

Questioning Method R

Cary Millsap (cary.millsap@hotsos.com)
Hotsos Enterprises, Ltd.
Northern California Oracle Users Group / San Francisco
11:00am–12:00n Thursday 2 November 2006



Agenda

- Introduction to Method R and profiling
- Questioning Method R

Method R and Profiling

*The only thing worse than having a slow system
is not being able to figure out why.*

Our goal was a problem diagnosis method with high scores in eight measurement categories.

- Diagnostic method goal categories
 - Relevance
 - Efficiency
 - Measurability
 - Predictive capacity
 - Reliability
 - Determinism
 - Finiteness
 - Practicality

Method R steps, take 1...

1. Select the user actions for which the business needs improved performance.
2. Collect properly scoped diagnostic data that will allow you to identify the causes of response time consumption for each selected user action while it is performing sub-optimally.
3. Execute the candidate optimization activity that will have the greatest net payoff to the business. If even the best net-payoff activity produces insufficient net pay-off, then suspend your performance improvement activities until something changes.
4. Go to step 1.

This is what's in the book.



Two minor adjustments have allowed Method R to apply directly in a preemptive context.

Reactive

- ...needs improved performance
- ...while it is performing sub-optimally

Preemptive

- ...that are critical
- ...while it is exhibiting the behavior you want to record

Method R steps, take 2...

1. Target the tasks that are critical to the business.
2. Collect properly scoped, un-aggregated profile data for each task while the task is exhibiting the behavior you want to record.
3. React with the candidate repair that will have the greatest net payoff to the business
 - a. Stop if the cost of the repair exceeds the cost of the problem.
4. Go to step 1.

What is a profile?

Call	Duration (seconds)	Call count	Avg. dur/call	
read	22.719	69.6%	2,198	0.010
write (async)	3.248	9.9%	1,348	0.002
compute_line	1.864	5.7%	1,348	0.001
other	4.820	14.8%		
Total	32.651	100.0%		

- Full accounting for the end-user response time
 - Spanning
 - Non-overlapping
- Partitioned along any dimension
 - By subroutine, SQL statement, dbcall, business process phase, ...

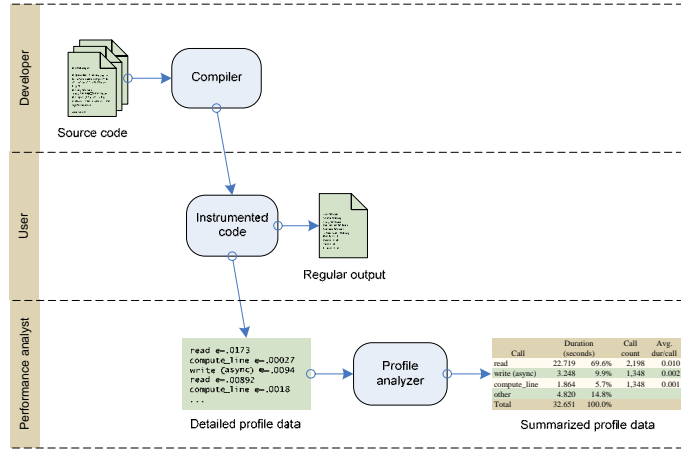
Why do you need a profile?

- Shows exactly where your time went
 - End-user experience is the **it**
 - Easy to see where the problem is (and is not)
 - Measurement is objective
 - Easy “what-if” analysis
 - Problems can’t hide from it
 - No guesswork
 - Shows when you can’t go any further
 - Fast and easy to make if you have the right tools

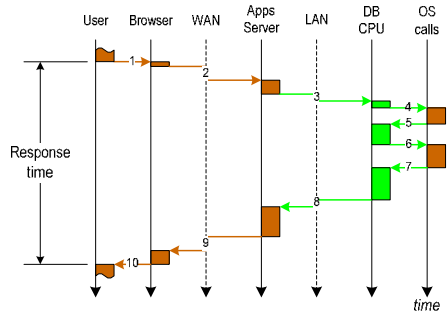
Profiling is a three-step process.

