Method – GAPP

Predicting and Profiling of End-User Performance while Focussing on AWR Data

a Method to Diagnose and Predict Performance in Complex Architectures

Gerwin Hendriksen
IT Architect
gerwin.hendriksen@gmail.com

http://method-gapp.com

V1.2

Who I am…

- Started as DBA at Dedicate (later called Parity) for 1.5 years
- Worked for 5 years at Oracle The Netherlands
  - First 1.5 year as RDBMS analyst
  - Rest worked as ARE EBS Performance and RAC
    - This included Content Lead of EBS with RAC world wide
- Worked for 2 years at IBM
  - Business Consultancy Services
- Worked for 3 years at AMIS Services BV
- Currently working for IT-eye as IT Architect

Agenda…

- What is Method-GAPP
- When to use Method-GAPP
- End-User processes and infrastructure
- The Five D’s from Method-GAPP
- AWR data usage
- Conclusions
- Q/A

What is Method-GAPP

What is Method-GAPP…

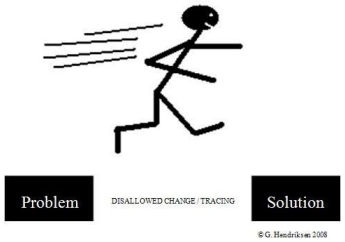

Method-GAPP is a performance method which makes smart use of underlying queuing models and data mining to find bottlenecks in complex architectures, for specific end-user processes within an enterprise.

"GAPP" is an abbreviation of:

"**General Approach Performance Profiling**"

The Five D's from Method-GAPP…

- Data Collection

- Data Synchronization

- Data Modeling

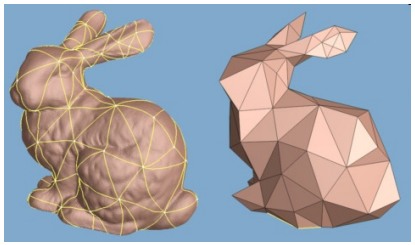

Mixed a bit !!

- Data Mining

- Data Interpretation

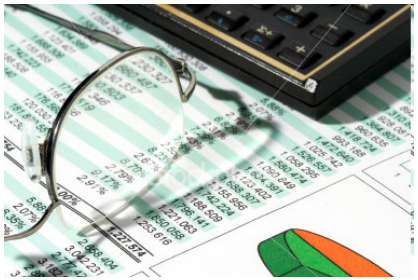

When to use Method-GAPP

When…

- End user processes are sometimes too slow and the cause hard to be determined.
- The technical infrastructure has shared components like a SAN, Server, etc.
- The application is hard to be debugged, because it is a third party application.

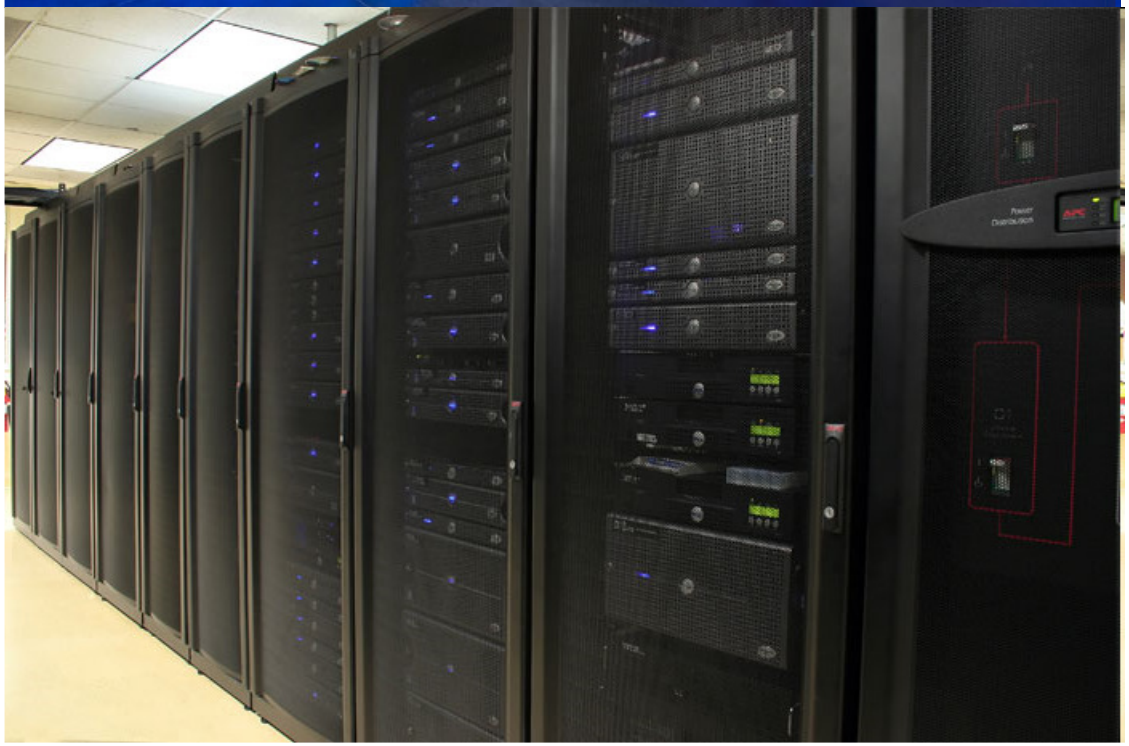

- Virtualization
- Shared Components
- Many applications, sharing the same technical infra structure
  - 10's - 1000's
- Performance problems occur in production can't be all reproduced in test due to unknown inter-application influences

It would be great if…

- You could find the cause of such problems without
  - Changing code (hooks)
  - Technology dependency
- You could even find out the influence of other applications on our end user processes.
- You could predict what the performance impact would be before doing investments.

End-User processes and infrastructure

Method-R and the sequence diagram…

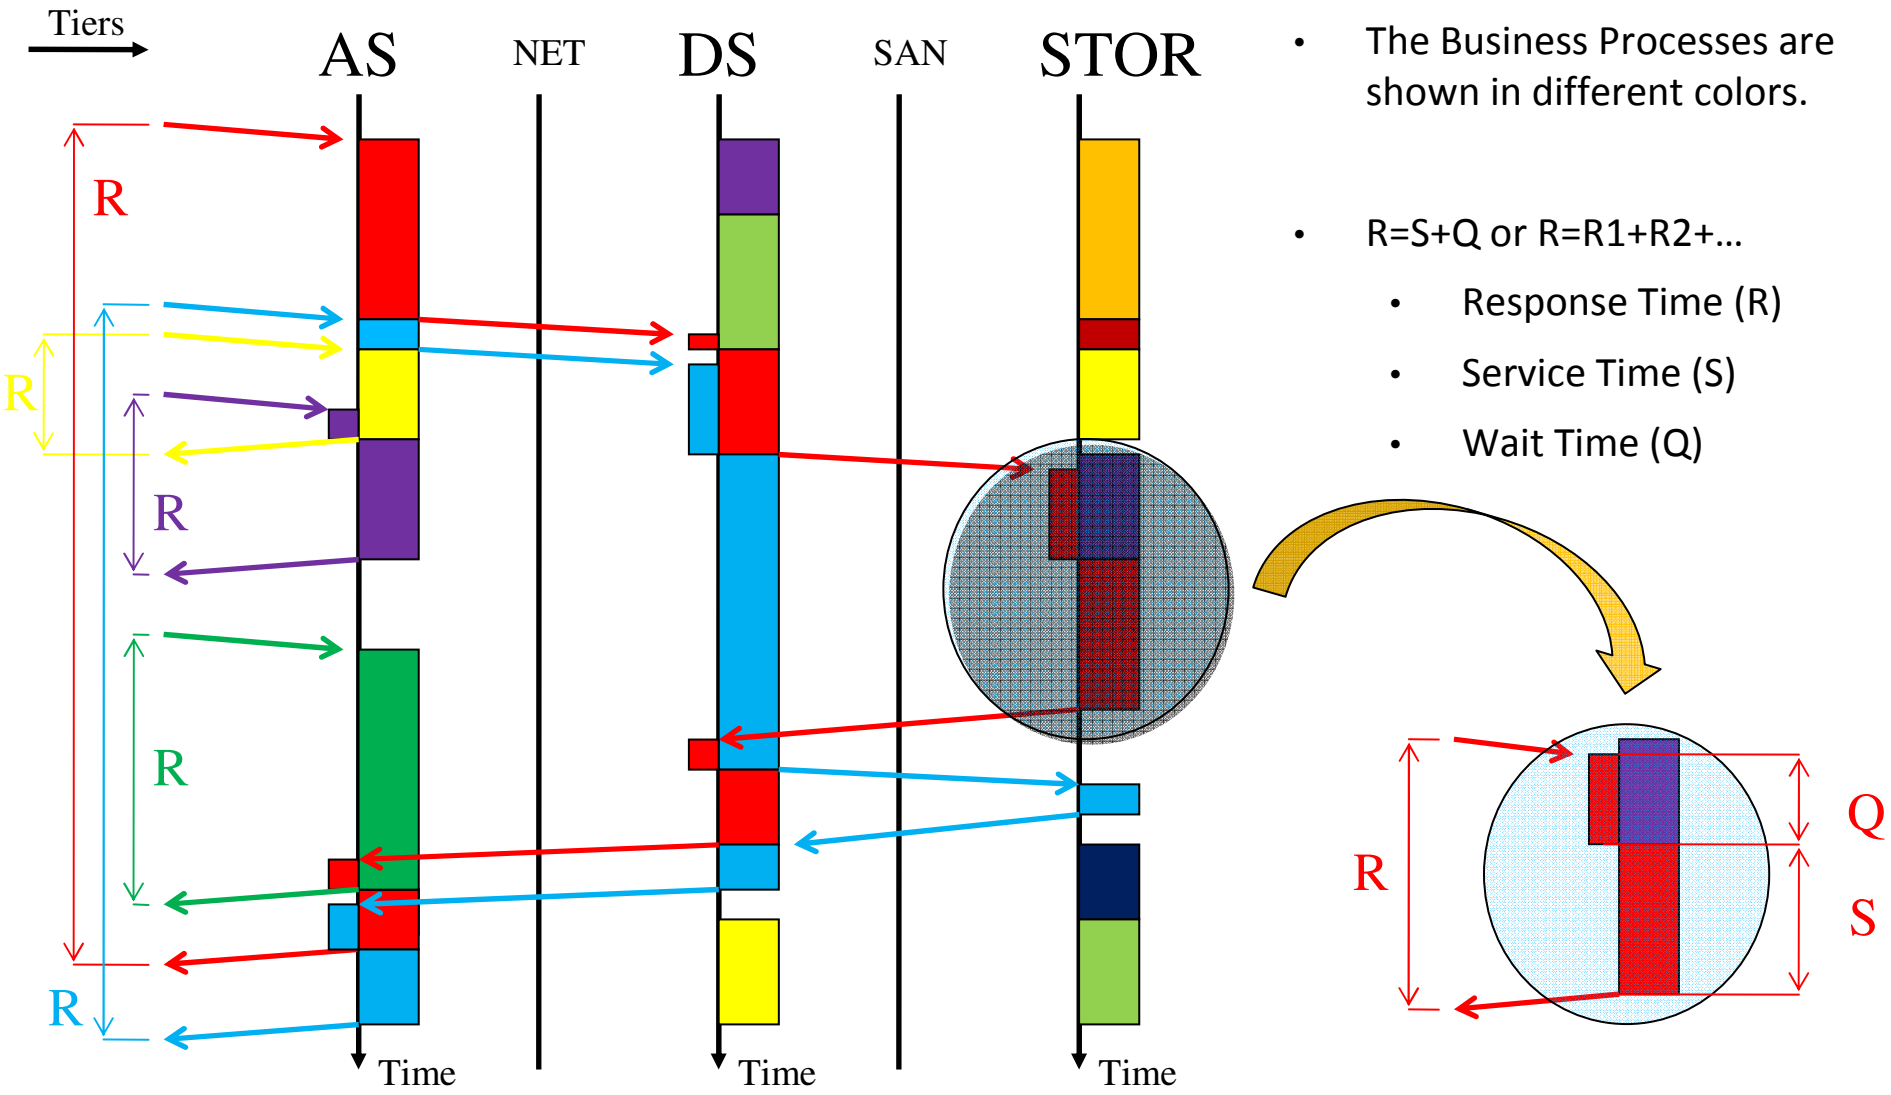

- The Business Processes are shown in different colors.
- R=S+Q or R=R1+R2+…
  - Response Time (R)
  - Service Time (S)
  - Wait Time (Q)

Highlighted the “Red” process…

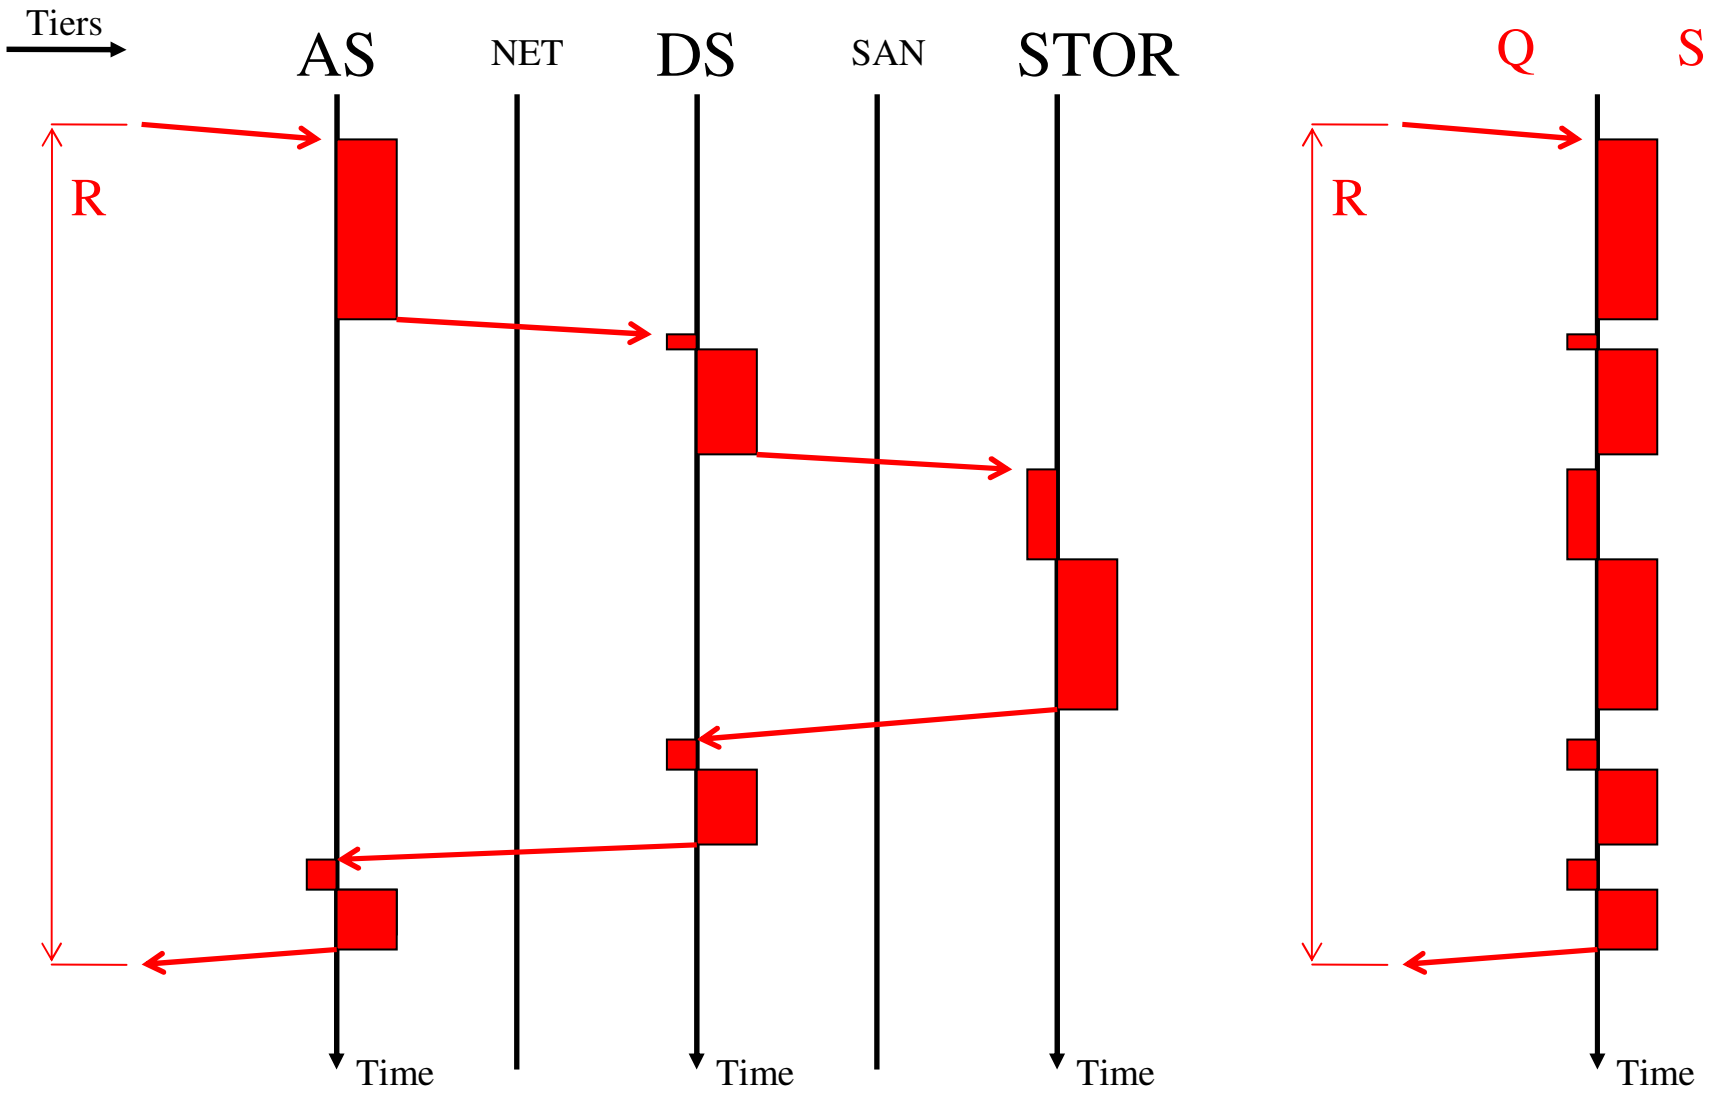

The First Two D's

Data Collection & Data Synchronization

Primary and Secondary Components…

- Primary components are (System resources):
  - CPU
  - I/O (also network)
  - Memory
- Secondary components are (rest) e.g.:
  - SQL response time
  - Java threads
  - Virtual Machine CPU data
  - Etc.

Data Collection and Data Synchronization…

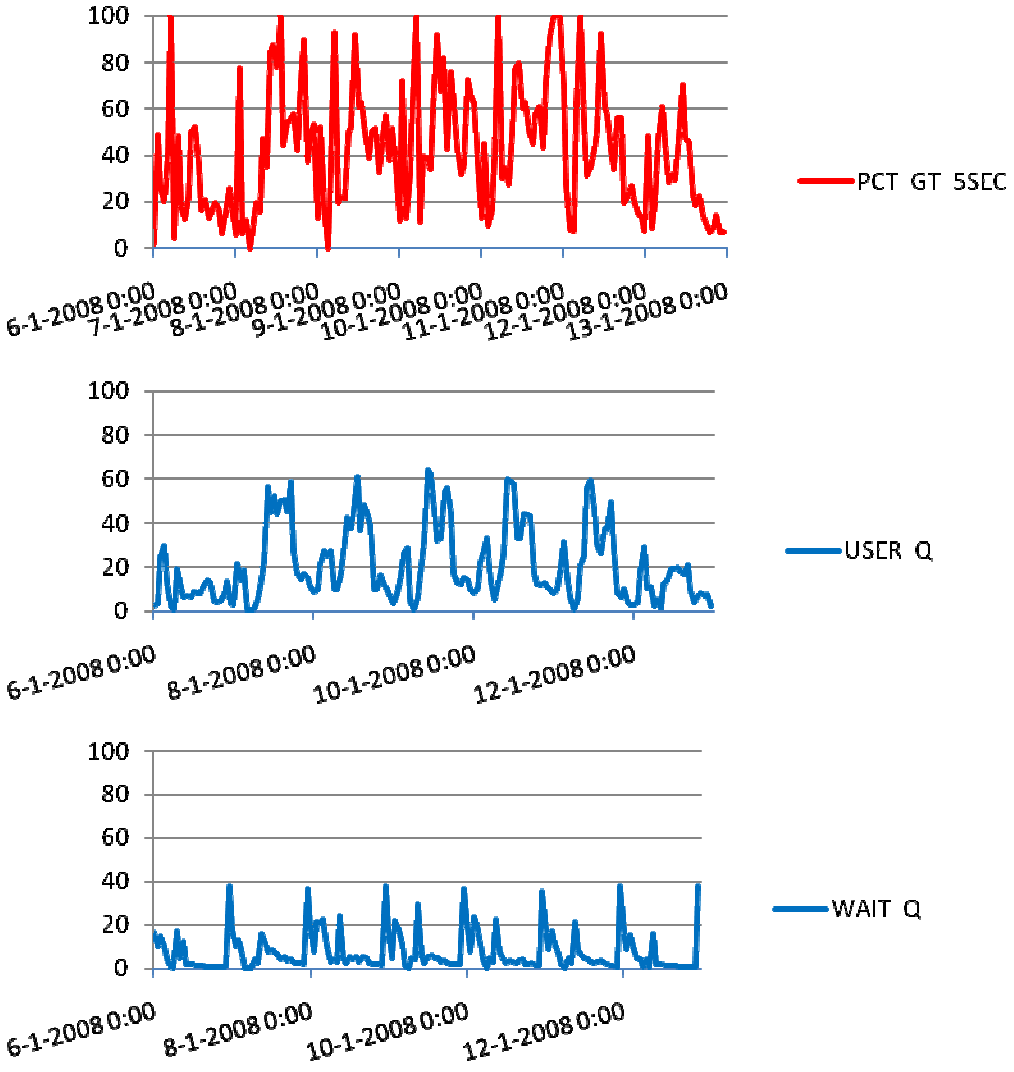

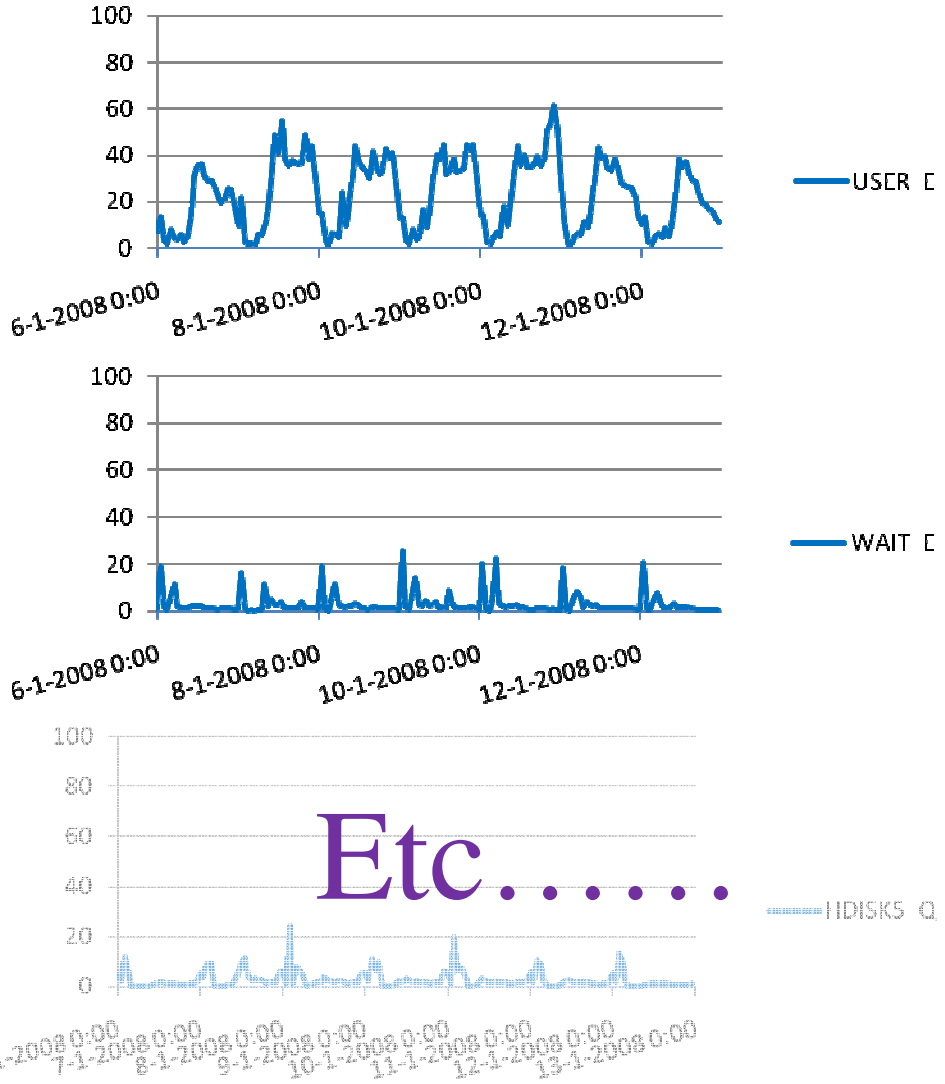

Data Synchronization…

- Aggregate data on time stamps
- Be aware of time differences
- Make sure that big workload changes are identified in your data
- Values examination
- Etc.

