Using the usl package

Analyze System Scalability in R with the Universal Scalability Law

Stefan Möding

February 29, 2020

The Universal Scalability Law is used to quantify the scalability of hardware or software systems. It uses sparse measurements from an existing system to predict the throughput for different loads and can be used to learn more about the scalability limitations of the system. This document introduces the usl package for R and shows how easily it can be used to perform the relevant calculations.

Contents

1 Version

This document describes version 3.0.0 of the usl package.

2 Introduction

Every system architect faces the challenge to deliver an application system that meets the requirements. A critical point during the design is the scalability of the system.

Informally scalability can be defined as the ability to support a growing amount of work. A system is said to scale if it handles the changing demand or hardware environment in a reasonable efficient and practical way.

Scalability can have two facets with respect to a computer system. On the one hand, there is software scalability where the focus is about how the system behaves when the demand increases, i.e., when more users are using it or more requests need to be handled. On the other hand, there is hardware scalability where the behavior of an application system running on larger hardware configurations is investigated.

The Universal Scalability Law (USL) has been developed by Dr. Neil J. Gunther to allow the quantification of scalability for the purpose of capacity planning. It provides an analytic model for the scalability of a computer system.

A comprehensive introduction to the Universal Scalability Law including the mathematical grounding has been published by Dr. Gunther (cf. [1]).

3 Background

Dr. Gunther shows how the scalability of every computer system can be described by a common rational function. This function is *universal* in the sense that it does not assume any specific type of software, hardware or system architecture.

Equation (1) illustrates the original Universal Scalability Law where $C(N) = X(N)/X(1)$ is the relative capacity given by the ratio of the measured throughput $X(N)$ for load N to the throughput $X(1)$ for load 1.

$$C(N) = \frac{N}{1 + \alpha(N-1) + \beta N(N-1)} \tag{1}$$

The denominator consists of three terms that all have a specific physical interpretation:

Concurrency: The first term models linear scalability that would exist if the different parts of the system (processors, threads …) could work without any interference caused by their interaction.

Contention: The second term of the denominator refers to the contention between different parts of the system. Most common are issues caused by serialization or queueing effects.

Coherency: The last term represents the delay induced by keeping the system in a coherent and consistent state. This is necessary when writable data is shared in different parts of the system. Predominant factors for such a delay are caches implemented in software and hardware.

In other words: α and β represent two concrete physical issues that limit the achievable speedup for parallel execution. Note that the contention and coherency terms grow linearly respectively quadratically with N. As a consequence their influence becomes larger with an increasing N.

Due to the quadratic characteristic of the coherency term there will be a point where the throughput of the system will start to go retrograde, i.e., will start to decrease with further increasing load.

Dr. Gunther proved that eq. (1) is reduced to Amdahl's Law for $\beta = 0$ (cf. [1]). Therefore the Universal Scalability Law can be seen as a generalization of Amdahl's Law for speedup in parallel computing.

We can solve this nonlinear equation to estimate the coefficients α and β using a sparse set of measurements for the throughput X_i at different loads N_i. Initially the Unversal Scalabilty Law normalized the measurements using the throughput $X(1)$ for load 1 as the point of reference. This placed an additional burden on the performance analyst who needed to either measure the value or find a smart way to estimate it.

It was then discovered that a third coefficient γ can be added to the formula to represent the missing point of reference (cf. [3]). Equation (2) shows the updated Universal Scalability Law using three parameters.

$$X(N) = \frac{\gamma N}{1 + \alpha(N-1) + \beta N(N-1)} \tag{2}$$

Nonlinear regression will solve this equation and find best-fit values for the three parameters α, β and γ.

The usl package has been created to subsume the computation into one simple function call. This greatly reduces the manual work that previously was needed to perform the scalability analysis.

