# Commit Write Facility; dml1.sql

Author: Craig Shallahamer (craig@orapub.com), Version 1f, 17-Feb-2012.

## Background and Purpose

*The purpose of this notepad is to see if using Oracle's commit write facility increases performance measured by a decrease in response time (ms/commits) and increase in arrival rate.*

Because it can be difficult to compare the response times when the arrival rate changes, at the bottom of the notebook I complete the relationship between work and time, and then set the work to be same then derived the response time. This allows a pretty good apples-to-apples comparision.

The data used in the notebook is based on the dml1.sql script, which creates a heavy CPU load and a light IO load on the system. But the top wait event during the wait,immediate option is still log file sync.

The terminalolgy used is a mix of queury theory and Oracle database-ease. Without getting into details, CPU time per unit of work is equal to the service time, the non-idle wait time per unit of work is the queue time, and respons time is service time plus queue time. Key to understanding this is knowing the time is based on a single unit of work. For this experiment the chosen unit of work was the Oracle statistic "user commit". Every commit issued by an Oracle client/user process will tick-up the user commit statistics. As a result, the user commit is a very good unit of work when understanding the commit rate of an application or workload.

## Experimental Data

Below is all the experimental data. The experiment was run on a Dell single four-core CPU, Oracle 11.2G. According to "cat /proc/version": Linux version 2.6.18-164.el5PAE (mockbuild@ca-build10.us.oracle.com) (gcc version 4.1.2 20080704 (Red Hat 4.1.2-46)) #1 SMP Thu Sep 3 02:28:20 EDT 2009.
There was an intense DML update load by five sessions. **The DML script is dml1.sql**, which in the experimental enviroment cause a severe CPU bottleneck and then some log file sync time. The CPU consumption was many, many times more than the wait time. This can be seen far below, by comparing the average CPU per commmit compared to the average wait time per commit.

After each update and commit, there was a subsecon lognomal distirbuted sleep time. During each of the data collection periods, the follow data was gathered. The percentage figures were visually observed and recorded.

```
Grid[{
  {"Setting", "Commit/sec", "CPU %",
   "log file sync\nWait Time %", "log file par write\nWait Time %"},
  {"wait immediate", "8.32", "98%", "48%", "38%"},
  {"wait batch", "8.46", "98%", "53%", "37%"},
  {"nowait batch", "8.47", "99%", "0%", "70%"},
  {"nowait immediate", "8.47", "99%", "0%", "13%"}
 }, Frame → All
]
```

| Setting | Commit/sec | CPU % | log file sync Wait Time % | log file par write Wait Time % |
|---|---|---|---|---|
| wait immediate | 8.32 | 98% | 48% | 38% |
| wait batch | 8.46 | 98% | 53% | 37% |
| nowait batch | 8.47 | 99% | 0% | 70% |
| nowait immediate | 8.47 | 99% | 0% | 13% |

As you can see below, for each of the four commit write facility options, there were 32, 90 second samples taken.

(sec), and instance CPU consumption (sec).

```
ssWaitImmediate = {1, 90.016833, 771, 19.99, 354.364125, 2, 90.033525, 749, 17.87, 354.691076, 3,
   90.008822, 773, 21.36, 350.727679, 4, 90.013937, 751, 22.03, 346.919255, 5, 90.018566, 731,
   20.81, 350.842663, 6, 90.011928, 758, 20.78, 354.665081, 7, 90.018995, 703, 18.66, 354.482108,
   8, 90.01774, 760, 21.02, 354.027181, 9, 90.009464, 748, 21.74, 354.479108, 10, 90.009953, 731,
   20.36, 350.538708, 11, 90.014985, 765, 22.67, 354.848052, 12, 90.011934, 749, 20.76, 347.112231,
   13, 90.127603, 734, 21.82, 355.003026, 14, 90.034552, 782, 18.93, 358.275535, 15, 90.017466, 760,
   21.59, 354.611091, 16, 90.007765, 725, 18.74, 349.996791, 17, 90.181153, 737, 21.2, 354.792059,
   18, 90.016256, 733, 20.15, 346.241361, 19, 90.015182, 753, 19.65, 353.121318, 20, 90.015107, 734,
   18.79, 352.353429, 21, 90.007711, 739, 23.59, 352.740372, 22, 90.012739, 719, 19.95, 352.312437,
   23, 90.013051, 758, 21.24, 348.47902, 24, 90.028347, 765, 20.94, 354.810059, 25, 90.005068, 757,
   20.21, 346.68229, 26, 90.015379, 786, 22.92, 354.465113, 27, 90.013851, 751, 23.76, 354.618089,
   28, 90.016416, 747, 19.47, 358.364517, 29, 90.015136, 723, 20, 355.046019, 30, 90.027755, 757,
   21.76, 354.118158, 31, 90.018465, 736, 19.13, 347.09023, 32, 90.015684, 783, 20.79, 350.503711};
ssWaitBatch = {1, 90.012865, 774, 22.6, 358.332522, 2, 90.008947, 800, 21.06, 355.045021, 3,
   90.017088, 758, 21.58, 355.118006, 4, 90.127519, 766, 20.55, 354.750064, 5, 90.102755, 758,
   19.97, 362.216932, 6, 90.018209, 780, 22.16, 350.464718, 7, 90.009452, 752, 20.37, 346.508324,
   8, 90.014387, 745, 19.89, 354.683074, 9, 90.005624, 787, 21.59, 354.466105, 10, 90.013743, 734,
   20.27, 354.671084, 11, 90.25017, 779, 19.93, 355.018023, 12, 90.014355, 740, 21.23, 347.152222,
   13, 90.007112, 763, 21.49, 354.397123, 14, 90.045861, 755, 19.64, 362.34791, 15, 90.096495, 750,
   19.92, 350.322742, 16, 90.009461, 783, 22.79, 354.400121, 17, 90.005002, 752, 19.87, 354.659082,
   18, 90.178983, 770, 22.42, 354.747066, 19, 90.008779, 742, 21.12, 351.146615, 20, 90.032213, 754,
   21.07, 354.689078, 21, 90.013209, 752, 20.24, 354.948034, 22, 90.013705, 738, 19.6, 350.747672,
   23, 90.05598, 770, 20.74, 354.680078, 24, 90.017555, 740, 19.09, 354.166153, 25, 90.130722, 769,
   21.93, 359.919283, 26, 90.043504, 749, 22.23, 354.775064, 27, 90.008103, 760, 20.05, 346.497318,
   28, 90.036155, 802, 21.51, 354.282141, 29, 90.010146, 743, 18.66, 355.071018, 30, 90.14741, 784,
   20.61, 354.73207, 31, 90.071883, 790, 20.73, 358.664473, 32, 90.040811, 734, 19.8, 354.707074};
ssNowaitImmediate = {1, 90.009553, 767, .81, 356.866744, 2, 90.019994, 783, .51, 357.144702,
   3, 90.019171, 770, .64, 357.324679, 4, 90.021457, 796, .45, 357.037716, 5, 90.008423, 770,
   .53, 361.16009, 6, 90.008982, 752, .54, 356.940735, 7, 90.019636, 822, 15.07, 355.901891,
   8, 90.010252, 752, .44, 356.987728, 9, 90.011047, 789, .62, 353.240297, 10, 90.009254, 751,
   .51, 357.439661, 11, 90.019374, 734, 3.09, 360.610173, 12, 90.015876, 734, .38, 355.347976,
   13, 90.01989, 739, .48, 352.75137, 14, 90.010219, 746, .77, 346.632299, 15, 90.008739, 753,
   .45, 346.246359, 16, 90.019661, 751, .36, 350.176761, 17, 90.019855, 818, .56, 354.17415,
   18, 90.01924, 760, .47, 357.274682, 19, 90.007998, 762, .51, 357.428658, 20, 90.010544, 778,
   .5, 357.219693, 21, 90.009052, 785, .44, 357.237689, 22, 90.021595, 764, .57, 356.870745,
   23, 90.01029, 777, .5, 357.269683, 24, 90.019274, 747, .57, 356.839753, 25, 90.008304, 739,
   .39, 352.970341, 26, 90.009707, 756, .42, 349.240906, 27, 90.008482, 753, .4, 350.541707, 28,
   90.012067, 755, .49, 353.878196, 29, 90.017928, 761, .61, 349.937801, 30, 90.008913, 744,
   .65, 349.2599, 31, 90.009847, 750, .35, 349.267899, 32, 90.008737, 750, .52, 355.206999};
ssNowaitBatch = {1, 90.020026, 770, .1, 357.219694, 2, 90.012683, 774, .1, 356.704768, 3,
   90.019968, 774, .13, 357.184694, 4, 90.020716, 736, .18, 357.260685, 5, 90.010322, 764,
   .1, 352.88535, 6, 90.009205, 721, .07, 349.190912, 7, 90.010882, 764, .16, 354.696071, 8,
   90.020235, 781, .07, 357.380669, 9, 90.008057, 765, .15, 356.88574, 10, 90.00892, 759, .22,
   356.866741, 11, 90.008881, 747, .16, 357.053721, 12, 90.019757, 768, .19, 357.088711, 13,
   90.019939, 753, .25, 357.274684, 14, 90.020031, 738, .2, 357.050717, 15, 90.019802, 787, .1,
   356.964727, 16, 90.008867, 778, .08, 357.258683, 17, 90.009095, 773, .17, 357.196697, 18,
   90.021109, 759, .22, 352.730374, 19, 90.009436, 761, .13, 357.308677, 20, 90.008822, 768,
   .13, 357.176699, 21, 90.008863, 753, .13, 357.222692, 22, 90.009543, 735, .09, 356.835751,
   23, 90.009349, 793, .21, 356.587787, 24, 90.020508, 767, .16, 356.848753, 25, 90.009135, 761,
   .13, 360.342215, 26, 90.008955, 787, .11, 354.730069, 27, 90.008543, 758, .14, 352.981338,
   28, 90.010361, 753, .2, 356.89174, 29, 90.019722, 748, .09, 357.274682, 30, 90.020663, 791,
   .38, 357.702619, 31, 90.009318, 742, .06, 353.872199, 32, 90.008455, 769, .2, 349.396881};
```

## Data Loading

All the data sets are contained in the above section.

```
ssNum = 4;
sampleNum = 32;
ss[1] = ssWaitImmediate; ssName[1] = "Wait Immediate";
ss[2] = ssWaitBatch; ssName[2] = "Wait Batch";
ss[3] = ssNowaitImmediate; ssName[3] = "Nowait Immediate";
ss[4] = ssNowaitBatch; ssName[4] = "Nowait Batch";

ncols = 5;
sampleCol = 1;
elapsedTCol = 2;
workCol = 3;
waitTCol = 4;
cpuTCol = 5;

Do[
  ssWork[ssidx] = {}; ssL[ssidx] = {}; ssSt[ssidx] = {}; ssQt[ssidx] = {}; ssRt[ssidx] = {};
  theSS = ss[ssidx];
  Table[
   elapsedT = theSS[[ncols sampleidx + elapsedTCol]];
   workTot = theSS[[ncols sampleidx + workCol]];
   cpuSecTot = theSS[[ncols sampleidx + cpuTCol]];
   waitSecTot = theSS[[ncols sampleidx + waitTCol]];
   (*Print[ssidx," ",sampleidx," elapsedT=",elapsedT];*)

   λ = workTot / (1000 elapsedT);
   St = (cpuSecTot 1000) / workTot;
   Qt = (waitSecTot 1000) / workTot;
   Rt = St + Qt;

   AppendTo[ssWork[ssidx], workTot];
   AppendTo[ssL[ssidx], λ];
   AppendTo[ssSt[ssidx], St];
   AppendTo[ssQt[ssidx], Qt];
   AppendTo[ssRt[ssidx], Rt];

   , {sampleidx, 0, sampleNum - 1}
  ];
  , {ssidx, ssNum}
 ];
ssName[1]
ss[1]
Length[ssWork[1]]
Take[ssWork[1], 5]
N[Mean[ssWork[1]]]
```

```
Wait Immediate
```

```
{1, 90.0168, 771, 19.99, 354.364, 2, 90.0335, 749, 17.87, 354.691, 3, 90.0088, 773, 21.36, 350.728, 4,
 90.0139, 751, 22.03, 346.919, 5, 90.0186, 731, 20.81, 350.843, 6, 90.0119, 758, 20.78, 354.665, 7,
 90.019, 703, 18.66, 354.482, 8, 90.0177, 760, 21.02, 354.027, 9, 90.0095, 748, 21.74, 354.479, 10,
 90.01, 731, 20.36, 350.539, 11, 90.015, 765, 22.67, 354.848, 12, 90.0119, 749, 20.76, 347.112, 13,
 90.1276, 734, 21.82, 355.003, 14, 90.0346, 782, 18.93, 358.276, 15, 90.0175, 760, 21.59, 354.611,
 16, 90.0078, 725, 18.74, 349.997, 17, 90.1812, 737, 21.2, 354.792, 18, 90.0163, 733, 20.15,
 346.241, 19, 90.0152, 753, 19.65, 353.121, 20, 90.0151, 734, 18.79, 352.353, 21, 90.0077, 739,
 23.59, 352.74, 22, 90.0127, 719, 19.95, 352.312, 23, 90.0131, 758, 21.24, 348.479, 24, 90.0283,
 765, 20.94, 354.81, 25, 90.0051, 757, 20.21, 346.682, 26, 90.0154, 786, 22.92, 354.465, 27,
 90.0139, 751, 23.76, 354.618, 28, 90.0164, 747, 19.47, 358.365, 29, 90.0151, 723, 20, 355.046, 30,
 90.0278, 757, 21.76, 354.118, 31, 90.0185, 736, 19.13, 347.09, 32, 90.0157, 783, 20.79, 350.504}
```

```
32
```

```
{771, 749, 773, 751, 731}
```

```
749.
```

In this section I calculate the basic statistics, such as the mean and median. My objective is to ensure the data has been collected and entered correctly and also to compare the two datasets to see if they appear to be different.

```
myData = Table[
   {
    ssName[ssidx], N[Mean[ssWork[ssidx]]],
    Mean[ssL[ssidx]], Mean[ssSt[ssidx]], Mean[ssQt[ssidx]], Mean[ssRt[ssidx]],
    Length[ssWork[ssidx]], N[StandardDeviation[ssL[ssidx]]], N[StandardDeviation[ssSt[ssidx]]],
    N[StandardDeviation[ssQt[ssidx]]], N[StandardDeviation[ssRt[ssidx]]]
   }, {ssidx, 1, ssNum}
  ];
toGrid = Prepend[myData, {"Settings", "Avg Work\n(cmt)", "Avg L\n(cmt/ms)",
    "Avg CPUt\n(ms/cmt)", "Avg Wt\n(ms/cmt)", "Avg Rt\n(ms/cmt)", "Samples",
    "Stdev L\n(cmt/ms)", "Stdev CPUt\n(ms/cmt)", "Stdev Wt\n(ms/cmt)", "Stdev Rt\n(ms/cmt)"}];
Grid[
 toGrid,
 Frame →
  All]
```

| Settings | Avg Work (cmt) | Avg L (cmt/ms) | Avg CPUt (ms/cmt) | Avg Wt (ms/cmt) | Avg Rt (ms/cmt) | Samples | Stdev L (cmt/ms) | Stdev CPUt (ms/cmt) | Stdev Wt (ms/cmt) | Stdev Rt (ms/cmt) |
|---|---|---|---|---|---|---|---|---|---|---|
| Wait Immediate | 749. | 0.00831995 | 470.976 | 27.6505 | 498.626 | 32 | 0.0002117009 | 12.4358 | 1.80699 | 12.6189 |
| Wait Batch | 761.656 | 0.00845821 | 465.452 | 27.2719 | 492.724 | 32 | 0.0002110098 | 11.4239 | 1.19211 | 11.4096 |
| Nowait Immediate | 762.75 | 0.00847372 | 465.186 | 1.33771 | 466.523 | 32 | 0.0002141156 | 12.6153 | 3.1666 | 11.7575 |
| Nowait Batch | 762.406 | 0.00846991 | 467.24 | 0.196875 | 467.436 | 32 | 0.0001889893 | 10.2195 | 0.0844592 | 10.2097 |

