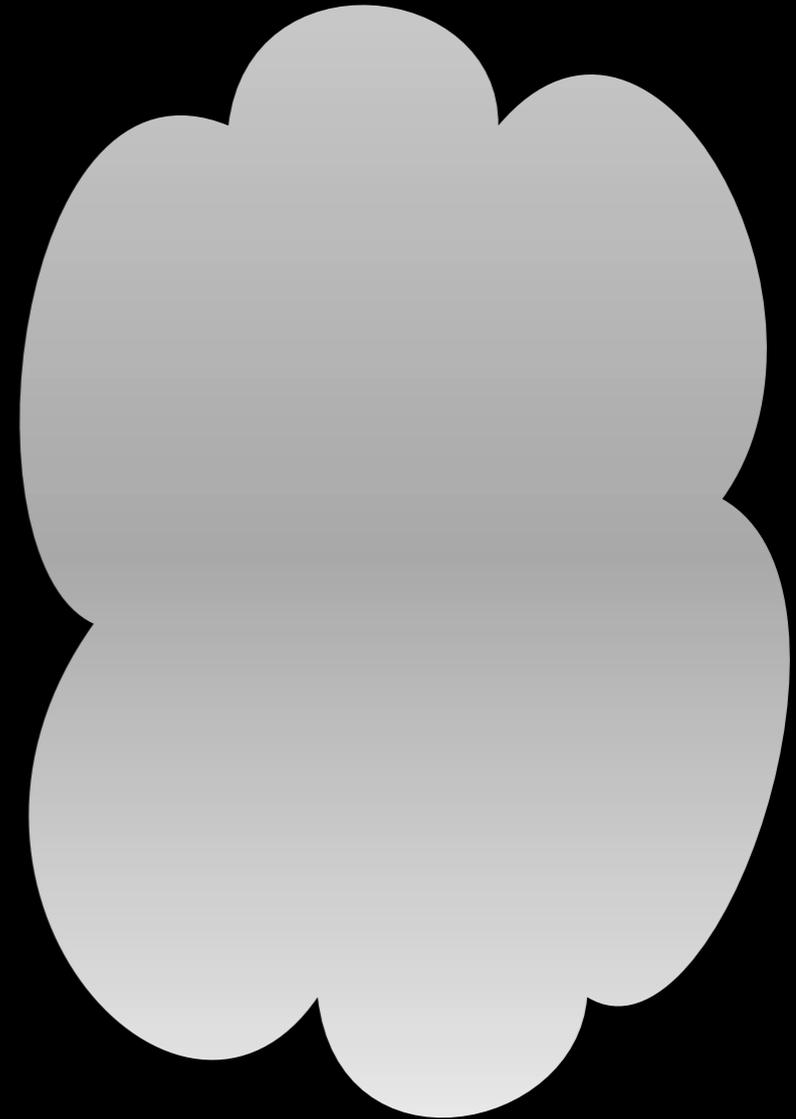
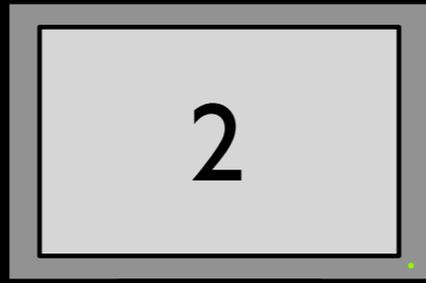
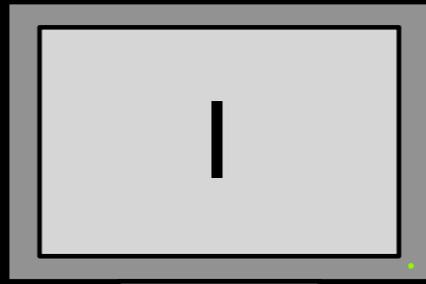
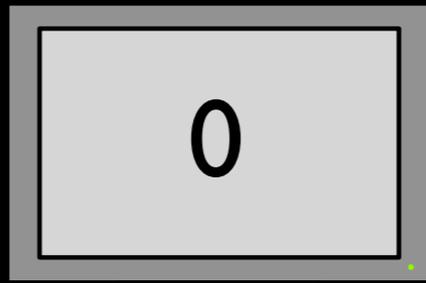
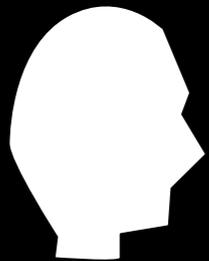
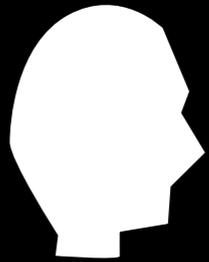
Epoch: 0



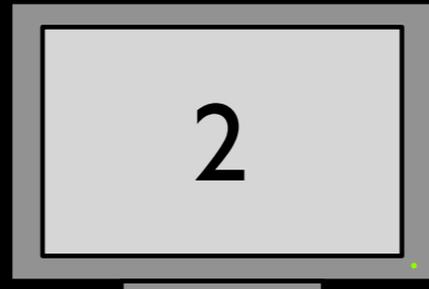
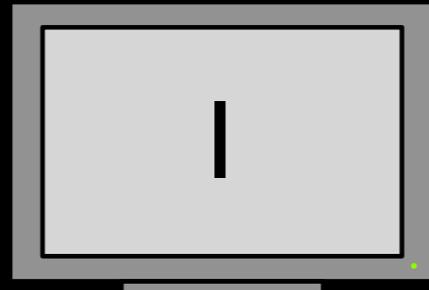
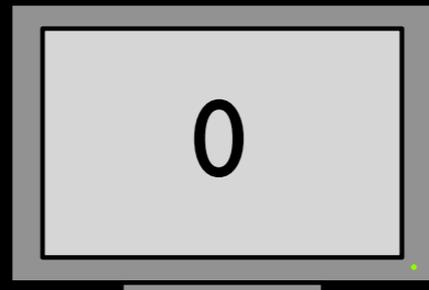
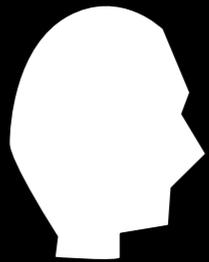
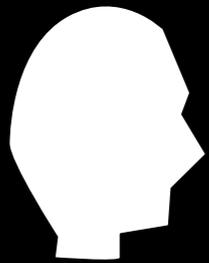
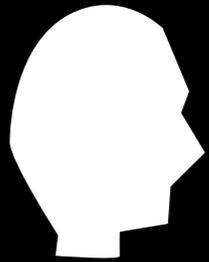
Epoch: 0



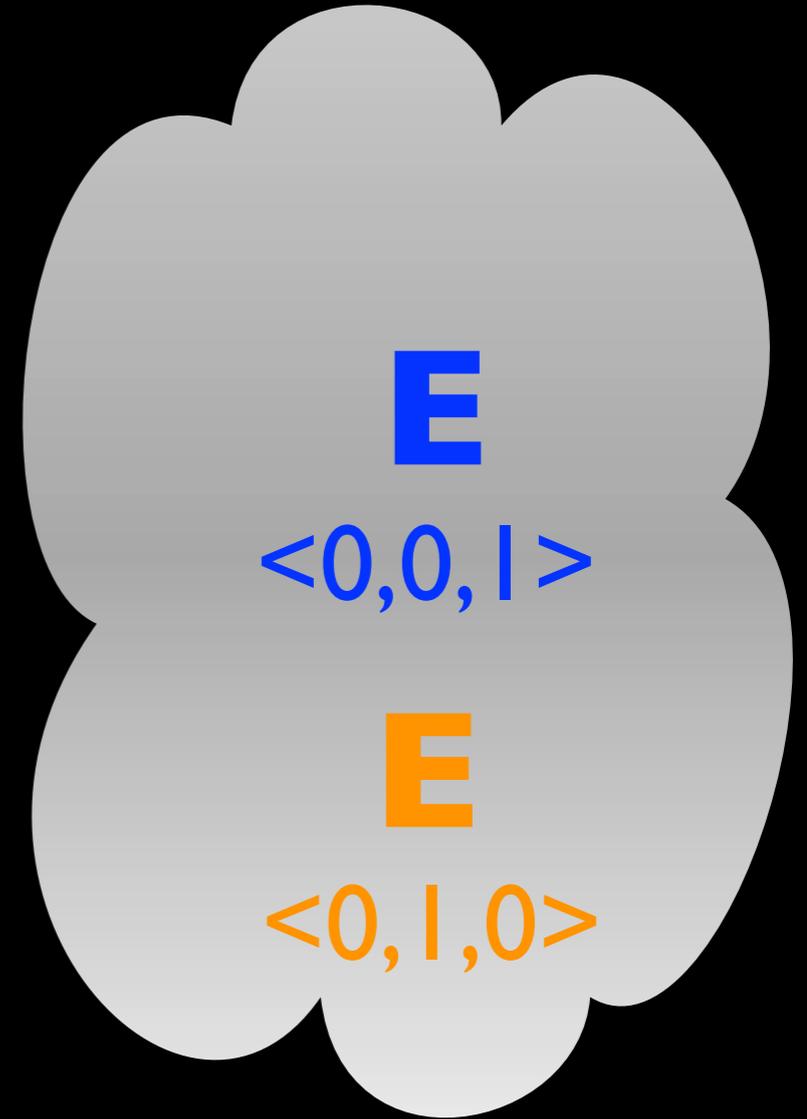
Epoch: 0



Epoch: 0



Epoch: 0





E
<0,1,3>



Map₂[1,3] += 2
<0,2,4>



E
<0,1,3>



E
<0,1,3>



E
<0,1,3>



E
<0,1,3>



Map₂[...] +=
Map₃[...] * Map₄[...]

E
<0,1,3>



Map₂[...] +=
Map₃[...] * Map₄[...]



Map₃[...]
Map₄[...]

E
<0,1,3>



Map₂[...] +=
Map₃[...] * Map₄[...]



$\Sigma\delta$

Map₃[...]
Map₄[...]



Map₂[...] += ...
<0,1,3>



Map₃[...] * Map₄[...]

E
<0,1,3>



Map₂[...] +=
Map₃[...] * Map₄[...]



$\Sigma\delta$

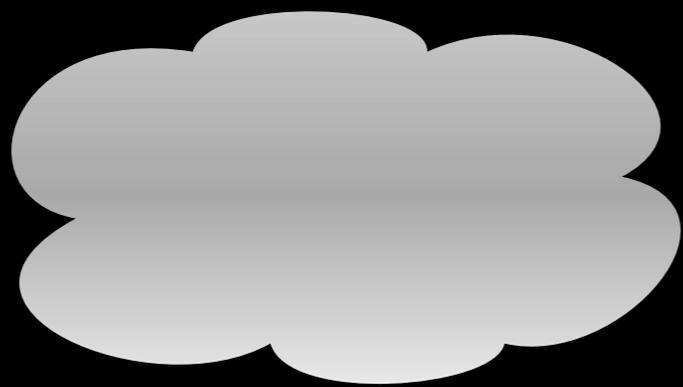
Map₃[...]
Map₄[...]

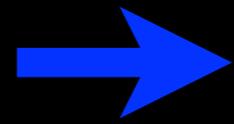
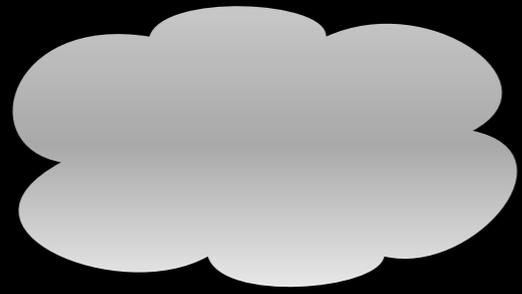


Map₃[...] * Map₄[...]

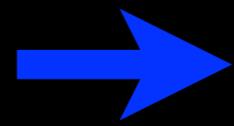
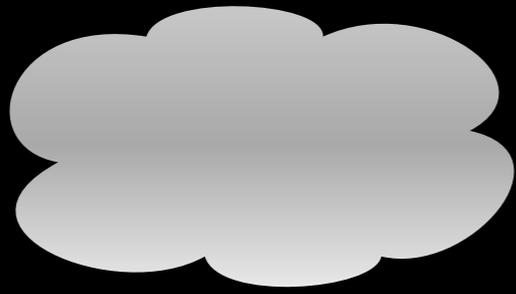


Map₂[...] += ...
<0,1,3>

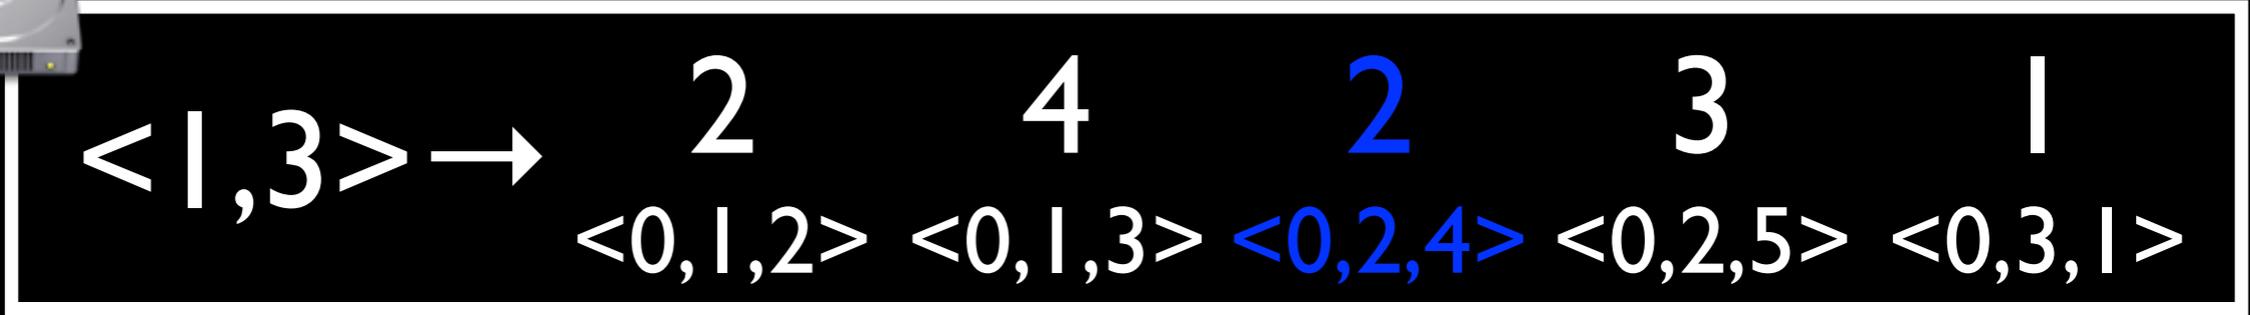


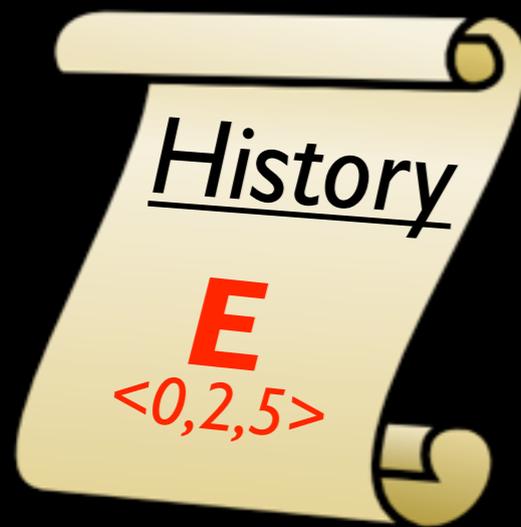
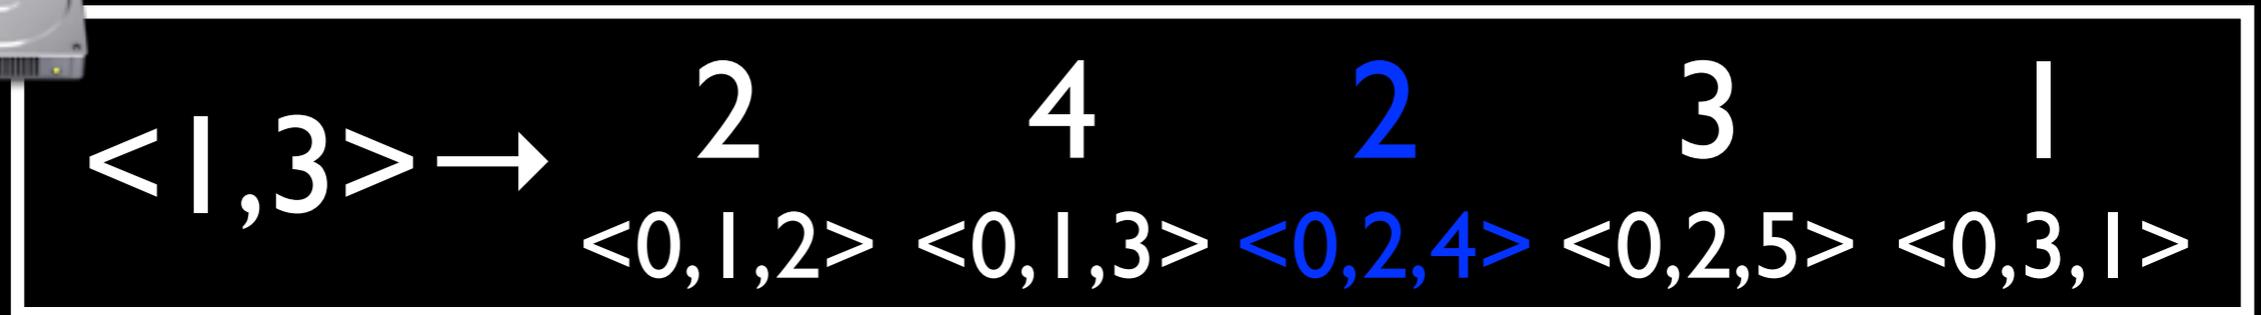