Initially the coefficients are called σ and κ when hardware scalability is evaluated but α and β when software scalability is analyzed. Up to version 1.8.0 the usl package always used `sigma` and `kappa` as coefficients. This has been changed starting with the 2.0.0 release of the usl package. Now the package will also use `alpha` and `beta` as coefficient names to follow Dr. Gunther's nomenclature. Additionally the `gamma` parameter was added to implement the Universal Scalability Law with three parameters.

4 Examples of Scalability Analysis

The following sections present some examples of how the usl package can be used when performing a scalability analysis. They also explain typical function calls and their arguments.

4.1 Case Study: Hardware Scalability

The usl package contains a demo dataset with benchmark measurements from a raytracer software[1]. The data was gathered on an SGI Origin 2000 with 64 R12000 processors running at 300 MHz.

A number of reference images with different levels of complexity were computed for the benchmark. The measurements contain the average number of calculated ray-geometry intersections per second for the number of used processors.

It is important to note that with changing hardware configurations the relative number of *homogeneous* application processes per processor is to be held constant. So when k application processes were used for the N processor benchmark then $2k$ processes must be used to get the result for $2N$ processors.

Start the analysis by loading the usl package and look at the supplied dataset.

```
R> library(usl)
R> data(raytracer)
R> raytracer

   processors throughput
1           1         20
2           4         78
3           8        130
4          12        170
5          16        190
6          20        200
7          24        210
8          28        230
9          32        260
10         48        280
11         64        310
```

The data shows the throughput for different hardware configurations covering the available range from one to 64 processors. We can easily see that the benefit for switching from one processor to four processors is much larger than the gain for upgrading from 48 to 64 processors.

Create a simple scatterplot to visualize the raw data.

---

[1]http://sourceforge.net/projects/brlcad/

```
R> plot(throughput ~ processors, data = raytracer)
```

Figure 1 shows the throughput of the system for the different number of processors. This plot is a typical example for the effects of *diminishing returns,* because it clearly shows how the benefit of adding more processors to the system gets smaller for higher numbers of processors.


Figure 1: Measured throughput of a raytracing software in relation to the number of available processors

Our next step builds the USL model from the dataset. The `usl()` function creates an S4 object that encapsulates the computation.

The first argument is a formula with a symbolic description of the model we want to analyze. In this case we would like to analyze how the "throughput" changes with regard to the number of "processors" in the system. The second argument is the dataset with the measured values. Note how this call matches the syntax of the `plot()` function.

```
R> usl.model <- usl(throughput ~ processors, data = raytracer)
```

The model object can be investigated with the `summary()` function.

```
R> summary(usl.model)


Call:
usl(formula = throughput ~ processors, data = raytracer)
```

```
Efficiency:
   Min     1Q Median     3Q    Max
 0.222  0.374  0.458  0.696  0.915

Residuals:
   Min     1Q Median     3Q    Max
-15.18  -5.30   2.71   7.07   9.69

Coefficients:
        Estimate  Std. Error  t value   Pr(>|t|)
alpha   0.057771    0.013293     4.35     0.0025  **
beta    0.000000    0.000118     0.00     1.0000
gamma  21.848843    2.196165     9.95  0.0000088  ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 9.34 on 8 degrees of freedom

Scalability bounds:
limit: throughput 378 (Amdahl asymptote)
peak:  none (beta=0)
opt:   throughput 195 at processors 17.3
```

The output of the `summary()` function shows different types of information.

- First of all it includes the call we used to create the model.

- The efficiency tells us something about the ratio of useful work that is performed per processor. It is obvious that two processors might be able to handle twice the work of one processor but not more. Calculating the ratio of the workload per processor should usually be less or equal to 1. But often the maximum of the distribution is not exactly 1. This is caused by the regression model when the expected value differs slightly from the measured value.

- We are performing a regression on the data to calculate the coefficients and therefore we determine the residuals for the fitted values. The distribution of the residuals is also given as part of the summary.