In this section I calculate the key parameters for understanding the change. The columns heading are a mix of queuing theory and Oracle-ease centric.

```
myData = Table[
   {
    ssName[ssidx],
    N[Mean[ssWork[ssidx]]], N[100 * (Mean[ssWork[ssidx]] - Mean[ssWork[1]]) / Mean[ssWork[1]]],
    Mean[ssL[ssidx]], 100 * (Mean[ssL[ssidx]] - Mean[ssL[1]]) / Mean[ssL[1]],
    Mean[ssSt[ssidx]], 100 * (Mean[ssSt[ssidx]] - Mean[ssSt[1]]) / Mean[ssSt[1]],
    Mean[ssQt[ssidx]], 100 * (Mean[ssQt[ssidx]] - Mean[ssQt[1]]) / Mean[ssQt[1]],
    Mean[ssRt[ssidx]], 100 * (Mean[ssRt[ssidx]] - Mean[ssRt[1]]) / Mean[ssRt[1]],
    Length[ssWork[ssidx]]
   }, {ssidx, 1, ssNum}
  ];
toGrid = Prepend[myData,
   {"Settings", "Avg Work\n(cmt)", "%\nChange", "Avg L\n(cmt/ms)", "%\nChange", "Avg CPUt\n(ms/cmt)",
    "%\nChange", "Avg Wt\n(ms/cmt)", "%\nChange", "Avg Rt\n(ms/cmt)", "%\nChange", "Samples"}];
Grid[
 toGrid,
 Frame →
  All]
```

| Settings | Avg Work (cmt) | % Change | Avg L (cmt/ms) | % Change | Avg CPUt (ms/cmt) | % Change | Avg Wt (ms/cmt) | % Change | Avg Rt (ms/cmt) | % Change | Samples |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Wait Immediate | 749. | 0. | 0.0083199 5 | 0. | 470.976 | 0. | 27.6505 | 0. | 498.626 | 0. | 32 |
| Wait Batch | 761.656 | 1.689 75 | 0.0084582 1 | 1.661 79 | 465.452 | -1.17 276 | 27.2719 | -1.36 94 | 492.724 | -1.18 367 | 32 |
| Nowait Immediate | 762.75 | 1.835 78 | 0.0084737 2 | 1.848 25 | 465.186 | -1.22 937 | 1.33771 | -95.1 621 | 466.523 | -6.43 826 | 32 |
| Nowait Batch | 762.406 | 1.789 89 | 0.0084699 1 | 1.802 49 | 467.24 | -0.79 327 6 | 0.196 875 | -99.2 88 | 467.436 | -6.25 515 | 32 |

In this section I'm showing change and if that change is statictially significant, all compared to our baseline sample set, wait,immediate. This is pretty cool: As mentioned in the Normality Tests section below, distributions in a ttest must be normal. If not, then a location test can be performed. If you look closely at the code segment, below I tested and adjusted for this.

```
myData = Table[
   {
    ssName[ssidx],
    100 * (Mean[ssL[ssidx]] - Mean[ssL[1]]) / Mean[ssL[1]],
    If[DistributionFitTest[ssL[ssidx]] ≥ 0.050,
     TTest[{ssL[1], ssL[ssidx]}],
     LocationEquivalenceTest[{ssL[1], ssL[ssidx]}]
    ],
    100 * (Mean[ssRt[ssidx]] - Mean[ssRt[1]]) / Mean[ssRt[1]],
    If[DistributionFitTest[ssRt[ssidx]] ≥ 0.050,
     TTest[{ssRt[1], ssRt[ssidx]}],
     LocationEquivalenceTest[{ssRt[1], ssRt[ssidx]}]
    ]
   }, {ssidx, 1, ssNum}
  ];
toGrid = Prepend[myData, {"Settings", "Avg L\n% Change",
    "Avg L\nTtest pvalue", "Avg Rt\n% Change", "Avg Rt\nTtest pvalue"}];
Grid[
 toGrid,
 Frame →
  All]
```

| Settings | Avg L % Change | Avg L Ttest pvalue | Avg Rt % Change | Avg Rt Ttest pvalue |
|---|---|---|---|---|
| Wait Immediate | 0. | 1. | 0. | 1. |
| Wait Batch | 1.66179 | 0.0119672 | -1.18367 | 0.0541915 |
| Nowait Immediate | 1.84825 | 0.0245961 | -6.43826 | $1.95955 \times 10^{-15}$ |
| Nowait Batch | 1.80249 | 0.00458214 | -6.25515 | $5.38924 \times 10^{-16}$ |

## Sample Set Normality Tests

Before we can perform a standard t-test hypothesis tests on our data, we need to ensure it is normally distributed...because that is one of the underlying assumptions and requirements for properly performing a t-test.

### Statistical and vIsual normality test

Our alpha will be 0.05, so if the distribution fit test results in a value greater than 0.05 then we can assume the data set is indeed normally distributed.

The first test is just to double check to make sure my thinking is correct. Since I creating a normal distribution based on a mean and standard deviation (just happens to be based on the my sample set data), I would expect a p-value (the result) to greatly exceed 0.05. Notice that the more samples I have created (the final number), the closer the p-value approaches 1.0.

```
check = DistributionFitTest[
   RandomVariate[NormalDistribution[Mean[ssSt[1]], StandardDeviation[ssSt[1]]], 10000]];
Print["This number should be much greater than 0.05: ", check,
  " If not try again by re-evaluating."];
Do[
  Print["=============================================================================="];
  Print[ssName[i], ", ", Length[ssSt[i]], " sample values"];
  pValueWork = DistributionFitTest[ssWork[i]];
  pValueL = DistributionFitTest[ssL[i]];
  pValueSt = DistributionFitTest[ssSt[i]];
  pValueQt = DistributionFitTest[ssQt[i]];
  pValueRt = DistributionFitTest[ssRt[i]];
  Print["Work pvalue=", pValueWork];
  Print[
   Histogram[ssWork[i], PlotLabel → "Occurances vs Work", AxesLabel → {"Sample value", "Occurs"}]];
  Print["----------------------------"];
  Print["L pvalue=", pValueL];
  Print[Histogram[ssL[i], PlotLabel → "Occurances vs L", AxesLabel → {"Sample value", "Occurs"}]];
  Print["----------------------------"];
  Print["St pvalue=", pValueSt];
  Print[Histogram[ssSt[i], PlotLabel → "Occurances vs St", AxesLabel → {"Sample value", "Occurs"}]];
  Print["----------------------------"];
  Print["Qt pvalue=", pValueQt];
  Print[Histogram[ssQt[i], PlotLabel → "Occurances vs Qt", AxesLabel → {"Sample value", "Occurs"}]];
  Print["----------------------------"];
  Print["Rt pvalue=", pValueRt];
  Print[Histogram[ssRt[i], PlotLabel → "Occurances vs Rt", AxesLabel → {"Sample value", "Occurs"}]];
  , {i, 1, ssNum}
 ];
```

```
This number should be much greater than 0.05: 0.0889051 If not try again by re-evaluating.

==============================================================================

Wait Immediate, 32 sample values

Work pvalue=0.742428
```

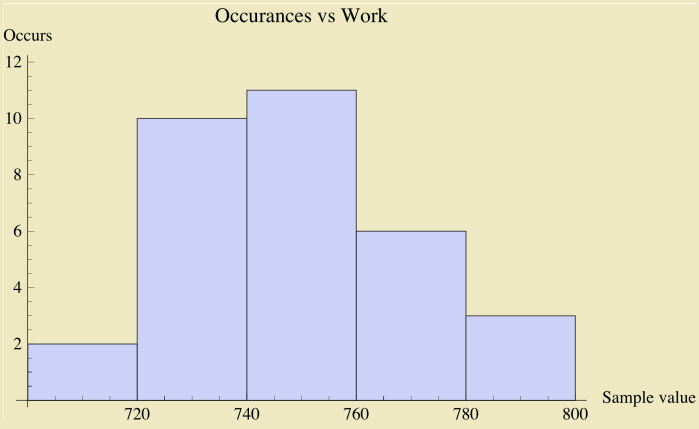


```
----------------------------

L pvalue=0.687486
```

