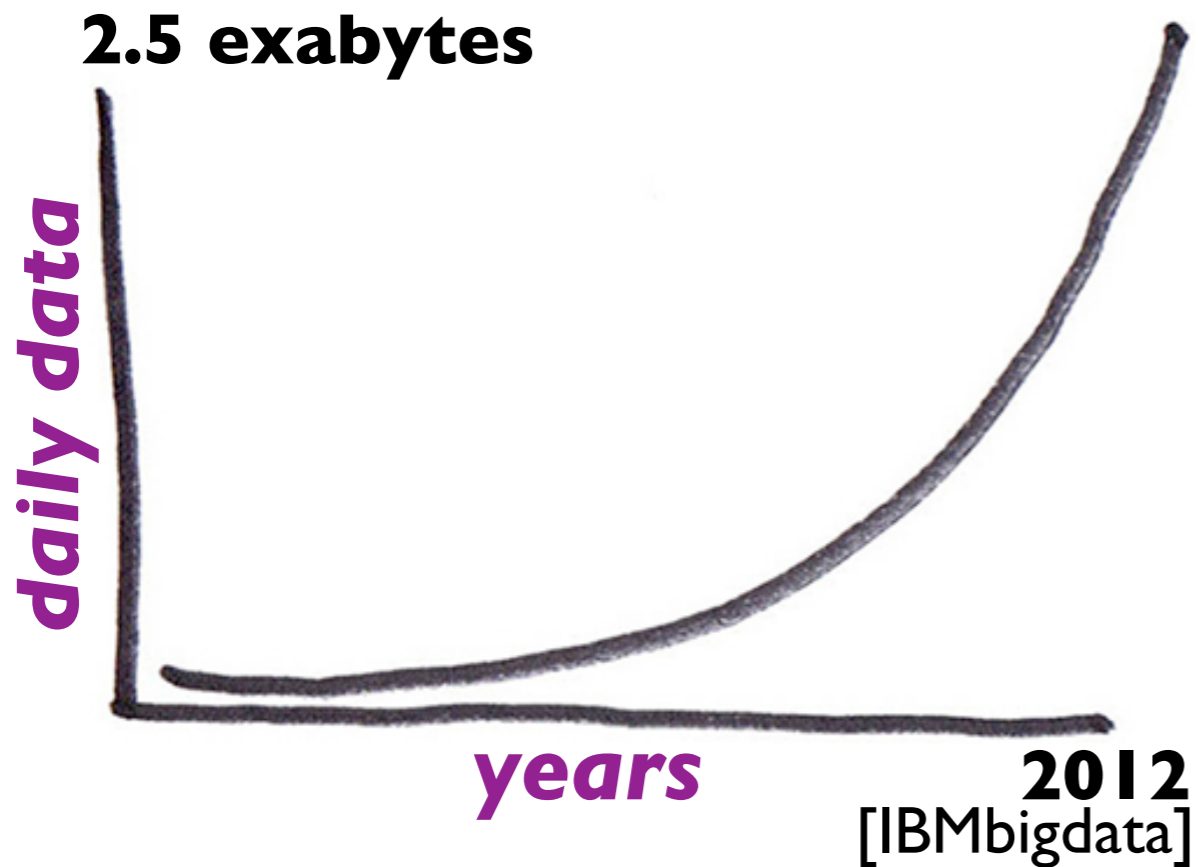


# ***Database Cracking***

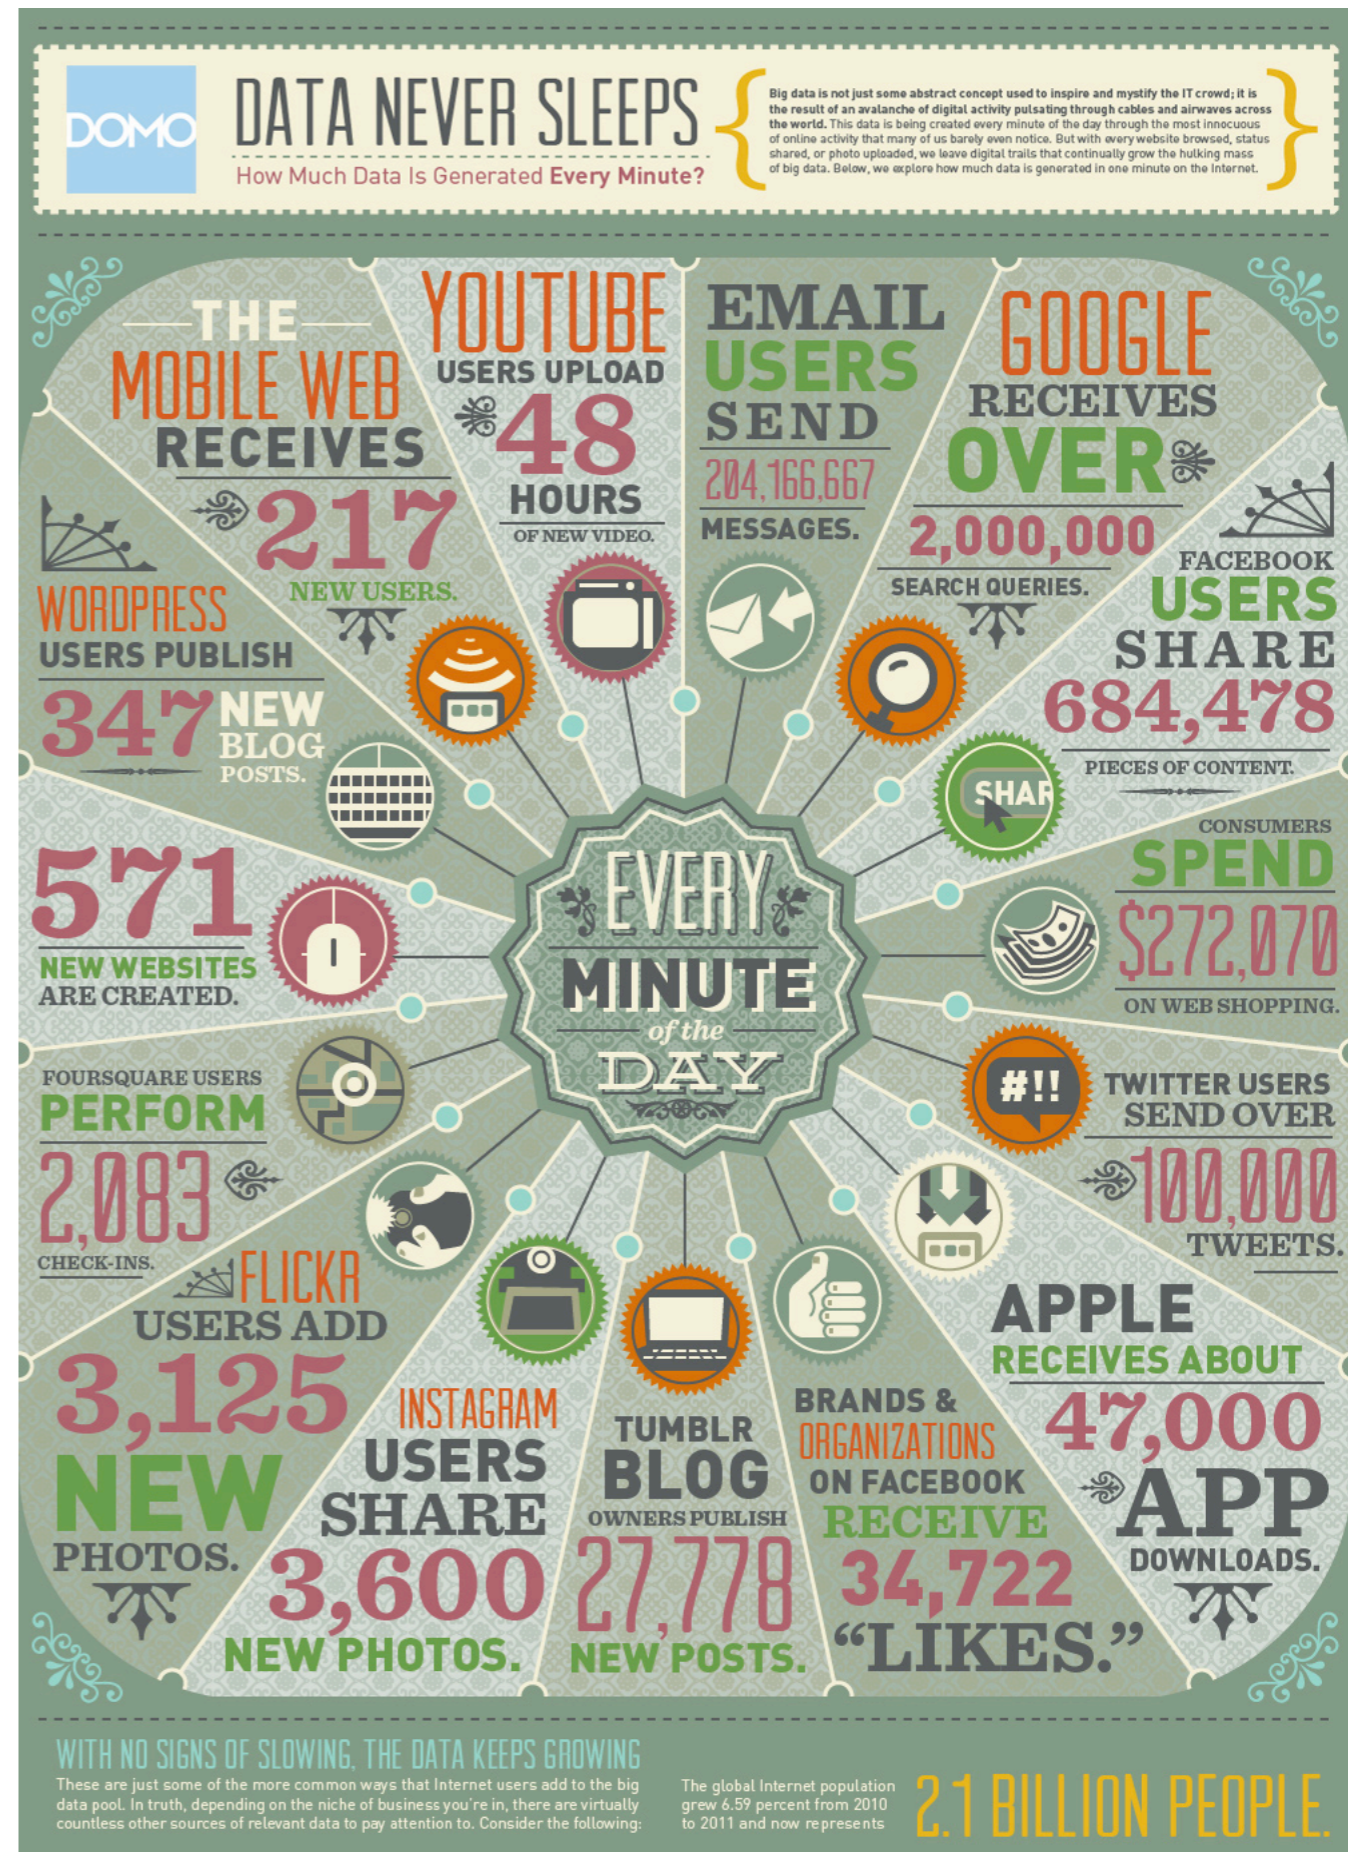
*and the path towards auto-tuning database kernels*

***Stratos Idreos***

Dutch National Research Center  
for Mathematics and Computer Science



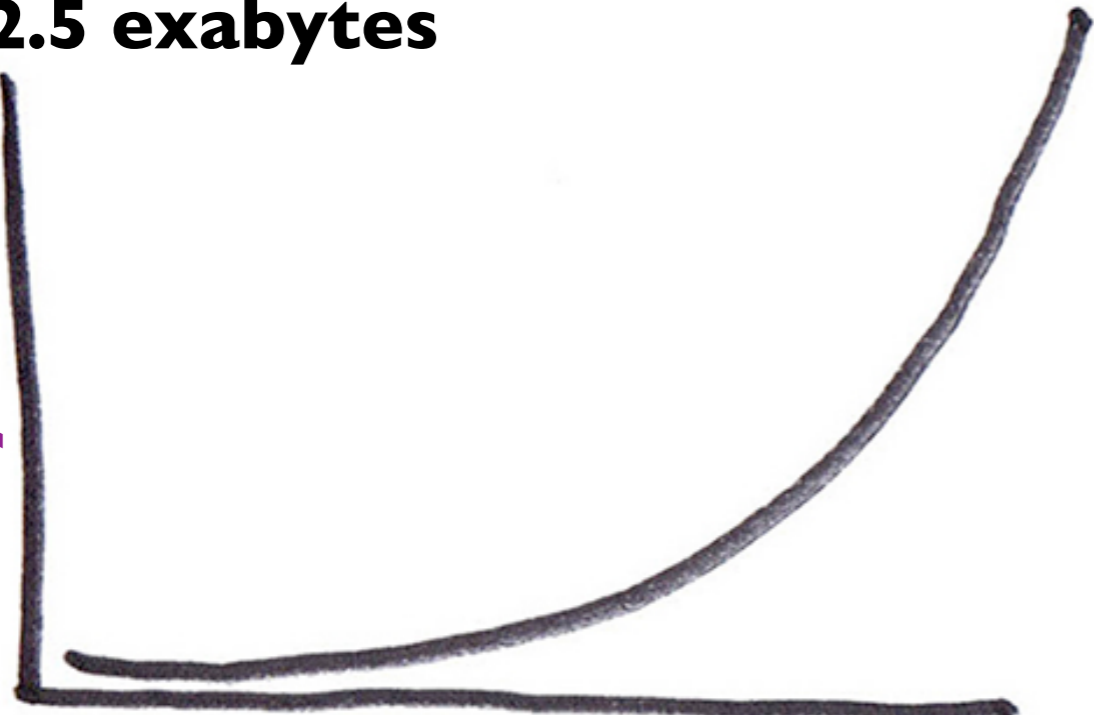
Every two days we create as much data as much we did from dawn of humanity to 2003  
[Eric Schmidt]



# easy data management

2.5 exabytes

daily data




years

2012  
[IBMbigdata]

Every two days we create as much data as much we did from dawn of humanity to 2003

[Eric Schmidt]



[Domo]

# data exploration

not always sure what we are looking for (until we find it)

## Big Data V's

volume

velocity

variety

veracity



# databases

5 decades of research

>\$100B industry, growing 10% every year

[Economist, “Data, data everywhere”]



twitter



facebook



# databases

5 decades of research

>\$100B industry, growing 10% every year

[Economist, "Data, data everywhere"]

SQL queries



**database kernel**



twitter



facebook



WIRED MAGAZINE: 16.07

SCIENCE : DISCOVERIES 

## The End of Theory: The Data Deluge Makes the Scientific Method Obsolete


By Chris Anderson  06.23.08



Illustration: Marian Bantjes

THE PETABYTE AGE:

"All models are wrong, but some are useful."

**it is time for a **paradigm shift**  
in how we design databases systems**


***database systems great...***

declarative processing, back-end to numerous apps

***database systems great...***

declarative processing, back-end to numerous apps


***but databases have become too heavy and blind!***



***database systems great...***

declarative processing, back-end to numerous apps

***but databases have become too heavy and blind!***




timeline →

***database systems great...***

declarative processing, back-end to numerous apps

***but databases have become too heavy and blind!***



**load**




timeline →

***database systems great...***

declarative processing, back-end to numerous apps

***but databases have become too heavy and blind!***



**load**


**tune**

timeline →

***database systems great...***

declarative processing, back-end to numerous apps

***but databases have become too heavy and blind!***



**load**


**tune**

**query**

**timeline** →

**expert users - idle time - workload knowledge**

*but databases have become too heavy and blind!*



**load**

**tune**

**query**

**timeline** →





**data systems tailored for data exploration**

**no workload knowledge**

**no installation steps**



**data systems tailored for data exploration**

**no workload knowledge**

**no installation steps**



**minimize data-to-query time**

**data systems tailored for data exploration**

## adaptive indexing



7 years, 7 papers

## adaptive loading



3 years, 3 papers

## dbTouch



0.5 year, 1 paper

Martin Kersten, Stefan Manegold, Felix Halim, Panagiotis Karras, Roland Yap, Goetz Graefe, Harumi Kuno, Eleni Petraki, Anastasia Ailamaki, Ioannis Alagiannis, Renata Borovica, Miguel Branco, Ryan Johnson, Erietta Liarou



ÉCOLE POLYTECHNIQUE  
FÉDÉRALE DE LAUSANNE



NUS

National University  
of Singapore

**load**

**tune**

**query**

**adaptive  
loading**

**adaptive  
indexing**

**dbTouch**

Martin Kersten, Stefan Manegold, Felix Halim, Panagiotis Karras, Roland Yap, Goetz Graefe, Harumi Kuno, Eleni Petraki, Anastasia Ailamaki, Ioannis Alagiannis, Renata Borovica, Miguel Branco, Ryan Johnson, Erietta Liarou