- The coefficients α and β are the result that we are essentially interested in. They tell us the magnitude of the contention and coherency effects within the system.

- The third coefficient γ estimates the throughput for a single processor. In this case we actually have a measurement for the single processor case but for a real-world analysis this value if often unknown or can't be measured. The difference between the value of γ and the measurement is also caused by the regression. The difference should be small if the regression finds a reasonable model.

- The residual standard error estimates how well the model fits the data. We can see that the difference between the model prediction and the measured values is typically within 9.34 ray-tracing operations per second.

6

- Finally the scalability bounds are printed. These bounds state important limits of the scalability as defined by the model. We will look at the meaning of these values shortly.

The function `efficiency()` extracts the efficiency values from the model and allows us to have a closer look at the specific efficiencies of the different processor configurations.

```
R> efficiency(usl.model)


     1      4      8     12     16     20     24     28     32     48     64
0.9154 0.8925 0.7437 0.6484 0.5435 0.4577 0.4005 0.3760 0.3719 0.2670 0.2217
```

A bar plot is useful to visually compare the decreasing efficiencies for the configurations with an increasing number of processors. Figure 2 shows the output diagram.

```
R> barplot(efficiency(usl.model), ylab = "efficiency / processor", xlab = "processors")
```


Figure 2: Rate of efficiency per processor for different numbers of processors running the raytracing software

The efficiency can be used for a first validation and sanity check of the measured values. Values larger than 1.0 usually need a closer investigation. It is also suspicious if the efficiency gets bigger when the load increases.

The model coefficients alone can be retrieved with the `coef()` function.

```
R> coef(usl.model)


   alpha     beta    gamma
 0.05777  0.00000 21.84884
```

We can see that the model estimates the contention parameter α to be about 5.8 percent and the coherency parameter β to be 0.

With $\beta = 0$ the model indicates no coherency penalty for the system. In this case the throughput will asymptotically approach an upper bound as the number of processors is increased. The summary output above indicates a limit of 378 as the upper bound for the maximal achievable throughput. There is no way to increase the throughput beyond that number, no matter how many processors are available whatsoever.

For $\beta = 0$ the predicted throughput never reaches a maximum. Theoretically the throughput could always be increased by adding more processors. At the same time it is limited as stated in the previous paragraph. So the improvement becomes infinitesimal for more and more processors.

To get an impression of the scalability function we can use the `plot()` function and create a combined graph with the original data as dots and the calculated scalability function as a solid line. The scalability bounds are added as dotted lines to the output. Figure 3 has the result of that plot.

```
R> plot(throughput ~ processors, data = raytracer, pch = 16, ylim = c(0, 400))
R> plot(usl.model, add = TRUE, bounds = TRUE)
```


Figure 3: Throughput of a raytracing software using different numbers of processors

The plot has dotted lines for the upper limit defined by Amdahl (here for a throughput of 378) and the linear scalability we strive to see in all our applications. Unfortunately the throughput falls away from the linear path rather quickly.

8

Optimal scalability is reached at 17.3 processors performing 195 ray-tracing operations per second. The optimum is defined as the point on the x-axis below the intersection of the linear scalability bound and the scalability limit defined by Amdahl's Law.

Running a system below this point will underutilize already purchased capacity. But pushing the system above this point leads to *diminishing returns*. The point of optimal scalability is also given in the output of the `summary()` function as shown before.

The confidence intervals for the model coefficients are returned by calling the `confint()` function.

```
R> confint(usl.model, level = 0.95)


           2.5 %      97.5 %
alpha   0.0336773   0.0818642
beta   -0.0002137   0.0002137
gamma  17.8683796  25.8293061
```

The confidence interval shows for a given confidence level the expected range that the true value lies in. We can see that the true β might in fact not really be zero. One reason is that all measurements contain measurement and observation errors. The other reason is that a very small β might only cause a measurable effect when a much higher number of processors are used.