The Fourth D

Data Mining – Explain

(The third D will be explained later)

Doing a small test in Swingbench…

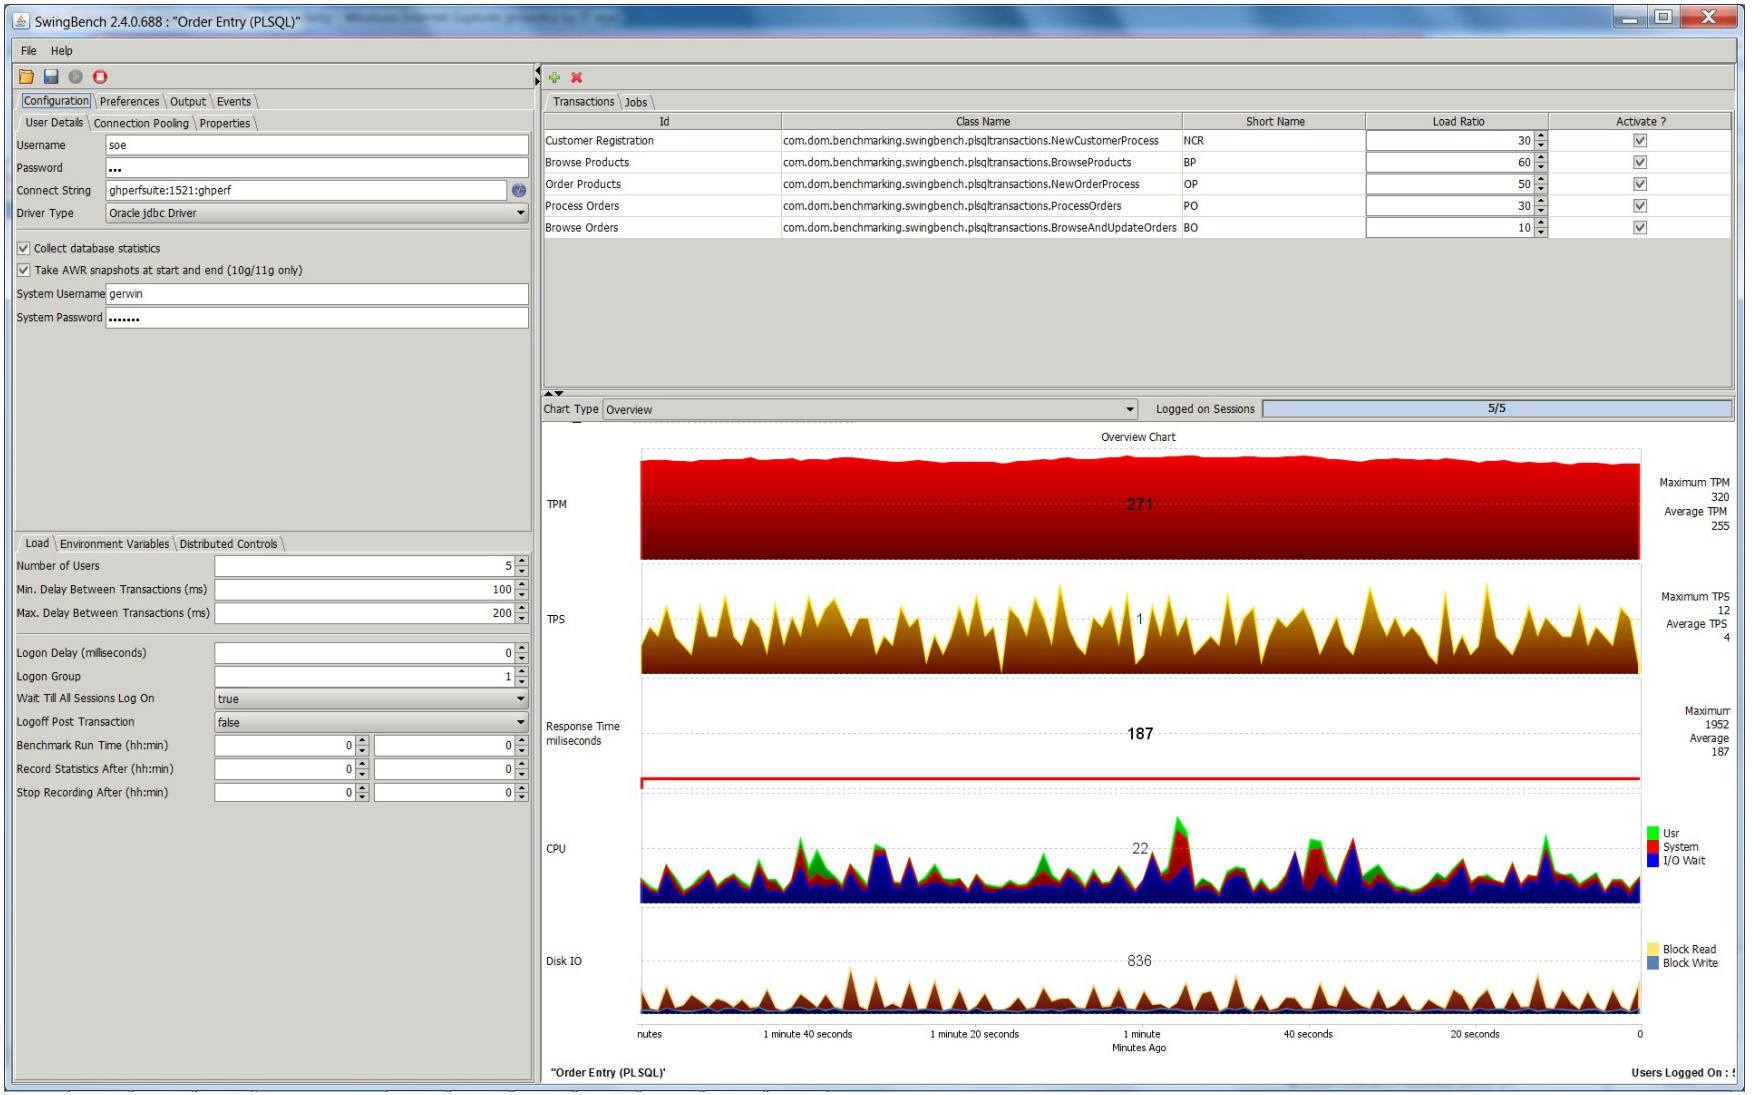

Gathered data in the test…

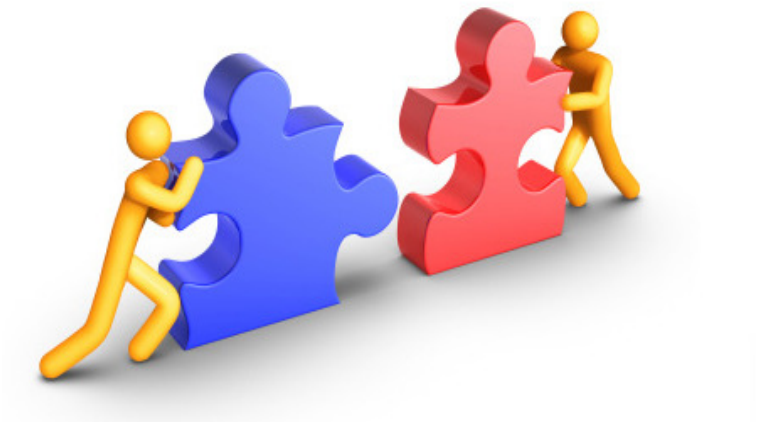

- Client data (Swingbench)
- SAR data on the linux virtual machine
- AWR data on the 11.2.0.3 database

Do the data synchronization and…

Factorial Analysis on the SAR data…

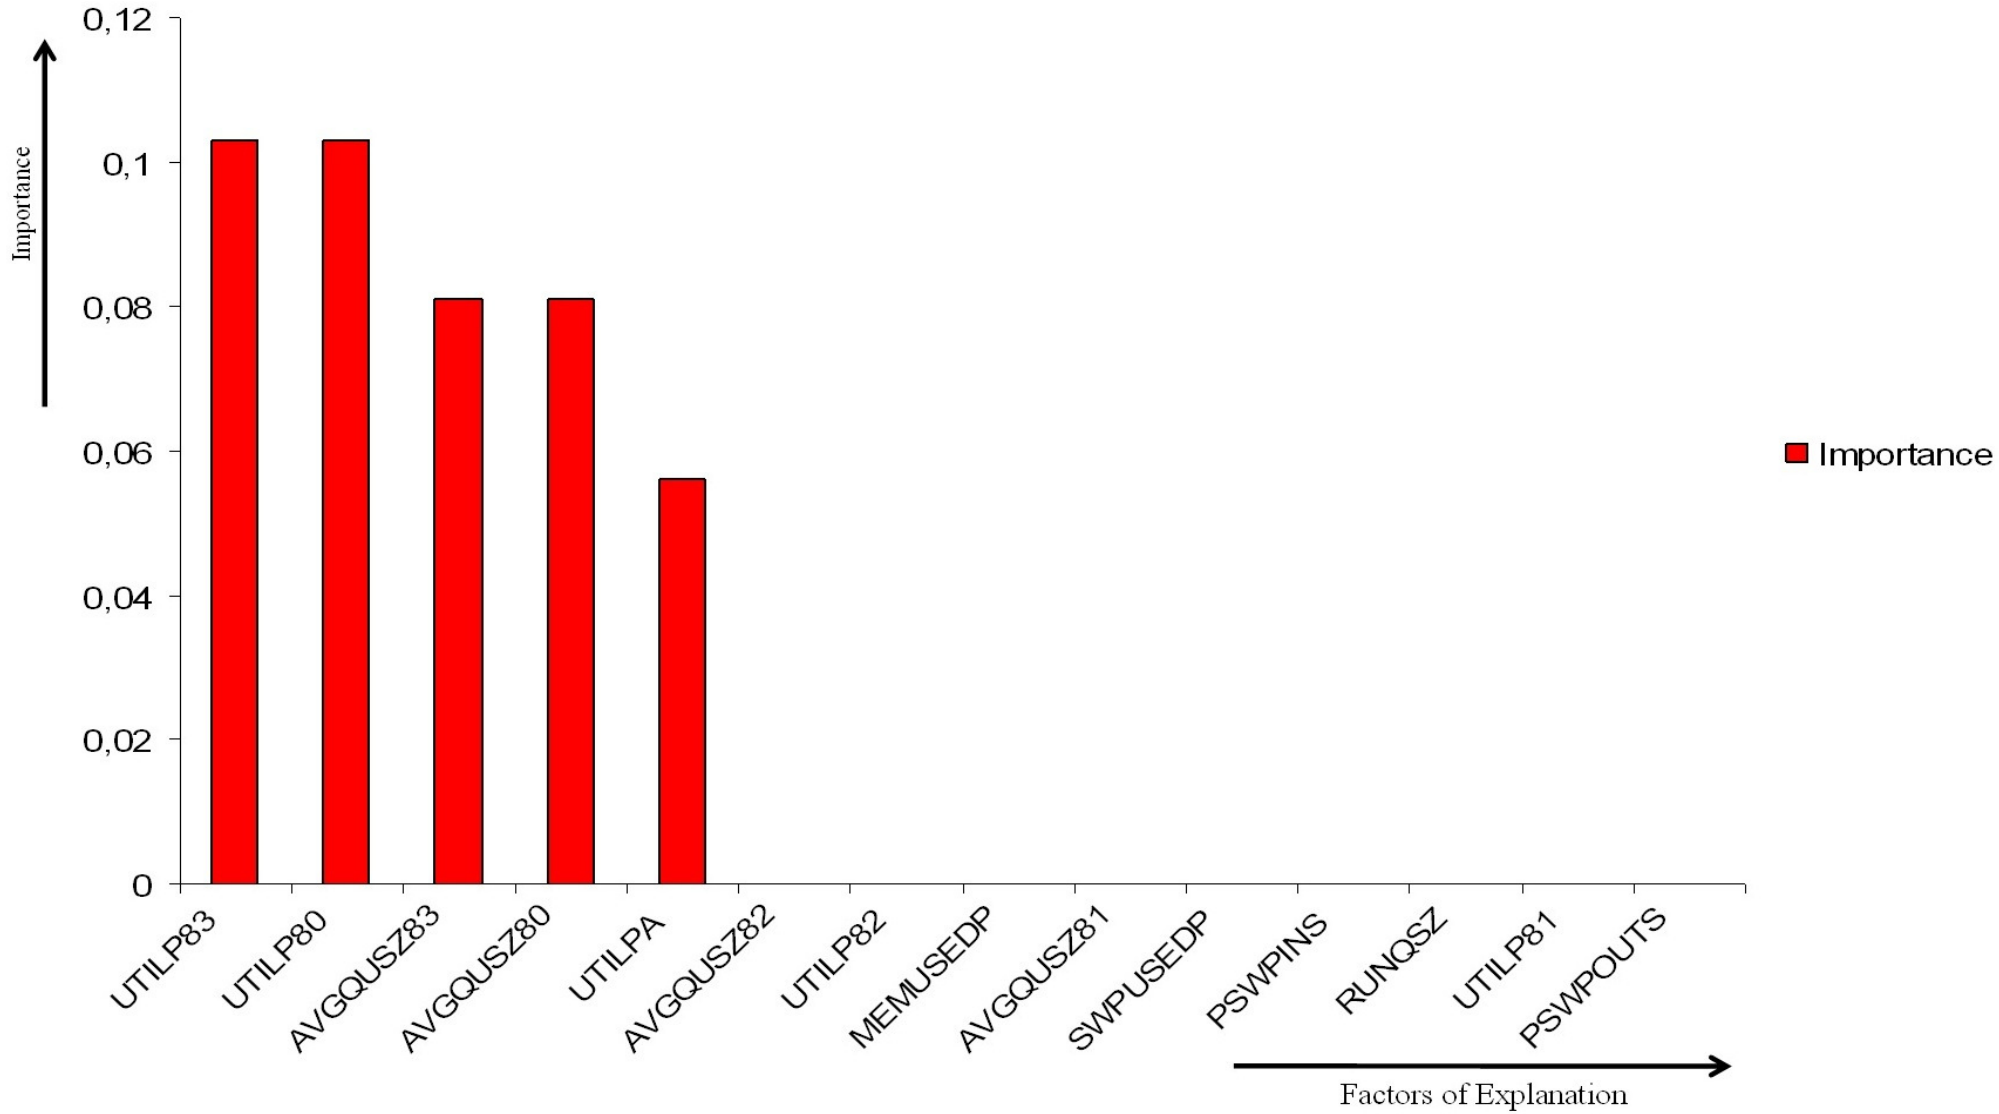

The Importance / Index number…

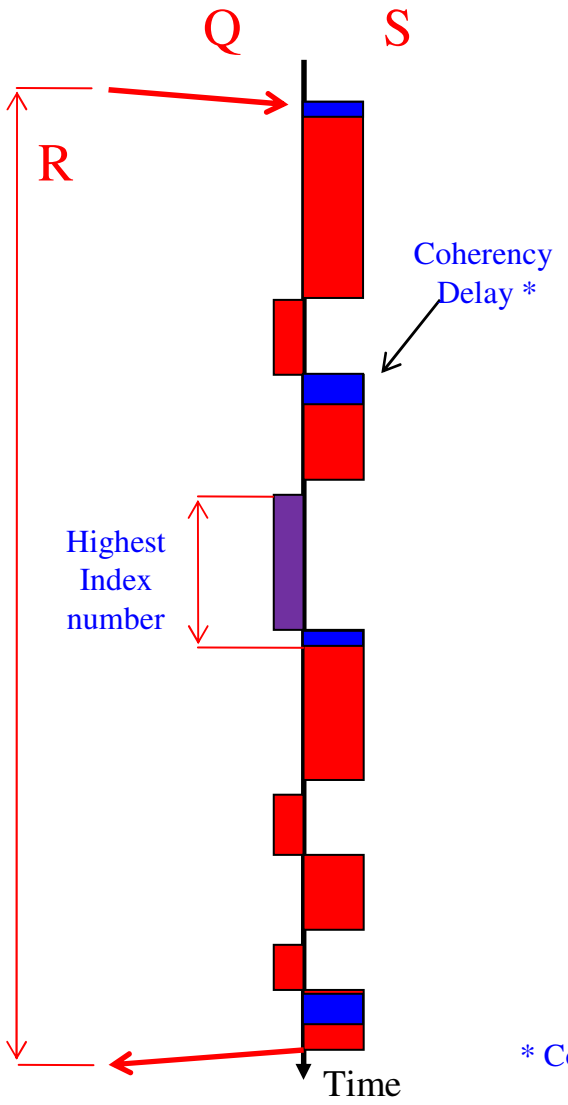

- Basically we search for the block with the biggest variance
- The factor with the biggest variance in the total of "Wait Time" + "Stretched Service Time", will have the highest index number.
- So the higher the Index number, the higher the impact on the total variance of R, and so worth investigation.

