

Row versus Set Processing Experiment.

31-May-2012 by Craig Shallahamer, craig@orapub.com

Set processing was performed using Oracle SQL and the row processing was performed using Oracle's plsql. The experimental details and analysis is posted on Craig Shalahamer's (craig@orapub.com) blog (A Wider View) in May of 2012.

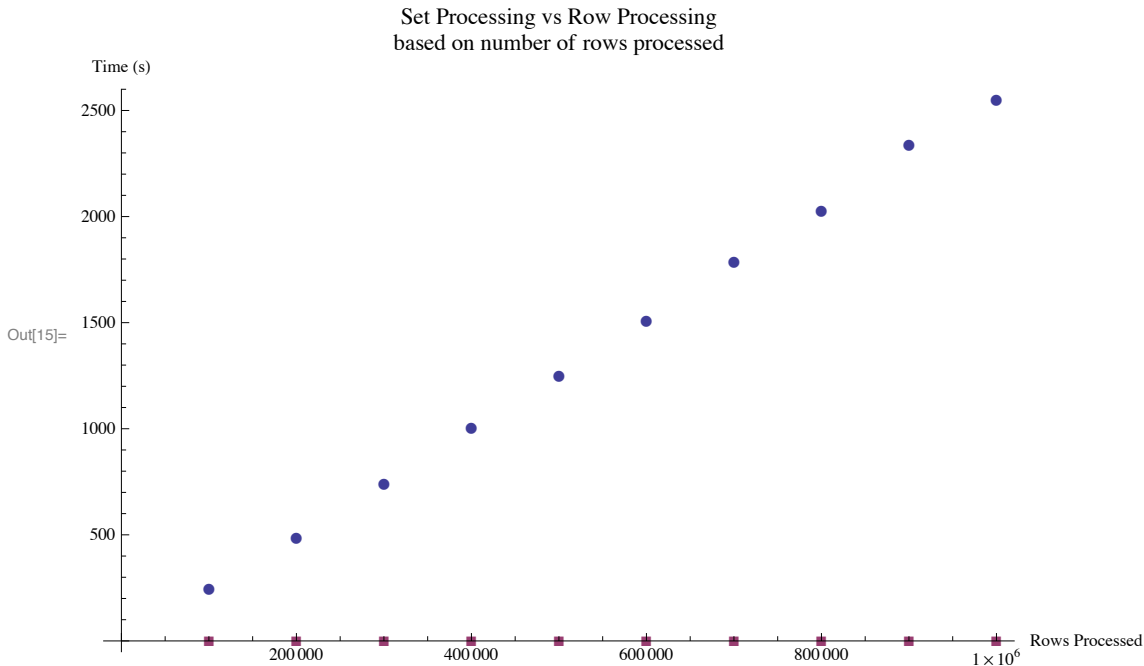
The data below is summarized from the raw experimental data. Each of the below 10 row and set processing samples is the average from 12 actual samples. You can see the samples are grouped; (rows processed, elapsed time in seconds).

```
In[13]:= rowProcessing = {{100 000, 244.827484}, {200 000, 488.268893}, {300 000, 740.67795},
      {400 000, 1002.74595}, {500 000, 1250.8051}, {600 000, 1507.10145}, {700 000, 1789.35509},
      {800 000, 2026.92326}, {900 000, 2339.14553}, {1 000 000, 2548.92394}}
      setProcessing = {{100 000, .065524417}, {200 000, .131115083}, {300 000, .197437},
      {400 000, .231783}, {500 000, .30142925}, {600 000, .360746583}, {700 000, .428058917},
      {800 000, .489497917}, {900 000, .56059475}, {1 000 000, .628519167}}
```


```
Out[13]= {{100 000, 244.827}, {200 000, 488.269}, {300 000, 740.678},
      {400 000, 1002.75}, {500 000, 1250.81}, {600 000, 1507.1}, {700 000, 1789.36},
      {800 000, 2026.92}, {900 000, 2339.15}, {1 000 000, 2548.92}}
```

```
Out[14]= {{100 000, 0.0655244}, {200 000, 0.131115}, {300 000, 0.197437},
      {400 000, 0.231783}, {500 000, 0.301429}, {600 000, 0.360747}, {700 000, 0.428059},
      {800 000, 0.489498}, {900 000, 0.560595}, {1 000 000, 0.628519}}
```


```
In[15]:= ListPlot[{rowProcessing, setProcessing},
      PlotLabel -> "Set Processing vs Row Processing\nbased on number of rows processed",
      AxesLabel -> {"Rows Processed", "Time (s)"}, AxesOrigin -> {0, 0}, PlotMarkers -> Automatic]
      ListPlot[rowProcessing, PlotLabel -> "Row Processing\nTime vs Rows Processed",
      AxesLabel -> {"Rows Processed", "Time (s)"},
      AxesOrigin -> {0, 0}, PlotMarkers -> Automatic, Filling -> Axis]
      ListPlot[setProcessing, PlotLabel -> "Set Processing\nTime vs Rows Processed",
      AxesLabel -> {"Rows Processed", "Time (s)"},
      AxesOrigin -> {0, 0}, PlotMarkers -> Automatic, Filling -> Axis]
```



Row Processing
Time vs Rows Processed




Set Processing
Time vs Rows Processed



```
In[18]:= rLm = LinearModelFit[rowProcessing, x, x]
Show[ListPlot[rowProcessing, PlotLabel -> "Row Processing\nTime vs Rows Processed",
  AxesLabel -> {"Rows Processed", "Time (s)"}, AxesOrigin -> {0, 0},
  PlotMarkers -> Automatic], Plot[rLm[x], {x, 0, 1 000 000}]]
```