Map₂[1,3] += 2
<0,2,4>

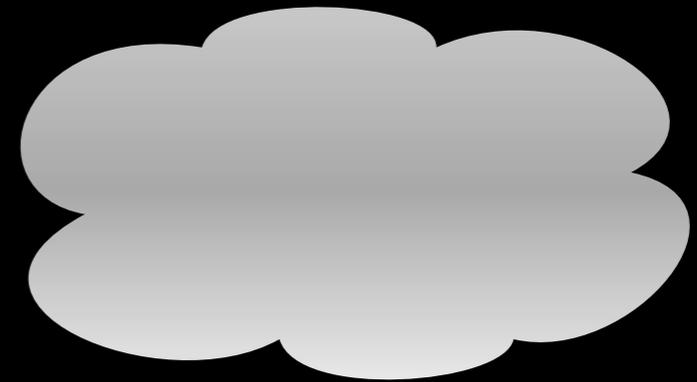


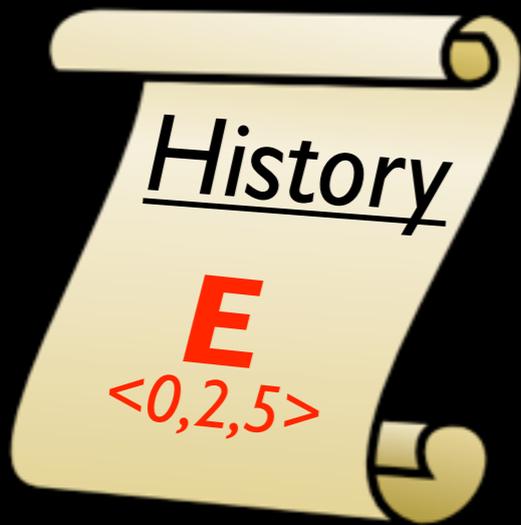
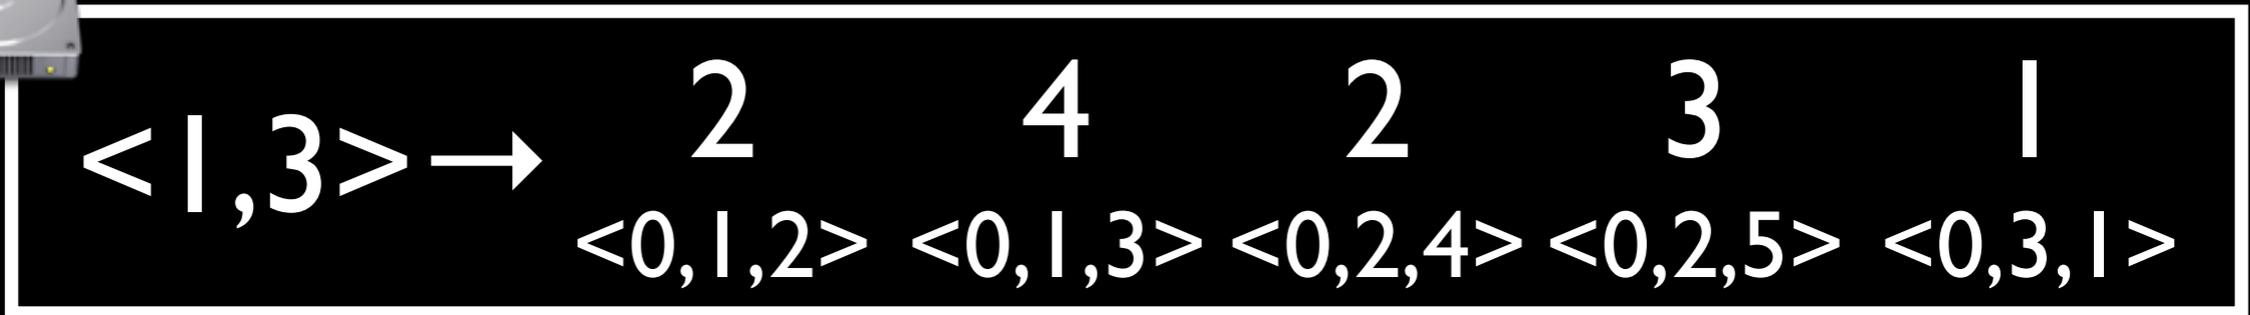
Map₂[1,3] += 2
<0,2,4>

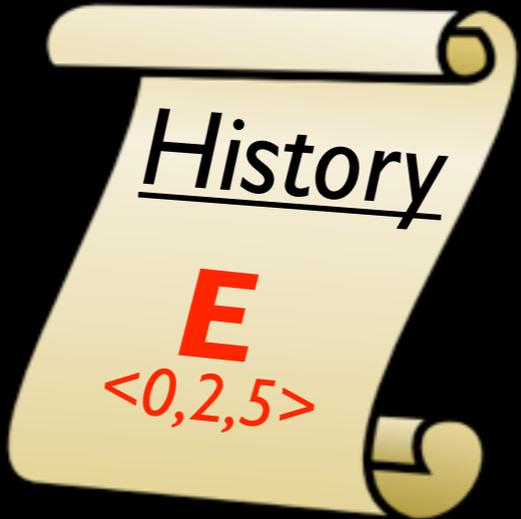
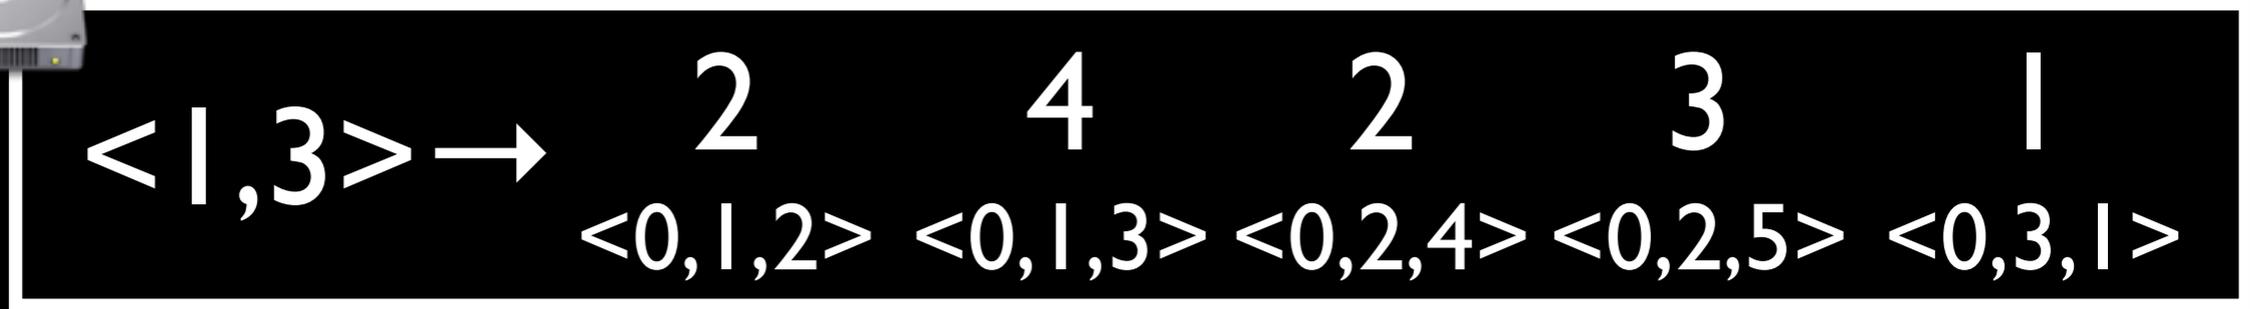




δ
<0,2,5>



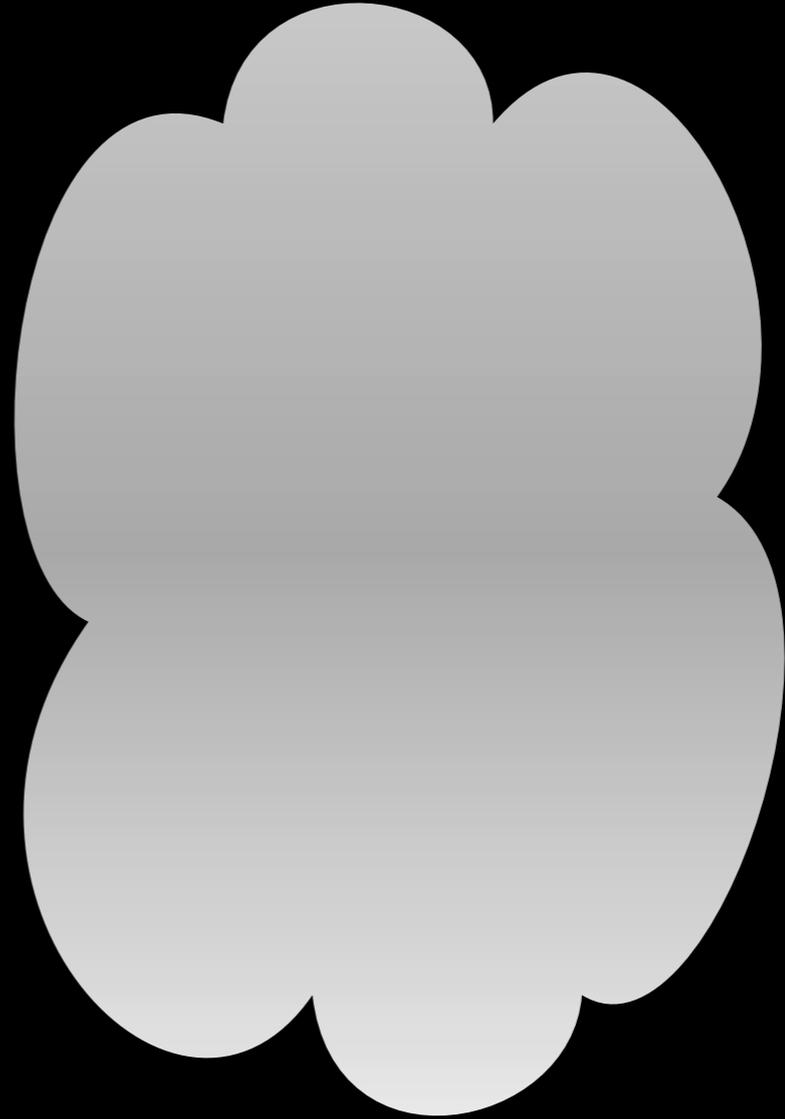




Coord.

Coord.

Coord.



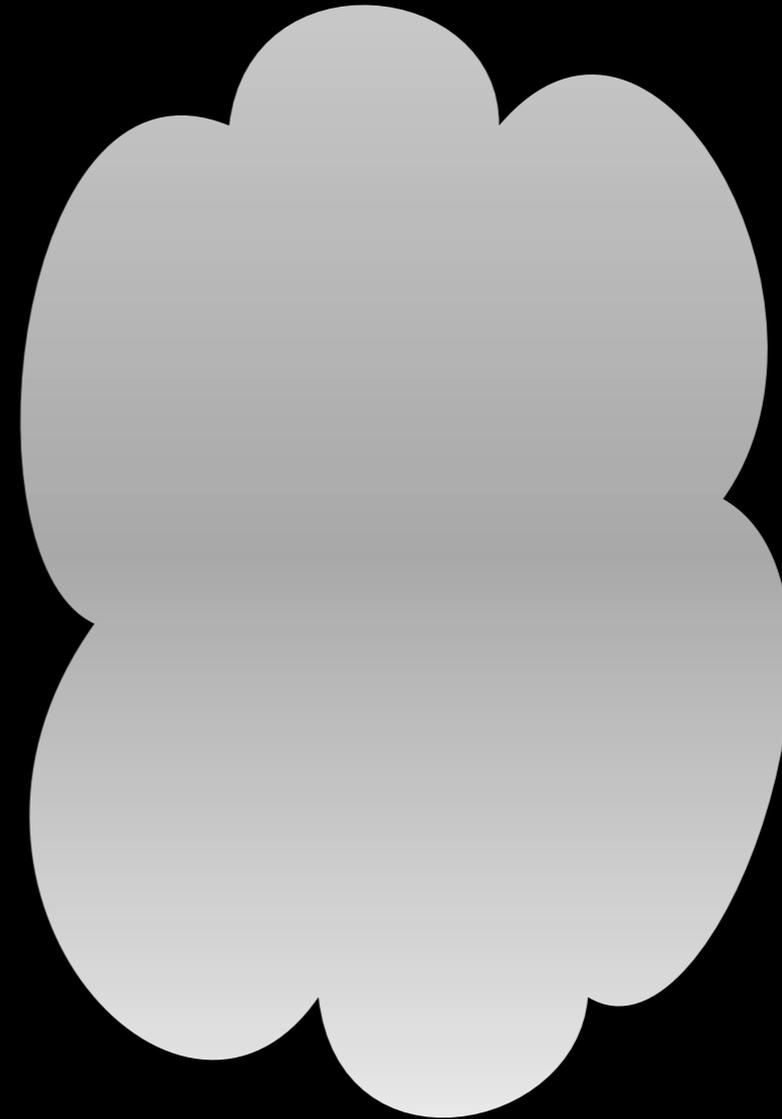
Coord.