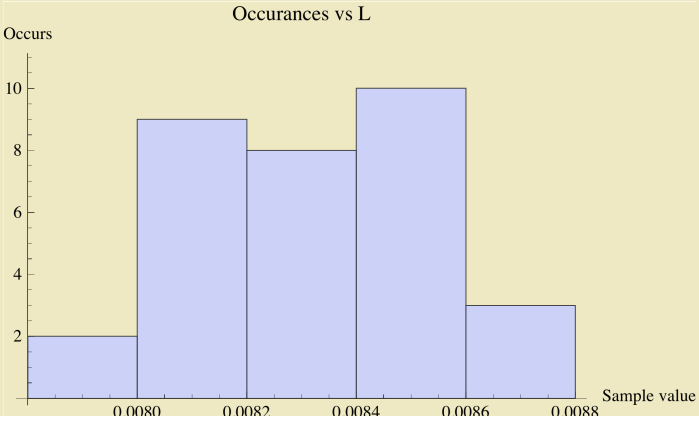

---

St pvalue=0.892997

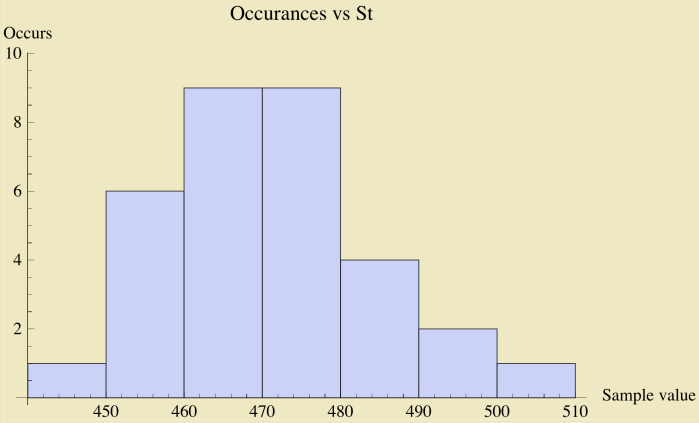


---

Qt pvalue=0.607359

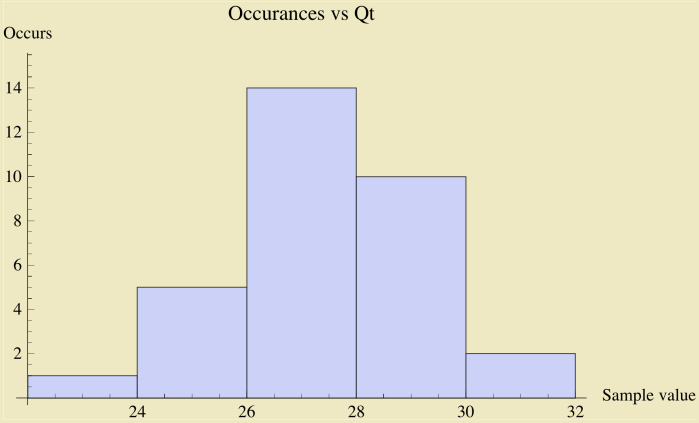


---

Rt pvalue=0.823046

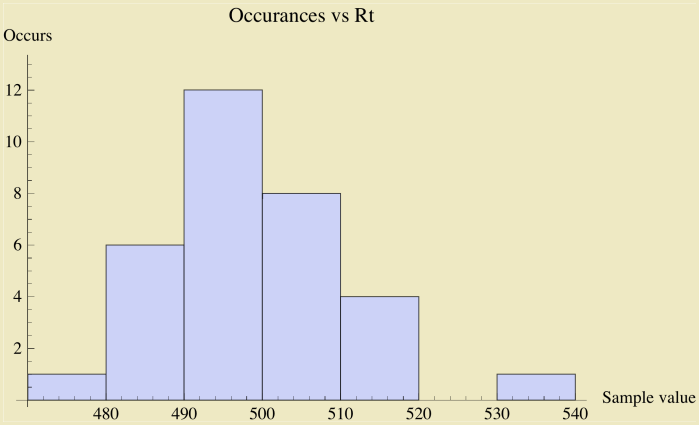


==========

Wait Batch, 32 sample values

Work pvalue=0.273505

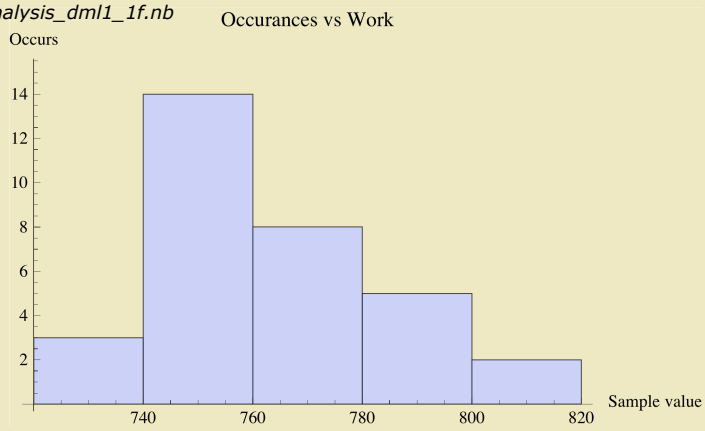


---

L pvalue=0.30573

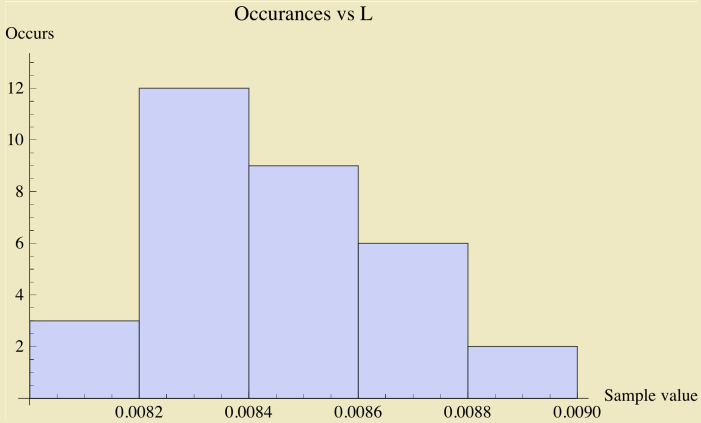


---

St pvalue=0.448043

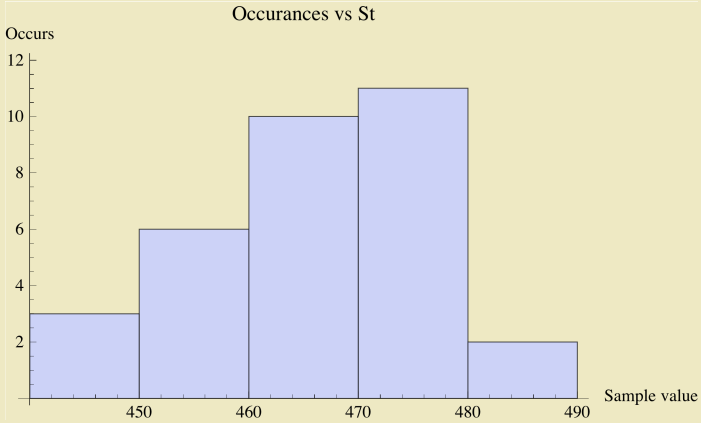


---

Qt pvalue=0.023525

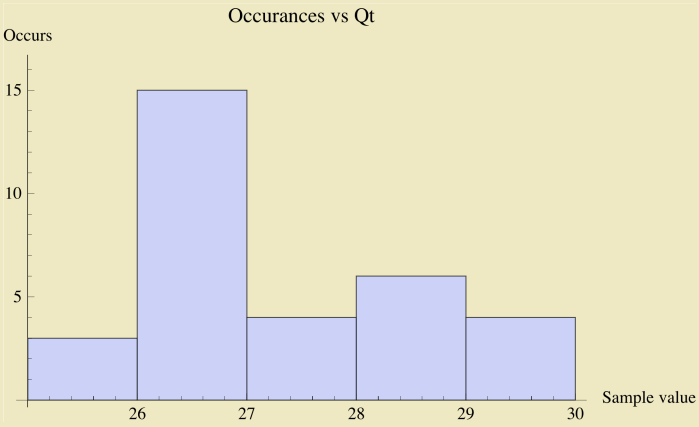


---

Rt pvalue=0.140102

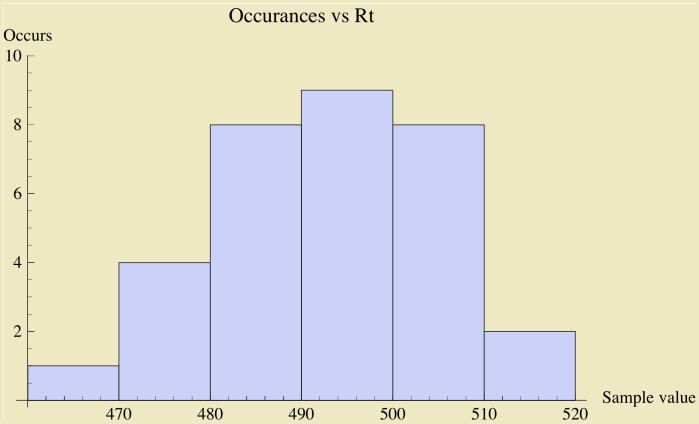


==========

Nowait Immediate, 32 sample values

Work pvalue=0.0117987

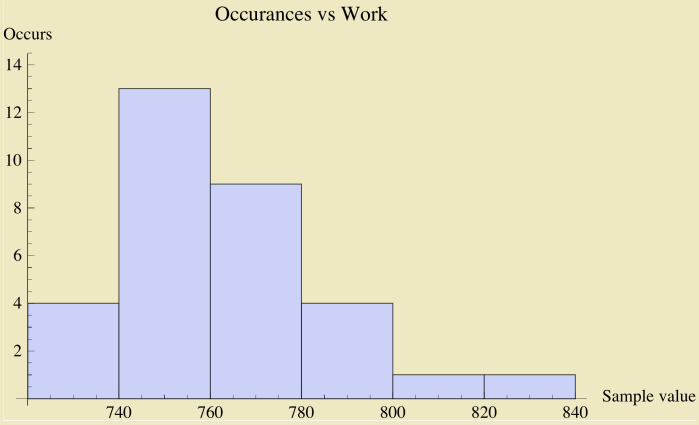


---

L pvalue=0.0119974

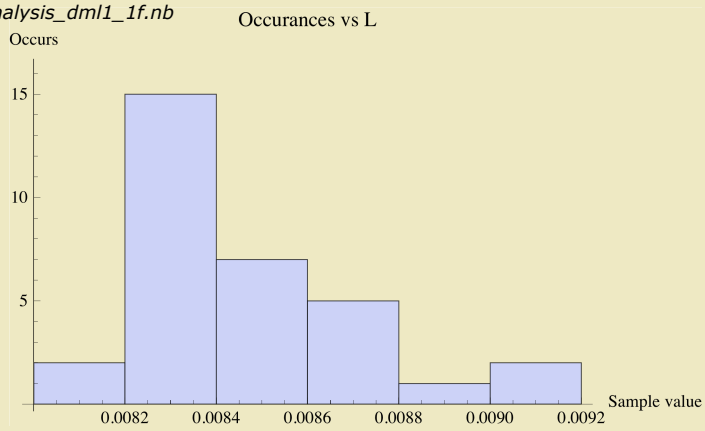


---

St pvalue=0.140427

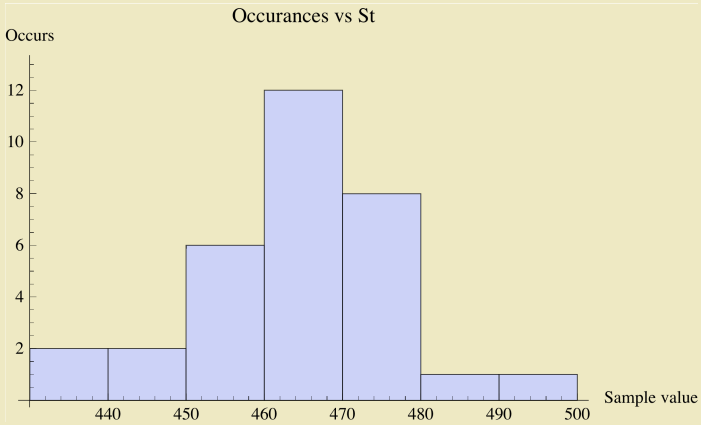


---

Qt pvalue=0

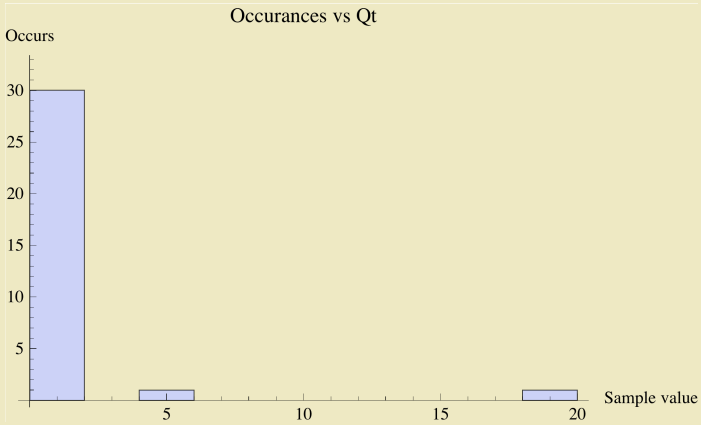


---

Rt pvalue=0.473645

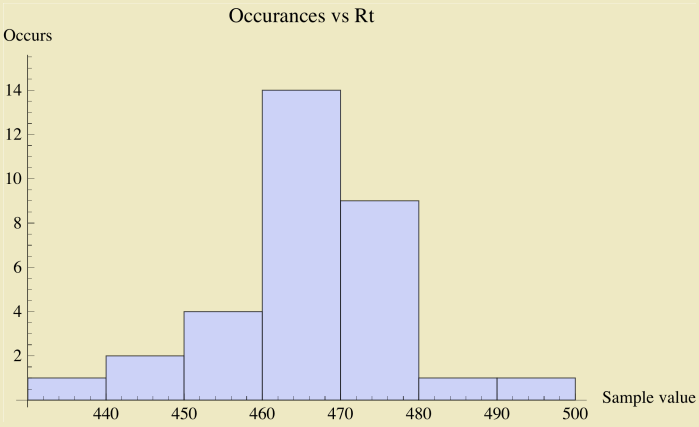


==========

Nowait Batch, 32 sample values

Work pvalue=0.849278

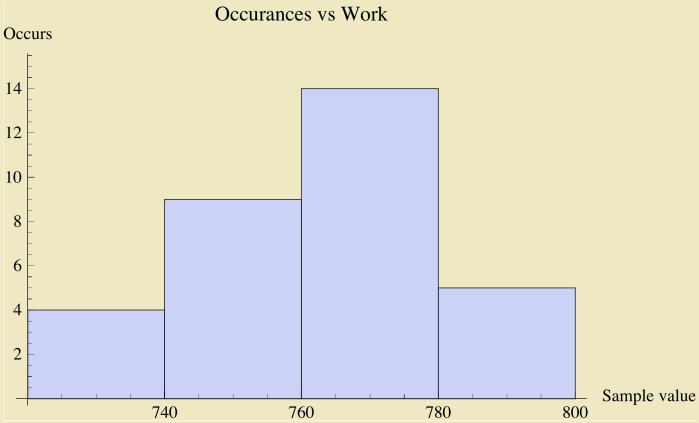


---

L pvalue=0.849148

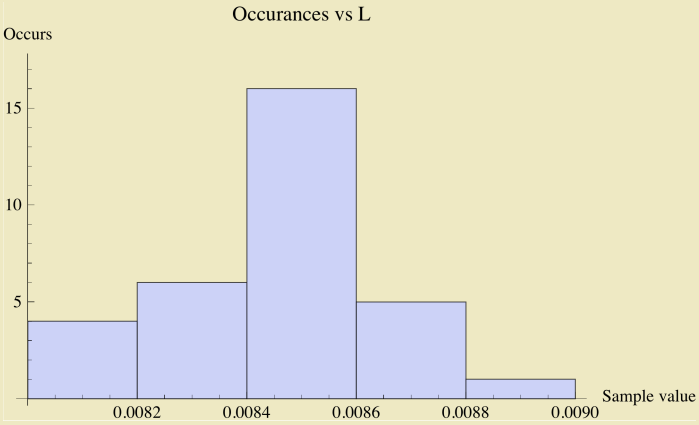


---

St pvalue=0.374021

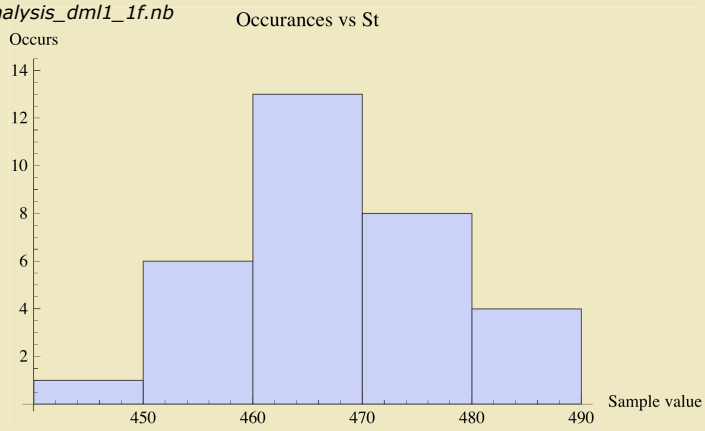


```
----------------------------

Qt pvalue=0.251681
```

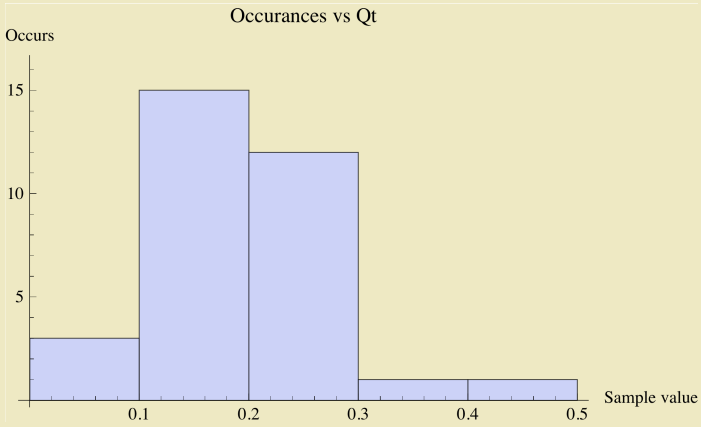


```
----------------------------

Rt pvalue=0.380833
```

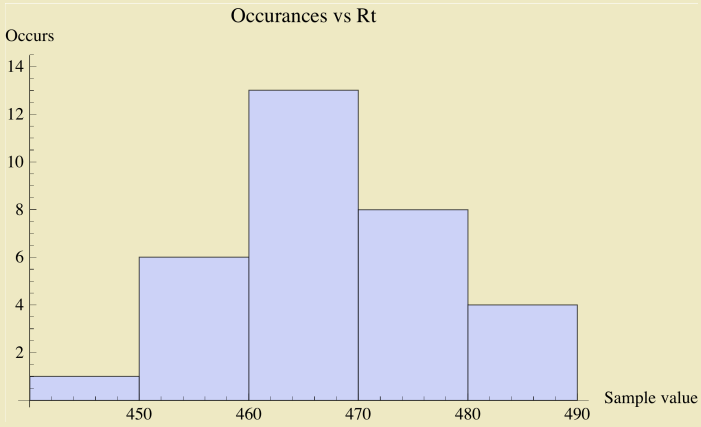


## Sample Comparison Tests (when normality exists)

Assuming our samples **are normally distributed**, now it's time to see if they are significantly different. If so, then we know changing the commit write optoins indeed makes a significant performance difference...at least statistically.

The null hypothesis is; there is no real difference between our samples sets. We need to statistically prove that any difference is the result of randomness; like we just happened to pick poor set of samples and it makes their difference look much worse than it really is.

A t-test will produce a statistic p. The p value is a probability, with a value ranging from zero to one. It is the answer to this question: If the populations really have the same mean overall, what is the probability that random sampling would lead to a difference between sample means larger than observed?

For example, if the p value is 0.03 we can say a random sampling from identical populations would lead to a difference smaller than you observed in 97% of the experiments and larger than you observed in 3% of the experiments.

Said another way, suppose I have a single sample set and I copy it, resultling in two identical sample sets. Now suppose we perform a significance test on these two identical sample sets. The resuting p-value will be 1.0 because they are exactly the same. We are essentially doing the same thing here except we have to different sample sets... but we still want to see if they "like" each other..and in our case we hope they are NOT like each other, which means the p-value will low... below our cut off value of 0.05.

For our analysis we choose alpha of 0.05. To accept that our two samples are statistically similar the p value would need to be less than 0.05 (our alpha).

Good reference about the P-Value and significance testing: http://www.graphpad.com/articles/pvalue.htm

Here we go (assuming our samples are normally distributed):

1. Our P value threshold is 0.05, which is our alpha.
2. The null hypothesis is the two populations have the same mean. (Remember we have to sample sets, which not the population.)
3. Do the statistical test to compute the P value.
4. Compare the result P value to our threshold alpha value. If the P value is less then our threshold, we will reject the null hypothesis and say the difference between our samples is significant. However, if the P value is greater than the threshold, we cannot reject the null hypothesis and any difference between our samples are not statistically significant.

```
Print["P-values assumming normality exists."];
Table[
  Table[
    If[i ≠ j,
      Print["================================================================"];
      Print[ssName[i], "(", Length[ssL[i]], ") and ", ssName[j], " (", Length[ssL[j]], ")"];
      pValueWork = TTest[{ssWork[i], ssWork[j]}];
      Print["Work: ", pValueWork];
      pValueL = TTest[{ssL[i], ssL[j]}];
      Print["L: ", pValueL];
      pValueSt = TTest[{ssSt[i], ssSt[j]}];
      Print["St: ", pValueSt];
      pValueQt = TTest[{ssQt[i], ssQt[j]}];
      Print["Qt: ", pValueQt];
      pValueRt = TTest[{ssRt[i], ssRt[j]}];
      Print["Rt: ", pValueRt];
     ];
    ,
    {j, 1, ssNum}
   ];
  , {i, 1, ssNum}
 ];
```

```
P-values assumming normality exists.

================================================================

Wait Immediate(32) and Wait Batch (32)

Work: 0.0108505

L: 0.0119672

St: 0.0690354

TTest::nortst : At least one of the p-values in {0.607359, 0.023525}, resulting from
      a test for normality, is below 0.025`. The tests in {T} require that the data is normally distributed. ≫

Qt: 0.326288

Rt: 0.0541915

================================================================

Wait Immediate(32) and Nowait Immediate (32)
```

TTest::nortst : At least one of the p-values in {0.742428, 0.0117987}, resulting from a test for normality, is below 0.025`. The tests in {T} require that the data is normally distributed. ≫

Work: 0.00980034

TTest::nortst : At least one of the p-values in {0.687486, 0.0119974}, resulting from a test for normality, is below 0.025`. The tests in {T} require that the data is normally distributed. ≫

General::stop : Further output of TTest::nortst will be suppressed during this calculation. ≫

L: 0.00938465

St: 0.0692285

Qt: $1.21047 \times 10^{-39}$

Rt: $1.95955 \times 10^{-15}$

==========

Wait Immediate(32) and Nowait Batch (32)

Work: 0.00481099

L: 0.00458214

St: 0.194013

Qt: $1.45446 \times 10^{-38}$

Rt: $5.38924 \times 10^{-16}$

==========

Wait Batch(32) and Wait Immediate (32)

Work: 0.0108505

L: 0.0119672

St: 0.0690354

Qt: 0.326288

Rt: 0.0541915

==========

Wait Batch(32) and Nowait Immediate (32)

Work: 0.831121

L: 0.784703

St: 0.929672

Qt: $5.28216 \times 10^{-35}$

Rt: $6.18455 \times 10^{-13}$

==========

Wait Batch(32) and Nowait Batch (32)

Work: 0.868893

L: 0.81588

St: 0.511946

Qt: $3.6158 \times 10^{-44}$

Rt: $1.92621 \times 10^{-13}$

==========

Nowait Immediate(32) and Wait Immediate (32)

Work: 0.00980034

L: 0.00938465

St: 0.0692285

Qt: $1.21047 \times 10^{-39}$

Rt: $1.95955 \times 10^{-15}$

==========

Nowait Immediate(32) and Wait Batch (32)

Work: 0.831121

L: 0.784703

St: 0.929672

Qt: $5.28216 \times 10^{-35}$

Rt: $6.18455 \times 10^{-13}$

==========

Nowait Immediate(32) and Nowait Batch (32)

Work: 0.944143

L: 0.944292

St: 0.476899

Qt: 0.0502241

Rt: 0.741235

==========

Nowait Batch(32) and Wait Immediate (32)

Work: 0.00481099

L: 0.00458214

St: 0.194013

Qt: $1.45446 \times 10^{-38}$

Rt: $5.38924 \times 10^{-16}$

==========

Nowait Batch(32) and Wait Batch (32)

Work: 0.868893

L: 0.81588

St: 0.511946

Qt: $3.6158 \times 10^{-44}$

Rt: $1.92621 \times 10^{-13}$

==========

Nowait Batch(32) and Nowait Immediate (32)

Work: 0.944143

L: 0.944292

St: 0.476899

Qt: 0.0502241

Rt: 0.741235

If the above T-Test results (p value) are less then our threshold we can say there is a significant difference between the two sample sets.

## Sample Comparison Tests (when normality may NOT exist)

http://www.statsoft.com/textbook/nonparametric-statistics

The paragraph below (which is from the reference above) is a key reference to what we're doing here:

...the need is evident for statistical procedures that enable us to process data of "low quality," from small samples, on variables about which nothing is known (concerning their distribution). Specifically, nonparametric methods were developed to be used in cases when the researcher knows nothing about the parameters of the variable of interest in the population (hence the name nonparametric). In more technical terms, nonparametric methods do not rely on the estimation of parameters (such as the mean or the standard deviation) describing the distribution of the variable of interest in the population. Therefore, these methods are also sometimes (and more appropriately) called parameter-free methods or distribution-free methods.

Being that I'm not a statistician but still need to determine if these sample sets are significant different, I let *Mathematica* determine the appropriate test. Notice that one of the above mentioned tests will probably be the test *Mathematica* chooses.