* Coherency Delay simplified

The Third D

Data Modeling

Normalized R based on Erlang-C…

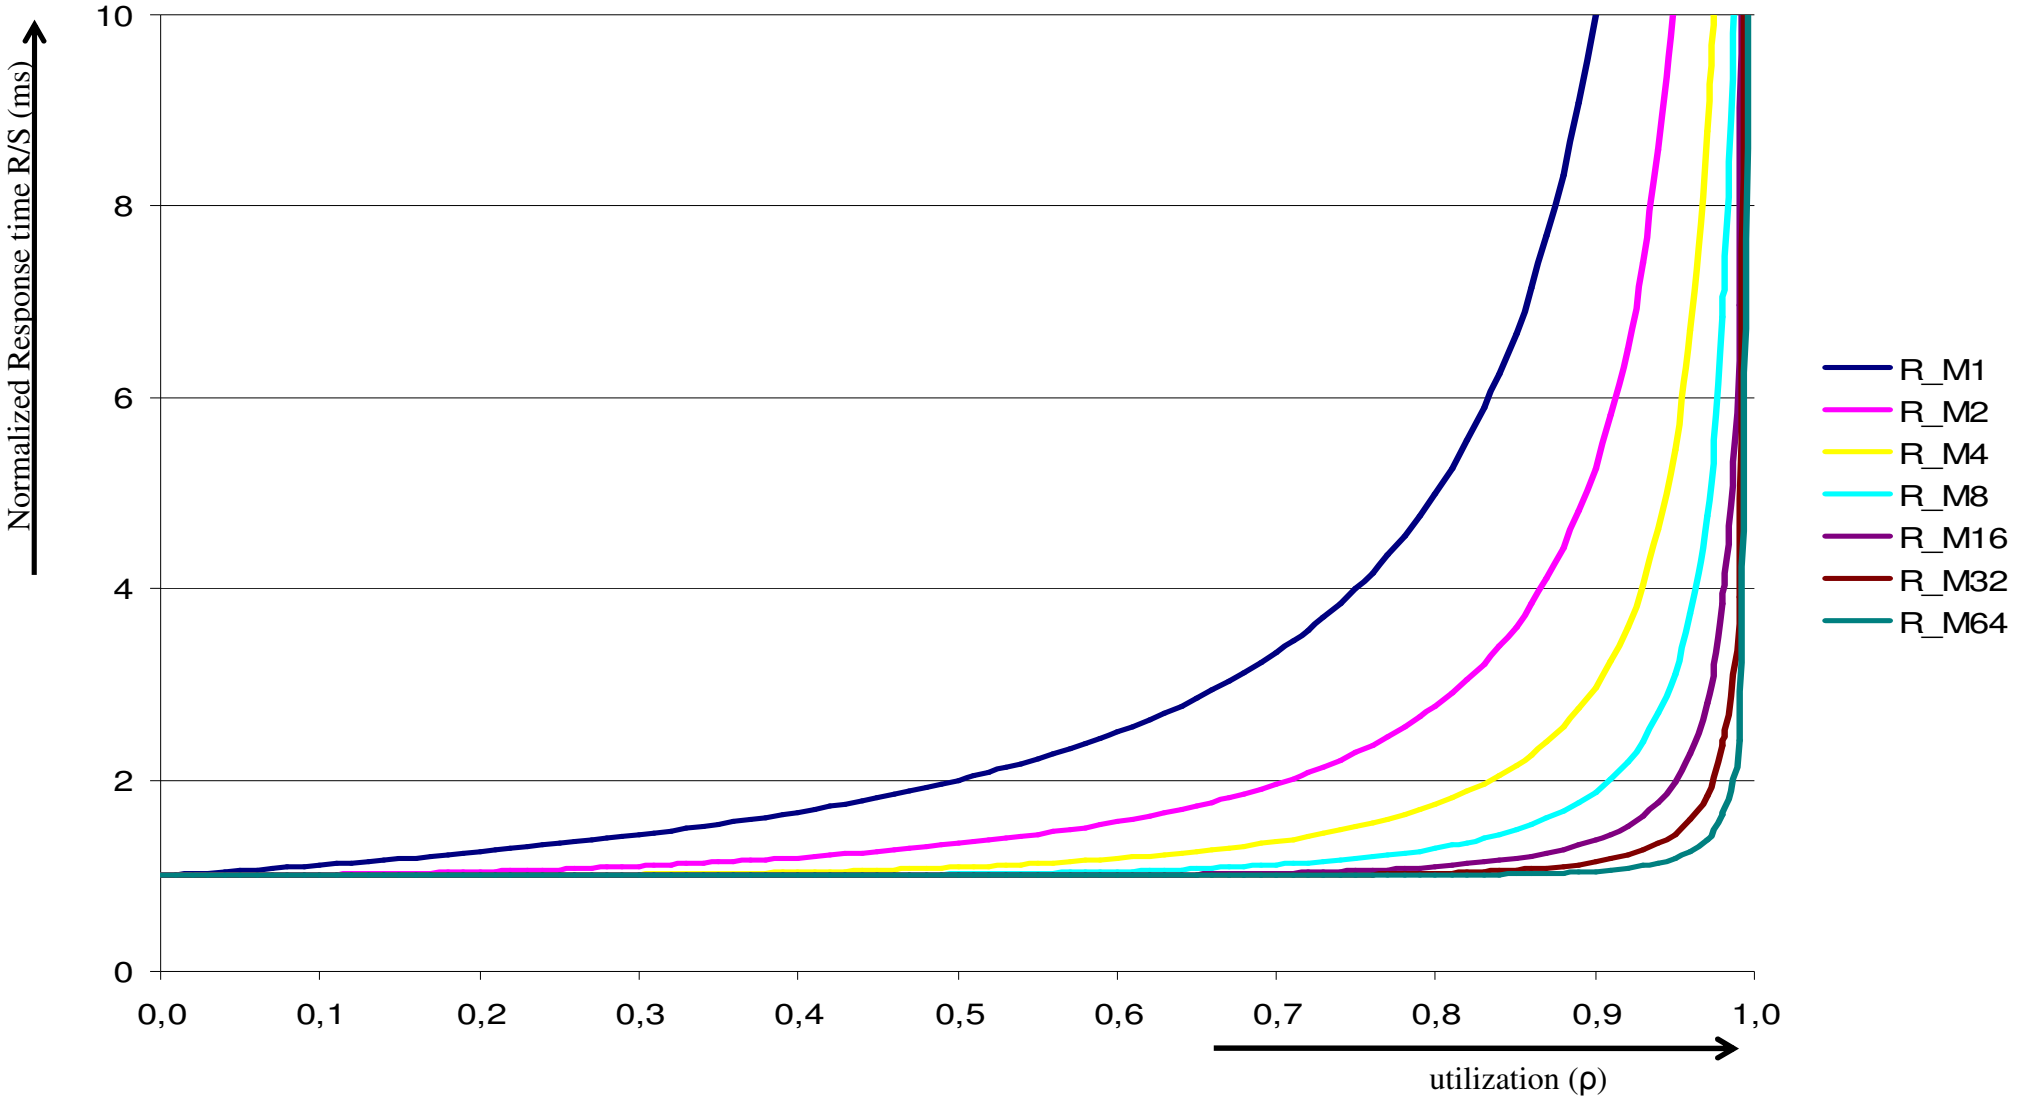

Erlang-C formula and R calculation…

$$C(m,\rho) = \frac{\frac{(m\rho)^m}{m!}}{(1-\rho)\sum_{n=0}^{m-1}\frac{(m\rho)^n}{n!} + \frac{(m\rho)^m}{m!}}$$

$$R = \frac{C(m,\rho)S}{m(1-\rho)} + S$$

Why doing data modeling…

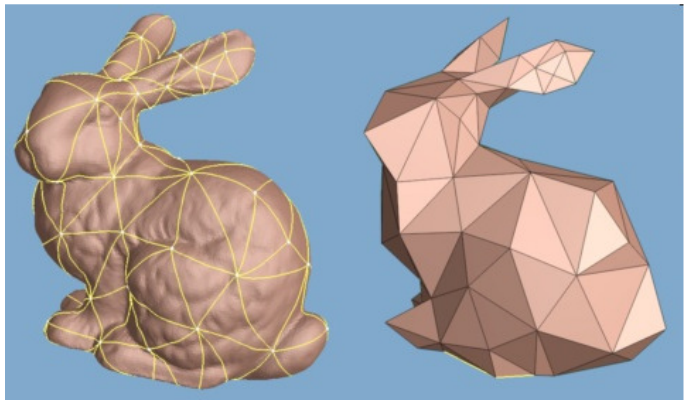

- The data modeling is done to be able to do:
  - Multi Linear Regression Data Mining
  - Able to create a linear function to describe the end-user response time (R)
  - When having a linear function we can determine the service time (S) and queuing time (Q) per component

M1 Curve not corrected yet…

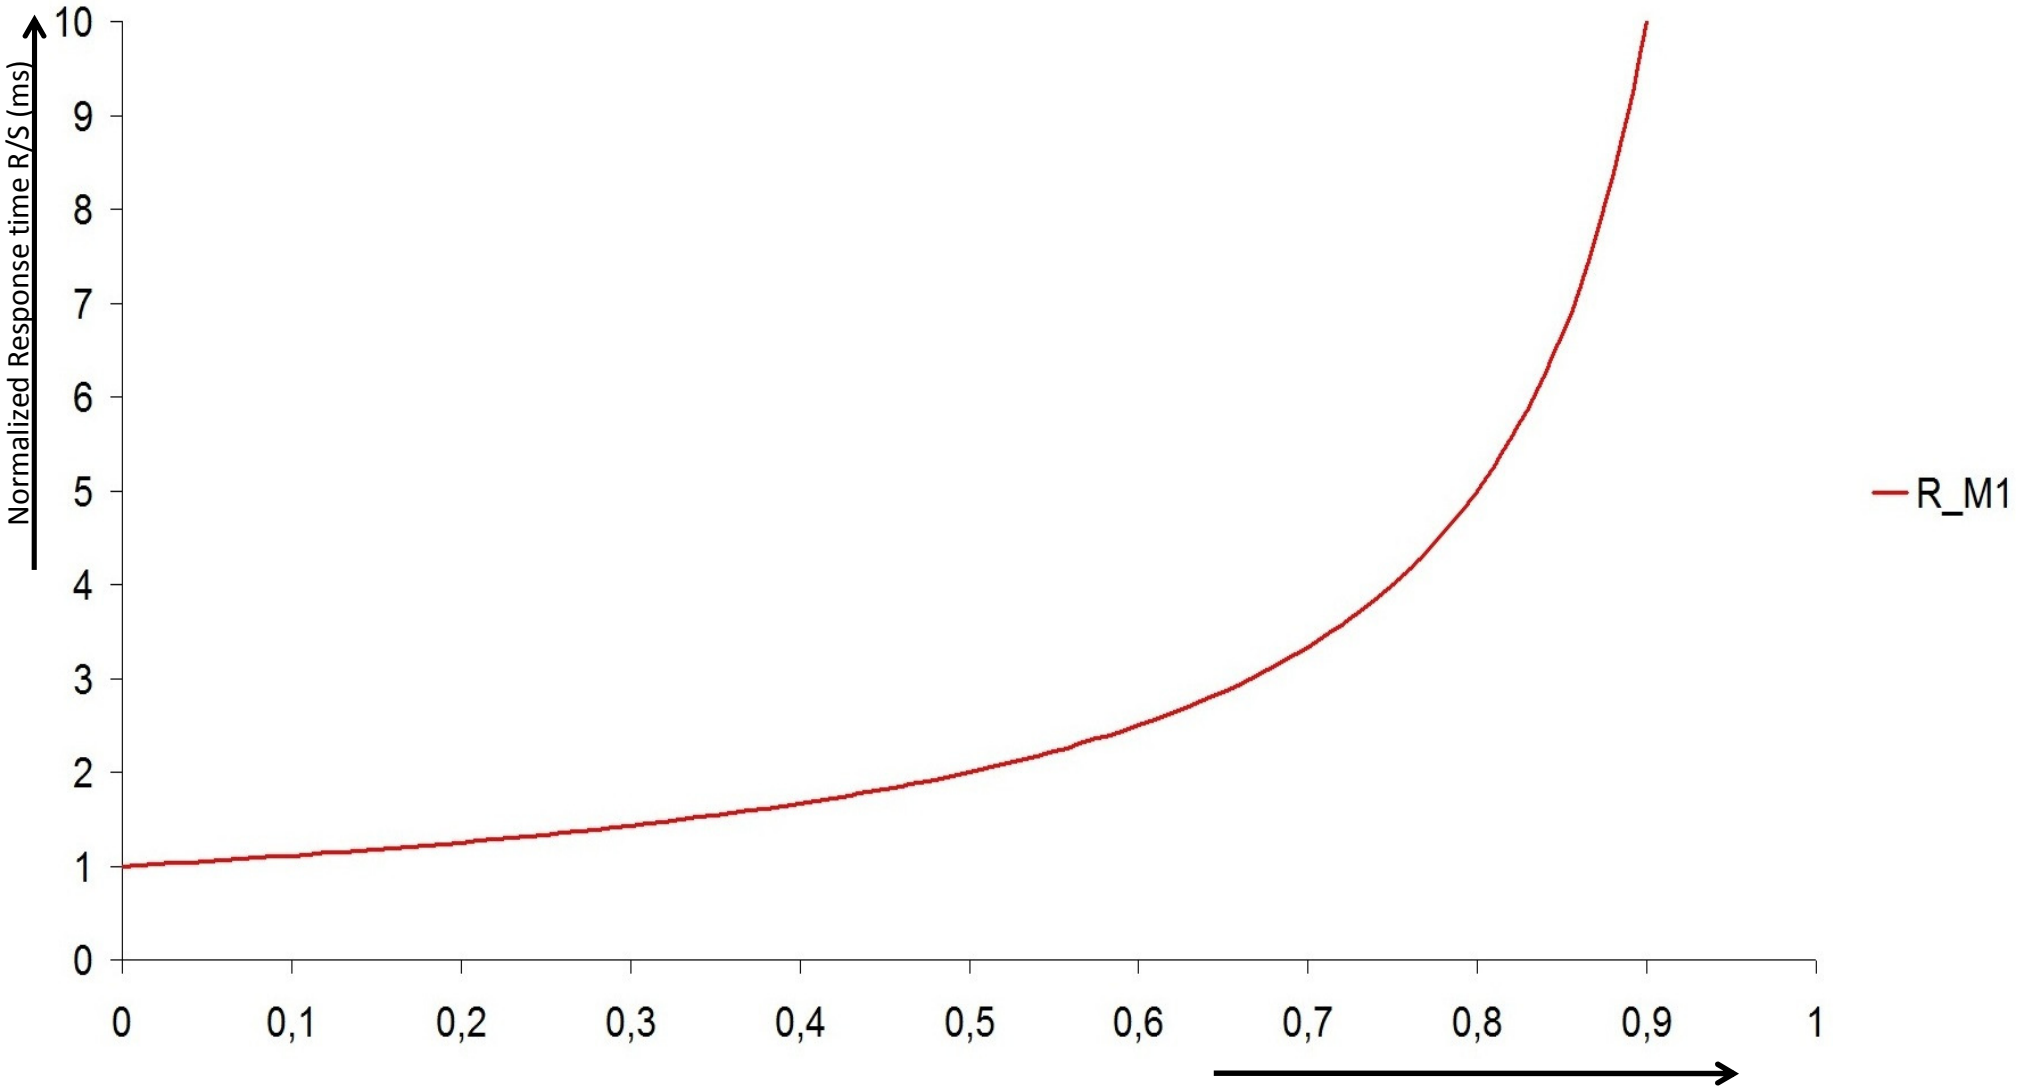

M1 Curve corrected by the modelling…

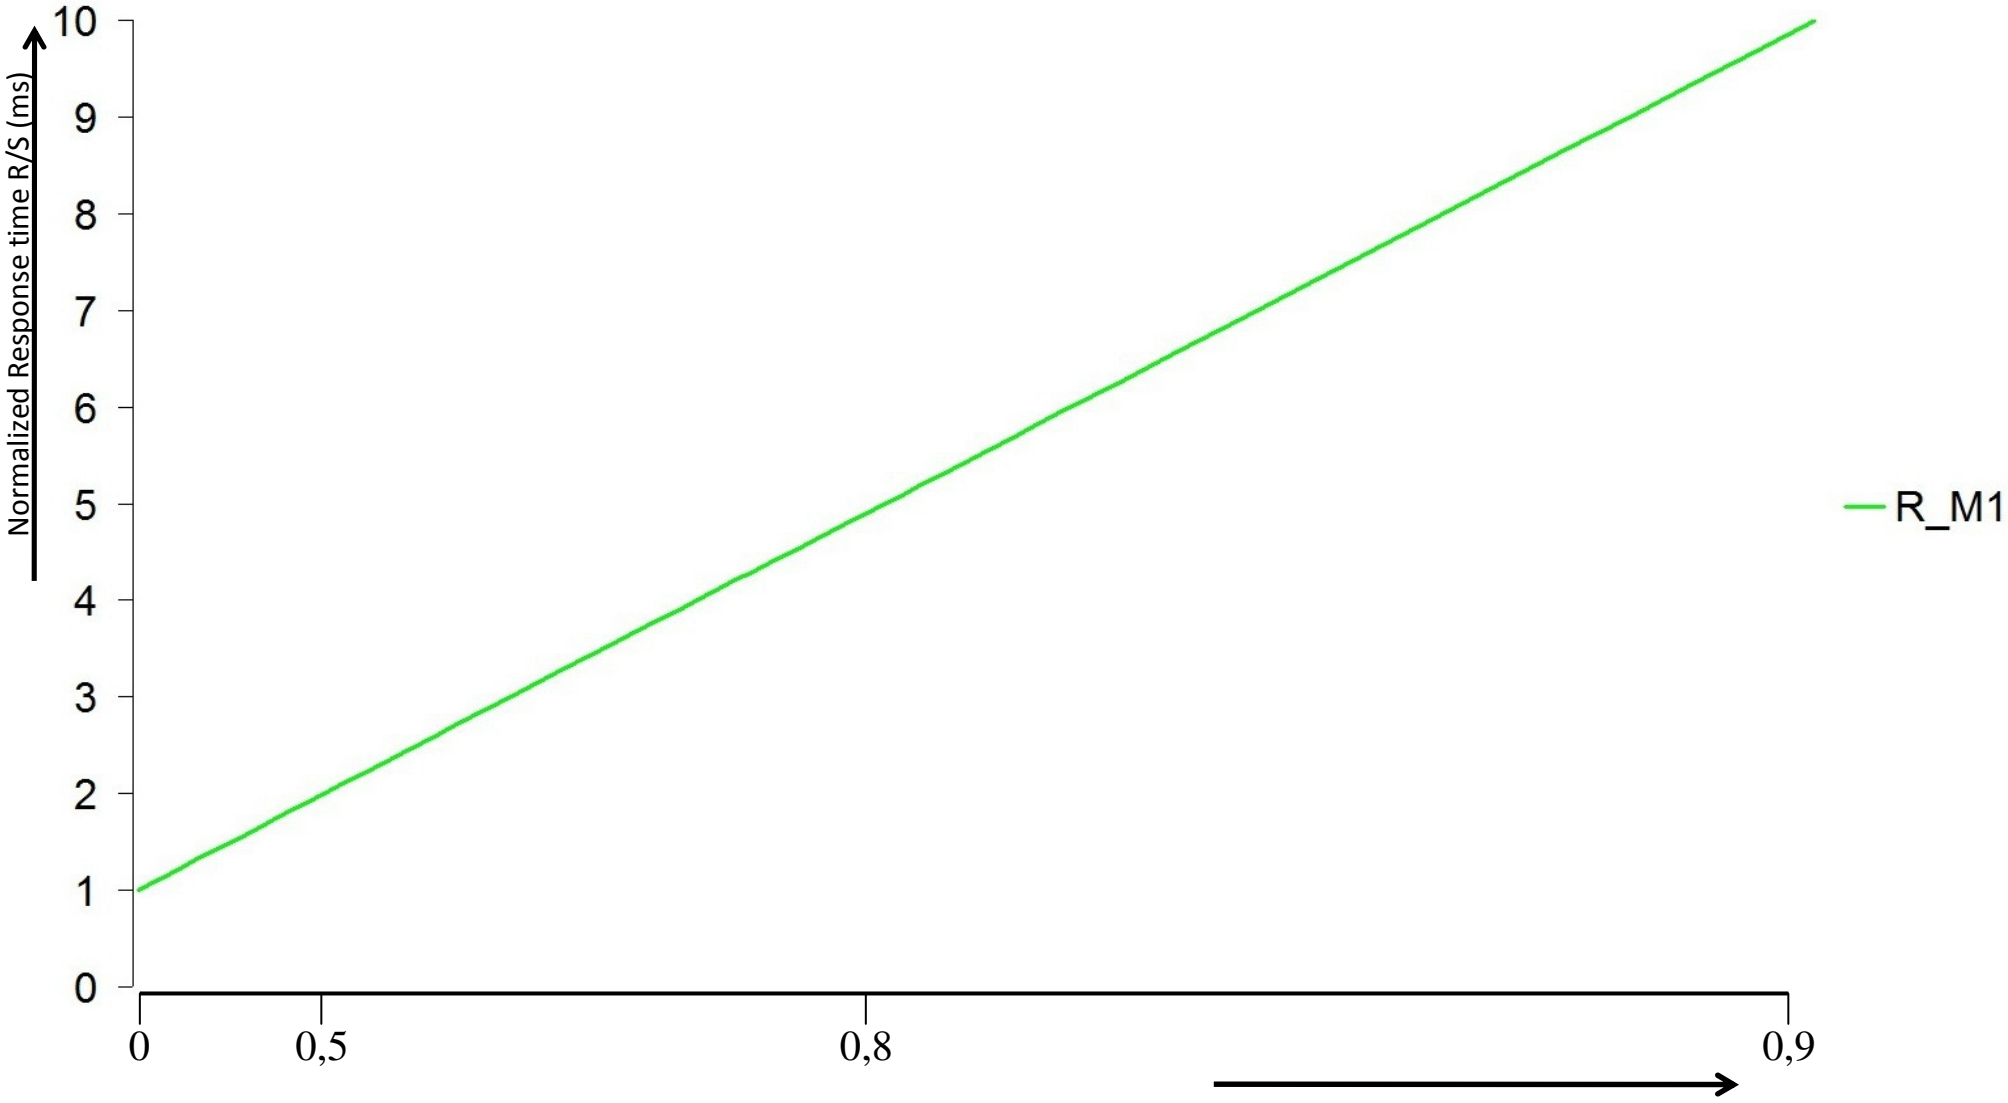

Response time versus the I/O util…

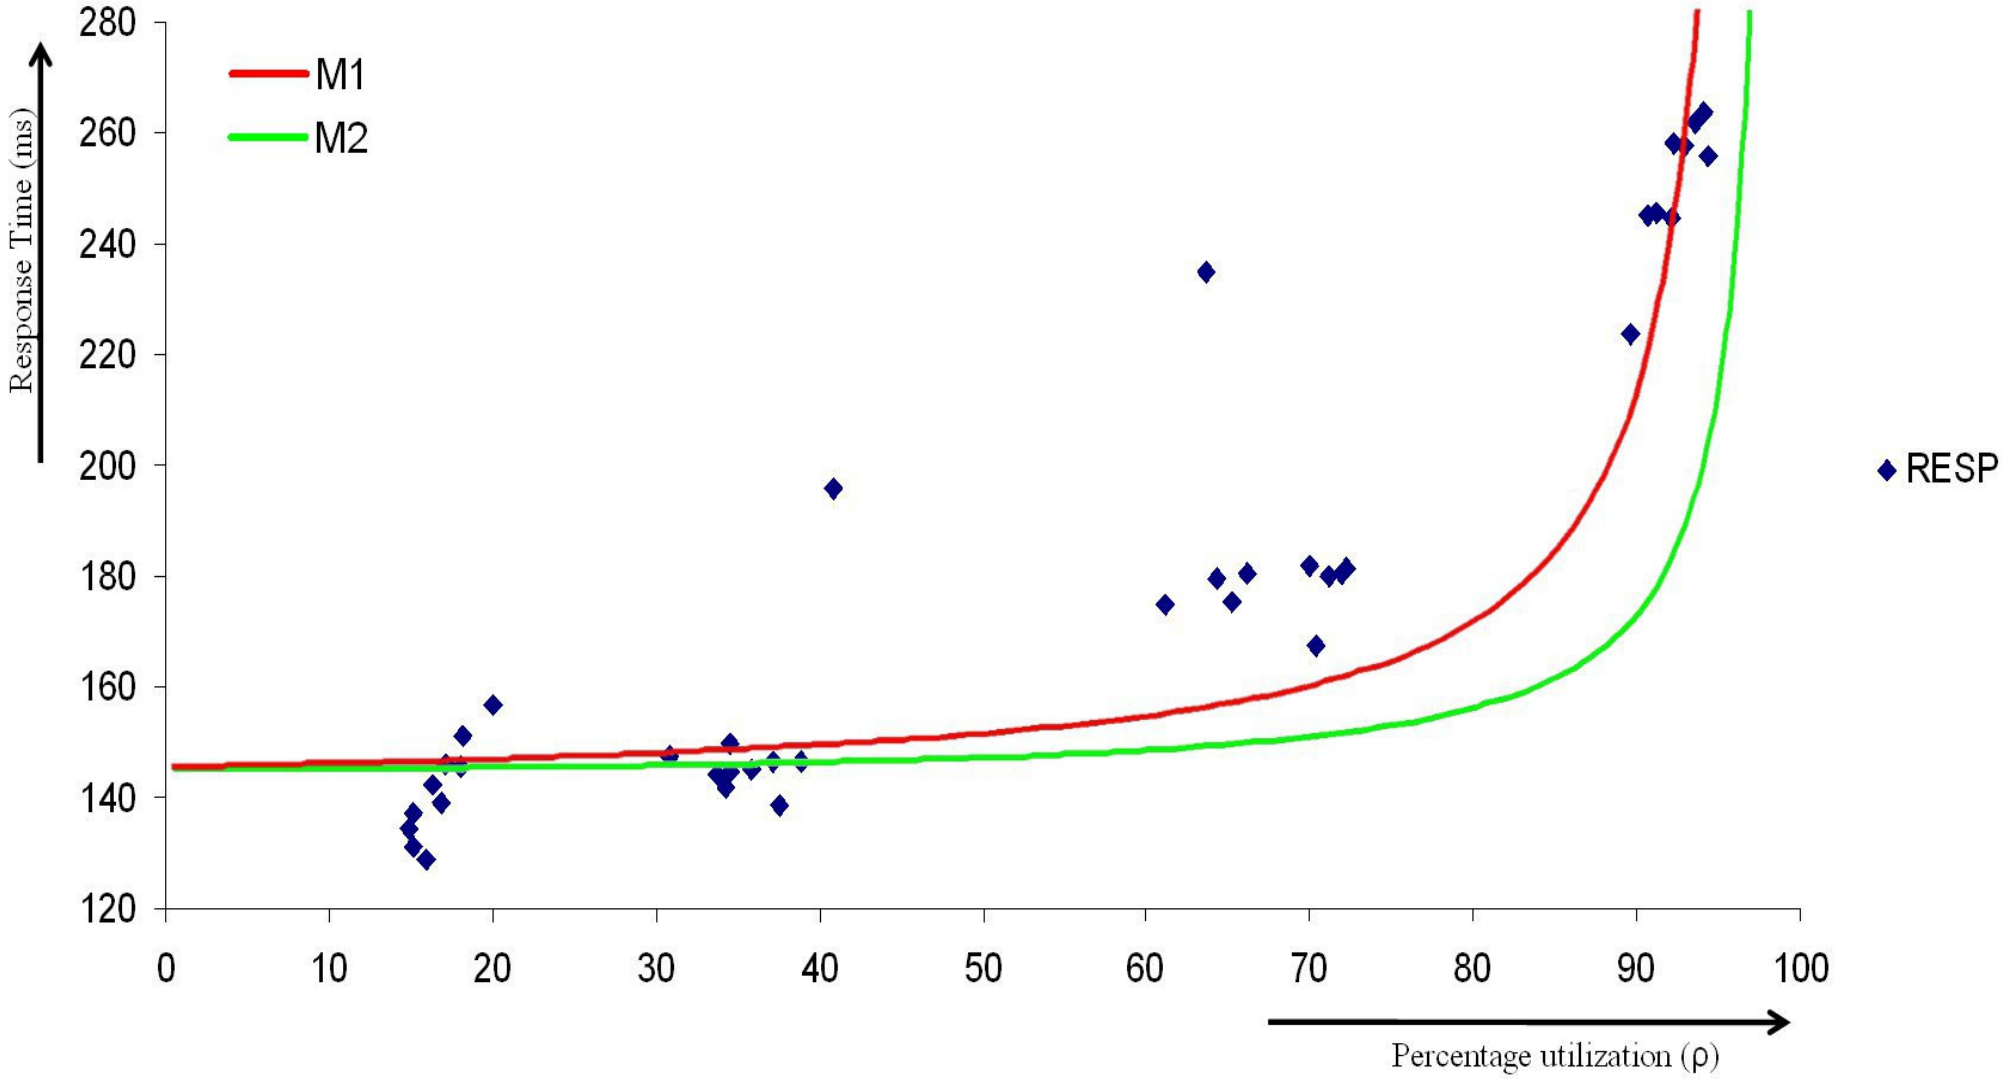

The Fourth D

Data Mining – Model

(Second part of the fourth D)