SGI marketed the Origin 2000 with up to 128 processors. Let's assume that going from 64 to 128 processors does not introduce any additional limitations to the system architecture. Then we can use the existing model and forecast the system throughput for other numbers like 96 and 128 processors using the `predict()` function.

```
R> predict(usl.model, data.frame(processors = c(96, 128)))


    1     2
323.3 335.5
```

We can see from the prediction that the throughput will reach 335.5 for 128 processors.

4.2 Case Study: Software Scalability

In this section we will perform an analysis of a SPEC benchmark. A SPARCcenter 2000 with 16 CPUs was used in October 1994 for the SDM91 benchmark[2]. The benchmark simulates a number of users working on a UNIX server (editing files, compiling ...) and measures the number of script executions per hour.

First, select the demo dataset with the data from the SPEC SDM91 benchmark.

---

[2] http://www.spec.org/osg/sdm91/results/results.html

```
R> library(usl)
R> data(specsdm91)
R> specsdm91


  load throughput
1    1        64.9
2   18       995.9
3   36      1652.4
4   72      1853.2
5  108      1828.9
6  144      1775.0
7  216      1702.2
```

The data provides the measurements made during the benchmark. The column "load" shows the number of virtual users that were simulated by the benchmark and the column "throughput" has the measured number of script executions per hour for that load.

Next we create the USL model for this dataset by calling the usl() function. Again we specify a symbolic description of the model and the dataset with the measurements. But this time we choose a different method for the analysis.

```
R> usl.model <- usl(throughput ~ load, specsdm91, method = "nls")
```

There are currently three possible values for the method parameter:

nls: This method uses the nls() function of the stats package for a nonlinear regression model. It estimates the coefficients α, β and γ. The nonlinear regression uses constraints for its parameters which means the "port" algorithm is used internally to solve the model. So all restrictions of the "port" algorithm apply.

nlxb: A nonlinear regression model is also used in this case. But instead of the nls() function it uses the nlxb() function from the nlsr package (cf. [4]). It is expected to be more robust than the nls method.

default: The default method traditionally used a transformation into a 2nd degree polynomial and required the data to be normalized. This was the original algorithm for the two-parameter Universal Scalability Law (cf. chapter 5.2.3 of [1]). Starting with version 2.0.0 of the usl package this algorithm is no longer available. The nlxb() function is used instead.

We also use the summary() function to look at the details for the analysis.

```
R> summary(usl.model)
```

```
Call:
usl(formula = throughput ~ load, data = specsdm91, method = "nls")

Efficiency:
   Min     1Q Median     3Q    Max
0.0876 0.1626 0.2860 0.5624 0.7211

Residuals:
   Min     1Q Median     3Q    Max
 -81.7  -48.3  -25.1   29.5  111.1

Coefficients:
        Estimate  Std. Error  t value  Pr(>|t|)
alpha   0.0277285   0.0091217     3.04    0.0384  *
beta    0.0001044   0.0000199     5.25    0.0063  **
gamma  89.9952384  14.2134906     6.33    0.0032  **
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 82.8 on 4 degrees of freedom

Scalability bounds:
limit: throughput 3250 (Amdahl asymptote)
peak:  throughput 1880 at load 96.5
opt:   throughput 1540 at load 36.1
```

Looking at the coefficients we notice that α is about 0.028 and β is about 0.0001. The parameter α indicates that about 2.8 percent of the execution time is strictly serial. As a rule of thumb we can say that 3 percent is an excellent value for a reasonable complex application. Note that this serial fraction is also recognized in Amdahl's Law.

We hypothesize that a proposed change to the system — maybe a redesign of the cache architecture or the elimination of a point to point communication — could reduce β by half. Now we want to predict how the scalability of the system would change.