Note: If we run our normally distributed data through this analysis (speically, the "LocationEquivalenceTest"), *Mathematica* should detect this and use a more appropriate significant test, like a t-test.
Here we go with the hypothesis testing (assuming our sample sets are not normally distributed):

1. Our P value threshold is 0.05, which is our alpha.
2. The null hypotheses is the two populations have the same mean. (Remember we have to sample sets, which is not the population.)
3. Do the statistical test to compute the P value.
4. Compare the result P value to our threshold alpha value. If the P value is less then our threshold, we will reject the null hypothesis and say the difference between our samples is significant. (Which is what I'm hoping to see.) However, if the P value is greater than the threshold, we cannot reject the null hypothesis and any difference between our samples are not statistically significant; randomness, picked the "wrong" samples, etc.

```
Print["P-values assumming normality MAY not exist."];
Table[
  Table[
    If[i ≠ j,
      Print["================================================================"];
      Print[ssName[i], "(", Length[ssL[i]], ") and ", ssName[j], " (", Length[ssL[j]], ")"];

      test1 = MannWhitneyTest[{ssWork[i], ssWork[j]}];
      test2 = LocationEquivalenceTest[{ssWork[i], ssWork[j]}, {"TestDataTable", "AutomaticTest"}];
      Print["Work: Test1=", test1, " Test2=", test2];
      Print[SmoothHistogram[{ssWork[i], ssWork[j]}]];

      Print["------------"];
      test1 = MannWhitneyTest[{ssL[i], ssL[j]}];
      test2 = LocationEquivalenceTest[{ssL[i], ssL[j]}, {"TestDataTable", "AutomaticTest"}];
      Print["L: Test1=", test1, " Test2=", test2];
      Print[SmoothHistogram[{ssL[i], ssL[j]}]];

      Print["------------"];
      test1 = MannWhitneyTest[{ssSt[i], ssSt[j]}];
      test2 = LocationEquivalenceTest[{ssSt[i], ssSt[j]}, {"TestDataTable", "AutomaticTest"}];
      Print["St: Test1=", test1, " Test2=", test2];
      Print[SmoothHistogram[{ssSt[i], ssSt[j]}]];

      Print["------------"];
      test1 = MannWhitneyTest[{ssQt[i], ssQt[j]}];
      test2 = LocationEquivalenceTest[{ssQt[i], ssQt[j]}, {"TestDataTable", "AutomaticTest"}];
      Print["Qt: Test1=", test1, " Test2=", test2];
      Print[SmoothHistogram[{ssQt[i], ssQt[j]}]];

      Print["------------"];
      test1 = MannWhitneyTest[{ssRt[i], ssRt[j]}];
      test2 = LocationEquivalenceTest[{ssRt[i], ssRt[j]}, {"TestDataTable", "AutomaticTest"}];
      Print["St: Test1=", test1, " Test2=", test2];
      Print[SmoothHistogram[{ssRt[i], ssRt[j]}]];

     ];
    ,
    {j, 1, ssNum}
   ];
  , {i, 1, ssNum}
 ];
```

P-values assumming normality MAY not exist.

==========

Wait Immediate(32) and Wait Batch (32)

Work: Test1=0.0180978 Test2={ | Statistic | P-Value | , KSampleT}

| K-Sample T | 6.8993 | 0.0108505 |

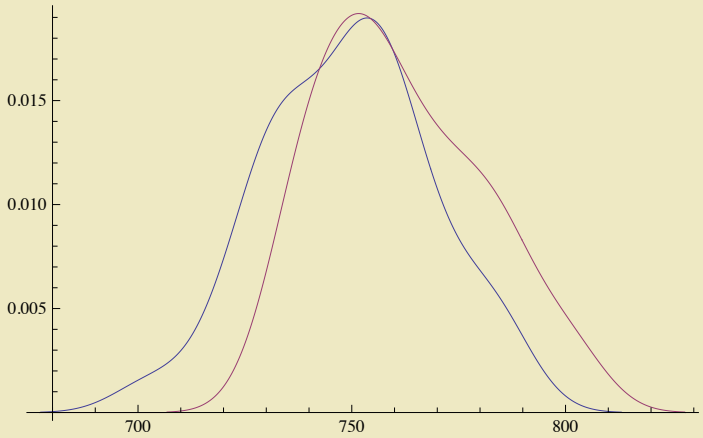


---

L: Test1=0.0191274 Test2={ | Statistic | P-Value | , KSampleT}

| K-Sample T | 6.70476 | 0.0119672 |

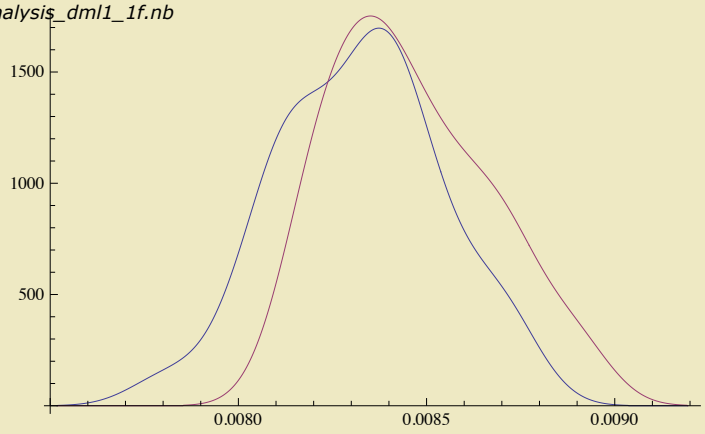


---

St: Test1=0.130902 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 3.42364 | 0.0690354 |

, KSampleT}

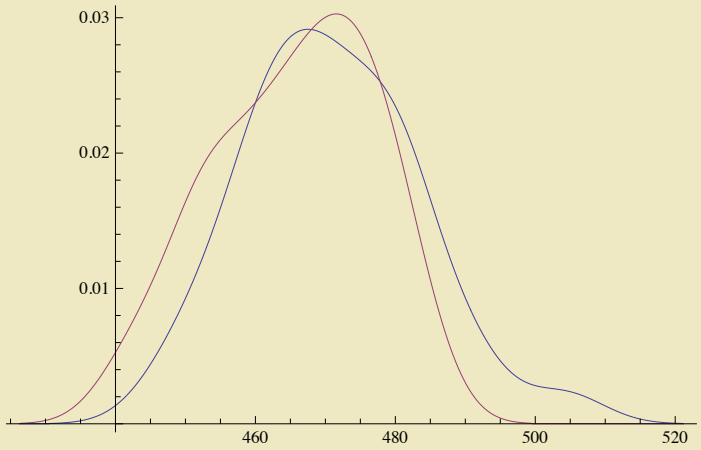


---

Qt: Test1=0.386462 Test2={

| | Statistic | P-Value |
|---|---|---|
| Kruskal-Wallis | 0.761719 | 0.387064 |

, KruskalWallis}

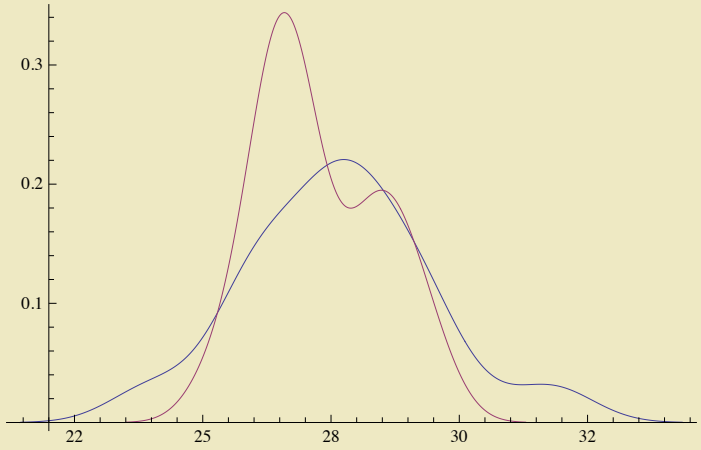


---

St: Test1=0.114637 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 3.85155 | 0.0541915 |

, KSampleT}

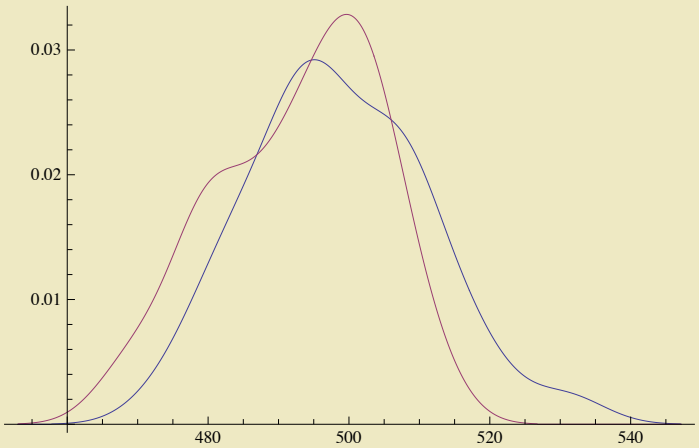


==========

Wait Immediate(32) and Nowait Immediate (32)

Work: Test1=0.0236441 Test2={

| | Statistic | P-Value |
|---|---|---|
| Kruskal-Wallis | 5.09219 | 0.0227979 |

, KruskalWallis}

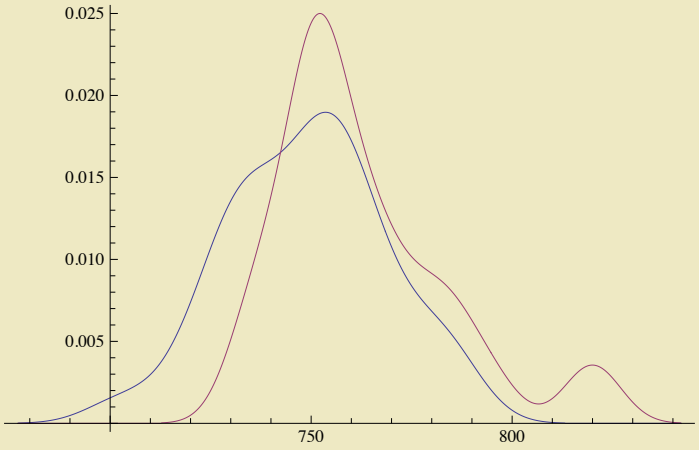


---

L: Test1=0.0253766 Test2={

| | Statistic | P-Value |
|---|---|---|
| Kruskal-Wallis | 4.96803 | 0.0245961 |

, KruskalWallis}

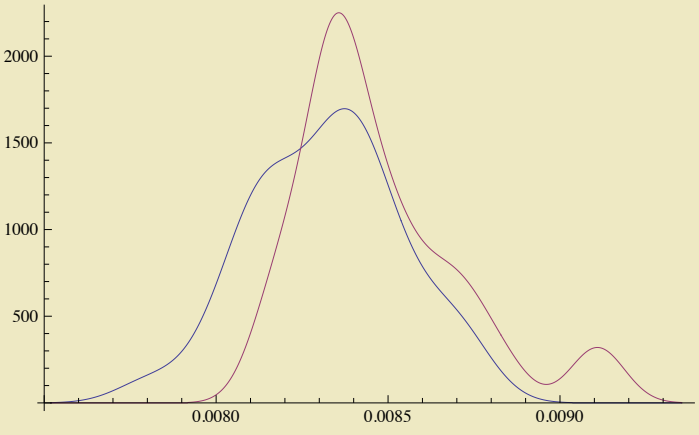


---

St: Test1=0.172928 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 3.41876 | 0.0692285 |

, KSampleT}

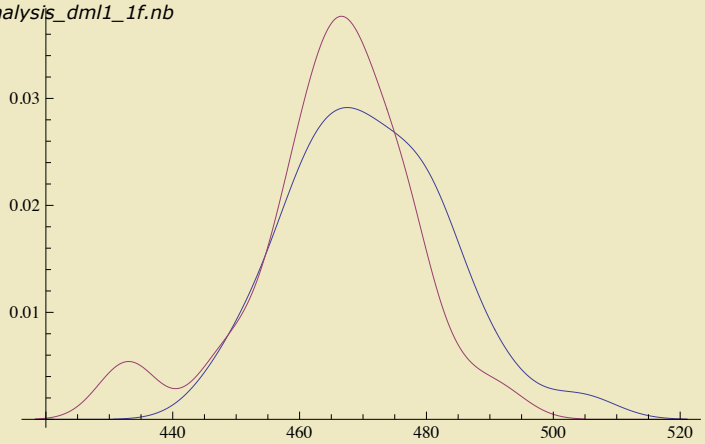


---

Qt: Test1=$6.51131 \times 10^{-12}$ Test2={

| | Statistic | P-Value |
|---|---|---|
| Kruskal-Wallis | 47.2615 | $2.45726 \times 10^{-20}$ |

, KruskalWallis}

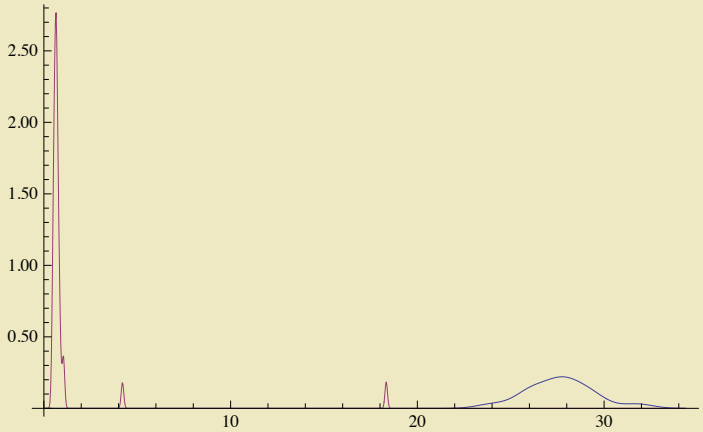


---

St: Test1=$7.08328 \times 10^{-11}$ Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 110.862 | $1.95955 \times 10^{-15}$ |

, KSampleT}

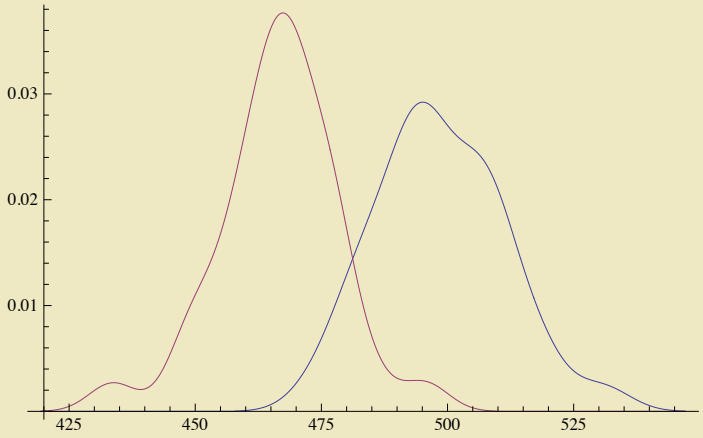


==========

Wait Immediate(32) and Nowait Batch (32)

Work: Test1=0.00388399 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 8.55364 | 0.00481099 |

, KSampleT}

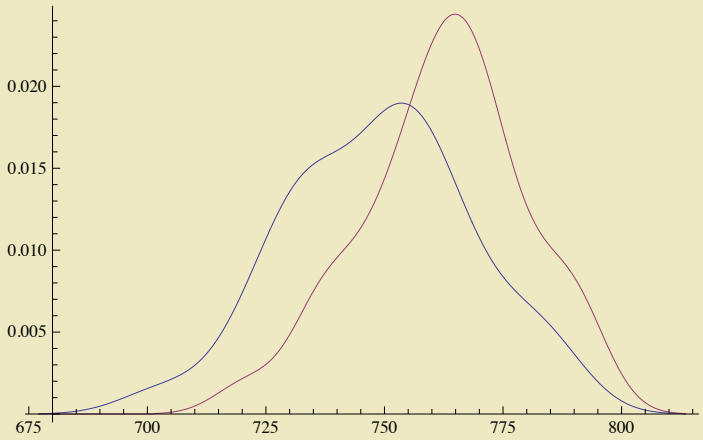


---

L: Test1=0.00349576 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 8.65493 | 0.00458214 |

, KSampleT}

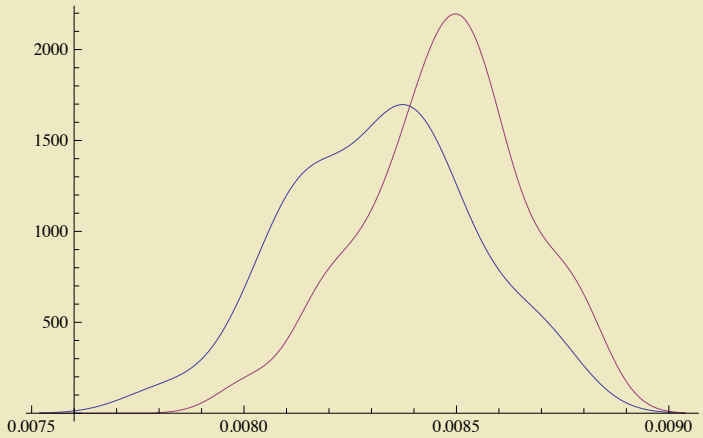


---

St: Test1=0.273818 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 1.72404 | 0.194013 |

, KSampleT}

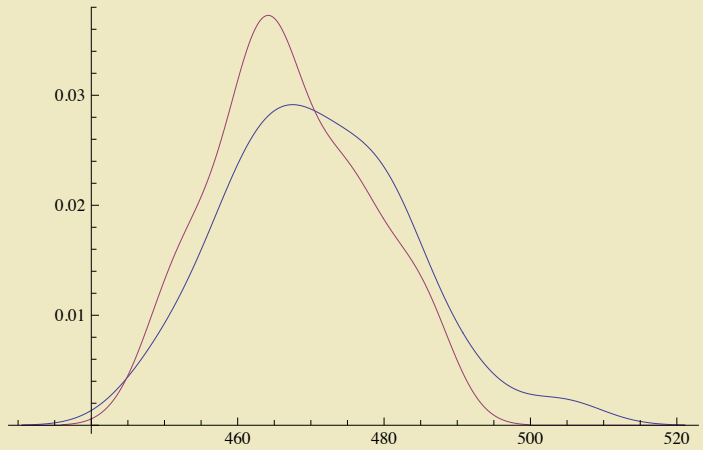


---

Qt: Test1=$6.50414 \times 10^{-12}$ Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 7370.44 | $3.67139 \times 10^{-66}$ |