Data Mining and data modeling…

For Example:

$$R_{tot} = c_1 R_{1,n=8} + c_2 R_{2,n=4} + c_3 R_{3,n=1} + etc.$$

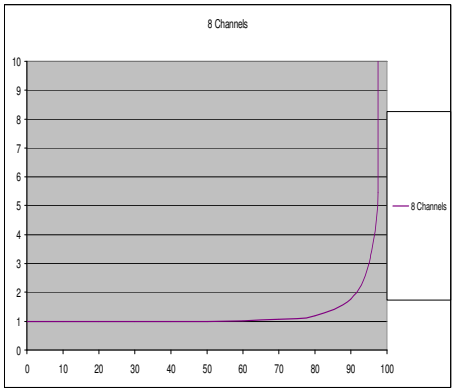

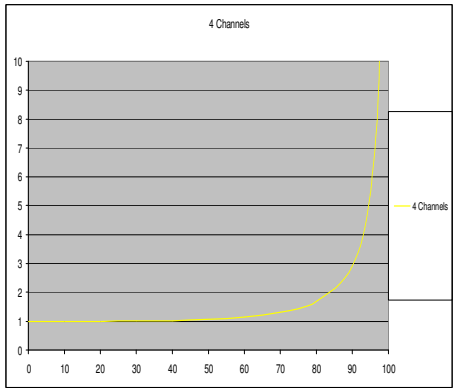

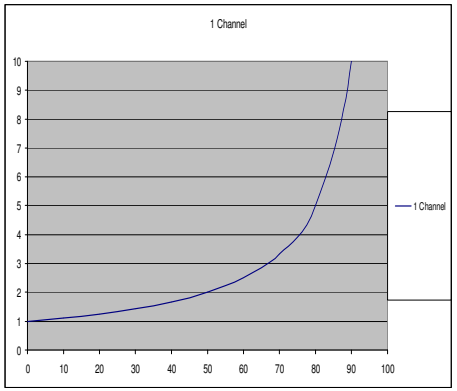

Gives the best fit, determined by data mining…

Creating the linear formula…

- Based on the coefficients of the calculated model and the “intercept row” of the model we can create the formula for the response time prediction. (below model only with I/O)

$$R_{RESP} \approx 7.52 \times M_{IO\,n=1} + 145.41$$

$$R_{RESP} \approx R_{UTILR80} + R_{REST}$$

Finding Service Time (S)…

$$R_{RESP} = c_1 R_{Normalized\_UTILRAU} + c_2 R_{Normalized\_UTILR80} + R_{REST}$$

- When the corrected factors of the components are normalized, that means R/S (M-Curve fitting)
  - Than c1 is the Service time (S) for CPU (UTILRAU)
  - Than c2 is the Service time (S) for I/O util. (UTILR80)