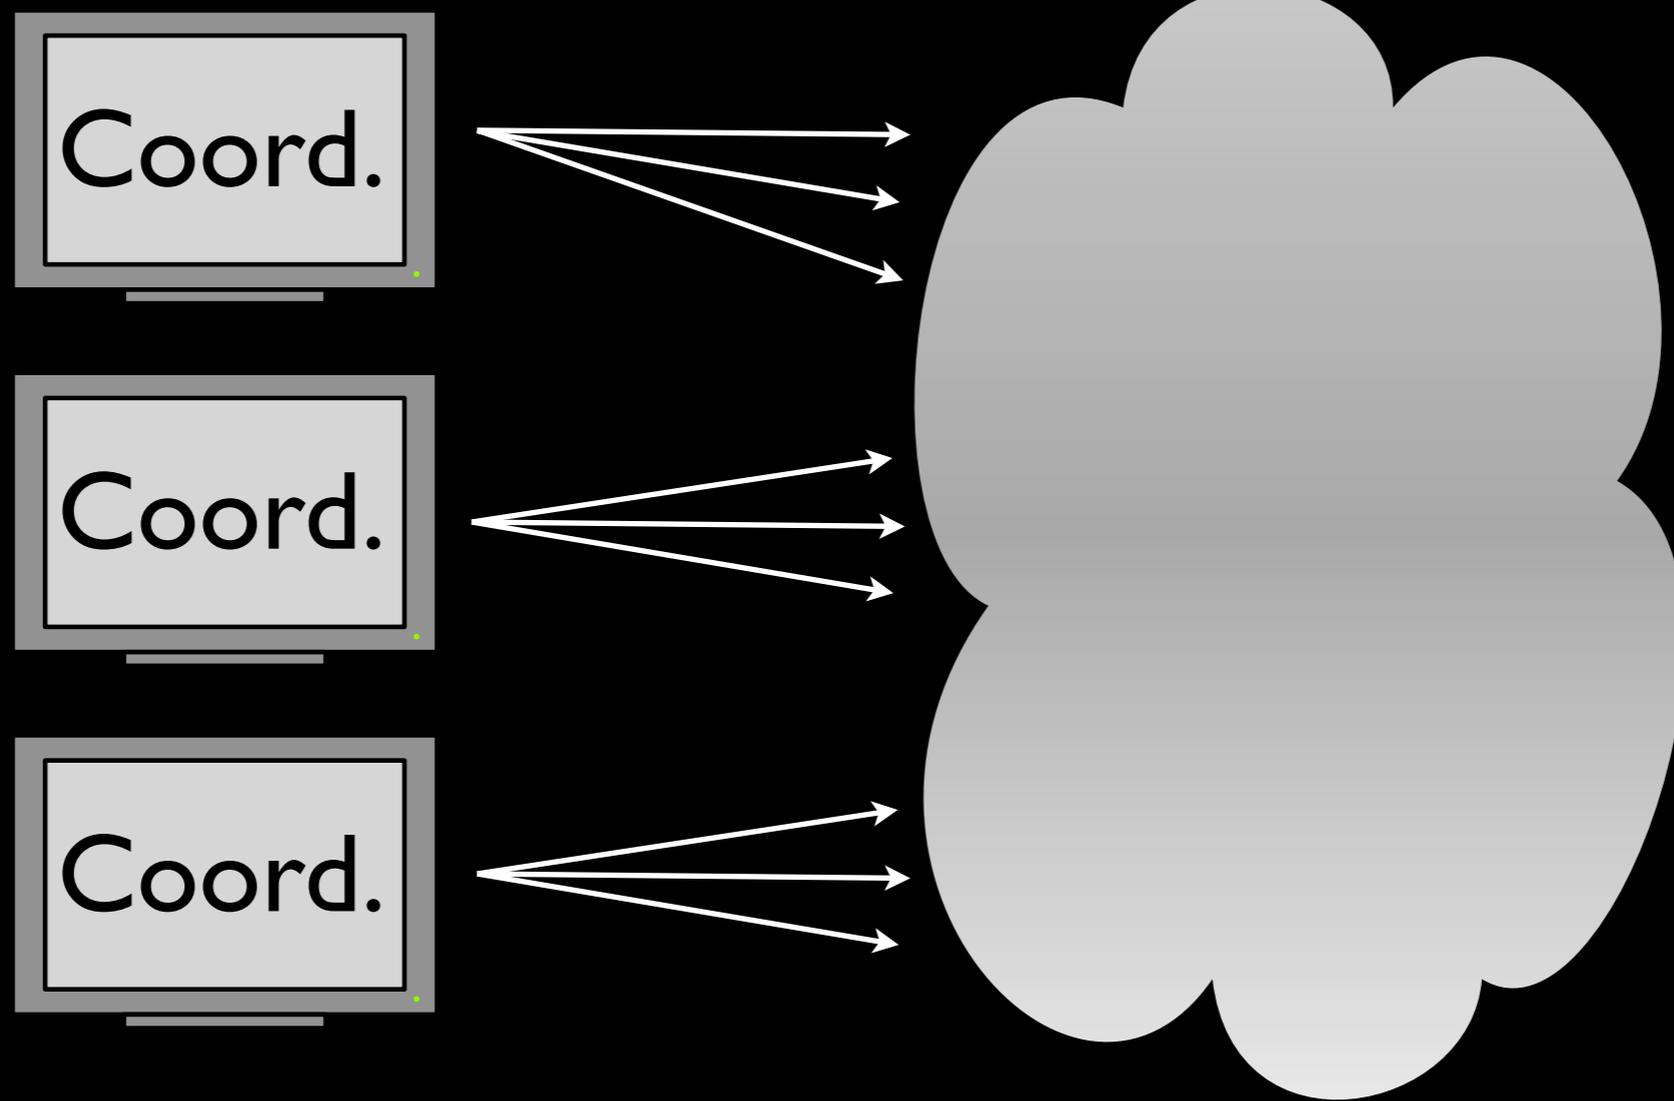


Coord.

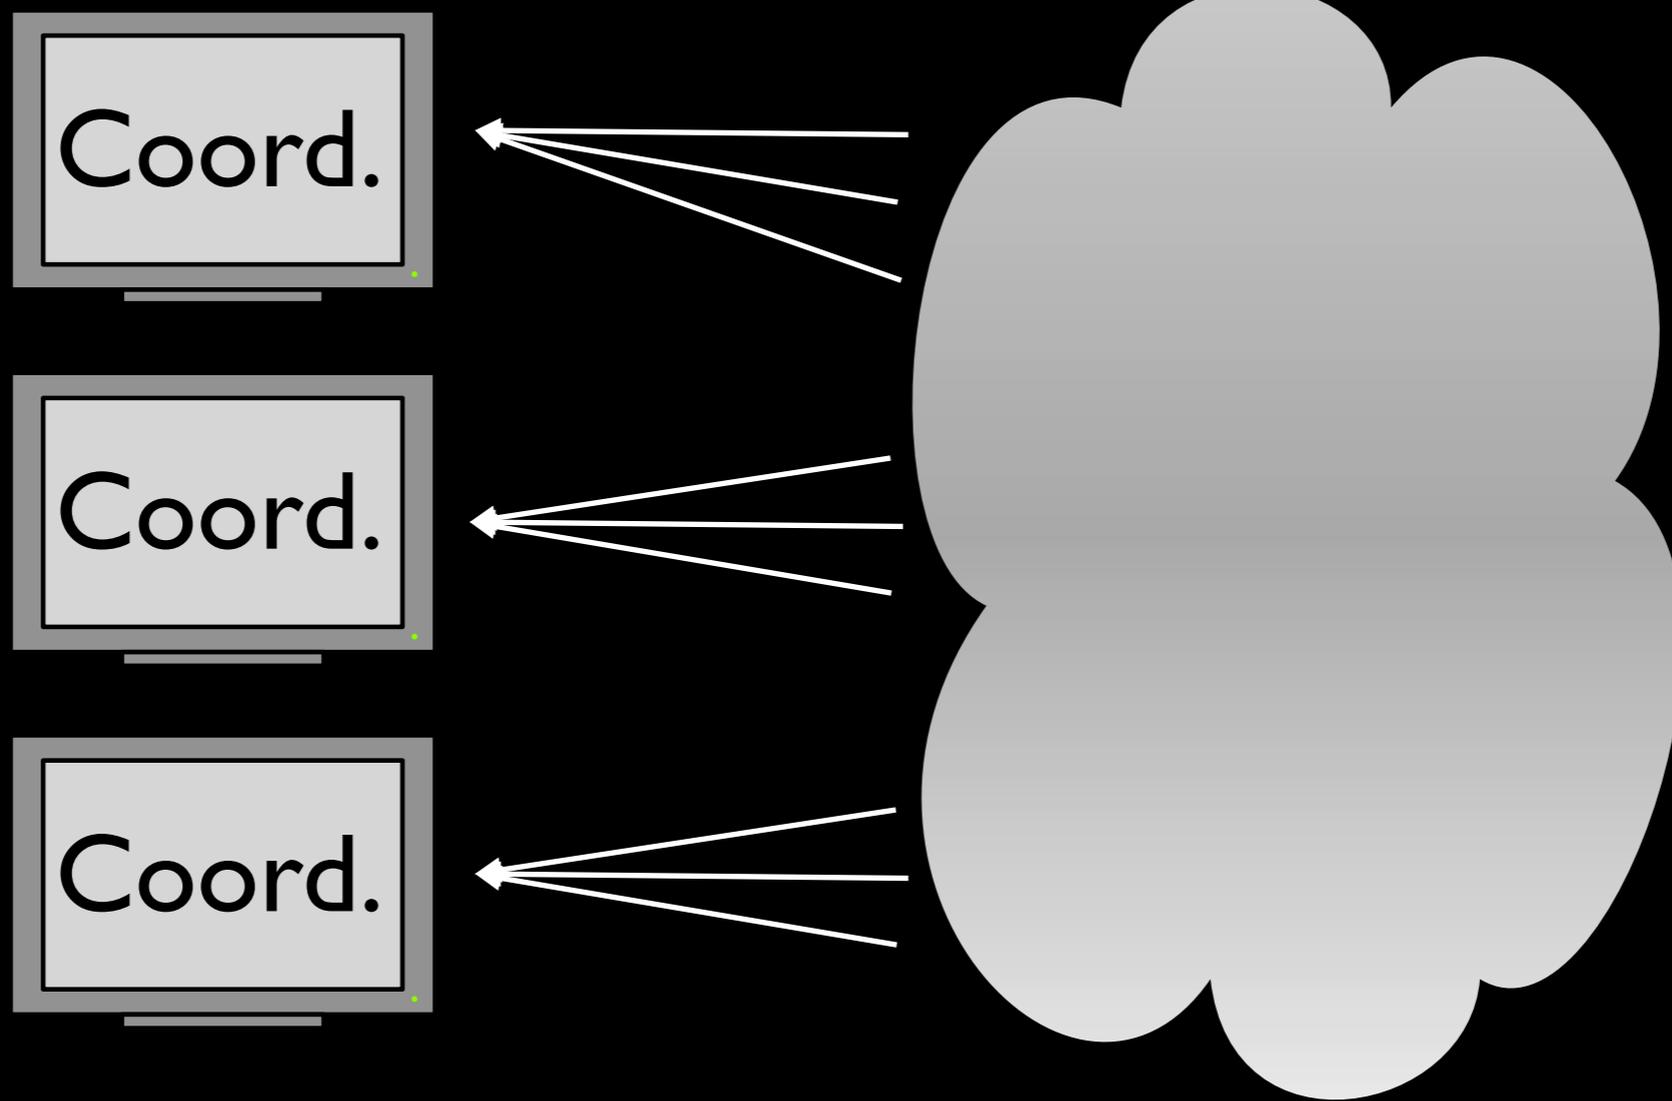


Coord.

