

DB Performance with Graphics

With Kyle Hailey

Product Manager for DB Optimizer

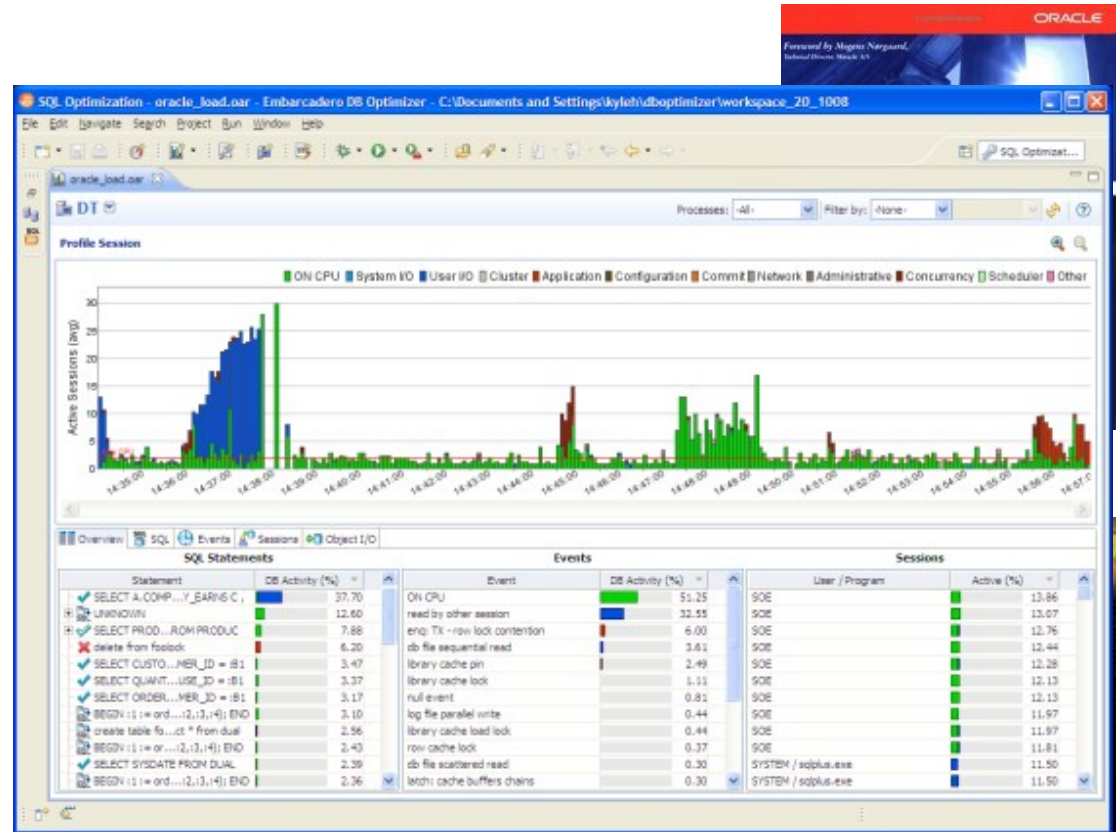


OracleMonitor.com



Who is Kyle Hailey

- 1990 Oracle
 - 90 support
 - 92 Ported v6
 - 93 France
 - 95 Benchmarking
 - 98 ST Real World Performance
- 2000 Dot.Com
- 2001 Quest
- 2002 Oracle OEM 10g
- 2006 Independent
- 2008 Embarcadero
 - DB Optimizer



Foreword by Magnus Nergaard, Technical Director, Oracle
Kyle Hailey, Arjo Kolk, Thomas Kyte, Jonathan Lewis, Connor McDonald, Cary V. Millsap, James Morie, David Ruthven, Gaja Krishna Vaidyanatha
Apress

Launch: Pressure

Midnight before
January 28, 1986
Lives are on the line



Thanks to Edward Tufte

13 Pages Faxed

HISTORY OF O-RING DAMAGE ON SRM FIELD JOINTS

CONCLUSIONS :

- ° TEMPERATURE OF O-RING IS NOT ONLY PARAMETER CONTROLLING BLOW-BY

SRM 15 WITH BLOW-BY HAD AN O-RING TEMP AT 53°F
SRM 22 WITH BLOW-BY HAD AN O-RING TEMP AT 75°F
FOUR DEVELOPMENT MOTORS WITH NO BLOW-BY WERE TESTED AT O-RING TEMP OF 47° TO 52°F

DEVELOPMENT MOTORS HAD PUTTY PACKING WHICH RESULTED IN BETTER PERFORMANCE
- ° AT ABOUT 50°F BLOW-BY COULD BE EXPERIENCED IN CASE JOINTS
- ° TEMP FOR SRM 25 ON 1-28-86 LAUNCH WILL BE 29°F 9AM
38°F 2PM
- ° HAVE NO DATA THAT WOULD INDICATE SRM 25 IS DIFFERENT THAN SRM 15 OTHER THAN TEMP


RECOMMENDATIONS :

- ° O-RING TEMP MUST BE $\geq 53^{\circ}\text{F}$ AT LAUNCH

DEVELOPMENT MOTORS AT 47° TO 52°F WITH PUTTY PACKING HAD NO BLOW-BY
SRM 15 (THE BEST SIMULATION) WORKED AT 53°F
- ° PROJECT AMBIENT CONDITIONS (TEMP & WIND) TO DETERMINE LAUNCH TIME

SRM-22 FORWARD FIELD JOINT HAD PUTTY PATH TO PRIMARY O-RING, BUT NO O-RING EROSION AND NO SOOT BLOWBY. OTHER SRM-22 FIELD JOINTS HAD NO BLOWHOLES IN PUTTY.

Original Engineering data




“damages at the hottest and coldest temperature”
- *management*

only showed damage


Congressional Hearings Evidence

History of O-Ring Damage in Field Joints (Cont)



INFORMATION ON THIS PAGE WAS PREPARED TO SUPPORT AN ORAL PRESENTATION AND CANNOT BE CONSIDERED COMPLETE WITHOUT THE ORAL DISCUSSION

Clearer



1. Include successes
2. Mark Differences
3. Normalize same temp
4. Scale known vs unknown

Difficult

- NASA Engineers Fail
- Congressional Investigators Fail
- Data Visualization is Difficult

But ...

Lack of Clarity can be devastating


Solutions

- **Clear Identification**
 - Know how to identify problems and issues
- **Access to details**
 - Provide solutions and/or information to address the issues
- **Graphics**
 - Easy understanding, effective communication and discussion

First Step: Graphics

“The humans ... are exceptionally good at parsing *visual* information, especially when that information is coded by **color** and/or motion _____.”

Knowledge representation in cognitive science. Westbury, C. & Wilensky, U. (1998)



Why Use Graphics

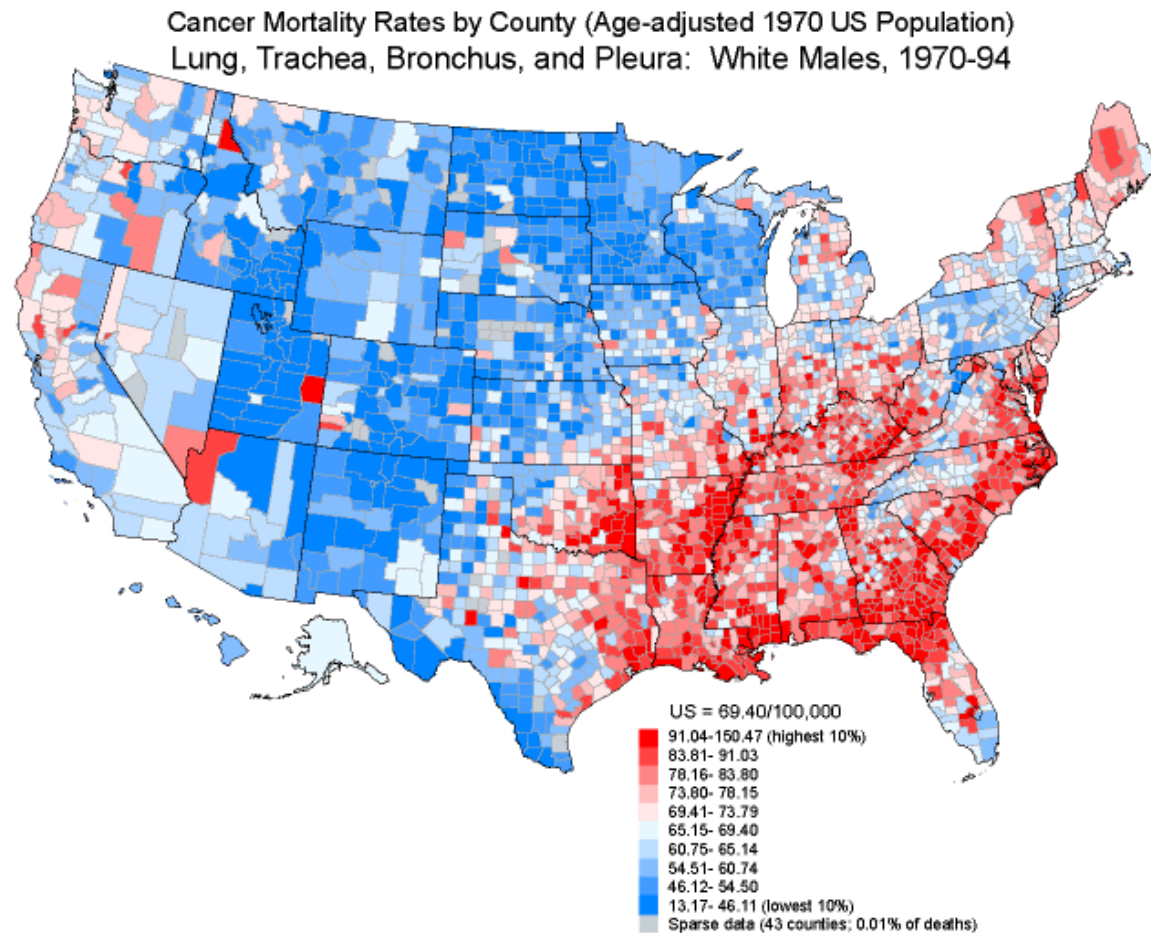
You can't imagine how many times I was told that nobody wanted or would use graphics ...

-- Jef Raskin, the creator of the Macintosh

Infocus – (overhead projectors) cited a study that humans can parse graphical information 400,000 times faster than textual data

Counties in US

- 3101 Counties in US
- 50 pages







“If I can't picture it, I can't understand it”

- Albert Einstein

Anscombe's Quartet

I		II		III		IV	
x	y	x	y	x	y	x	y
10	8.04	10	9.14	10	7.46	8	6.58
8	6.95	8	8.14	8	6.77	8	5.76
13	7.58	13	8.74	13	12.74	8	7.71
9	8.81	9	8.77	9	7.11	8	8.84
11	8.33	11	9.26	11	7.81	8	8.47
14	9.96	14	8.1	14	8.84	8	7.04
6	7.24	6	6.13	6	6.08	8	5.25
4	4.26	4	3.1	4	5.39	19	12.5
12	10.84	12	9.13	12	8.15	8	5.56
7	4.82	7	7.26	7	6.42	8	7.91
5	5.68	5	4.74	5	5.73	8	6.89
Average	9 7.5	9 7.5	9 7.5	9 7.5	9 7.5	9 7.5	9 7.5
Standard Deviation	3.31 2.03	3.31 2.03	3.31 2.03	3.31 2.03	3.31 2.03	3.31 2.03	3.31 2.03
Linear Regression	1.33	1.33	1.33	1.33	1.33	1.33	1.33

Graphics for Anscombe's Quartet



Tuning the Database

Complex Averages

Anscombe's Quartet

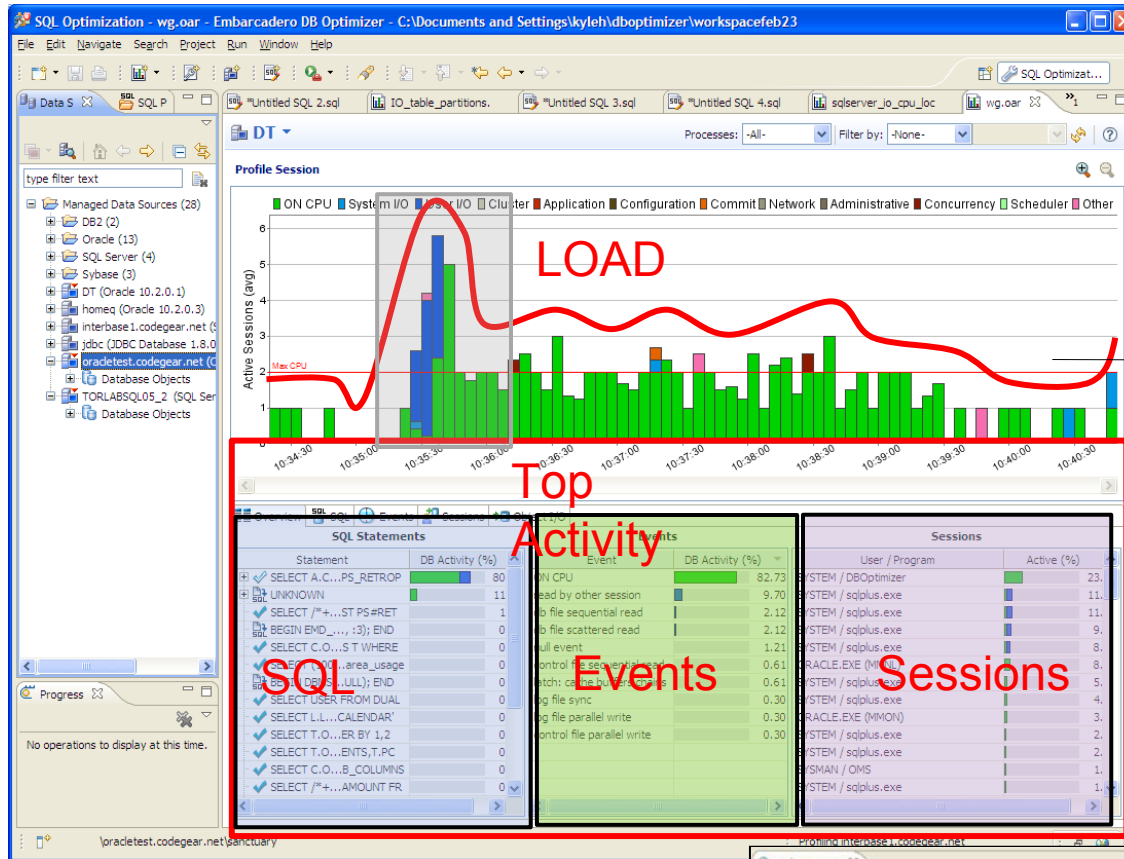
Average

Standard Deviation

Linear Regression

	I		II		III		IV	
	x	y	x	y	x	y	x	y
Average	9	7.5	9	7.5	9	7.5	9	7.5
Standard Deviation	3.31	2.03	3.31	2.03	3.31	2.03	3.31	2.03
Linear Regression	1.33		1.33		1.33		1.33	

How Can We Open the Black Box?



Max CPU
(yard stick)


Click here

Get Details

The 'Profiling Details' window shows the following table:

SQL Identification	Optimizer and Outline	Parsing Statistics	Execution Statistics (total)	Execution Statistics (per execution)
SQL ID 1143219105	Optimizer Mode ALL_ROWS	Memory 2793096	Fetches 0.00	Fetches 0.00
SQL Address 69877FD4	Parsing User ID 5	Loads 4	Executions 1	Executions 0.00
Child Address 698E0B98	Outline Category	Invalidations 0	Sorts 0	Sorts 0.00
Children 1	Outline STD 0		Disk Reads 0	Disk Reads 0.00
Plan Hash Value 2069026503			Buffer Gets 0	Buffer Gets 0.00
Module Executor.exe			Rows Processed 0	Rows Processed 0.00
Action			CPU Time 0.00	CPU Time 0.00
SQL Operation Code 3			Elapsed Time 0.00	Elapsed Time 0.00
Program ID 101646				
Program Line# 56				

How do we get our data?



Graph represents # of sessions active, but also represents amount of time active in the database
For Every Active Session there is a user (or application) waiting

Sample v\$session : fast and light weight

We collect

- session
- SQL
- State (CPU, IO, LOCK, Other Waits)
- and many other columns

AWR collects

SQL (v\$sqlstats)

F1qcyh20550cf
fj6gjjgsshtxyx
0cjsxw5ndqdbc
8t8as9usk11qw
dr1rkrznhh95b
10dkqv3kr8xa5
38zhkf4jdyff4
298wmz1kxjs1m

Waits (v\$system_event)

CPU
Enq: TX – row lock cont
SQL*Net break/reset to
db file scattered read
db file sequential read

Sessions (v\$sesstat)

25
34
36
38
45
63

Which session executed which SQL?

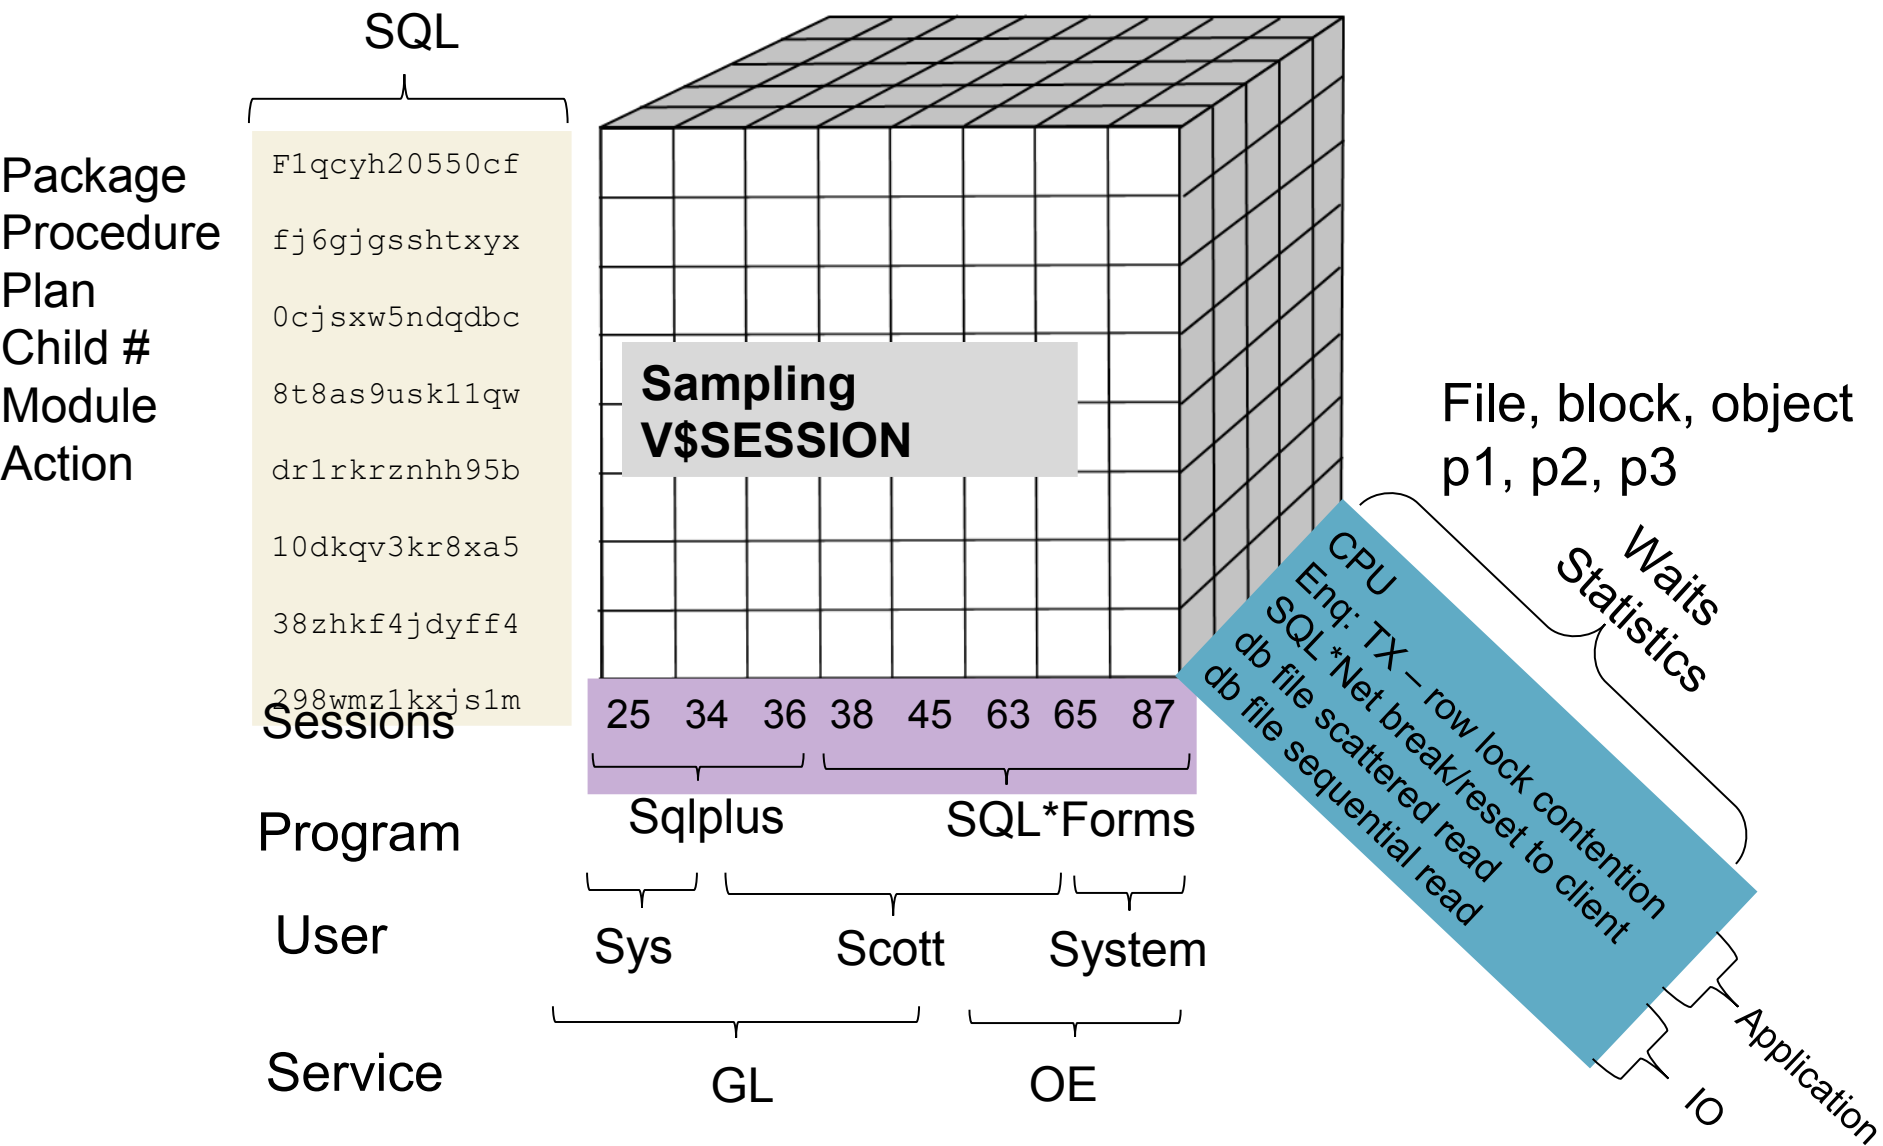
Which SQL was block on a wait event ?