We can calculate the point of maximum scalability for the current system and for the hypothetical system with the `peak.scalability()` function.

```
R> peak.scalability(usl.model)


[1] 96.52


R> peak.scalability(usl.model, beta = 0.00005)


[1] 139.4
```

The function accepts the optional arguments `alpha` and `beta`. They are useful to do a what-if analysis. Setting these parameters override the calculated model parameters and show how the system would behave with a different contention or coherency coefficient.

In this case we learn that the point of peak scalability would move from around 96.5 to about 139 virtual users if we would be able to actually build the system with the assumed optimization.

Both calculated scalability functions can be plotted using the `plot()` or `curve()` functions. The following commands create a graph of the original data points and the derived scalability functions. To completely include the scalability of the hypothetical system, we have to increase the range of the plotted values with the first command.

```
R> plot(specsdm91, pch = 16, ylim = c(0,2500))
R> plot(usl.model, add = TRUE)
R>
R> # Create function cache.scale to perform calculations with the model
R> cache.scale <- scalability(usl.model, beta = 0.00005)
R> curve(cache.scale, lty = 2, add = TRUE)
```

We used the function `scalability()` here. It is a higher order function returning a function and not just a single value. The return value is assigned to define `cache.scale`. That makes it possible to use the `curve()` function to plot the values over the specific range.

Figure 4 shows the measured throughput in scripts per hour for a given load, i.e., the number of simulated users. The solid line indicates the derived USL model while the dashed line resembles our hypothetical system using the proposed optimization.

From the figure we can see that the scalability really peaks at one point. Increasing the load beyond that point leads to retrograde behavior, i.e., the throughput decreases again. As we have calculated earlier, the measured system will reach this point sooner than the hypothetical system.

We can combine the `cache.scale()` (defined before) and `peak.scalability()` functions to get the predicted throughput values for the peak values.

```
R> scalability(usl.model)(peak.scalability(usl.model))


[1] 1884


R> # Use cache.scale function defined before
R> cache.scale(peak.scalability(usl.model, beta = 0.00005))


[1] 2162
```

Figure 4: The result of the SPEC SDM91 benchmark for a SPARCcenter 2000 (dots) together with the calculated scalability function (solid line) and a hypothetical scalability function (dashed line)

This illustrates how the Universal Scalability Law can help to decide if the system currently is more limited by contention or by coherency issues and also what impact a proposed change would have.

The `predict()` function can also be used to calculate a confidence bands for the scalability function at a specified level. To get a smoother graph it is advisable to predict the values for a higher number of points. Let's start by creating a data frame with the required load values.

```
R> load <- with(specsdm91, expand.grid(load = seq(min(load), max(load))))
```

We use the data frame to determine the fitted values and also the upper and lower confidence bounds at the requested level. The result will be a matrix with column names `fit` for the fitted values, `lwr` for the lower and `upr` for the upper bounds.

```
R> fit <- predict(usl.model, newdata = load, interval = "confidence", level = 0.95)
```

The matrix is used to define the coordinates of a polygon containing the area between the lower and the upper bounds. The polygon connects the points of the lower bounds from lower to higher values and then back using the points of the upper bounds.

```
R> usl.polygon <- matrix(c(load[, 1], rev(load[, 1]), fit[, 'lwr'], rev(fit[, 'upr'])),
+                        nrow = 2 * nrow(load))
```

13

The plot is composed from multiple single plots. The first plot initializes the canvas and creates the axis. Then the polygon is plotted using a gray area. In the next step the measured values are added as points. Finally a solid line is plotted to indicate the fitted scalability function. See fig. 5 for the entire plot.

```
R> # Create empty plot (define canvas size, axis, ...)
R> plot(specsdm91, xlab = names(specsdm91)[1], ylab = names(specsdm91)[2],
+       ylim = c(0, 2000), type = "n")
R>
R> # Plot gray polygon indicating the confidence interval
R> polygon(usl.polygon, border = NA, col = "gray")
R>
R> # Plot the measured throughput
R> points(specsdm91, pch = 16)
R>
R> # Plot the fit
R> lines(load[, 1], fit[, 'fit'])
```


Figure 5: The result of the SPEC SDM91 benchmark with confidence bands for the scalability function at the 95% level

Another way to illustrate the impact of the parameters α and β on the scalability is by looking at the achievable speedup when a fixed load is parallelized. A naive estimation would be that doubling the degree of parallelization should cut the execution time in halve.