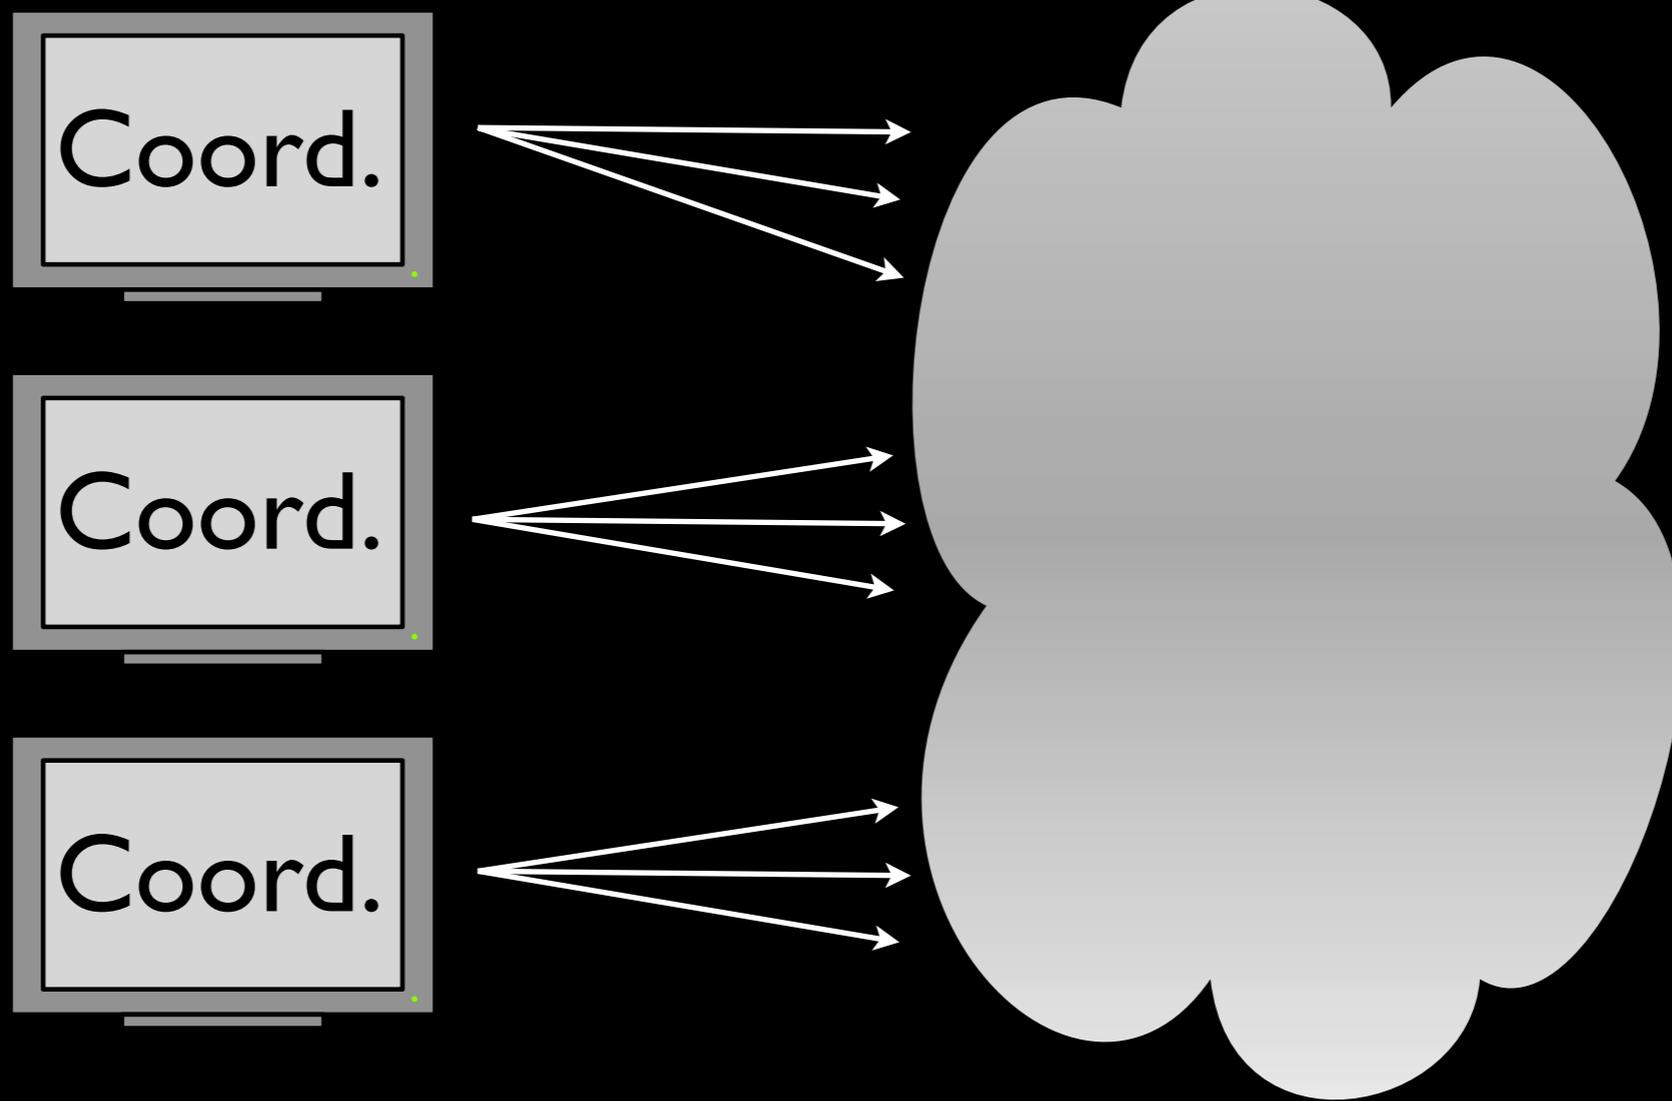




Epoch Commit



Commit OK

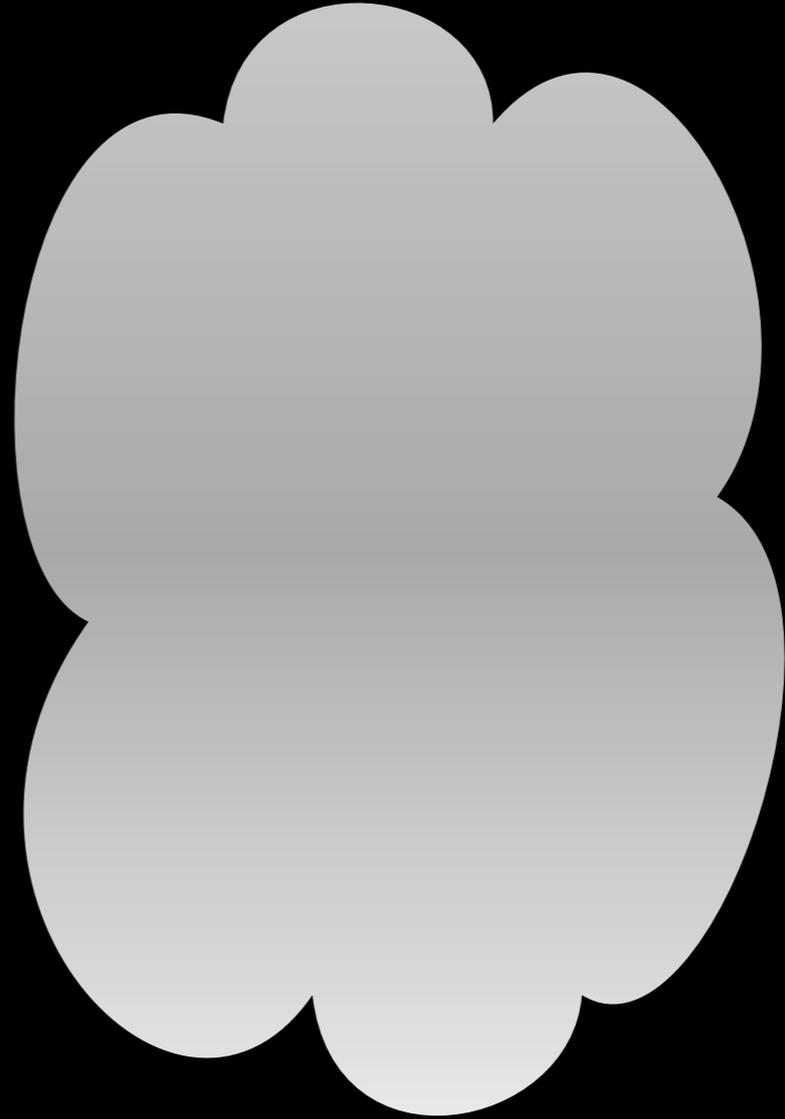


Epoch End

Coord.

Coord.

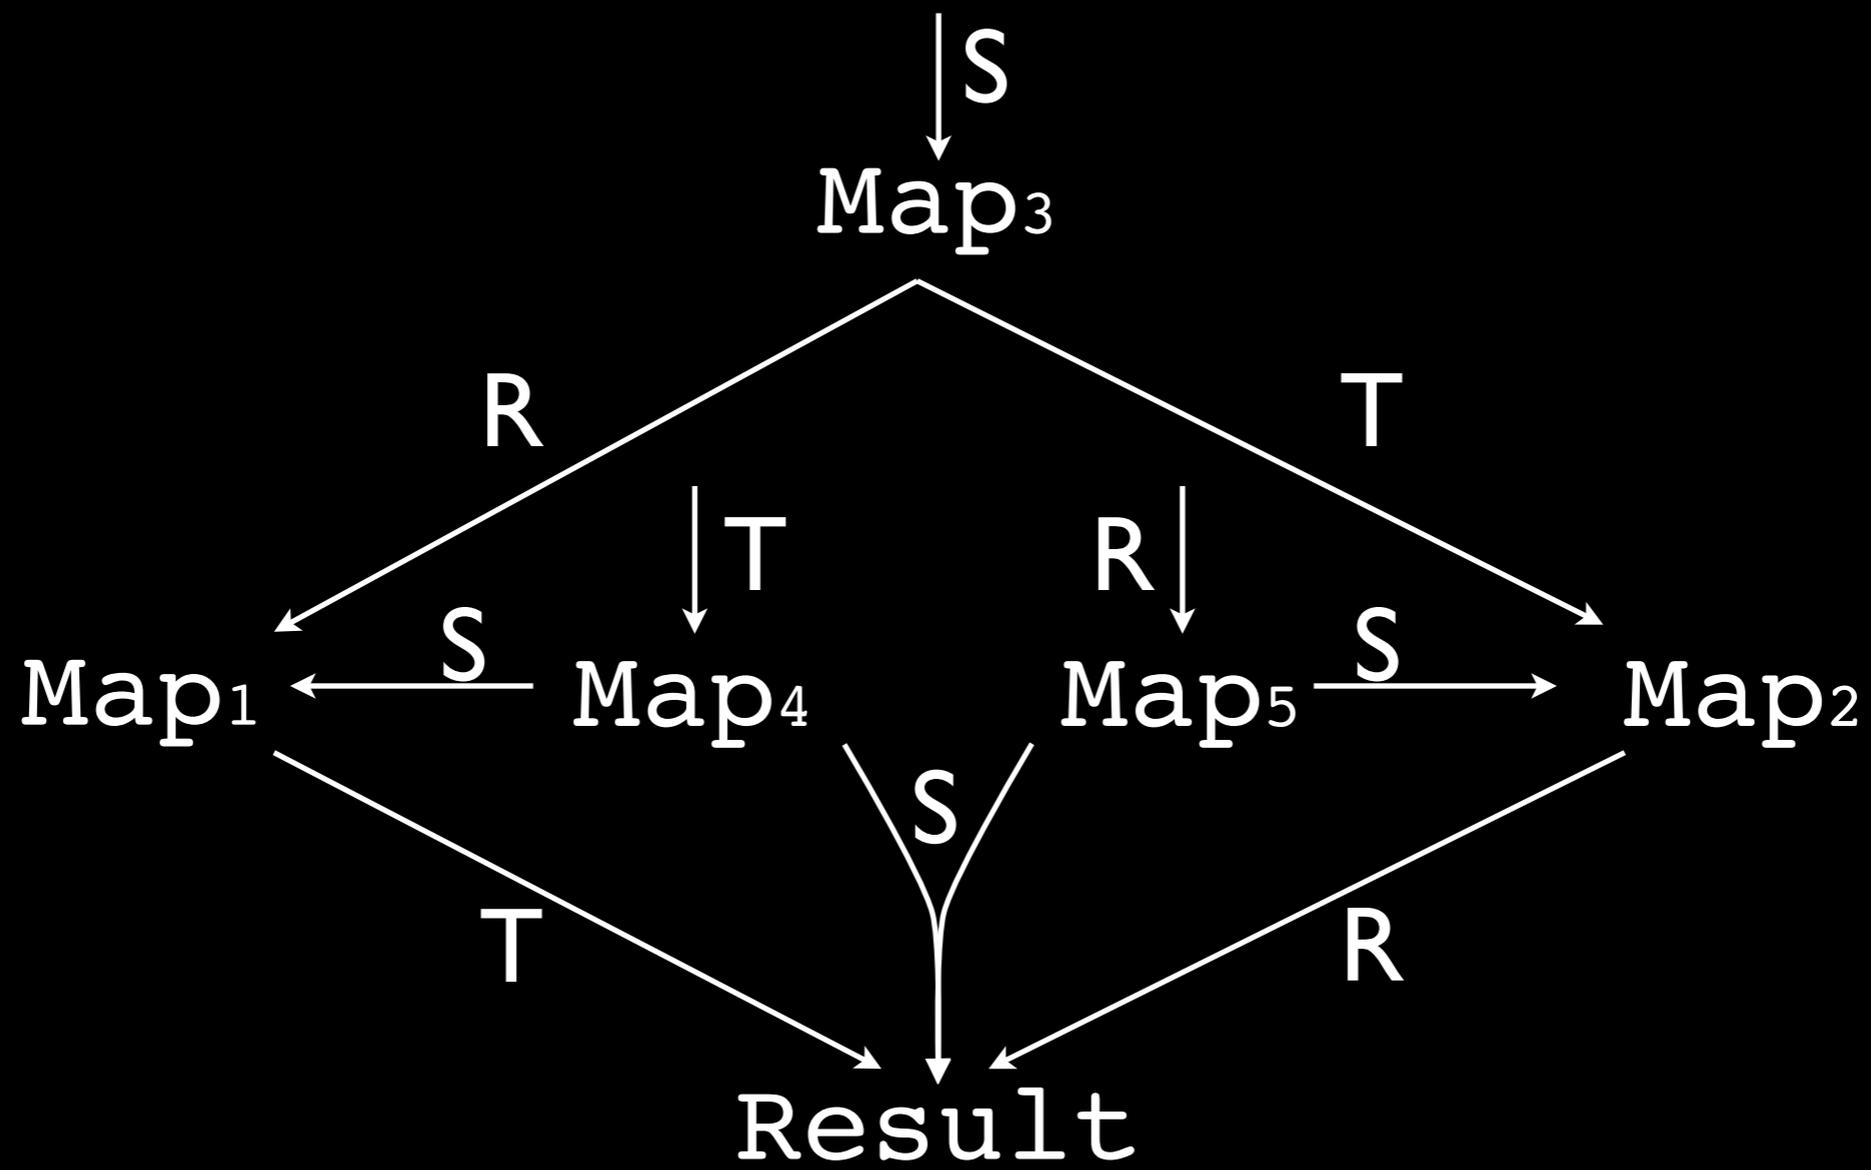
Coord.