AWR (and Statspack)
Can't answer!

Old Method




After sampling Multi-dimensional



DB Optimizer shows it all

Profile Session



SQL Statements		Events		Sessions	
Statement	DB Activity (%)	Event	DB Activity (%)	User / Program	Active (%)
INSERT INTO ...ALUES ('a')	88.52	buffer busy waits	60.74	SYSTEM	49.2
begin for i ...nd loop; end	7.41	ON CPU	30.74	SYSTEM	49.2
Non-SQL Activity	4.07	log file sync	6.30	SYSTEM	49.2
		log file parallel write	2.22	SYSTEM	47.8
				SYSTEM	

Event: buffer busy waits						
Count	Object Name	Object Type	SQL ID	Block Type	Tablespace Name	ASSM
164	FOO	TABLE	2745032144	data block	SYSTEM	MANUAL

• TABLE: Then the insert block is hot, so potential solutions are:

- Put the object in an ASSM tablespace.
- Use free lists.
- Hash partition the object.

Average Active Sessions (AAS)

Use CPU count as yardstick:

✓ $AAS < 1$

Database is not blocked


✓ $AAS \approx 0$

Database basically idle

Problems are in the APP not DB

❖ $AAS \gg \# \text{ of CPUs}$

There is a bottleneck



Interpreting the Load Chart

- 1. Application
 - Code inefficient?
- 1. Database
 - Configured correctly?
- 1. Machine
 - Is the machine undersized?
- 1. SQL
 - Inefficient SQL?

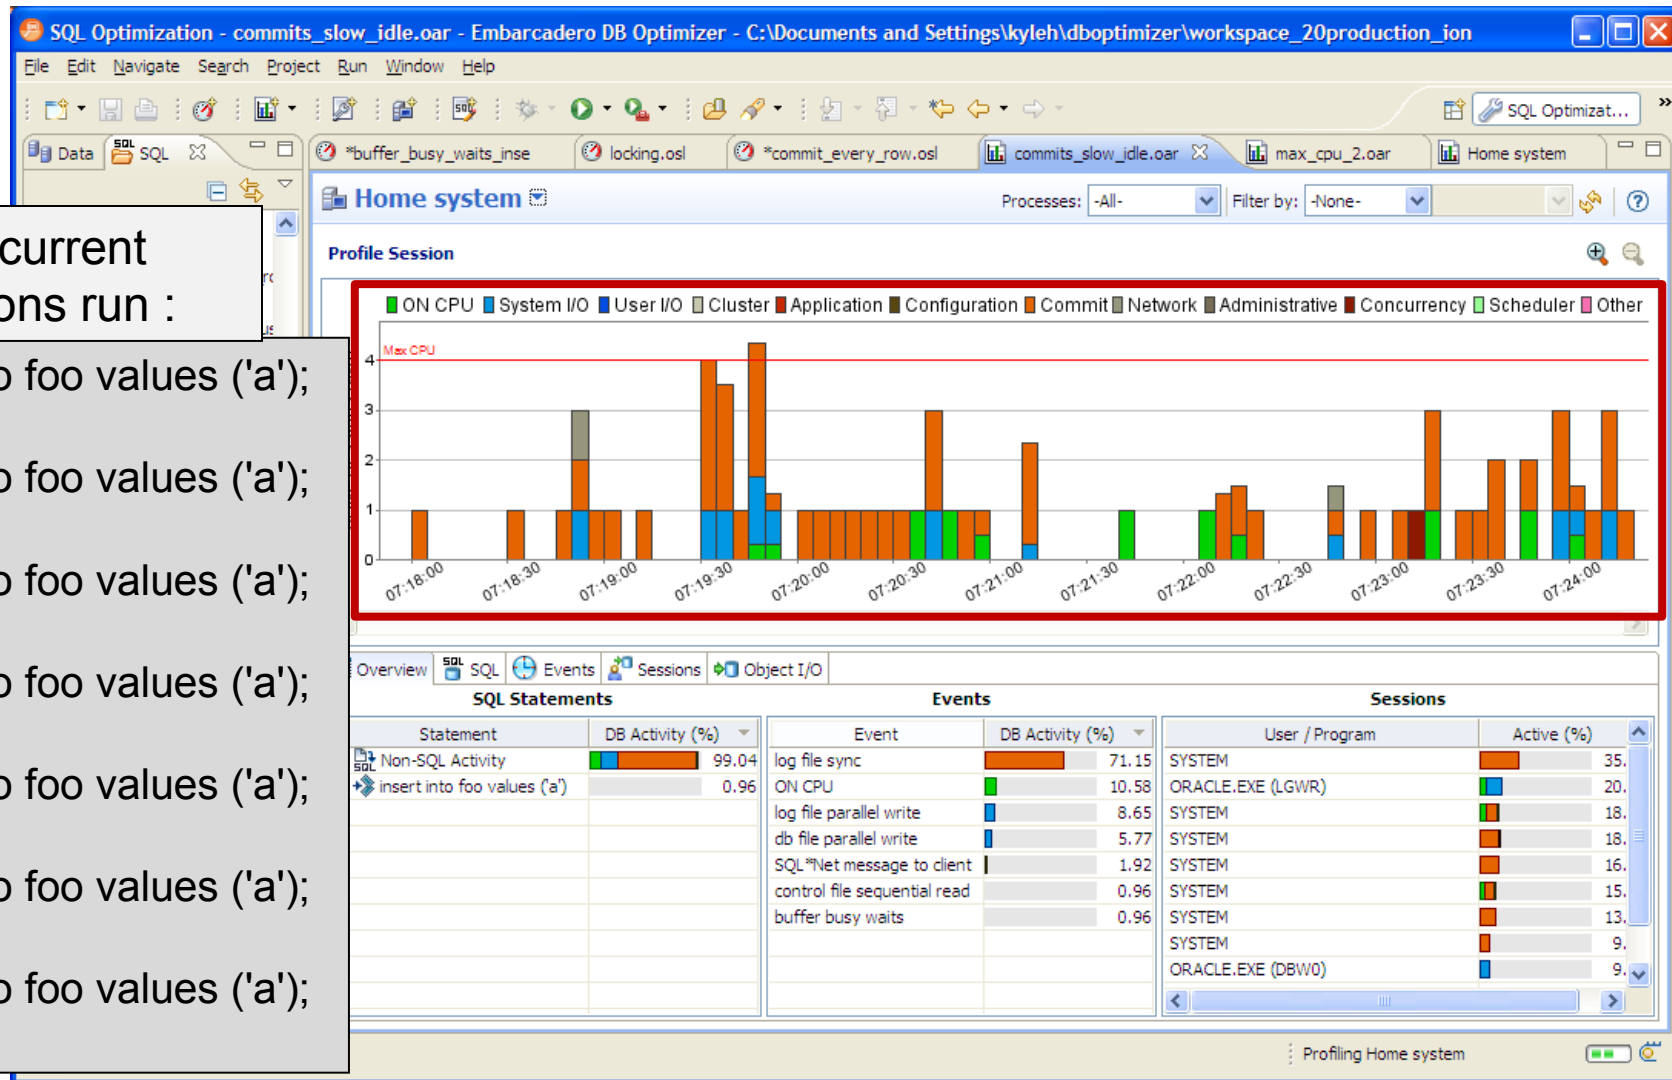
Who's problem is it ?

Chart answers this question and assists with solutions

1. Application Issues

4 concurrent sessions run :

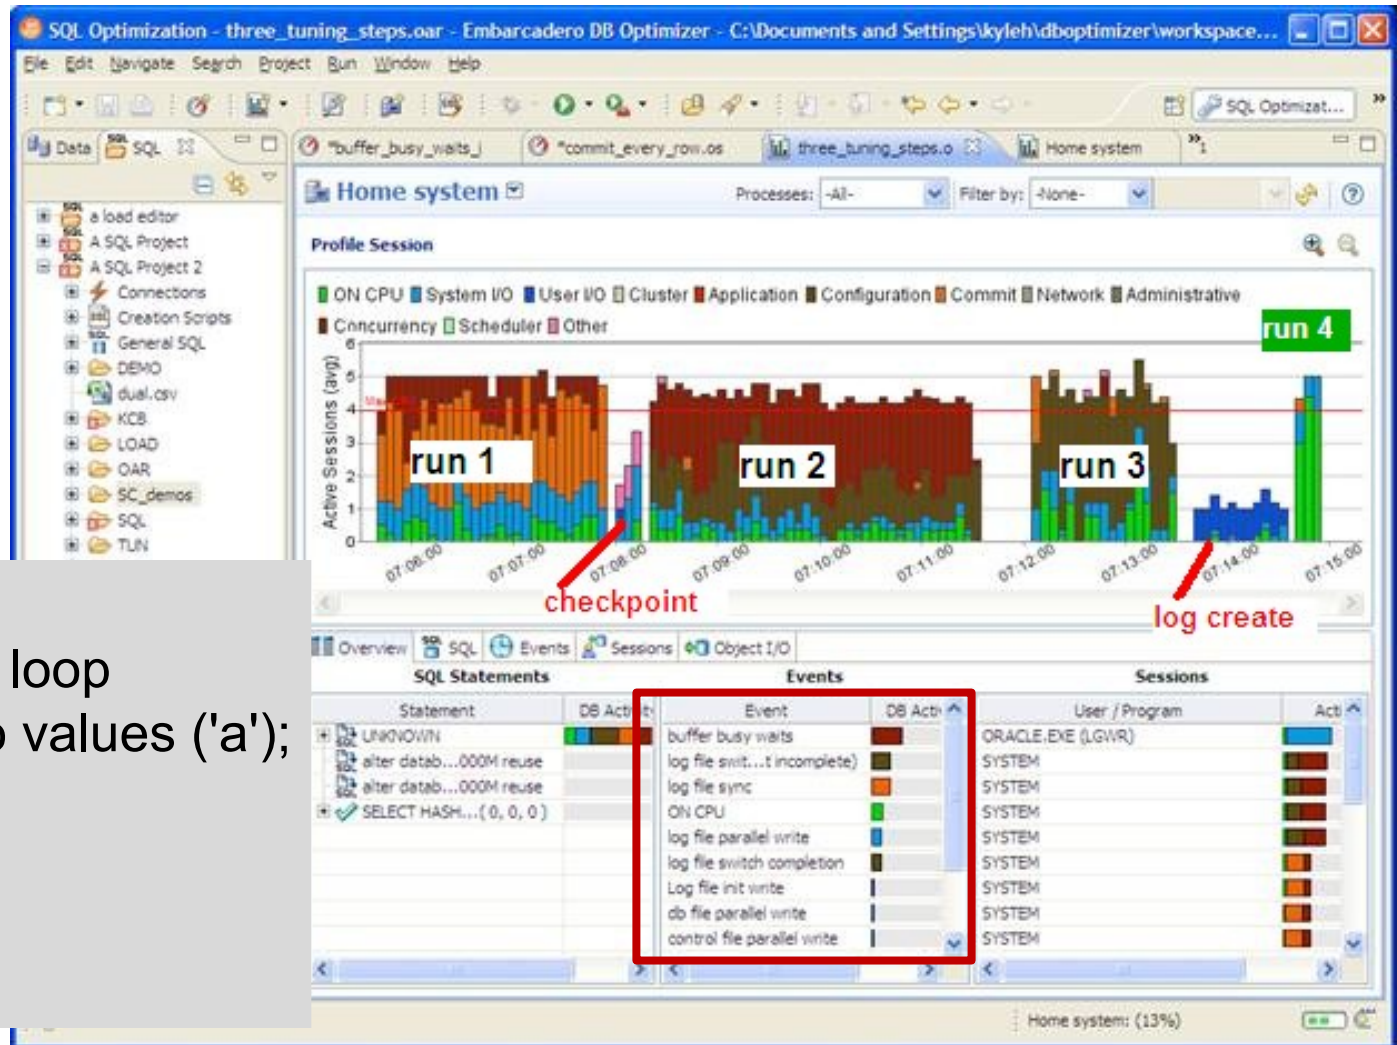
insert into foo values ('a');
commit;
insert into foo values ('a');
commit;
insert into foo values ('a');
commit;
insert into foo values ('a');
commit;
insert into foo values ('a');
commit;
insert into foo values ('a');
commit;
insert into foo values ('a');
commit;



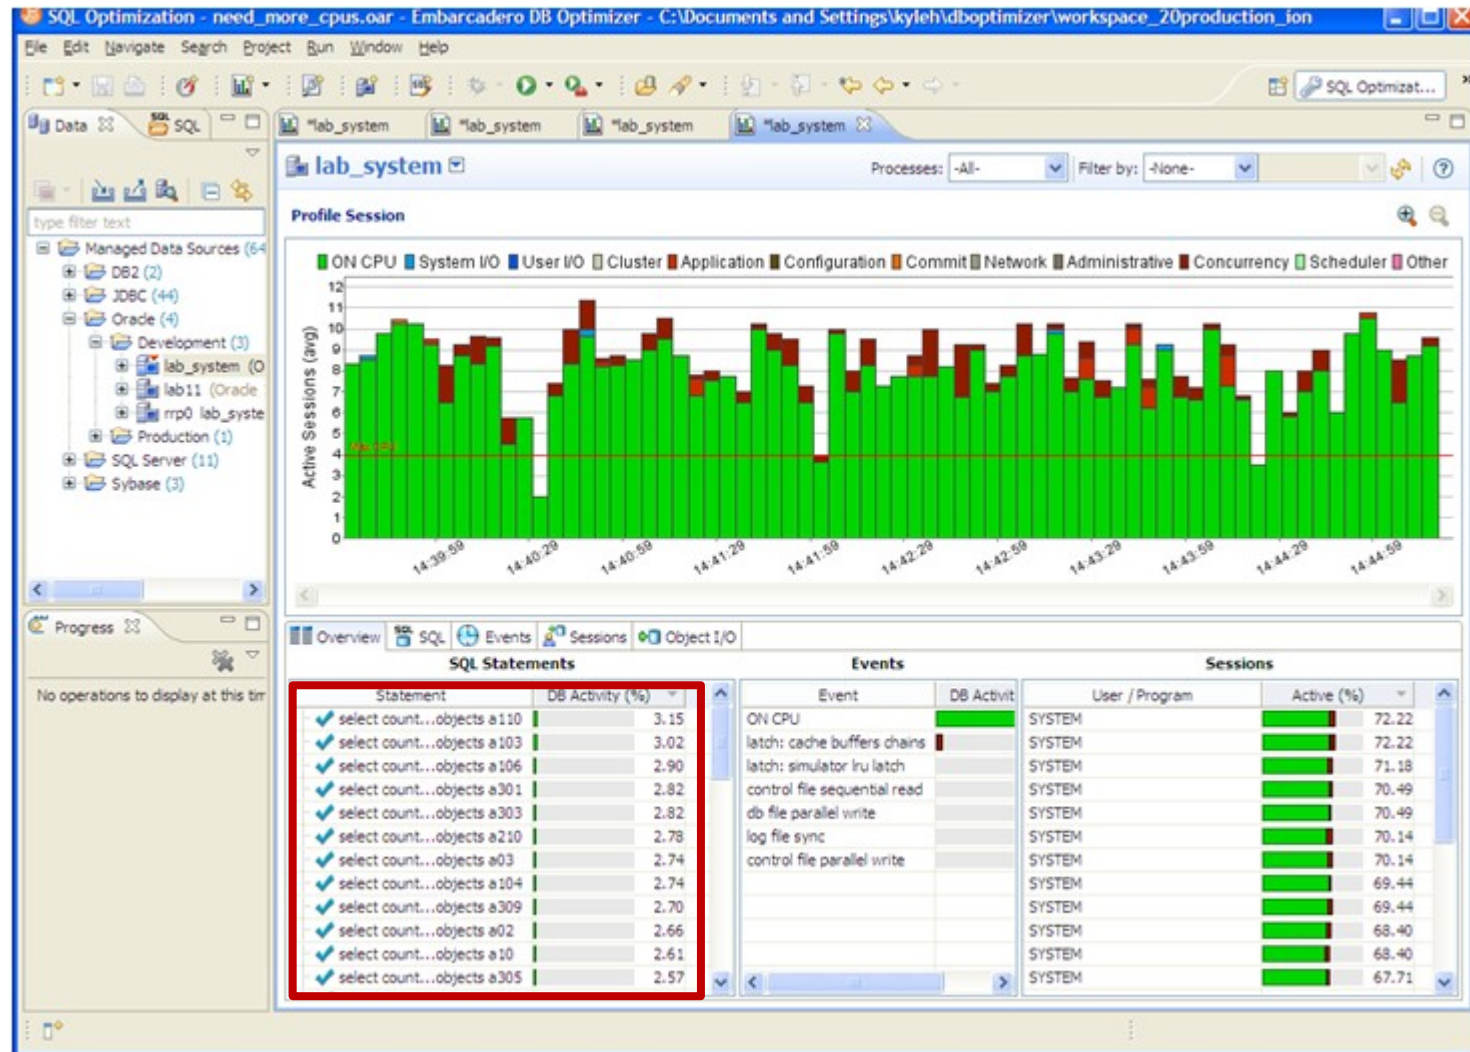
2. Database issue

4 concurrent sessions run :

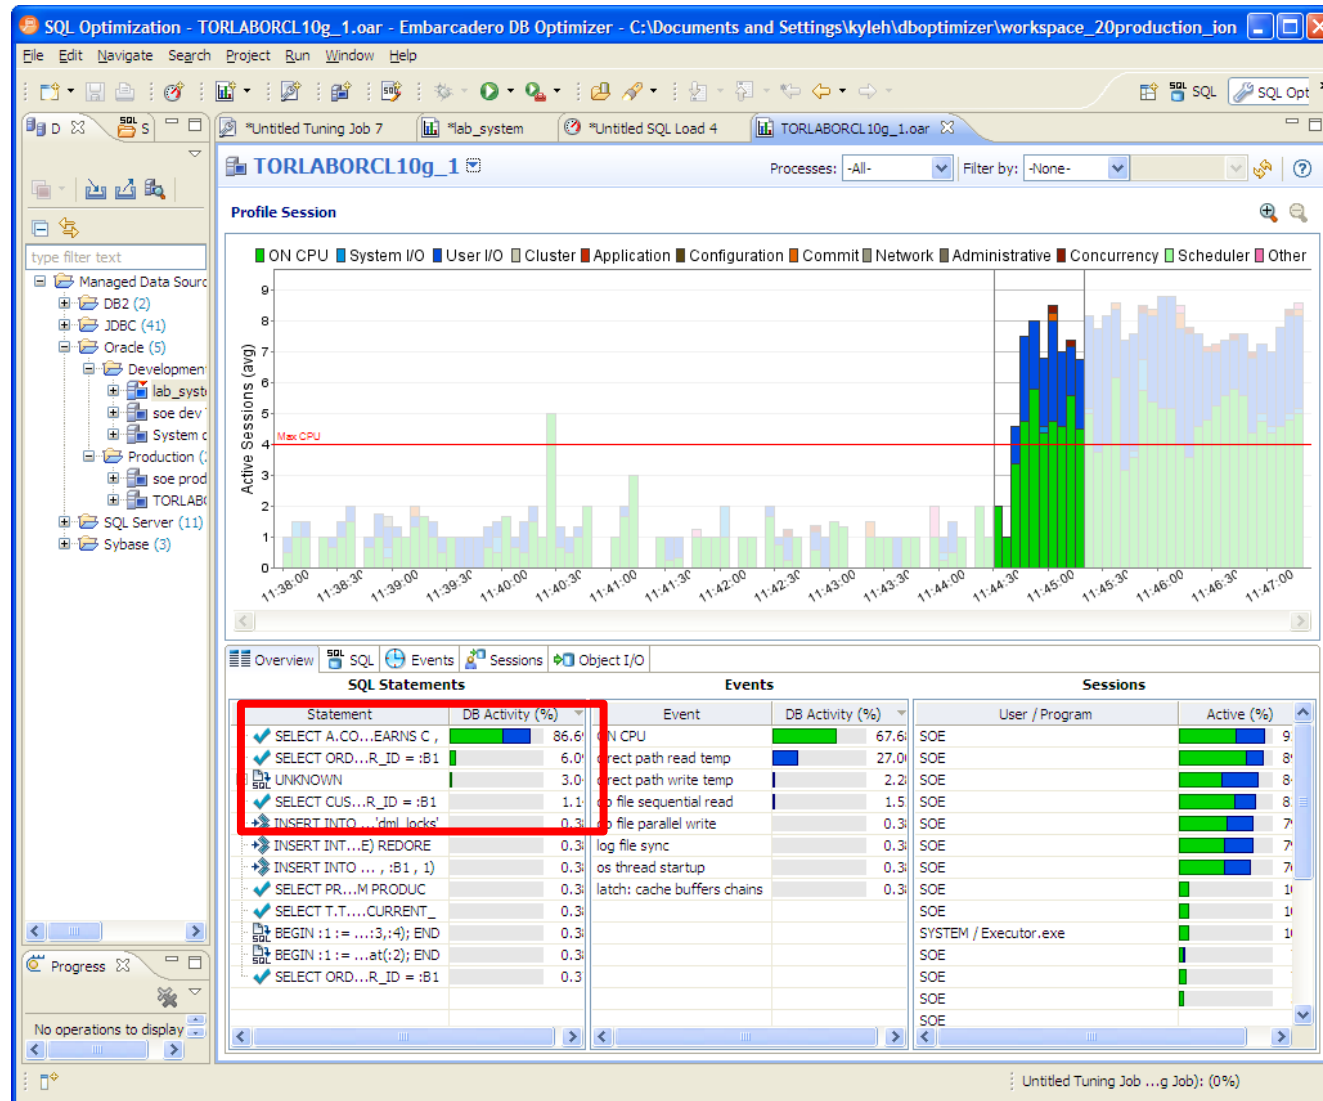
```
begin
  for i in 1..1000 loop
    insert into foo values ('a');
  end loop;
end;
/
Commit;
```