The modeled I/O utilization with M1…

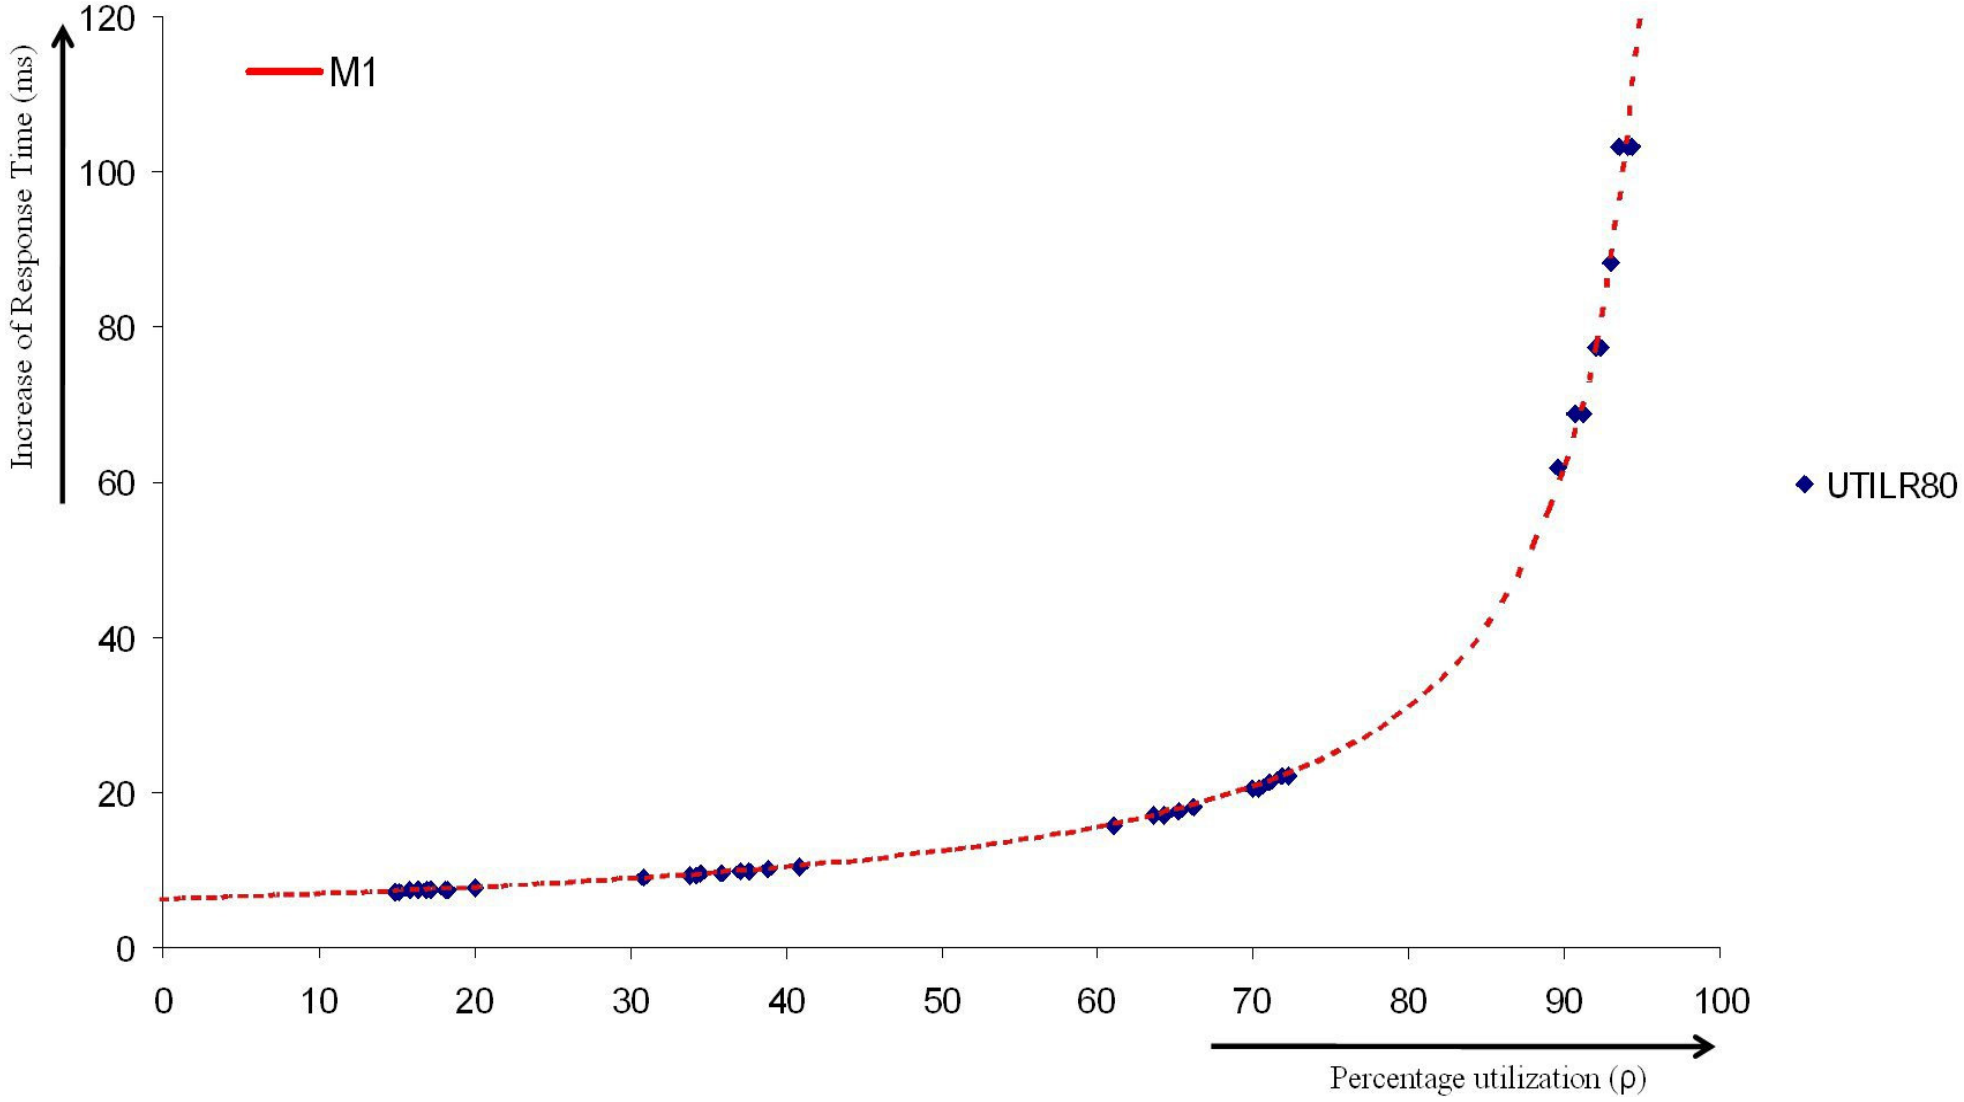

The modeled CPU utilization with M4…

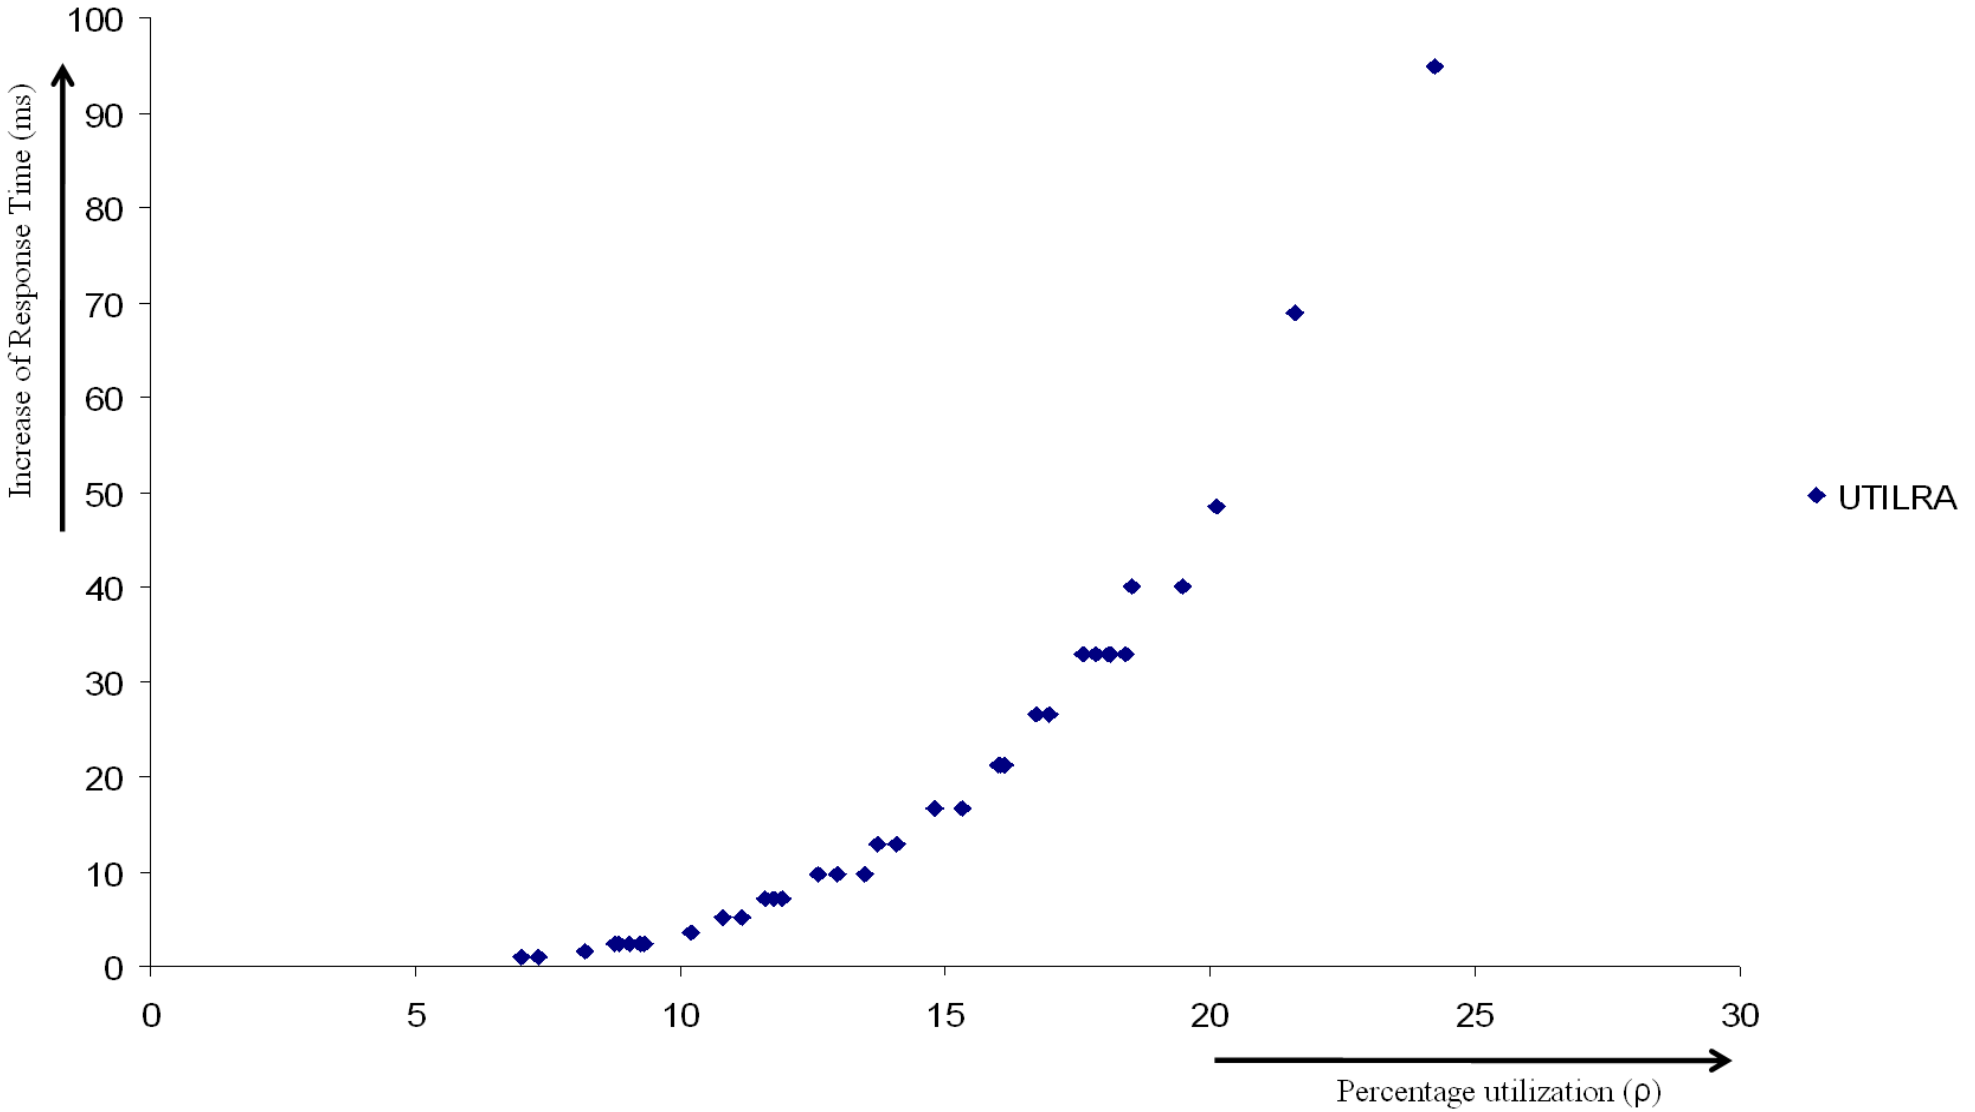

The model with the breakdown…

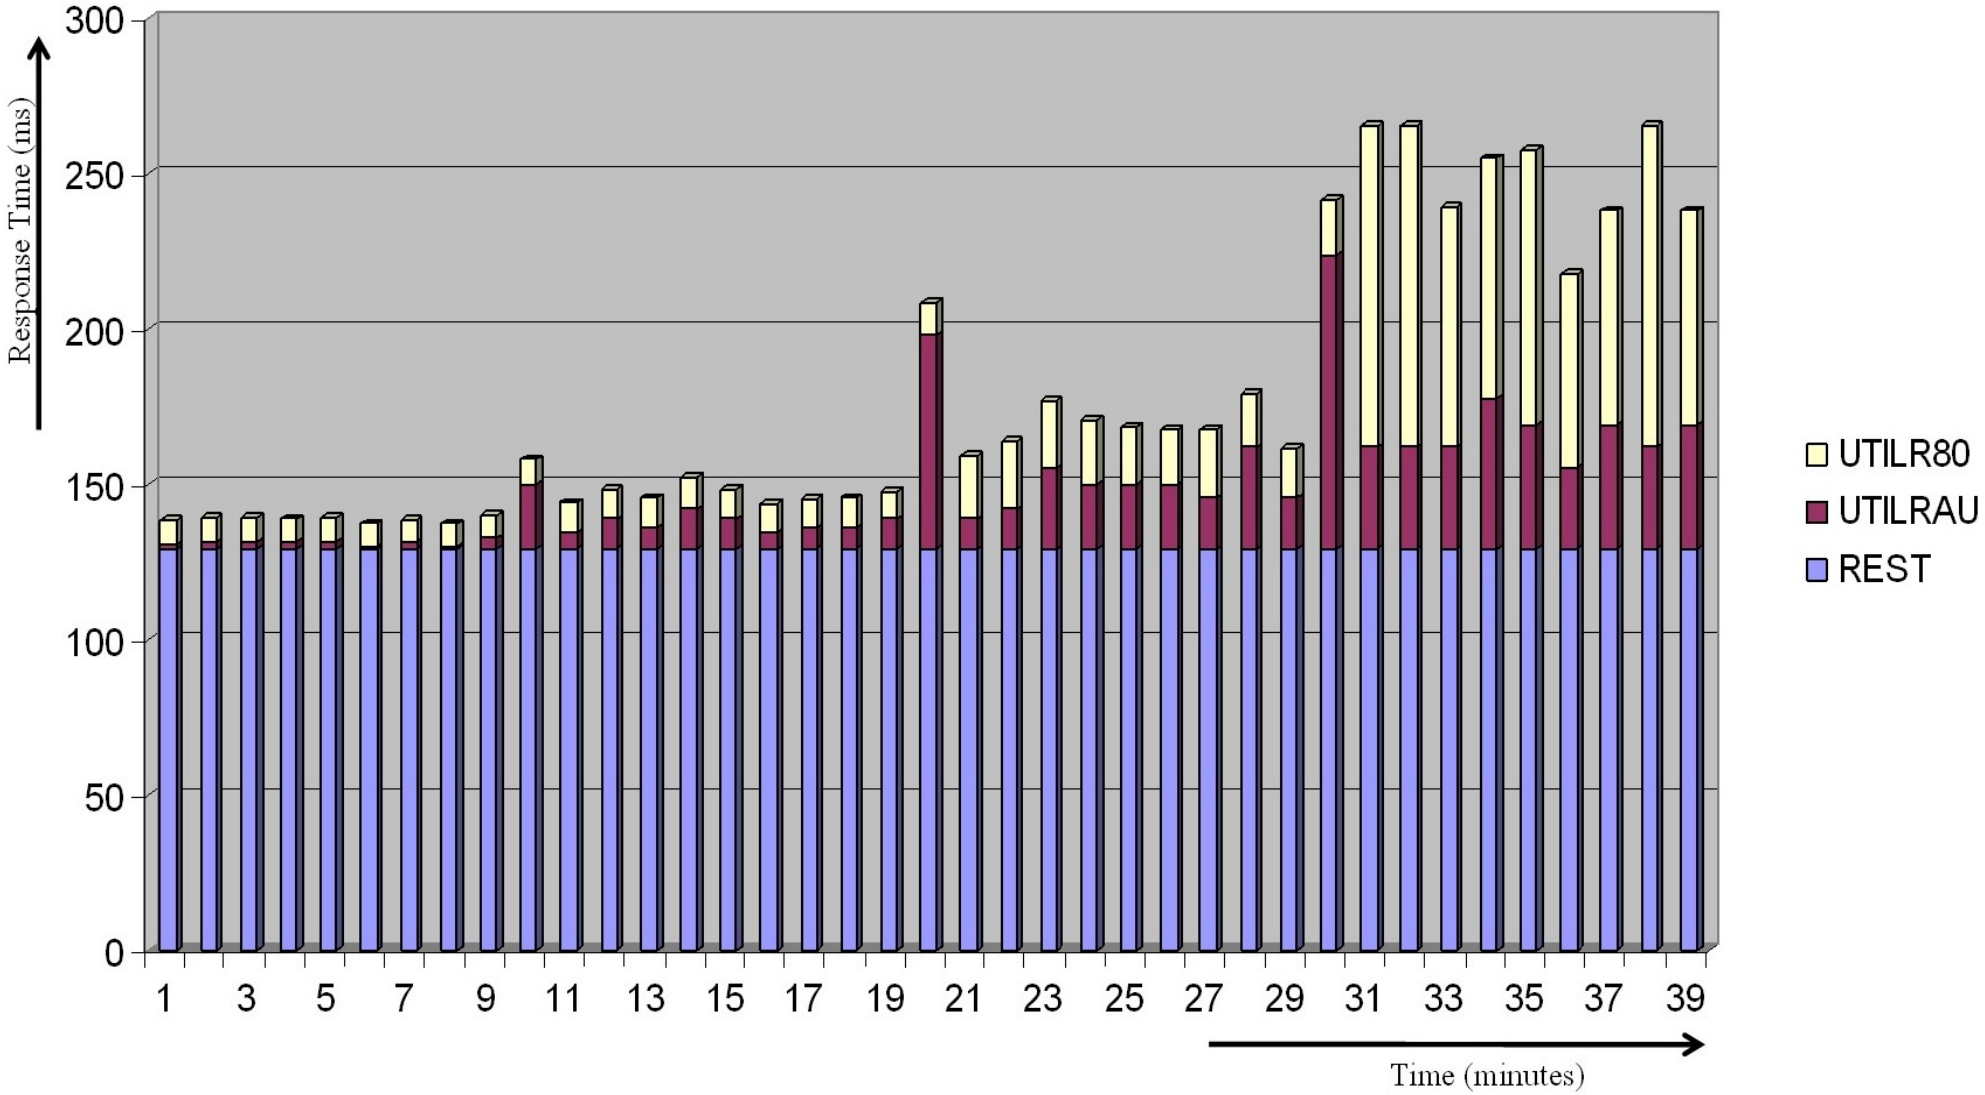

Models in compare with R measured…

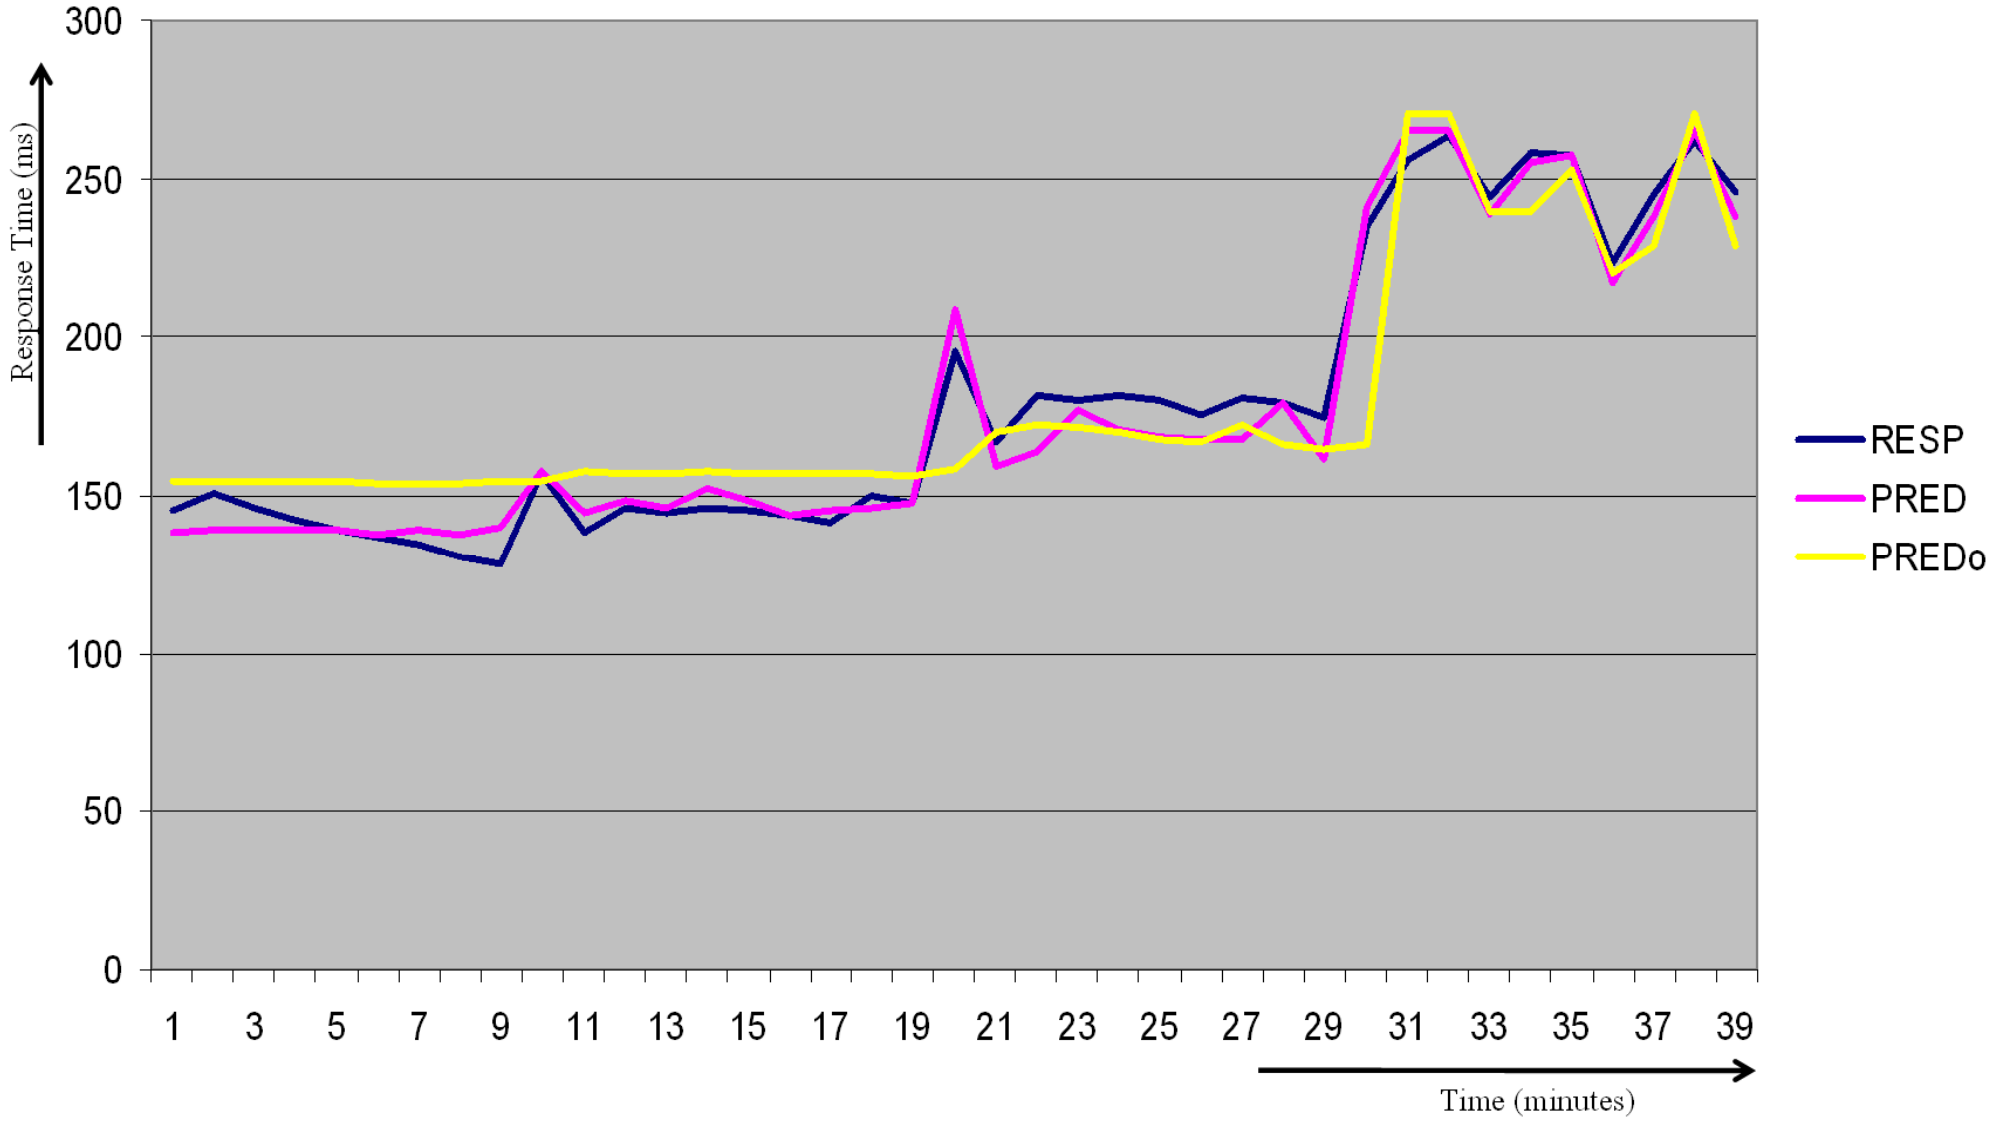

AWR Data usage

Fourth D Data Mining – Explain using secondary component data

How to retrieve the used AWR data…

- To get the AWR SQL data response times:
  - elapsed_time_delta / executions_delta
    - From DBA_HIST_SQLSTAT
- To get the AWR events data:
  - TIME_WAITED_MICRO_FG
    - From DBA_HIST_SYSTEM_EVENT
    - Calculate the delta between snapshots

SQL Response Time Data from AWR…

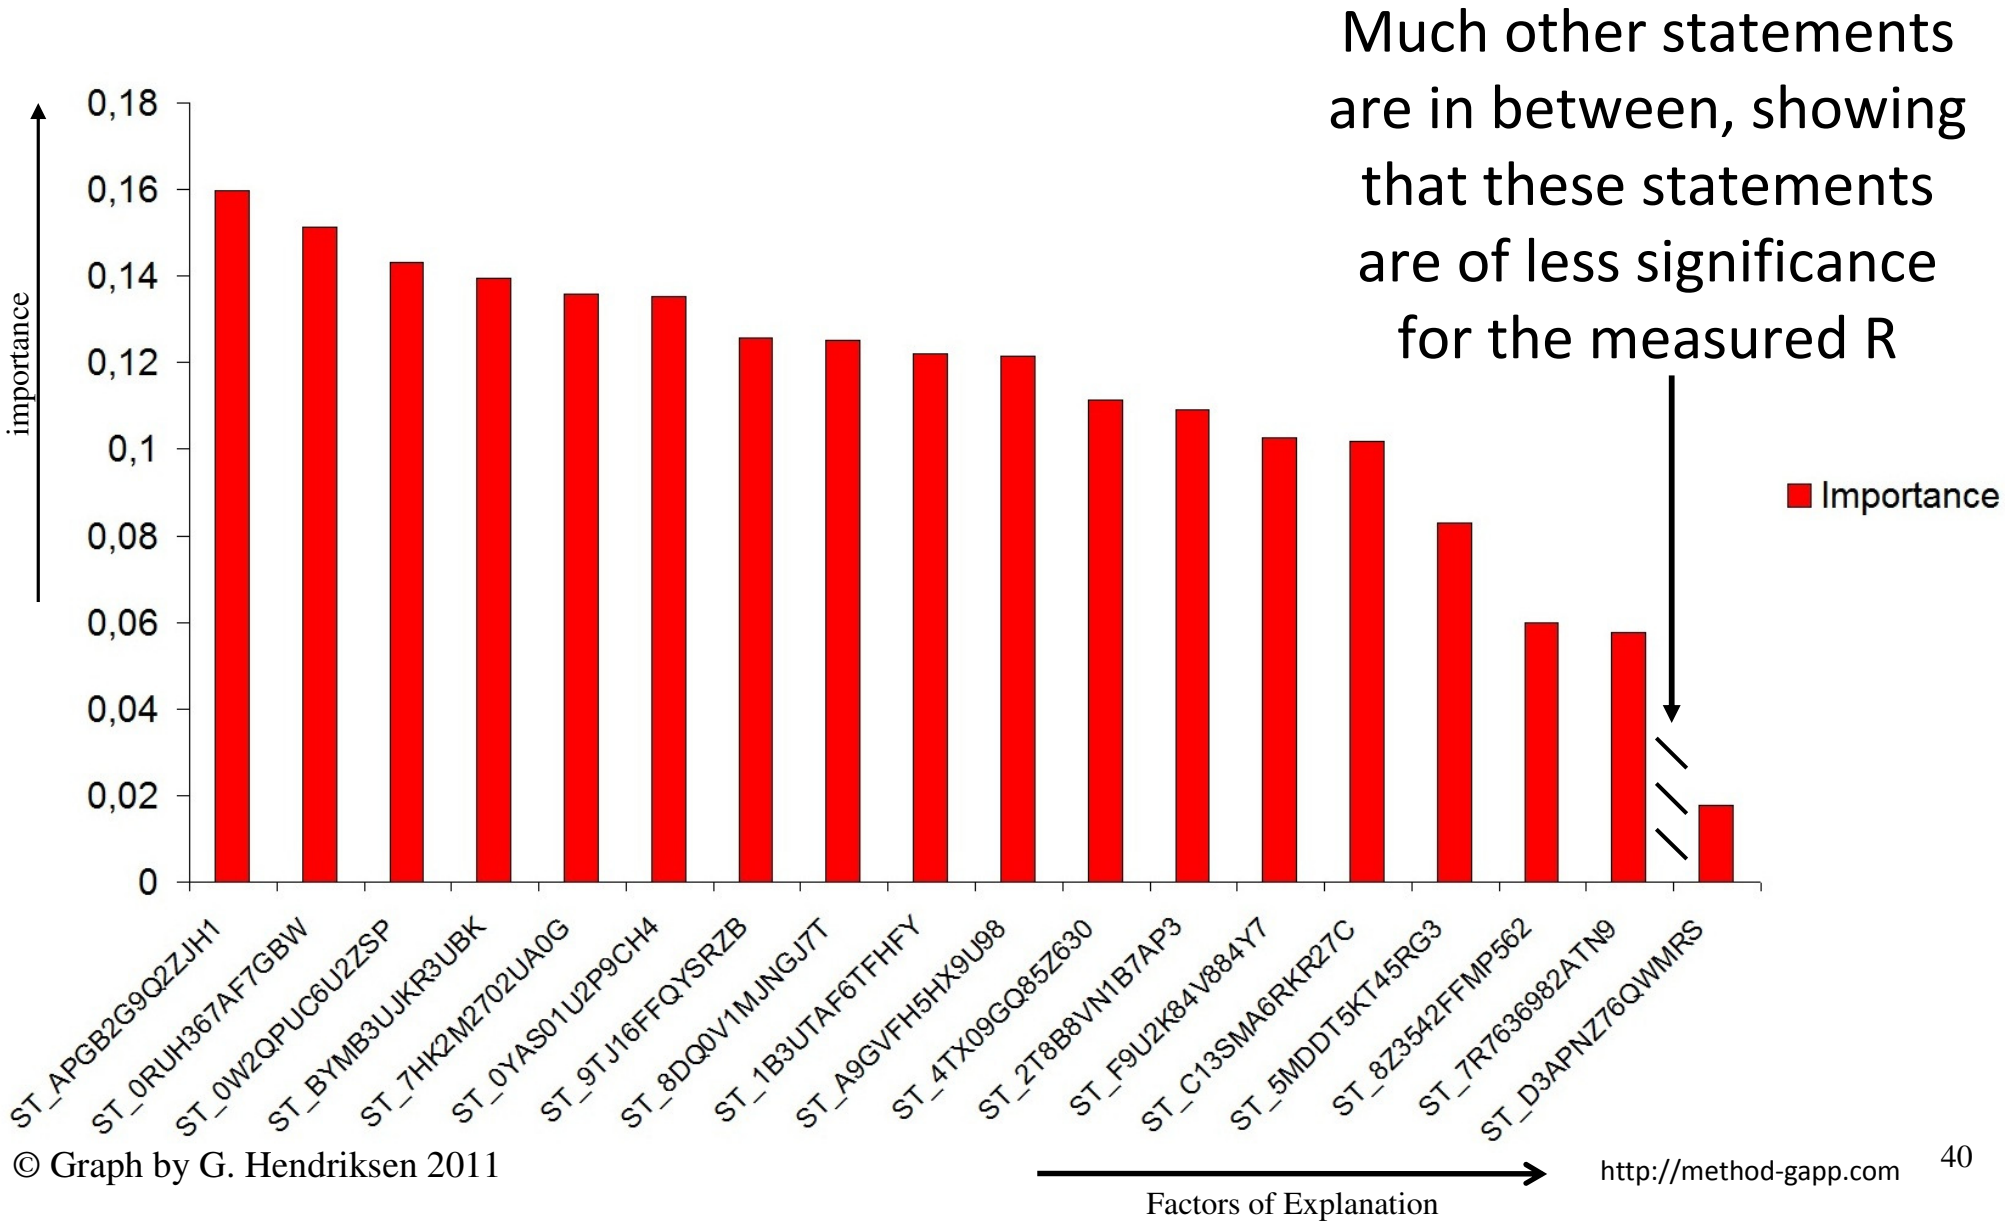

Involved important Statements…

- **'apgb2g9q2zjh1'**
  - BEGIN :1 := orderentry.browseandupdateorders(:2 ,:3 ,:4); END;
- **'0ruh367af7gbw'**
  - SELECT ORDER_ID, ORDER_MODE, CUSTOMER_ID, ORDER_STATUS, ORDER_TOTAL, SALES_REP_ID, PROMOTION_ID FROM ORDERS WHERE CUSTOMER_ID = :B2 AND ROWNUM < :B1
- **'0w2qpuc6u2zsp'**
  - BEGIN :1 := orderentry.neworder(:2 ,:3 ,:4); END;
- **'bymb3ujkr3ubk'**
  - INSERT INTO ORDERS(ORDER_ID, ORDER_DATE, CUSTOMER_ID, WAREHOUSE_ID) VALUES (ORDERS_SEQ.NEXTVAL + :B3 , SYSTIMESTAMP , :B2 , :B1) RETURNING ORDER_ID INTO :O0
- **'7hk2m2702ua0g'**
  - WITH NEED_TO_PROCESS AS (SELECT ORDER_ID, CUSTOMER_ID FROM ORDERS WHERE ORDER_STATUS <= 4 AND WAREHOUSE_ID = :B1 AND ROWNUM < 10) …
- **'0yas01u2p9ch4'**
  - INSERT INTO ORDER_ITEMS(ORDER_ID, LINE_ITEM_ID, PRODUCT_ID, UNIT_PRICE, QUANTITY) VALUES (:B4 , :B3 , :B2 , :B1 , 1)
- **Etc.**

- **Not significant statement**: **'d3apnz76qwmrs'**
  - select pol#, usr_name, usr_labels, package, privs from lbac$user_logon where usr_name = :username