ÉCOLE POLYTECHNIQUE  
FÉDÉRALE DE LAUSANNE



# indexing



*tune = create proper indices offline*

*performance 10-100X*

# indexing



*tune= create proper indices offline*

*performance 10-100X*

*but it depends on workload!*

*which indices to build?  
on which data parts?  
and when to build them?*



**load**

**tune**

**query**



**load**

**tune**

**query**



timeline  
→

**load**

**tune**



**query**






**sample workload**




timeline  
→









**complex and time consuming process**

# human administrators + auto-tuning tools

**sample workload**

**analyze**

**create indices**

**query**

timeline



**complex and time consuming process**

# Big Data V's

volume

velocity

variety

veracity

## *what can go wrong?*

not enough space to index all data

not enough idle time to finish proper tuning

by the time we finish tuning, the workload changes

not enough money - energy - resources

# Big Data V's

volume

velocity

variety

veracity

## *what can go wrong?*

not enough space to index all data

not enough idle time to finish proper tuning

by the time we finish tuning, the workload changes

not enough money - energy - resources

# ***database cracking***

# *database cracking*



idle time



workload  
knowledge



external  
tools




human  
control


# *database cracking*

*auto-tuning database kernels*


*incremental, adaptive, partial indexing*



initialization



querying




## ***database cracking***

***auto-tuning database kernels***


*incremental, adaptive, partial indexing*



initialization



querying




## ***database cracking***

***auto-tuning database kernels***


*incremental, adaptive, partial indexing*



initialization



querying



## ***database cracking***


***auto-tuning database kernels***

*incremental, adaptive, partial indexing*

**every query is treated as an advice  
on how data should be stored**

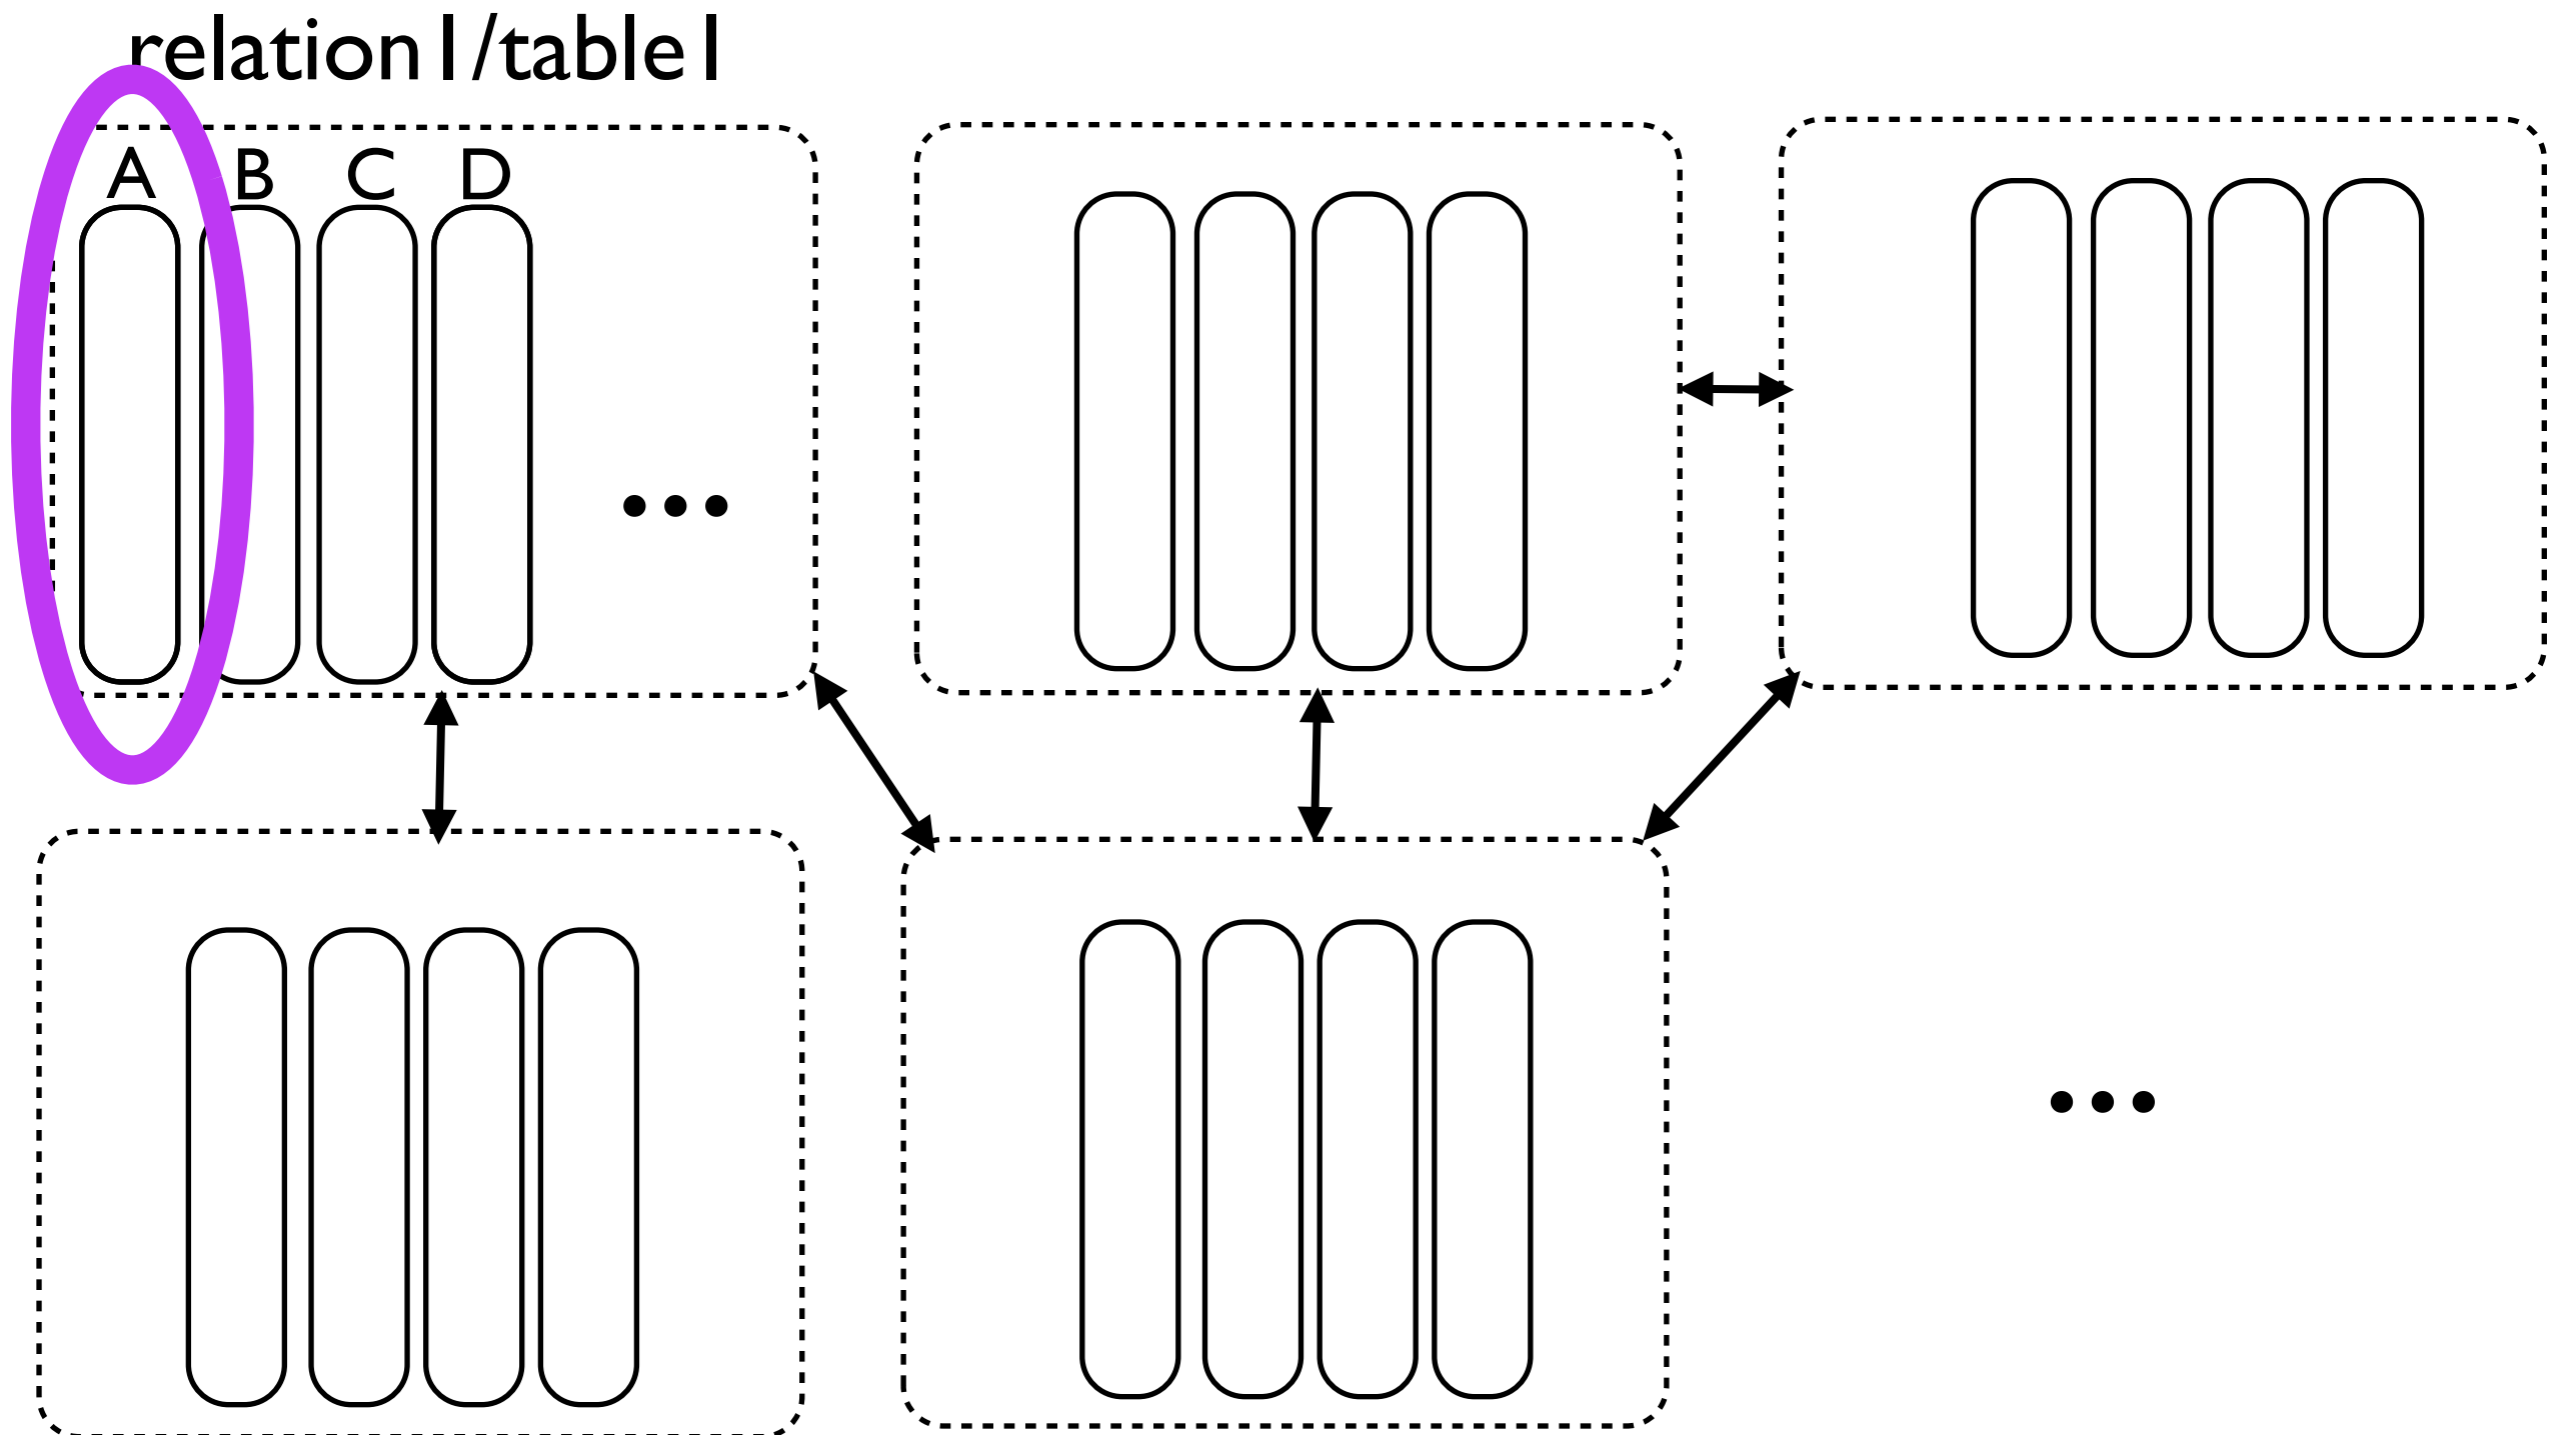
# column-store database

a fixed-width and dense array per attribute



# column-store database

a fixed-width and dense array per attribute



# full indexing example

column A

Q1:  
select R.A  
from R  
where R.A > 10  
and R.A < 14

13  
16  
4  
9  
2  
12  
7  
1  
19  
3  
14  
11  
8  
6

# full indexing example

column A

Q1:  
select R.A  
from R  
where R.A > 10  
and R.A < 14

13  
16  
4  
9  
2  
12  
7  
1  
19  
3  
14  
11  
8  
6

# full indexing example

column A

Q1:  
select R.A  
from R  
where R.A > 10  
and R.A < 14

13  
16  
4  
9  
2  
12  
7  
1  
19  
3  
14  
11  
8  
6

**sort**

1  
2  
3  
4  
6  
7  
8  
9  
11  
12  
13  
14  
16  
19

# full indexing example

Q1:  
select R.A  
from R  
where R.A > 10  
and R.A < 14

column A

13  
16  
4  
9  
2  
12  
7  
1  
19  
3  
14  
11  
8  
6

**sort** →

**binary  
search**

1  
2  
3  
4  
6  
7  
8  
9  
11  
12  
13  
14  
16  
19

# full indexing example

Q1:  
select R.A  
from R  
where R.A > 10  
and R.A < 14

column A

13  
16  
4  
9  
2  
12  
7  
1  
19  
3  
14  
11  
8  
6

**sort**

**binary  
search**

1  
2  
3  
4  
6  
7  
8  
9  
11  
12  
13  
14  
16  
19

result

# full indexing example

Q1:  
select R.A  
from R  
where R.A > 10  
and R.A < 14

column A

13  
16  
4  
9  
2  
12  
7  
1  
19  
3  
14  
11  
8  
6

**sort**

**binary  
search**

1  
2  
3  
4  
6  
7  
8  
9  
11  
12  
13  
14  
16  
19

**time  
+  
knowledge**

**result**

# cracking example

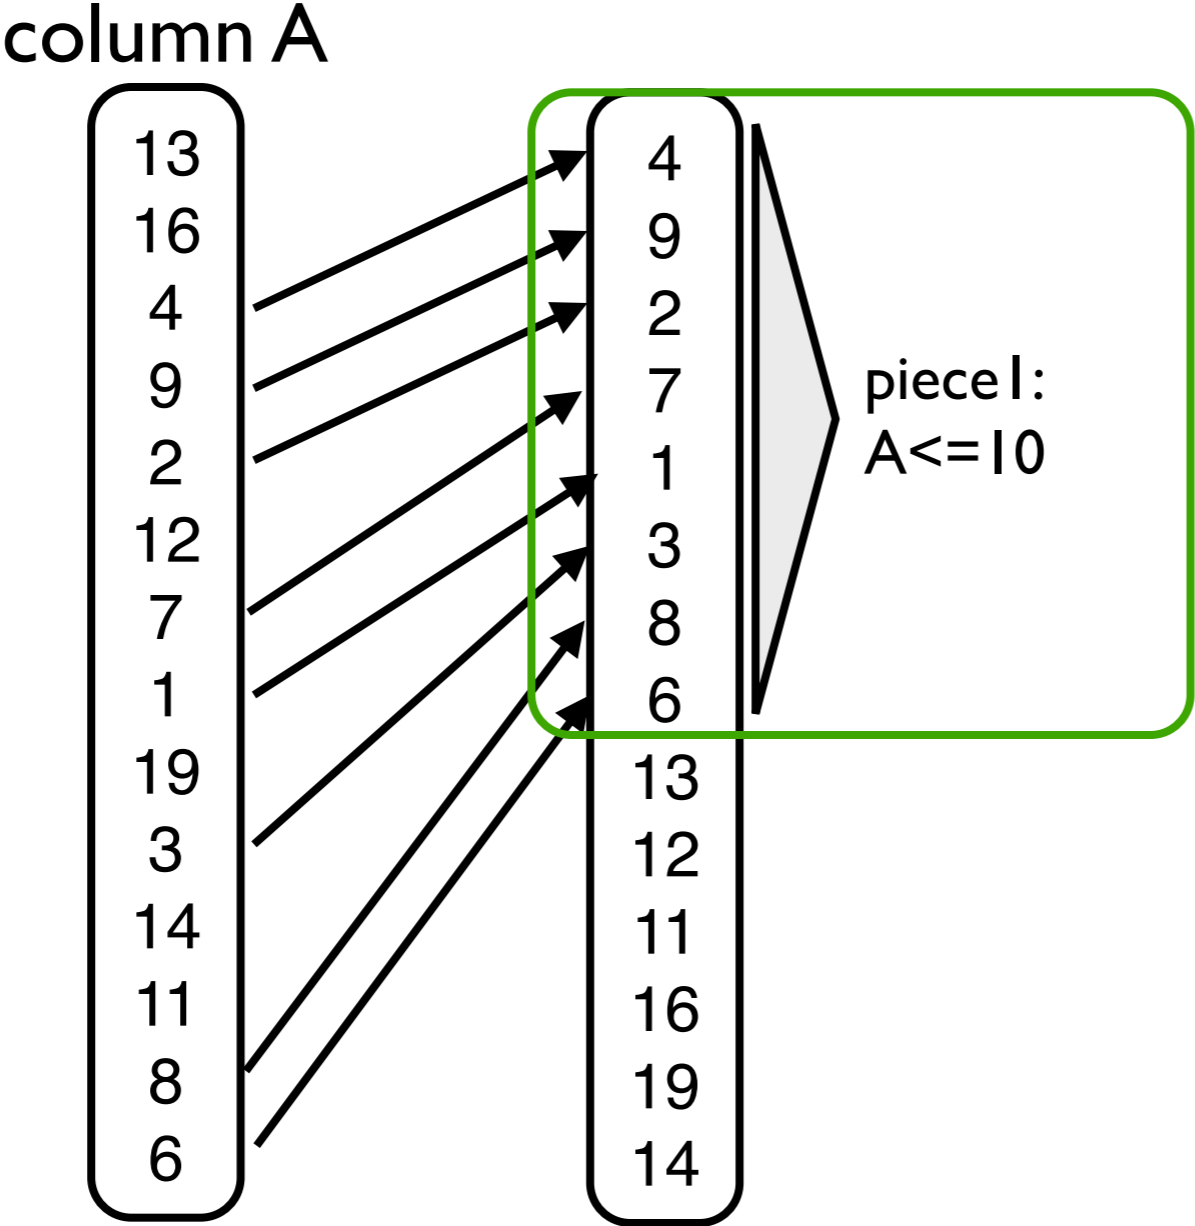
Q1:  
select R.A  
from R  
where R.A > 10  
and R.A < 14