3. Machine Undersized



4. SQL needs Tuning



Tuning SQL Complexity

Trace file

```
PARSING IN CURSOR #2 len=53 dep=0 uid=61 oct=2  
lid=61 tim=1151519905950403 hv=2296704914  
ad='4e50010c'  
SELECT 'Hello, world; today is ' || SYSDATE FROM  
END OF STMT  
PARSE  
#2:c=4000,e=1540,p=0,cr=0,cu=0,mis=1,r=0,dep=0,  
tim=1151519905950397  
BINDS #2:  
EXEC  
#2:c=0,e=58,p=0,cr=0,cu=0,mis=0,r=0,dep=0,og=1,t  
1151519906034782  
WAIT #2: nam='SQL*Net message to client' ela= 2  
id=1650815232 #bytes=1 p3=0 obj#=-1  
tim=1151519906034809  
FETCH  
#2:c=0,e=29,p=0,cr=0,cu=0  
1151519906034864  
WAIT #2: nam='SQL*Net message to client' ela= 1  
driver id=1650815232 #bytes=1 p3=0 obj#=-1  
tim=1151519906035165  
FETCH  
#2:c=0,e=1,p=0,cr=0,cu=0,mis=0,r=0,dep=0,og=1,t  
1151519906035165  
WAIT #2: nam='SQL*Net message to client' ela= 1  
id=1650815232 #bytes=1 p3=0 obj#=-1  
tim=1151519906035400  
STAT #2 id=1 cnt=1 pid=0 pos=1 obj=0 op='FAST D  
(cr=0 pr=0 pw=0 time=3 us)
```

Execution Plan

Id	Operation	Name	Starts	E-Rows	A-Rows
1	HASH GROUP BY		1	1	1
* 2	FILTER		1		1909
* 3	TABLE ACCESS		1	1	3413
4			1	165	6827
			1	165	3413
			1	165	3624
	INDEX RANGE SCAN	WB_JOB	1	242	2895
			1	233	2897
	TABLE ACCESS BY INDEX ROWID	PS_PAY_CALENDAR	1	1	1
	INDEX RANGE SCAN	PS0PAY_CALENDAR	1	1	1
11	INDEX RANGE SCAN	WBBJOB_B	1	286	2895
* 12	TABLE ACCESS FULL	WB_RETROPAY_EARNS	1	27456	
13	TABLE ACCESS FULL	PS_RETROPAY_RQST	1	13679	13679
* 14	INDEX RANGE SCAN	PS#RETROPAYPGM_TBL	3413	1	3413
15	SORT AGGREGATE		1791	1	1791
16	FIRST ROW		1791	1	1579
* 17	INDEX RANGE SCAN (MIN/MAX)	WB_JOB_F	1791	1	1579
18	SORT AGGREGATE		1539	1	1539
19	FIRST ROW		1539	1	1539
* 20	INDEX RANGE SCAN (MIN/MAX)	WB_JOB_G	1539	1	1539

Complexity!

Visual SQL Tuning

SQL Optimization - A SQL Project/index_recommendation_3.tun - Embarcadero DB Optimizer - C:\Documents and Settings\kyleh\dboptimizer\workspace_20_1008

File Edit Navigate Search Project Run Window Help

SQL Optimizat...

index_recommendation_3.tun *riyaj.tun

Oracle lab_system (10.2.0.1)

Input Overview Analysis

SQL Analysis

Select statement of interest: SELECT 4

Operation

SELECT STATEMENT

NESTED LOOPS

NESTED LOOPS

TABLE ACCESS - SYS...IENT_TRANSACTI

INDEX - SYSTEM.CLIENT_PK

TABLE ACCESS - SYSTEM.INVESTMENT

INDEX - SYSTEM.INVESTMENT_PK

TABLE ACCESS - SYSTEM.INVESTMENT_TYPE

INDEX - SYSTEM.INVESTMENT_TYPE_PK

type filter text

Plan <

SELECT

ct.action,
c.client_id,
i.investment_unit,
it.investment_type_name

FROM

client_transaction ct,
client c,
investment_type it,
investment i

WHERE

ct.client_id = c.client_id AND
ct.investment_id = i.investment_id AND
i.investment_type_id = it.investment_type_id
ct.broker_commission = 100

CLIENT_TRANSACTION (ct)

CLIENT (c)

INVESTMENT (i)

INVESTMENT_TYPE (it)

INVESTMENT_TYPE (it)

INVESTMENT_TYPE_ID: NUMBER

INVESTMENT_TYPE_NAME: VARCHAR2

INVESTMENT_TYPE_PK

Index Analysis Table Statistics Column Statistics And Histograms Outlines

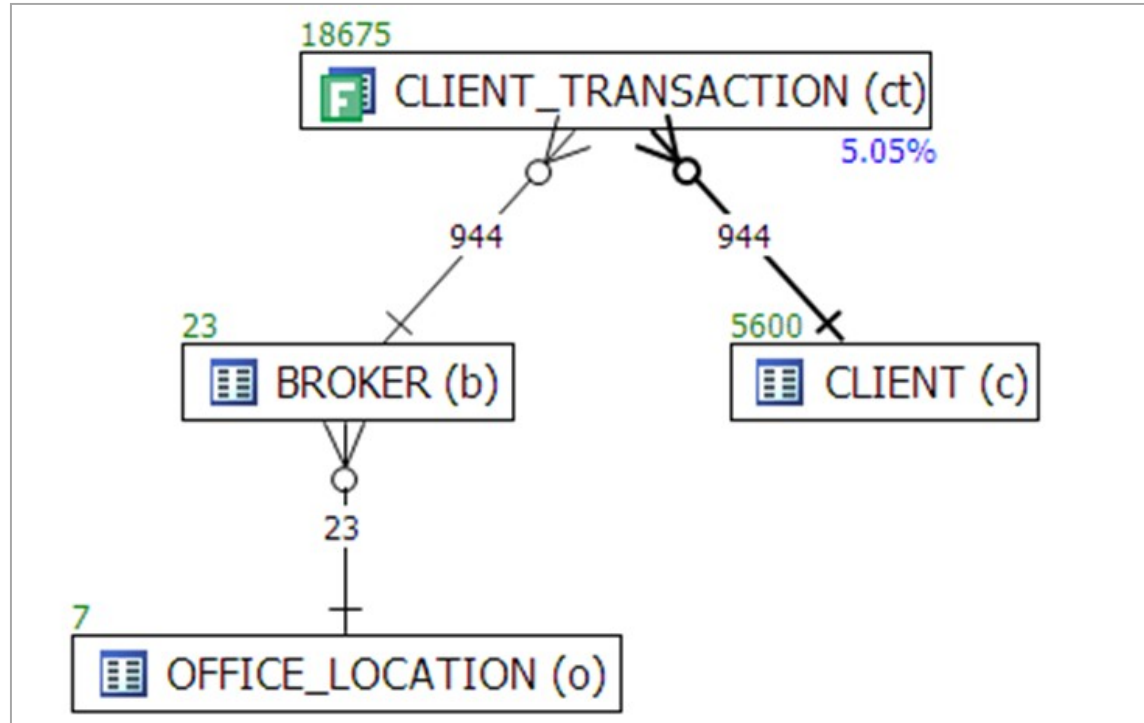
Collect and create indexes

Index Name	Table Owner	Table Name	Column Name	Index Type
IDX_CLIENT...SACTION_0	SYSTEM	CLIENT...ACTION	BROKER_COMMISSION	Normal
CLIENT_PK	SYSTEM	CLIENT	CLIENT_ID	Unique
INVESTMENT_PK	SYSTEM	INVESTMENT	INVESTMENT_ID	Unique
INVESTMENT_TYPE_PK	SYSTEM	INVESTMENT_TYPE	INVESTMENT_TYPE_ID	Unique
CLIENT_TRA...ION_CLIENT	SYSTEM	CLIENT...ACTION	CLIENT_ID	Normal
CLIENT_TRA...INVESTMENT	SYSTEM	CLIENT...ACTION	INVESTMENT_ID	Normal
INVESTMENT...TMENT_TYPE	SYSTEM	INVESTMENT	INVESTMENT_TYPE_ID	Normal
CLIENT_BROKER	SYSTEM	CLIENT	BROKER_ID	Normal
CLIENT_INCOME	SYSTEM	CLIENT	CLIENT_HOUSEHOLD_INCOME	Normal
CLIENT_MULTI	SYSTEM	CLIENT	CLIENT_FIRST...NT_LAST_NAME	Normal

Embarcadero

Tuning SQL

- Indexes
 - Missing indexes
- Table Stats
 - Stale statistics
- Histograms
 - Recommendations
- Execution Plan
 - DB Optimizer automatically finds better plans
- Visual SQL Tuning diagrams



VST Steps

1. Tables

- drawn as nodes

1. Joins

- drawn as connector lines

1. Filters

- mark on each table with filter in where clause

How to VST: Tables and Joins

```
SELECT C.Phone_Number, C.Honorific, C.First_Name, C.Last_Name,  
       C.Suffix, C.Address_ID, A.Address_ID, A.Street_Address_Li  
       A.Street_Address_Line2, A.City_Name, A.State_Abbreviation  
       A.ZIP_Code, OD.Deferred_Shipment_Date, OD.Item_Count,  
       ODT.Text, OT.Text, P.Product_Description, S.Shipment_Date  
FROM Orders O, Order_Details OD, Products P, Customers C, Shipme  
       Addresses A, Code_Translations ODT, Code_Translations OT  
WHERE UPPER(C.Last_Name) LIKE :Last_Name||'%'  
       AND UPPER(C.First_Name) LIKE :First_Name||'%'  
       AND OD.Order_ID = O.Order_ID  
       AND O.Customer_ID = C.Customer_ID  
       AND OD.Product_ID = P.Product_ID(+)  
       AND OD.Shipment_ID = S.Shipment_ID(+)  
       AND S.Address_ID = A.Address_ID(+)  
       AND O.Status_Code = OT.Code  
       AND OT.Code_Type = 'ORDER_STATUS'  
       AND OD.Status_Code = ODT.Code  
       AND ODT.Code_Type = 'ORDER_DETAIL_STATUS'  
       AND O.Order_Date > :Now - 366  
ORDER BY C.Customer_ID, O.Order_ID DESC, S.Shipment_ID, OD.Order
```

Tables

Orders O,
Order_Details OD,
Products P,
Customers C,
Shipments S,
Addresses A,
Code_Translations ODT,
Code_Translations OT

Joins

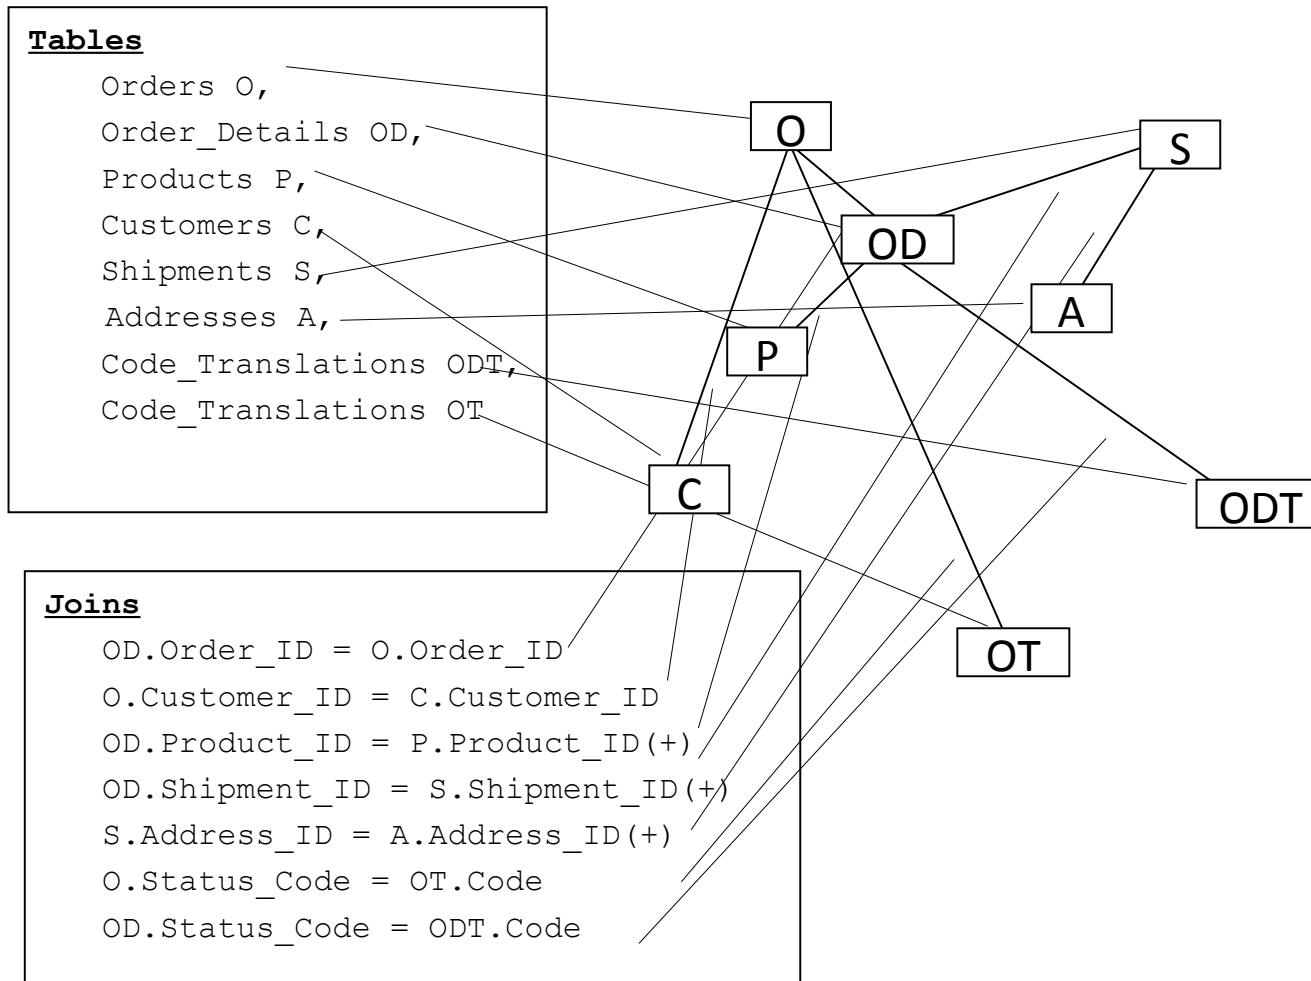
OD.Order_ID = O.Order_ID
O.Customer_ID = C.Customer_ID
OD.Product_ID = P.Product_ID(+)
OD.Shipment_ID =
S.Shipment_ID(+)
S.Address_ID = A.Address_ID(+)
O.Status_Code = OT.Code
OD.Status_Code = ODT.Code

Filters

WHERE UPPER(C.Last_Name) LIKE :Last_Name||'%'
 AND UPPER(C.First_Name) LIKE :First_Name||'%'
 AND OT.Code_Type = 'ORDER_STATUS'
 AND O.Order_Date > :Now - 366
 AND ODT.Code_Type = 'ORDER_DETAIL_STATUS'

Dan Tow – SQL TUNING


Layout tables and connections



Unstructured

Joins

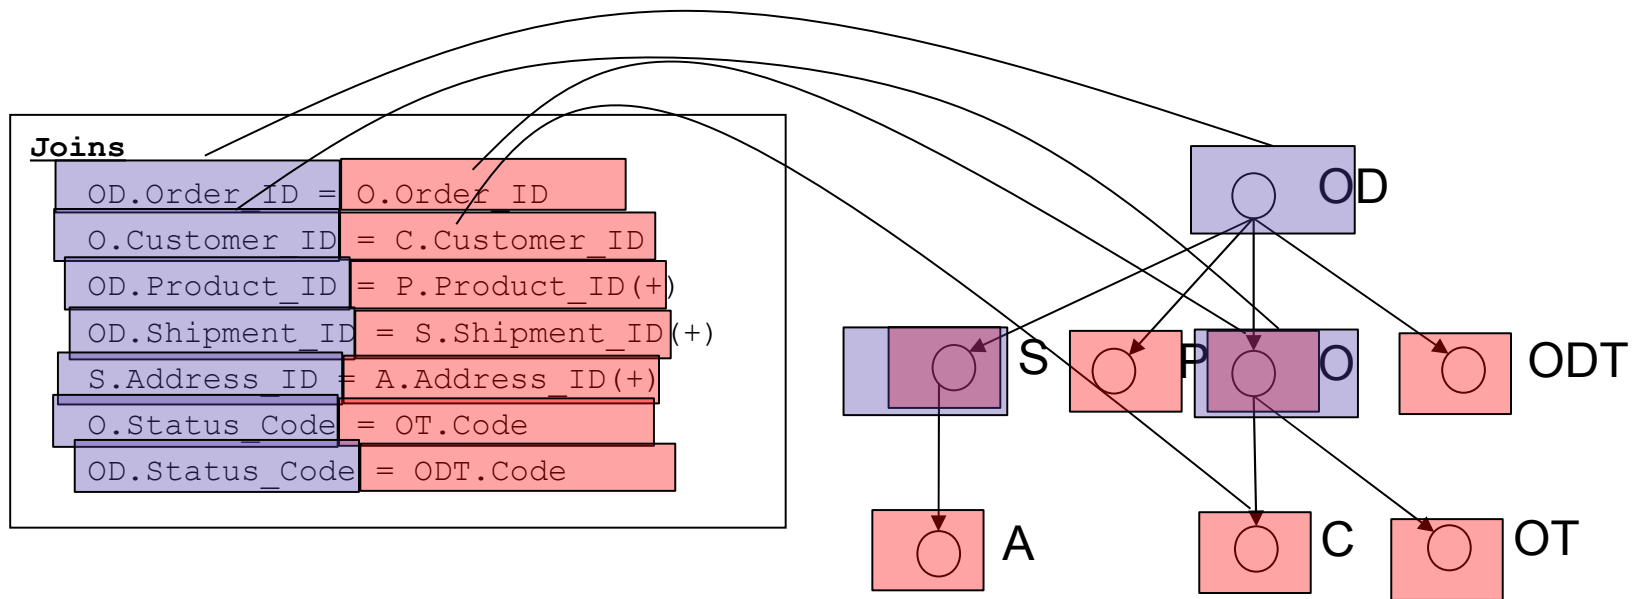

```
OD.Order_ID = O.Order_ID  
O.Customer_ID = C.Customer_ID  
OD.Product_ID = P.Product_ID(+)  
OD.Shipment_ID = S.Shipment_ID(+)  
S.Address_ID = A.Address_ID(+)  
O.Status_Code = OT.Code  
OD.Status_Code = ODT.Code
```



Neater, but can you do anything with it?
What's the optimal execution path?