Events Data from AWR…

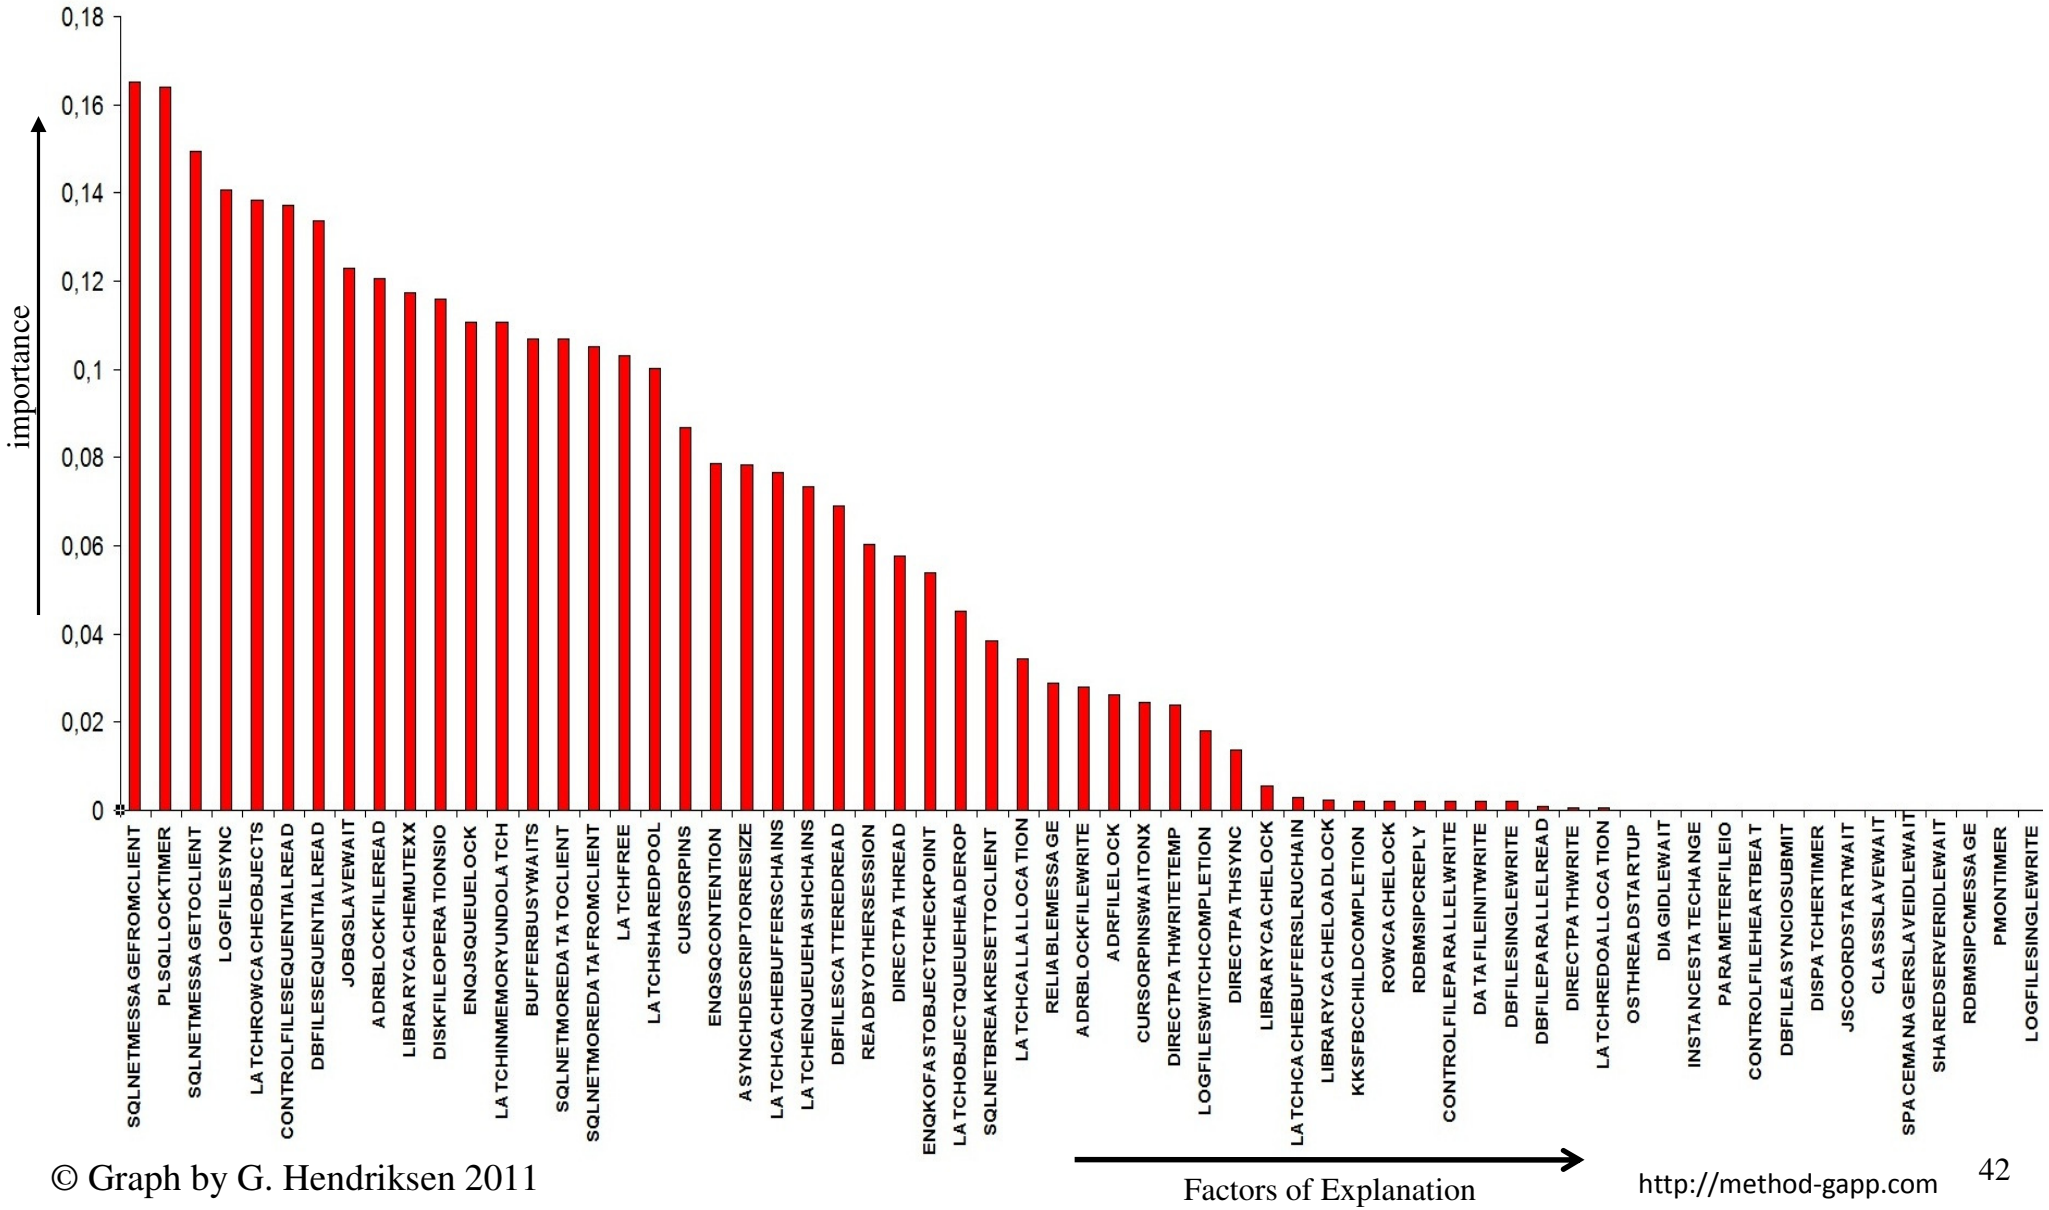

Involved wait events…

- SQL*Net message from client
- PL/SQL lock timer
- SQL*Net message to client
- log file sync
- latch: row cache objects
- control file sequential read
- db file sequential read
- Etc.

Top 5 Timed Foreground Events

Event	Waits	Time(s)	Avg wait (ms)	% DB time	Wait Class
log file sync	12,599	972	77	59.54	Commit
db file sequential read	29,465	379	13	23.19	User I/O
DB CPU		209		12.82	
resmgr:cpu quantum	2,475	15	6	0.91	Scheduler
log file switch (private strand flush incomplete)	28	14	509	0.87	Configuration

The data used in Method-GAPP was over a lot of hours the AWR report shown only over a 10 minute time slice.

Event	Waits	%Time -outs	Total Wait Time (s)	Avg wait (ms)	Waits /txn	% DB time
log file sync	12,599	0	972	77	0.43	59.54
db file sequential read	29,465	0	379	13	1.00	23.19
resmgr:cpu quantum	2,475	0	15	6	0.08	0.91
log file switch (private strand flush incomplete)	28	0	14	509	0.00	0.87
buffer busy waits	37	0	5	139	0.00	0.31
read by other session	10	0	0	33	0.00	0.02
SQL*Net message to client	19,922	0	0	0	0.68	0.02
library cache: mutex X	124	0	0	2	0.00	0.01
ADR block file read	4	0	0	38	0.00	0.01
control file sequential read	1,742	0	0	0	0.06	0.01
log file switch completion	2	0	0	30	0.00	0.00
latch free	11	0	0	4	0.00	0.00

The Fourth D

Data Mining - Predict

What will happen with 2 channels…

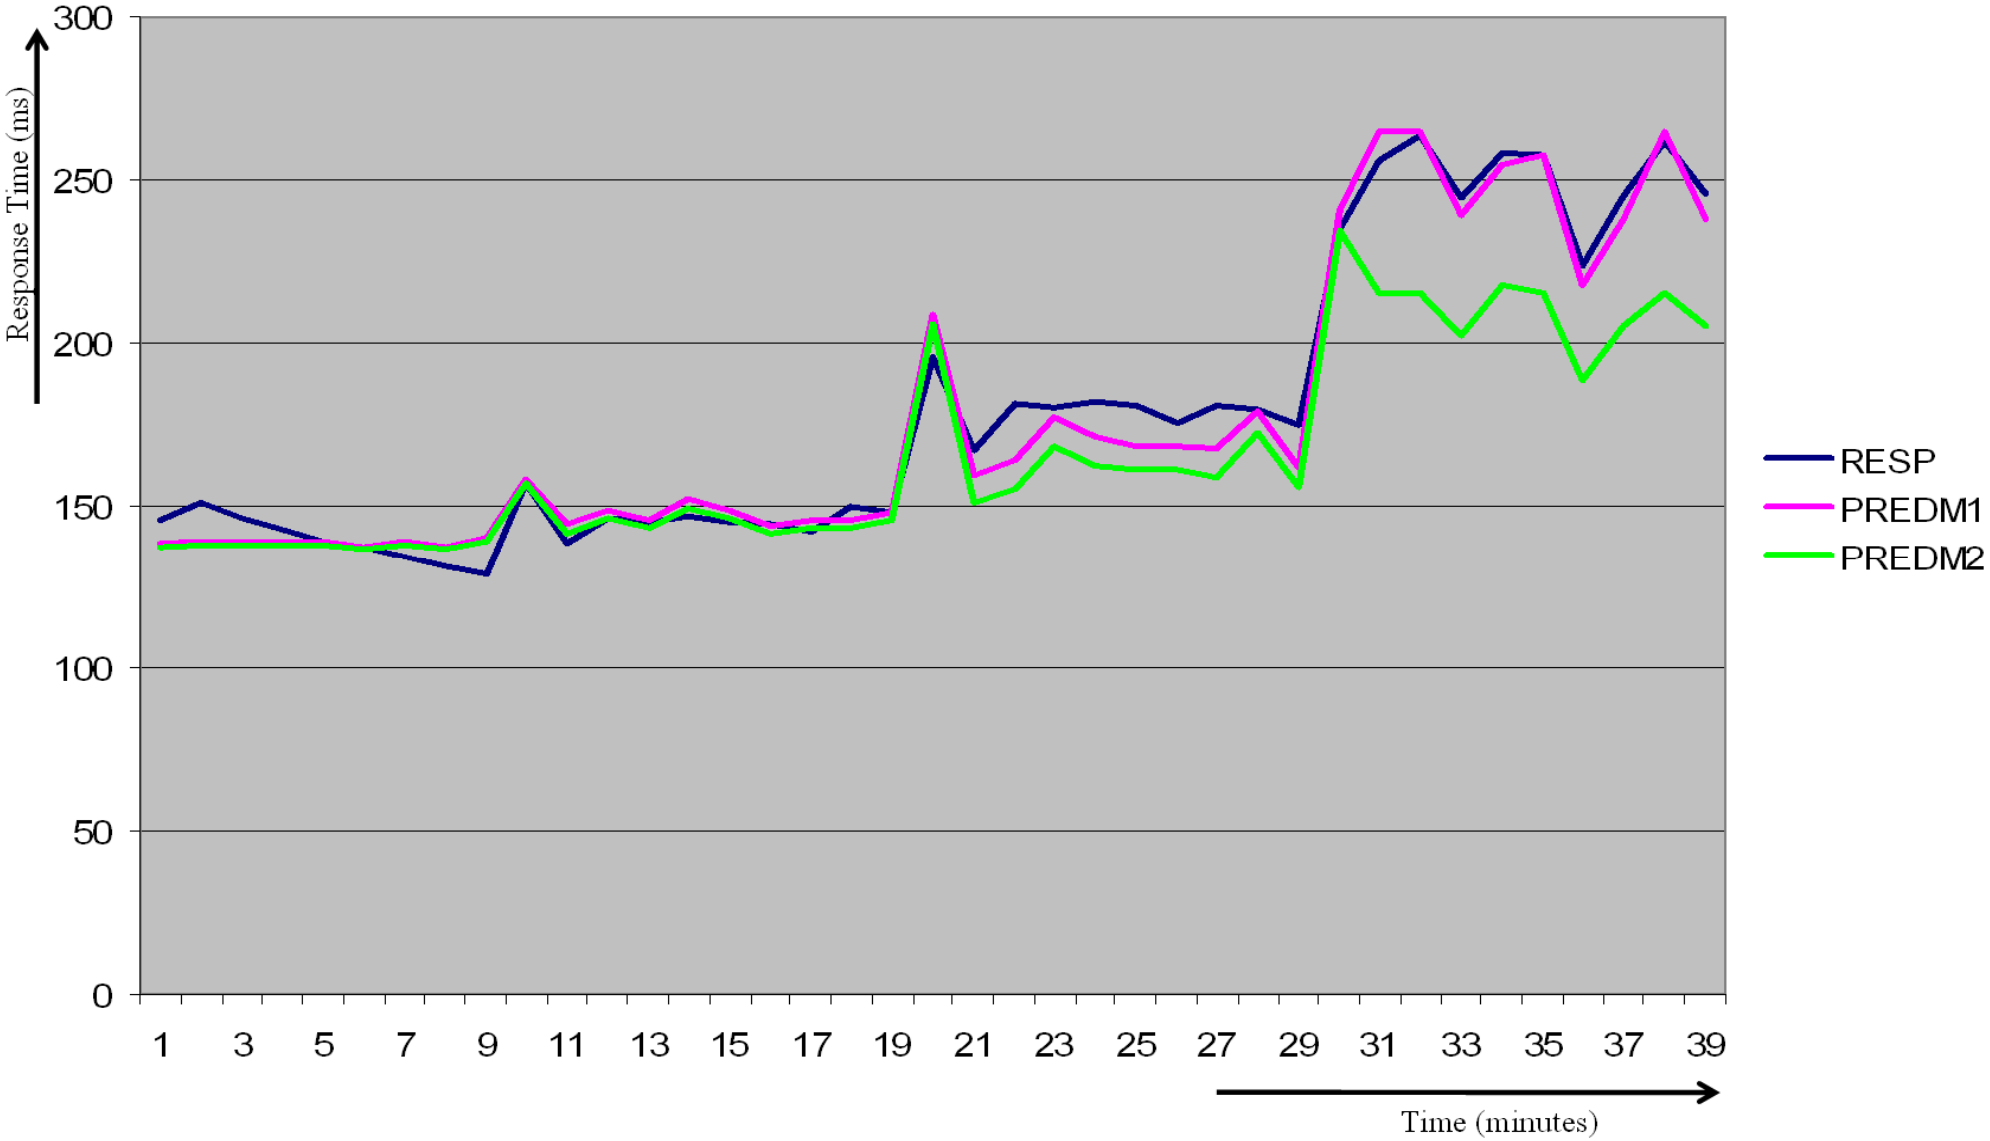

When 25% I/O Reduction…

Percentage logins (PCT_GT_5SEC) which took longer than 5 seconds:

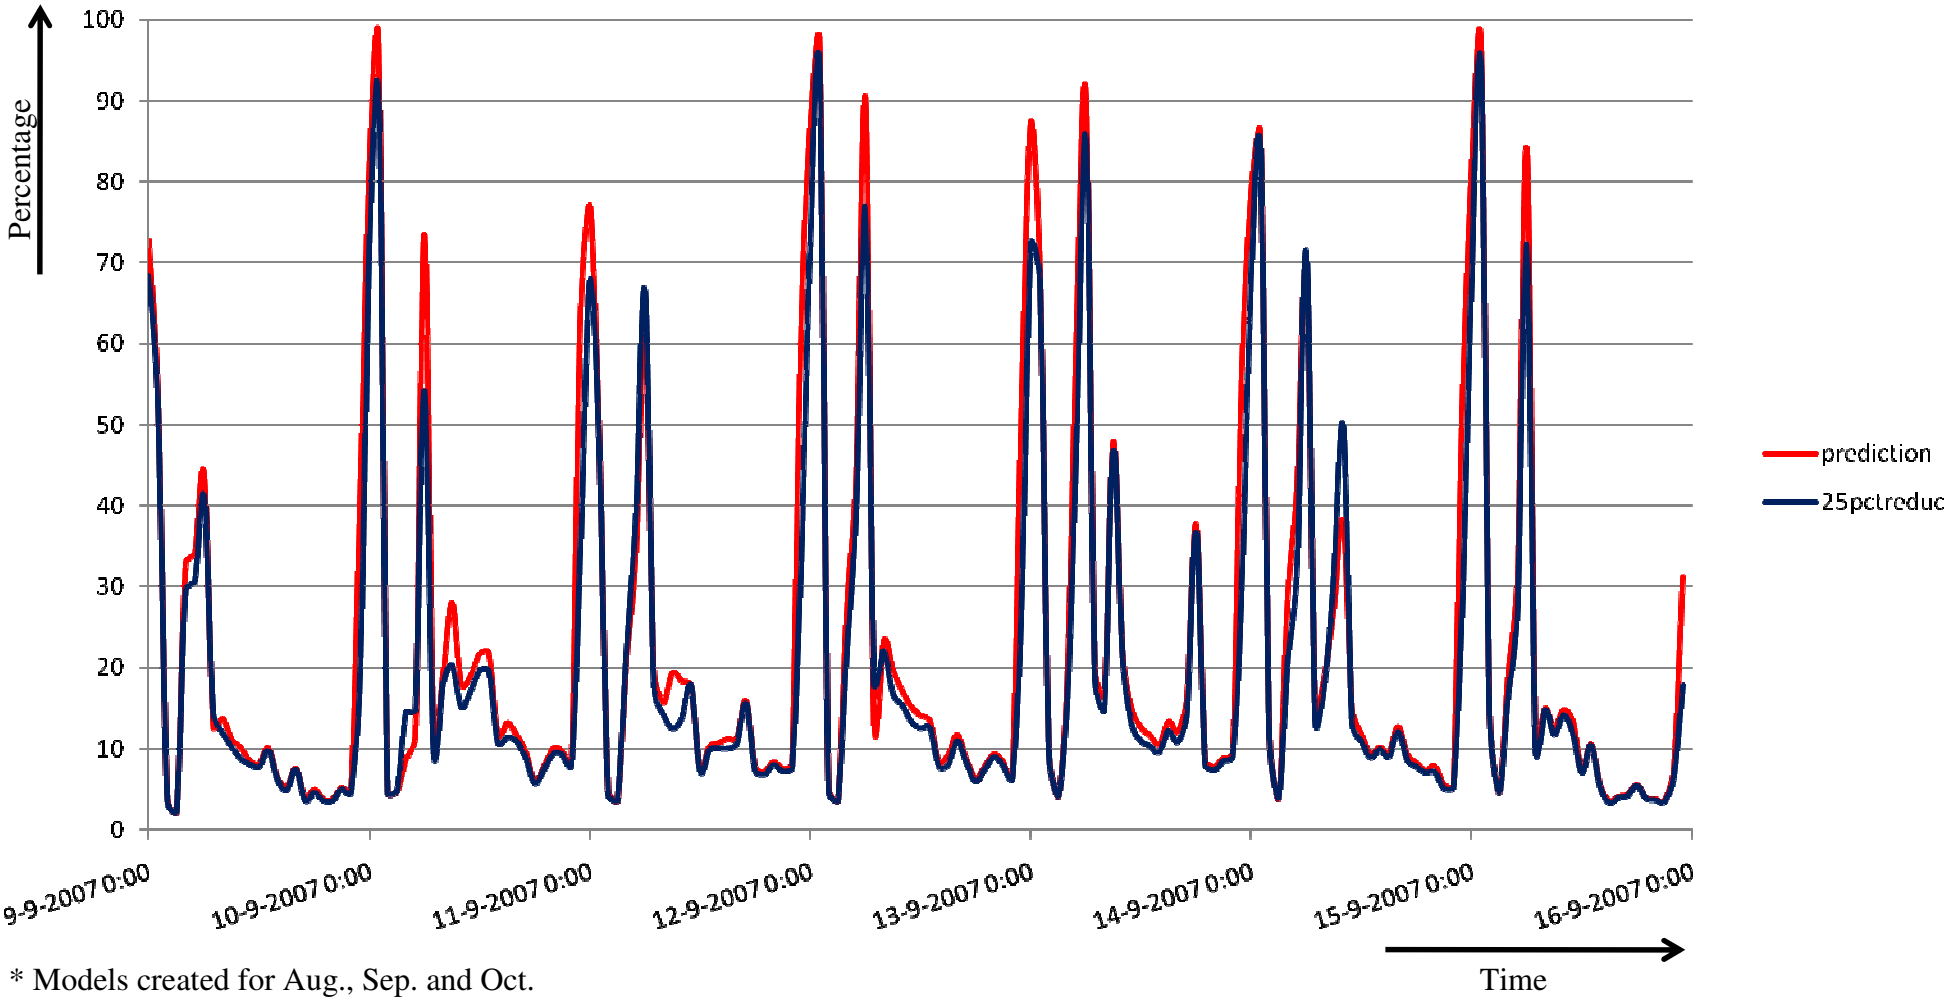

* Models created for Aug., Sep. and Oct.