- The steps
 1. Compile and link with profiling enabled.
 2. Execute your program to generate a profile data file.
 3. Analyze the profile data.
- Examples
 - Knuth’s FORTRAN profiler (1971)
 - GNU *gprof* (1988)
 - Oracle SQL trace data and *tkprof* (1988)
 - Oracle extended SQL trace data (1992)

Where do profiles come from?



The Oracle extended SQL trace feature gives you everything you need to profile the Oracle tier.



Profiling the Oracle tier shows all the time consumed by interactions up and down from the Oracle tier.

Recap

- Method R is about finding out where your time is going
- A profiler tells you exactly where a task's time has gone
- Oracle emits everything you need to begin profiling
- Profilers give two capabilities you're not accustomed to having
 - “What happened?” becomes very easy
 - “What if?” becomes very easy

**Once you have these capabilities,
it's hard to imagine life without them.**

Questioning Method R

*What have we learned about performance management—and what we've learned about you—since we published *Optimizing Oracle Performance in 2003*?*

Need for a preventive method...

Method R is reactive. I don't just want a problem diagnosis method; I want a proactive method that keeps me from having problems to begin with.

- Step 1: Target the tasks that are critical to the business.
 - Tasks that could directly impede your business
 - Tasks that consume a resource you're trying to conserve
 - Tasks that run a long time
 - Tasks that run extremely frequently

What good is profiling during normal operations?

- Reasons to profile during periods of calm
 - Waste detection
 - Baselines to make diagnosis easy later
 - SLA management (trending)

Can't prioritize...

Can I use Method R in situations where the business can not or will not prioritize business tasks?

- Your business has priorities
 - ...whether or not you know what they are
- Don't obsess over priorities
 - Don't let the apparent "competition" make people nervous
 - It's okay for lots of tasks to have the same priority
 - But which one will get fixed first?
 - Fixing top tasks isn't going to hurt your bottom ones
 - You'll fix everything soon

Oracle trace data...

Do you have to use Oracle extended SQL trace data to use Method R?

- On the Oracle tier, yes
 - Other data sources conceal skew
- On other tiers, no

Dependence upon Hotsos products...

Do you have to have the Hotsos Profiler to use
Method R on Oracle databases?

- I've found that I do...
- Several problems with *tkprof* led us to create Hotsos Profiler
 - No measure of end-user experience
 - Summarize by subroutine or by cursor
 - Time consumed between dbcalls
 - Similar but distinct SQL
 - ...

Example: There's no way to know by looking at *tkprof*
that this task took 11.601 seconds.

OVERALL TOTALS FOR ALL NON-RECURSIVE STATEMENTS

call	count	cpu	elapsed	disk	query	current	rows
total	10003	2.25	4.79	0	5073	25778	5000

Elapsed times include waiting on following events:

Event waited on	Times Waited	Max. Wait	Total Waited
SQL*Net message to client	20007	0.03	0.06
SQL*Net message from client	20007	0.06	2.37
latch free	19	0.00	0.00
log file sync	3	0.01	0.01

OVERALL TOTALS FOR ALL RECURSIVE STATEMENTS

call	count	cpu	elapsed	disk	query	current	rows
total	250	0.01	0.09	2	310	4	118

Example: There's no way to know by looking at *tkprof* that this task took 11.601 seconds.

OVERALL TOTALS FOR ALL NON-RECURSIVE STATEMENTS

call	count	cpu	elapsed	disk	query	current	rows
total	10003	2.25	4.79	0	5073	25778	5000

Elapsed times include waiting on following events:

Event waited on	Times Waited	Max. Wait	Total Waited	
SQL*Net message to client	20007	0.03	0.06	2.37
SQL*Net message from client	20007	0.06	0.06	2.37
latch free	19	0.00	0.00	0.01
log file sync	3	0.01	0.01	0.01
				0.09
				7.32

OVERALL TOTALS FOR ALL RECURSIVE STATEMENTS

call	count	cpu	elapsed	disk	query	current	rows
total	250	0.01	0.09	2	310	4	118

Multi-tier applications...