Parents and Children

Structure
the
tree




VST – filters and best path

➤ Filters help determine best path

Filters



```
WHERE UPPER(C.Last_Name) LIKE :Last_Name||'%'
      AND UPPER(C.First_Name) LIKE :First_Name||'%'
      AND OT.Code_Type = 'ORDER_STATUS'
      AND ODT.Code_Type = 'ORDER_DETAIL_STATUS'
      AND O.Order_Date > :Now - 366
```



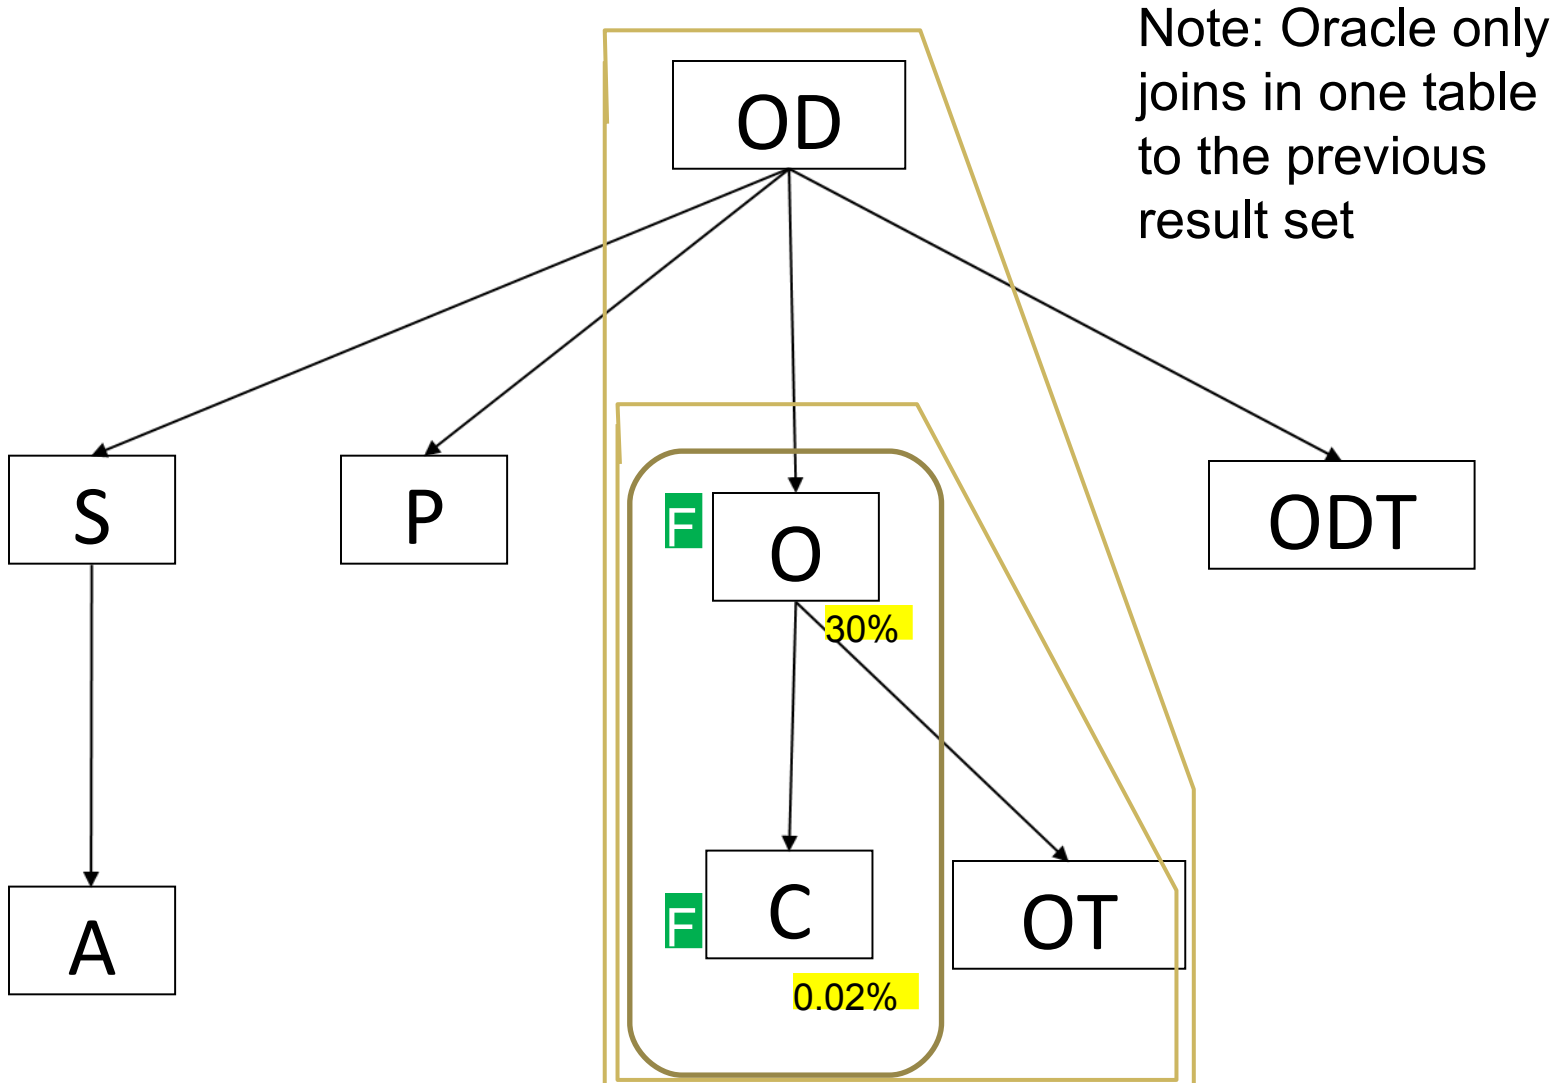
```
100% * (select count(*) from TAB where condition)
-----
(select count(*) from Tab)
```

Concept:

1. Start at most selective filter
2. Join down first, before joining upwards




VST – best path



Cartesian

```
SELECT
  A.BROKER_ID BROKER_ID,
  A.BROKER_LAST_NAME BROKER_LAST_NAME,
  A.BROKER_FIRST_NAME BROKER_FIRST_NAME,
  A.YEARS_WITH_FIRM YEARS_WITH_FIRM,
  C.OFFICE_NAME OFFICE_NAME,
  SUM (B.BROKER_COMMISSION)
TOTAL_COMMISSIONS
FROM
  BROKER A,
  CLIENT_TRANSACTION B,
  OFFICE_LOCATION C,
  INVESTMENT I
WHERE
  A.BROKER_ID = B.BROKER_ID AND
  A.OFFICE_LOCATION_ID =
C.OFFICE_LOCATION_ID
GROUP BY
  A.BROKER_ID,
  A.BROKER_LAST_NAME,
  A.BROKER_FIRST_NAME,
  A.YEARS_WITH_FIRM,
  C.OFFICE_NAME;
```

 INVESTMENT (I)



Implied Cartesian

```
select
  c.client_first_name, c.client_last_name,
  ct.action, ct.price,
  b.broker_last_name, b.broker_first_name,
  o.office_name
from
  client_transaction ct,
  client c,
  broker b,
  office_location o
where
  ct.price > 100
  and b.broker_id=ct.broker_id
  and c.broker_id = b.broker_id
  and o.office_location_id = b.office_location_id
```

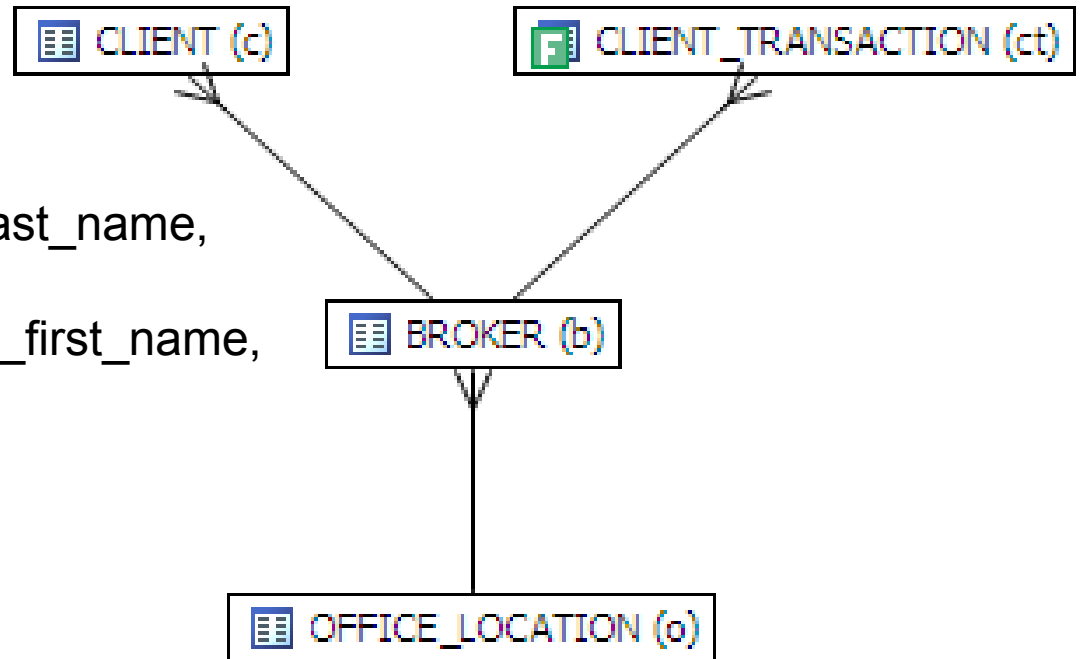


Diagram work for Many to One



One to many



Many to many





What about many to many?



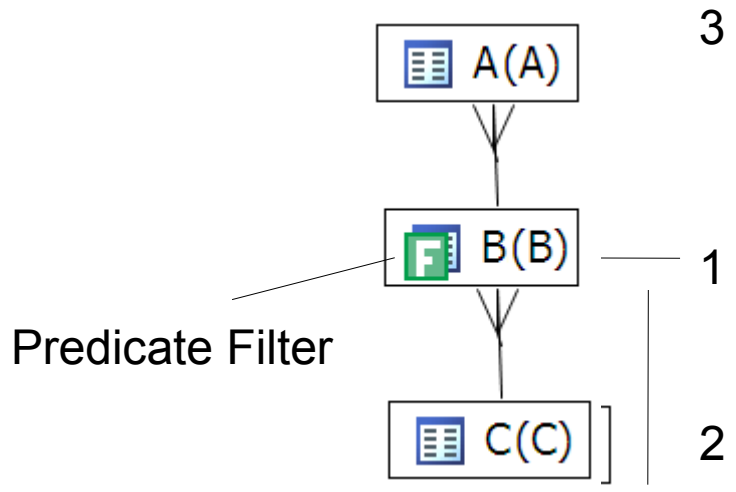
Unstructured

Joins


```
OD.Order_ID = O.Order_ID  
O.Customer_ID = C.Customer_ID  
OD.Product_ID = P.Product_ID(+)  
OD.Shipment_ID = S.Shipment_ID(+)  
S.Address_ID = A.Address_ID(+)  
O.Status_Code = OT.Code  
OD.Status_Code = ODT.Code
```



Many-to-One vs Many-to-Many



$B \rightarrow C \rightarrow A$



go to A or C?

Adding Constraints

```
SELECT COUNT (*)  
FROM
```

```
  b,
```

```
  c,
```

```
  a
```


```
WHERE
```

```
  b.val2 = 100 AND
```

```
  a.val1 = b.id AND
```


```
  b.val1 = c.id;
```

58 logical reads



```
alter table c add constraint c_pk unique (id);  
alter table b add constraint b_pk unique (id);
```


7 logical reads



Join sizes

Join Sizes

Use the Join sizes to determine path of least resistance



Look at 3 queries

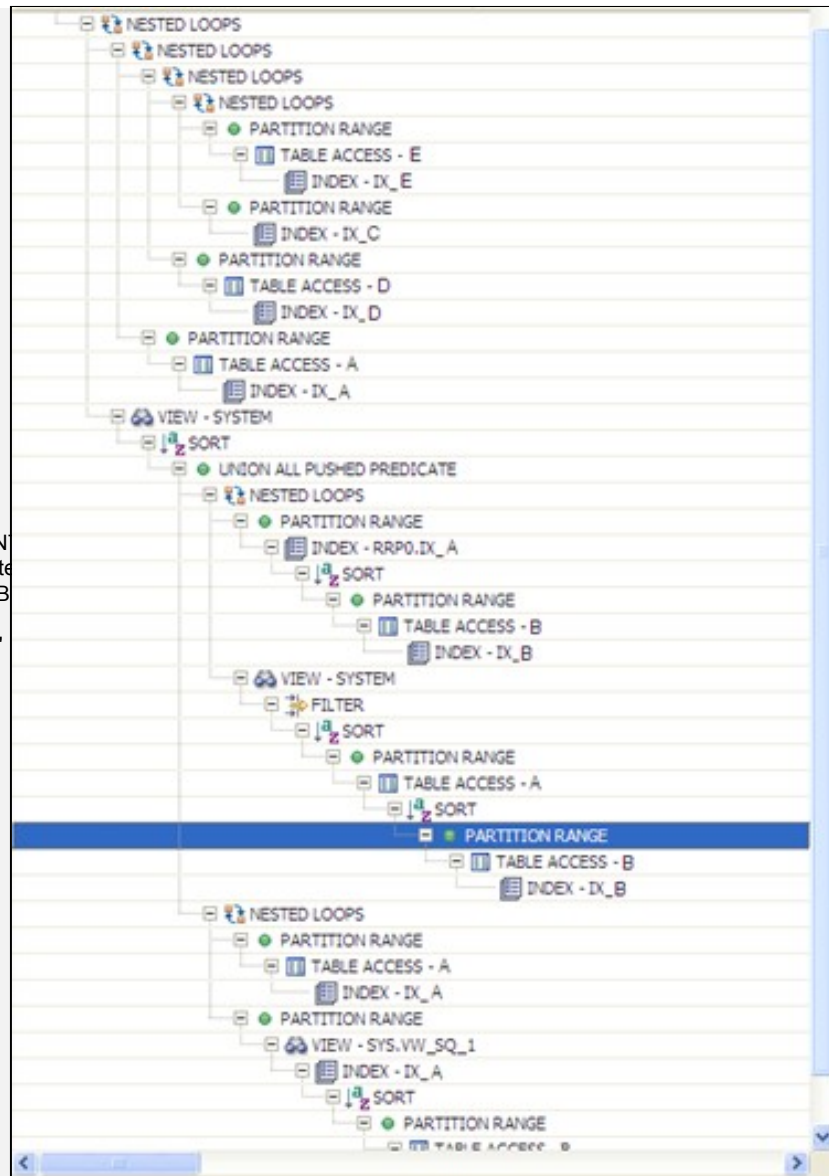
- Query 1 runs more than 24 hours
- Query 2 outer joins and scalar subqueries
- Query 3 create path not available to Oracle

Query 1 : Over 24 hours to run

```

SELECT
  A0.zuchinis,
  A0.brocoli,
  C0.Oranges
FROM
  (
    SELECT
      A1.planted_date,
      A1.pears,
      A1.zuchinis,
      A1.brocoli
    FROM
      FOO.A A1,
      (
        SELECT
          zuchinis,
          brocoli
        FROM FOO.A A2
        WHERE
          pears = 'M' AND
          planted_date + 0 >= ADD_MONTHS(
            MAX(planted_date),
            - 11)
        GROUP BY
          zuchinis,
          brocoli
        HAVING COUNT(*) = 12
      )
      i2
    WHERE
      A1.planted_date = (SELECT
        MAX(planted_date)
        FROM FOO.B B2
        WHERE
          pears = 'M'
      ) AND
      A1.pears = 'M' AND
      A1.zuchinis = i2.zuchinis (+) AND
      A1.brocoli = i2.brocoli (+)
    UNION
    SELECT
      A4.planted_date,

```

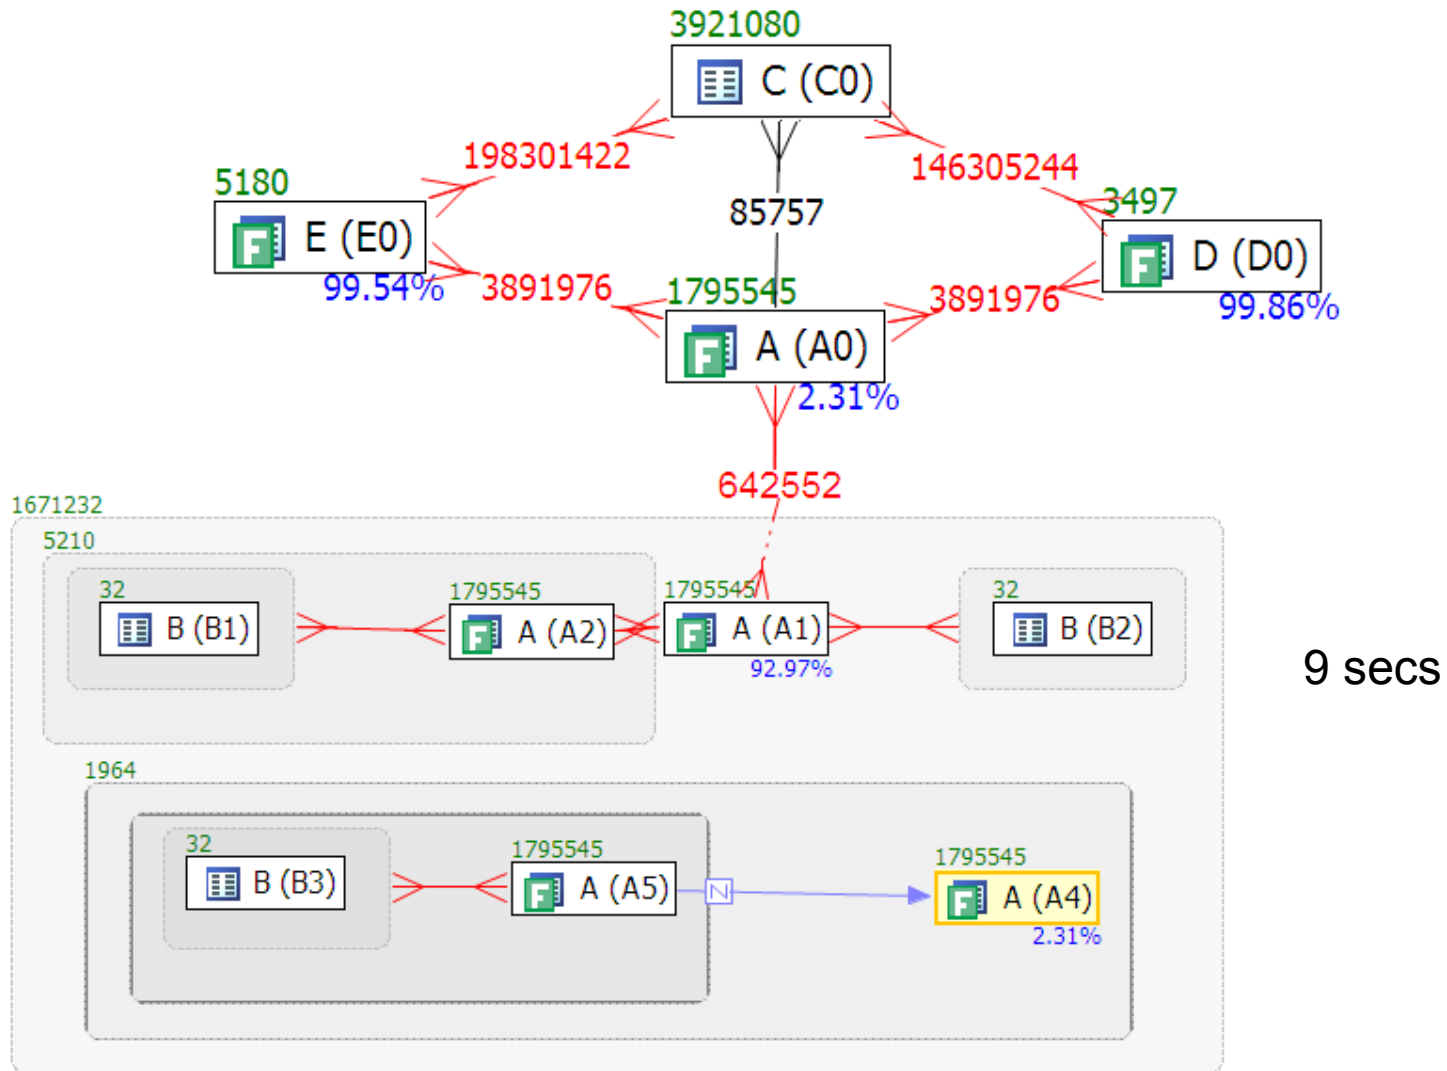


A4.planted_date <'03-OCT-08' AND

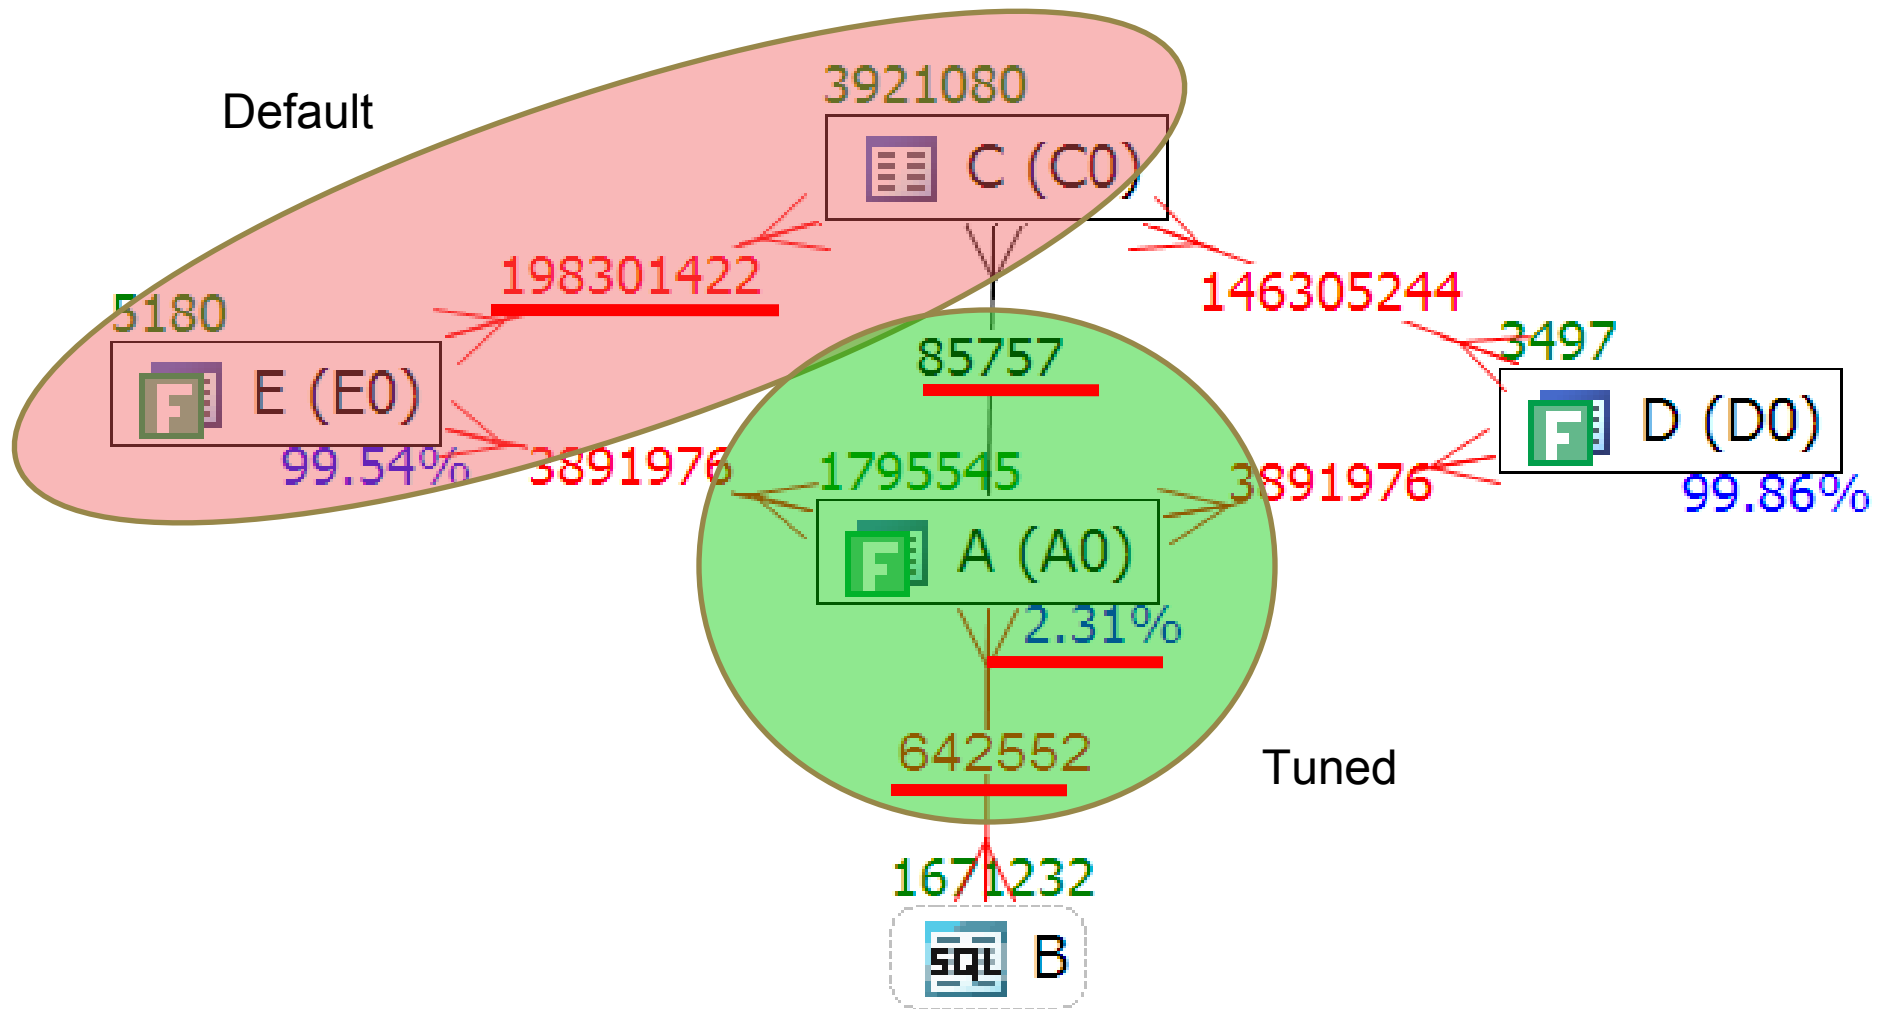
date)