, KSampleT}

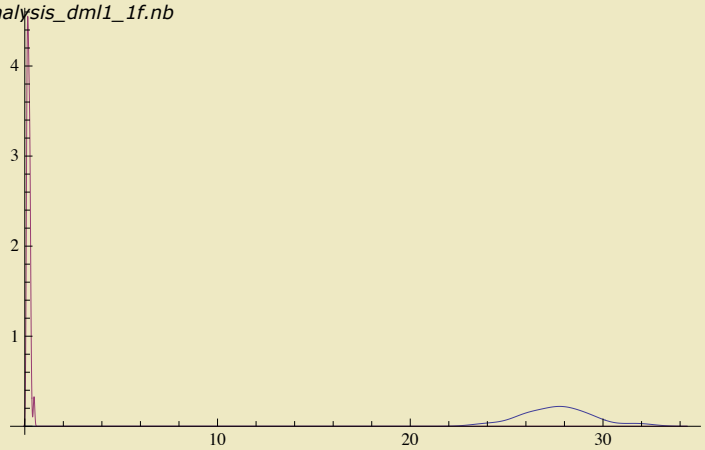


---

St: Test1=7.08328 × 10^-11 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 118.151 | 5.38924 × 10^-16 |

, KSampleT}

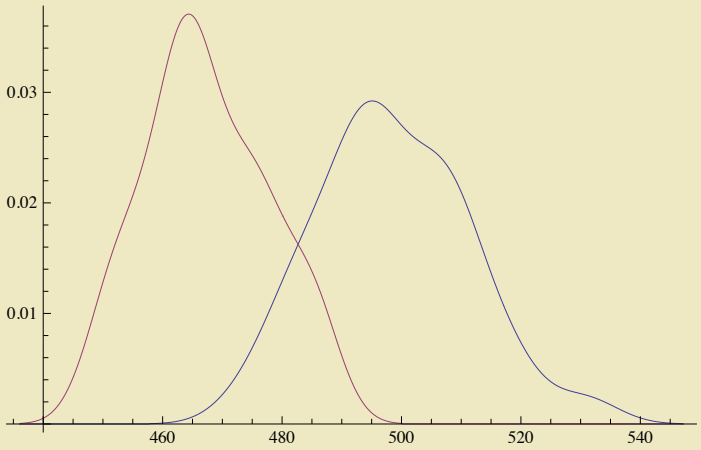


==========

Wait Batch(32) and Wait Immediate (32)

Work: Test1=0.0187642 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 6.8993 | 0.0108505 |

, KSampleT}

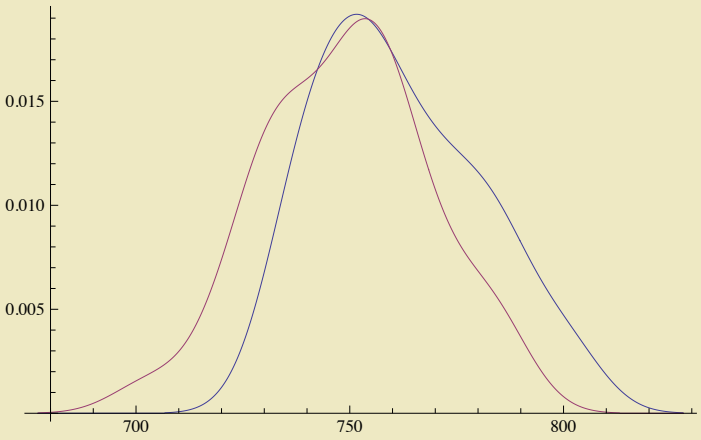


---

L: Test1=0.0198267 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 6.70476 | 0.0119672 |

, KSampleT}

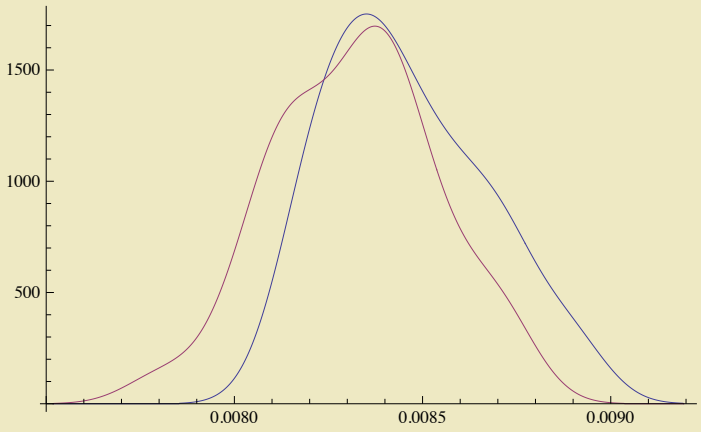


---

St: Test1=0.127513 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 3.42364 | 0.0690354 |

, KSampleT}

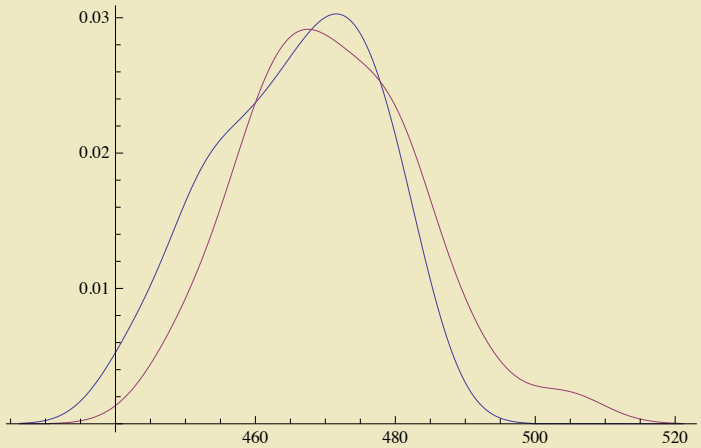


---

Qt: Test1=0.379142 Test2={

| | Statistic | P-Value |
|---|---|---|
| Kruskal-Wallis | 0.761719 | 0.387064 |

, KruskalWallis}

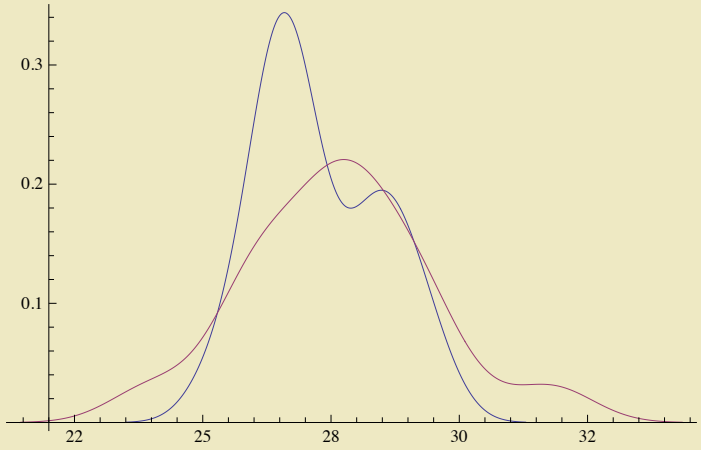


---

St: Test1=0.111583 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 3.85155 | 0.0541915 |

, KSampleT}

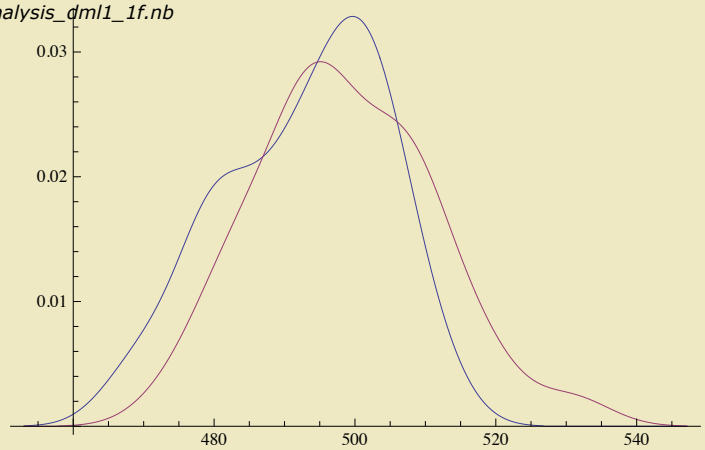


==========

Wait Batch(32) and Nowait Immediate (32)

Work: Test1=0.989285 Test2={

| | Statistic | P-Value |
|---|---|---|
| Kruskal-Wallis | 0.000406058 | 0.984115 |

, KruskalWallis}

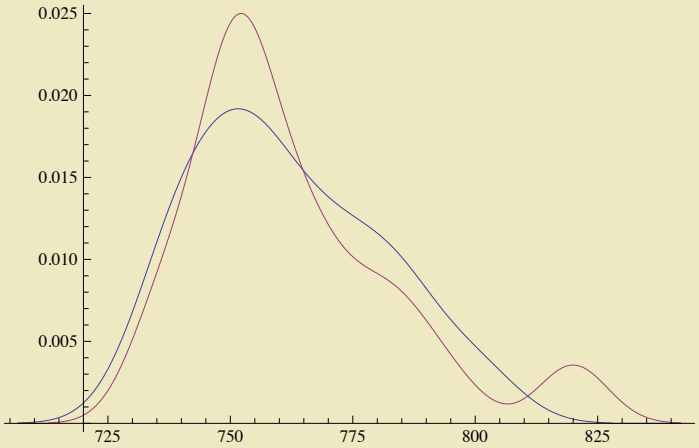


---

L: Test1=0.951819 Test2={

| | Statistic | P-Value |
|---|---|---|
| Kruskal-Wallis | 0.00288462 | 0.957679 |

, KruskalWallis}

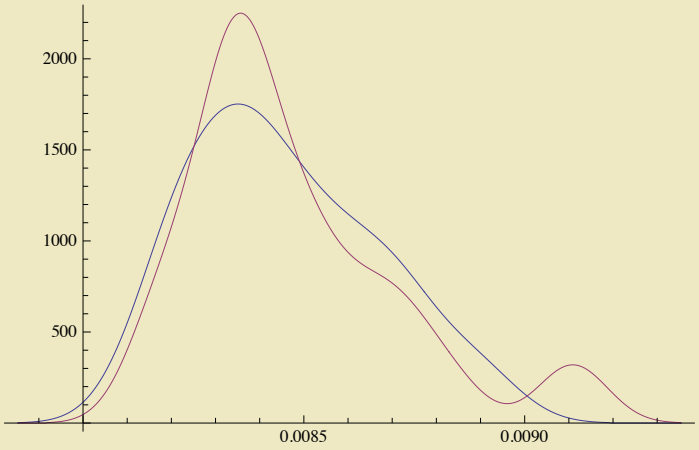


---

St: Test1=0.909134 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 0.00785287 | 0.929672 |

, KSampleT}

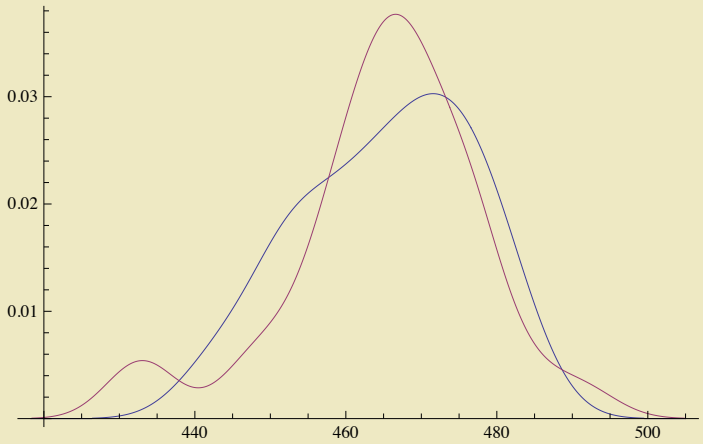


---

Qt: Test1=$6.51131 \times 10^{-12}$ Test2=

| | Statistic | P-Value |
|---|---|---|
| Kruskal-Wallis | 47.2615 | $2.45726 \times 10^{-20}$ |

, KruskalWallis

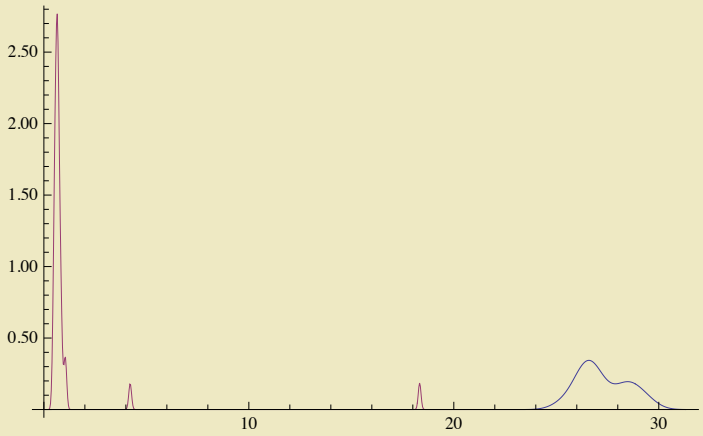


---

St: Test1=$6.83863 \times 10^{-10}$ Test2=

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 81.8401 | $6.18455 \times 10^{-13}$ |

, KSampleT

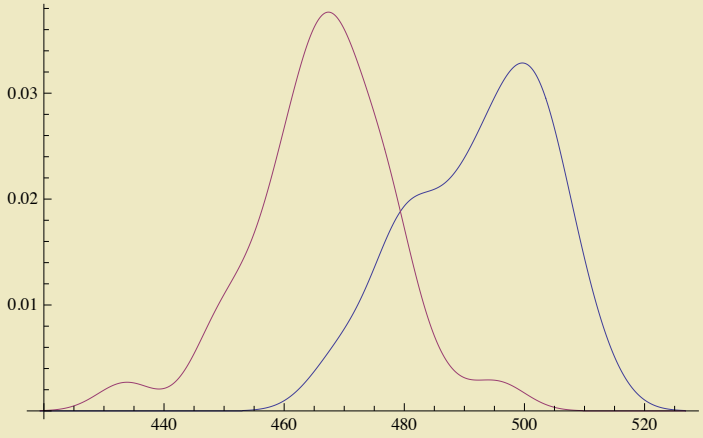


==========

Wait Batch(32) and Nowait Batch (32)

Work: Test1=0.609815 Test2=

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 0.0274729 | 0.868893 |

, KSampleT

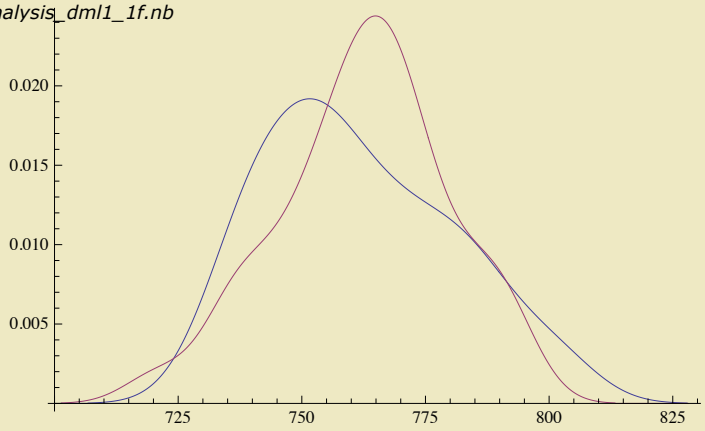


---

L: Test1=0.568234 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 0.0546811 | 0.81588 |

, KSampleT}

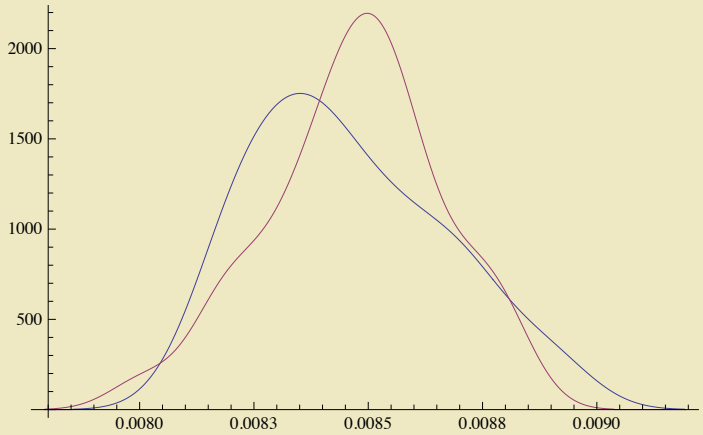


---

St: Test1=0.6146 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 0.435088 | 0.511946 |