column A

13  
16  
4  
9  
2  
12  
7  
1  
19  
3  
14  
11  
8  
6


# cracking example

Q1:  
select R.A  
from R  
where R.A > 10  
and R.A < 14



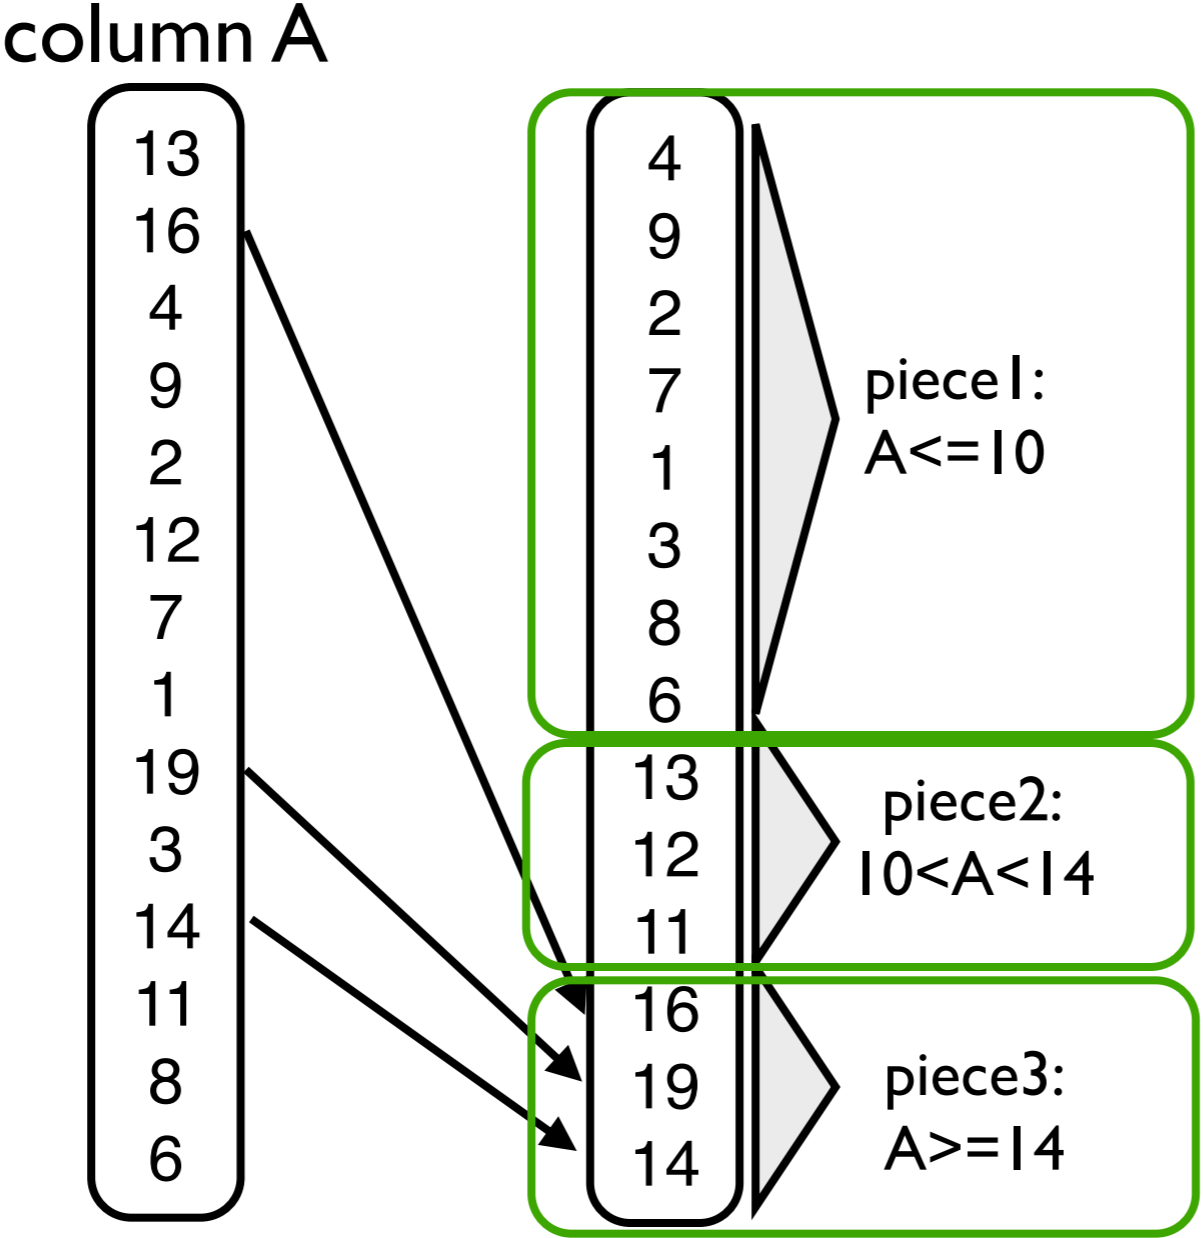
# cracking example

Q1:  
select R.A  
from R  
where R.A > 10  
and R.A < 14



# cracking example

Q1:  
select R.A  
from R  
where R.A > 10  
and R.A < 14




# cracking example

Q1:  
select R.A  
from R  
where R.A > 10  
and R.A < 14

column A

13  
16  
4  
9  
2  
12  
7  
1  
19  
3  
14  
11  
8  
6




# gain knowledge on how data is organized

Q1:  
select R.A  
from R  
where R.A > 10  
and R.A < 14

column A

13  
16  
4  
9  
2  
12  
7  
1  
19  
3  
14  
11  
8  
6




# gain knowledge on how data is organized

Q1:  
select R.A  
from R  
where R.A > 10  
and R.A < 14

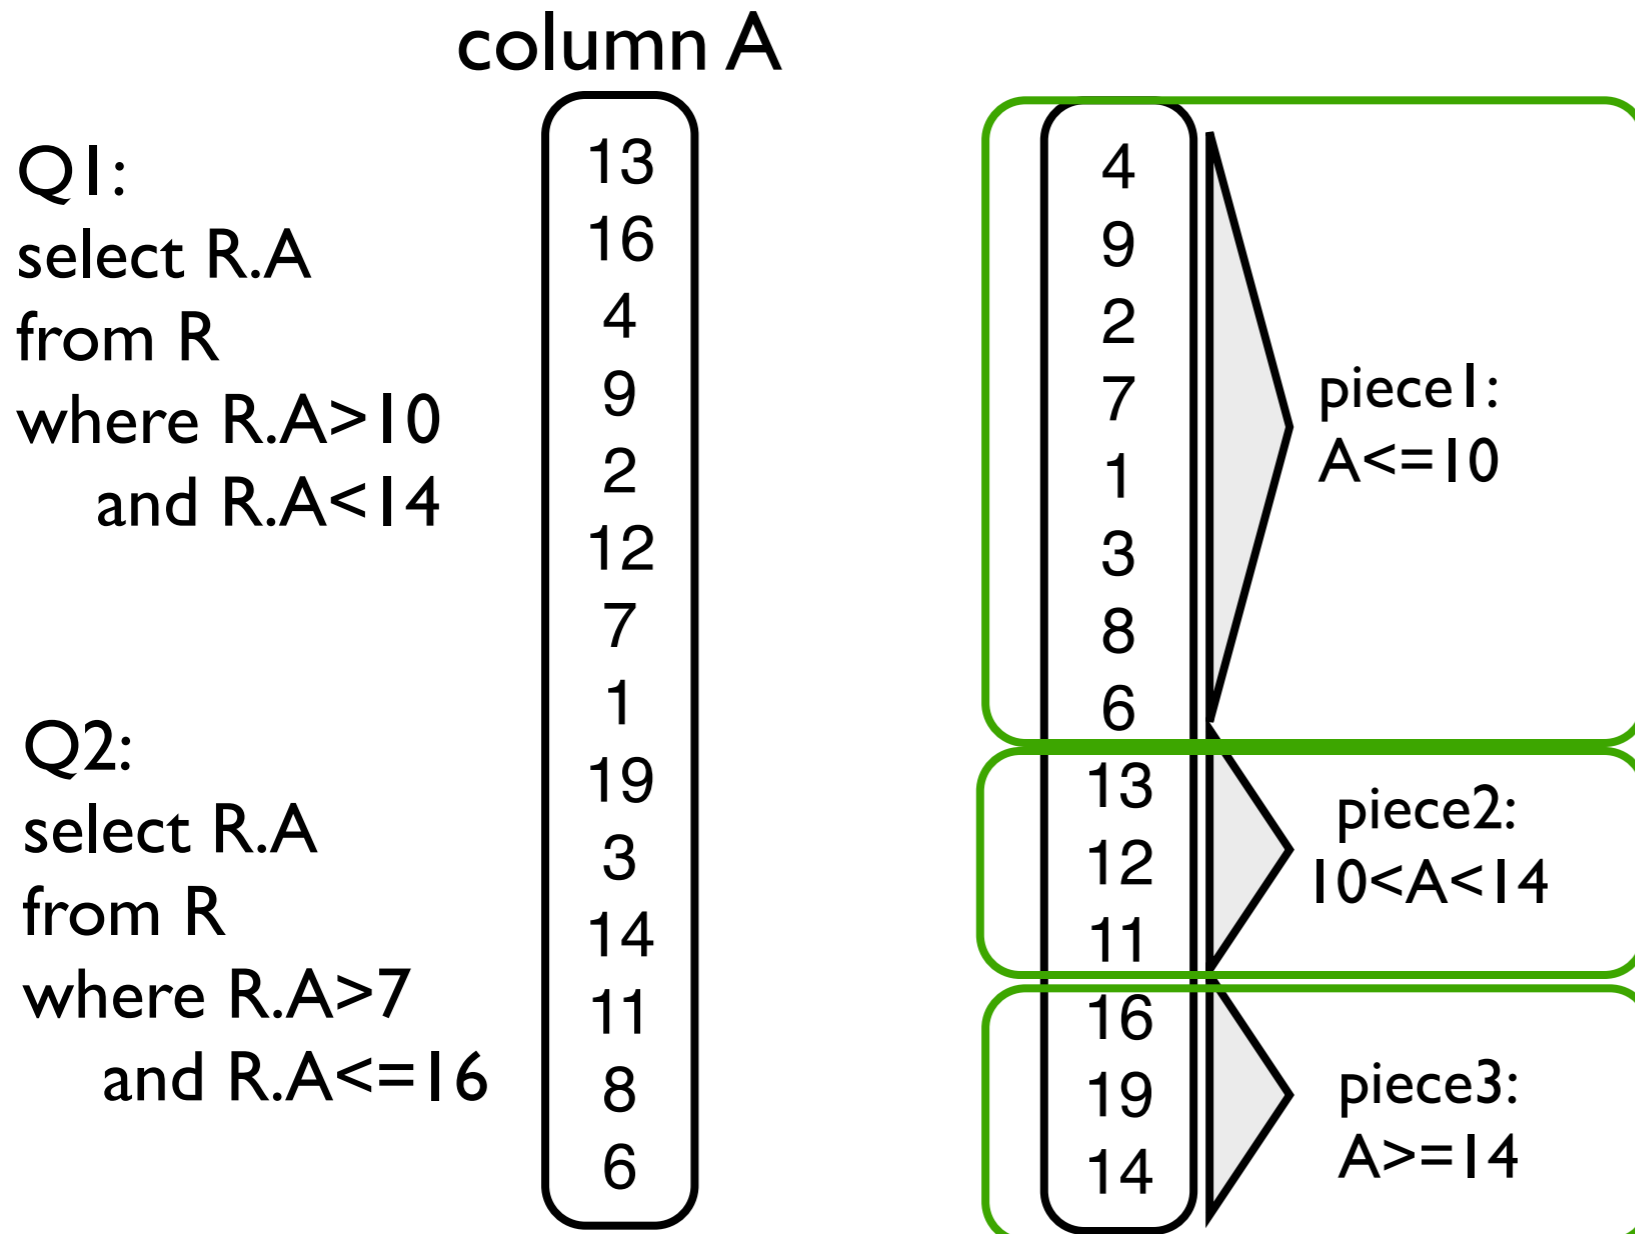
column A

13  
16  
4  
9  
2  
12  
7  
1  
19  
3  
14  
11  
8  
6



dynamically/on-the-fly within the select-operator

# cracking example



dynamically/on-the-fly within the select-operator


# cracking example

Q1:  
select R.A  
from R  
where R.A > 10  
and R.A < 14

Q2:  
select R.A  
from R  
where R.A > 7  
and R.A <= 16

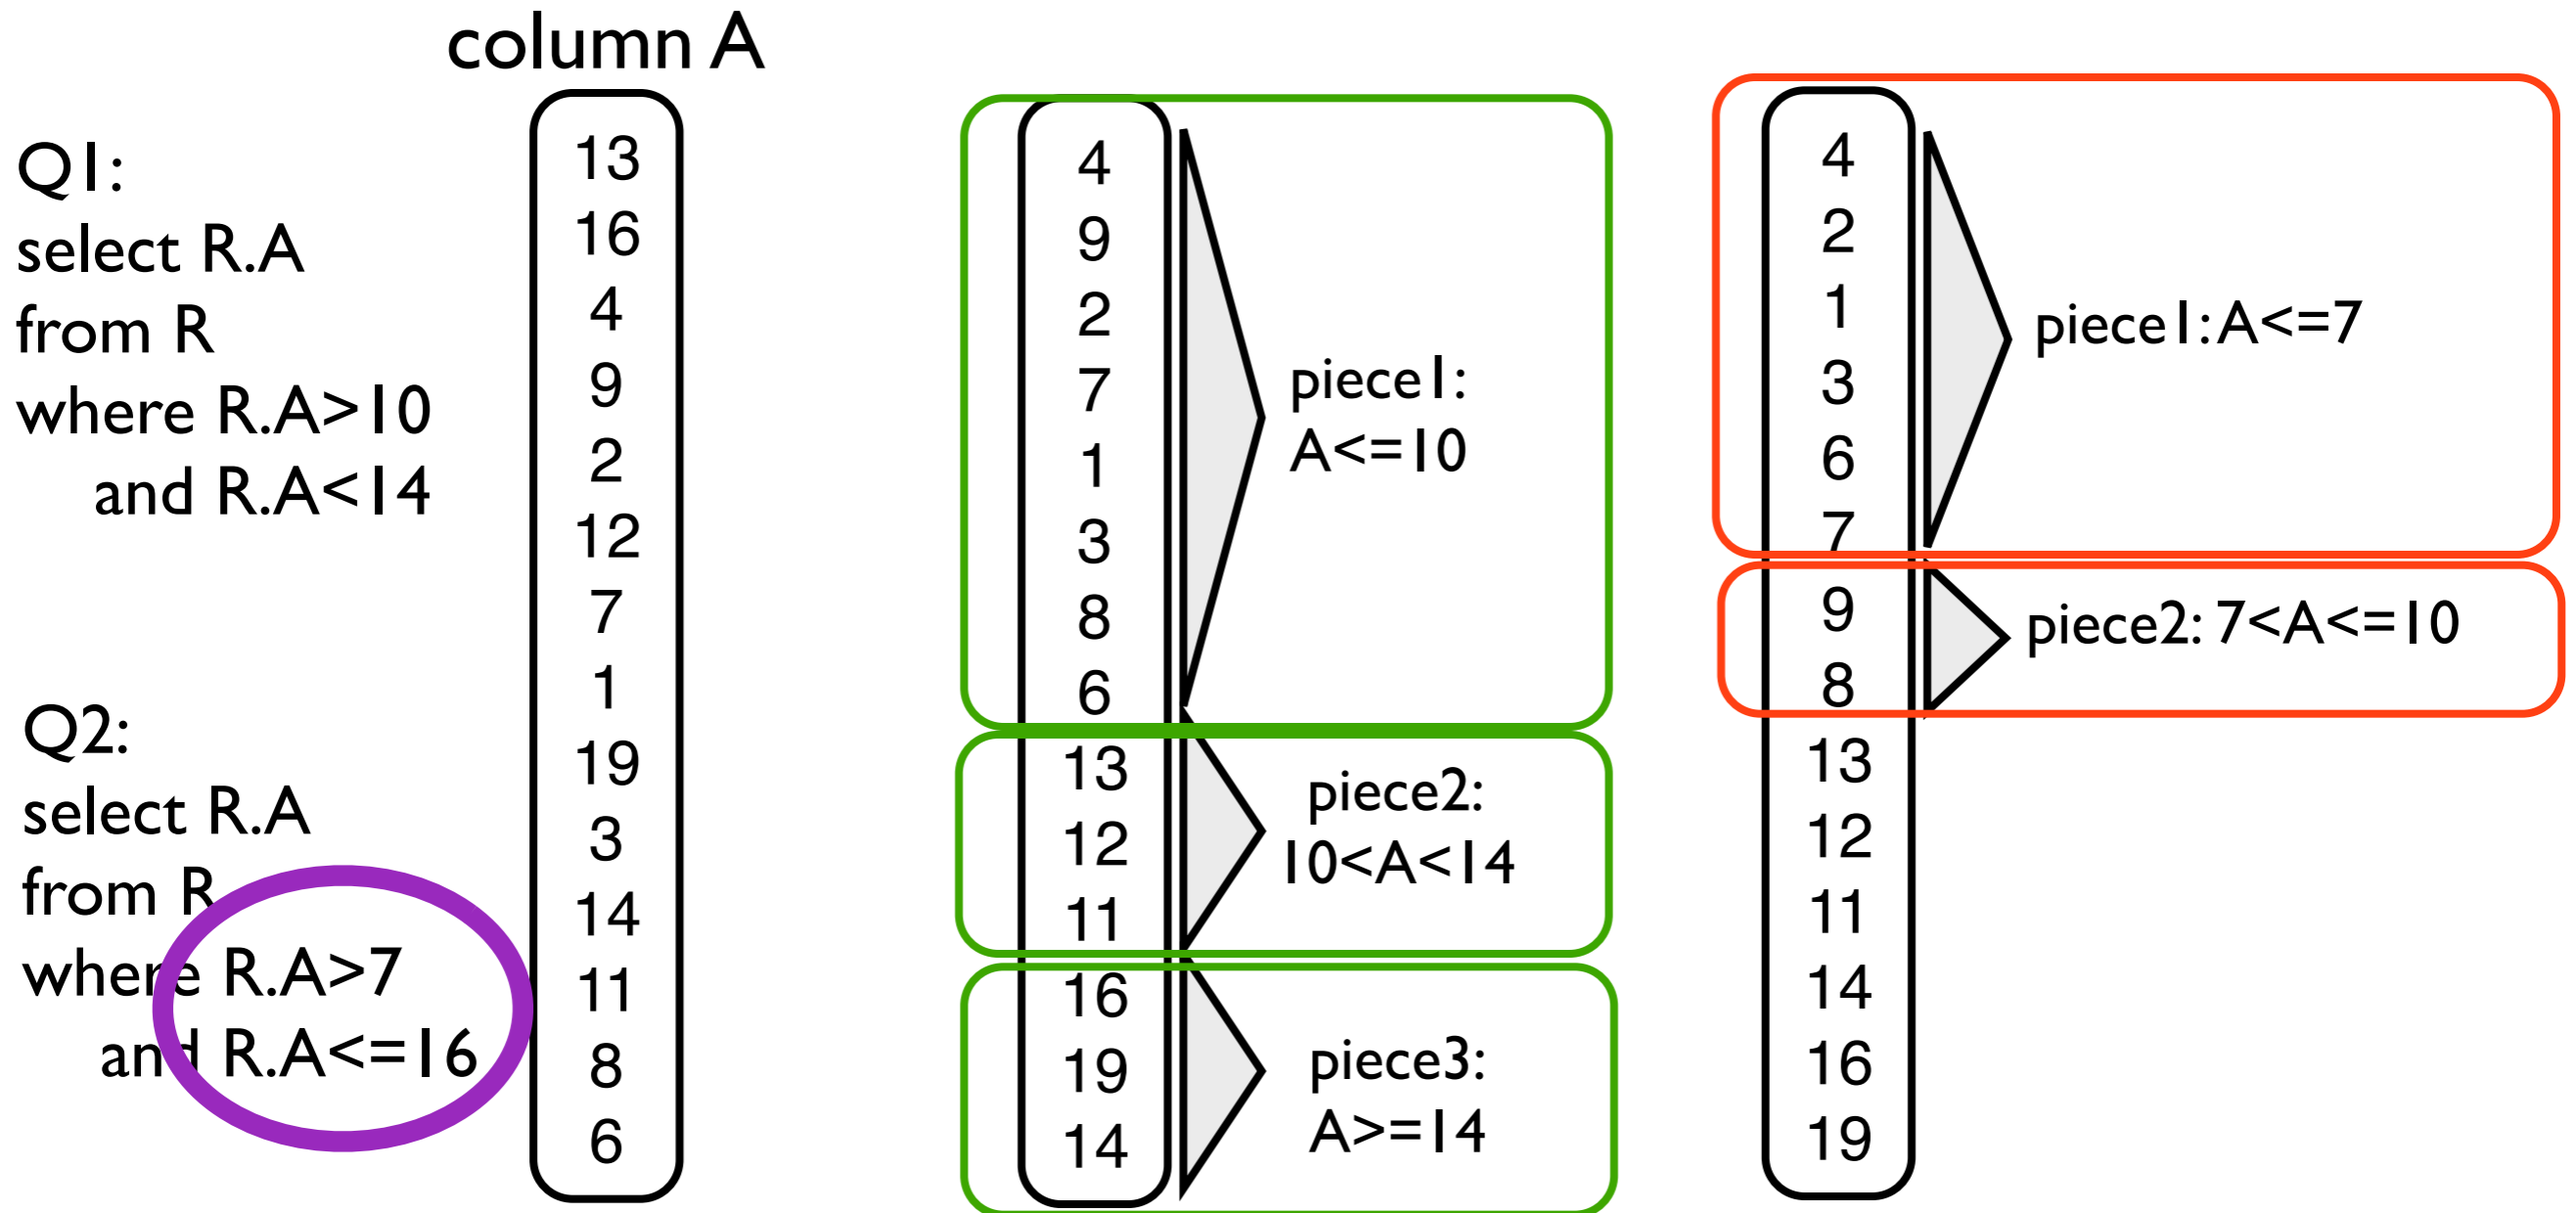
column A

13  
16  
4  
9  
2  
12  
7  
1  
19  
3  
14  
11  
8  
6



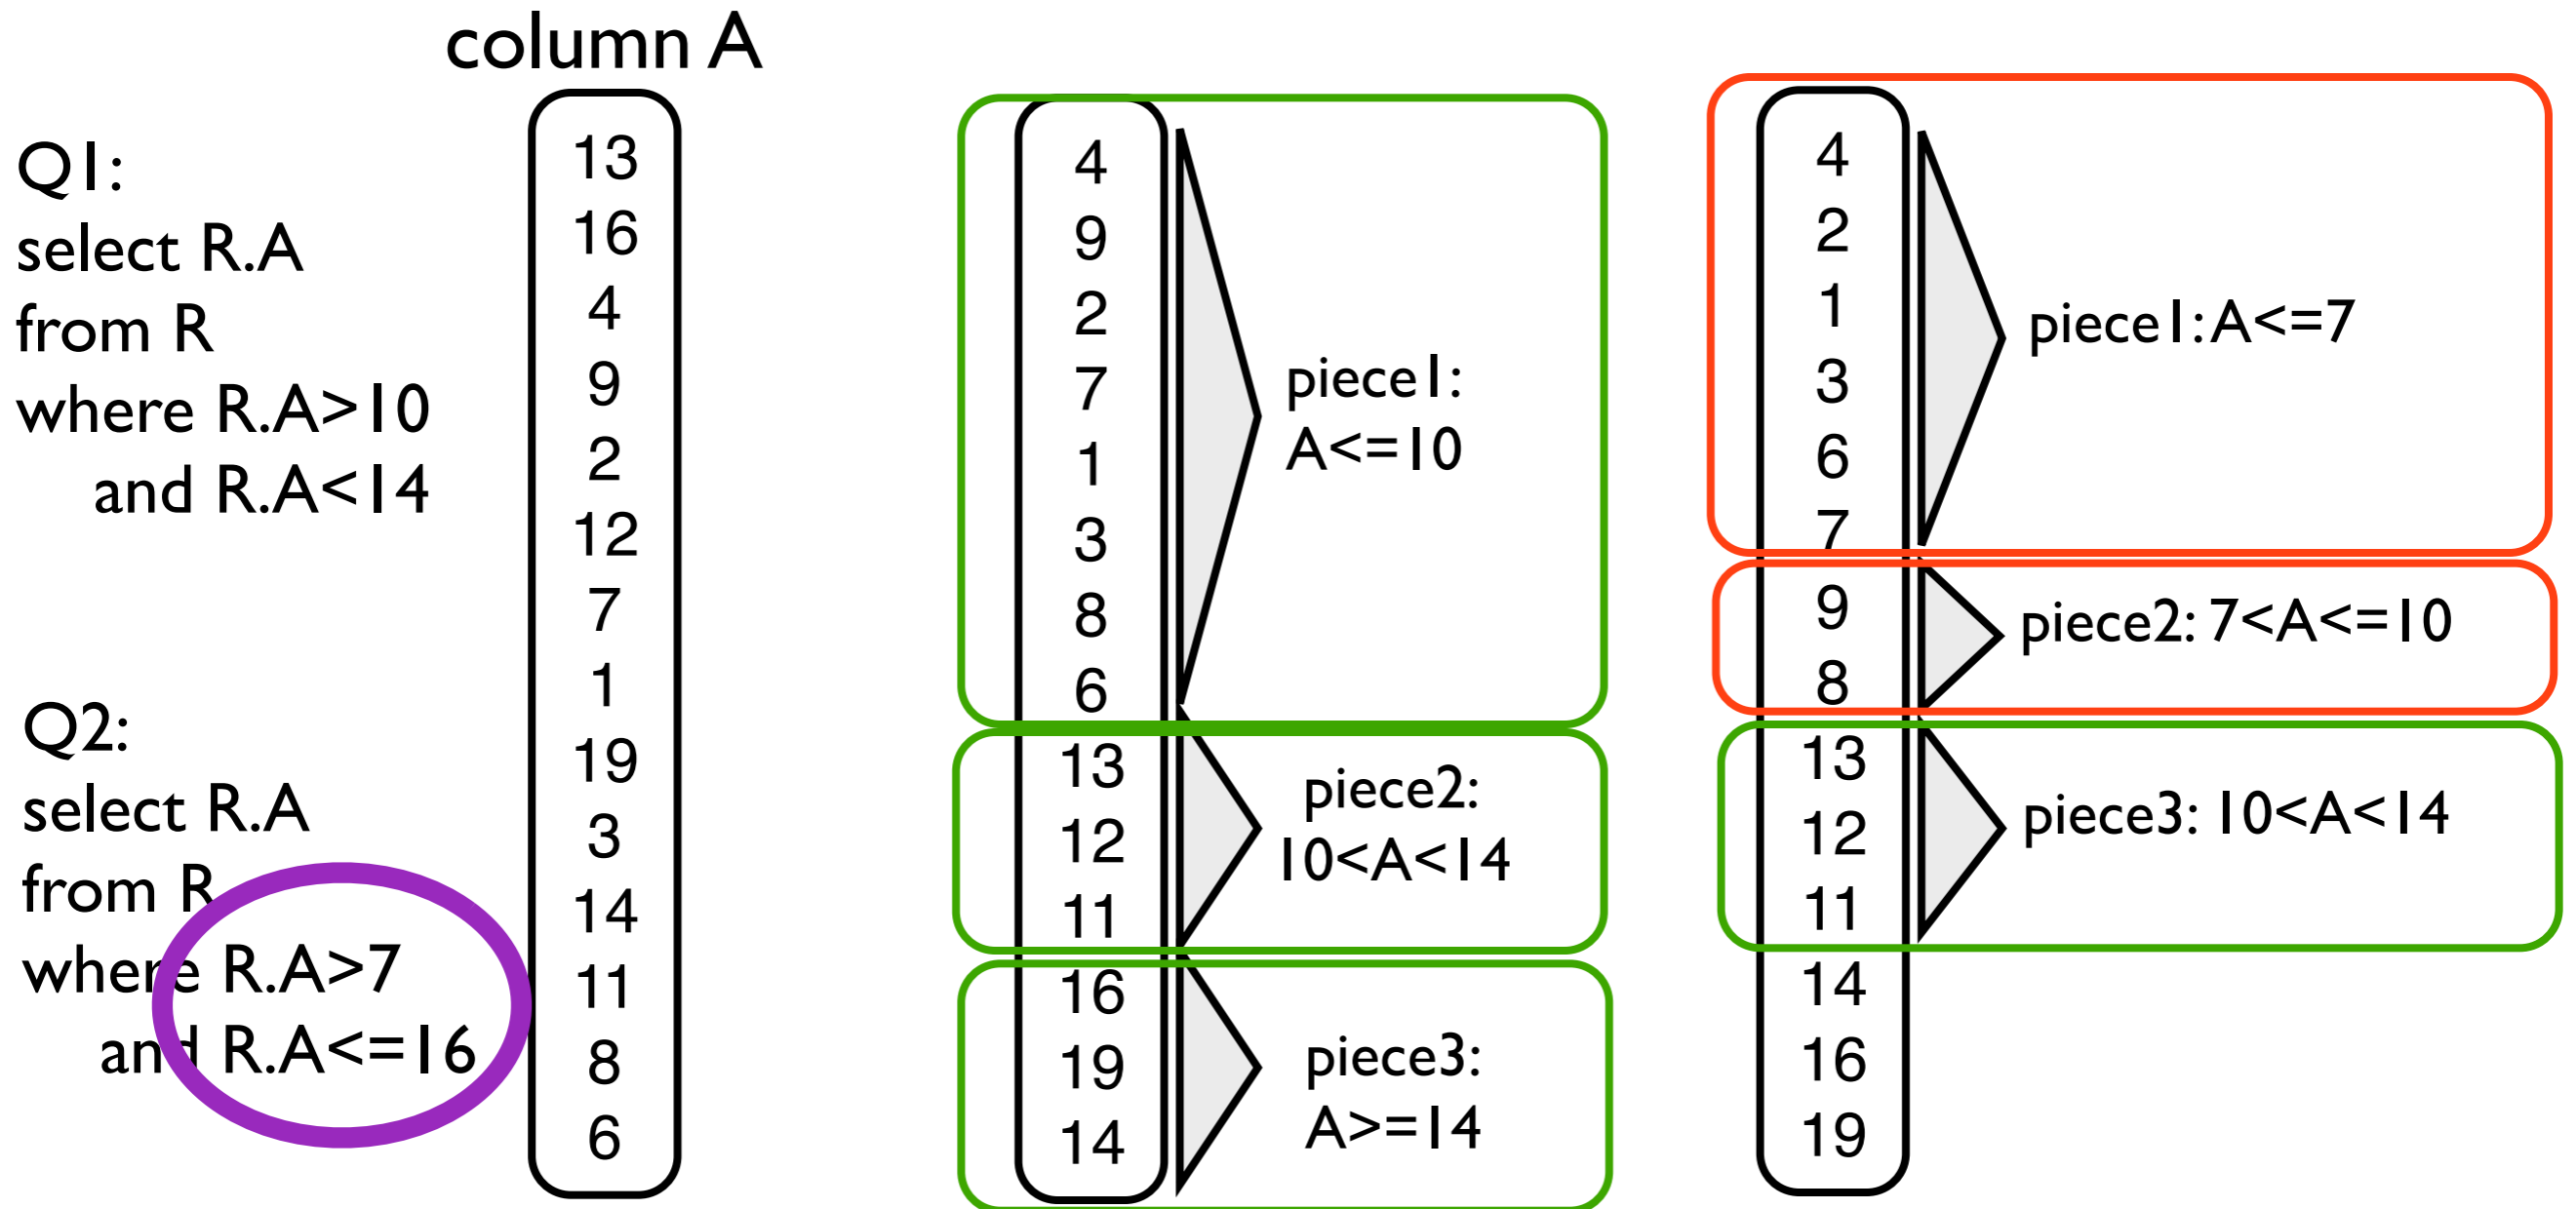
dynamically/on-the-fly within the select-operator

# cracking example



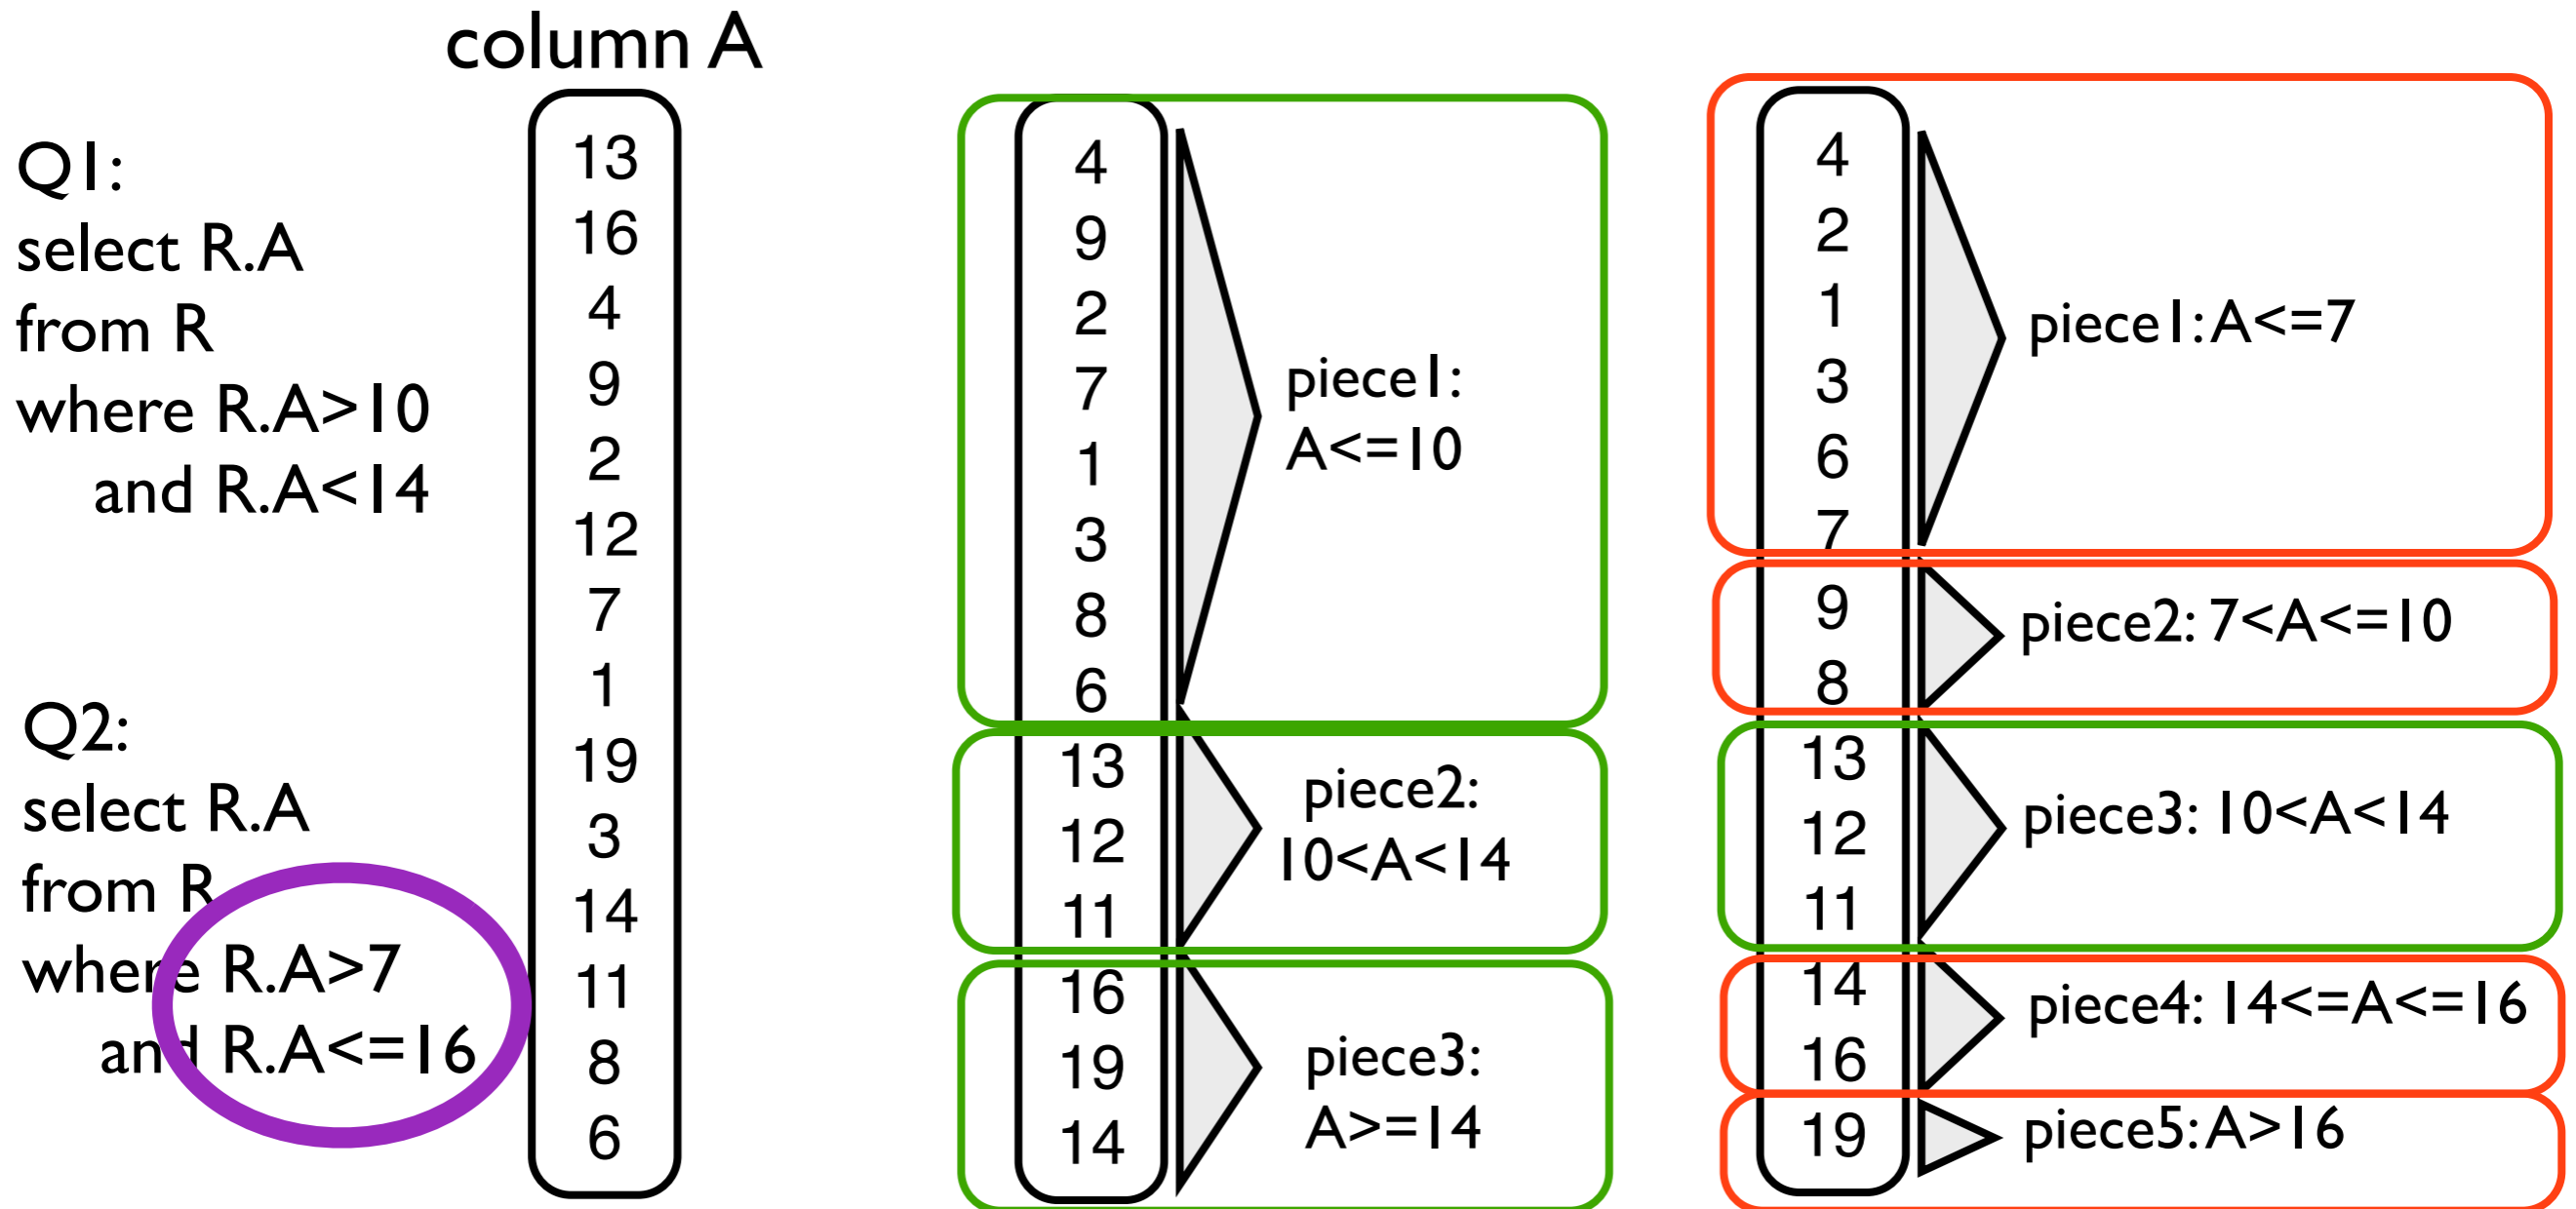
dynamically/on-the-fly within the select-operator

# cracking example




dynamically/on-the-fly within the select-operator

# cracking example




dynamically/on-the-fly within the select-operator

# cracking example



dynamically/on-the-fly within the select-operator

# the more we crack, the more we learn



dynamically/on-the-fly within the select-operator





**select [15,55]**



select [15,55]



select [15,55]

10

20


30

40

50

60






select [15,55]

10 20 30 40 50 60



select [15,55]



select [15,55]


10 20 30 40 50 60



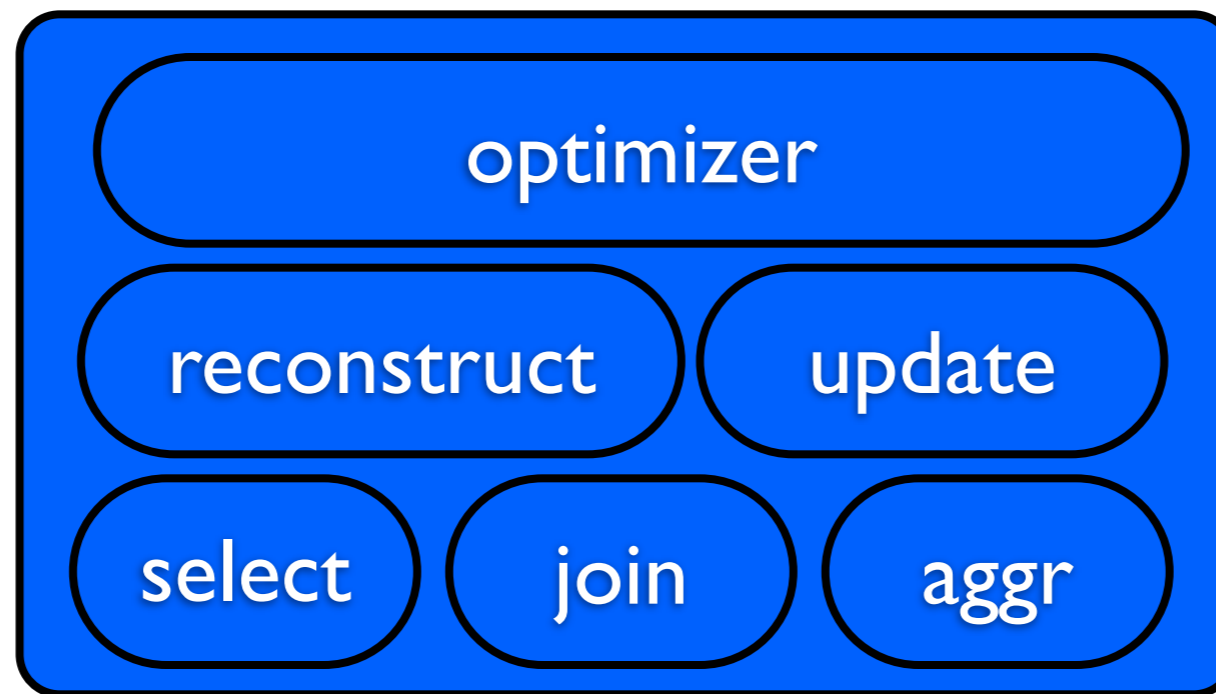
select [15,55]

**touch at most two pieces at a time**

**pieces become smaller and smaller**



*implemented in*  *monetdb*  
open-source column-store




**database kernel**

**code footprint  
monetdb 2M**

# continuous adaptation

## set-up

100K random selections  
 random selectivity  
 random value ranges  
 in a 10 million integer column




# continuous adaptation

## set-up

100K random selections  
 random selectivity  
 random value ranges  
 in a 10 million integer column

**almost no  
 initialization overhead**




# continuous adaptation

## set-up

100K random selections  
 random selectivity  
 random value ranges  
 in a 10 million integer column

**almost no  
 initialization overhead  
 continuous improvement**




# continuous adaptation

## set-up

100K random selections  
 random selectivity  
 random value ranges  
 in a 10 million integer column

**almost no  
 initialization overhead**


**continuous improvement**



# continuous adaptation

## set-up


10K random selections  
selectivity 10%  
random value ranges  
in a 30 million integer column



# continuous adaptation

## set-up


10K random selections  
selectivity 10%  
random value ranges  
in a 30 million integer column



# continuous adaptation

## set-up

10K random selections  
selectivity 10%  
random value ranges  
in a 30 million integer column




**10K queries later,  
Full Index still has not  
amortized the initialization costs**

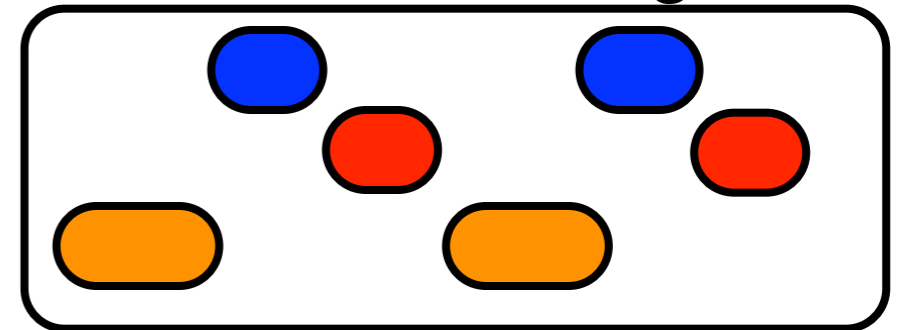
# traditional databases

monolithic/full indexing

*offline indexing*



*online indexing*




workload analysis  
index building  
query processing


# traditional databases

monolithic/full indexing

*offline indexing*



*online indexing*



workload analysis  
index building  
query processing

## database cracking

partial/adaptive/continuous indexing

*adaptive indexing*




table I






table I




select R.A from R where R.A>10 and R.A<14



select R.A from R where R.A>10 and R.A<14

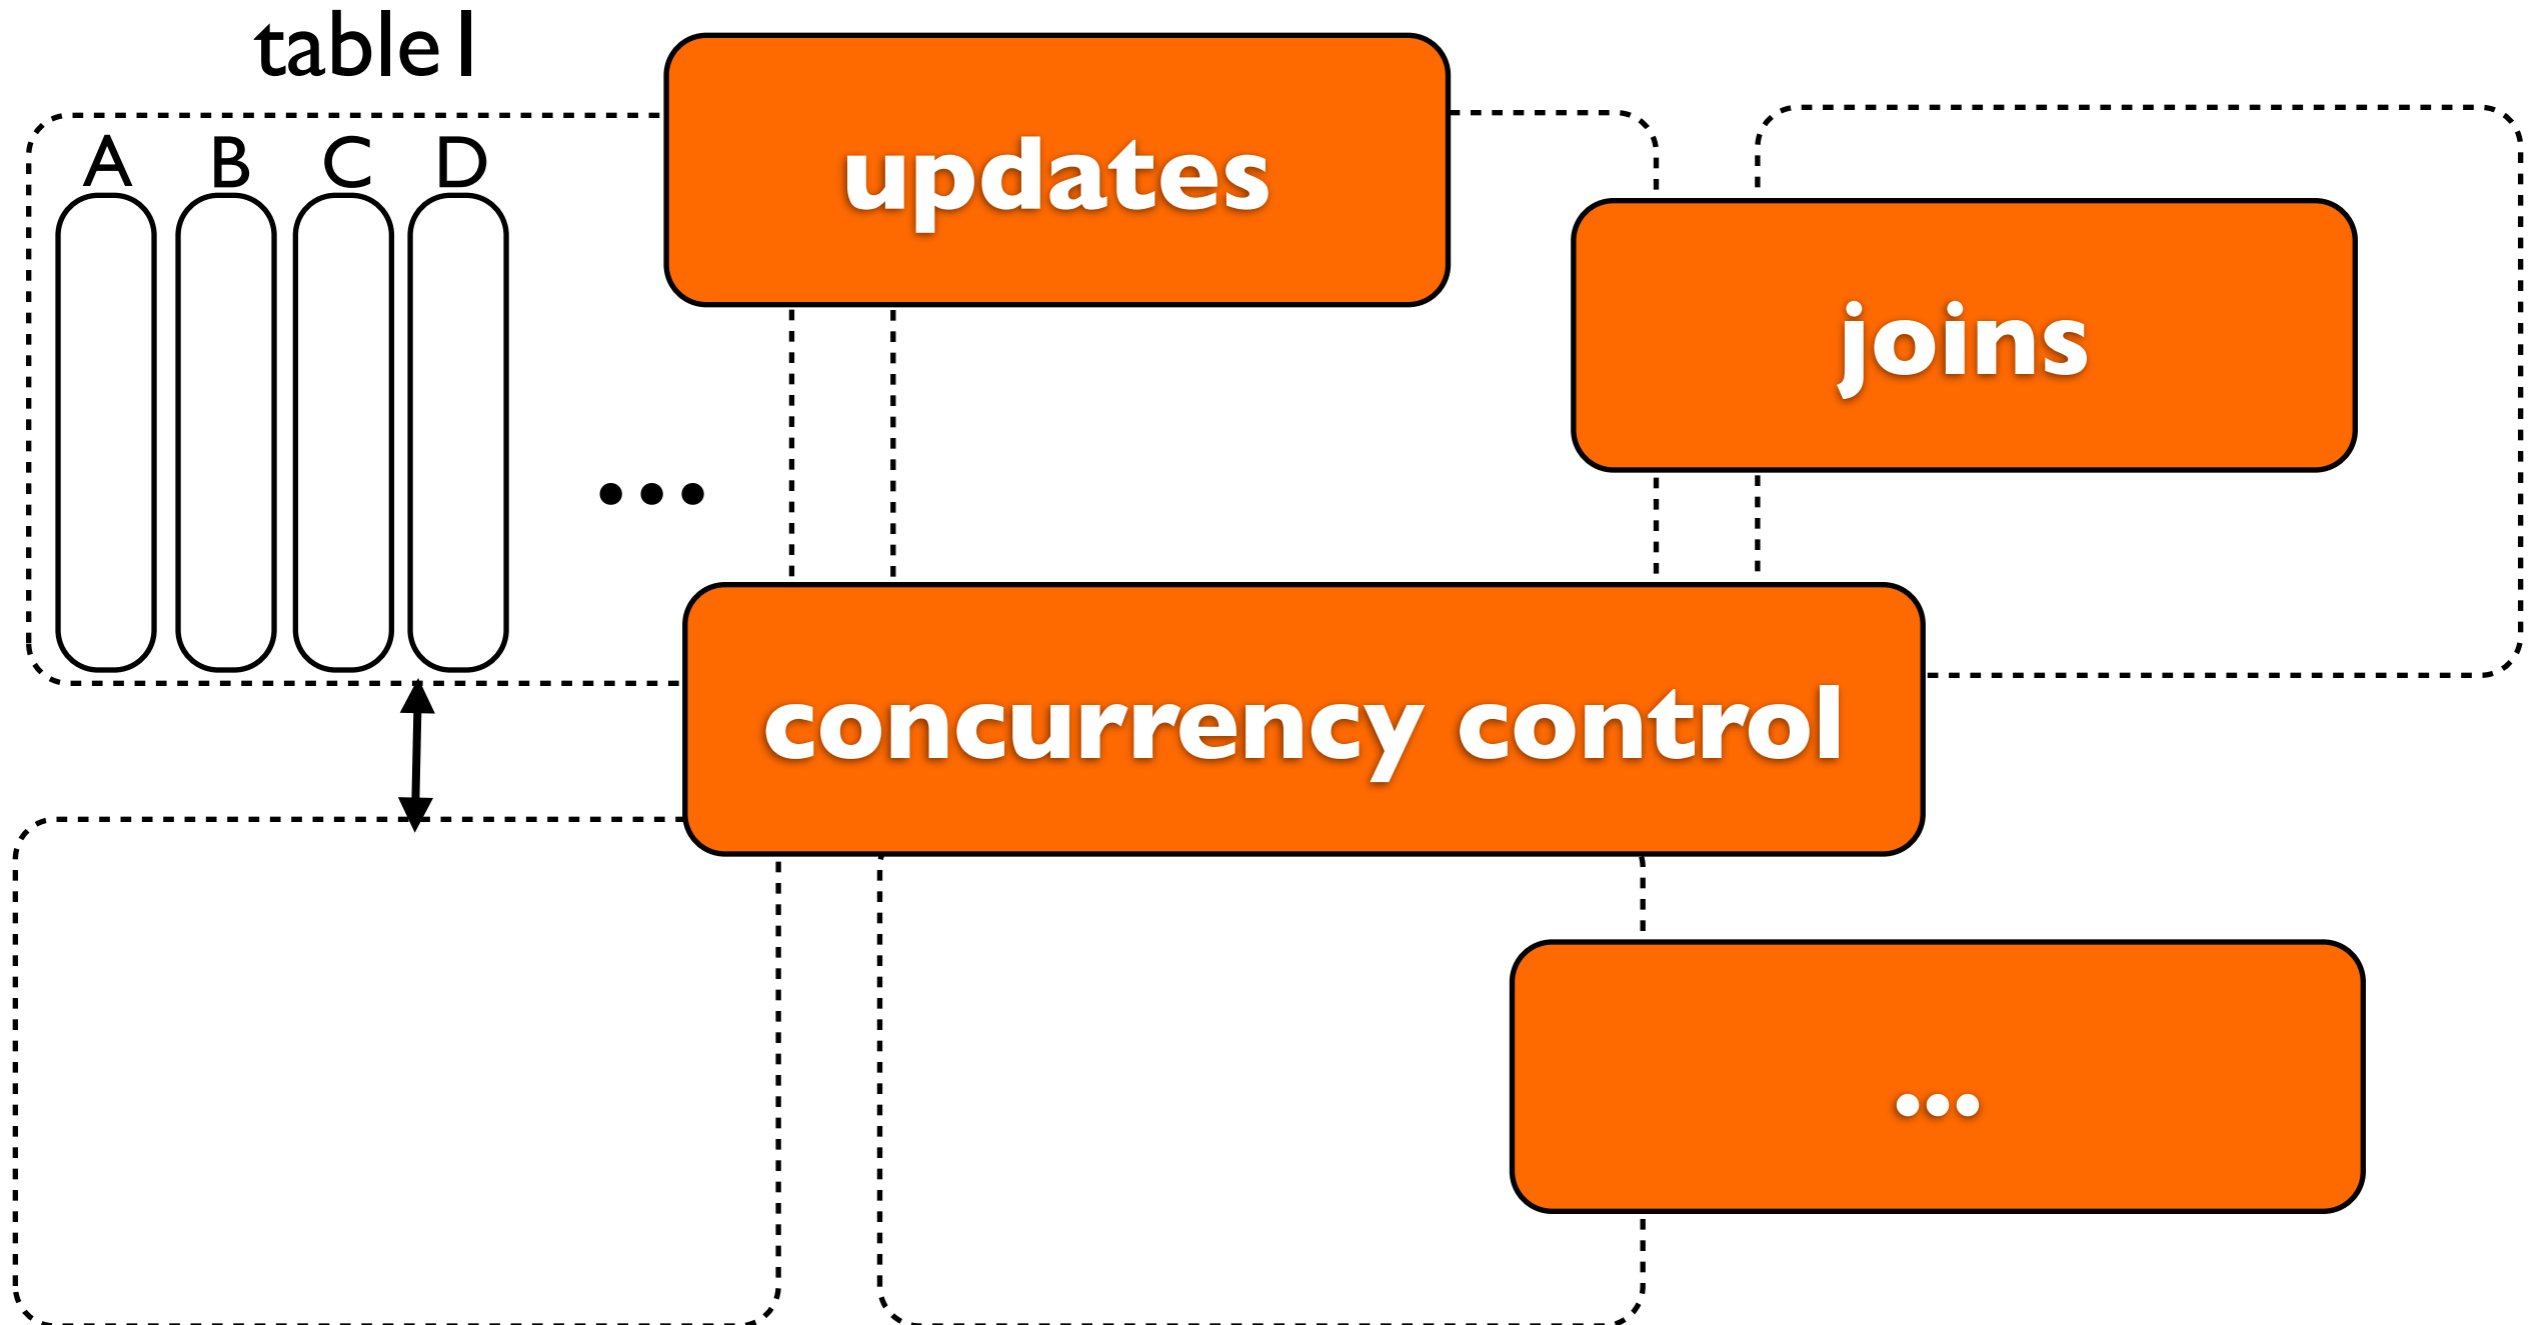
**select max(R.A),max(R.B),max(S.A),max(S.B) from R,S  
where v1 <R.C<v2 and v3 <R.D<v4  
and v5 <R.E<v6 and k1 <S.C<k2 and k3 <S.D<k4 and k5 <S.E<k6  
and R.F = S.F**

table I



select R.A from R where R.A>10 and R.A<14

select max(R.A),max(R.B),max(S.A),max(S.B) from R,S  
where v1 <R.C<v2 and v3 <R.D<v4  
and v5 <R.E<v6 and k1 <S.C<k2 and k3 <S.D<k4 and k5 <S.E<k6  
and R.F = S.F