Can I use Method R in my complicated multi-tier environment?

- To detect certain problem types, you must
- The trick is whether it will be easy to get the data
 - If you're writing the apps, you can make it easy
 - Some apps are easy (Oracle e-Business, PeopleSoft)
 - Excellent tools
 - *dbms_application_info*, *dbms_session*, *dbms_monitor*, *trcsess*
 - *strace*, *DTrace*, ... help

Data collection...

Is it practical to use detailed profile data as a primary diagnostic data source?

- You must choose some kind of compromise anyway
 - Trace data collected 24x7 compromises data management
 - Trace data collected intermittently compromises completeness
 - Aggregated data collected 24x7 compromises data quality
- Best compromise
 - Regular operational trace data for top priority tasks
 - Aggregations 24x7 too

Prepackaged applications...

Can I use Method R with prepackaged applications?
(If we have no control over the application source code, then why should we bother diagnosing it?)

- Your case is one of these
 - Your problem does not require the vendor's participation
 - The vendor needs to fix something
 - It's as fast as it's going to get until you upgrade something
- What category are you in?
 - You can't know unless you do a little diagnosis

Session interdependencies...

Does Method R ignore session interdependencies?

- Session interdependencies can't hide from a correct profile
 - Certain red-flag subroutines (*latch free, enqueue, ...*)
 - Unusually long latencies

Subroutine	Duration (seconds)	Call count	Average duration per call	
<i>latch: library cache</i>	4.827	99%	27	0.179
all other	0.042	1%	44	0.001
Total	4.869	100%		

Subroutine	Duration (seconds)	Call count	Average duration per call	
<i>db file sequential read</i>	4.000	98%	1	4.000
all other	0.1	2%	20	0.005
Total	4.1	100%		

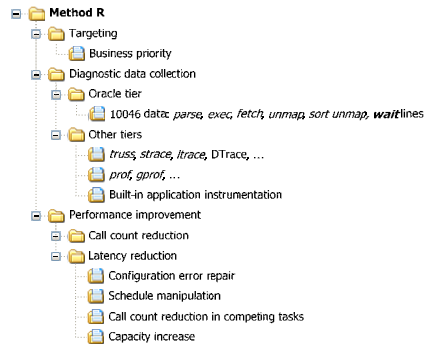
Non-Oracle problems...

Can Method R help me diagnose performance problems that are outside the Oracle tier?

- The trace data instruments the interactions between Oracle and its adjacent tiers
 - That's where many application performance problems show up

Oracle wait interface...

Is Method R just a different name for using the Oracle wait interface?



YAPP Method...

Is Method R another name for the YAPP method?

- YAPP is exclusively v\$ based
 - Oracle-specific
 - Susceptible to skew
- The lasting value of YAPP
 - First established response time as an analytical objective
 - Thus, Anjo Kolk is the father of modern Oracle performance analysis
 - “Breaking down wait time” is still the best there is

Bottleneck analysis...


Is Method R consistent with principles of bottleneck analysis?

- “Bottleneck is the component where a transaction spends most of its time” [Menascé and Almeida (2002), p10]
- Thus, a task’s bottleneck is its profile’s top line
 - Subroutine
 - Cursor
 - Dbcall
 - ...

Bottleneck analysis, system analyst’s perspective (1)...

- Definitions
 - Resource with the highest utilization is the bottleneck [Jain (1991), p34]
 - Utilization of a resource is measured as the fraction of time the resource is busy servicing requests [Jain (1991), p39]

Bottleneck analysis, system analyst's perspective (2)...




- So... what's the bottleneck here?
 - Red, by simple inspection
 - Blue, if Gavin over [1,4] is top priority
 - Gold, if Nikolas over [4,10] is top priority
 - Anything you want?

You don't get to choose the top-priority task and time interval; the business does.

References

- See also
 - *Optimizing Oracle Performance* (Millsap & Holt, O'Reilly 2003)
 - "Questioning Method R" at www.hotsos.com



Discussion

Thank you