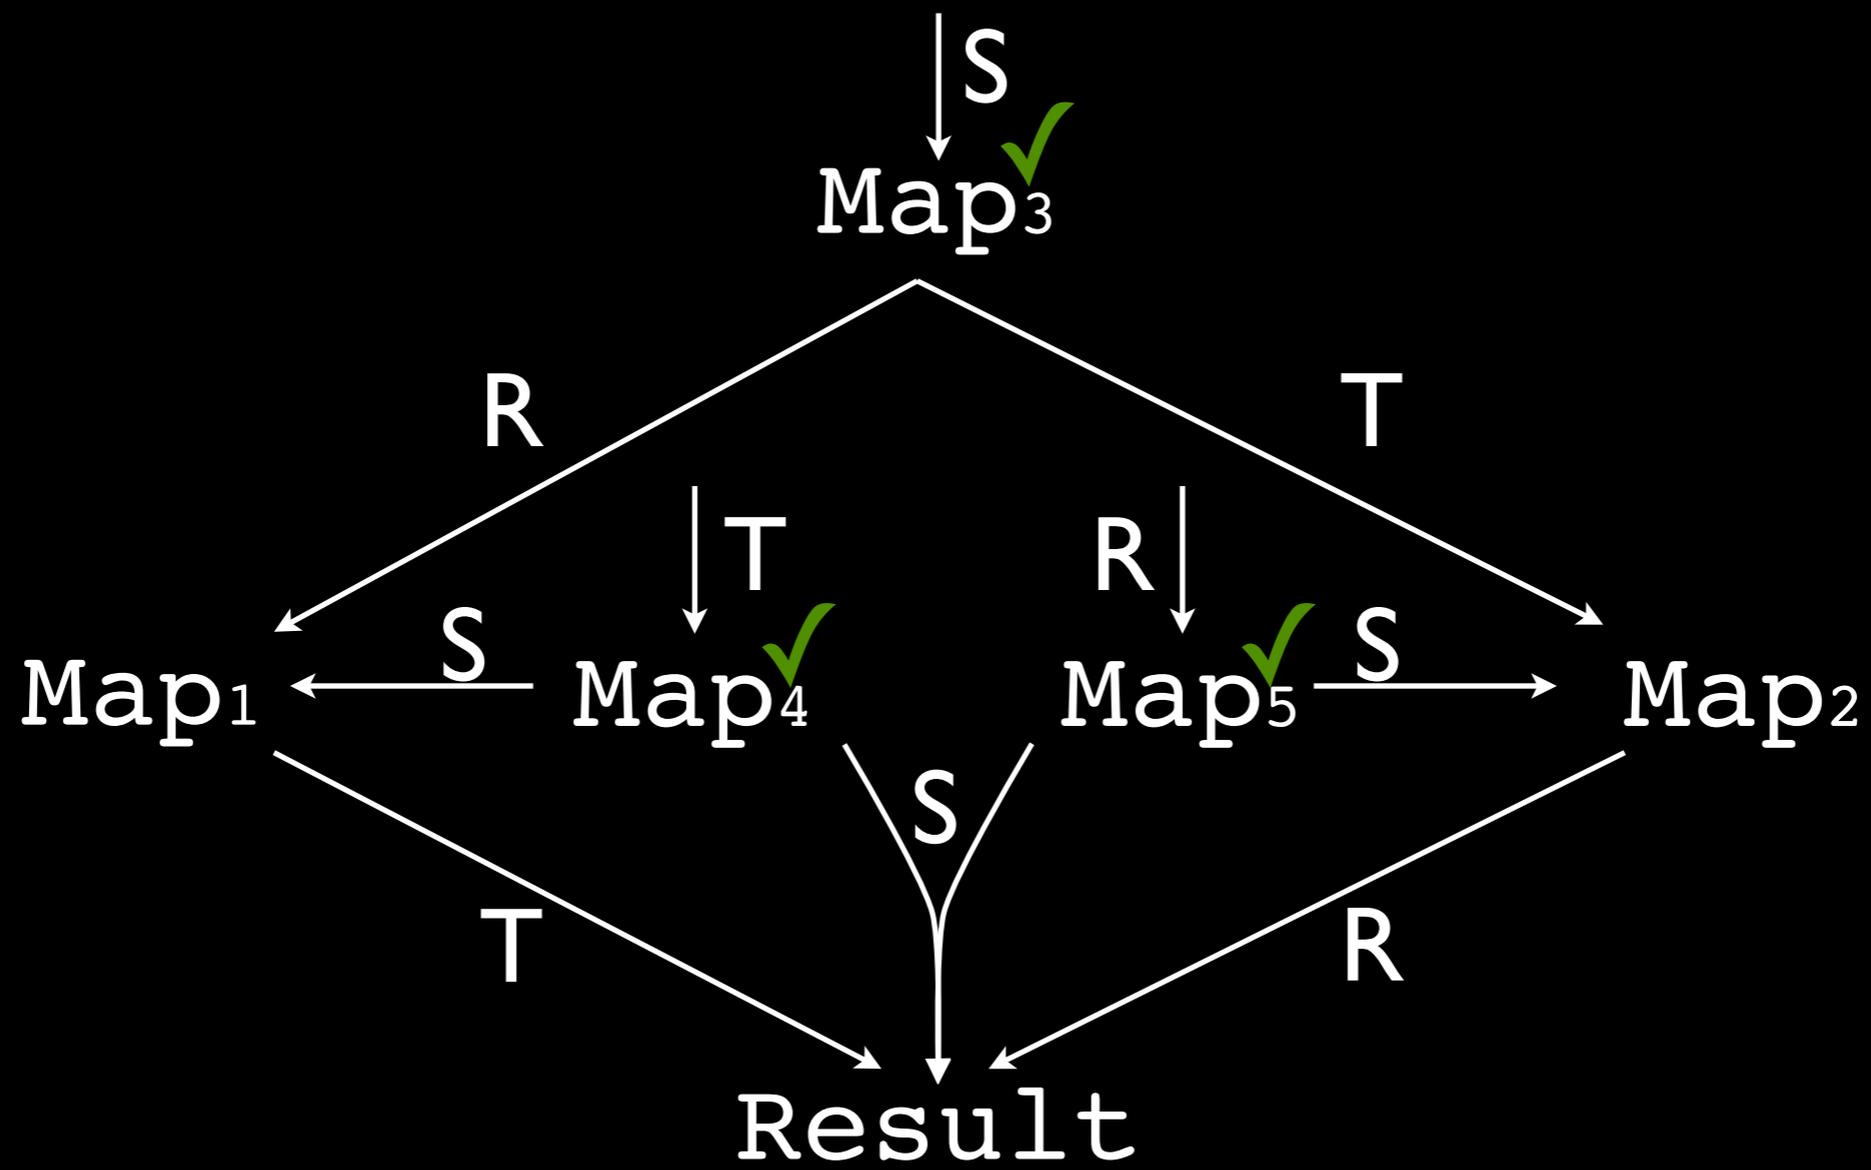


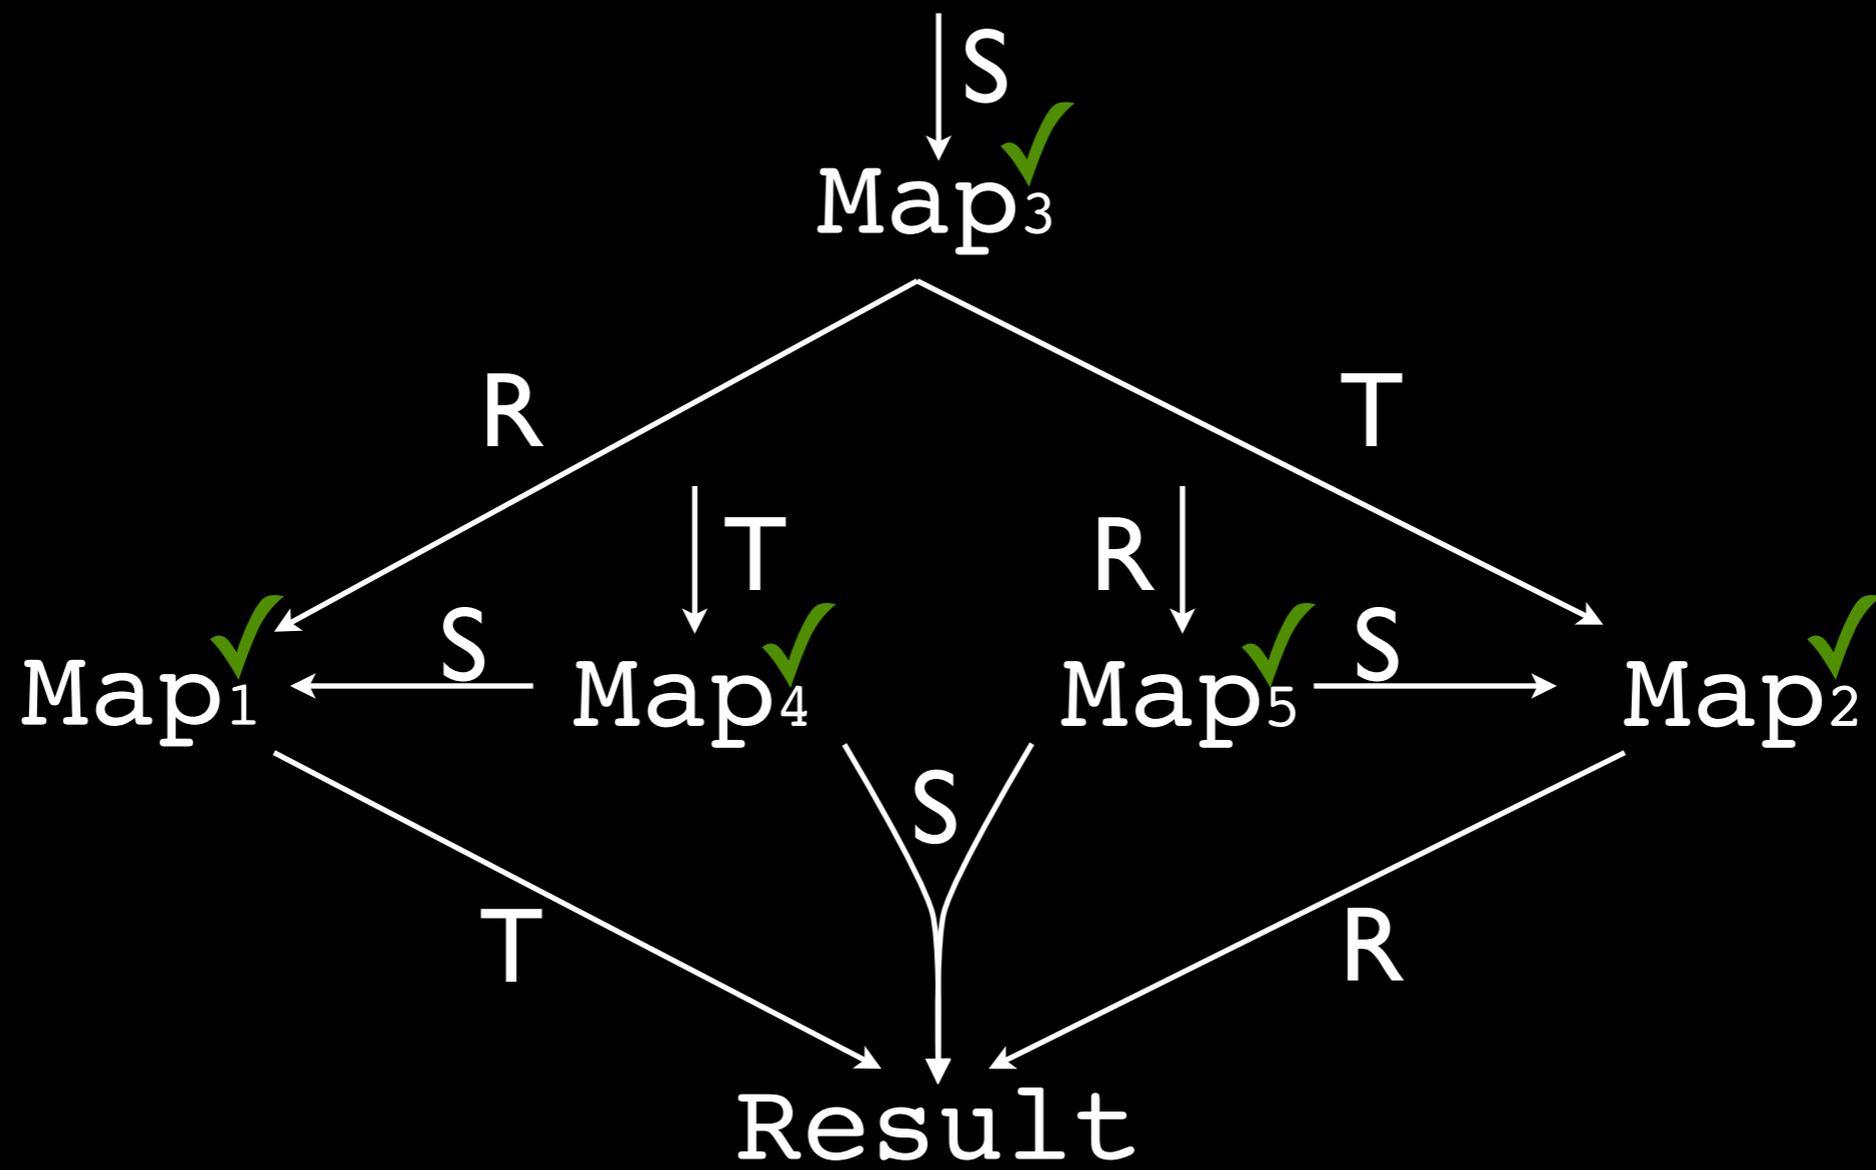
```
ON Insert_R(A,B) DO {  
  Result[] += A * Map2[B]  
  Map1[C] += A * Map3[B,C]  
  Map5[B] += A }
```

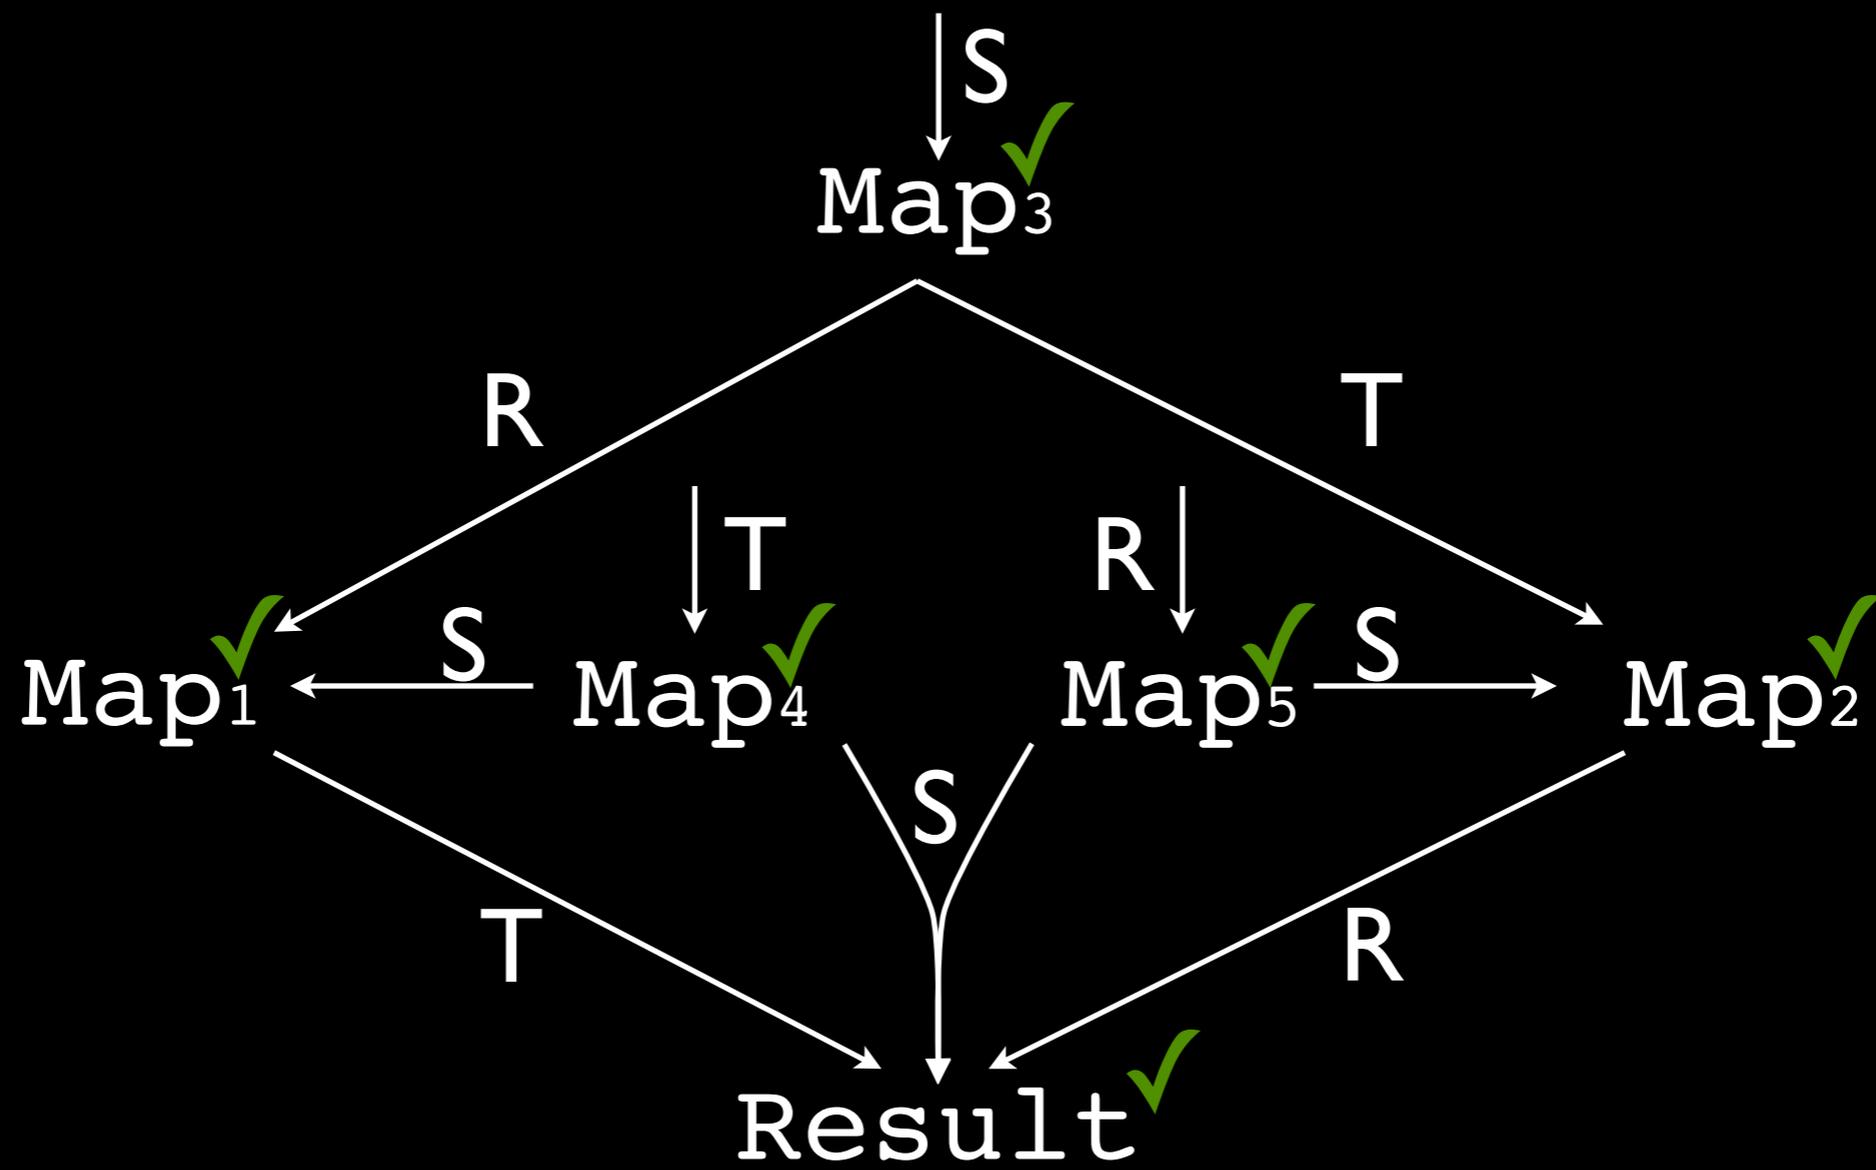
```
ON Insert_T(C,D) DO {  
  Result[] += Map1[C] * D  
  Map2[B] += D * Map3[B,C]  
  Map4[C] += D }
```

```
ON Insert_S(B,C) DO {  
  Result[] += Map5[B] * Map4[C]  
  Map1[C] += Map5[B]  
  Map2[B] += Map4[C]  
  Map3[B,C] += 1 }
```











$\langle 1, 3 \rangle \rightarrow$ 2 4 2 3 1
 $\langle 0, 1, 2 \rangle$ $\langle 0, 1, 3 \rangle$ $\langle 0, 2, 4 \rangle$ $\langle 0, 2, 5 \rangle$ $\langle 0, 3, 1 \rangle$



$\langle 1, 3 \rangle \rightarrow 12$
 $\langle \text{Epoch } 0 \rangle$



Cumulus



- Additive deltas streamline distribution of query processing.
- Key/Value stores and query processing engines compliment each other.
- Cumulus' hybrid consistency engine supports both low-latency and high-precision applications.