s AND

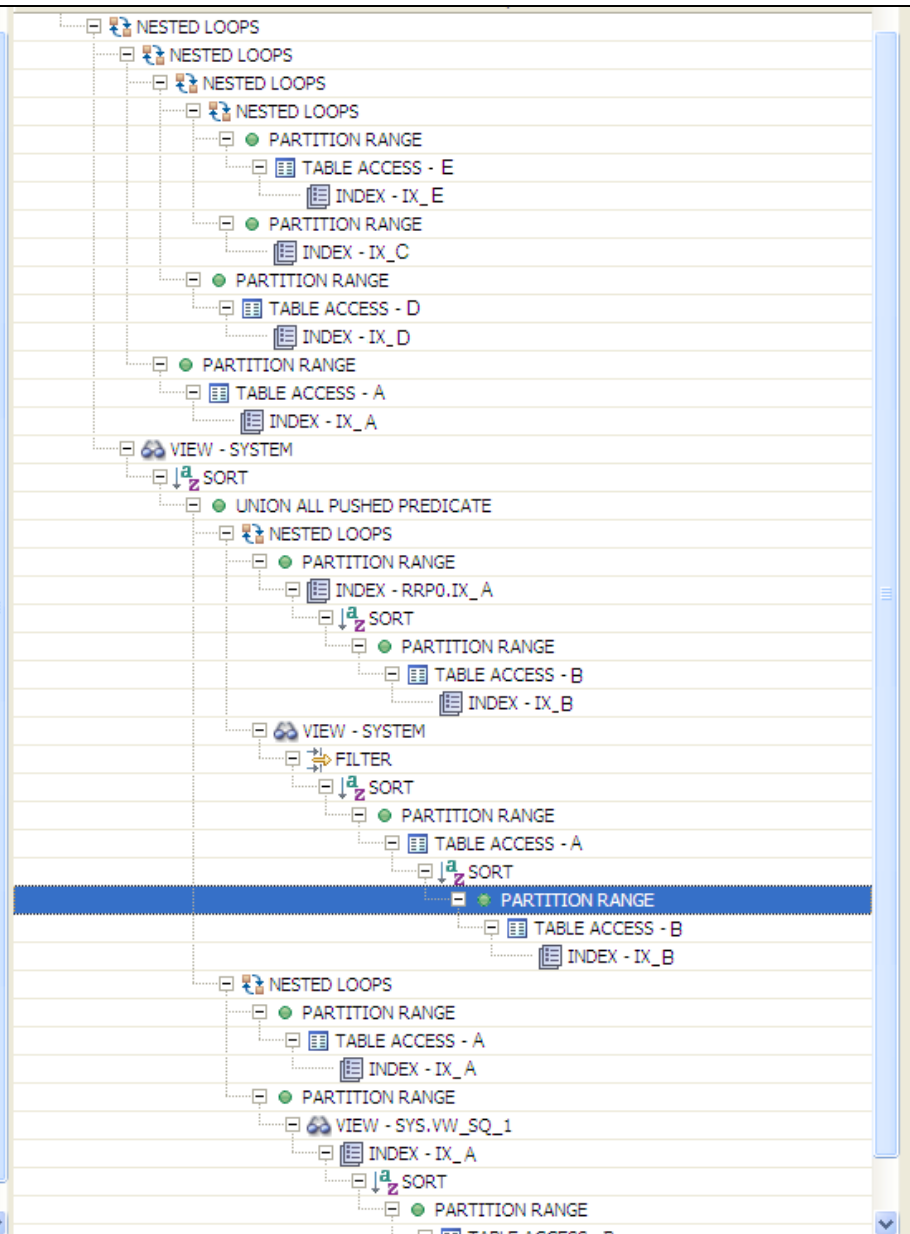

Visual SQL Diagram



Default vs Tuned



Comparing Plans : 24 hours to 5 mins



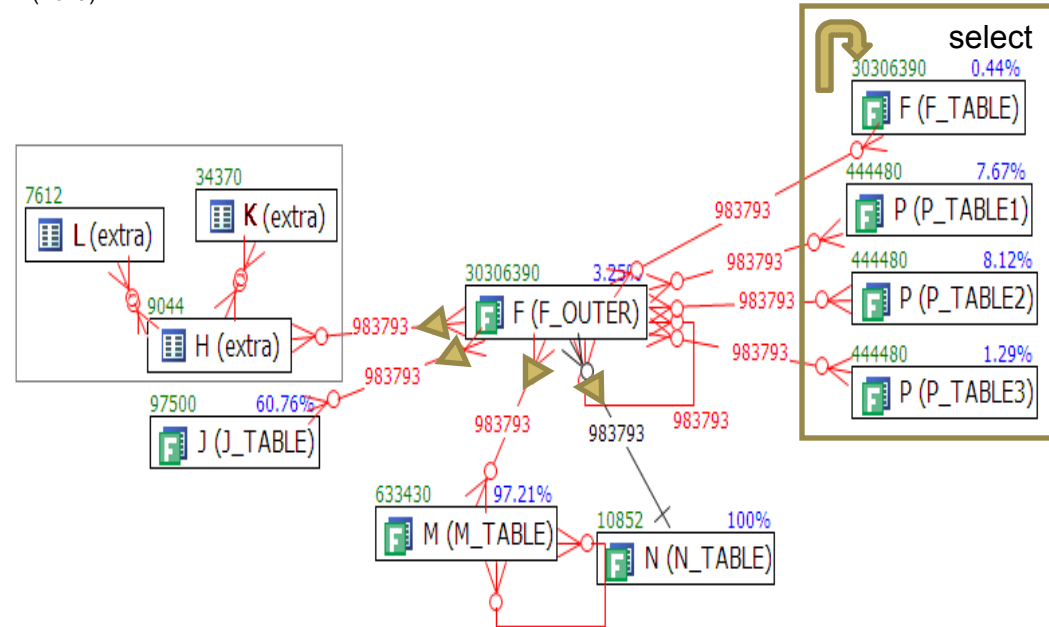
Q2

```
SELECT CASE WHEN M.NYC IS NULL THEN (SELECT /*+ qb_name(qb1) */ MAX (Kona)
FROM foo.F
```

```

WHERE harvest_date = to_date('08/10/2008','dd/mm/yyyy')
AND Argentina = TRIM('D') AND Norway = F_OUTER.Norway
ELSE M.NYC END AS NYC,
CASE WHEN F_OUTER.Perth IS NULL THEN NULL
ELSE (SELECT /*+ qb_name(qb2) */ Georgia FROM foo.P
WHERE harvest_date = to_date('08/10/2008','dd/mm/yyyy')
AND Argentina = TRIM('D') AND Paris = F_OUTER.Perth)
END AS richard,
CASE WHEN F_OUTER.Aruba IS NULL THEN NULL
ELSE (SELECT /*+ qb_name(qb3) */ Georgia FROM foo.P
WHERE harvest_date = to_date('08/10/2008','dd/mm/yyyy')
AND Argentina = TRIM('D') AND Paris = F_OUTER.Aruba)
END AS Jody,
CASE WHEN F_OUTER.Portland IS NULL THEN NULL
ELSE (SELECT /*+ qb_name(qb4) */ Georgia FROM foo.P
WHERE harvest_date = to_date('08/10/2008','dd/mm/yyyy')
AND Argentina = TRIM('D') AND Paris = F_OUTER.Portland)
END AS Tom
FROM foo.F F_OUTER, foo.M , foo.J , foo.N ,
(SELECT /*+ qb_name(qb5) */ H.SF, Oregon, H.Haiti, K.Bermuda, L.Denmark
FROM (foo.H LEFT OUTER JOIN foo.K
ON H.harvest_date = K.harvest_date
AND H.Argentina = K.Argentina AND H.SF = K.SF
AND K.Dallas = '001')
LEFT OUTER JOIN FOo.L
ON H.harvest_date = L.harvest_date
AND H.Argentina = L.Argentina AND H.SF = L.SF
WHERE H.harvest_date = to_date('08/10/2008','dd/mm/yyyy')
AND H.Argentina = TRIM('D')) extra
WHERE F_OUTER.harvest_date = M.harvest_date(+)
AND F_OUTER.Argentina = M.Argentina(+)
AND F_OUTER.Norway = M.Norway(+)
AND M.Norway(+) = M.Texas(+)
AND F_OUTER.harvest_date = to_date('08/10/2008','dd/mm/yyyy')
AND F_OUTER.Argentina = TRIM('D')
AND M.harvest_date(+) = to_date('08/10/2008','dd/mm/yyyy')
AND M.Argentina(+) = TRIM('D')
AND F_OUTER.Norway = F_OUTER.Hawaii
AND F_OUTER.harvest_date = J.harvest_date(+)
AND F_OUTER.Argentina = J.Argentina(+)
AND F_OUTER.Norway = J.Texas(+)
AND J.harvest_date(+) = to_date('08/10/2008','dd/mm/yyyy')
AND J.Argentina(+) = TRIM('D')
AND F_OUTER.Iraq = extra.SF(+)
AND F_OUTER.harvest_date = N.harvest_date(+)

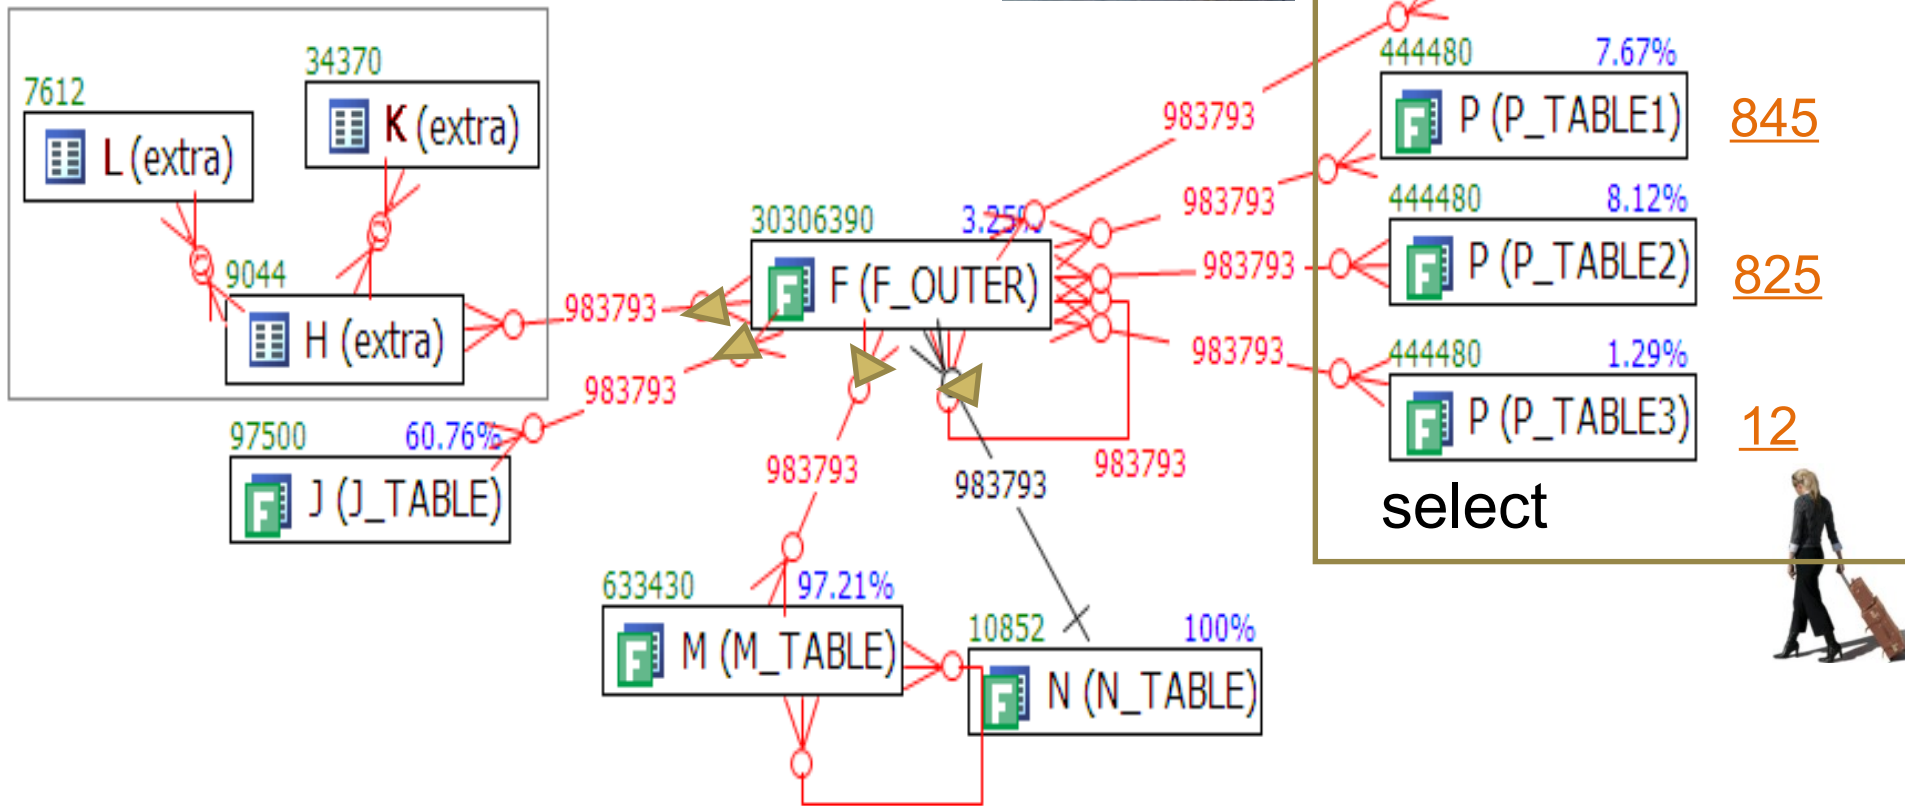
```



Two important qualities:

- All outer joins to F_OUTER
- 4 subqueries in select

Q2



Q2

The subqueries in the select clause look like


```
select CASE WHEN F.f1 IS NULL
      THEN NULL
      ELSE (SELECT X.f2
            FROM X
            WHERE code_vl = F.f2)
      END AS f0
from F;
```

and should be merged into the query like:


```
select CASE WHEN F.f1 IS NULL
      THEN NULL
      ELSE ( X.f2)
      END AS f0
from F , X
where code_vl(+) = F.f1;
```

Q3

```
SELECT DISTINCT *
FROM
  FOO.a a, FOO.c c, FOO.d d, FOO.g g
WHERE
  a.planted_date > '01-OCT-08' AND
  a.planted_date < '03-OCT-08' AND
  a.pears = 'D' AND a.green_beans = '1' AND
  a.planted_date = c.planted_date AND
  a.pears = c.pears AND
  a.zuchinis = c.zuchinis AND
  a.brocoli = c.brocoli AND
  a.planted_date = d.planted_date AND
  a.pears = d.pears AND
  a.harvest_size = d.harvest_size AND
  c.oranges = d.oranges AND
  c.apples = d.apples AND
  (d.lemons = 0 OR d.lemons IS NULL) AND
  a.planted_date = g.planted_date AND
  a.pears = g.pears AND
  a.harvest_size = g.harvest_size AND
  c.oranges = g.oranges AND
  c.apples = g.apples AND
  (g.lemons = 0 OR g.lemons IS NULL) AND
  a.zuchinis = '0236' AND
  d.apples = g.apples AND
  d.oranges = g.oranges
ORDER BY a.zuchinis, a.brocoli;
```



Q3: Transitivity



Q3

```
SELECT * FROM
```

```
(
  SELECT /*+ NO_MERGE */ c.apples, c.oranges, a.harvest_size
  FROM a, c
  WHERE
    a.planted_date = TO_DATE ('02/10/2008', 'dd/mm/yyyy') AND
    a.pears = 'D' AND
    a.green_beans = '1' AND
    a.planted_date = c.planted_date AND
    a.pears = c.pears AND
    a.zuchinis = c.zuchinis AND
    a.brocoli = c.brocoli AND
    a.zuchinis = '0236'
```

```
) X,
```

```
(
  SELECT /*+ NO_MERGE */ d.apples, d.oranges, d.harvest_size
  FROM d, g
  WHERE
    d.planted_date = TO_DATE ('02/10/2008', 'dd/mm/yyyy') AND
    g.planted_date = TO_DATE ('02/10/2008', 'dd/mm/yyyy') AND
    g.apples = d.apples AND
    d.oranges = g.oranges AND
    d.pears = 'D' AND
    g.pears = 'D' AND
    g.pears = d.pears AND
    g.harvest_size = d.harvest_size AND
    (d.lemons = 0 OR d.lemons IS NULL) AND
    (g.lemons = 0 OR g.lemons IS NULL)
```


```
) Y
```

```
WHERE
```

```
X.oranges = Y.oranges AND
```

```
X.apples = Y.apples AND
```

```
X.harvest_size = Y.harvest_size;
```



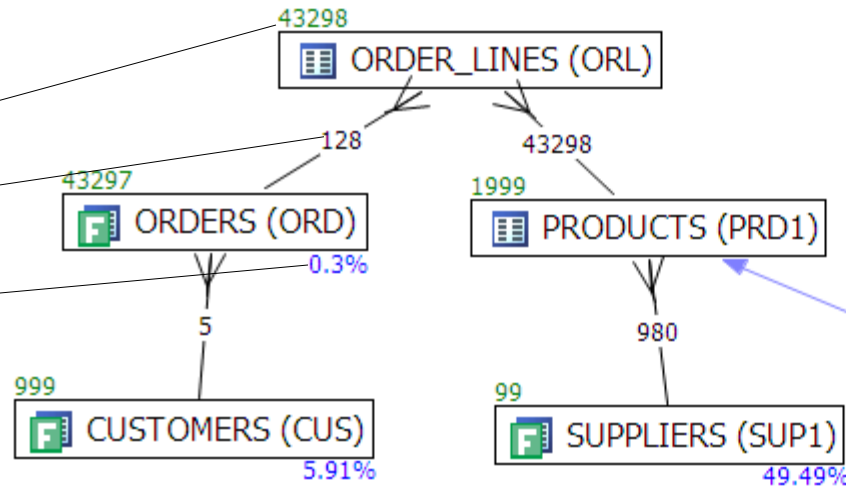
This final version runs in
elapsed **0.33 secs** and 12K logical reads
down from an original
elapsed **4.5 secs** and 1M logical reads

What tools to use to create VST diagrams?