The Fifth D

Data Interpretation

The Case…

- Response time of a very important business process in a time and labor system has sometimes very unpredictable performance

- The Used system is complex and the direct involved LPAR machines are:
  - HOT600 running several databases.
  - HOT720 running Portal and Oracle Internet Directory
  - HOT730 running Application Server for Time and Labor System

GAPP In A Complex Architecture…

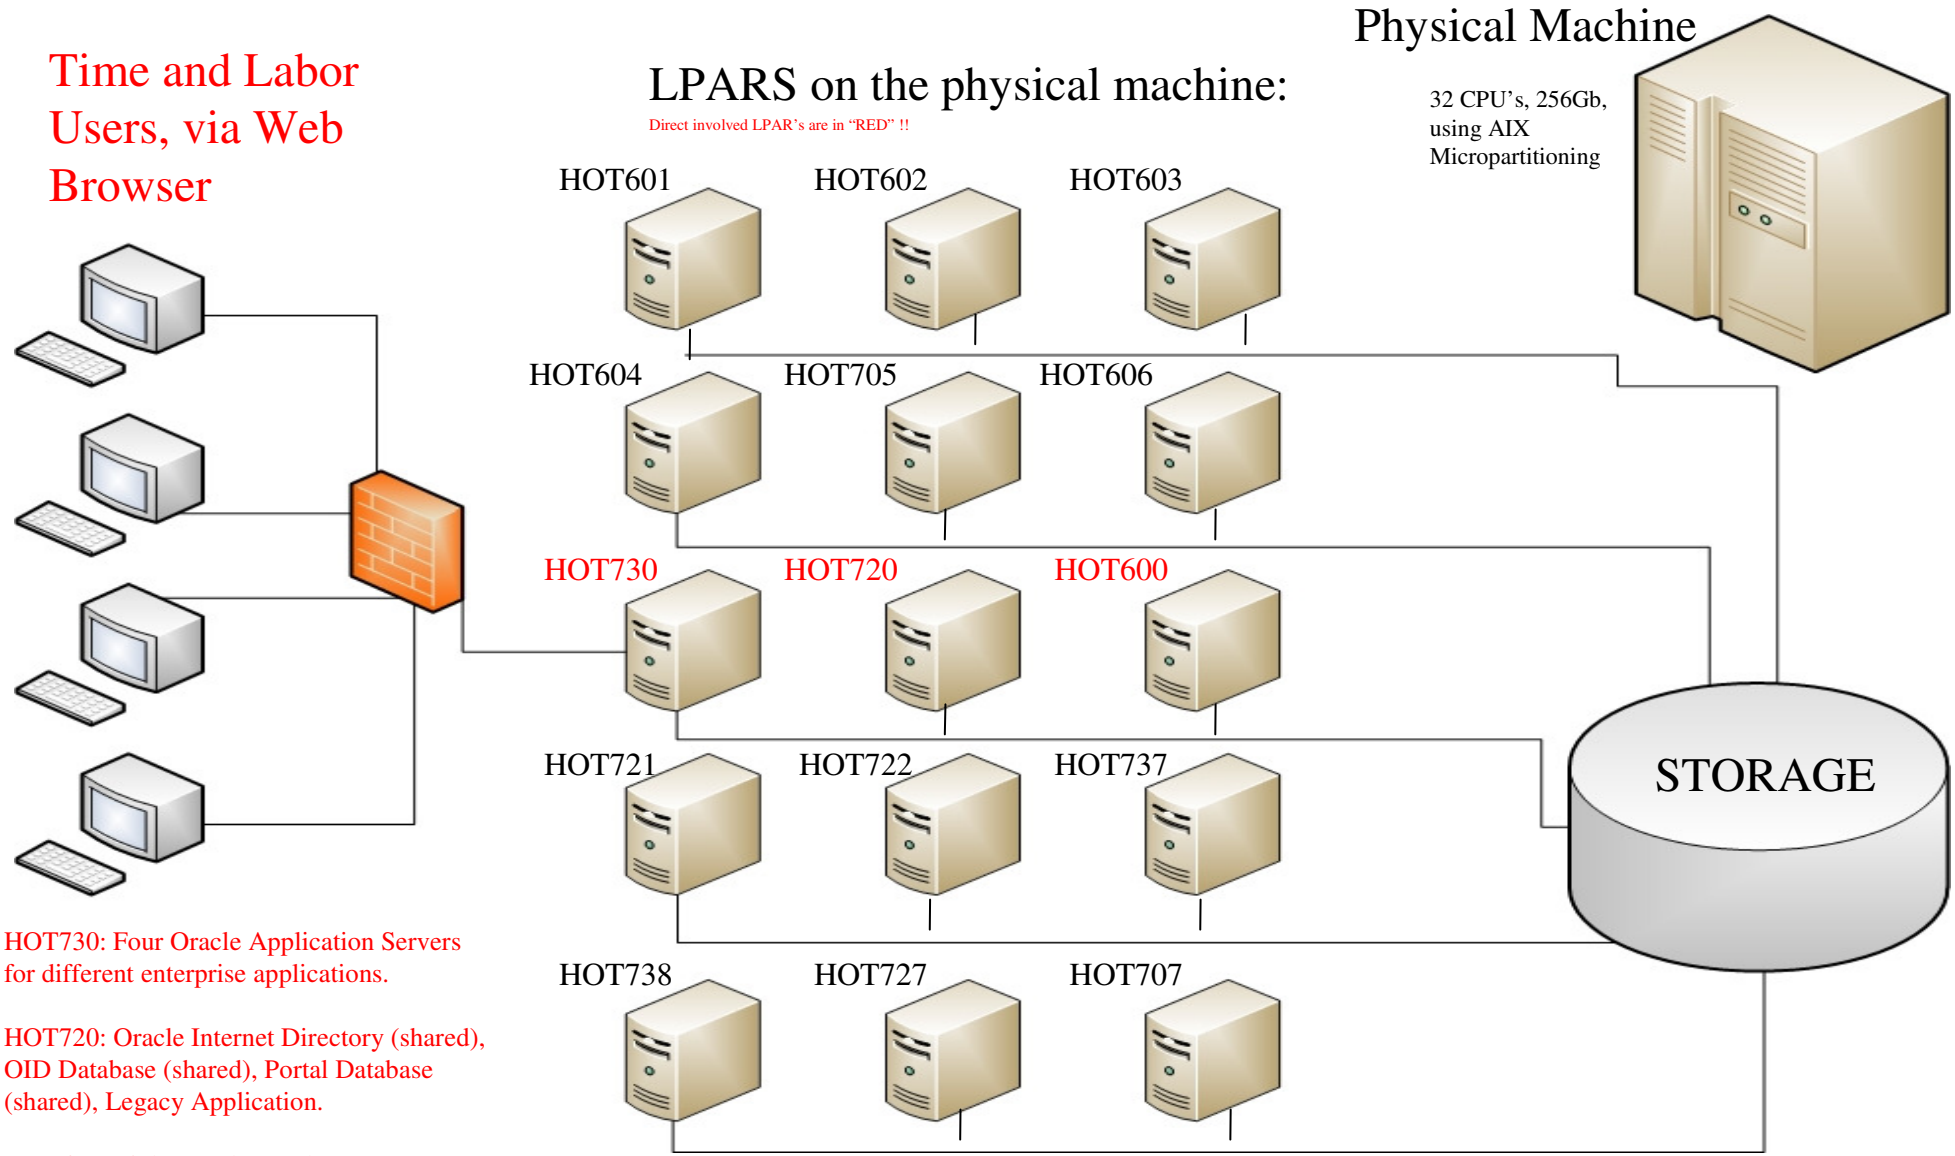

HOT730: Four Oracle Application Servers for different enterprise applications.

HOT720: Oracle Internet Directory (shared), OID Database (shared), Portal Database (shared), Legacy Application.