Laasie (and BarQL)

(with Luke Ziarek, Sumit Agarwal, and Daniel Bellinger)

Collaborative Web Applications

Logos are the property of their respective owners

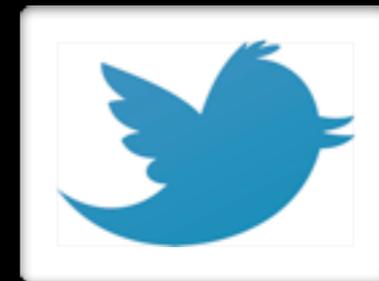
Collaborative Web Applications



Dropbox



Google Docs



Twitter



Office 365



Google Wave

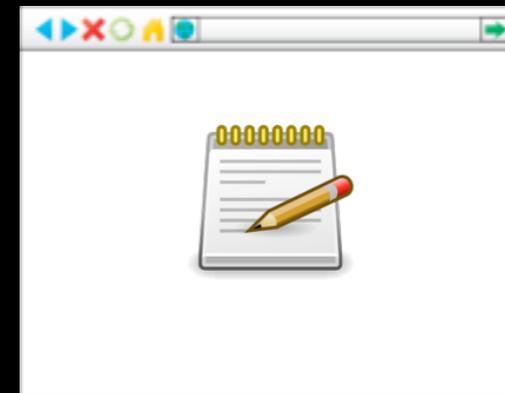
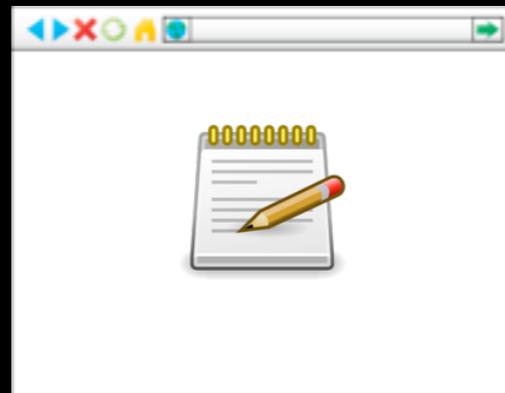
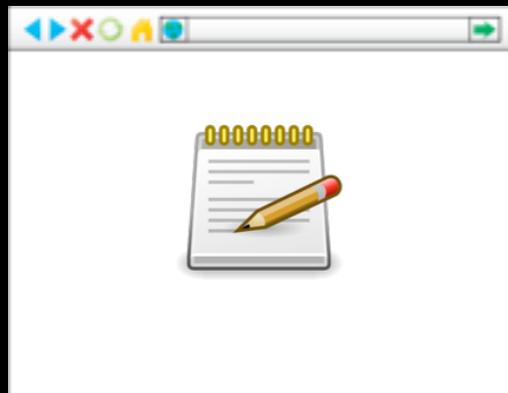


Facebook

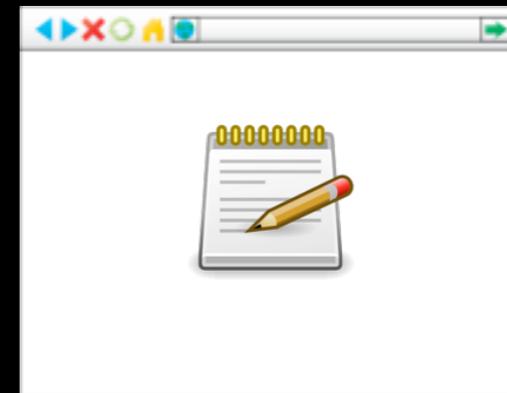
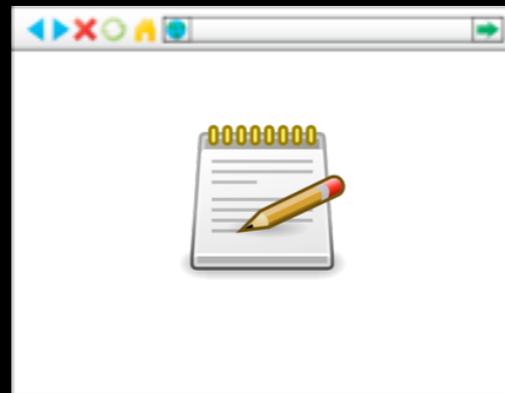
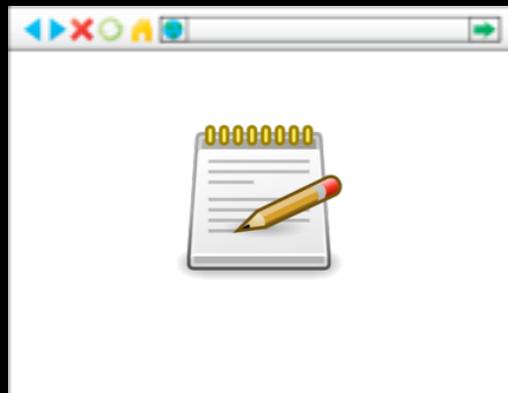


Logos are the property of their respective owners

Collaborative Web Applications

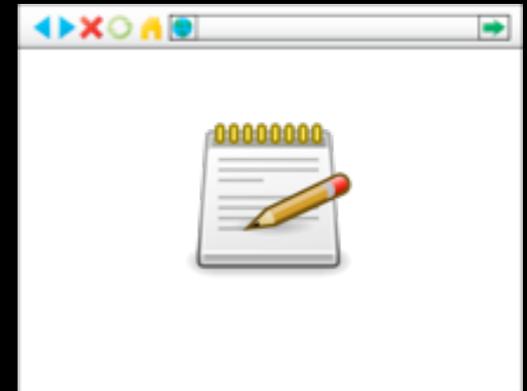
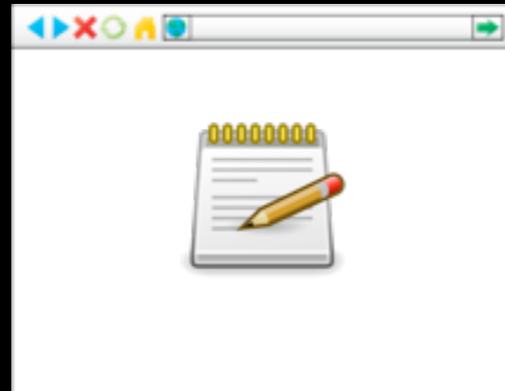
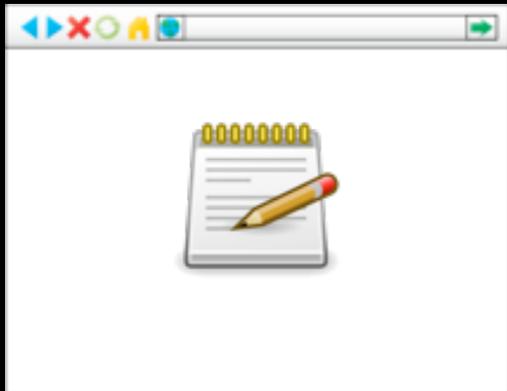


Collaborative Web Applications



The Application Lives in the Browser

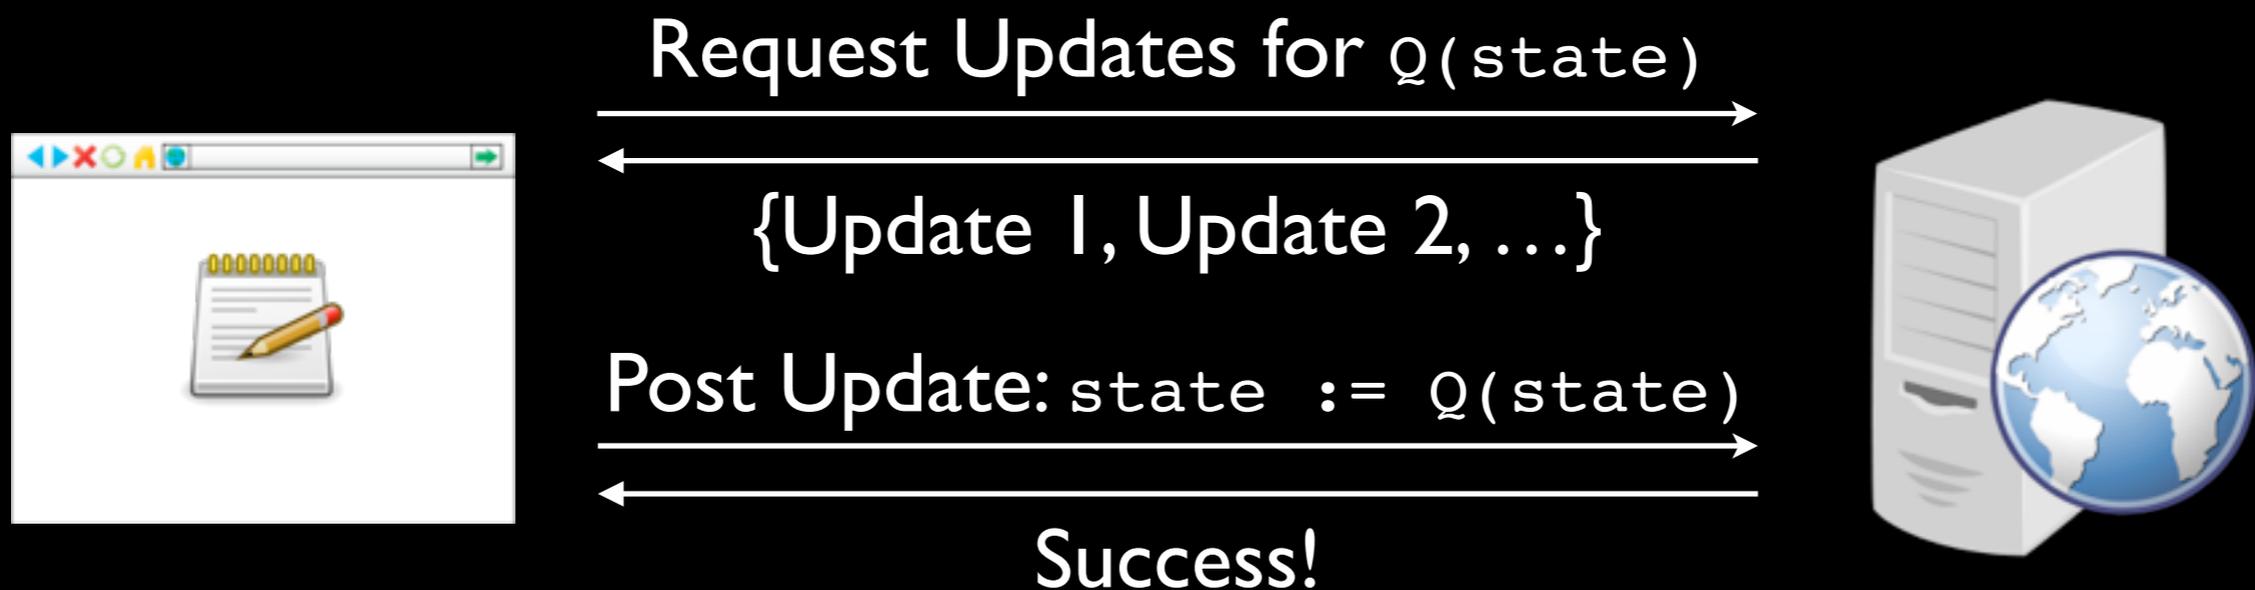
Collaborative Web Applications



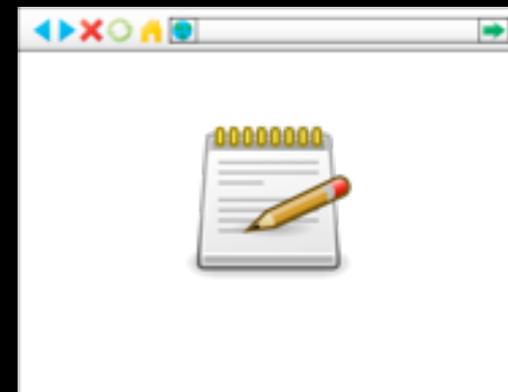
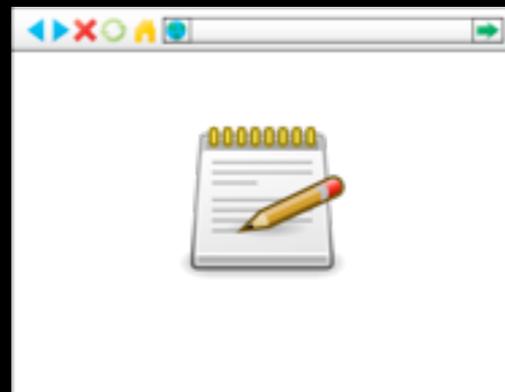
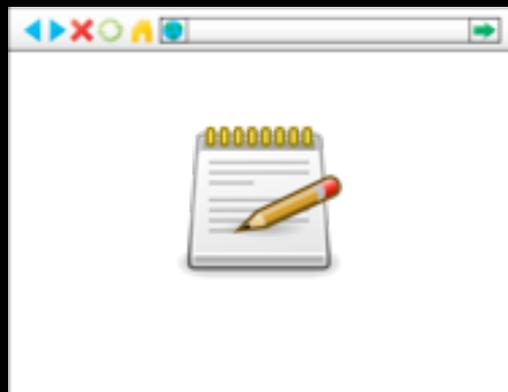
The Application Lives in the Browser
The Server Just Relays and Persists Application State

Image Source: openclipart.org, OpenIconLibrary

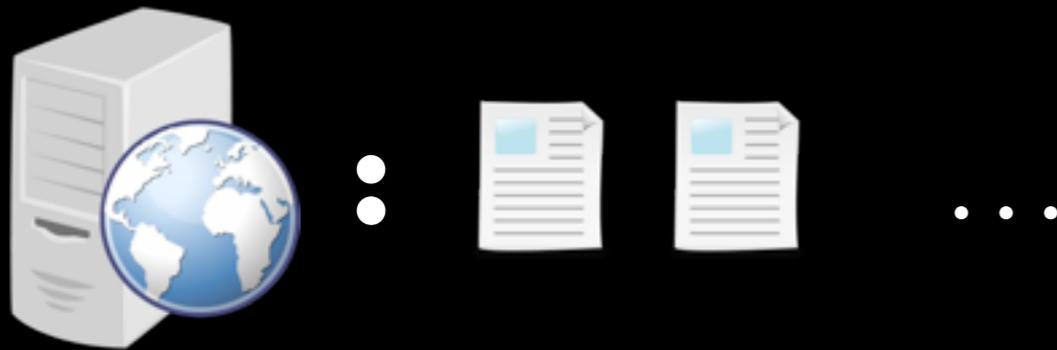
Collaborative Applications



State as a Log



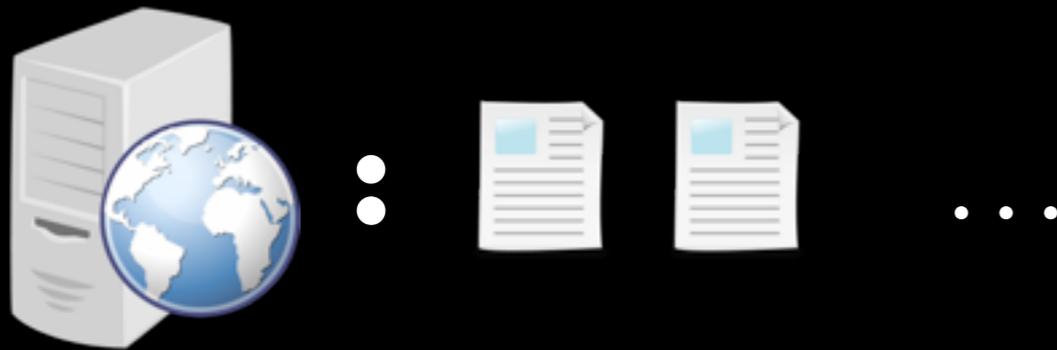
State as a Log



Clients Log State Updates
The Server Relays Log Entries

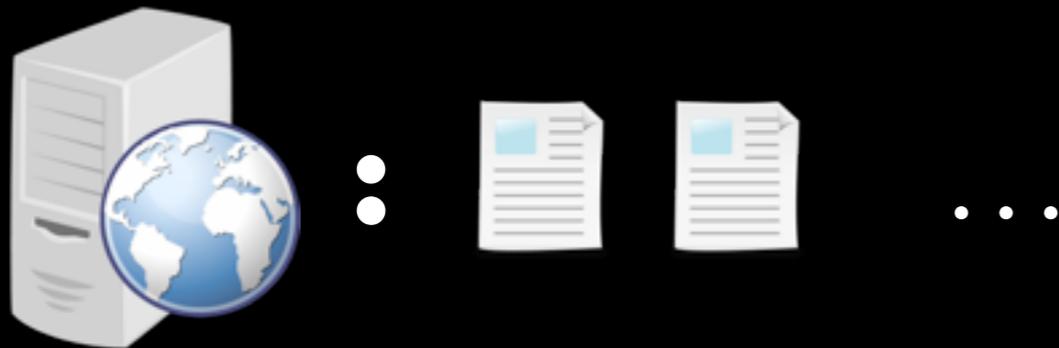
Image Source: openclipart.org, OpenIconLibrary