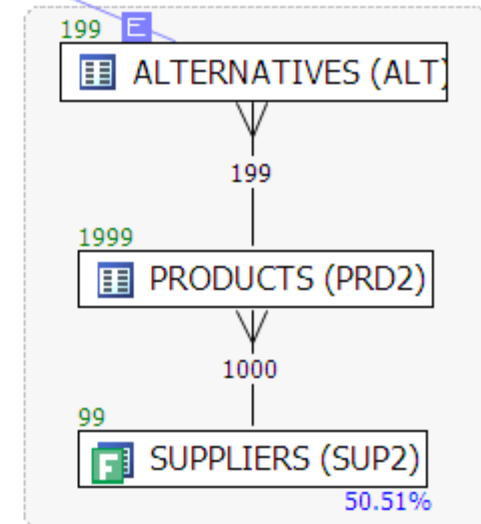
- Paper
 - Modifications messy and difficult
- PowerPoint
 - Can move things around
 - Sticky connectors
 - Easy to modify
- DB Optimizer
 - Automatic and still modifiable

Visual SQL Tuning (VST)

Table Sizes
Join Sizes
Filter Ratios



Exists



```
SELECT order_line_data
FROM      customers cus
  INNER JOIN orders ord ON ord.id_customer = cus.id
  INNER JOIN order_lines orl ON orl.id_order = ord.id
  INNER JOIN products prd1 ON prd1.id = orl.id_product
  INNER JOIN suppliers sup1 ON sup1.id = prd1.id_supplier
WHERE     cus.location = 'LONDON'
  AND ord.date_placed BETWEEN '04-JUN-10'
                        AND '11-JUN-10'
  AND sup1.location = 'LEEDS'
  AND EXISTS ( SELECT NULL
                FROM  alternatives      alt
                INNER JOIN      products prd2
                  ON prd2.id = alt.id_product_sub
                INNER JOIN      suppliers sup2
                  ON sup2.id = prd2.id_supplier
                WHERE           alt.id_product = prd1.id
                  AND sup2.location != 'LEEDS' )
```




VST Steps Summary

1. Diagram tables
2. Draw connectors for each join
3. Calculate filter ratios
4. Find the table sizes
5. Calculate two table join sizes

Execution Path

- Start at the most selective join filter
- Join to keep the running result set size small

Many to Many relationships = problems



Do You Want?

Engineering Data?

1161
05730, 1985

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1161

1

Do You Want?

Pretty Pictures

11.6.1
Oct 30, 1981
AFT

61A LH Center Field**
61A LH CENTER FIELD**
51C LH Forward Field**
51C RH Center Field (prim)***
51C RH Center Field (sec)***

41D RH Forward Field
41C LH Aft Field*
41B LH Forward Field


STS-2 RH Aft Field

*Hot gas path detected in putt
**Soot behind primary O-ring.
***Soot behind primary O-ring, h

Clocking location of leak che


OTHER SRM-15 FIELD JOINT
NEAR OR BEYOND THE PRIMA

SRM-22 FORWARD FIELD JO
AND NO SOOT BLOWBY, OTHER SRM-22 FIELD JOINTS HAD NO BLOWHOLES IN PUTTY.



Do You Want?

Clean and Clear

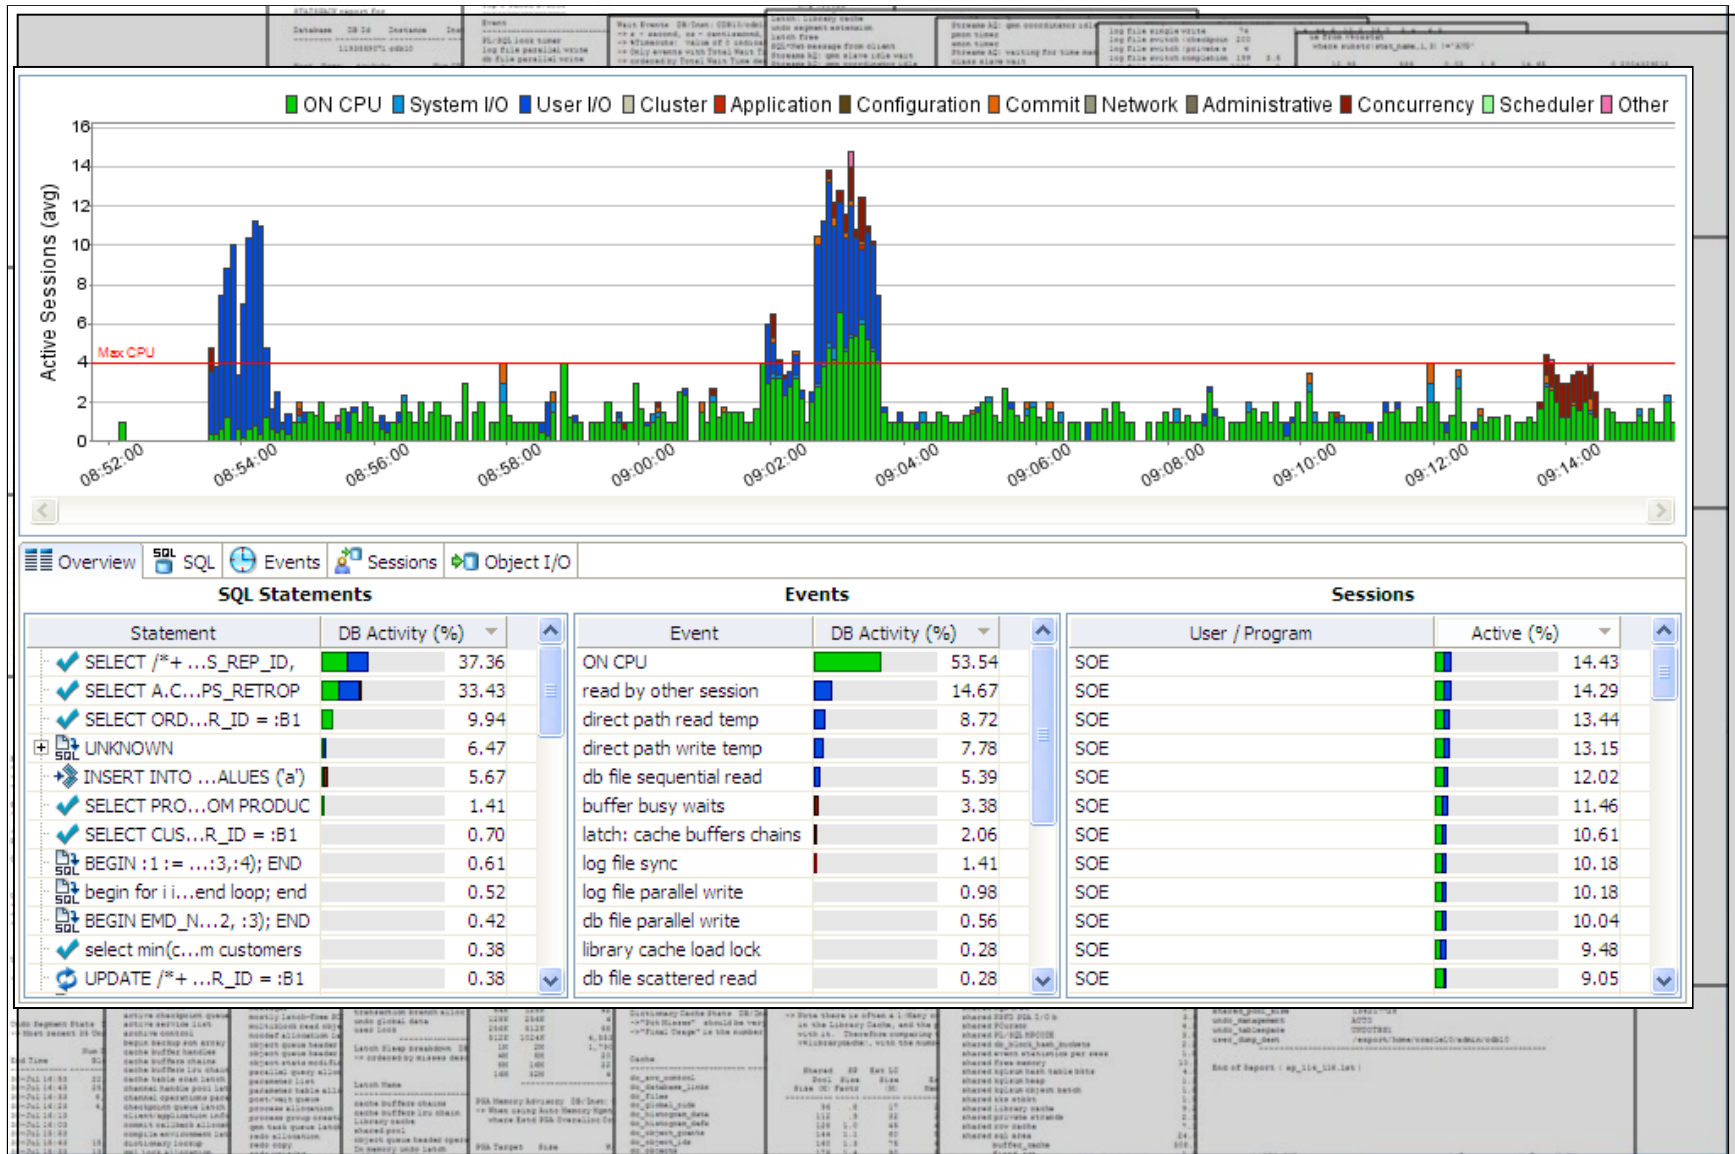


Would you want your dashboard to look like :

And is updated once an hour

Or would you like it to look ...

Or This



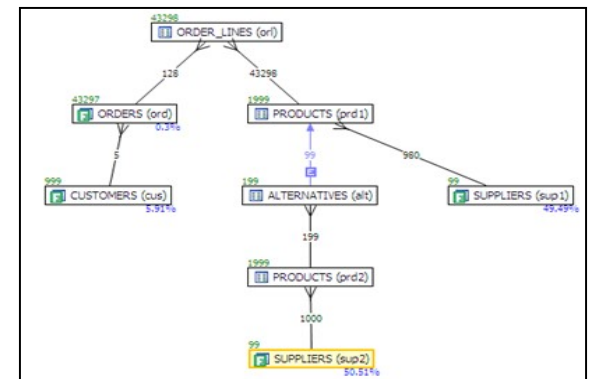
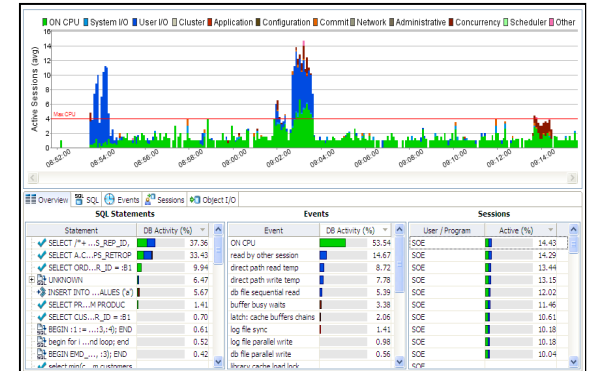
Summary

1.Database - AAS





- Profile database
- Use wait interface and graphics
- Identify machine, application, database or SQL

1.SQL - VST

- Indexes, stats, execution path
- Visual SQL Tuning



Bibliography



Refactoring SQL Applications – Stephane Faroult

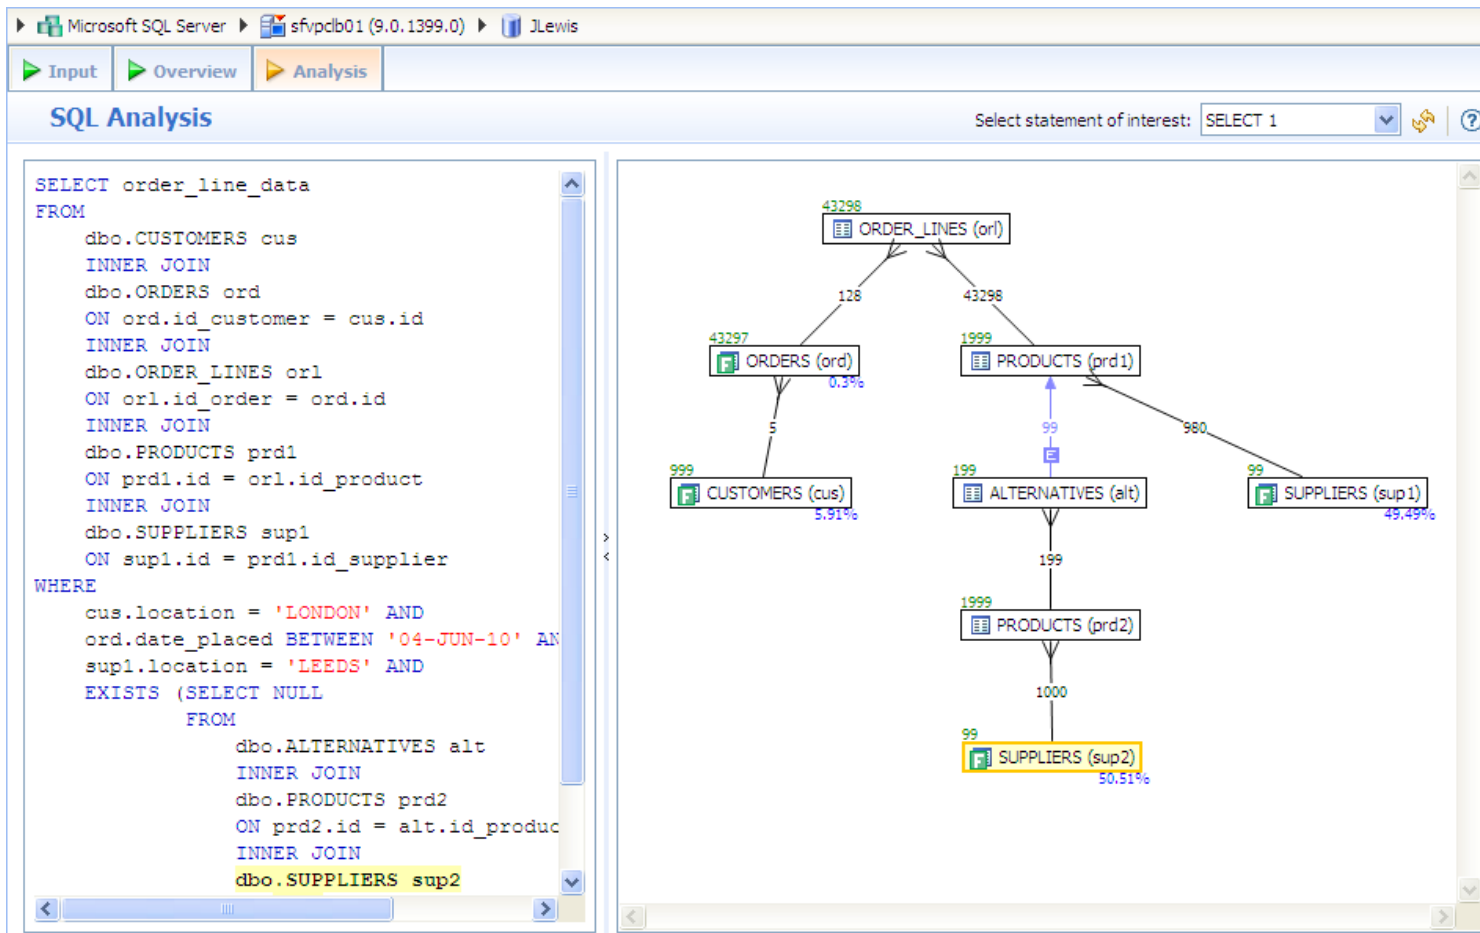
Troubleshooting Oracle Performance – Christian Antognini

SQL Tuning – Dan Tow

Cost-Based Oracle Fundamentals – Jonathan Lewis



<http://www.simple-talk.com/sql/performance/designing-efficient-sql-a-visual-approach/>

Visual SQL Tuning (VST)

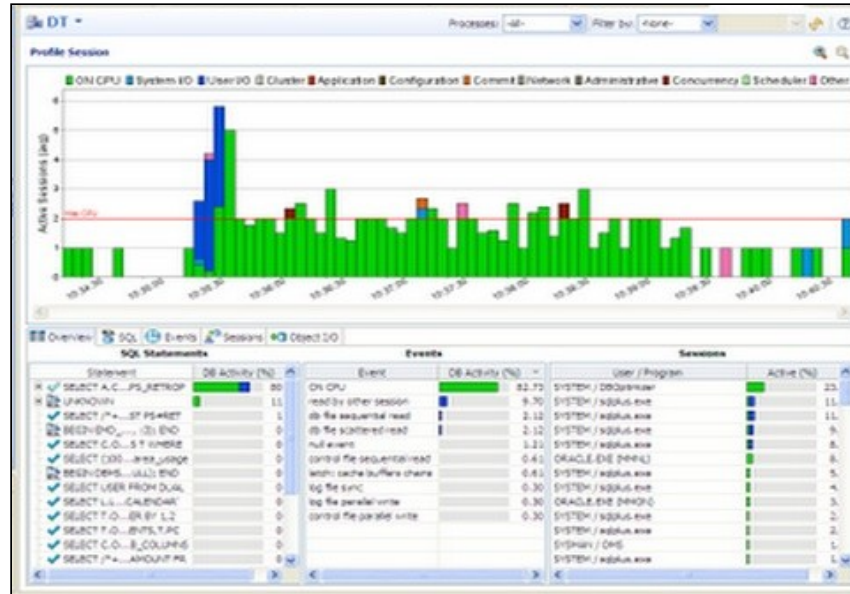


Industry Exclusive



DB Optimizer XE Key Features



PROFILE



DB Optimizer XE Key Features



PROFILE



TUNE

SQL Analysis

Select statement of interest: SELECT 4

```
SELECT
  ct.action,
  c.client_id,
  i.investment_unit,
  it.investment_type_name
FROM
  client_transaction ct,
  client c,
  investment_type it,
  investment i
WHERE
  ct.client_id = c.client_id AND
  ct.investment_id = i.investment_id AND
  i.investment_type_id = it.investment_type_id AND
  ct.broker_commission = 100
```



Query Plan Diagram:

```
graph TD
    CT[CLIENT_TRANSACTION (ct)] --> C[CLIENT (c)]
    CT --> I[INVESTMENT (i)]
    I --> IT[INVESTMENT_TYPE (it)]
```

Table Statistics:

Table Name	Table Owner	Table Name	Column Name	Index Type
INDEX_CLIENT_TRANSACTION_0	SYSTEM	CLIENT_TRANSACTION	BROKER_COMMISSION	Normal
CLIENT_PK	SYSTEM	CLIENT	CLIENT_ID	Unique
INVESTMENT_PK	SYSTEM	INVESTMENT	INVESTMENT_ID	Unique
INVESTMENT_TYPE_PK	SYSTEM	INVESTMENT_TYPE	INVESTMENT_TYPE_ID	Unique
CLIENT_TRANSACTION_CLIENT	SYSTEM	CLIENT_TRANSACTION	CLIENT_ID	Normal
CLIENT_TRANSACTION_INVESTMENT	SYSTEM	CLIENT_TRANSACTION	INVESTMENT_ID	Normal
INVESTMENT_INVESTMENT_TYPE	SYSTEM	INVESTMENT	INVESTMENT_TYPE_ID	Normal
CLIENT_BROKER	SYSTEM	CLIENT	BROKER_ID	Normal
CLIENT_INCOME	SYSTEM	CLIENT	CLIENT_HOUSEHOLD_INCOME	Normal
CLIENT_MULTI	SYSTEM	CLIENT	CLIENT_FIRST_NAME, LAST_NAME	Normal