, KSampleT}

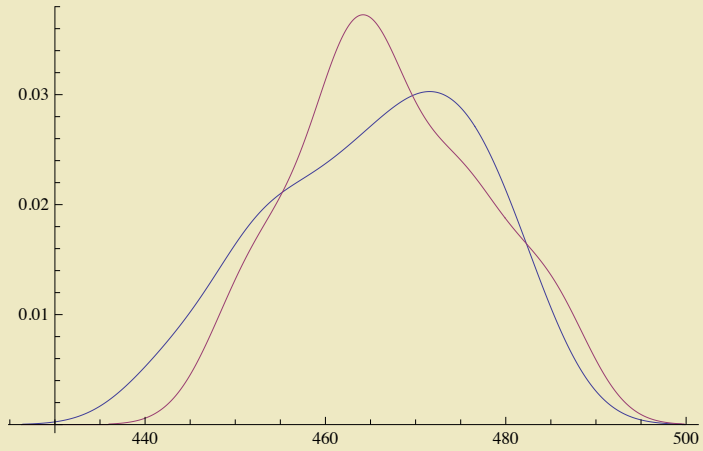


---

Qt: Test1=$6.50414 \times 10^{-12}$ Test2={

| | Statistic | P-Value |
|---|---|---|
| Kruskal-Wallis | 47.2615 | $2.45726 \times 10^{-20}$ |

, KruskalWallis}

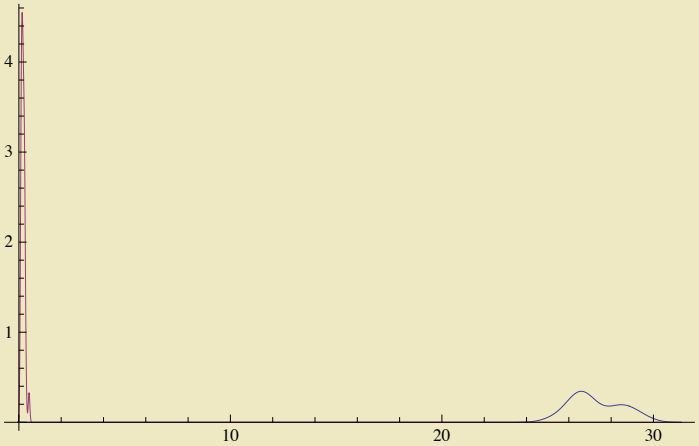


---

St: Test1=9.59121 × 10^-10 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 87.2935 | 1.92621 × 10^-13 |

, KSampleT}

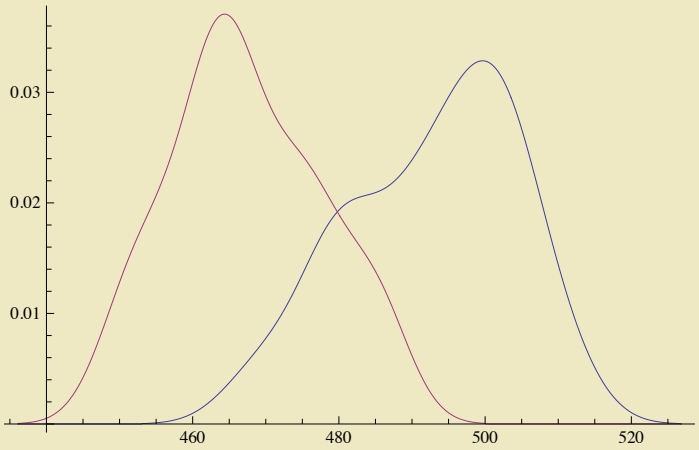


==========

Nowait Immediate(32) and Wait Immediate (32)

Work: Test1=0.0244849 Test2={

| | Statistic | P-Value |
|---|---|---|
| Kruskal-Wallis | 5.09219 | 0.0227979 |

, KruskalWallis}

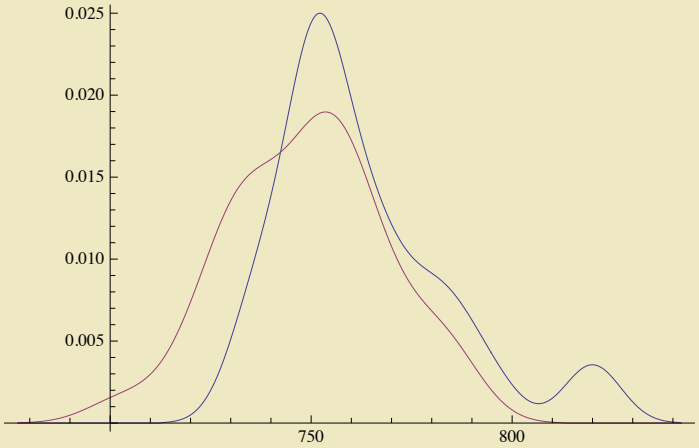


---

L: Test1=0.0262702 Test2={

| | Statistic | P-Value |
|---|---|---|
| Kruskal-Wallis | 4.96803 | 0.0245961 |

, KruskalWallis}

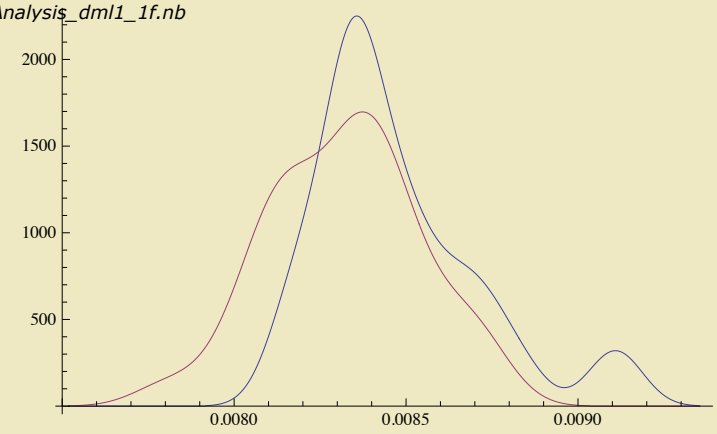


---

St: Test1=0.168734 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 3.41876 | 0.0692285 |

, KSampleT}

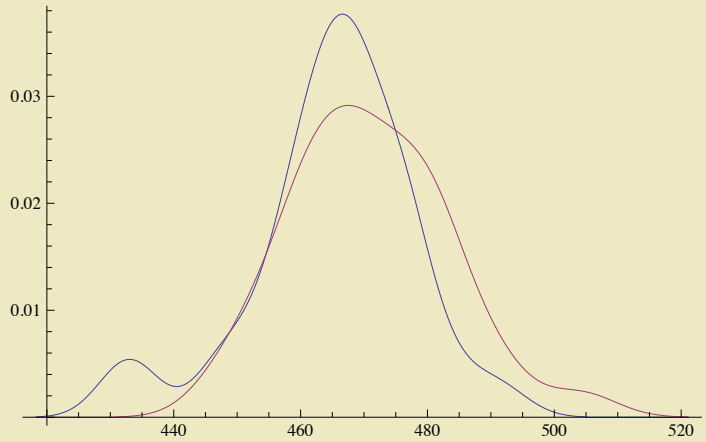


---

Qt: Test1=$5.92604 \times 10^{-12}$ Test2={

| | Statistic | P-Value |
|---|---|---|
| Kruskal-Wallis | 47.2615 | $2.45726 \times 10^{-20}$ |

, KruskalWallis}

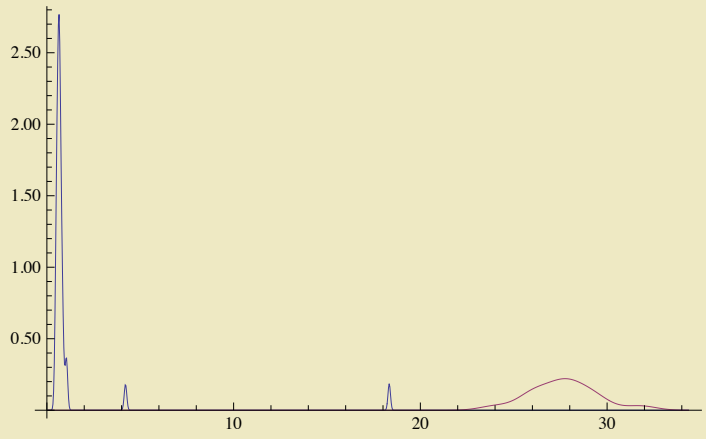


---

St: Test1=$6.47628 \times 10^{-11}$ Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 110.862 | $1.95955 \times 10^{-15}$ |

, KSampleT}

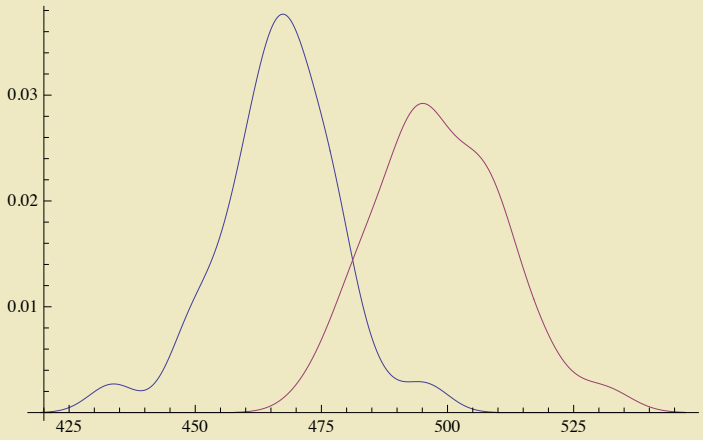


==========

Nowait Immediate(32) and Wait Batch (32)

Work: Test1=0.978572 Test2={

| | Statistic | P-Value |
|---|---|---|
| Kruskal-Wallis | 0.000406058 | 0.984115 |

, KruskalWallis}

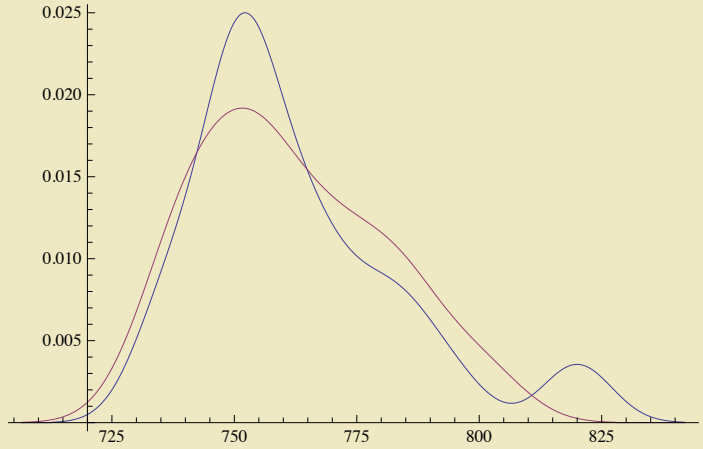


---

L: Test1=0.962517 Test2={

| | Statistic | P-Value |
|---|---|---|
| Kruskal-Wallis | 0.00288462 | 0.957679 |

, KruskalWallis}

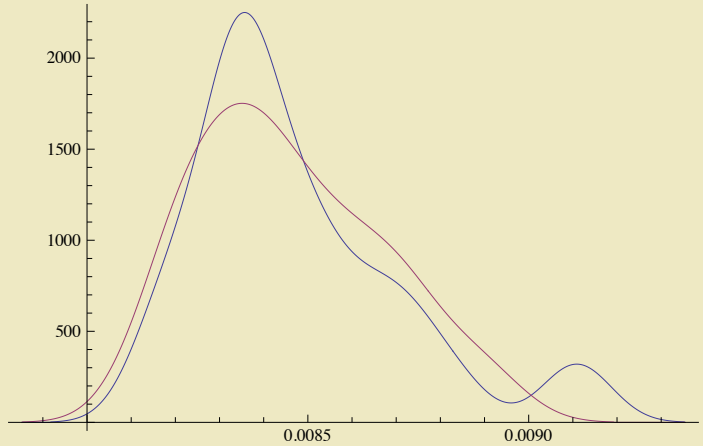


---

St: Test1=0.898499 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 0.00785287 | 0.929672 |

, KSampleT}

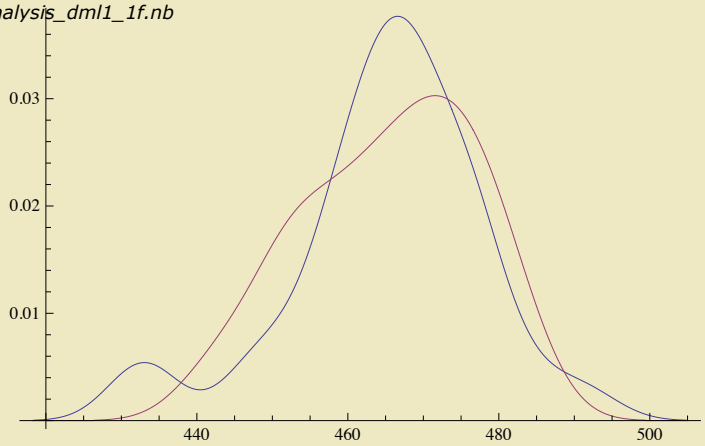


---

Qt: Test1=$5.92604 \times 10^{-12}$ Test2={

| | Statistic | P-Value |
|---|---|---|
| Kruskal-Wallis | 47.2615 | $2.45726 \times 10^{-20}$ |

, KruskalWallis}

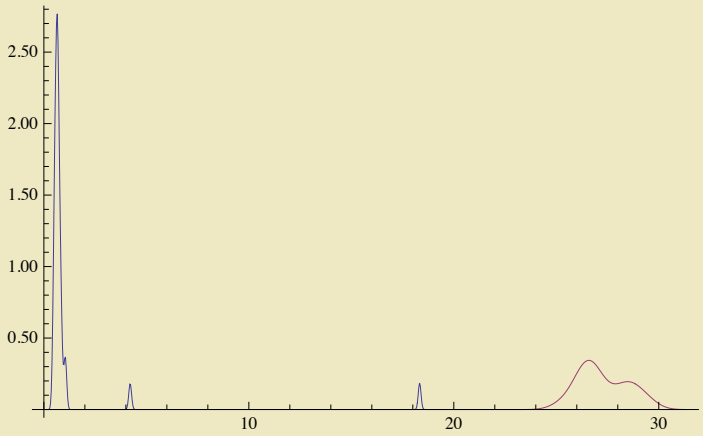


---

St: Test1=$6.28133 \times 10^{-10}$ Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 81.8401 | $6.18455 \times 10^{-13}$ |

, KSampleT}

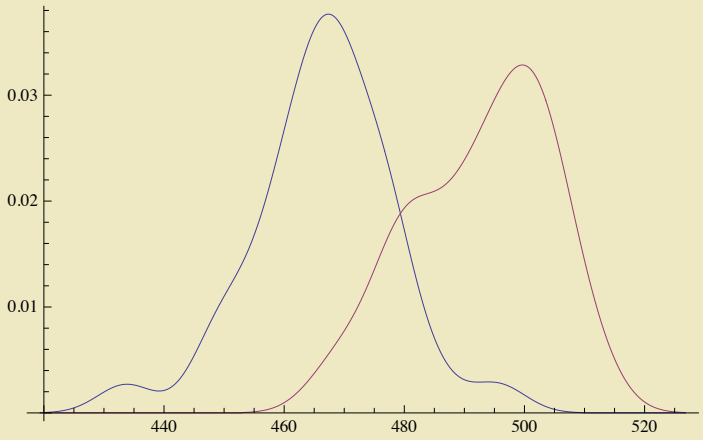


==========

Nowait Immediate(32) and Nowait Batch (32)

Work: Test1=0.48496 Test2={

| | Statistic | P-Value |
|---|---|---|
| Kruskal-Wallis | 0.478521 | 0.493481 |

, KruskalWallis}

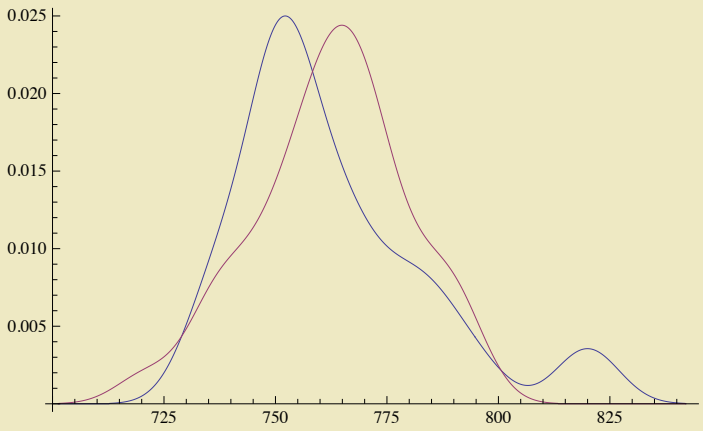


---

L: Test1=0.497727 Test2={

| | Statistic | P-Value |
|---|---|---|
| Kruskal-Wallis | 0.450721 | 0.506357 |

, KruskalWallis}

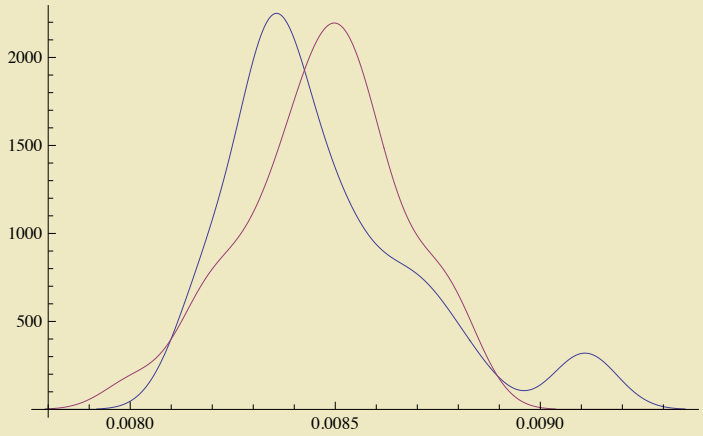


---

St: Test1=0.803823 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 0.512142 | 0.476899 |