Unfortunately it doesn't work this way. In general there is a range where doubling the parallelization will actually improve the execution time. But the improvement will get smaller and smaller when the degree of parallelism is increased further. This is also an effect of *diminishing returns* as already seen in section 4.1. The real execution time is in fact the sum

14

of the ideal execution time and the overhead for dealing with contention and coherency delays.

The total execution time of a parallelized workload depends on the degree of parallelism p and the coefficients α and β of the associated USL model.

The magnitude of the three components — given as fractions of the serial execution time T_1 — that account for the total execution time of the parallelized workload can be calculated as follows (cf. [2] eqn. 26).

$$T_{ideal} = \frac{1}{p} T_1 \tag{3}$$

$$T_{contention} = \alpha \left(\frac{p-1}{p} \right) T_1 \tag{4}$$

$$T_{coherency} = \beta \frac{1}{2} (p-1) T_1 \tag{5}$$

The function `overhead()` can be used to calculate the correspondent fractions for a given model. The function has the same interface as the `predict()` function. Calling it with only the model as argument will calculate the overhead for the fitted values. It can also be called with a data frame as second argument. Then the data frame will be used to determine the values for the calculation.

Let's use our current model to calculate the overhead for a load of 10, 20, 100 and 200 simulated users. We create a data frame with the number of users and use the `overhead()` function to estimate the overhead.

```
R> load <- data.frame(load = c(10, 20, 100, 200))
R> ovhd <- overhead(usl.model, newdata = load)
R> ovhd


  ideal contention coherency
1 0.100    0.02496 0.0004696
2 0.050    0.02634 0.0009915
3 0.010    0.02745 0.0051661
4 0.005    0.02759 0.0103844
```

We can see that the ideal execution time for running 10 jobs in parallel is 1/10 of the execution time of running the jobs unparallelized. To get the total fraction we have to add the overhead for contention (2.5%) and for coherency delays (0.047%). This gives a total of 12.54%. So with 10 jobs in parallel we are only about 8 times faster than running the same workload in a serial way.

Equation (4) shows that the percentage of time spent on dealing with contention will converge to the value of α. Equation (5) explains that coherency delays will grow beyond any limit if the degree of parallelism is large enough. This corresponds to the observation that adding more parallelism will sometimes make performance worse.

A stacked barplot can be used to visualize how the different effects change with an increasing degree of parallelism. Note that the result matrix must be transposed to match the format needed for for the `barplot()` command.

```
R> barplot(height = t(ovhd), names.arg = load[, 1],
+          xlab = names(load), legend.text = TRUE)
```


Figure 6: Decomposition of the execution time for parallelized workloads of the SPECSDM91 benchmark. The time is measured as a fraction of the time needed for serial execution of the workload.

Figure 6 shows the resulting plot. It clearly shows the decrease in ideal execution time when the degree of parallelism is increased. It also shows how initially almost only contention contributes to the total execution time. For higher degrees of parallelism the impact of coherency delays grows. Note how the difference in ideal execution time between 100 and 200 parallel jobs effectively has no effect on the total execution time.

4.3 Case Study: Multivalued Data

It is very common to use multivalued data for a scalability analysis. These measurements are often taken from a live system and may include many different data points for similar load values. This could be the result of a non-homogeneous workload and an analyst has to

decide how to take that into account. But for a production system there is usually no feasible way to create a homogeneous workload.