# cracking tanager

**base data**

*as queries arrive...*

table 1






table 2



# cracking tangram

**base data**

table 1



**as queries arrive...**

table 1




table 2





table 2



# cracking tangram

**base data**

**as queries arrive...**

***partial materialization***

table 1




table 1




table 2






table 2



# cracking tangram


**base data**

table 1



**as queries arrive...**

table 1



***partial materialization***

***partial indexing***

table 2






table 2



# cracking tangram

**base data**

table 1



**as queries arrive...**

table 1




table 2





table 2



***partial materialization***


***partial indexing***

***continuous adaptation***

# cracking tangram

**base data**

table 1



**as queries arrive...**

table 1




table 2





table 2



**partial materialization**

**partial indexing**


**continuous adaptation**

**storage adaptation**

# cracking tangram


**base data**

table 1



**as queries arrive...**

table 1



***partial materialization***

***partial indexing***

***continuous adaptation***

***storage adaptation***

table 2






table 2



# cracking tangram

**base data**

table 1



**as queries arrive...**

table 1




table 2





table 2



***partial materialization***

***partial indexing***

***continuous adaptation***


***storage adaptation***

***no tuple reconstruction***

# cracking tangram

**base data**

table 1



**as queries arrive...**

table 1




table 2





table 2



**partial materialization**

**partial indexing**

**continuous adaptation**

**storage adaptation**


**no tuple reconstruction**

**adaptive alignment**

# cracking tangram

**base data**

table 1



**as queries arrive...**

table 1




table 2


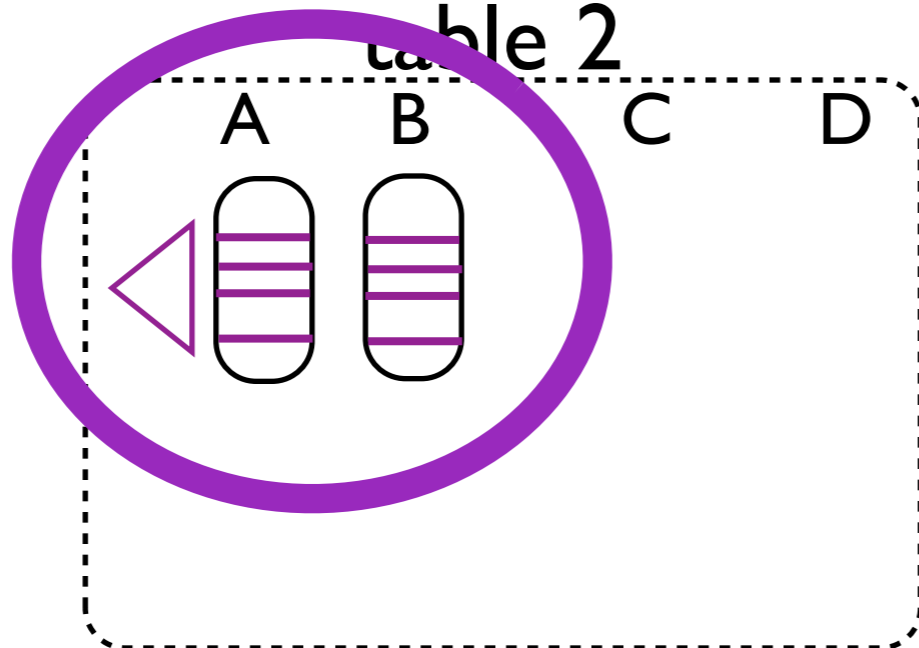


table 2



***partial materialization***

***partial indexing***

***continuous adaptation***

***storage adaptation***


***no tuple reconstruction***

***adaptive alignment***

# cracking tangram

**base data**

table 1