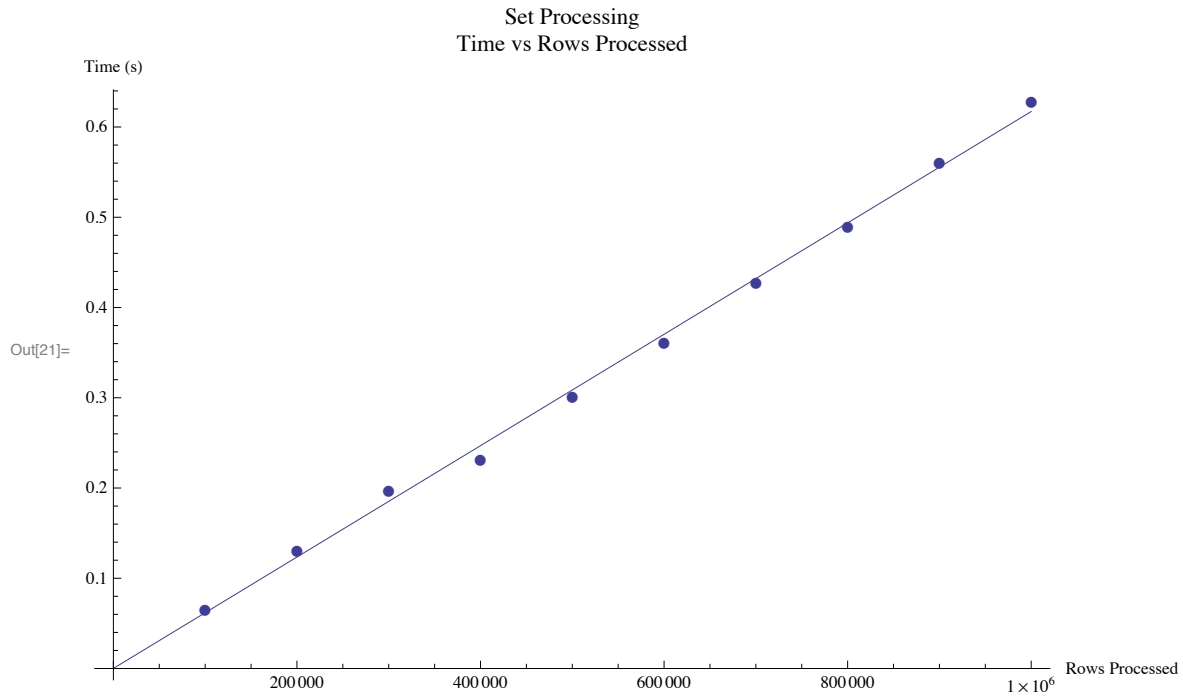
```
Out[18]= FittedModel[
$$-30.801 + 0.00259032 x$$
]
```



```
Out[19]=
```

```
In[20]:= sLm = LinearModelFit[setProcessing, x, x]
Show[ListPlot[setProcessing, PlotLabel -> "Set Processing\nTime vs Rows Processed",
  AxesLabel -> {"Rows Processed", "Time (s)"}, AxesOrigin -> {0, 0},
  PlotMarkers -> Automatic], Plot[sLm[x], {x, 0, 1 000 000}]]
```

```
Out[20]= FittedModel[ $0.0000786055 + 6.17076 \times 10^{-7} x$ ]
```



```
Out[21]=
```

```

In[22]:= setThroughput = {{100 000, 1 526 475.76},
  {200 000, 1 525 424.15}, {300 000, 1 519 560.23}, {400 000, 1 725 765.05},
  {500 000, 1 658 774.21}, {600 000, 1 663 226.12}, {700 000, 1 635 291.37},
  {800 000, 1 634 334.23}, {900 000, 1 605 444.99}, {1 000 000, 1 591 043.33}}
rowThroughput = {{100 000, 408.461317}, {200 000, 409.615077}, {300 000, 405.039052},
  {400 000, 398.911805}, {500 000, 399.751556}, {600 000, 398.117411}, {700 000, 391.215288},
  {800 000, 394.696477}, {900 000, 384.778072}, {1 000 000, 392.338925}}
ListPlot[{rowThroughput, setThroughput},
  PlotLabel -> "Throughput\nSet Processing vs Row Processing",
  AxesLabel -> {"Rows Processed", "Throughput (rows/s)"},
  AxesOrigin -> {0, 0}, PlotMarkers -> Automatic]

```

```

Out[22]= {{100 000, 1.52648 × 106}, {200 000, 1.52542 × 106},
  {300 000, 1.51956 × 106}, {400 000, 1.72577 × 106}, {500 000, 1.65877 × 106},
  {600 000, 1.66323 × 106}, {700 000, 1.63529 × 106}, {800 000, 1.63433 × 106},
  {900 000, 1.60544 × 106}, {1 000 000, 1.59104 × 106}}

```

```

Out[23]= {{100 000, 408.461}, {200 000, 409.615}, {300 000, 405.039},
  {400 000, 398.912}, {500 000, 399.752}, {600 000, 398.117}, {700 000, 391.215},
  {800 000, 394.696}, {900 000, 384.778}, {1 000 000, 392.339}}

```