The following data shows a subset of performance data gathered from an Oracle database system providing a login service for multiple web applications. For the analysis we focus on only two of the available metrics:

txn_rate: The average number of processed database transactions. This metric is given as transactions per second.

db_time: The average time spent inside the database either working on a CPU or waiting for resources (I/O, locks, buffers . . .). The time is expressed as seconds per second, so two sessions working for exactly one quarter of a second each will contribute a total of half a second for that second. Oracle has coined the term *Average Active Sessions* (AAS) for this metric.

Let's have a look at the first couple of data points in our data set. For each time interval of two minutes there is a corresponding value for the average database time per seconds and for the average number of transactions per second in this interval.

```
R> data(oracledb)
R> head(subset(oracledb, select = c(timestamp, db_time, txn_rate)))


            timestamp db_time txn_rate
1 2012-01-19 08:02:00  0.3120    2.205
2 2012-01-19 08:04:00  0.3224    2.574
3 2012-01-19 08:06:00  0.1918    1.790
4 2012-01-19 08:08:00  0.3136    2.587
5 2012-01-19 08:10:00  0.3584    2.321
6 2012-01-19 08:12:00  0.2354    1.958
```

A naive approch would be a plot of the data as a time series (see fig. 7). This plot shows the familiar pattern of an OLTP application that is mostly used during office hours. Unfortunately this type of plot is pretty much useless when performing a scalability analysis.

The Universal Scalability Law correlates a throughput with a load. In this case the throughput is clearly given by the transaction rate and the database time is a taken as the load metric. The definition above states that the total time spent — either running on a CPU or waiting — is a measurement for the average number of active sessions. So we use that to express the load on the database system.

As usual, we call the usl() function to carry out the analysis. See fig. 8 for the scatterplot of the data including the plot of the estimated scalability function.

```
R> plot(txn_rate ~ db_time, oracledb,
+       xlab = "Average active sessions", ylab = "Txn / sec")
R>
R> usl.oracle <- usl(txn_rate ~ db_time, oracledb)
R> plot(usl.oracle, add = TRUE)
```

Figure 7: Transaction rates of an Oracle database system during the day of January 19th, 2012


Figure 8: Relationship between the transaction rate and the number of average active sessions in an Oracle database system

Now we can retrieve the coefficients for this model.

```
R> coef(usl.oracle)


 alpha   beta  gamma
0.4414 0.0453 3.3861
```

Our α here is about an order of magnitude bigger than what we have seen in the previous sections. This indicates a major issue with some kind of serialization or queueing that severely limits the scalability. In fact it is so bad that the impact is already visible with only a few active sessions working at the same time: according to the model the peak throughput is reached at about 3.5 sessions.

```
R> peak.scalability(usl.oracle)


[1] 3.512
```

The confidence interval for α confirms that there is only a small uncertainty about the magnitude of the calculated coefficients.

```
R> confint(usl.oracle)


        2.5 %  97.5 %
alpha 0.36429 0.51845
beta  0.01861 0.07199
gamma 3.28529 3.48687
```

This analysis shows how we can use some of the metrics provided by a live Oracle database system to learn about the scalability. Note that neither the Oracle software nor the application needed any additional instrumentation to collect this data. Also the analysis was done without any internal knowledge about the way the application was using the database.

References

[1] Neil J. Gunther. *Guerrilla Capacity Planning: A Tactical Approach to Planning for Highly Scalable Applications and Services.* Springer, Heidelberg, Germany, 1st edition, 2007.

[2] Neil J. Gunther. A general theory of computational scalability based on rational functions. *CoRR*, abs/0808.1431, 2008.

[3] Neil J. Gunther. Usl scalability modeling with three parameters. `http://perfdynamics.blogspot.com/2018/05/usl-scalability-modeling-with-three.html`, 2018. Accessed: 2020-01-08.

[4] John C. Nash. *nlsr: Functions for nonlinear least squares solutions*, 2017. R package version 2017.6.18.