DB Optimizer XE Key Features



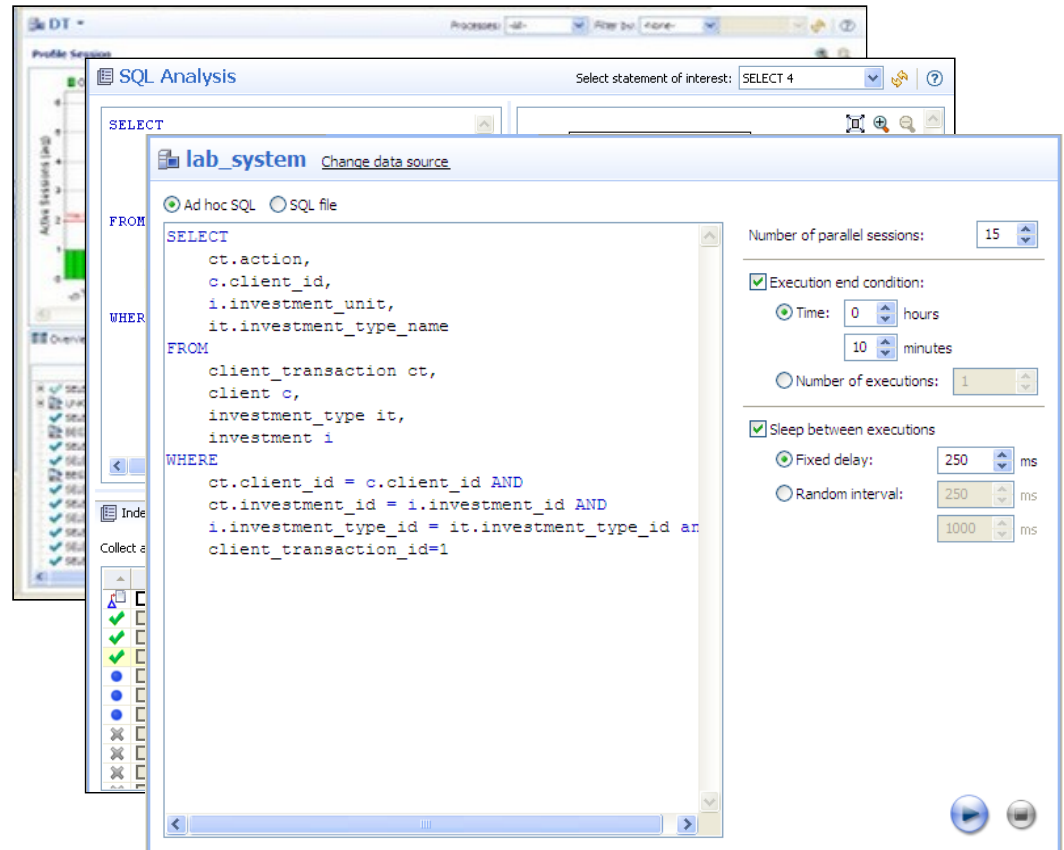
PROFILE





TUNE



LOAD TEST



DB Optimizer XE Key Features



PROFILE



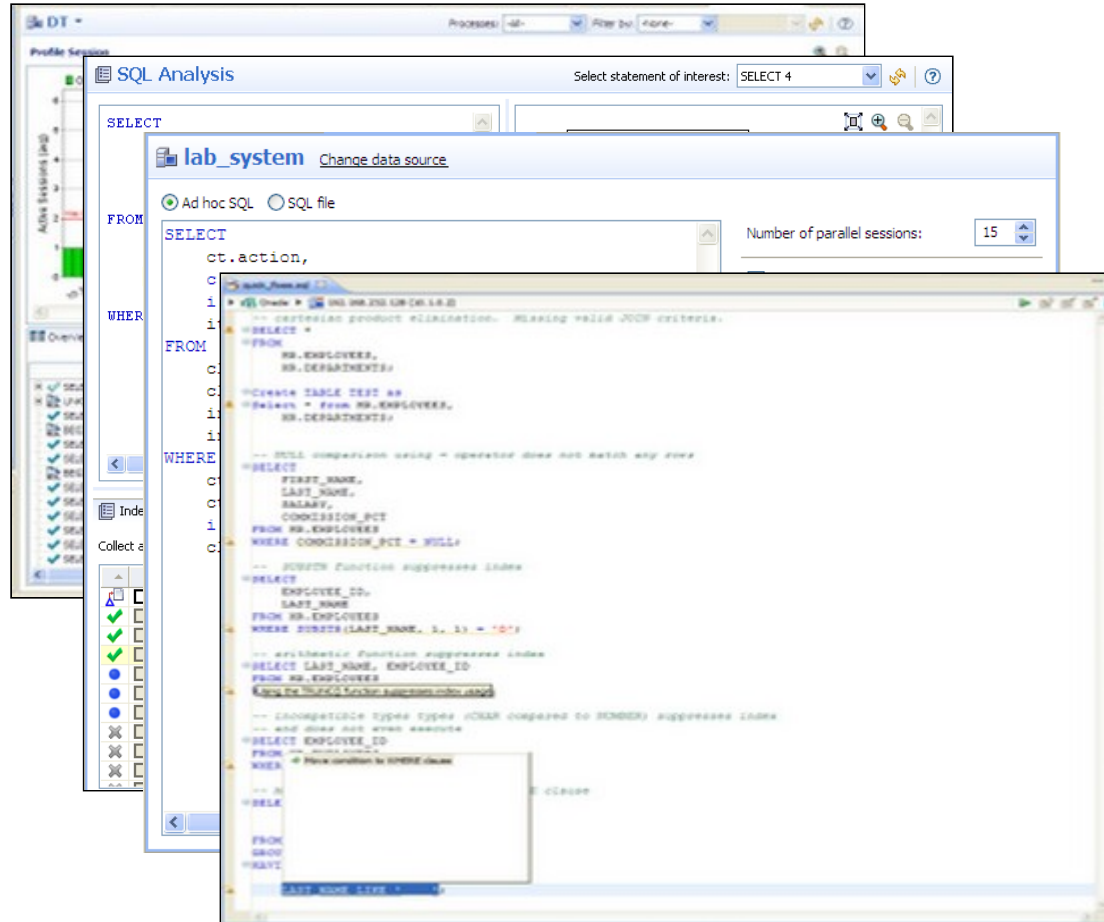
TUNE





LOAD TEST



SQL IDE



Review of Functionality



PROFILE



TUNE

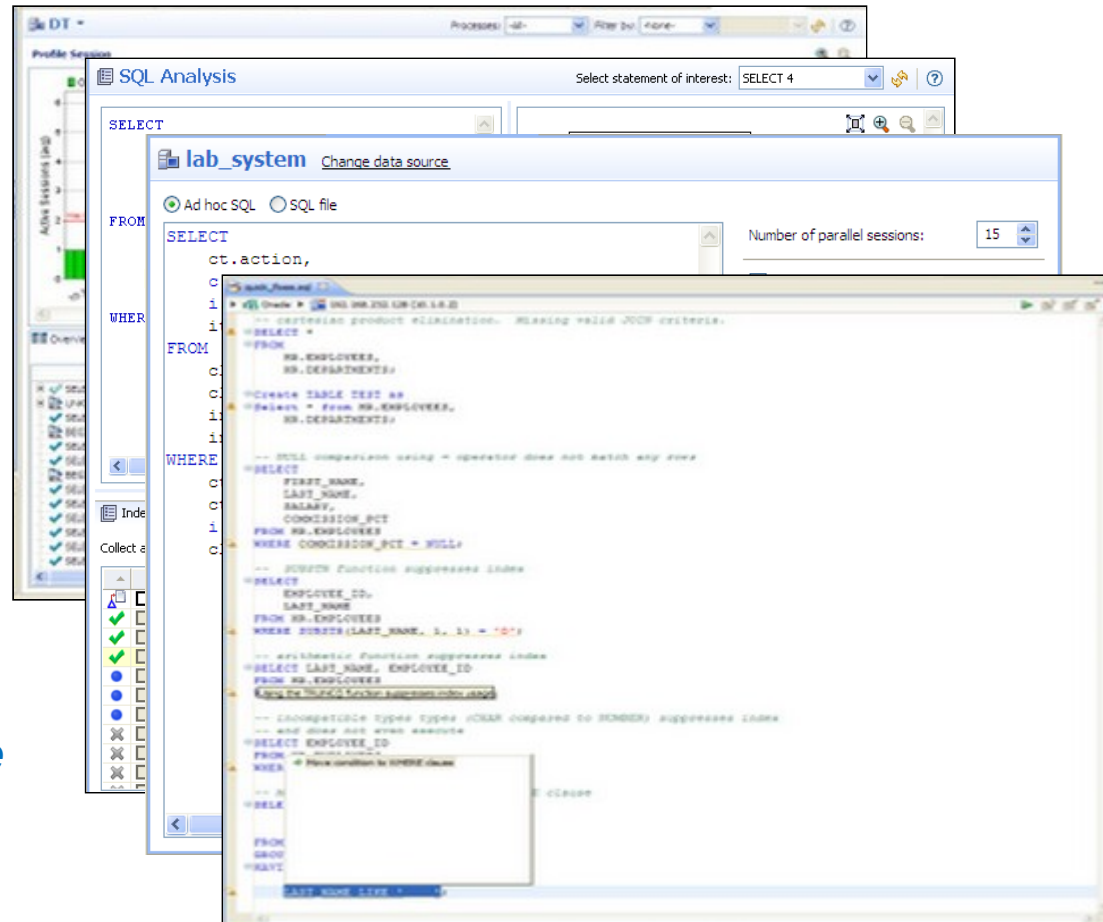


LOAD TEST



SQL IDE

Only tool on the market with these features integrated



Thank You

Kyle Hailey

<http://oraclemonitor.com>

Appendix

- Diagramming
 - Simple joins and inline views
 - Outer joins
 - Exists/not exists (in/not in)
 - Correlated aggregate sub-queries
- Hints
 - LEADING
 - USE_NL
 - USE_HASH
 - INDEX
 - NO_MERGE
- Execution Order
- Machine Health
- VST method with Statistics


Basic Joins



Many to
single value

One to one

One to many



Safe

Dangerous

Many to many

Cartesian



Simple queries and sub-queries

Simple join


```
select * from A, B where b.f1=a.f1
```

Non-correlated sub-query

```
select * from A, (select b.f1 from B) c  
where c.f1=a.f1
```



Unique index on b.f1




No unique indexes



Unique index on
Both b.f1 and a.f1

Special case: non correlated sub-query returns one row



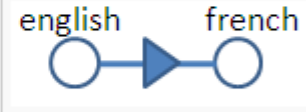
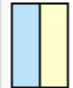
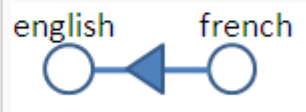

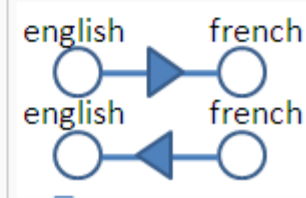

```
select * from A where a.f1 =  
      (select max(b.f1) from B )
```



Outer Joins

ENGLISH_TEXT	ORDINAL_ID	ORDINAL_ID	FRENCH_TEXT
One	1	1	Un
Two	2	3	Trois
Three	3	4	Quatre
Four	4	5	Cinq
Five	5	6	Six
Six	6	7	Sept
		8	Huit


Outer Joins

type	ANSI	ANSI 89 (Oracle)	type	type
inner join	english INNER JOIN french using (ordinal_id)	english e, french f where e.ordinal_id=f.ordinal_id		
left outer join	english LEFT JOIN french using (ordinal_id)	english e, french f where e.ordinal_id=f.ordinal_id(+)		
right outer join	english RIGHT JOIN french using (ordinal_id)	english e, french f where e.ordinal_id(+)=f.ordinal_id		
full join	english FULL JOIN french using (ordinal_id)	english e, french f where e.ordinal_id=f.ordinal_id(+) UNION english e, french f where e.ordinal_id(+)=f.ordinal_id		

Correlated scalar sub-queries

Correlated aggregate subquery:

```
select ename from emp a where a.sal >
      (select avg(sal)
       from emp b
       where a.deptno=b.deptno)
```



Scalar Subqueries


```
select ename, (select avg(sal)
               from emp b
               where a.deptno=b.deptno)
from emp a;
```

Exists and Not In

```
SELECT d.*  
FROM dept d WHERE exists (  
  SELECT null FROM emp e WHERE e.deptno=d.deptno);
```

```
SELECT d.*  
FROM dept d WHERE d.deptno in (  
  SELECT deptno FROM emp e );
```


```
select distinct d.* from dept1 d ,emp e  
where e.deptno = d.deptno;
```



```
SELECT d.*  
FROM dept d WHERE not exists (  
  SELECT null FROM emp e WHERE e.deptno=d.deptno);
```

```
SELECT d.*  
FROM dept d WHERE d.deptno not in (  
  SELECT deptno FROM emp e where e.deptno is not null )  
or d.deptno is null;
```

```
select d.* from dept1 d left outer join emp e  
on e.deptno = d.deptno where e.deptno is null;
```



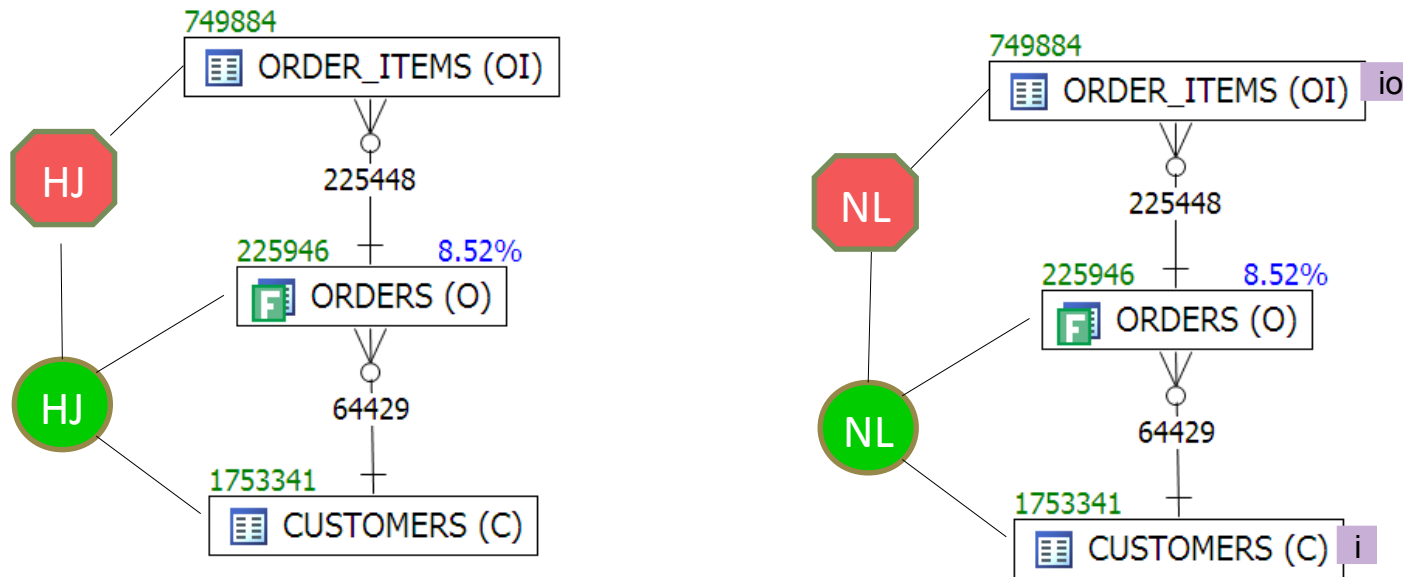
HINTS

- ORDERED - good on 9i
- Leading(tab_alias , table_alias ...) – 10g format
- USE_NL(table_alias) – Inner Table (not driving)
- USE_HASH(table_alias) – 2cd table, probe into
- INDEX(tab_alias index_name)
- NO_MERGE

Oracle first decides join order then join type

(example <http://www.adp-gmbh.ch/blog/2008/01/17.php>)

Visual SQL Tuning (VST) diagrams



Makes comparison of execution plans easy
Better yet, it will help us find the best execution path

Operation	Operation
SELECT STATEMENT	SELECT STATEMENT
TABLE ACCESS - SOE.ORDER_ITEMS	HASH JOIN
NESTED LOOPS	HASH JOIN
NESTED LOOPS	TABLE ACCESS - SOE.CUSTOMERS
PARTITION HASH	INDEX - SOE.CUSTOMERS_PK
TABLE ACCESS - SOE.ORDERS	PARTITION HASH
TABLE ACCESS - SOE.CUSTOMERS	TABLE ACCESS - SOE.ORDERS
INDEX - SOE.CUSTOMERS_PK	PARTITION HASH
INDEX - SOE.ORDER_ITEMS_PK	TABLE ACCESS - SOE.ORDER_ITEMS

VST vs Explain Plan

SELECT

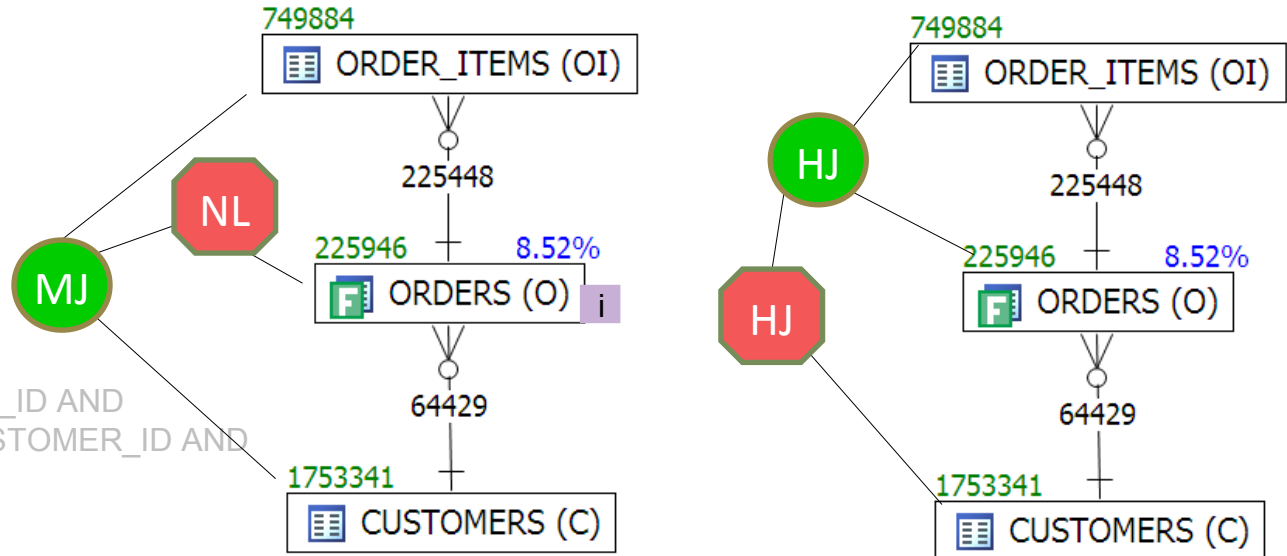
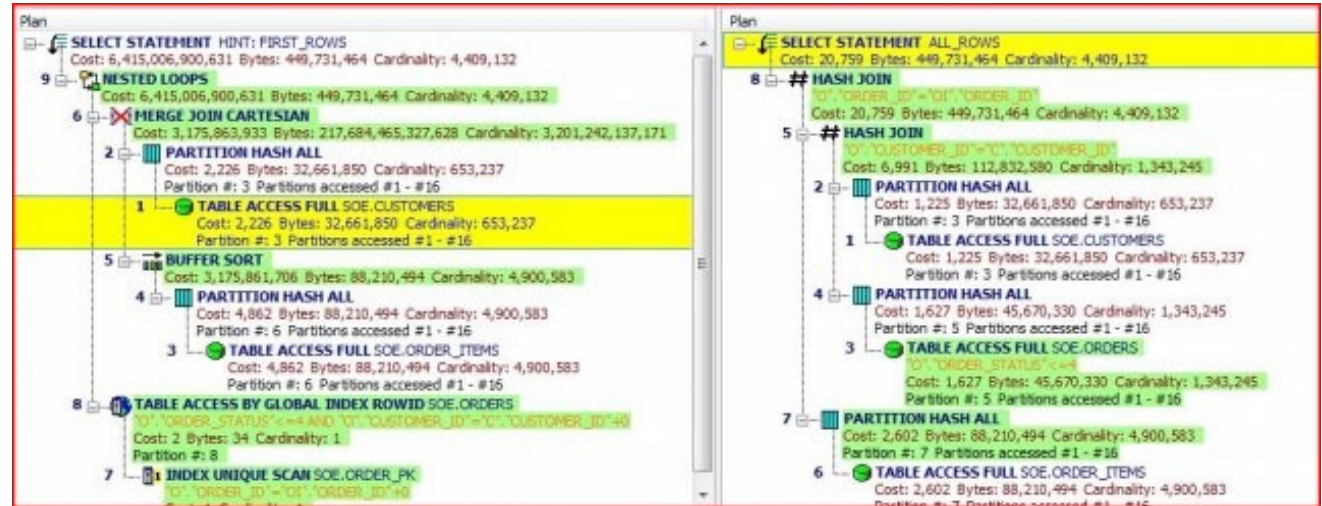
O.ORDER_ID,
LINE_ITEM_ID,
PRODUCT_ID,
UNIT_PRICE,
QUANTITY,
ORDER_MODE,
ORDER_STATUS,
ORDER_TOTAL,
SALES_REP_ID,
PROMOTION_ID,
C.CUSTOMER_ID,
CUST_FIRST_NAME,
CUST_LAST_NAME,
CREDIT_LIMIT,
CUST_EMAIL,
ORDER_DATE