, KSampleT}

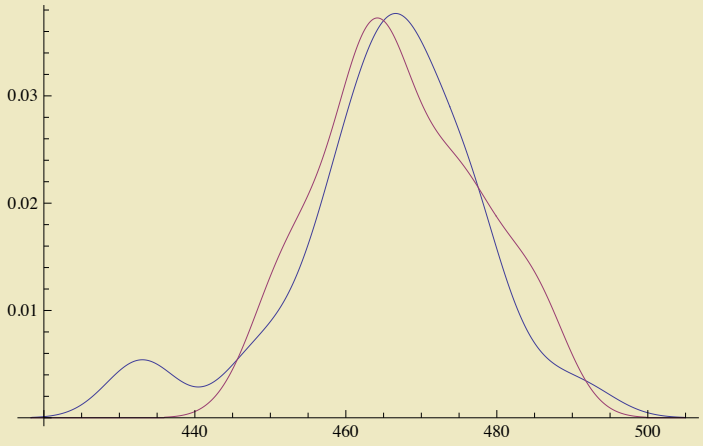


---

Qt: Test1=$7.84824 \times 10^{-12}$ Test2={

| | Statistic | P-Value |
|---|---|---|
| Kruskal-Wallis | 46.893 | $5.05453 \times 10^{-20}$ |

, KruskalWallis}

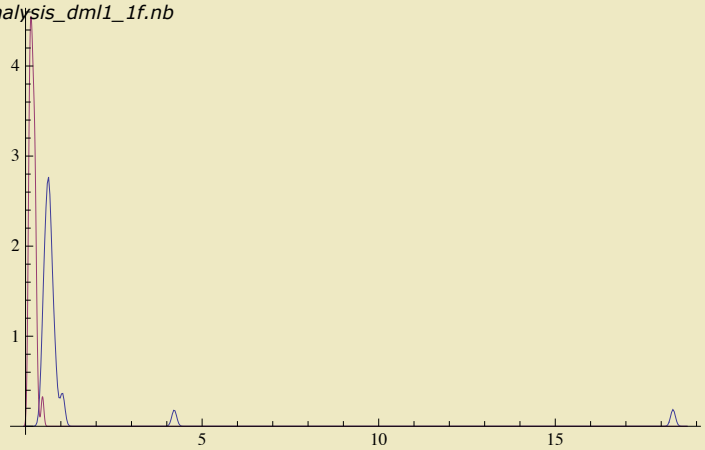


---

St: Test1=0.973222 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 0.110024 | 0.741235 |

, KSampleT}

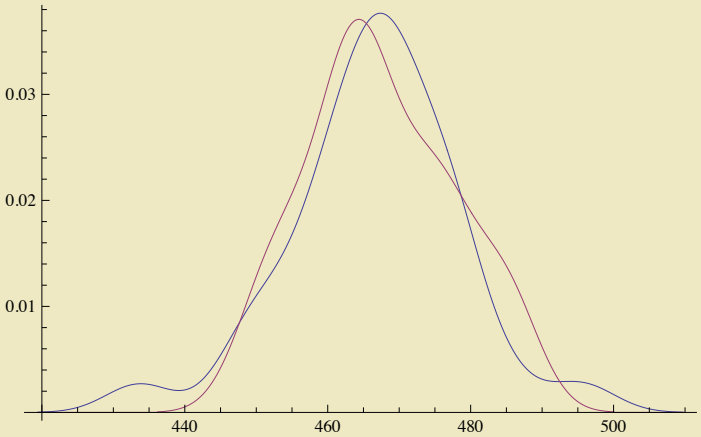


==========

Nowait Batch(32) and Wait Immediate (32)

Work: Test1=0.00405305 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 8.55364 | 0.00481099 |

, KSampleT}

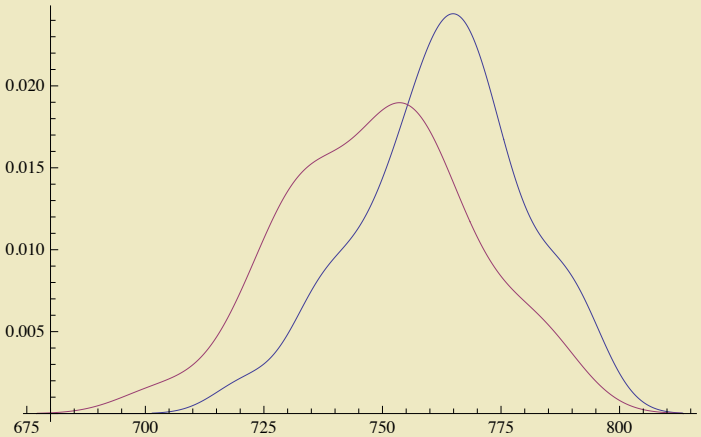


---

L: Test1=0.00364938 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 8.65493 | 0.00458214 |

, KSampleT}

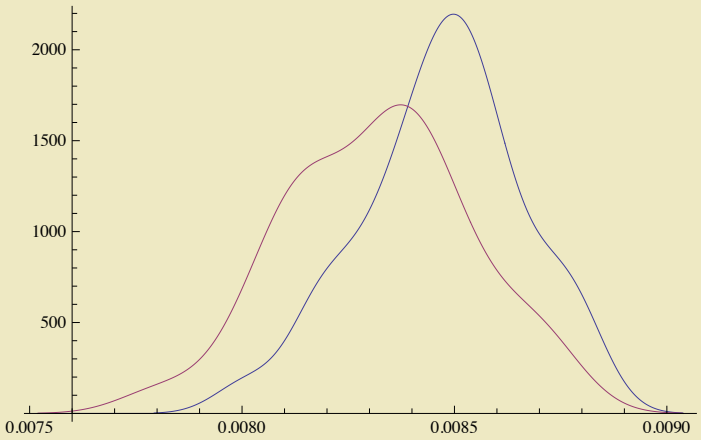


---

St: Test1=0.267974 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 1.72404 | 0.194013 |

, KSampleT}

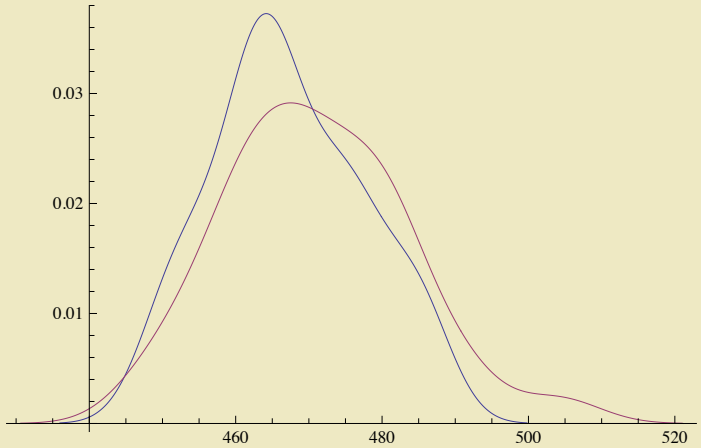


---

Qt: Test1=$5.91948 \times 10^{-12}$ Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 7370.44 | $3.67139 \times 10^{-66}$ |

, KSampleT}

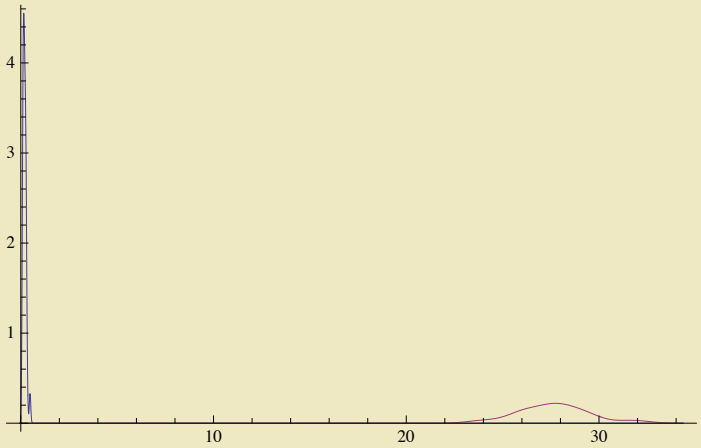


---

St: Test1=$6.47628 \times 10^{-11}$ Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 118.151 | $5.38924 \times 10^{-16}$ |

, KSampleT}

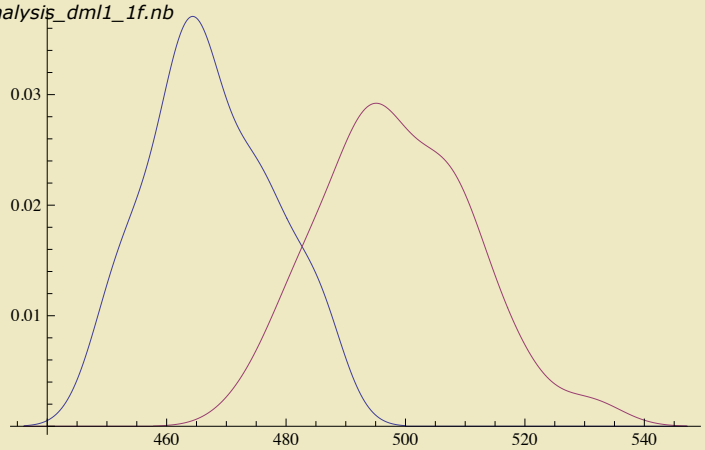


==========

Nowait Batch(32) and Wait Batch (32)

Work: Test1=0.619255 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 0.0274729 | 0.868893 |

, KSampleT}

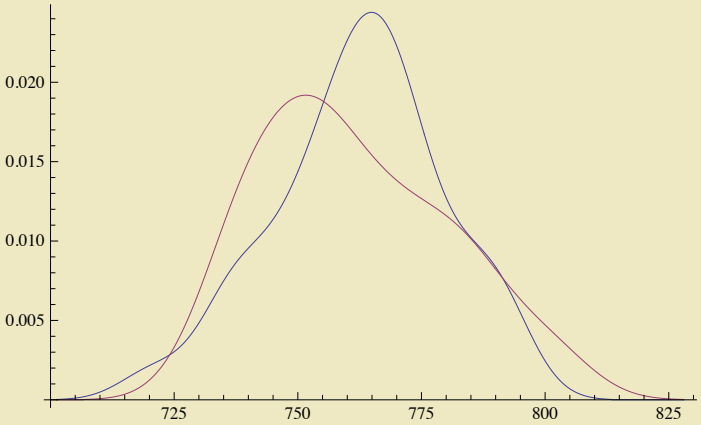


---

L: Test1=0.577372 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 0.0546811 | 0.81588 |

, KSampleT}

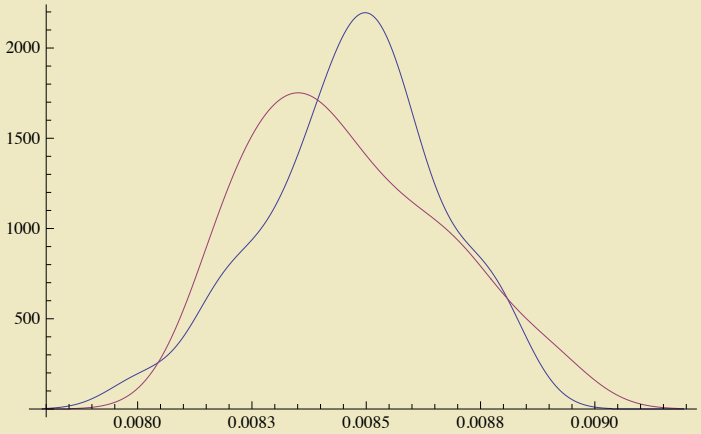


---

St: Test1=0.624069 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 0.435088 | 0.511946 |

, KSampleT}

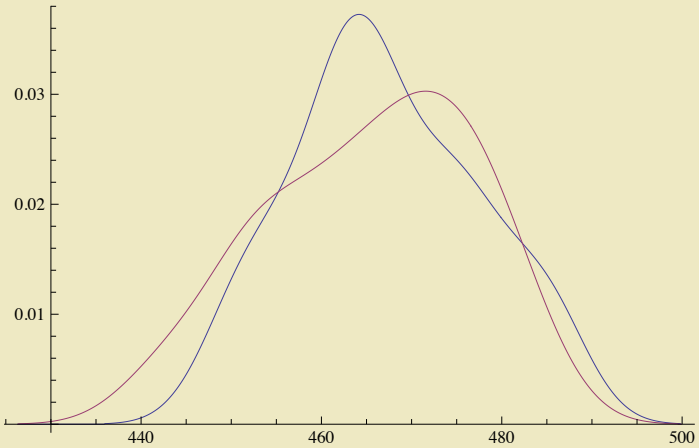


---

Qt: Test1=$5.91948 \times 10^{-12}$ Test2={

| | Statistic | P-Value |
|---|---|---|
| Kruskal-Wallis | 47.2615 | $2.45726 \times 10^{-20}$ |

, KruskalWallis}

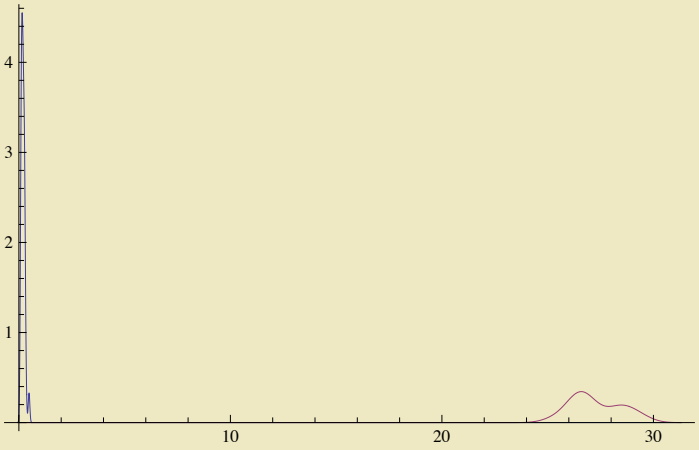


---

St: Test1=$8.8158 \times 10^{-10}$ Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 87.2935 | $1.92621 \times 10^{-13}$ |

, KSampleT}

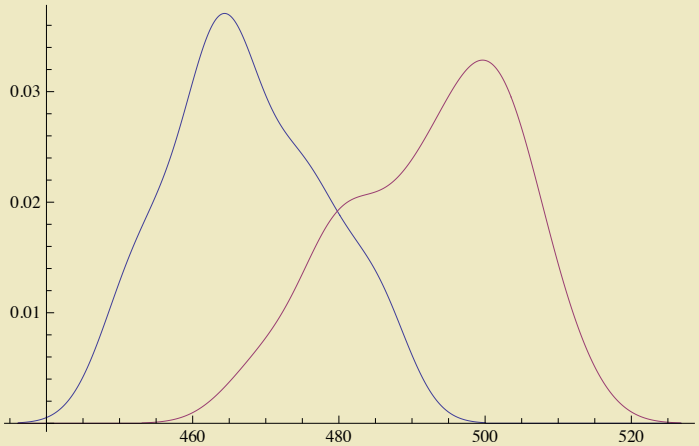


==========

Nowait Batch(32) and Nowait Immediate (32)

Work: Test1=0.493396 Test2={

| | Statistic | P-Value |
|---|---|---|
| Kruskal-Wallis | 0.478521 | 0.493481 |

, KruskalWallis}

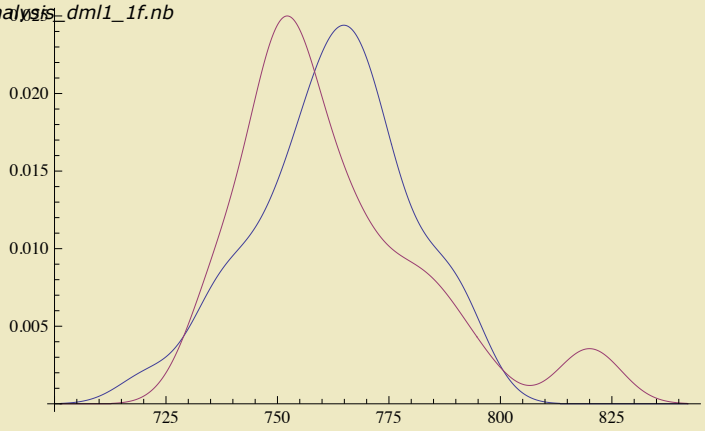


---

L: Test1=0.506278 Test2={

| | Statistic | P-Value |
|---|---|---|
| Kruskal-Wallis | 0.450721 | 0.506357 |

, KruskalWallis}

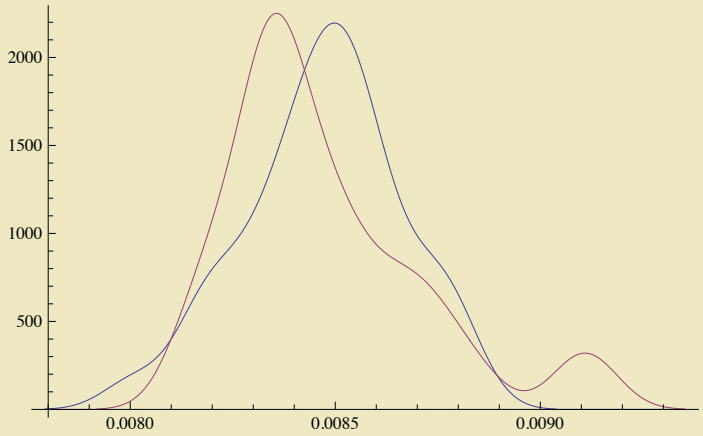


---

St: Test1=0.814228 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 0.512142 | 0.476899 |