**as queries arrive...**

table 1




table 2





table 2



**partial materialization**

**partial indexing**

**continuous adaptation**

**storage adaptation**

**no tuple reconstruction**


**adaptive alignment**

**sort in caches**

# cracking tangram

**base data**

table 1



**as queries arrive...**

table 1




table 2


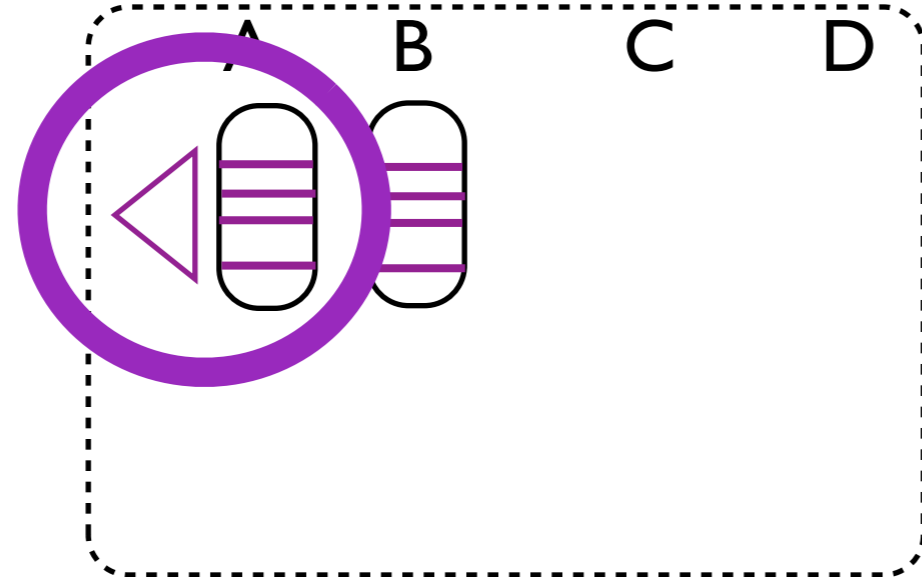


table 2



**partial materialization**

**partial indexing**

**continuous adaptation**

**storage adaptation**

**no tuple reconstruction**

**adaptive alignment**


**sort in caches**

**crack joins**

# cracking tangram

**base data**

table 1



**as queries arrive...**

table 1




table 2





table 2



**partial materialization**

**partial indexing**

**continuous adaptation**

**storage adaptation**

**no tuple reconstruction**

**adaptive alignment**

**sort in caches**


**crack joins**

**lightweight locking**

# cracking tangram

**base data**

table 1



**as queries arrive...**

table 1




table 2





table 2



**partial materialization**

**partial indexing**

**continuous adaptation**

**storage adaptation**

**no tuple reconstruction**

**adaptive alignment**

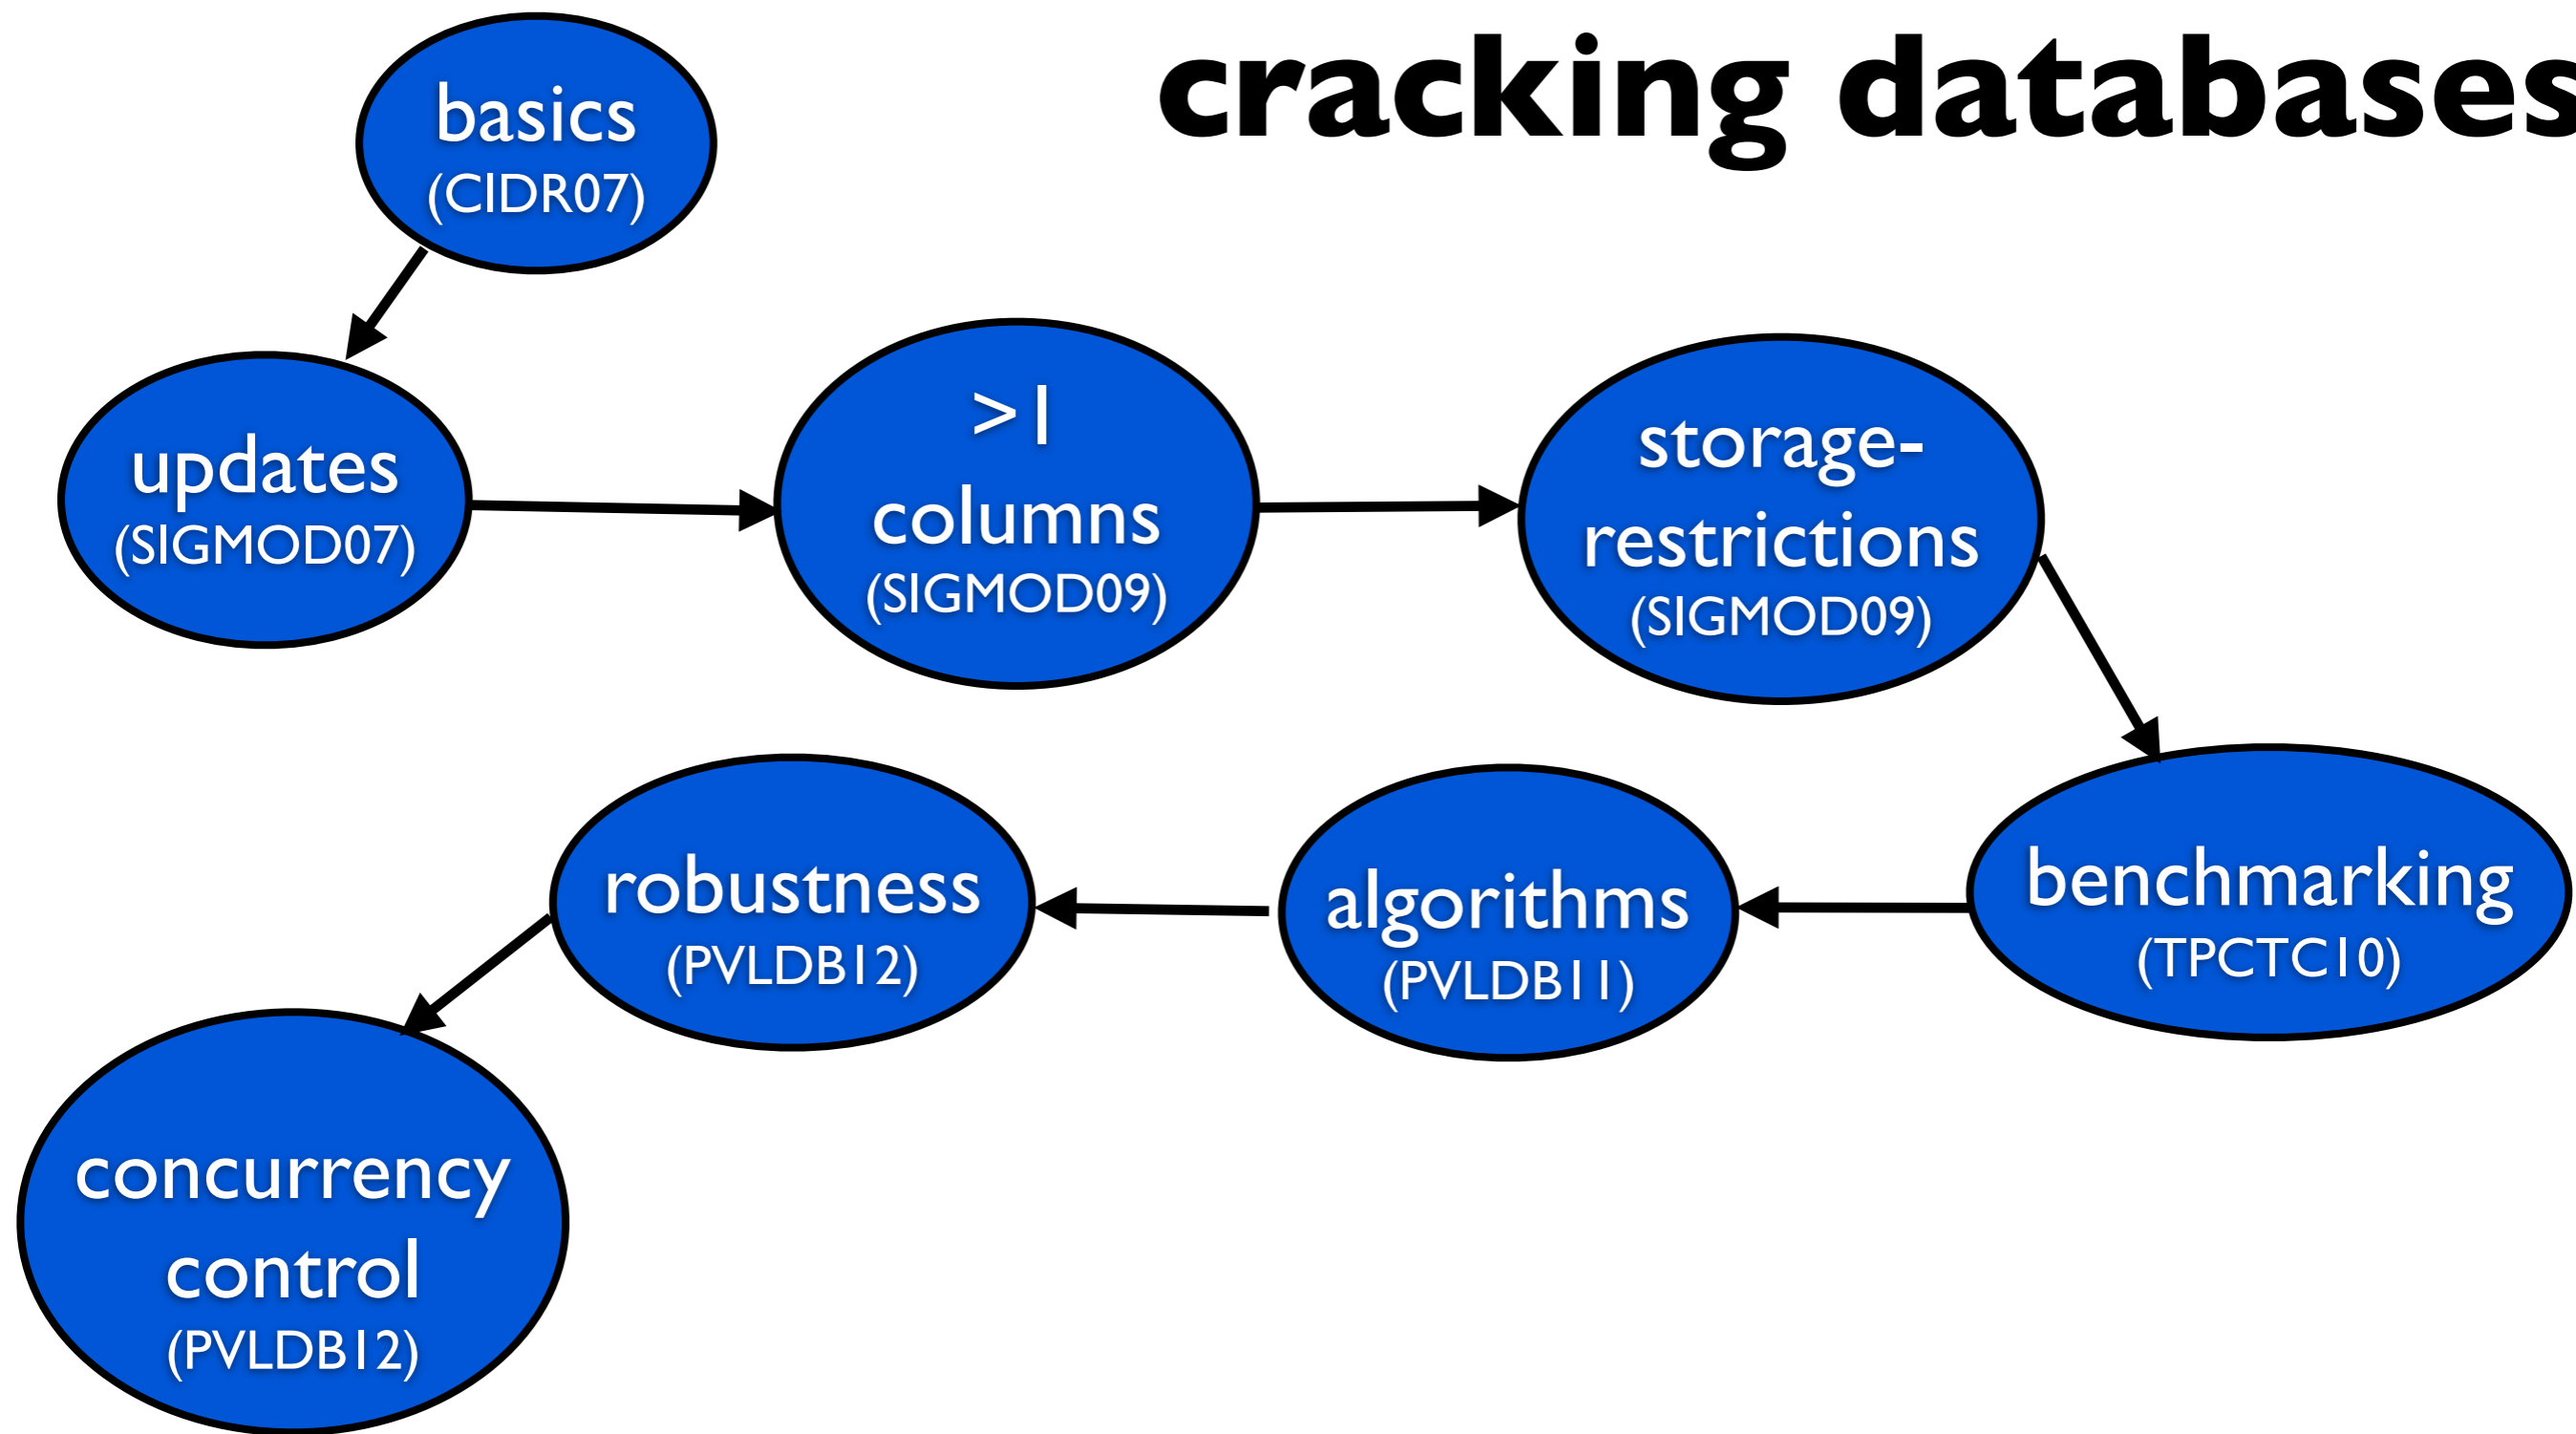
**sort in caches**

**crack joins**

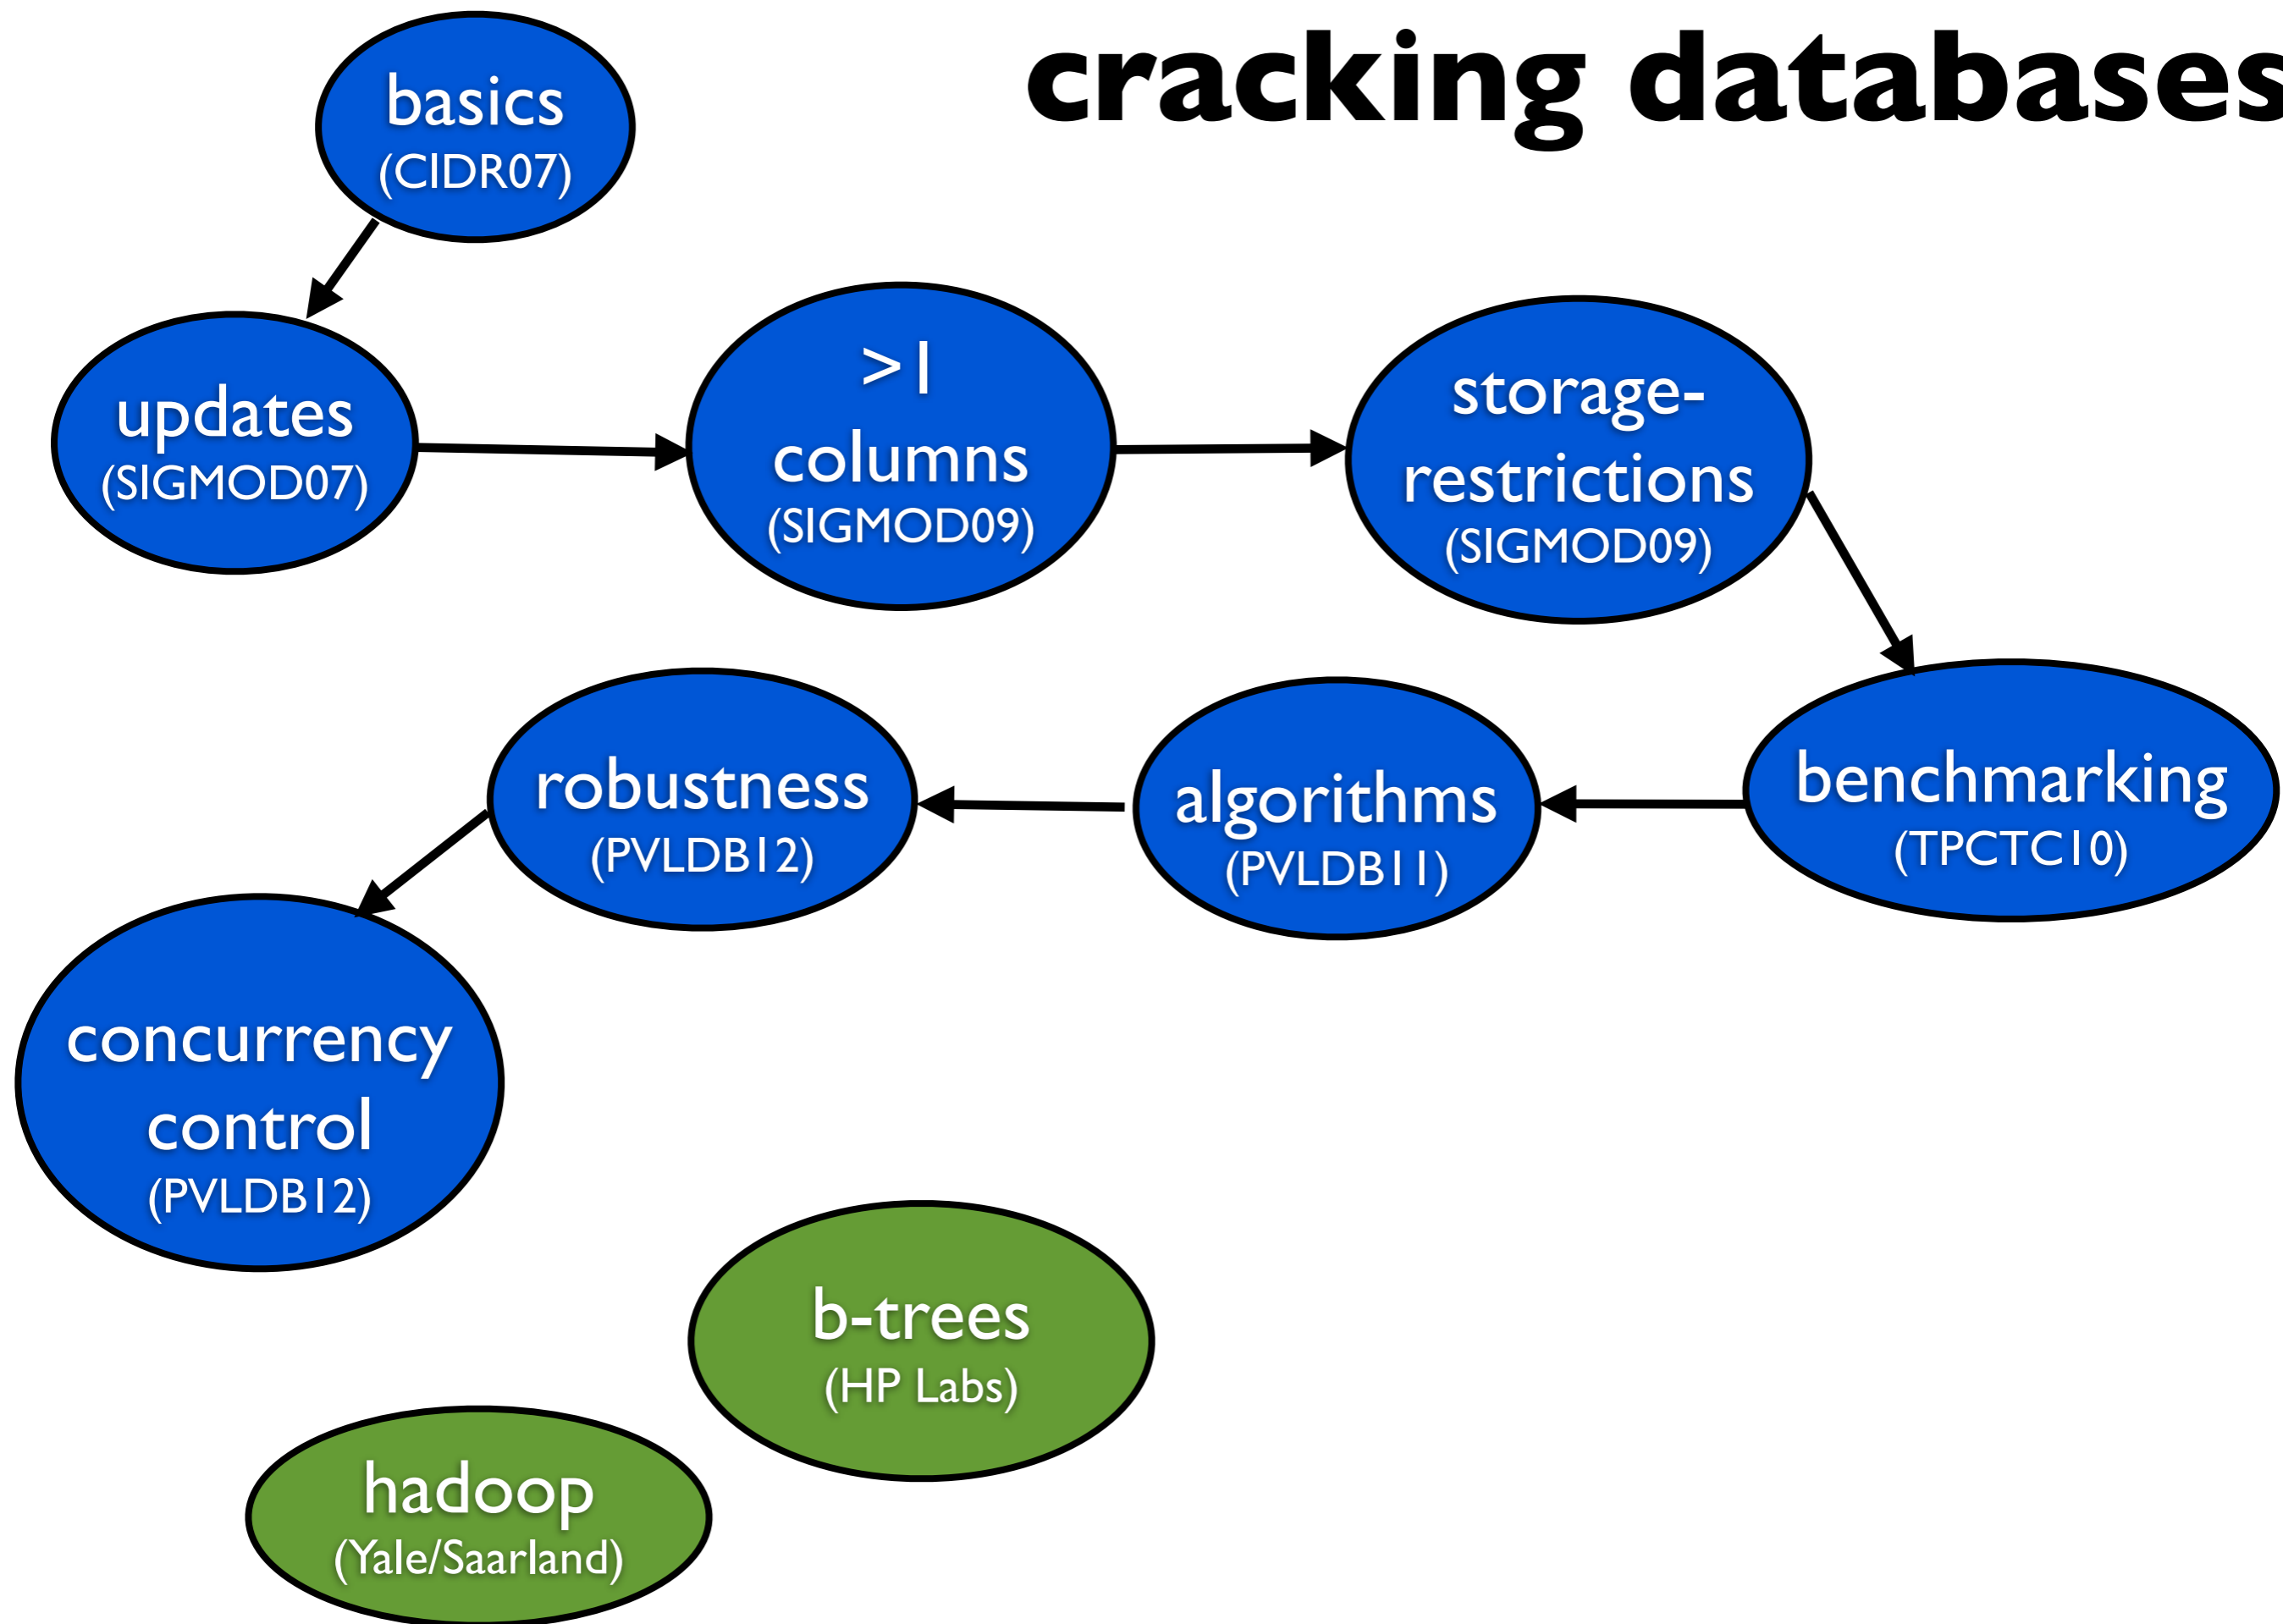
**lightweight locking**

**stochastic cracking**

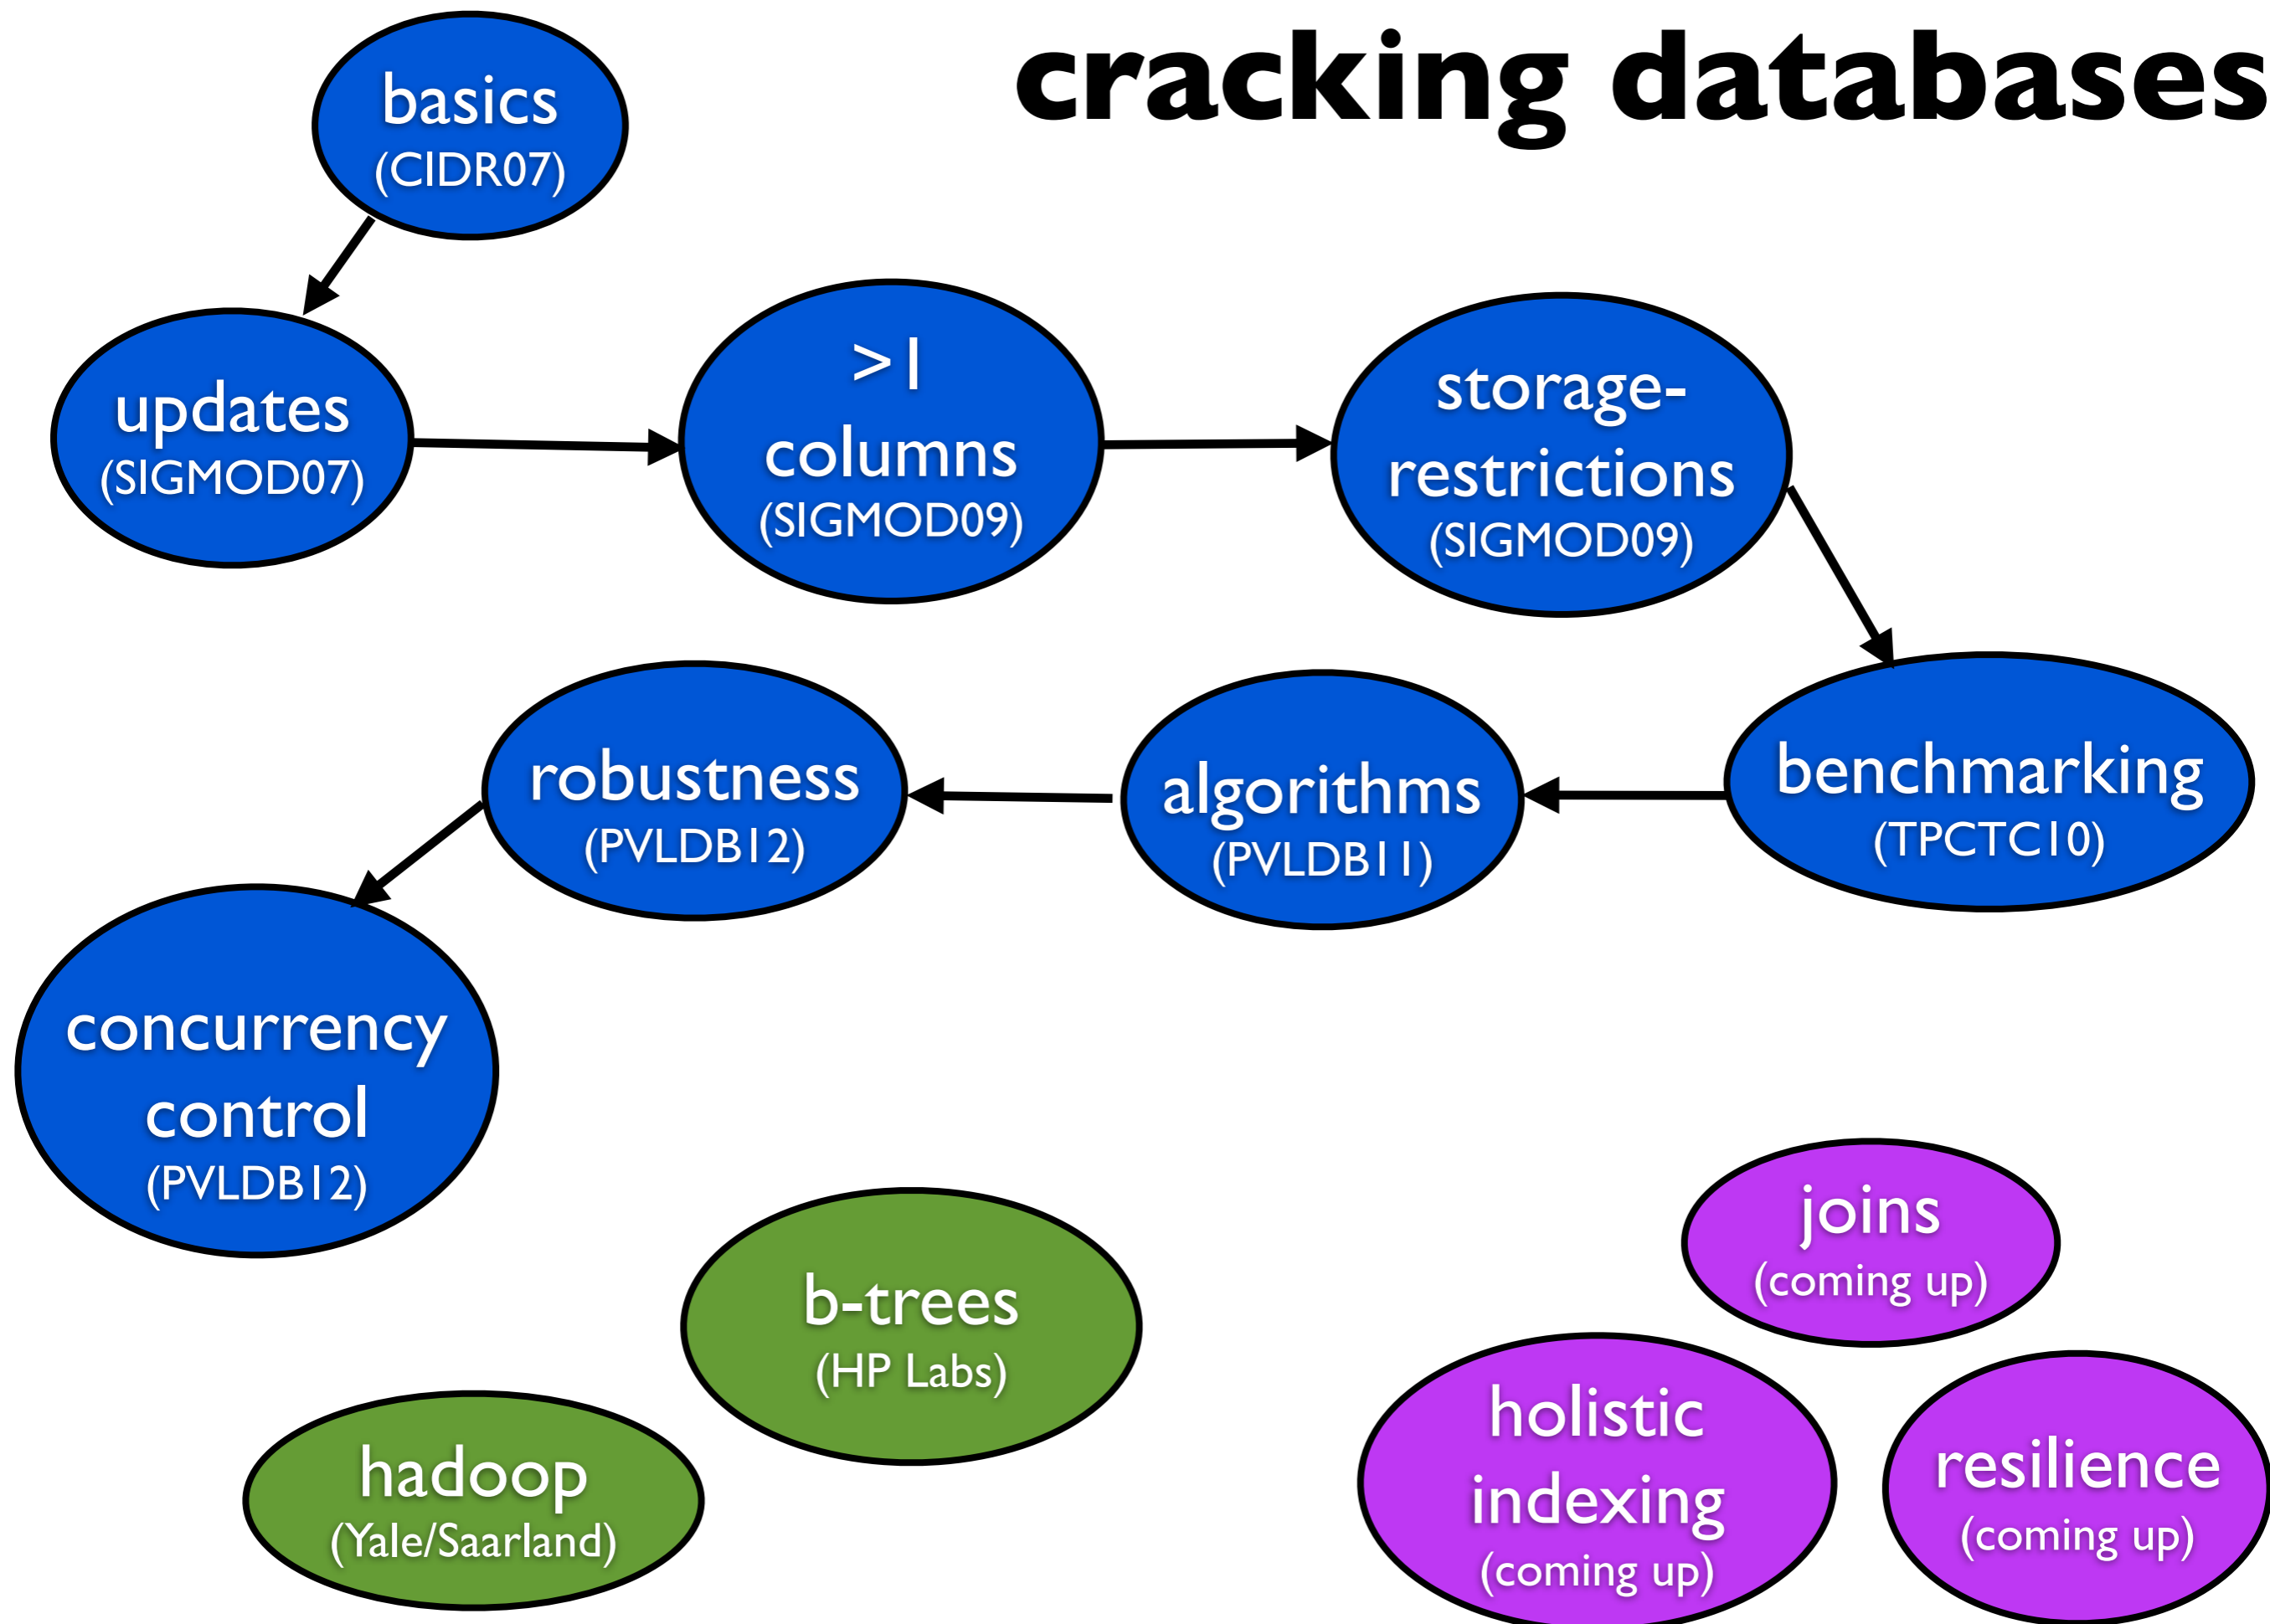
# cracking databases



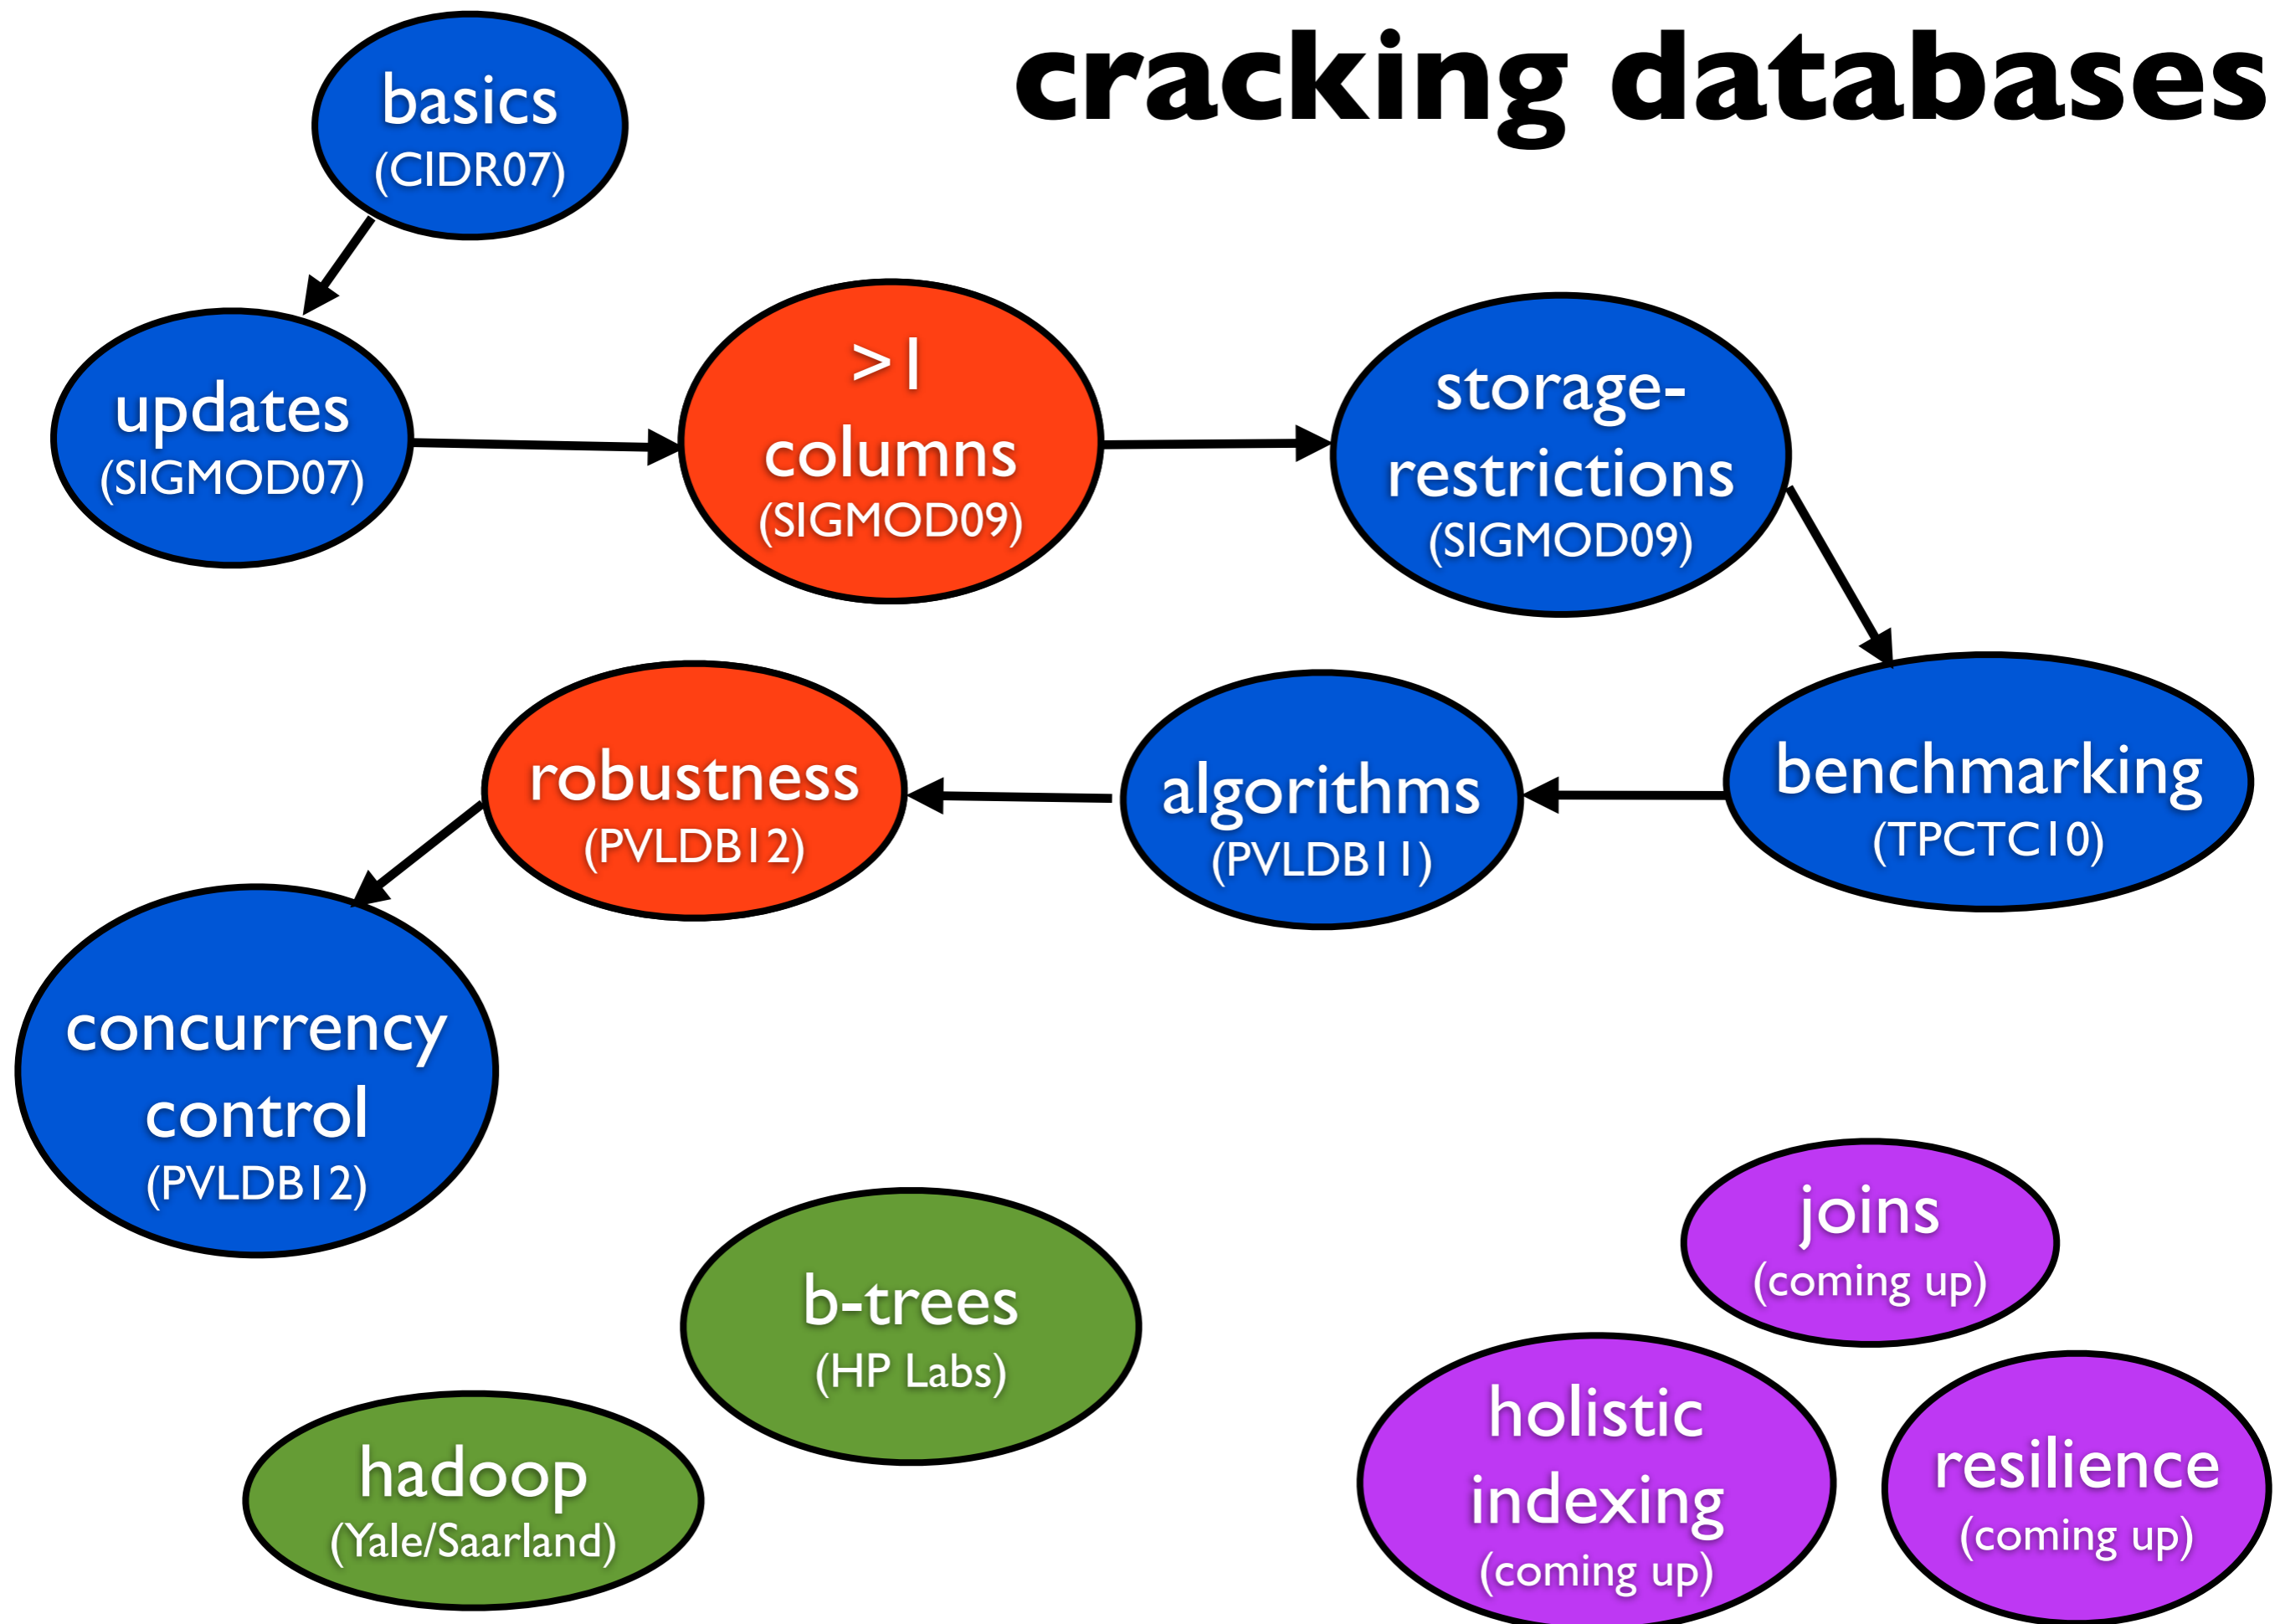
# cracking databases



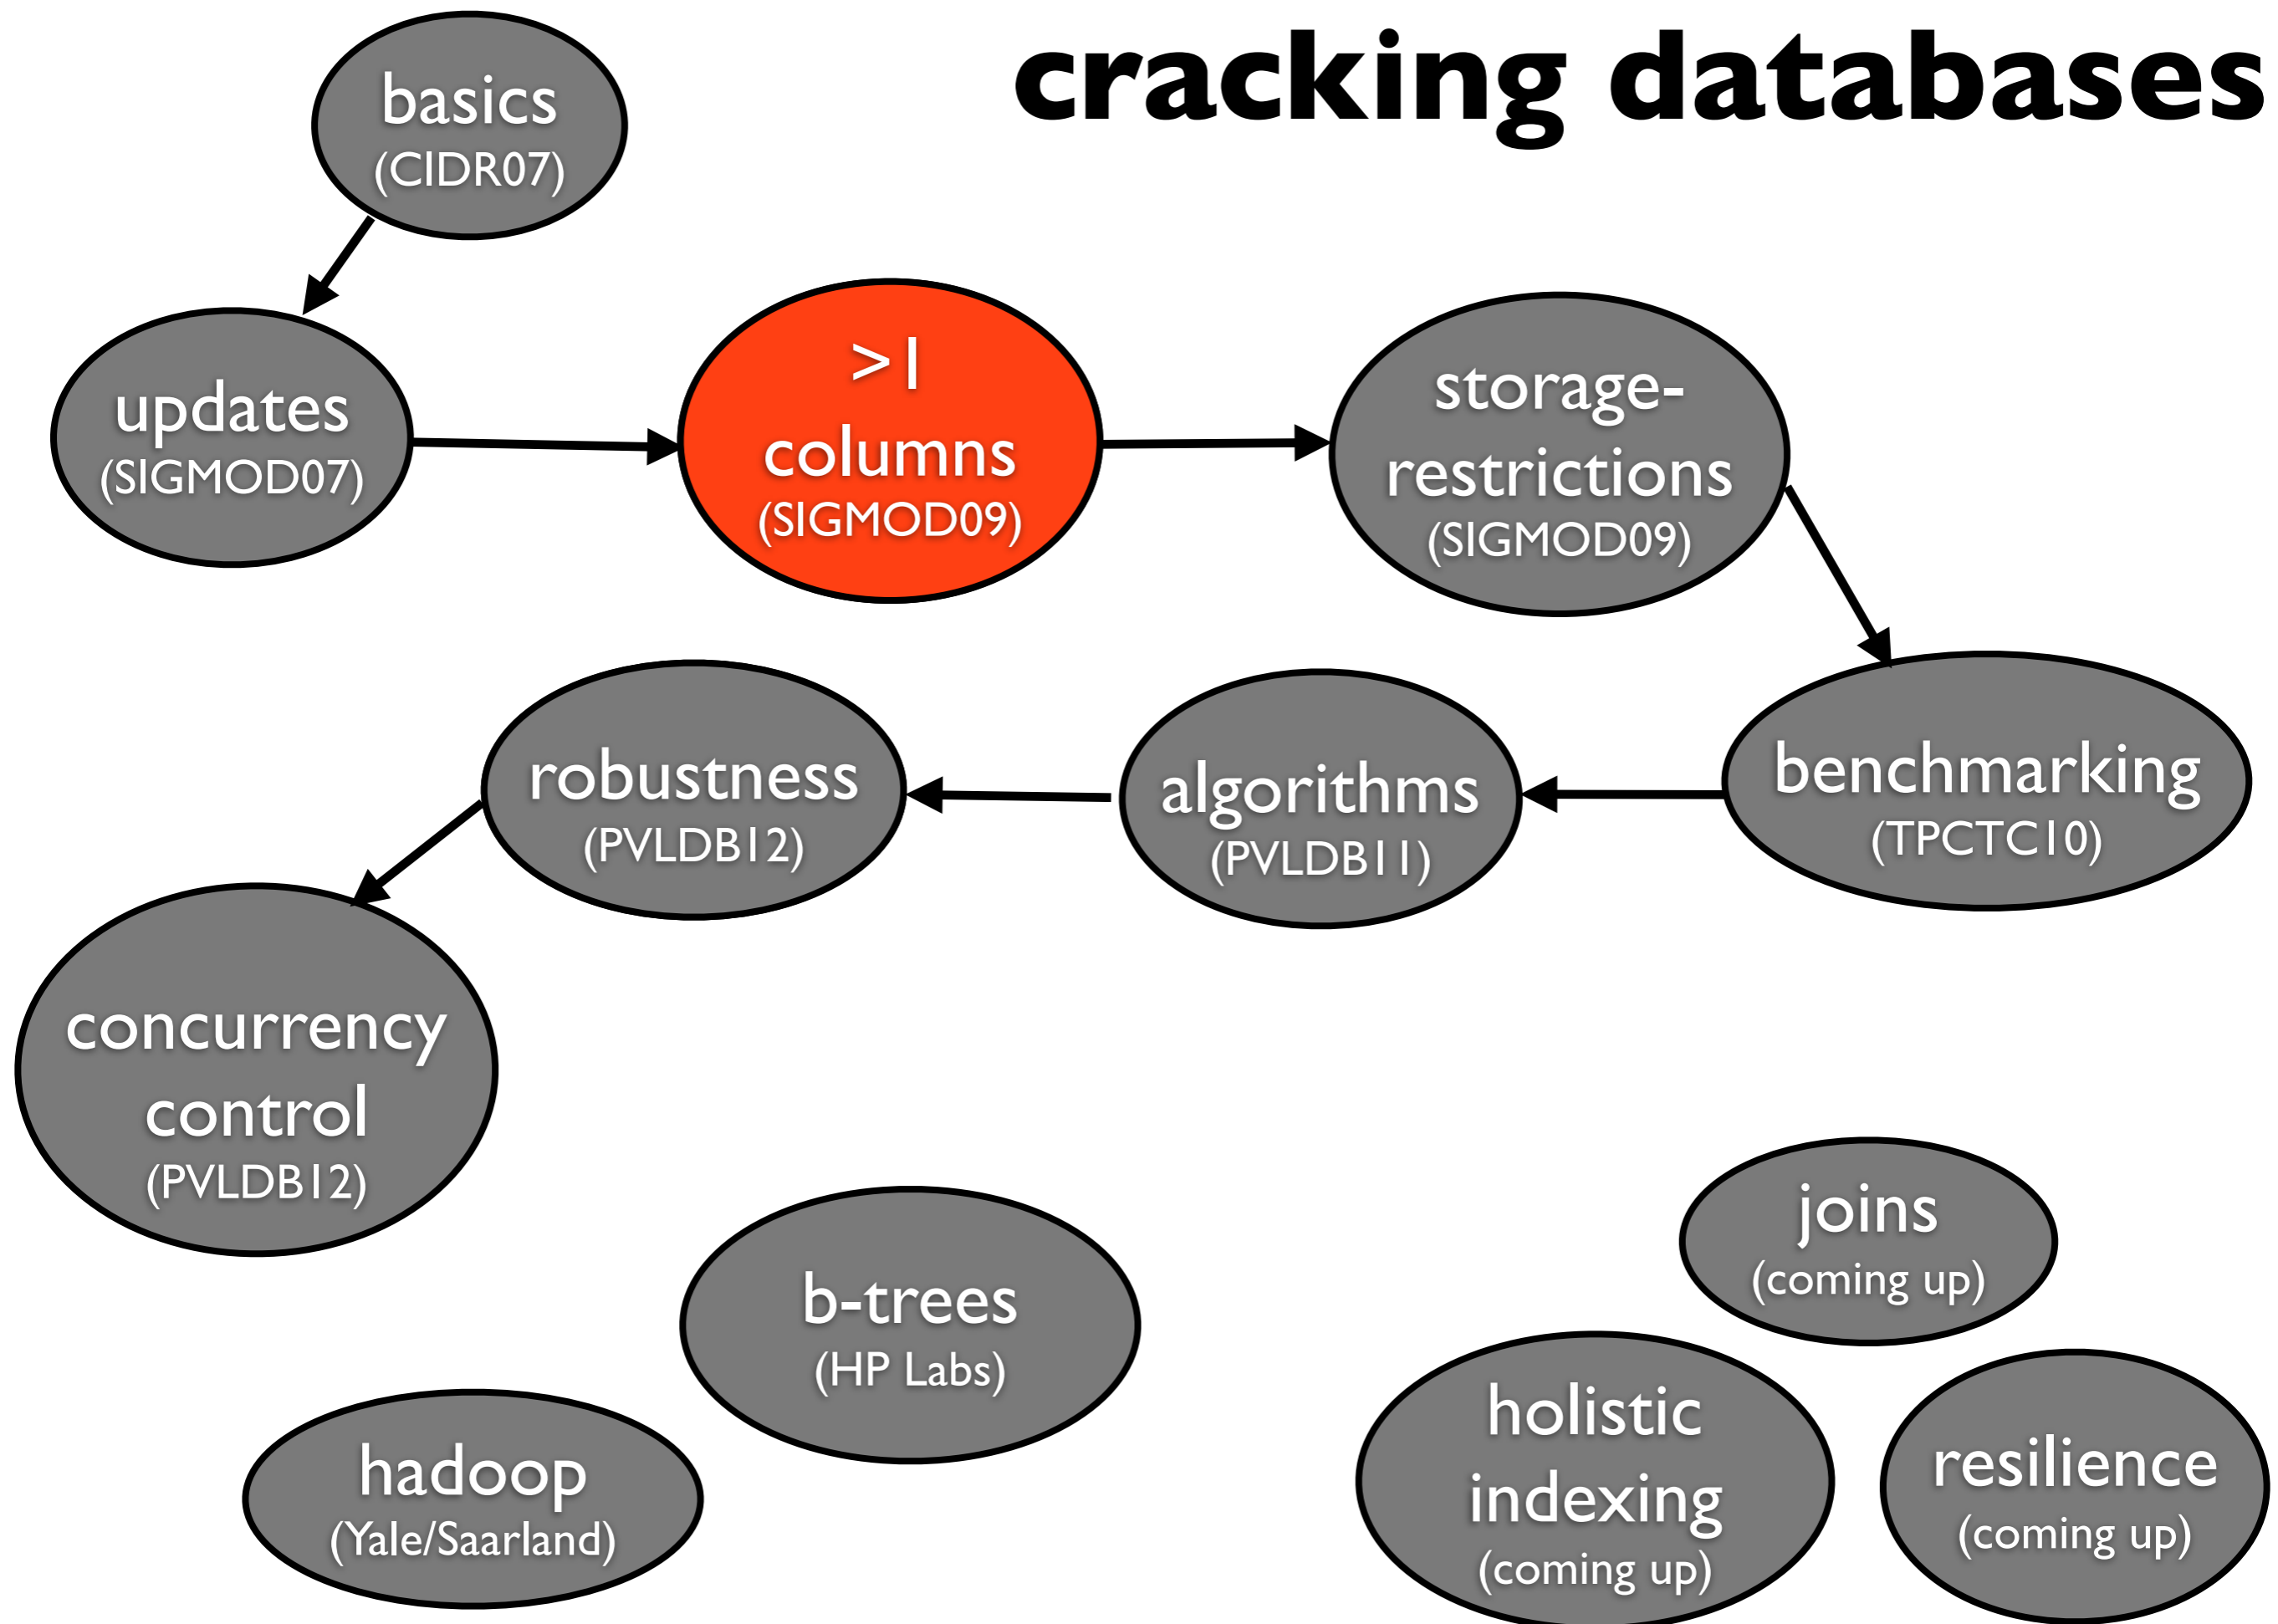
# cracking databases



# cracking databases




# cracking databases




# ***sideways cracking***

## ***tuple reconstruction***




**positional  
alignment**



**lookup**

$$A(i) = A + i * \text{width}(A)$$

tuple 1  
tuple 2  
tuple 3  
...



**A**

a1  
a2  
a3  
a4  
a5  
a6  
a7  
a8  
a9  
a10

**B**

b1  
b2  
b3  
b4  
b5  
b6  
b7  
b8  
b9  
b10


**positional  
alignment** →

**lookup**

$$A(i) = A + i * \text{width}(A)$$

**query**  
max(**B**) where **A** < 10

tuple 1 →  