State as a Log



Clients Can Leave Unexpectedly

State as a Log

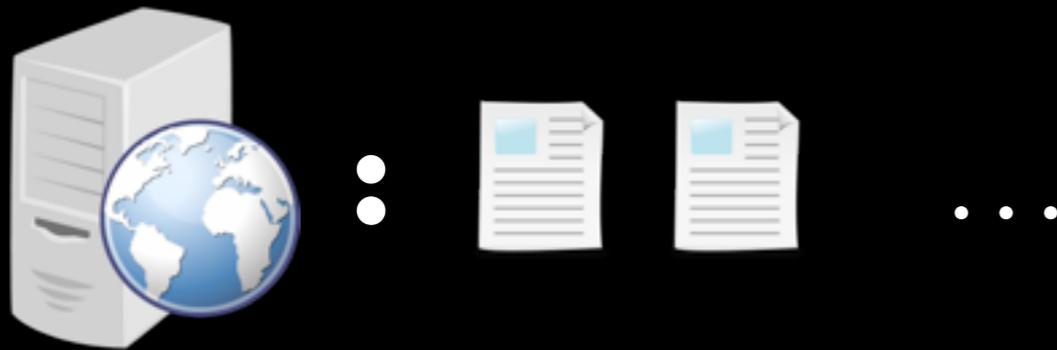


Clients Can Leave Unexpectedly

When Clients Arrive, the Server Can Restore Their State

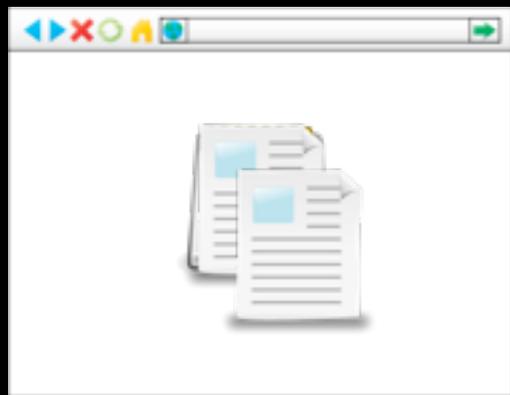
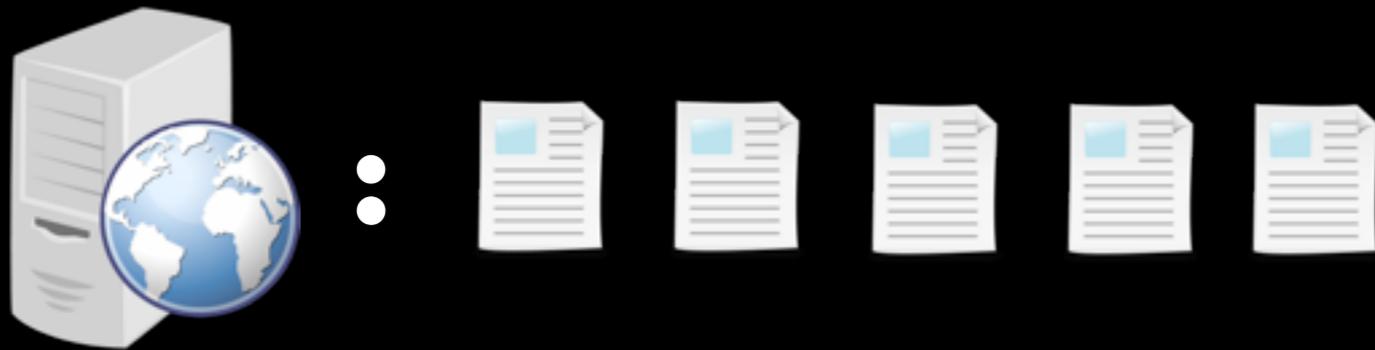
Image Source: openclipart.org, OpenIconLibrary

State as a Log



But Clients Can Leave For Other Reasons

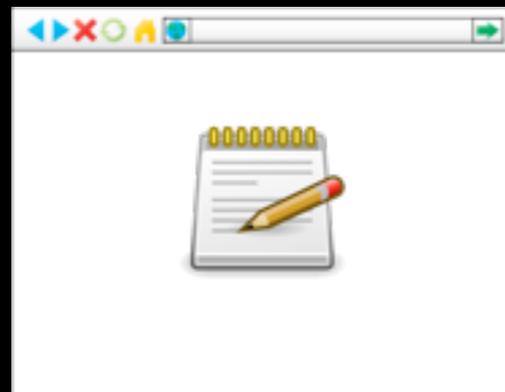
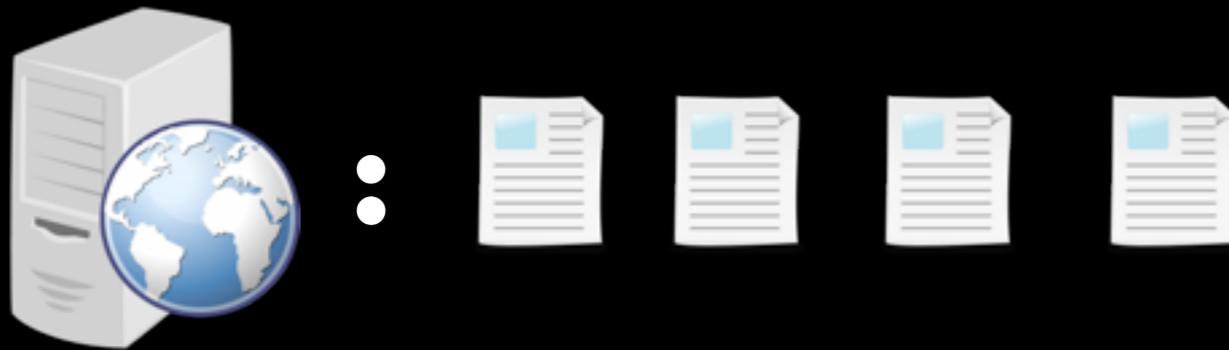
State as a Log



But Clients Can Leave For Other Reasons
...And Only Need The Latest Changes When It Returns

Image Source: openclipart.org, OpenIconLibrary

Log as a Service



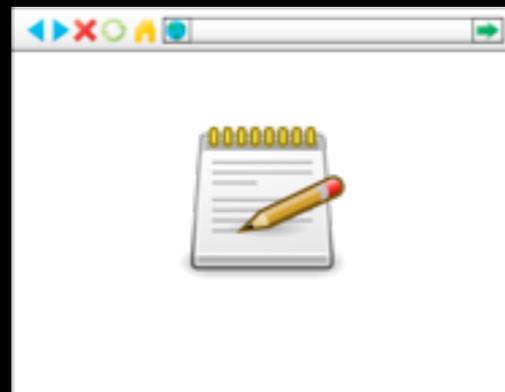
Log as a Service



:



I need all docs



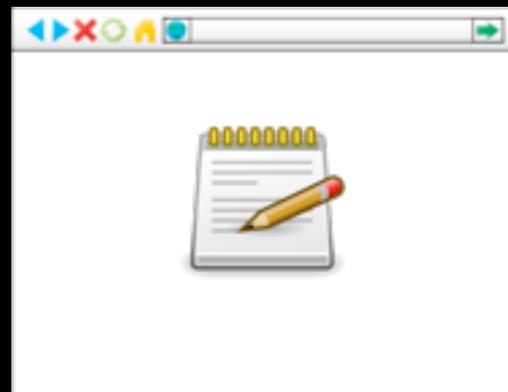
Log as a Service



:



I need all docs



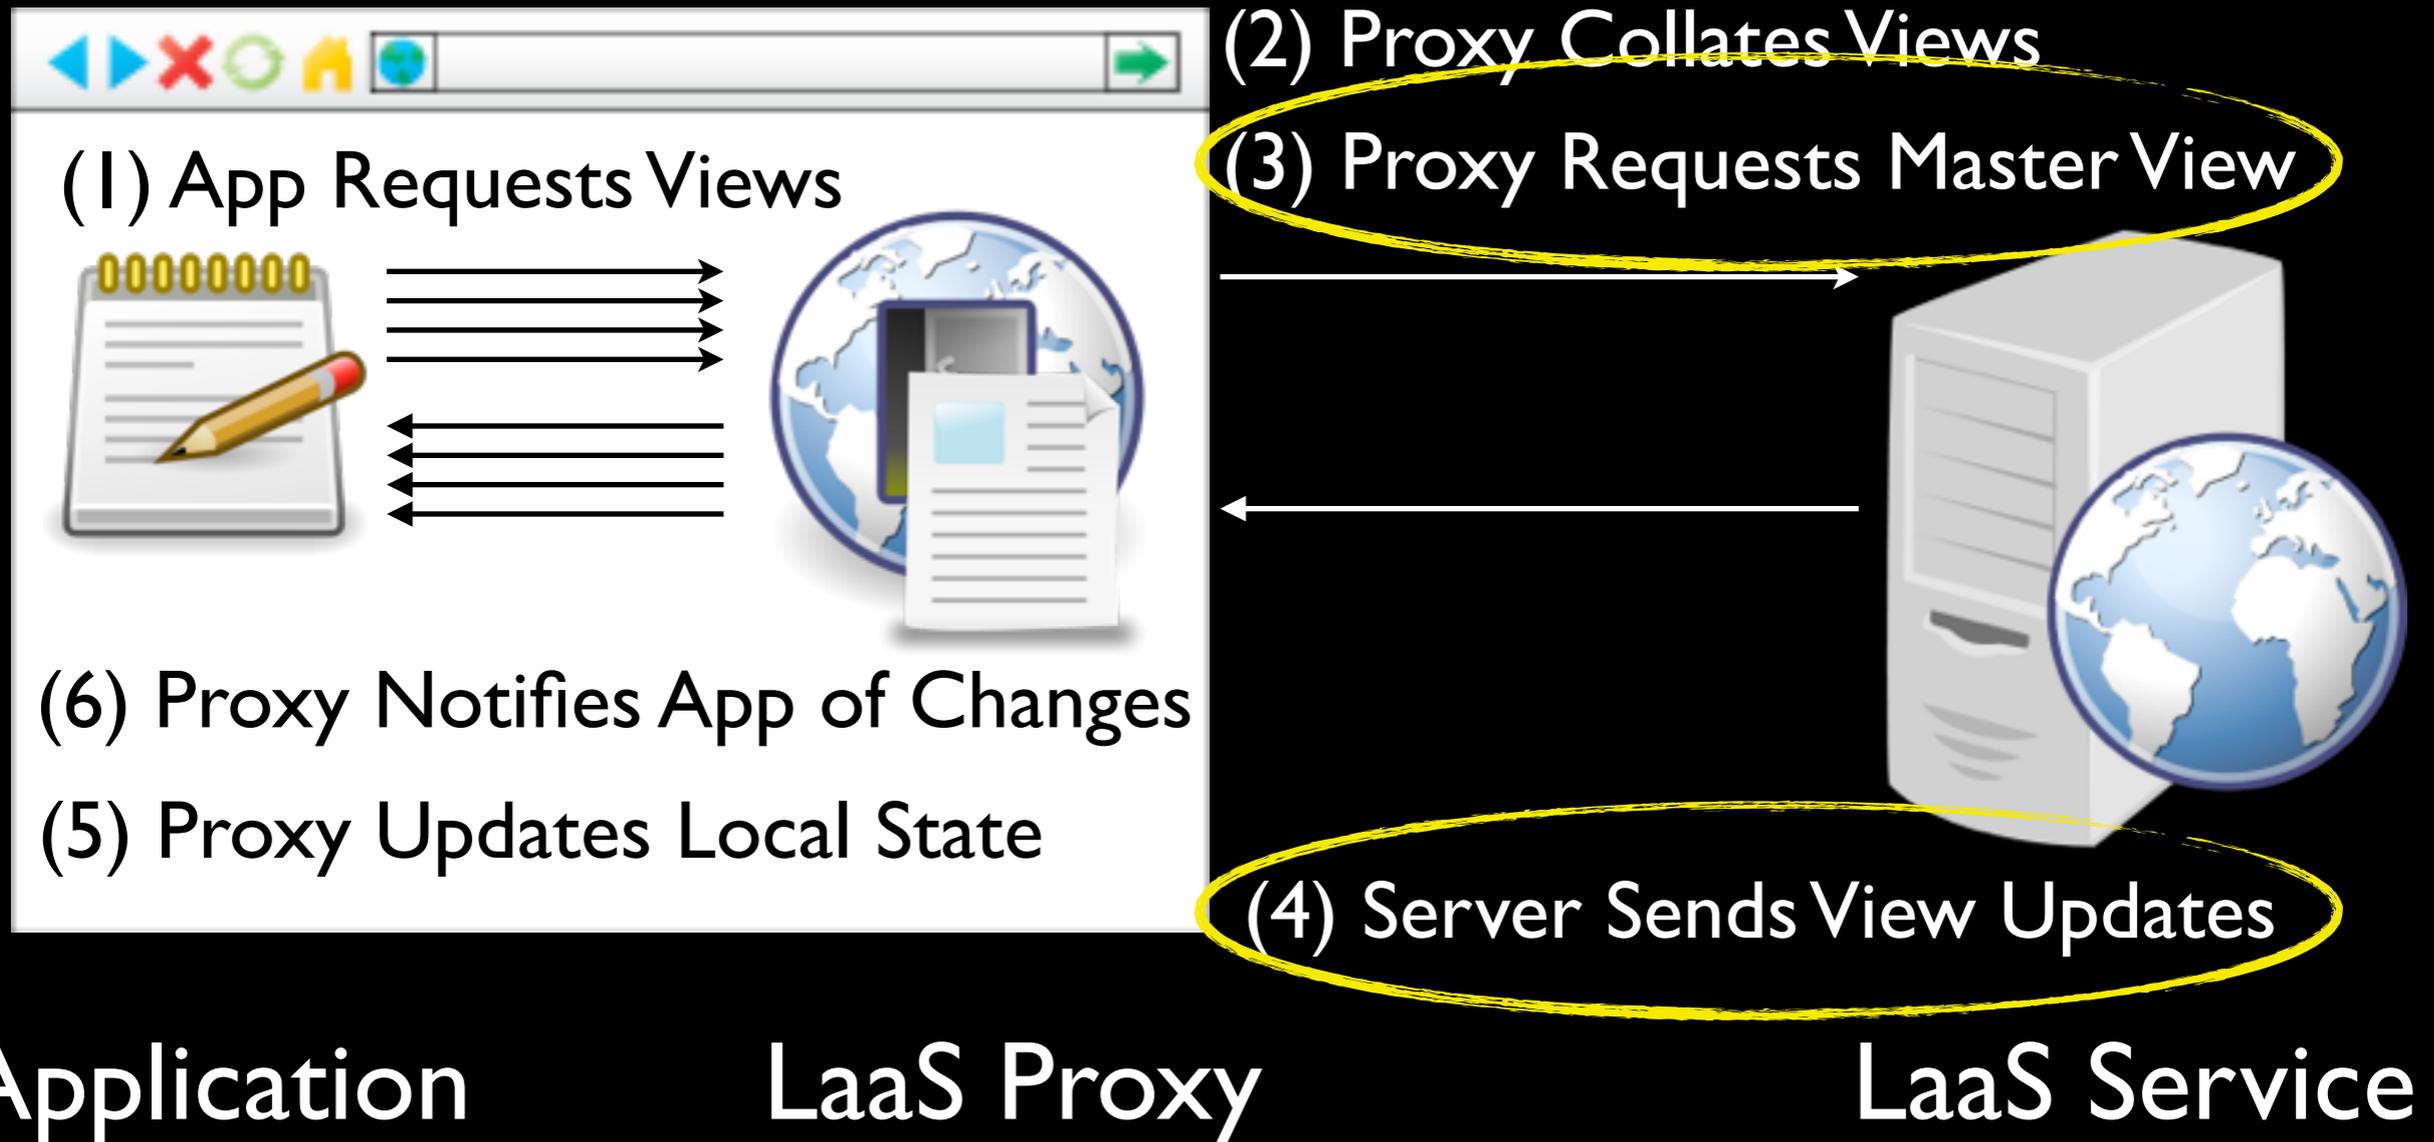
... about VLDB

... by Bob

... modified in the
last 2 days

Client state is a View!