HOT600: Eight Oracle Databases

The Variance in Response Time…

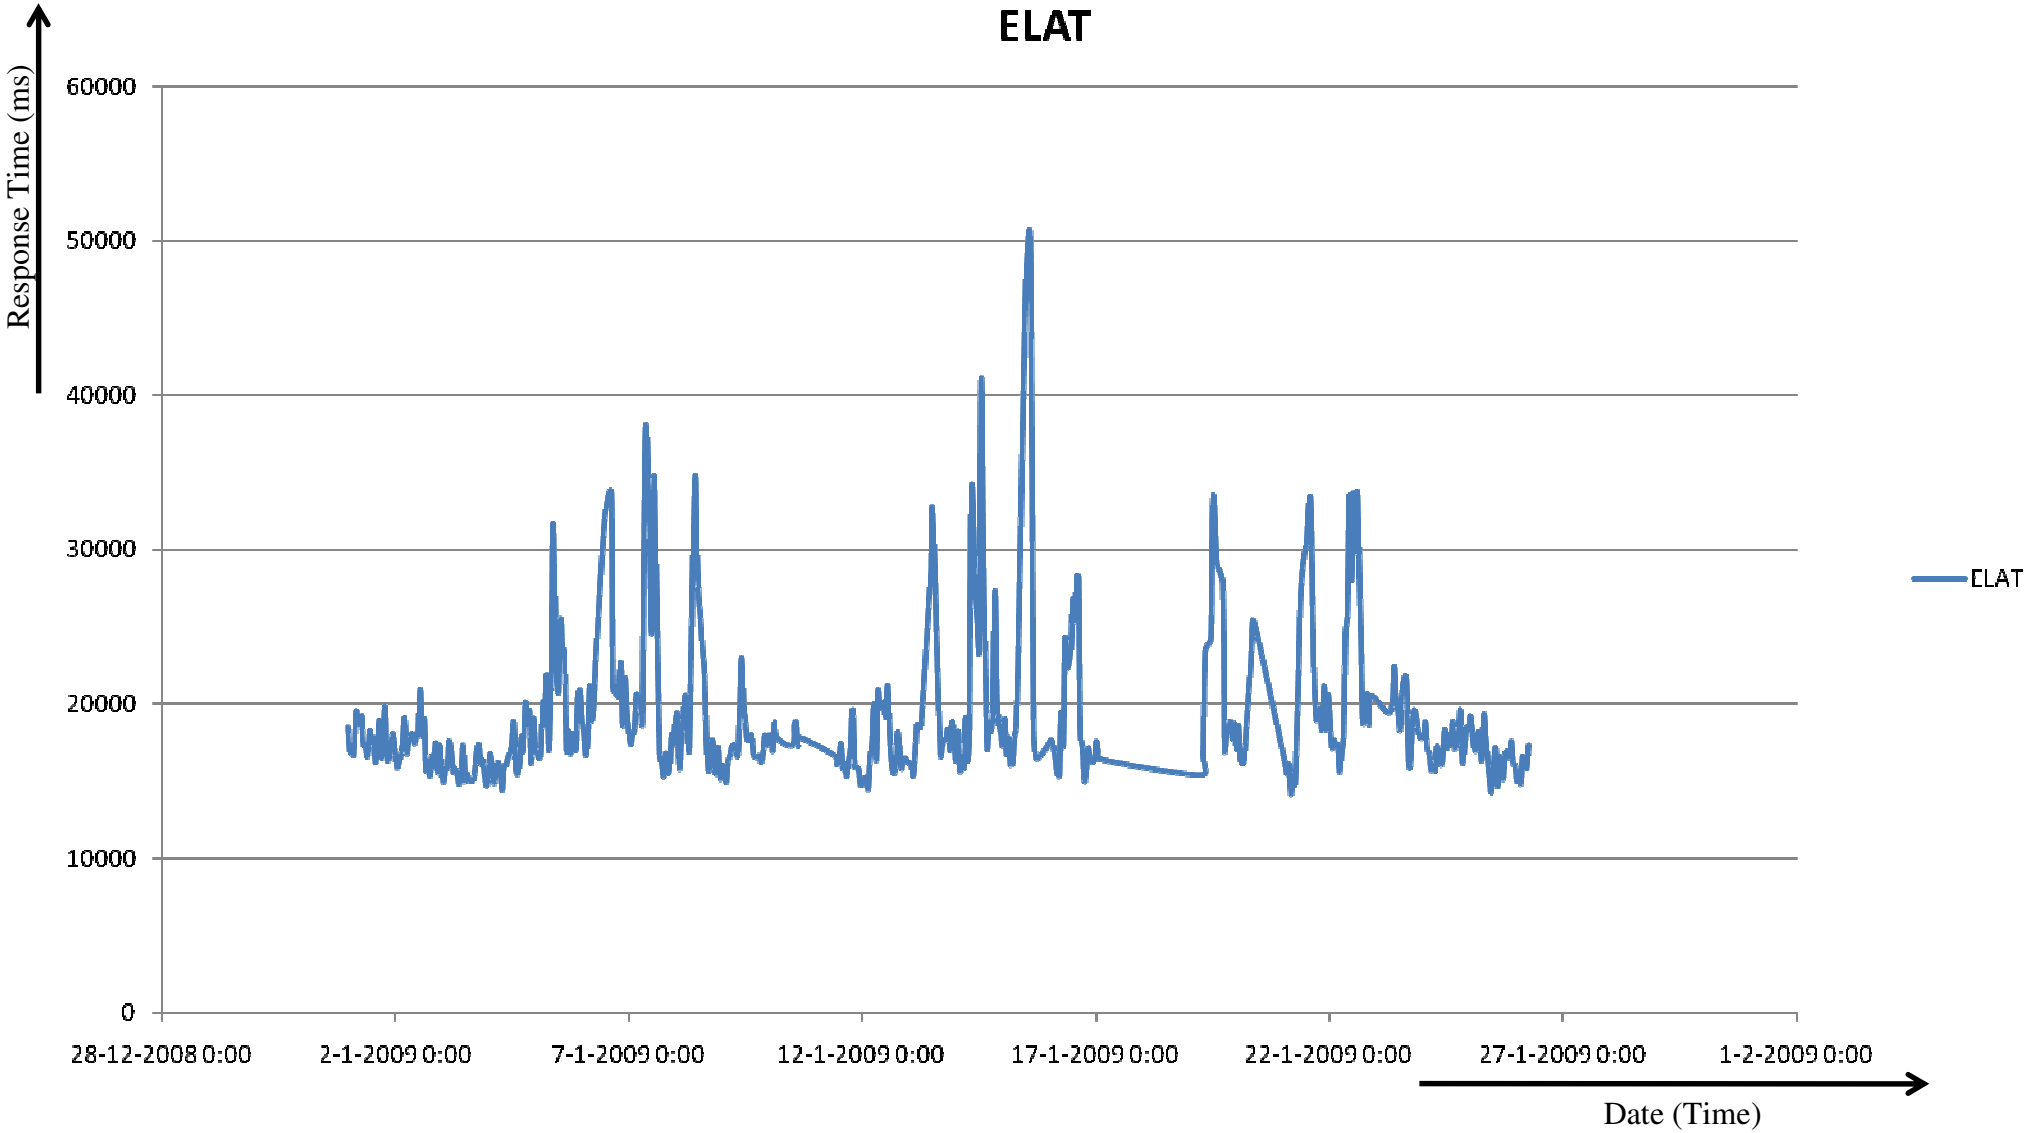

The Factorial Analyses (Explain)…

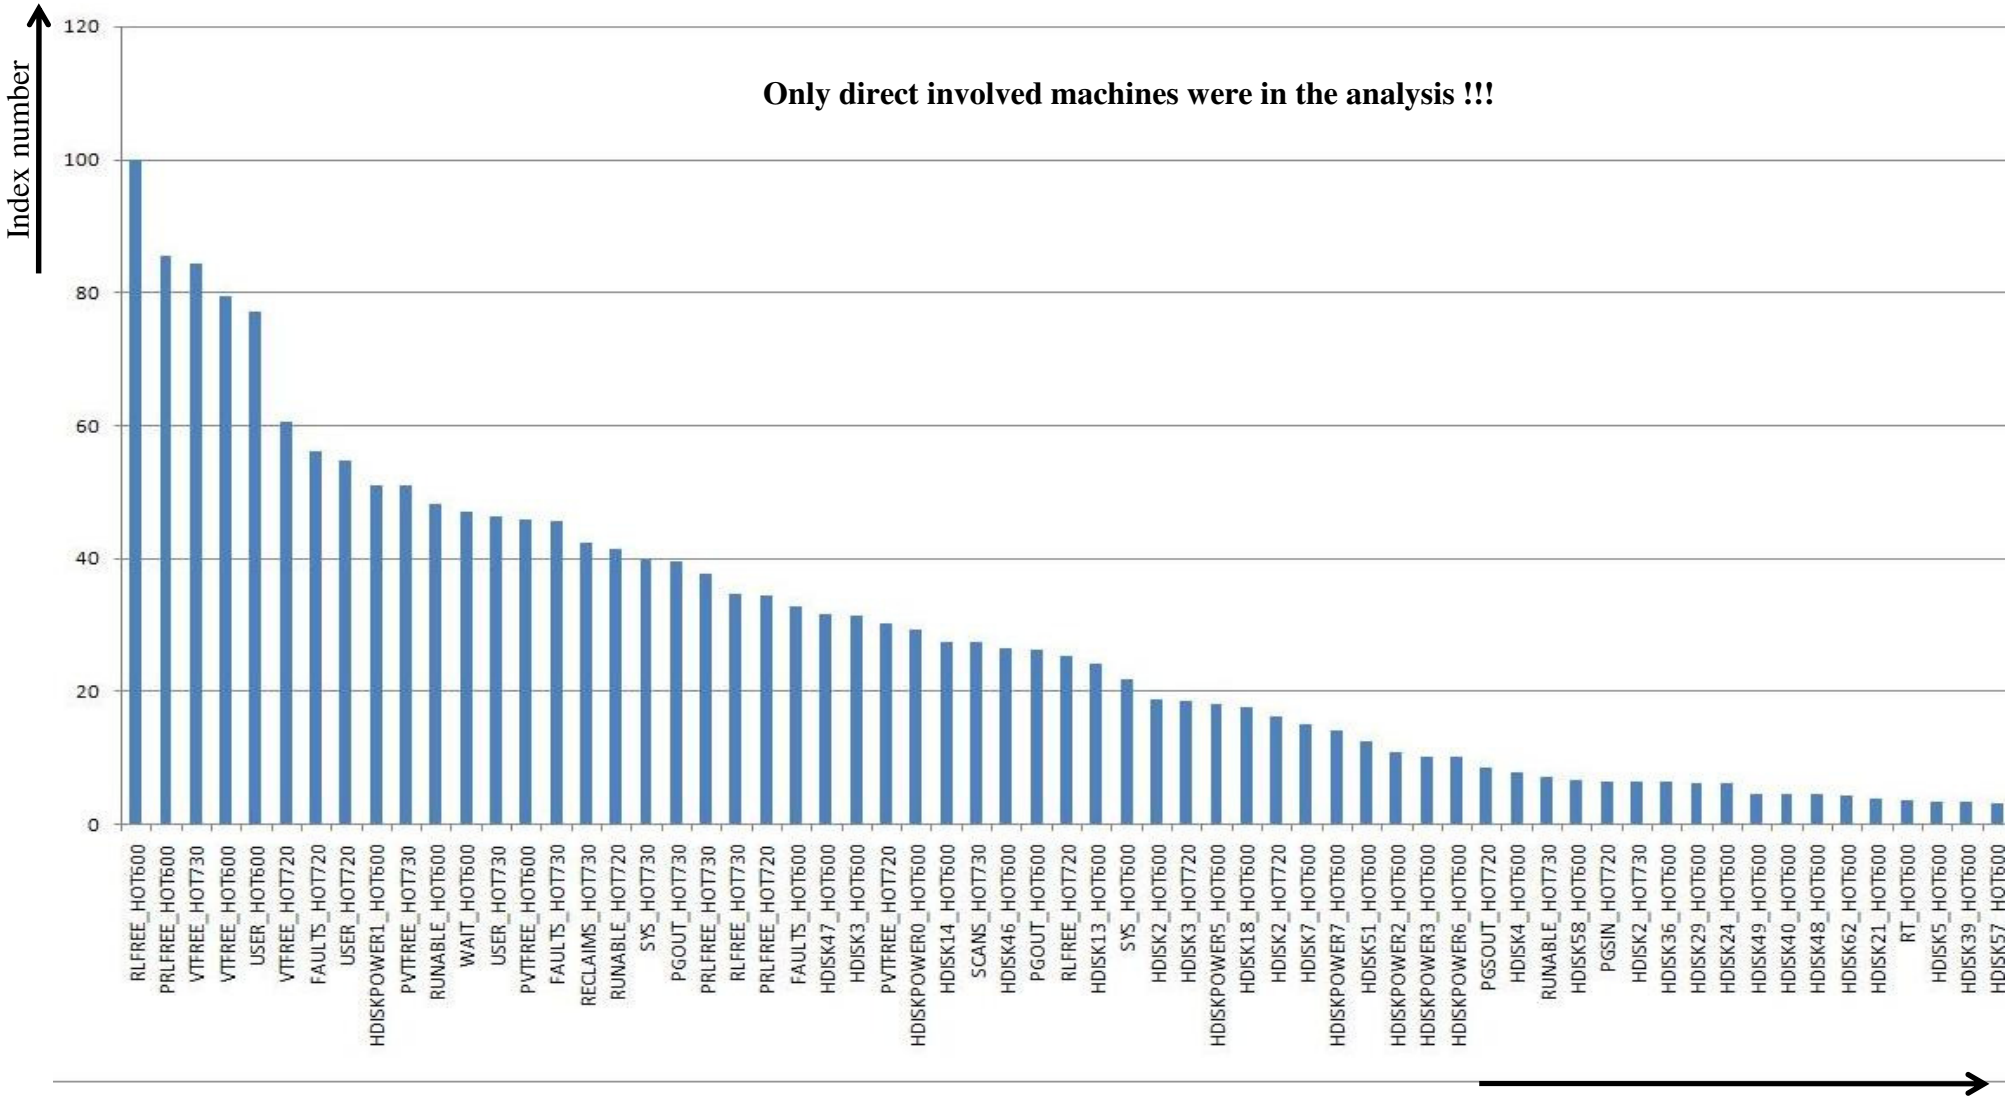

The Model (Predict)…

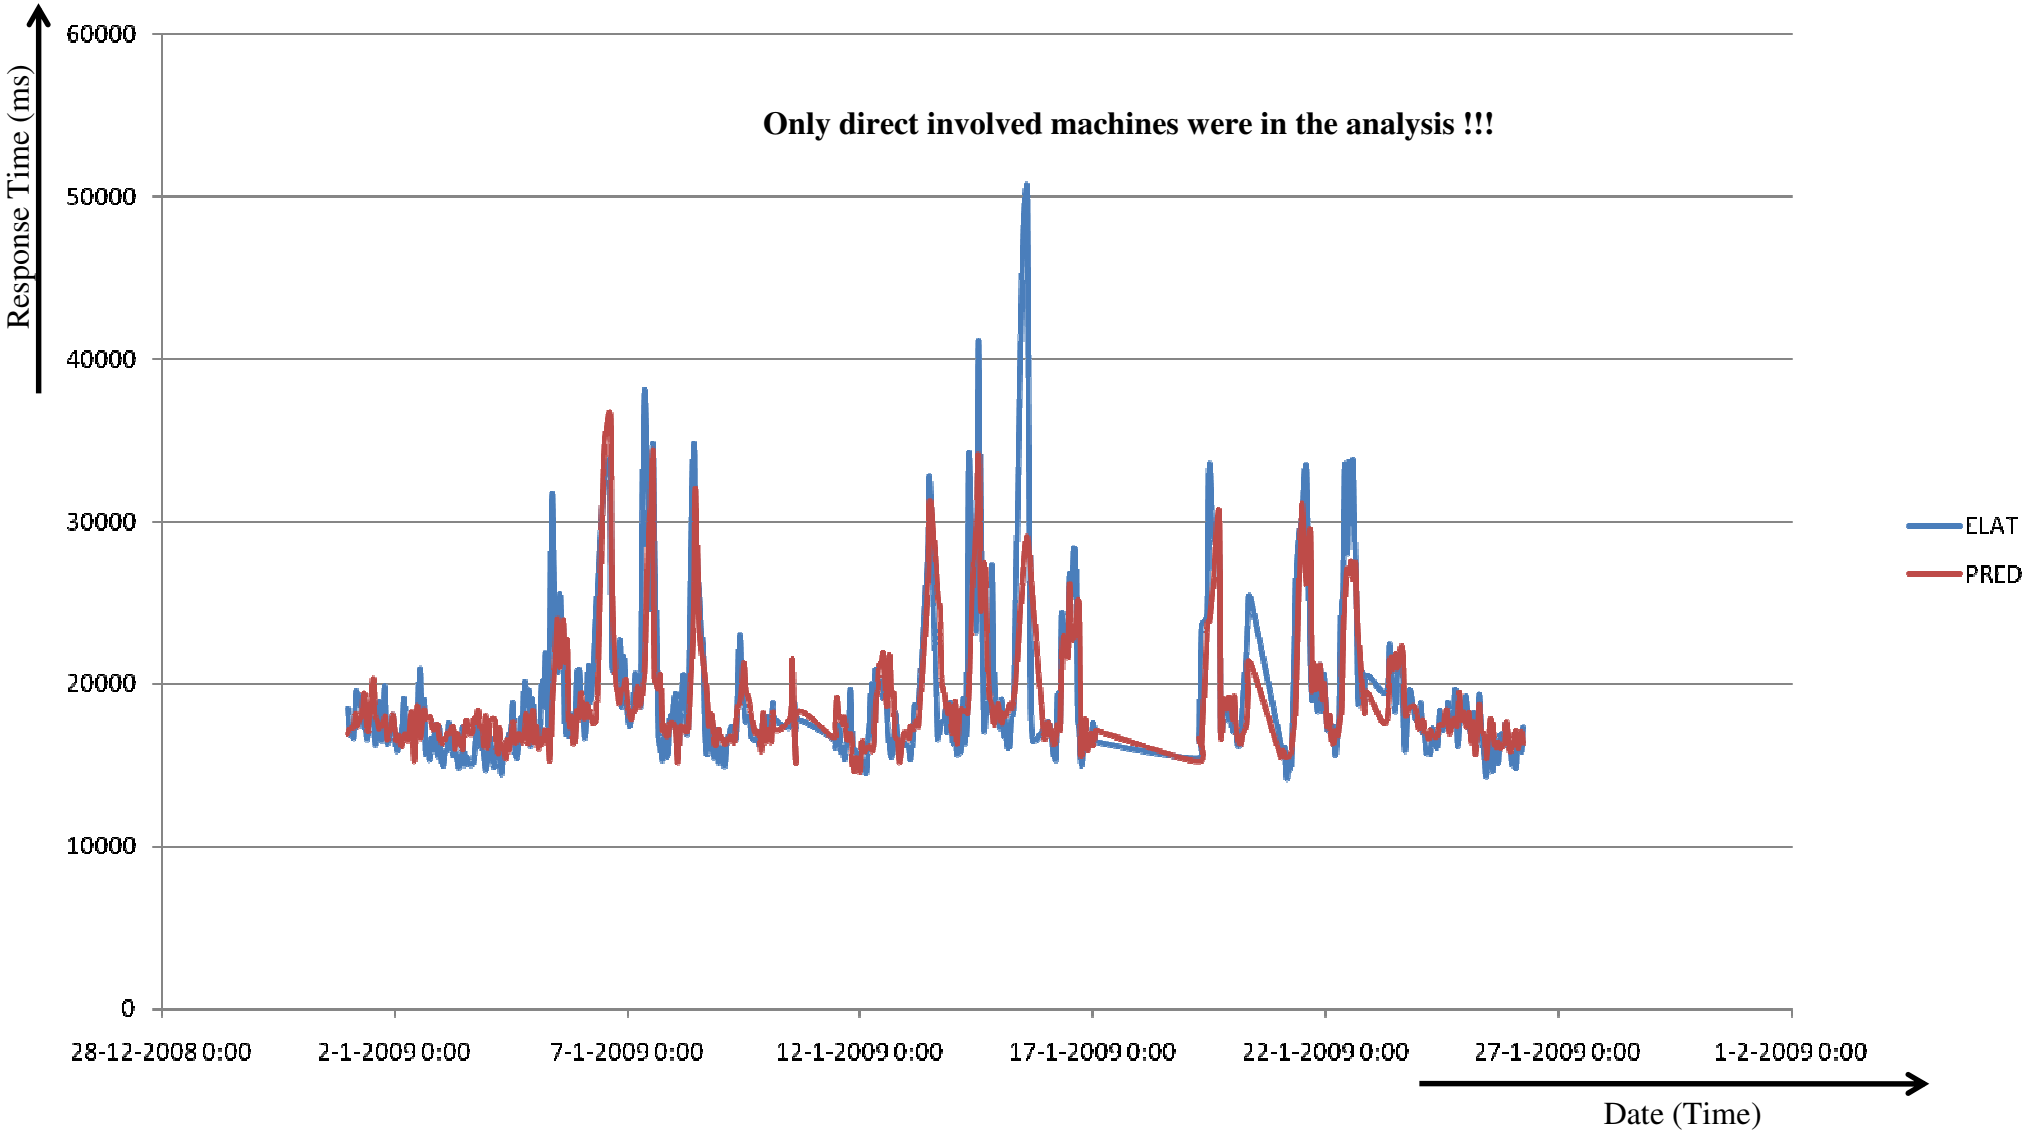

The Factorial Analyses (Explain)…

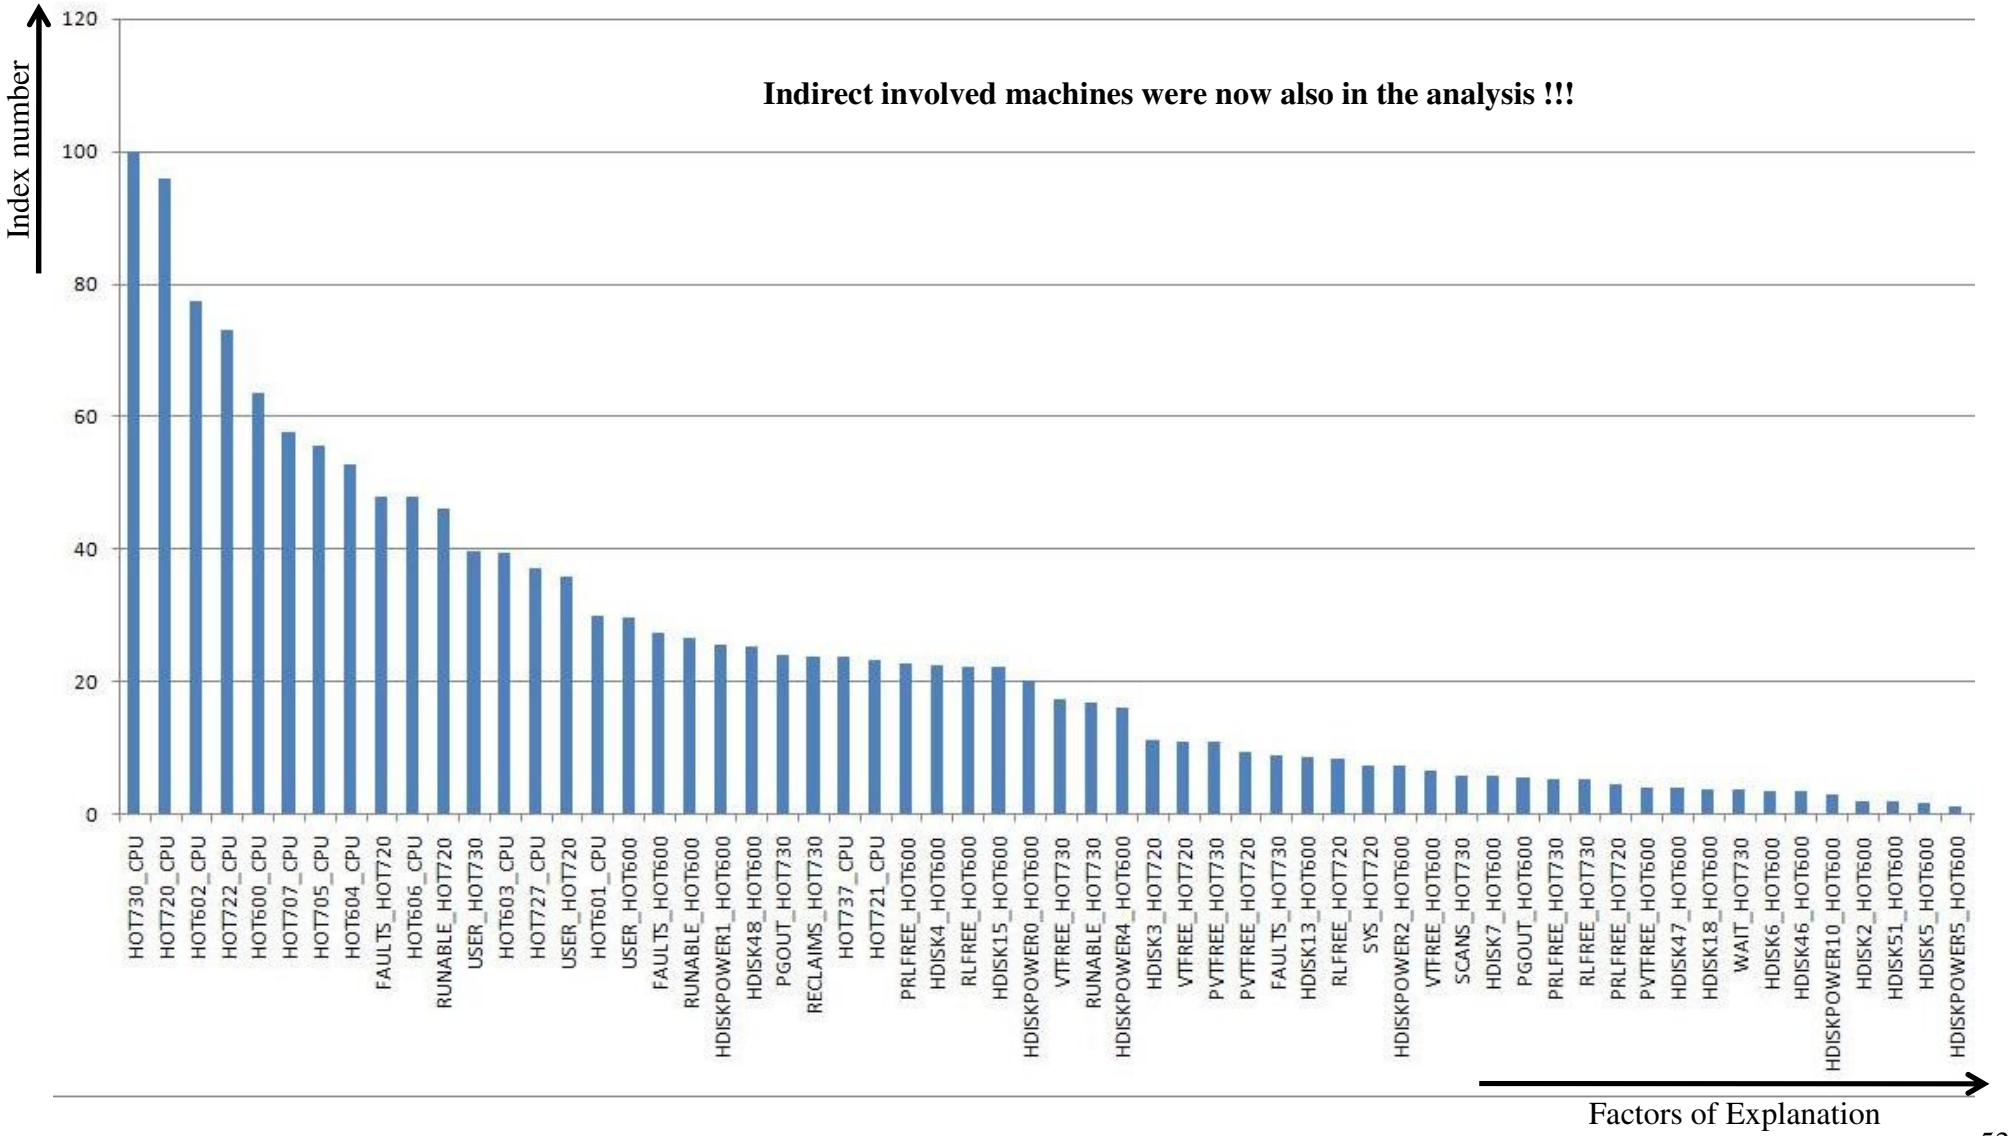

The Model (Predict)…

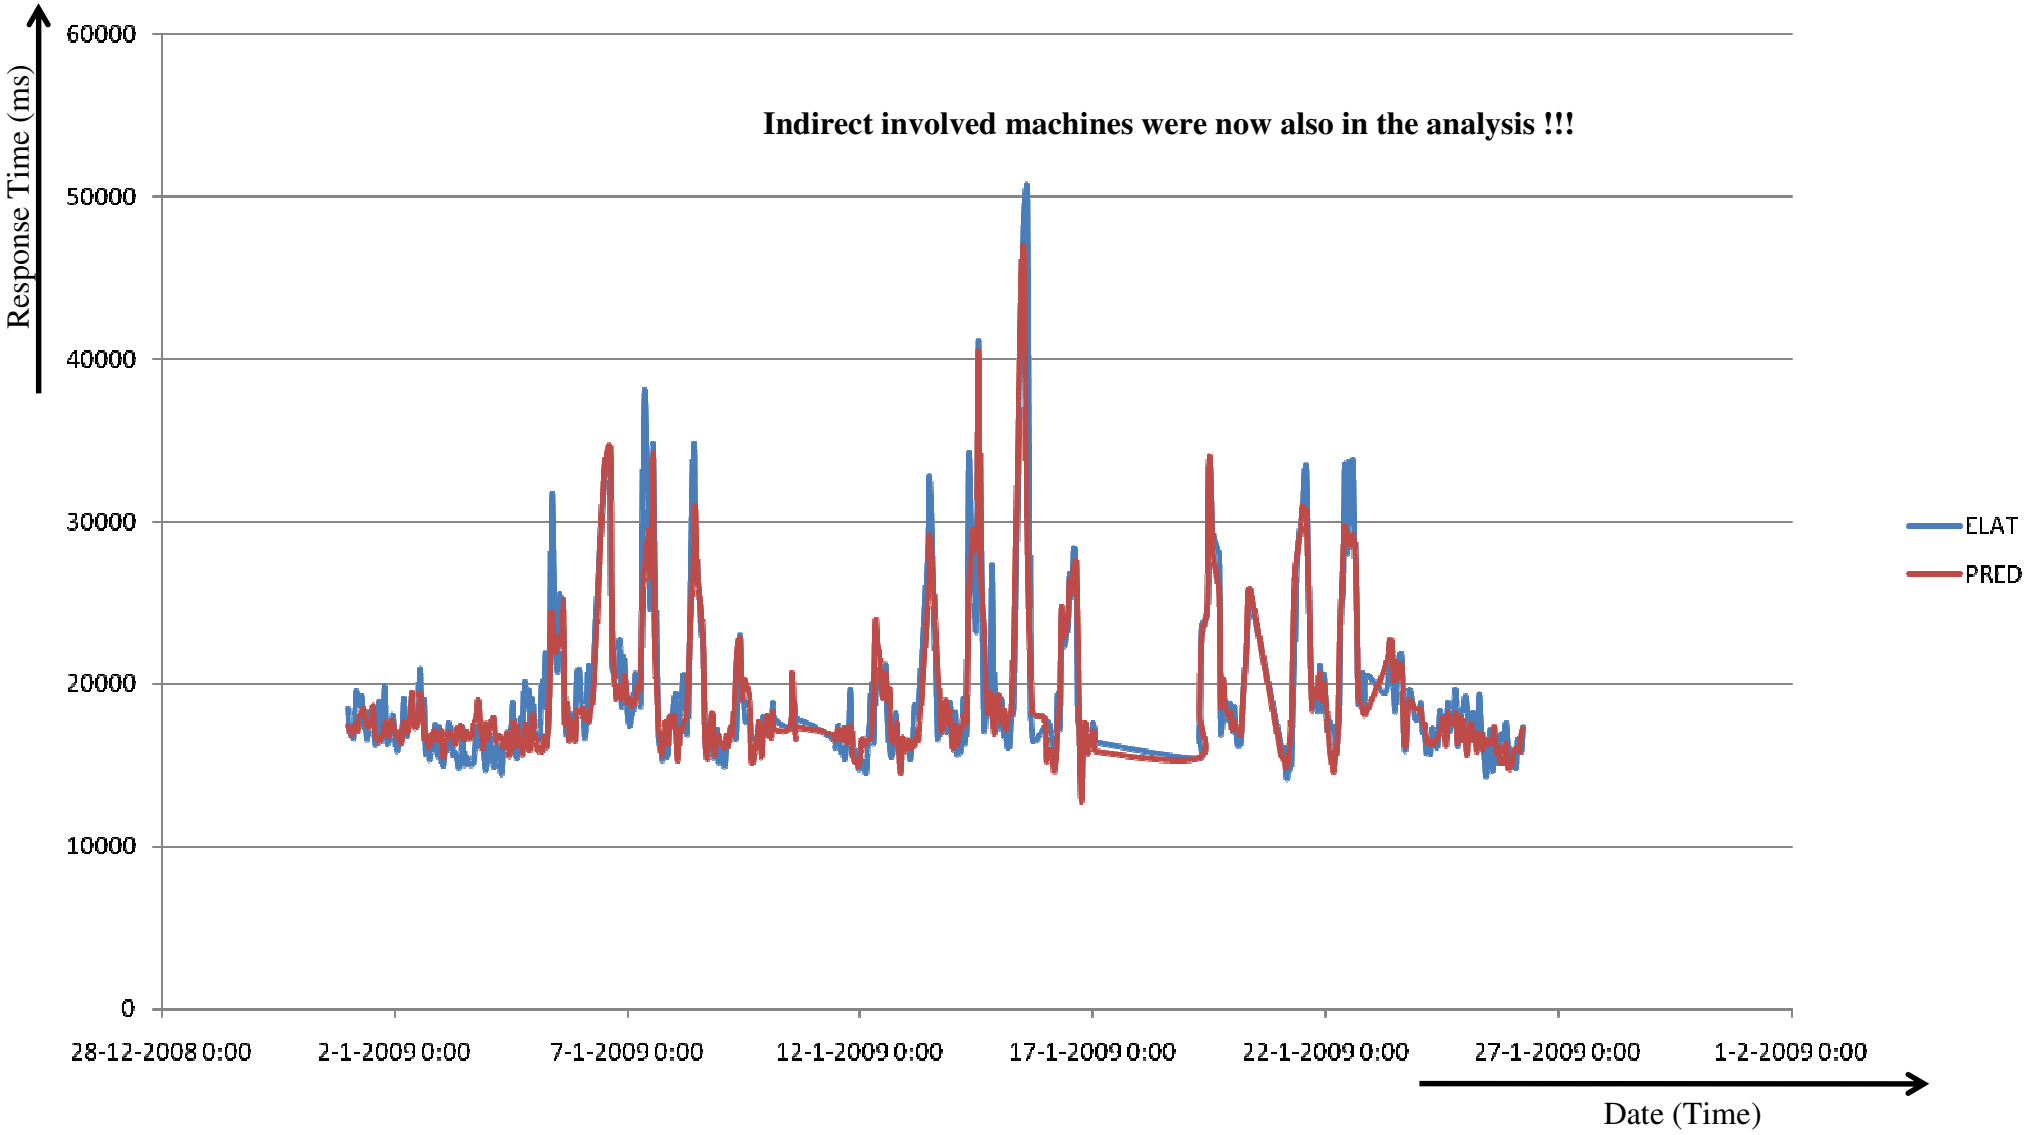

Interpretation…

- Most of the peaks are very well explained:
  - The performance of the involved business process was very depending on the machine's overall resources.
  - Beside the HOT730 and HOT720, the HOT602 and HOT722 had also much impact.

- The GAPP analyses showed that at the moment the physical machine gets short on physical Cpu all the different machines will have high impact on each other due to competition in hardware resources.
  - By Workload Scheduling the biggest problems could be solved

Conclusions

Conclusions…

- Although you didn't made "hooks" in the application you were able to find the problem, independent of technology.
- After your analyses was done you found out that also components in the infrastructure not part of the application were responsible for the encountered problem.
- The model you made, made it possible to do predictions of the impact of a possible investment on the end-user process performance.
- Adding secondary components data to the analyses can give a lot of extra detailed information regarding the faced performance problems.

Considerations…

- Adding as secondary components “variance” of different components can be used to detect skew with Method-GAPP

- Adding as secondary components “Business Activity Metrics” can make Method-GAPP the way to connect business data with technical data.

What could you use for the mining…

- Oracle Data Mining (ODM)
  - Cost 15K per CPU above EE
  - Makes method-GAPP dependent on ORACLE
  - Data can be stored in Oracle DB and be mined

- Project R:
  - Is Open Source
  - Makes method-GAPP completely platform independent
  - Needs Database (which could be open source)

Reference…

- Personal blog:
  - http://blog.gerwinhendriksen.com
  - http://method-gapp.com

- Method-R:
  - Oracle Performance by Cary Millsap (http://method-r.com/)

- Formula's:
  - Analyzing Computer Sys. Perf. By Dr. N.J. Gunther

- Oracle Data Mining:
  - http://www.oracle.com/technology/products/bi/odm/index.html

- Hadoop
  - http://hadoop.apache.org/

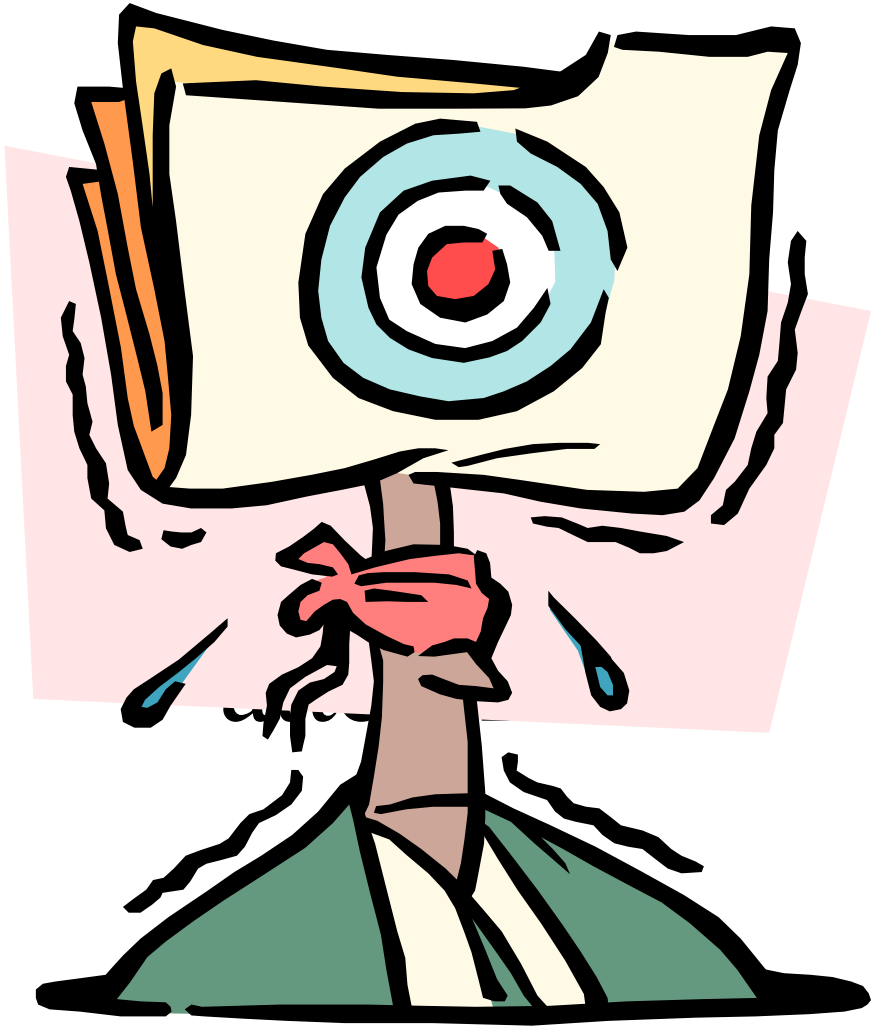

Q/A