tuple 2 →  
tuple 3 →  
...





**positional  
alignment** →

**lookup**

$$A(i) = A + i * \text{width}(A)$$

**query**  
max(**B**) where **A** < 10

tuple 1 →  
tuple 2 →  
tuple 3 →  
...





**positional alignment** →

**lookup**

$$A(i) = A + i * \text{width}(A)$$

**query**  
 $\max(\mathbf{B})$  where  $\mathbf{A} < 10$

tuple 1 →  
 tuple 2 →  
 tuple 3 →  
 ...




**positional alignment** →

**lookup**

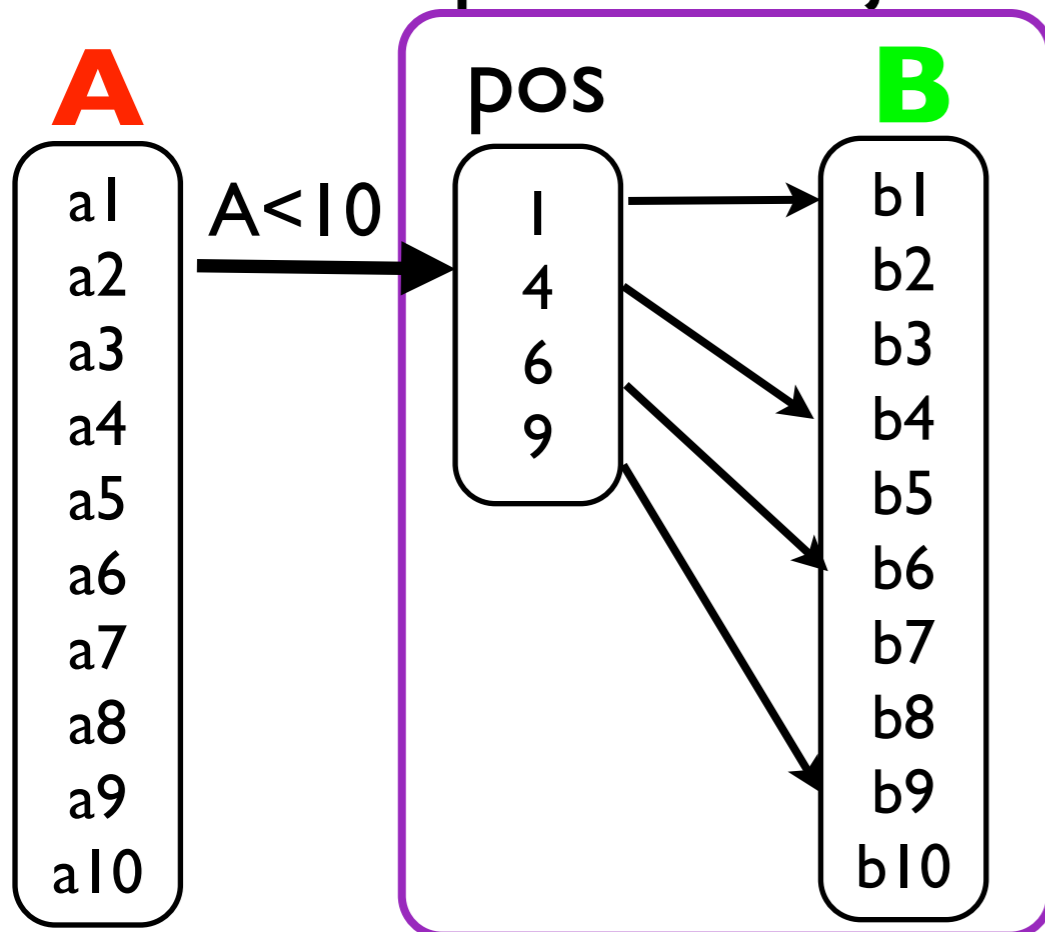
$$A(i) = A + i * \text{width}(A)$$

tuple 1 →  
 tuple 2 →  
 tuple 3 →  
 ...



**query**  
 max(**B**) where **A** < 10

**positional join**




**positional alignment** →

**lookup**


$$A(i) = A + i * \text{width}(A)$$

tuple 1 →  
 tuple 2 →  
 tuple 3 →  
 ...



**query**  
 max(**B**) where **A** < 10

**positional join**





**positional alignment** →

**lookup**


$$A(i) = A + i * \text{width}(A)$$

tuple 1 →  
 tuple 2 →  
 tuple 3 →  
 ...




**query**  
 max(**B**) where **A** < 10

**positional join**




**with cracking**




# sideways cracking

A

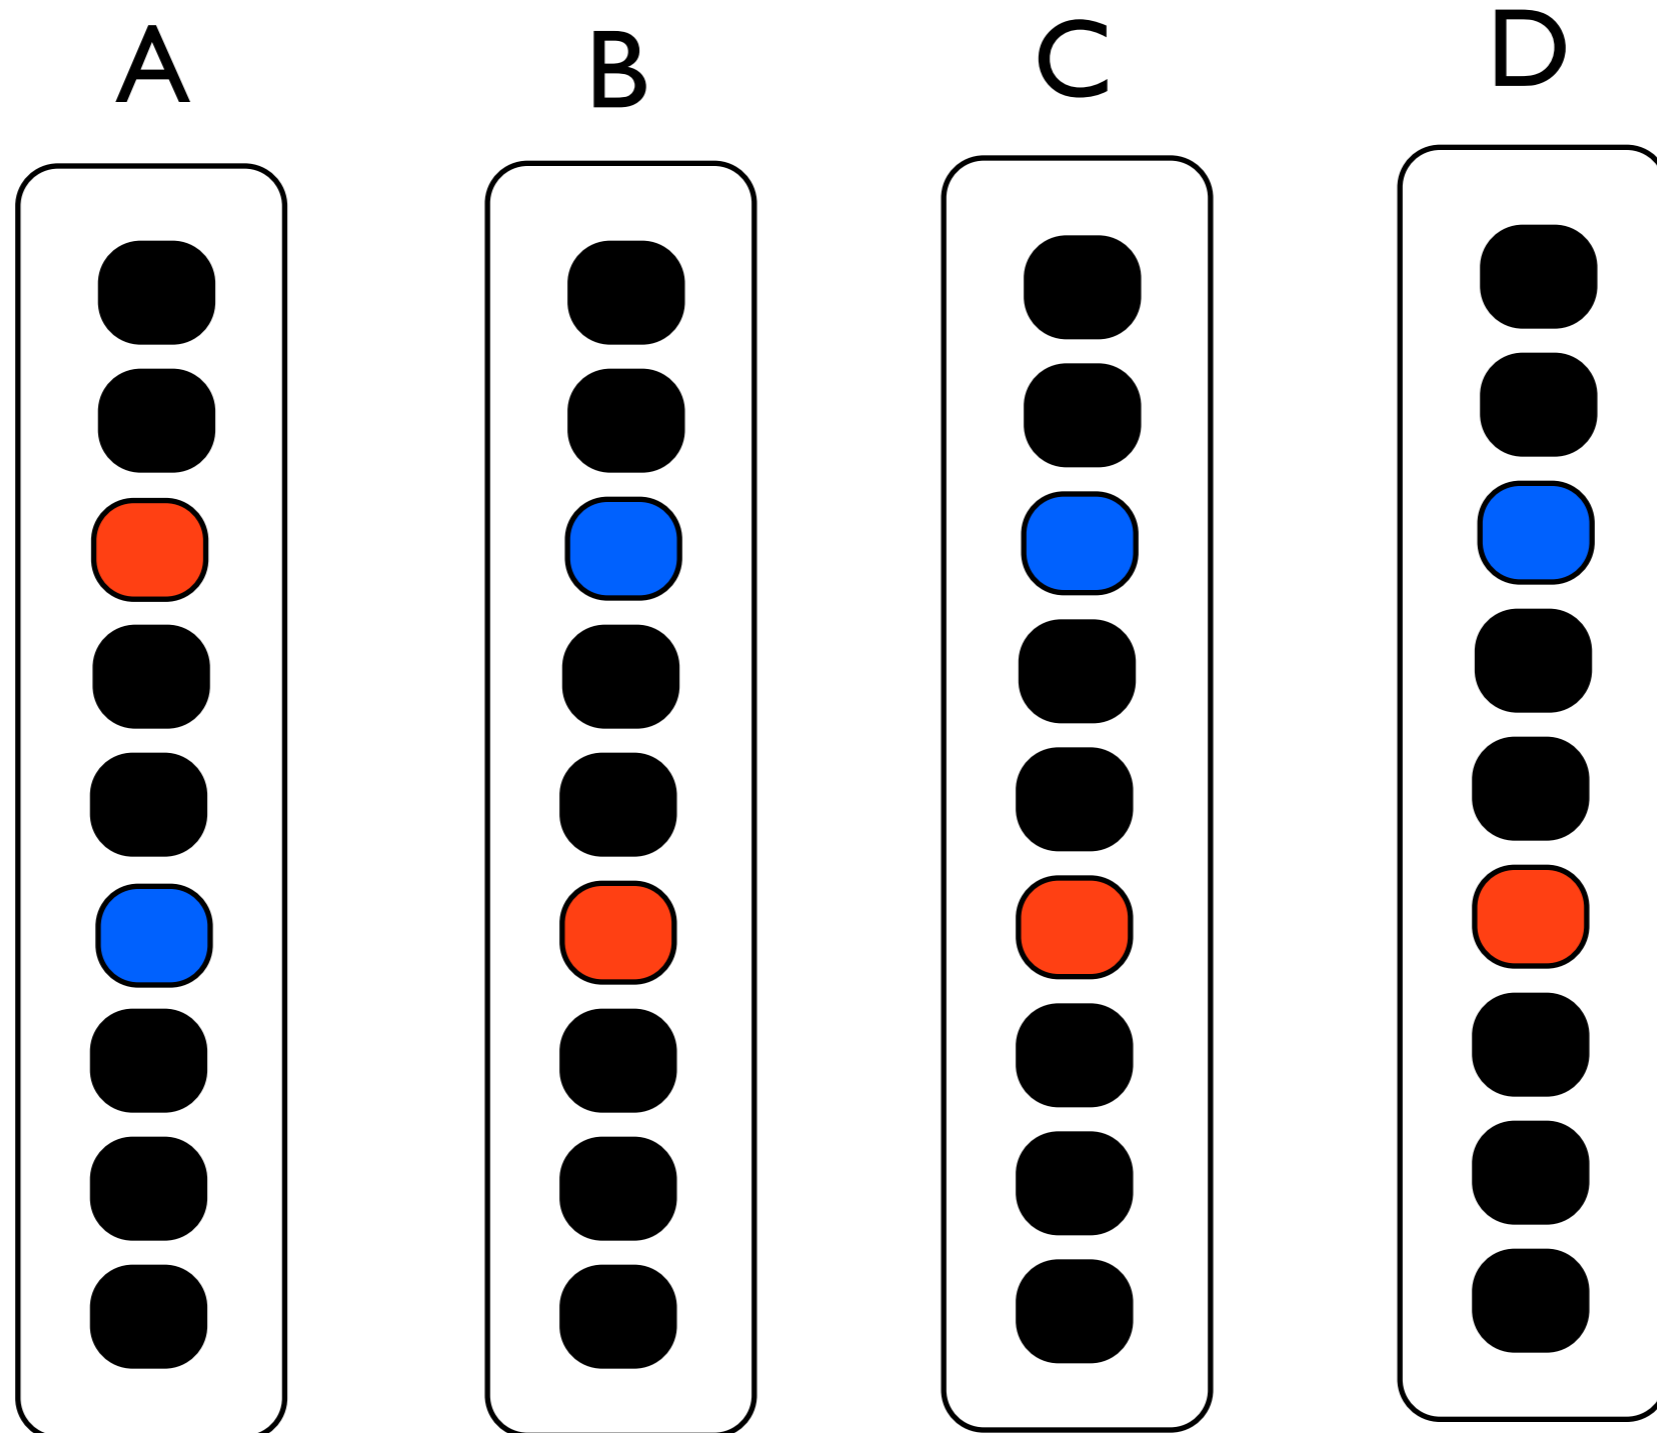


# sideways cracking

A




# sideways cracking




# sideways cracking

query




# sideways cracking

query




# sideways cracking

query




# sideways cracking

query




# sideways cracking


query




# sideways cracking




# sideways cracking



# sideways cracking




# sideways cracking


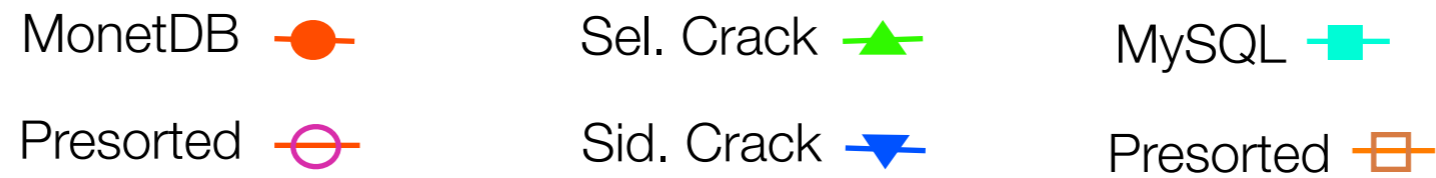


**replace tuple reconstruction with cracking**


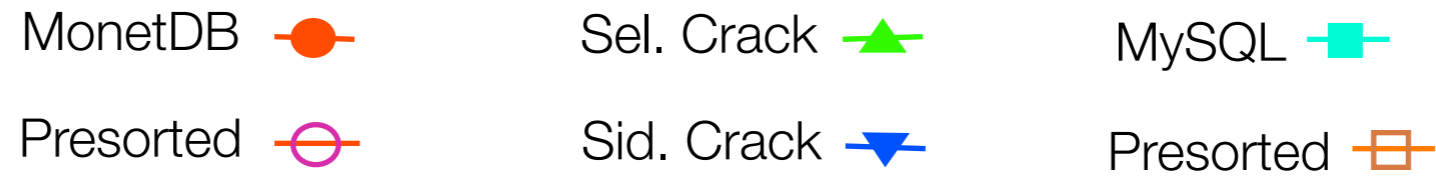
log crack actions and **replay** to align columns dynamically





# TPC-H



# TPC-H





# TPC-H





Fully tuned  
MonetDB  
Preparation cost  
~3 hours



# TPC-H



# TPC-H



# TPC-H





**Plain MonetDB**

**Selection cracking**

**MonetDB with sideways cracking**

**Fully tuned MonetDB**  
**Preparation cost ~3 hours**

# TPC-H



Plain MonetDB


Selection cracking

MonetDB with sideways cracking

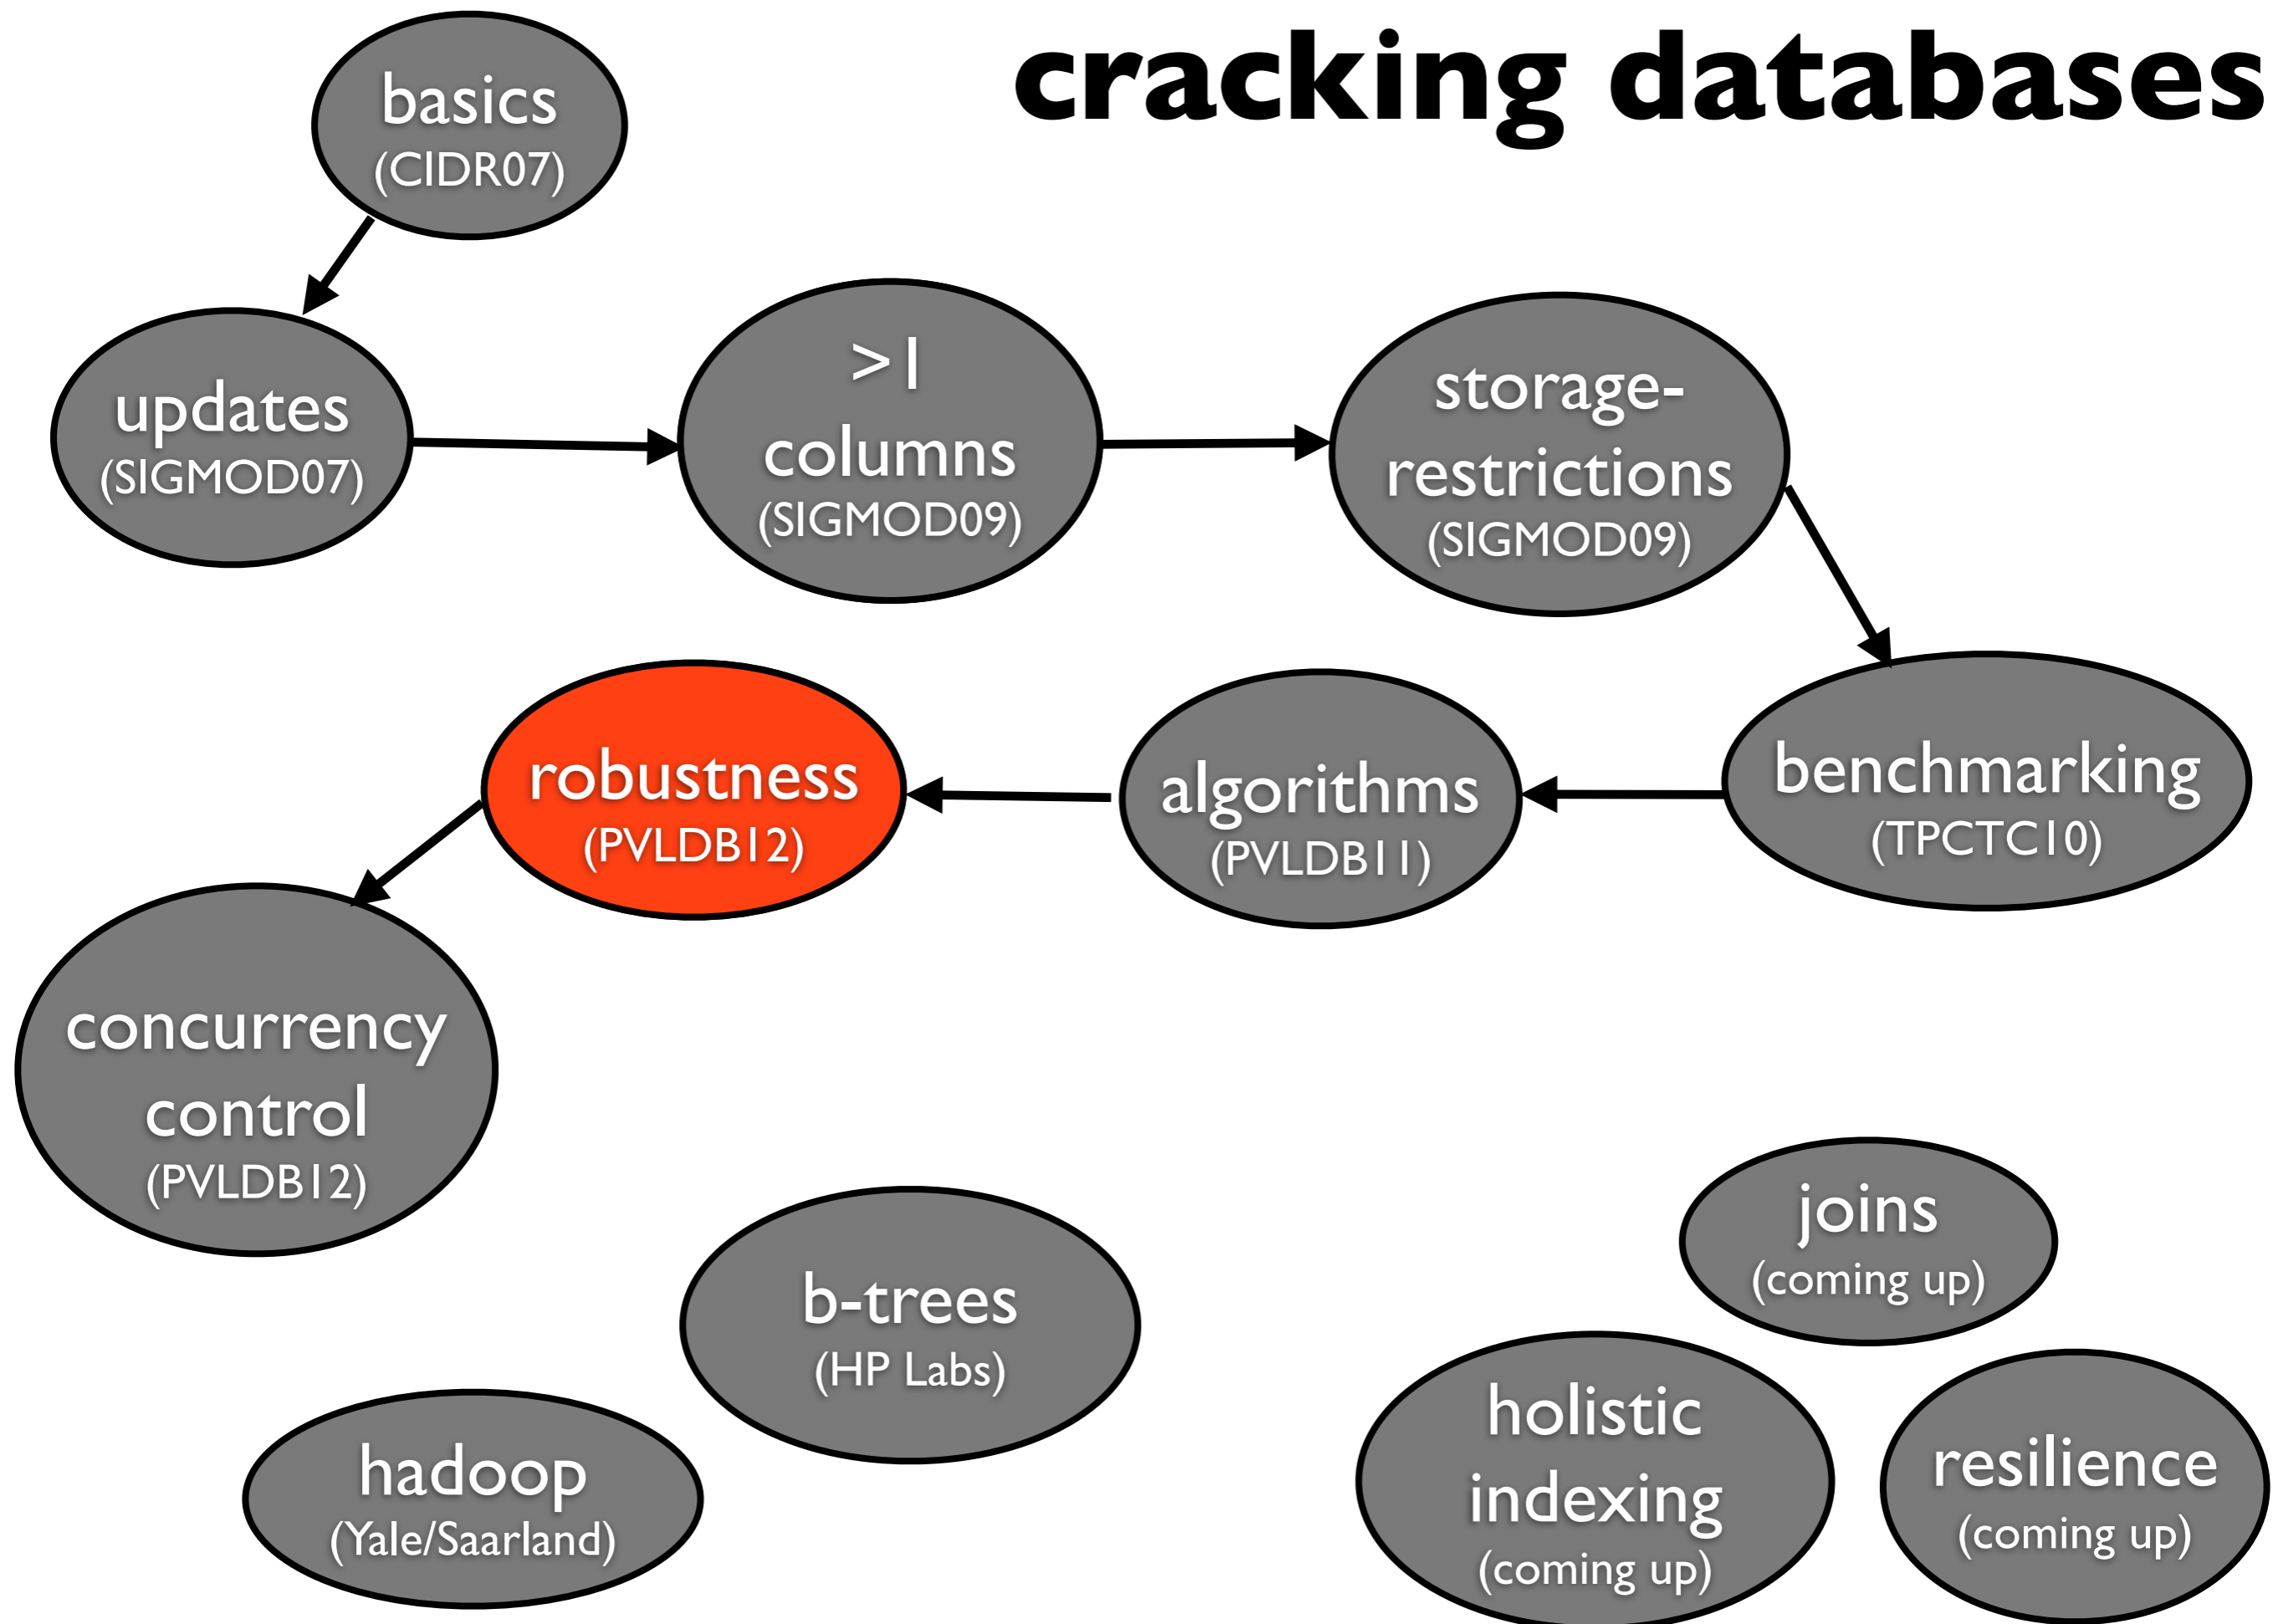
Fully tuned MonetDB  
Preparation cost ~3 hours