Laasie



Laasie

Making the Log Scale: Indexing

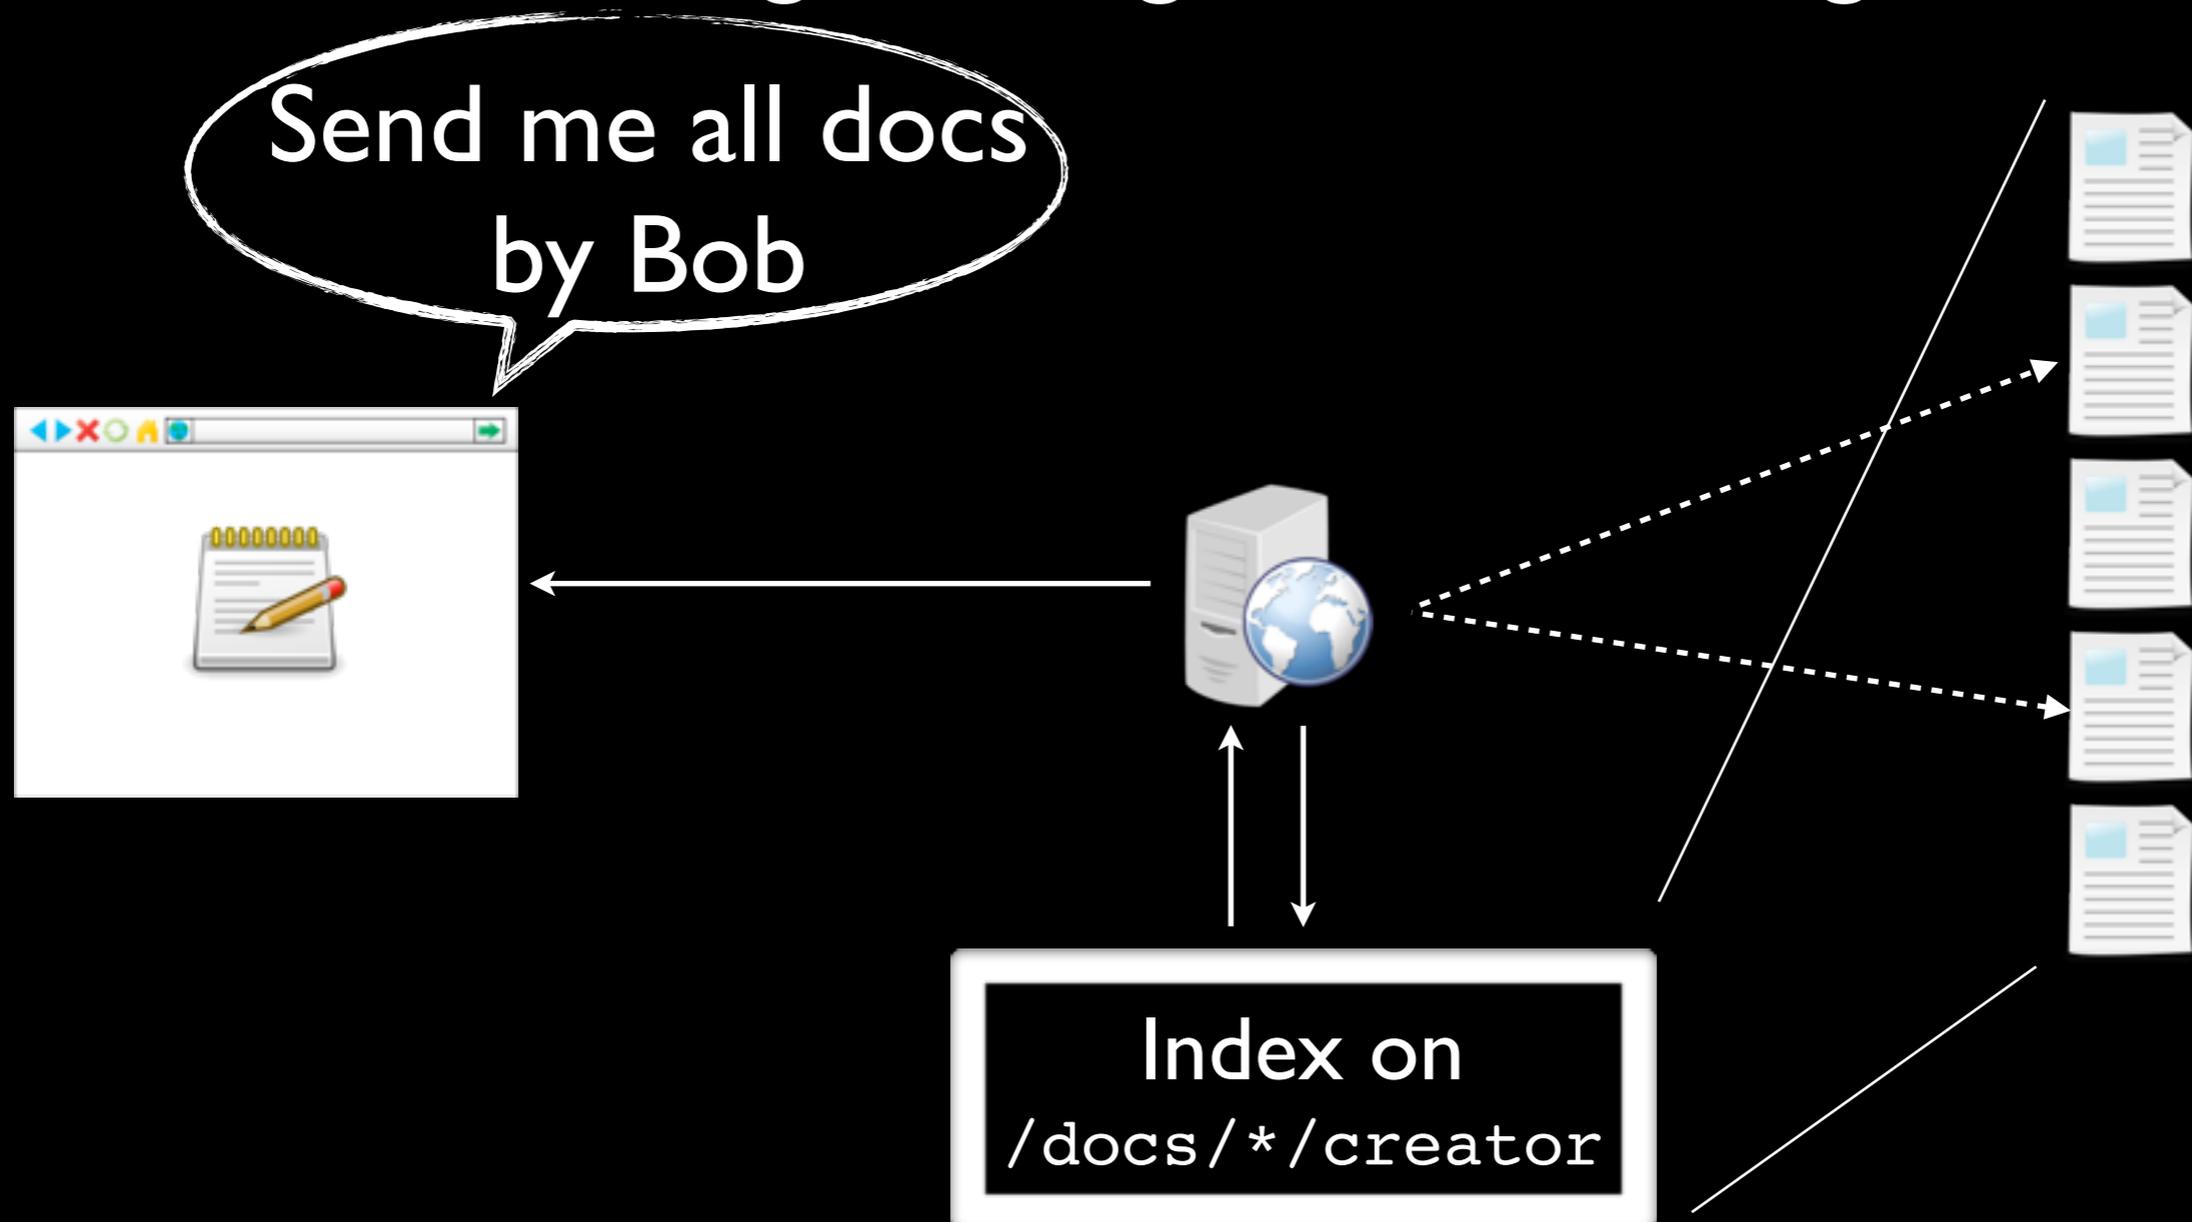
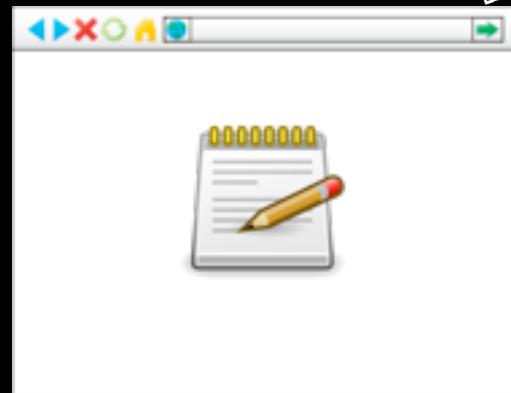


Image Source: openclipart.org, OpenIconLibrary

Laasie

Making the Log Scale: Materialization

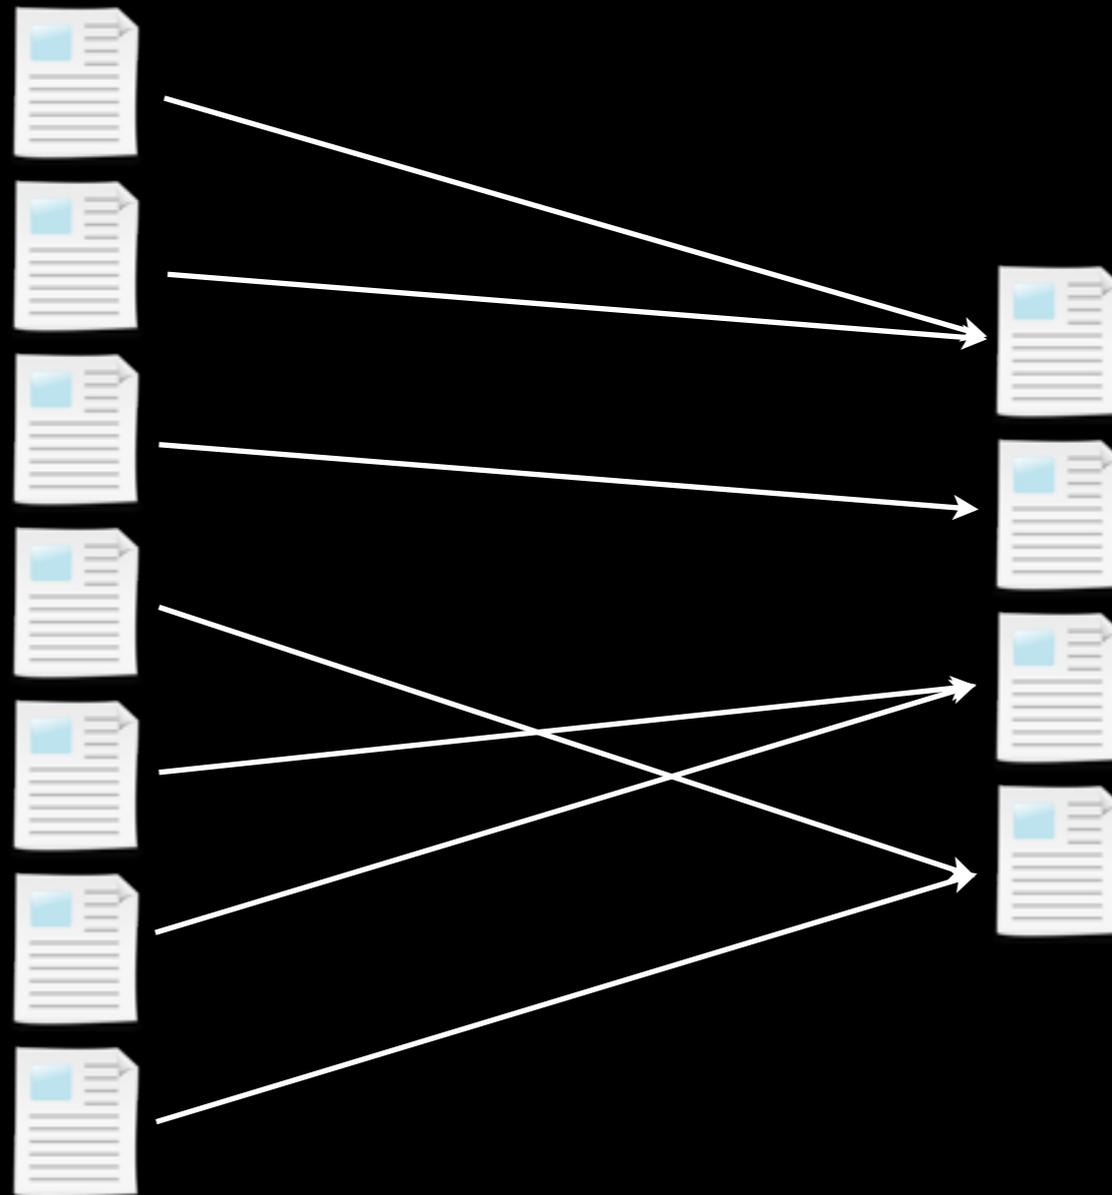
How many unread messages?



```
Materialized  
COUNT(filter(/msgs))
```

Laasie

Making the Log Scale: Log Rewrites



Active Research

- Cumulus -- Distributed Monitoring
- Laasie -- Collaborative Applications
- Mìmisbrunnr -- Managing Uncertain Knowledge
- GraphDBs -- Graphs as First Class Objects

<http://okennedy.xthemage.net/?page=research>