FROM

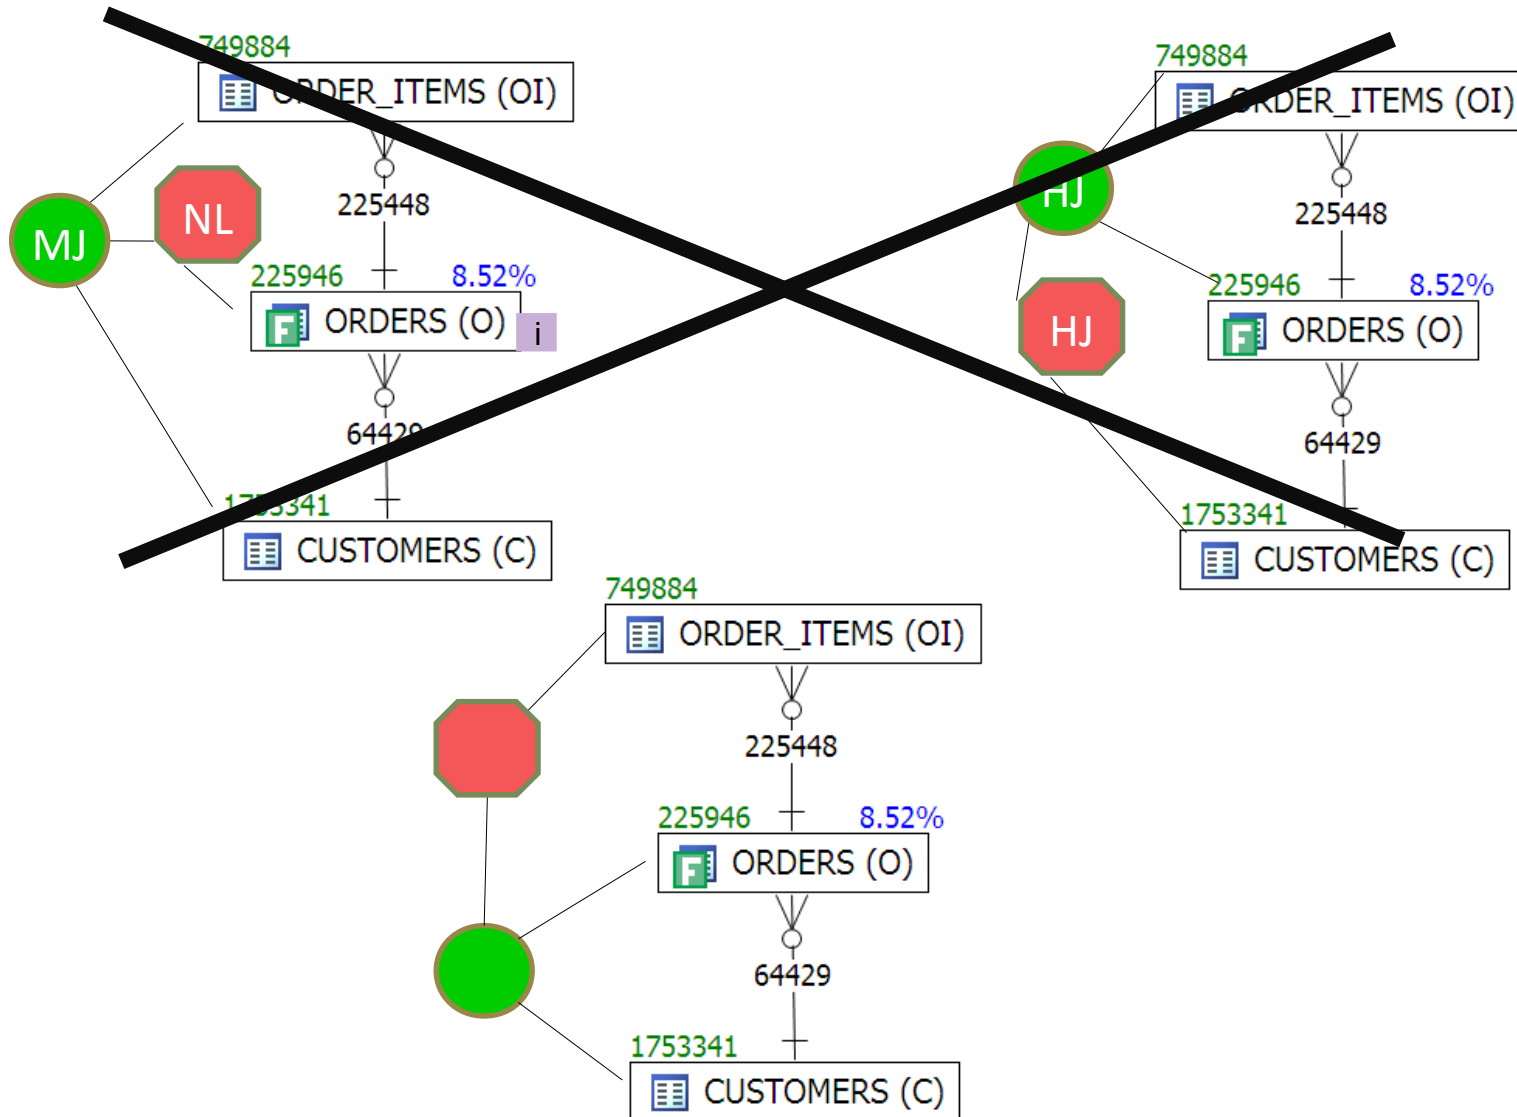
ORDERS O,
ORDER_ITEMS OI,
CUSTOMERS C

WHERE

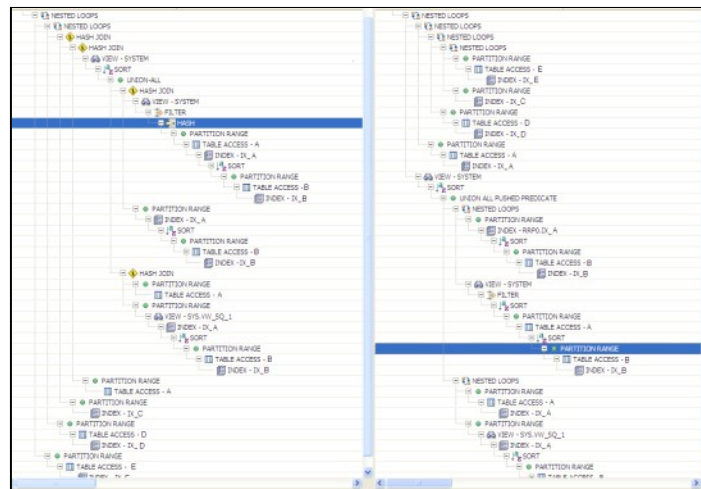
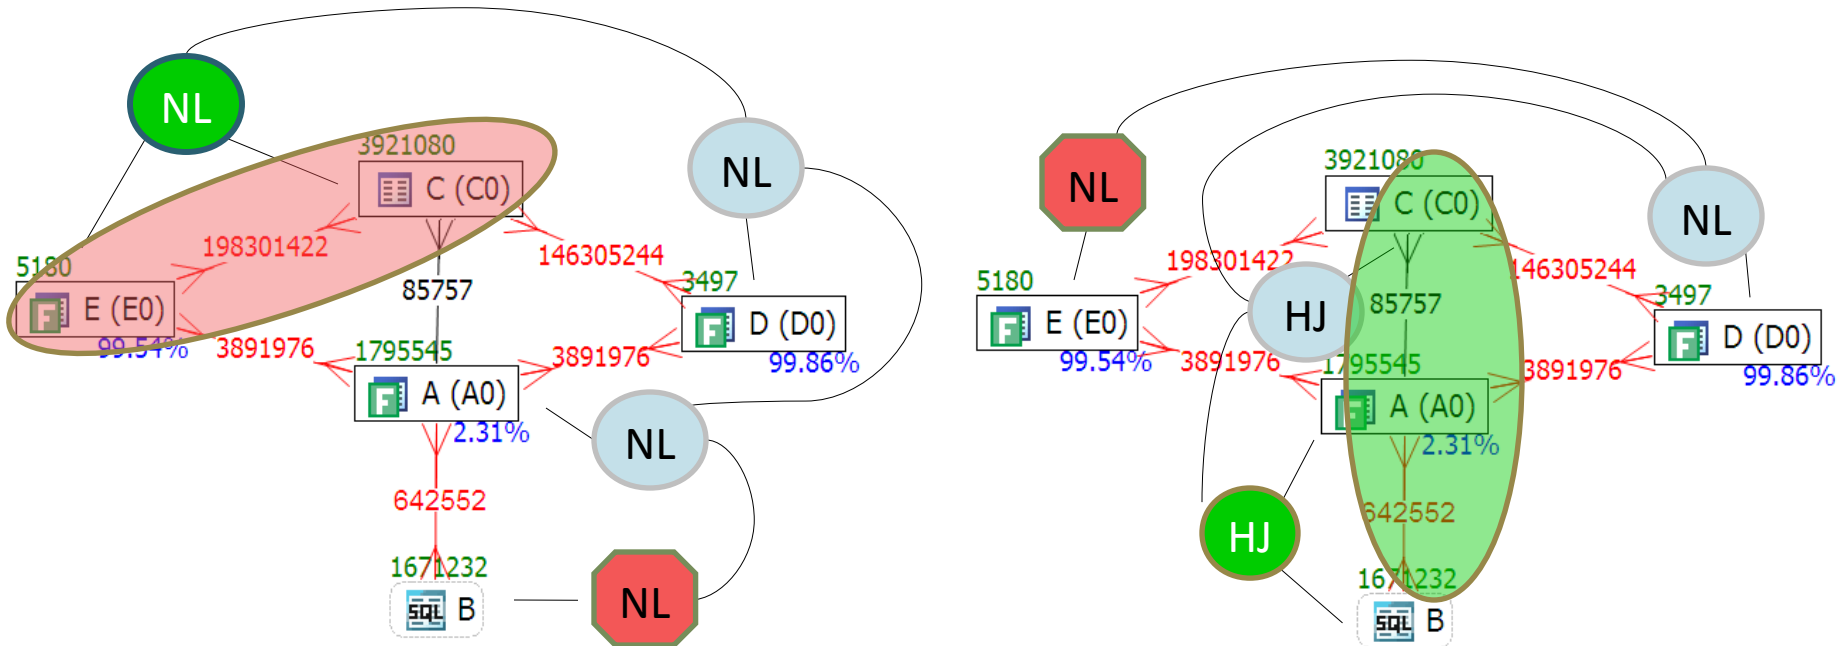
O.ORDER_ID = OI.ORDER_ID AND
O.CUSTOMER_ID = C.CUSTOMER_ID AND
O.ORDER_STATUS <= 4



Bad Plans vs Good Plan



Comparing Plans : 24 hours to 5 mins



When to Tune


1. Machine

a) CPU

- Response times skewed
- 100% CPU might be fine
- Users wait in queue (run queue) => machine under

a) Memory

- Paging
- Wait times skewed (ex : latch free)
- Erratic response times (ex : ls)



1. Oracle

1) Waits > CPU ?

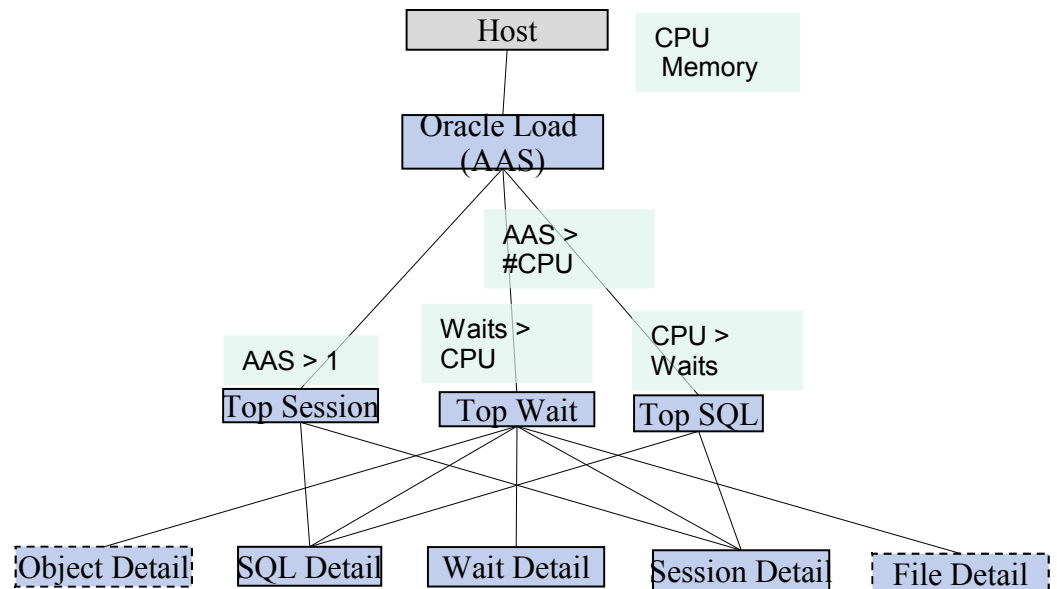
- tune waits

1) CPU > 100% ?

- tune top CPU SQL

1) Else

- It's the application



Machine

Make sure the machine is healthy before tuning Oracle

- CPU => use run queue, $< 2 * \text{\#CPU}$
- Memory => page out

VMSTAT

```

# vmstat 1

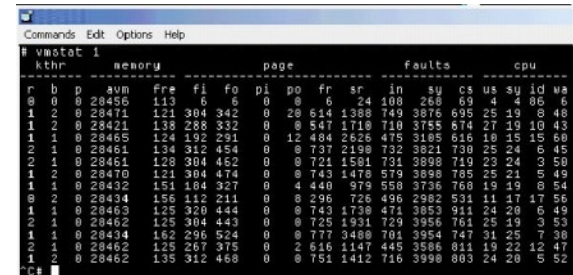
```

kthr			memory			page			faults				cpu				
r	b	p	avm	fre	fi	fo	pi	po	fr	sr	in	sy	cs	us	sy	id	wa
0	0	0	28456	113	6	6	0	0	6	24	108	268	69	4	4	86	6
1	2	0	28471	121	304	342	0	20	614	1388	749	3876	695	25	19	8	48
1	2	0	28421	138	288	332	0	0	547	1710	710	3755	674	27	19	10	43
1	1	0	28465	124	192	291	0	12	484	2626	475	3105	616	10	15	15	60
2	1	0	28461	134	312	454	0	0	737	2190	732	3821	730	25	24	6	45
2	1	0	28461	128	304	462	0	0	721	1501	731	3898	719	23	24	3	50
1	2	0	28470	121	304	474	0	0	743	1478	579	3898	785	25	21	5	49
1	1	0	28432	151	184	327	0	4	440	979	558	3736	768	19	19	8	54
0	2	0	28434	156	112	211	0	8	296	726	496	2982	531	11	17	17	56
1	1	0	28463	125	320	444	0	0	743	1730	471	3853	911	24	20	6	49
2	1	0	28462	125	304	443	0	0	725	1931	729	3956	761	25	19	3	53
1	1	0	28434	162	296	524	0	0	777	3480	701	3954	747	31	25	7	38
2	1	0	28462	125	267	375	0	2	616	1147	445	3586	811	19	22	12	47
1	2	0	28462	135	312	468	0	0	751	1412	716	3990	803	24	20	5	52

Summary

1. Machine - vmstat

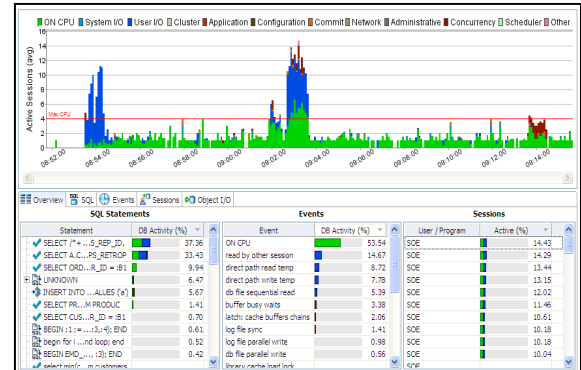
- Memory, CPU (we can see IO response in Oracle)



procs		memory		page		faults		cpu	
in	ou	fr	fi	po	fr	in	ou	us	sy
0	0	113	6	0	0	24	100	268	59
1	2	121	304	342	0	28	514	1388	749
1	2	138	288	332	0	547	1710	110	3755
1	1	124	192	291	0	12	484	2628	475
1	0	134	312	454	0	737	2190	732	3921
1	2	128	304	462	0	721	1501	731	3898
1	2	121	304	474	0	743	1478	579	3898
1	1	151	104	327	0	4	440	979	550
0	2	156	112	211	0	296	726	496	2982
1	1	125	320	444	0	743	1738	471	3853
2	1	125	304	443	0	725	1931	729	3956
1	1	162	296	524	0	777	3489	701	3954
1	1	125	267	375	0	2	616	1147	448
1	2	135	312	468	0	751	1412	710	3998

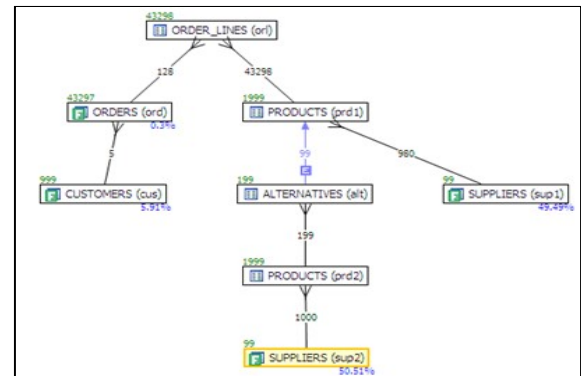
1. Database - AAS

- Use wait interface and graphics
- Identify machine, application, database or SQL



1. SQL - VST

- Indexes, stats, execution path
- Visual SQL Tuning

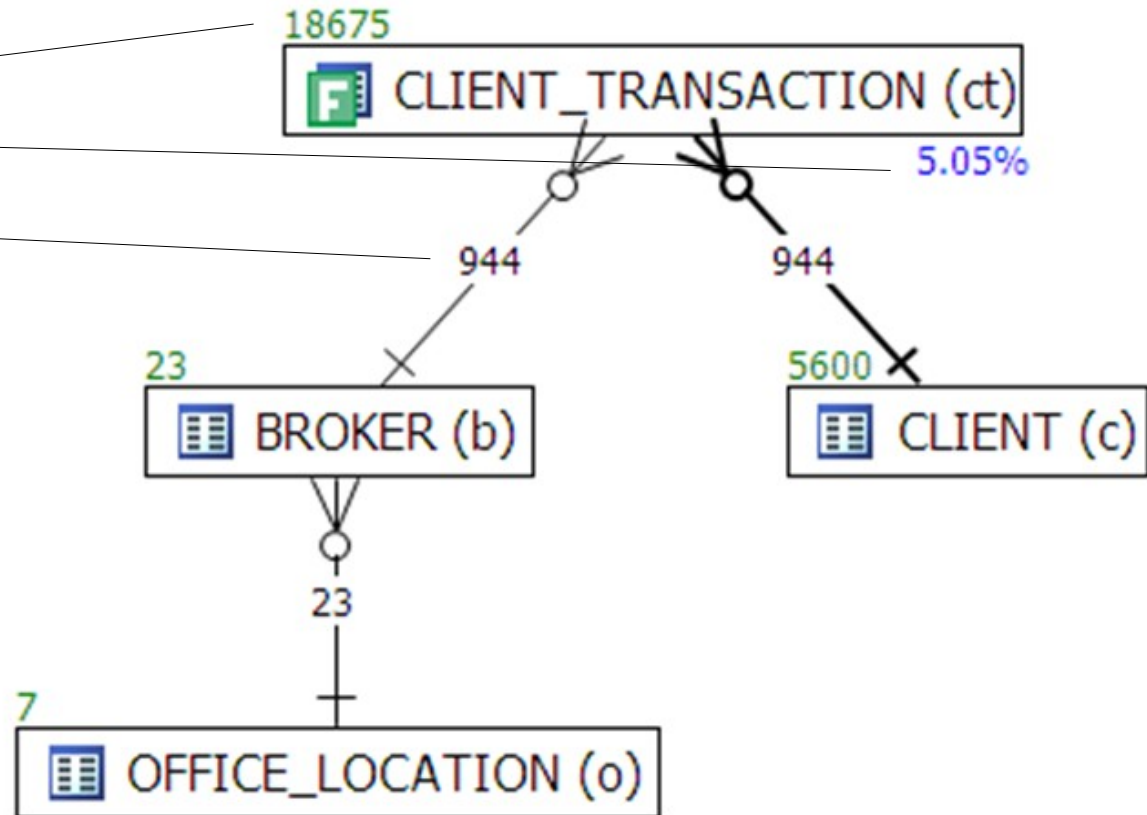


Solution to Many to Many

Table Sizes

Filter Ratios

Join Sizes



VST Steps

Objects:

1. Tables : drawn as nodes
2. Joins : drawn as connector lines
3. Filters : mark on each table with filter in where clause

Statistics:

1. Table sizes
2. Join sizes
3. Calculate filter percentages
 - $\text{filter ratio} = \text{number of rows returned with filter} / \text{number of rows}$

Google Directions

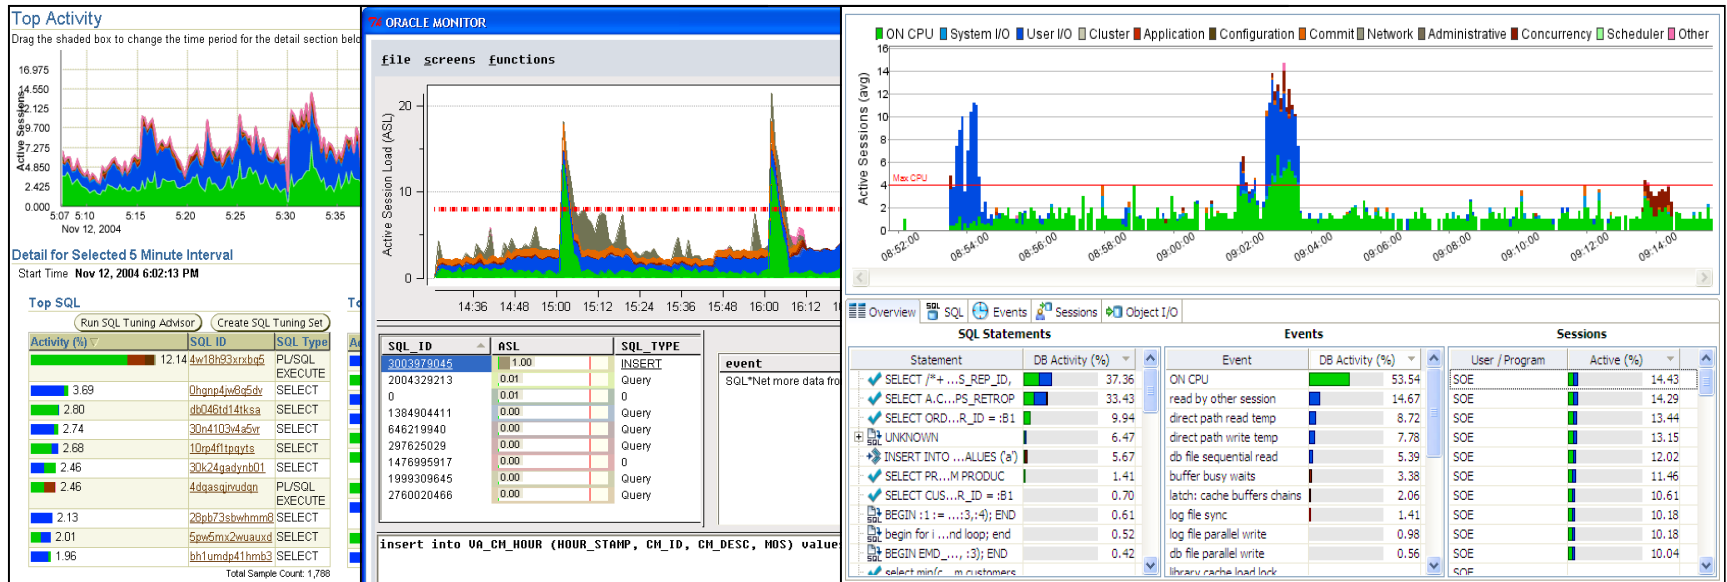
- Compare two sets of directions
- (optionally show graphical set of directions)
- Show map
- Show map with traffic flow, red, yellow, green

How Can We Open the Black Box?

OEM


ASHMON/SASH

DB Optimizer



- **Powerful** - Identifies issues quickly and powerfully
- **Interactive** - Allows exploring the data
- **Easy** - Understandable by everyone, DBA, Dev and Managers !

Clearer



1. Include successes
2. Remove Irrelevant
3. Normalize same temp
4. Scale known vs unknown