# TPC-H

**reducing data-to-query time**




# cracking databases



# ***stochastic cracking***

**robustness: maintain performance levels or have graceful degradation when input or status changes**

# adaptive indexing



# adaptive indexing




**Response time: X**

## adaptive indexing



**Response time: X**

# adaptive indexing



**Response time: X**



**Response time: 1000X**



column with 100 unique integers [1,100]

## **bad query pattern**



column with 100 unique integers [1,100]

# bad query pattern

<2

q1




column with 100 unique integers [1,100]

# bad query pattern


 $<2$ **q1****N**

column with 100 unique integers [1, 100]

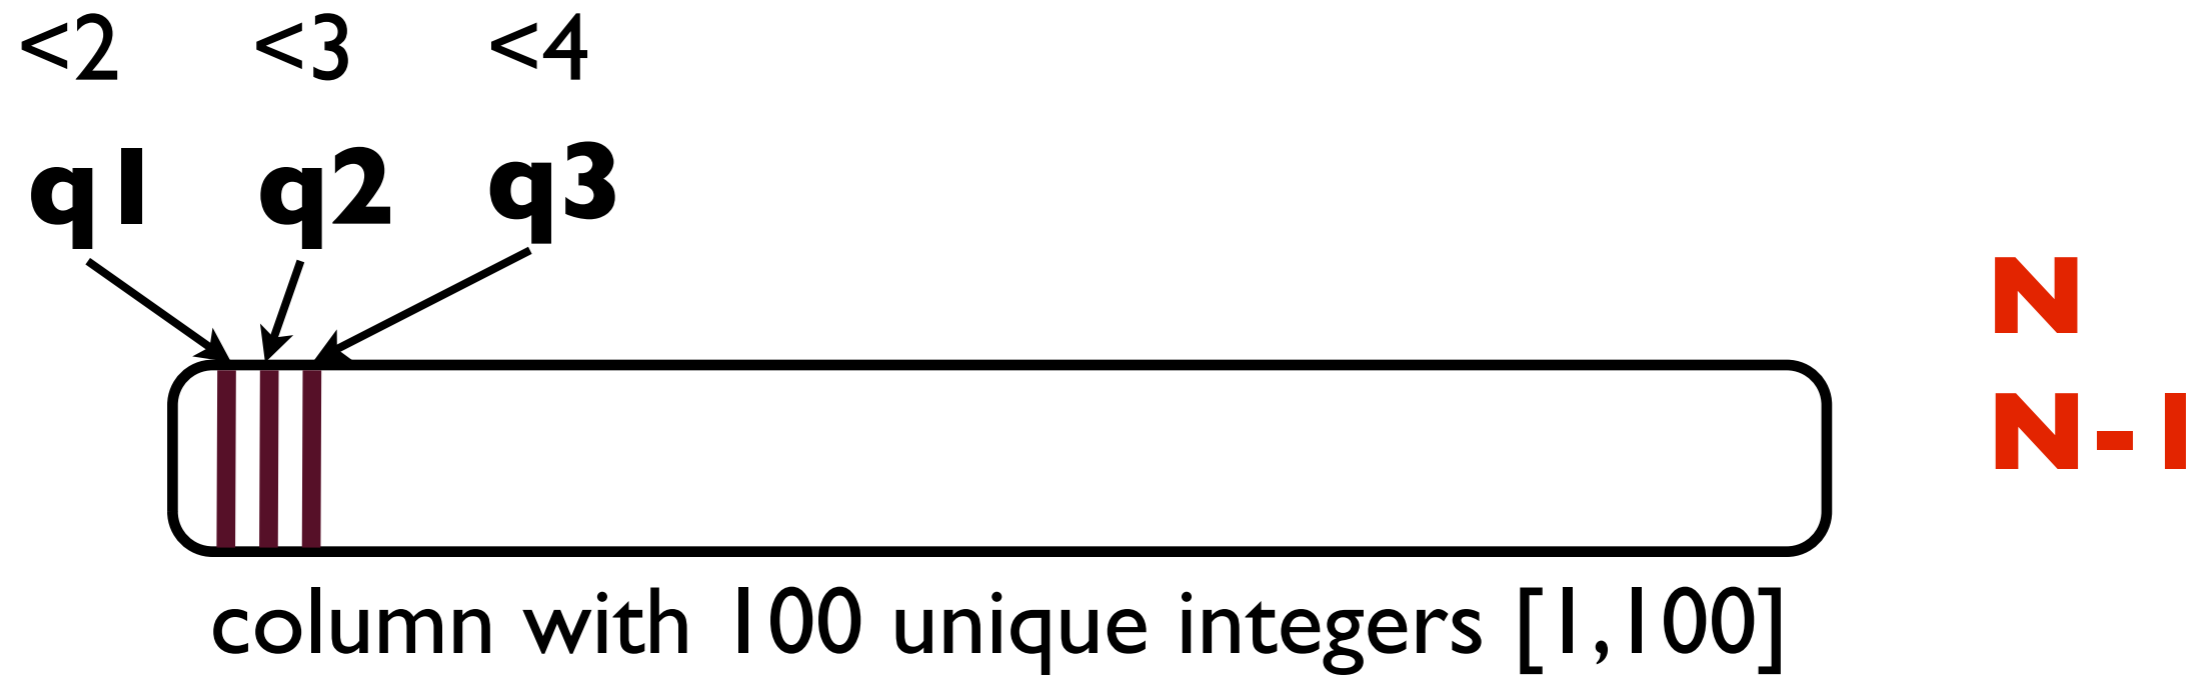
## bad query pattern




## bad query pattern




## bad query pattern



## bad query pattern




## bad query pattern




**blindly adapting to queries  
is not always a good idea**

**query driven + stochastic**


**query driven + stochastic**



# query driven + stochastic




# query driven + stochastic



**progressive cracking**

# query driven + stochastic




## progressive cracking

$q_l: <v_l$



# query driven + stochastic



## progressive cracking


$q_l: <v_l$



crack + filter  $<v_l$




# query driven + stochastic




## progressive cracking

$q_l: <v_l$

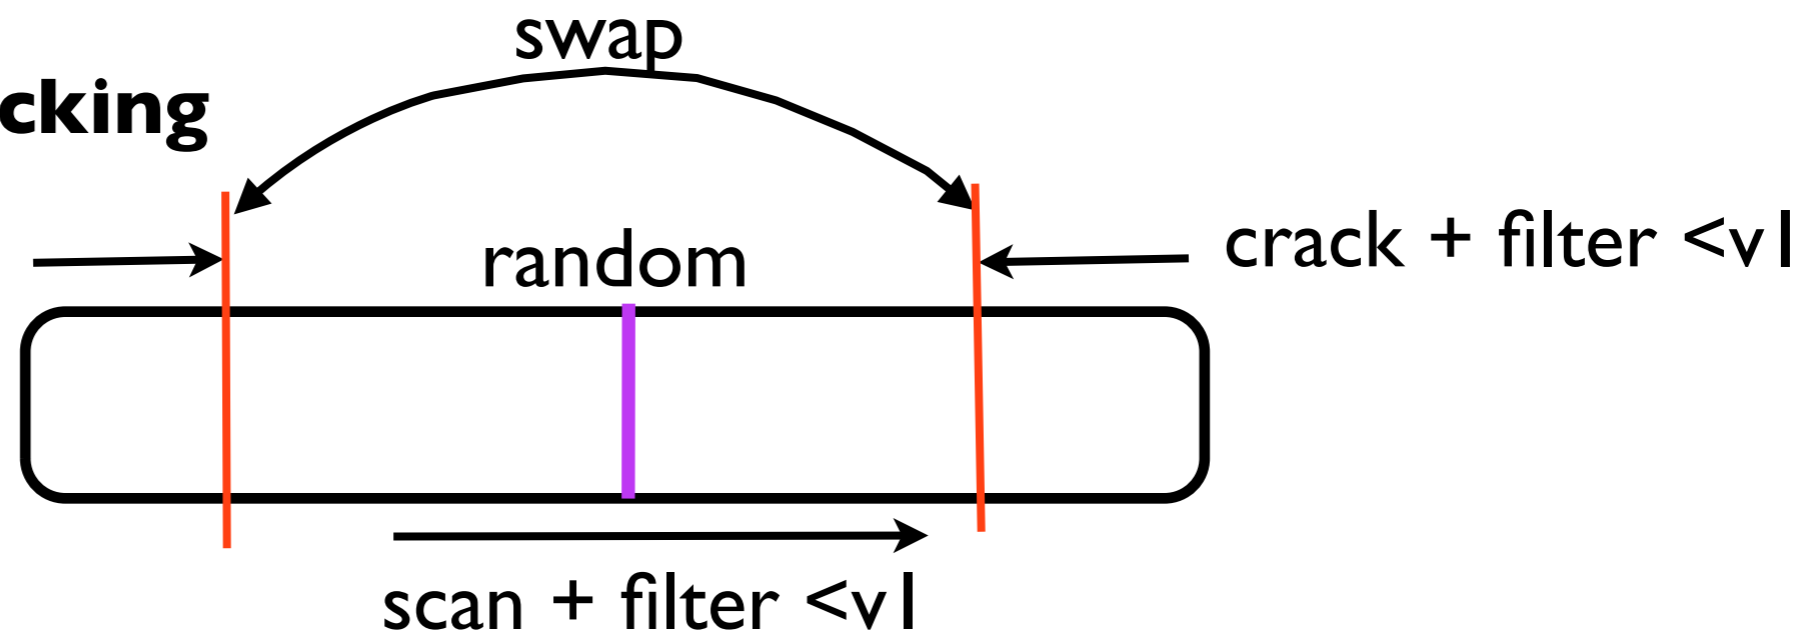


# query driven + stochastic




## progressive cracking

$q_l: <v_l$




# query driven + stochastic




## progressive cracking

q1: <v1  
q2: <v2




# query driven + stochastic




## progressive cracking

q1: <v1  
q2: <v2

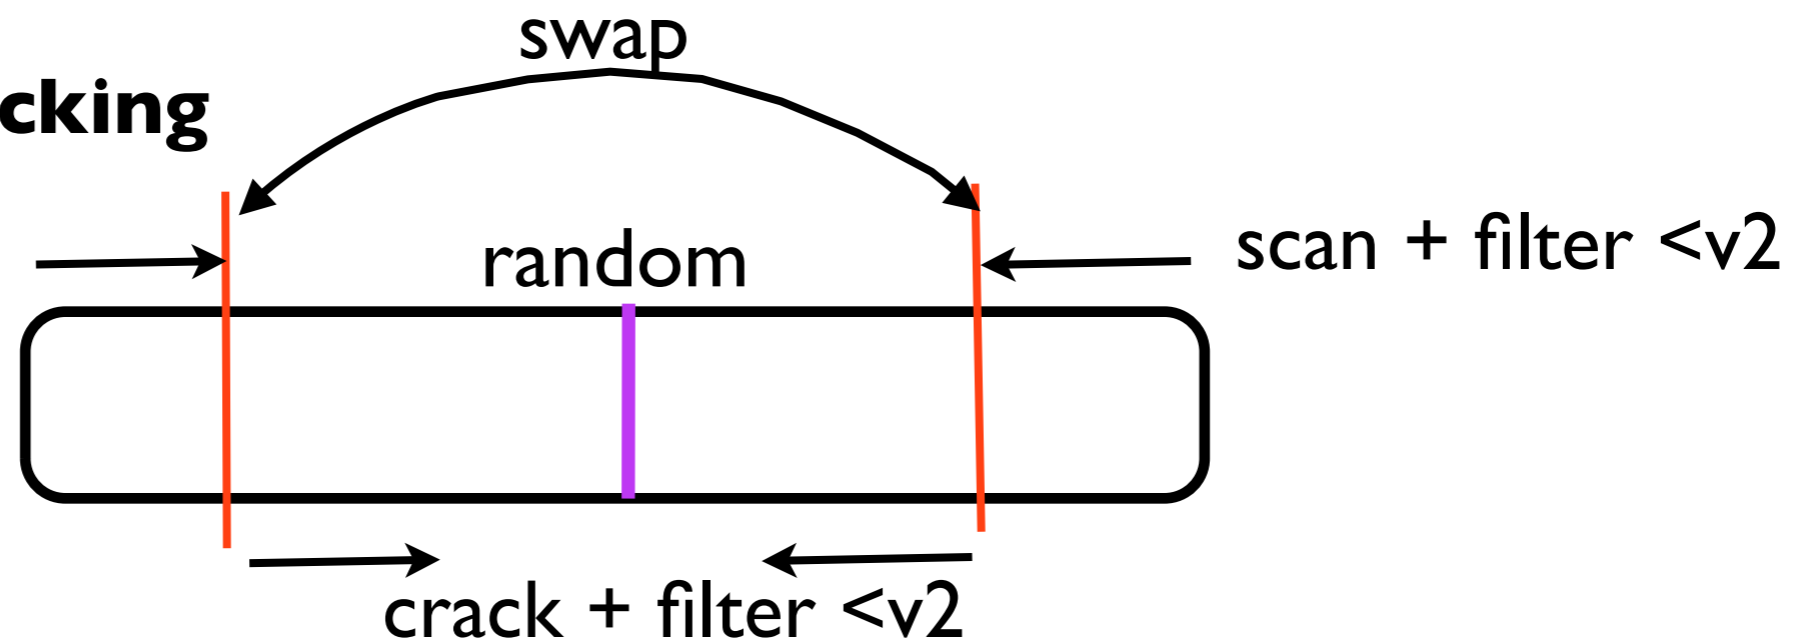


# query driven + stochastic




## progressive cracking

q1: <v1  
q2: <v2




# query driven + stochastic



## progressive cracking

q1: <v1  
q2: <v2



# **cracking on Skyserver (4TB)**

(Sloan Digital Sky Survey, [www.sdss.org](http://www.sdss.org))

**cracking answers 160.000 queries while full indexing is still half way creating one index**

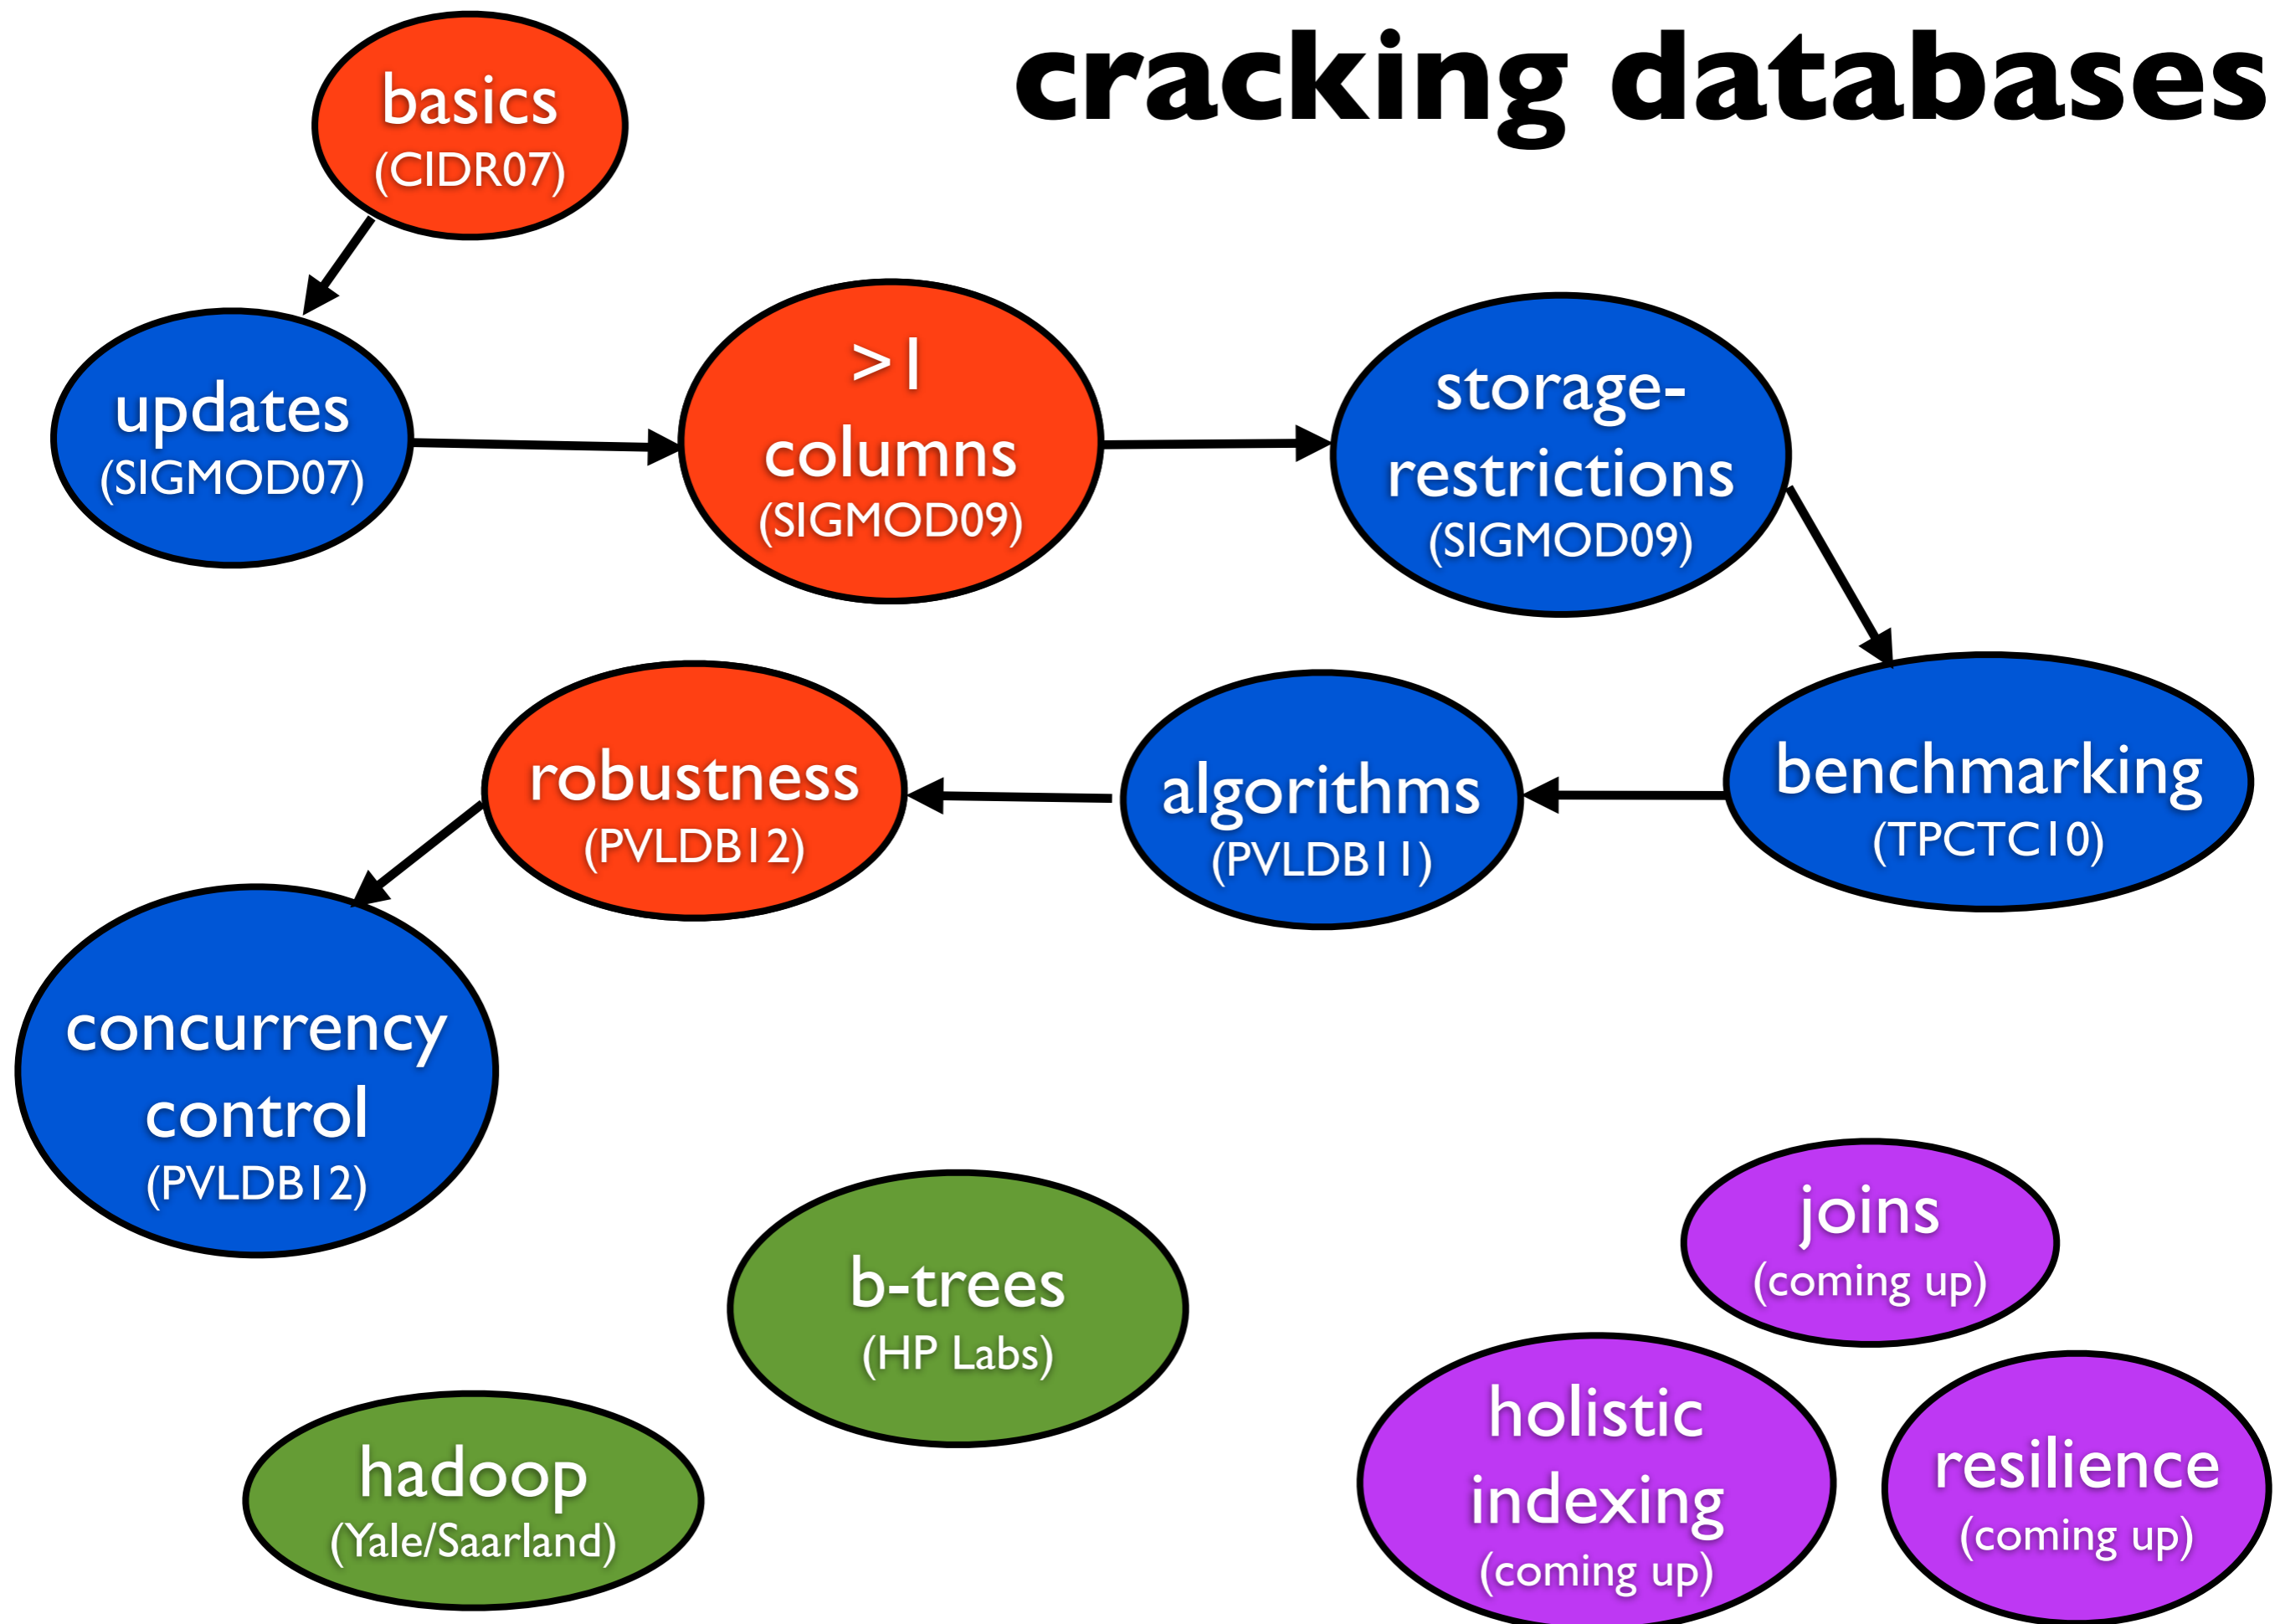
**reducing data-to-query time**

**cracking on Skyserver (4TB)**

(Sloan Digital Sky Survey, [www.sdss.org](http://www.sdss.org))

**cracking answers 160.000 queries while full indexing is still half way creating one index**

# cracking databases



# loading



# loading



***copy data inside the database***  
***database now has full control***

# loading




*copy data inside the database*  
*database now has full control*

**slow process...not all data might be  
needed all the time**


# database vs. unix tools

1 file, 4 attributes,  
1 billion tuples




# database vs. unix tools

1 file, 4 attributes,  
1 billion tuples




## break down db cost

- Loading
- Query Processing




# database vs. unix tools

1 file, 4 attributes,  
1 billion tuples



## break down db cost


- Loading
- Query Processing



**loading is a major bottleneck**


# database vs. unix tools

1 file, 4 attributes,  
1 billion tuples



## break down db cost

- Loading
- Query Processing



**loading is a major bottleneck**

... but writing/maintaining scripts is hard too

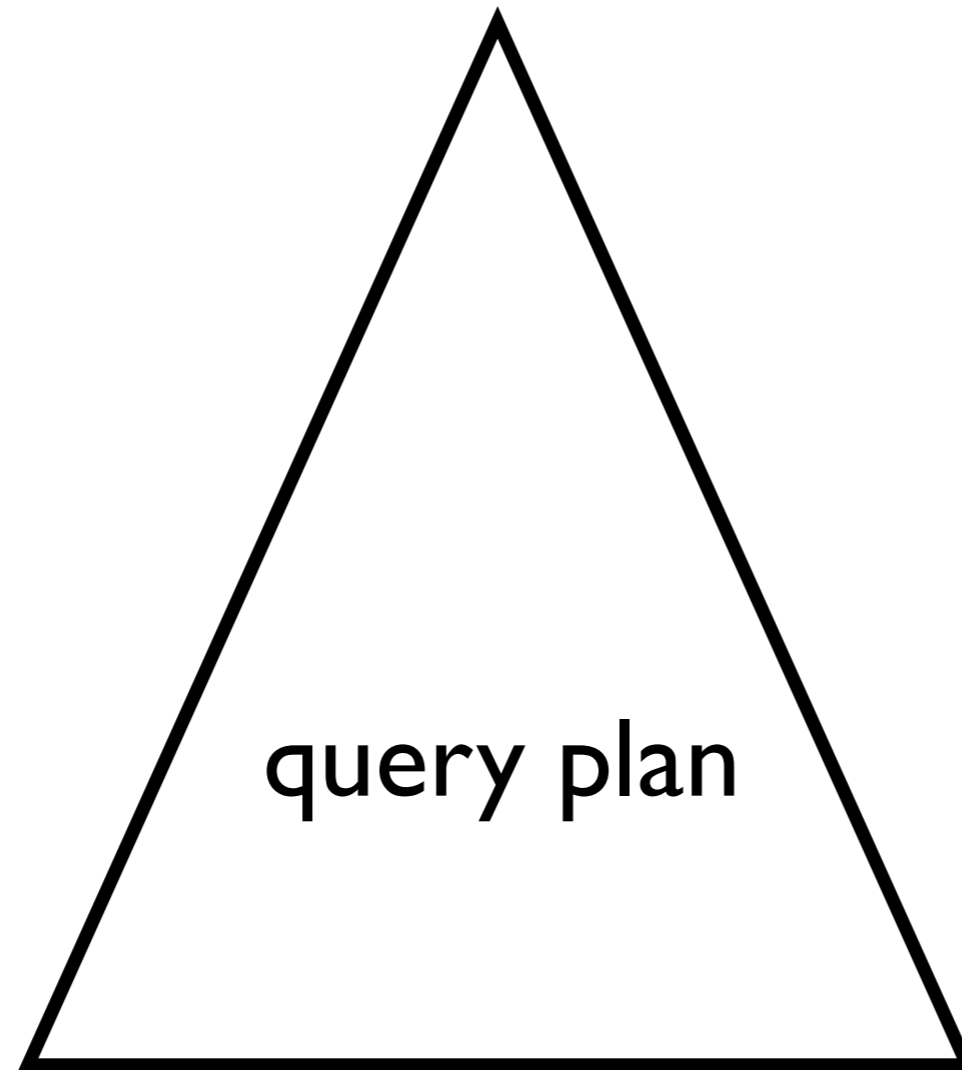
# **adaptive loading**


**load/touch only what is needed  
and only when it is needed**


**but raw data access is expensive**


tokenizing - parsing - no indexing - no statistics


**challenge: fast raw data access**



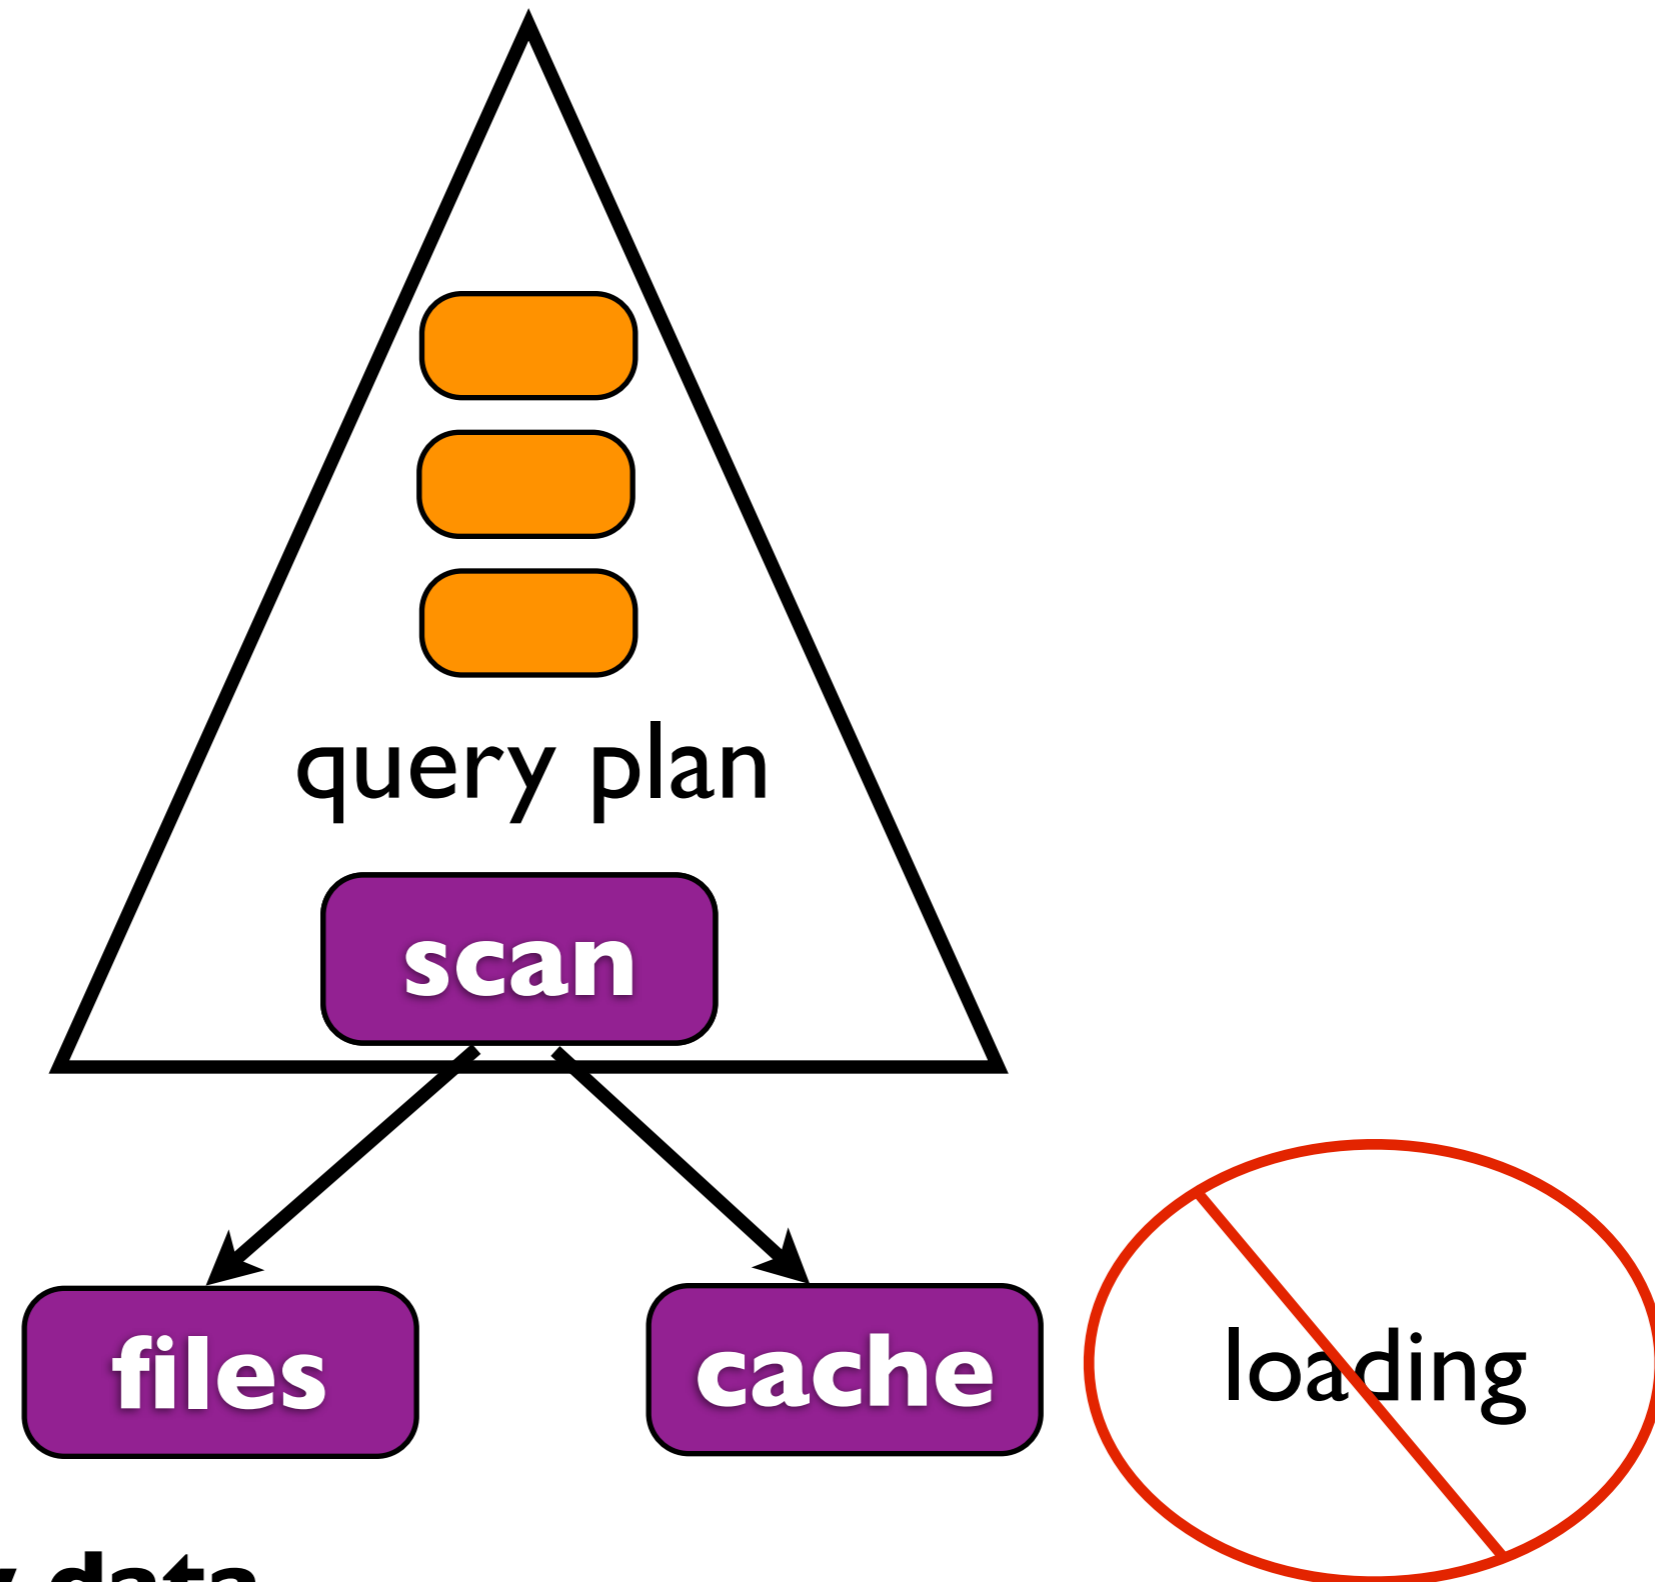




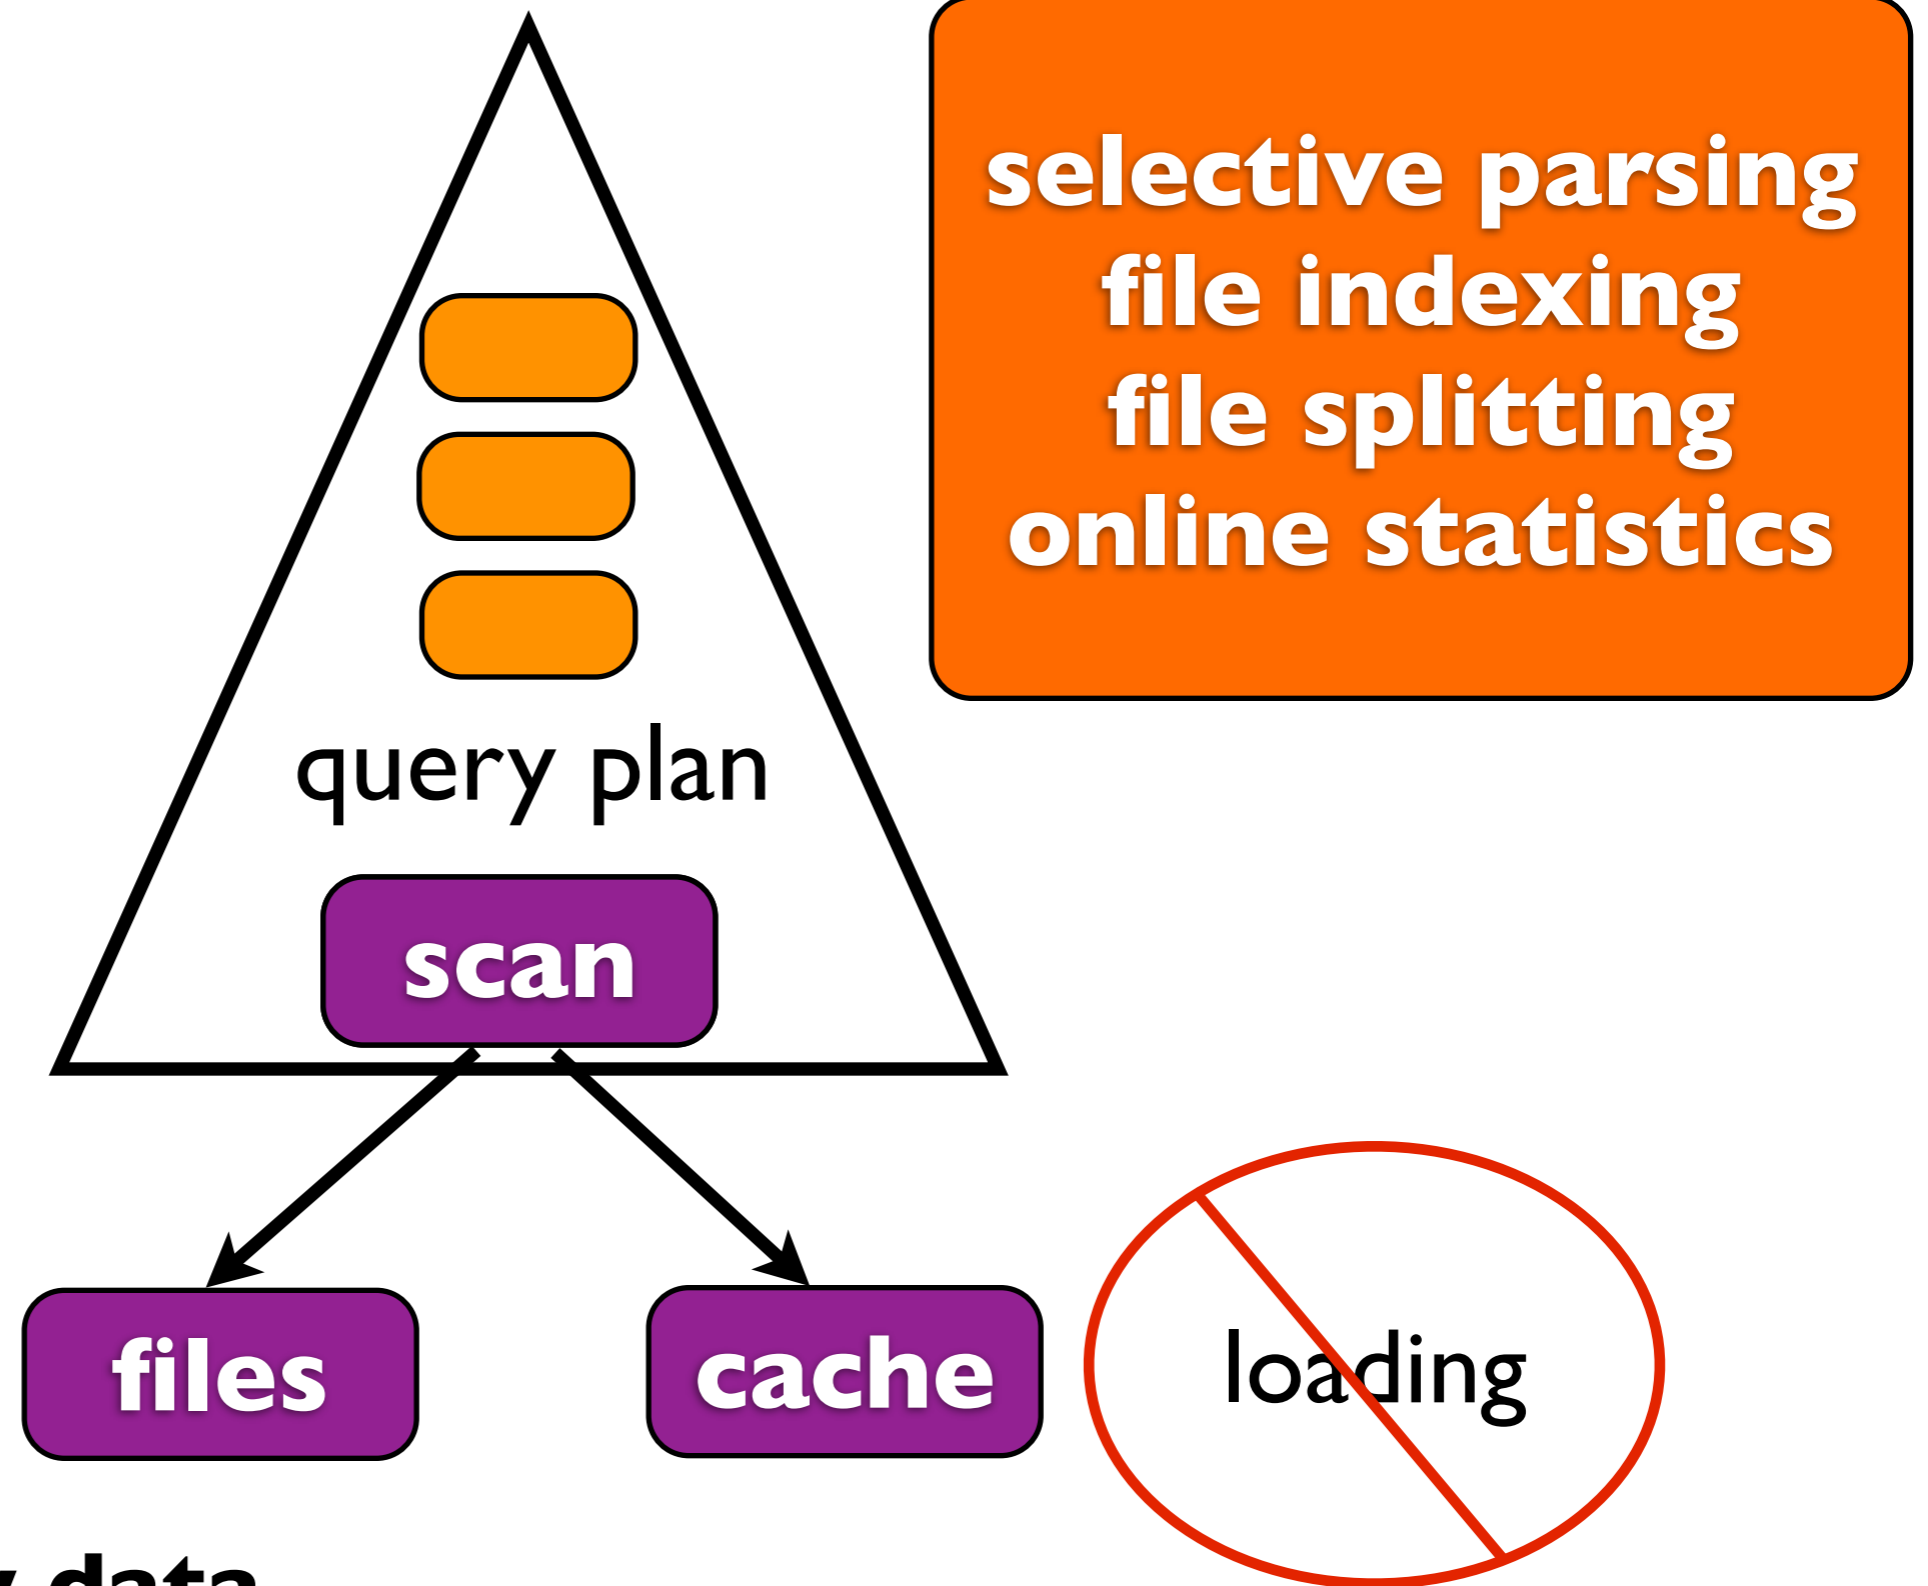




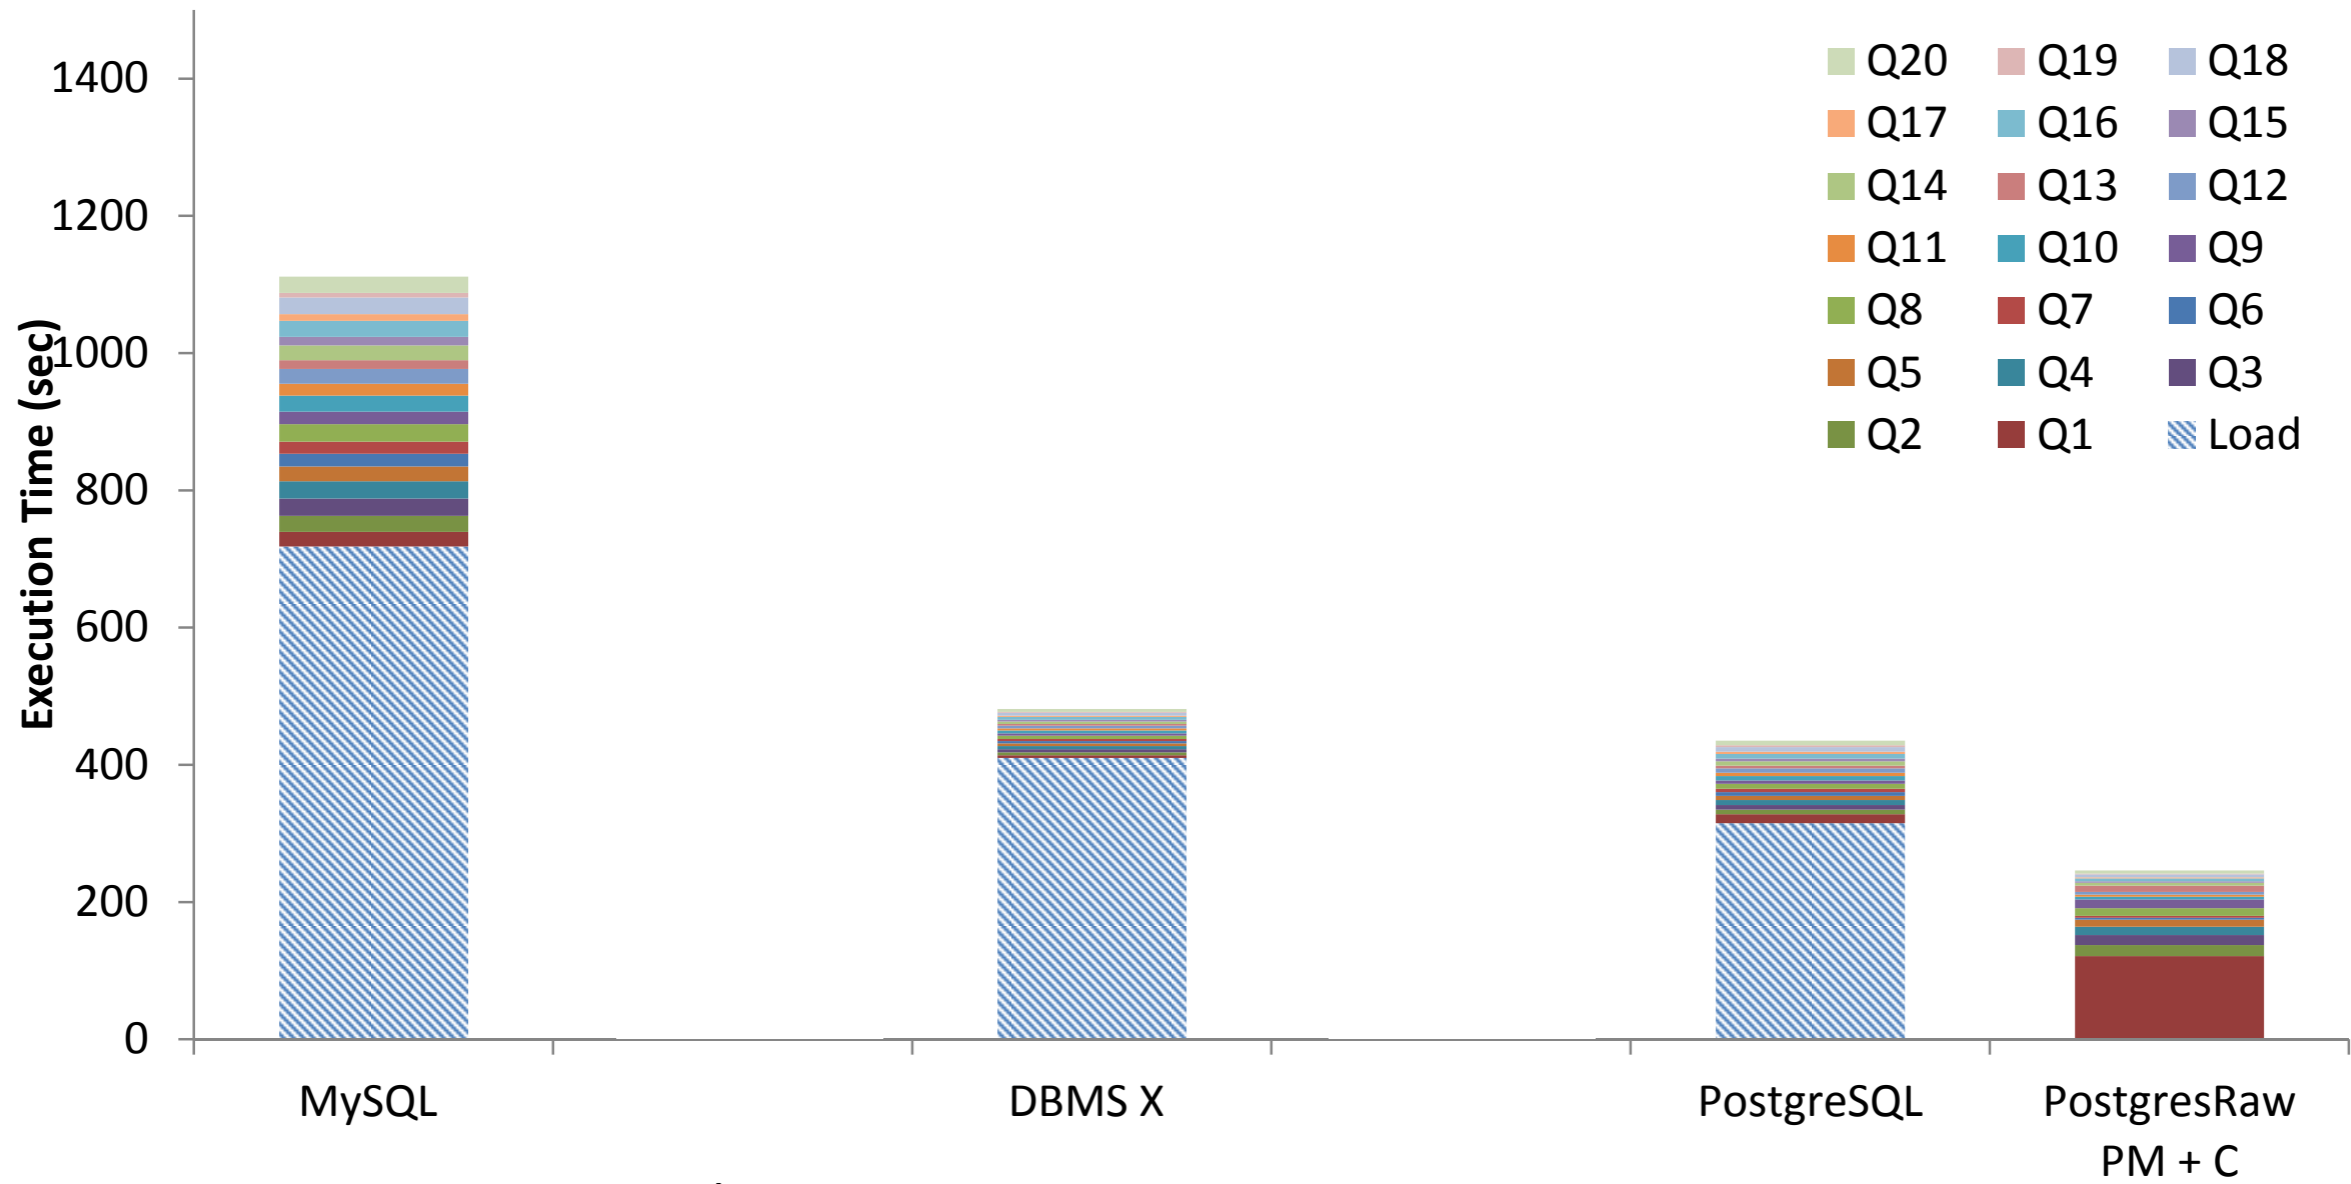
**access raw data  
adaptively on-the-fly**





**access raw data  
adaptively on-the-fly**





**access raw data  
adaptively on-the-fly**












# reducing data-to-query time



# towards auto-tuning data kernels

**load**

**tune**

**query**

# towards auto-tuning data kernels

**load**

**tune**

**query**

## so what's next?

adaptive (**load-store-execute**)

**cracking(+ AI, + OS, +ML)**

compression

*disk based cracking*

*multidimensional cracking*

multi-core cracking

row-store cracking

aggregations

***...and many more...***

# ***interactive data systems***

# querying



# querying

load

tune

query

*SQL interface*



# querying

load

tune

query

***SQL interface***

***correct and complete answers***




# querying

**complex and slow - not fit for exploration**

***SQL interface***

***correct and complete answers***






***just touch the data you need***




***just touch the data you need***

**this is not about query building  
it is about query processing**

# dbTouch





# dbTouch






# dbTouch

**column 1**

64000






**data**





# dbTouch

**column 1**


**touch/  
query**



**data**



# dbTouch




**table**



59428, 40571, 40572, 9429

**touch/  
query**

# ***HCI + databases***




+



**what does this mean for db kernels?**


**db**



select R.a from R

**what does this mean for db kernels?**

**db**



select R.a from R

**what does this mean for db kernels?**


**dbTouch**

**56 38 45 2**



**process only  
what you touch**

## **explore: touch, observe and react**



**the system does not have control of the data flow  
the user dictates which is the next tuple**

**hierarchies of samples  
incremental and adaptive operators  
adaptive indexing - adaptive loading**



## an **exploration** tool

get a quick feeling about your data  
focus on interesting areas



***HCI + databases***

**come play with the dbTouch demo on iPad!**

**adaptive  
indexing**

**dbTouch**

**adaptive  
loading**

## ***3 Ideas for Big Data Exploration***

**adaptive systems - tailored for exploration**

*it is not the strongest species that survive, nor the most intelligent, but the ones most responsive to change*  
[Darwin, Megginson]

***Stratos Idreos***  
CWI, Amsterdam

**adaptive  
indexing**

**dbTouch**

**adaptive  
loading**

## ***3 Ideas for Big Data Exploration***

**adaptive systems - tailored for exploration**

*it is not the strongest species that survive, nor the most intelligent, but the ones most responsive to change*  
[Darwin, Megginson]

***Stratos Idreos***  
CWI, Amsterdam

**online  
analytics**

**query  
morphing**

**linked  
views**

**linked  
history**

**1 minute  
db kernels**

**queries as  
answers**

**adaptive  
indexing**

**dbTouch**

**adaptive  
loading**

## ***3 Ideas for Big Data Exploration***

**adaptive systems - tailored for exploration**

*it is not the strongest species that survive, nor the most intelligent, but the ones most responsive to change*  
[Darwin, Megginson]

**Thank you!**

**Stratos Idreos**  
CWI, Amsterdam

**online  
analytics**

**query  
morphing**

**linked  
views**

**linked  
history**

**1 minute  
db kernels**

**queries as  
answers**