, KSampleT}

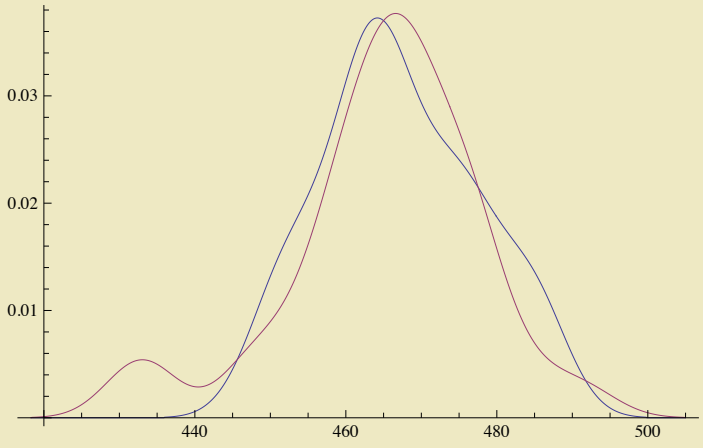


---

Qt: Test1=$7.14528 \times 10^{-12}$ Test2={

| | Statistic | P-Value |
|---|---|---|
| Kruskal-Wallis | 46.893 | $5.05453 \times 10^{-20}$ |

, KruskalWallis}

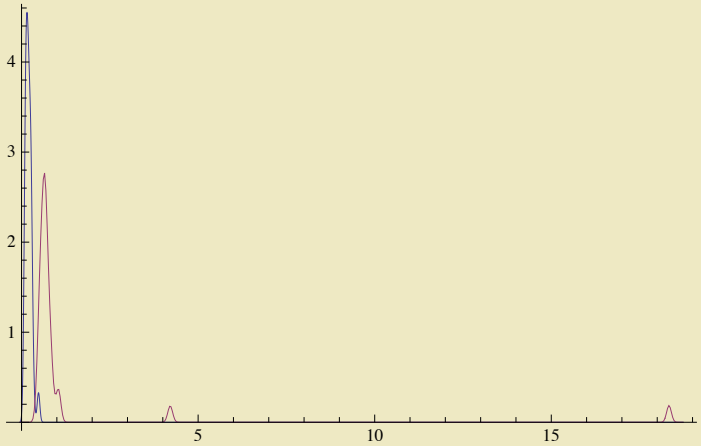


---

St: Test1=0.962517 Test2={

| | Statistic | P-Value |
|---|---|---|
| K-Sample T | 0.110024 | 0.741235 |

, KSampleT}

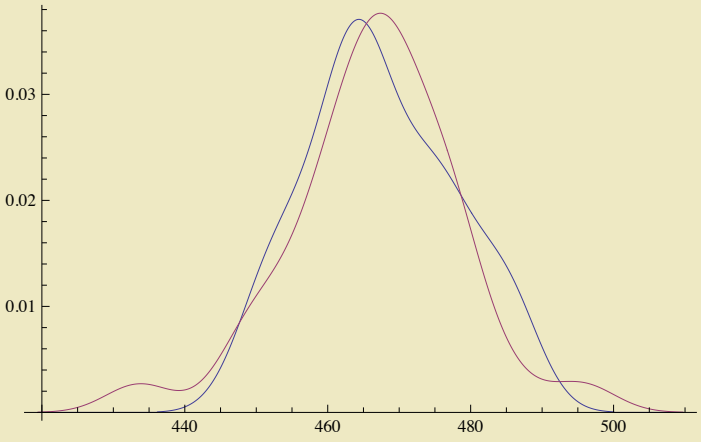


## Visually Comparing All Samples

I also wanted to get a nice visual picture of my sample sets...together. Sometimes I include all the sample sets and sometimes I don't. It's just based on what I want to convey. Sometimes you get a more appropriate view if all the data is not included.

Here is the colors in order of sample set; blue, red, yellow, green.

```
gset = {};
Table[
  AppendTo[gset, ssWork[i]];
  , {i, 1, ssNum}
 ];
SmoothHistogram[gset,
 PlotLabel → "Occurances vs Work (cmt)", AxesLabel → {"cmt", "Occurs"}]
```

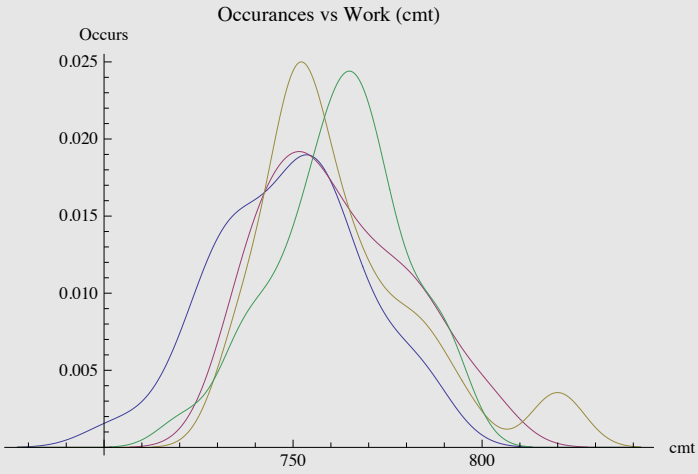


```
gset = {};
Table[
  AppendTo[gset, ssL[i]];
  , {i, 1, ssNum}
 ];
SmoothHistogram[gset,
 PlotLabel → "Occurances vs Workload (cmt/ms)", AxesLabel → {"cmt/workload", "Occurs"}]
```

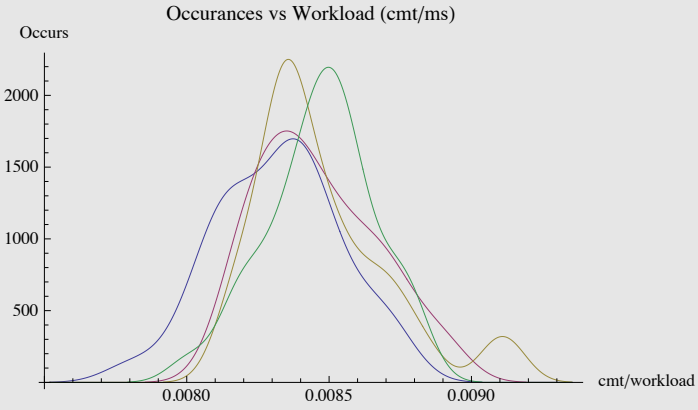

```
gset = {};
Table[
  AppendTo[gset, ssSt[i]];
  , {i, 1, ssNum}
 ];
SmoothHistogram[gset,
 PlotLabel → "Occurances vs CPU(ms)/cmt", AxesLabel → {"CPU(ms)/cmt", "Occurs"}]
```

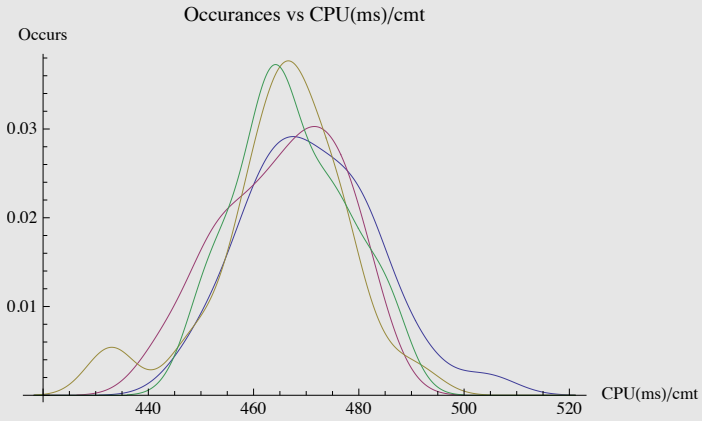


```
gset = {};
Table[
  AppendTo[gset, ssQt[i]];
  , {i, 1, ssNum}
 ];
SmoothHistogram[gset,
 PlotLabel → "Occurances vs Wait(ms)/cmt", AxesLabel → {"Wait(ms)/cmt", "Occurs"}]
```

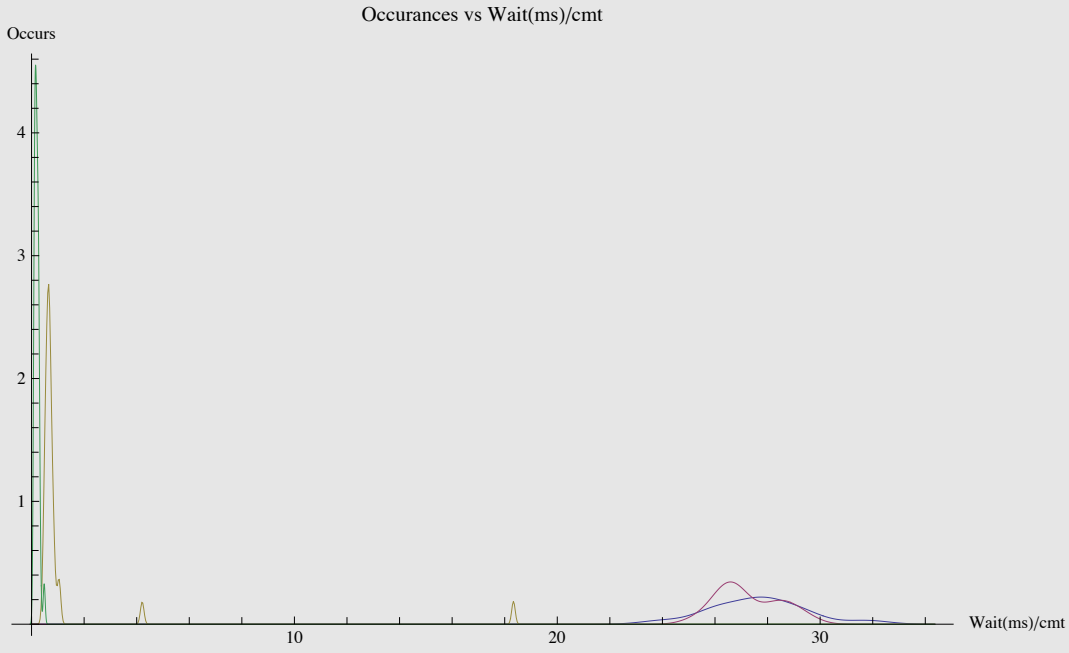

```
gset = {};
Table[
  AppendTo[gset, ssRt[i]];
  , {i, 1, ssNum}
 ];
SmoothHistogram[gset,
 PlotLabel → "Occurances vs CPU+Wait(ms)/cmt", AxesLabel → {"CPU+Wait(ms)/cmt", "Occurs"}]
```

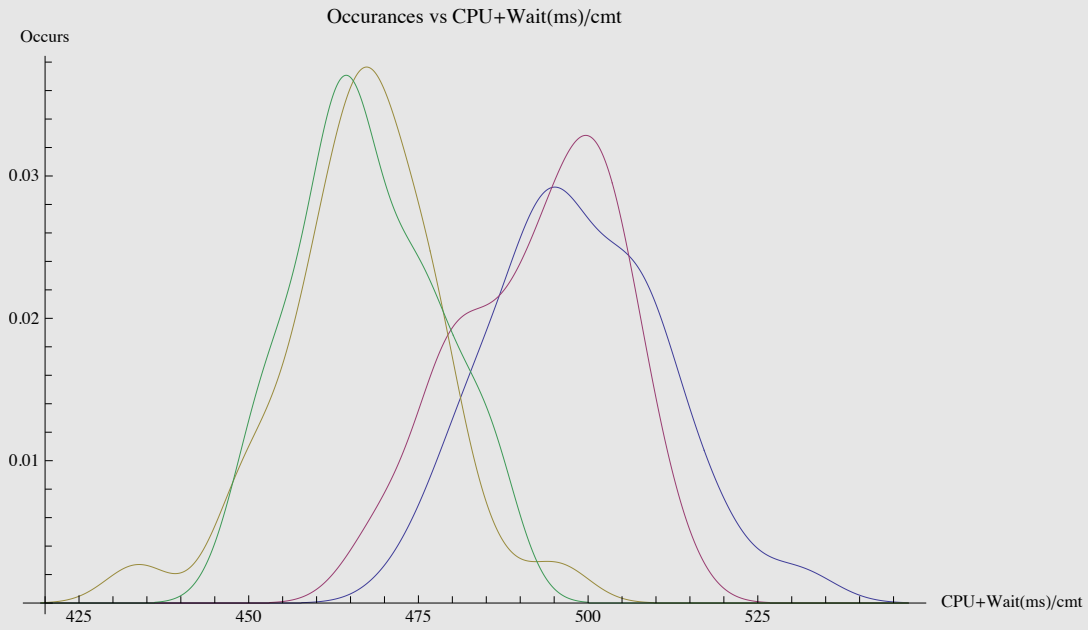


## Fair: Comparing situations fairly

As workload increases, we can expect response time to also increase. If the workload is decreases response time *may* decrease. When comparing two sample sets, if their arrival rates are different, comparing their queue times and response times using statitcal significant tests is problematic and downright unfiar.

To get around this problem, we need to do the testing at the same arrival rate. However, our sample data may not have been collected at the same arrival rate! That's a problem.

To get this problem, we need to develop an equation for each of our four sample sets relating the arrival rate to the response time. Fortunately, there already exists a formal mathematical equation related the arrival rate, service time, queue time, the number of transaction processors (think: CPU cores), and the response time.

For a CPU constained system, the equation is r = s / (1-(sL/m)^m)

For an IO constrained system, the equation is r = s / (1-(sL/m))

There are many ways to go about this process. Here is who I'm choosing to do it. Within in each of the four sample sets (ssNum), for each sample value (sampleIdx), I derive the missing variable M and store these in mList. Then I take the average M and save that in bestMList. I could have weighted the average or taken the median, but I just kept it simple. You see the details in the code segment below.

Note that this can take awhile to run. The "Solve" is cpu intensive.

```
Clear[s, q, r, l, sol, bestM, mList, bestMList];
bestMList = {};
Table[
  mList = {};
  Table[
   s = ssSt[ssidx][[sampleIdx]];
   q = ssQt[ssidx][[sampleIdx]];
   r = s + q;
   l = ssL[ssidx][[sampleIdx]];
   sol = Solve[s / (1 - (s * l / m) ^m) == r && m > 0, m];
   {mm} = m /. sol;
   (*Print[ssidx," ",sampleIdx," mm=",mm];*)
   AppendTo[mList, mm];
   , {sampleIdx, 1, sampleNum}
  ];
  bestM = Mean[mList];
  (*Print[ssidx," bestM=",bestM];*)
  AppendTo[bestMList, bestM];
  , {ssidx, 1, ssNum}
 ];
```

Solve::ratnz : Solve was unable to solve the system with inexact coefficients.
The answer was obtained by solving a corresponding exact system and numericizing the result. ≫

Solve::ratnz : Solve was unable to solve the system with inexact coefficients.
The answer was obtained by solving a corresponding exact system and numericizing the result. ≫

Solve::ratnz : Solve was unable to solve the system with inexact coefficients.
The answer was obtained by solving a corresponding exact system and numericizing the result. ≫

General::stop : Further output of Solve::ratnz will be suppressed during this calculation. ≫

```
Print["The best M values are:"];
Table[
  Print[ssName[ssidx], " M=", bestMList[[ssidx]]];
  , {ssidx, 1, ssNum}
 ];
```

```
The best M values are:

Wait Immediate M=6.23085

Wait Batch M=6.25177

Nowait Immediate M=8.41341

Nowait Batch M=9.24755
```

The baseline arrival rate will be the average from our first sample set, (nowait, immediate). I expect the other three options to provide better performance and the nowwait,immediate is the Oracle default, hence this is my chosen baseline.

```
baselineL = Mean[ssL[1]]
```

```
0.00831995
```

Now that we have a good M for each of our four sample sets (bestMList) along with the standard arrival rate (baselineL), and the observed service time for each individual sample, we will derive the queue time and the response time for each indivisual sample (within each of our four sample sets). All the inputs and derived queue time and response time are stored in the standardized lists, stndL, stndM, stndSt, stndQt, and stndRt. At this point, we essentially have an entirely new (i.e., standardized) set of experimental data. This allows us to run this data through the same statistical analysis as I did above!

```
Clear[stndRt, stndSt, stndL, stndQt, stndM];
Clear[s, l, m, r, q];
Table[
  stndRt[ssidx] = {};
  stndQt[ssidx] = {};
  stndSt[ssidx] = {};
  stndL[ssidx] = {};
  stndM[ssidx] = {};
  Table[
   s = ssSt[ssidx][[sampleIdx]];
   l = baselineL;
   m = bestMList[[ssidx]];
   r = s / (1 - (s l / m)^m);
   q = r - s;
   (*Print[ssidx," ",sampleIdx," s=",s," l=",l," m=",m," r=",r];*)
   AppendTo[stndRt[ssidx], r];
   AppendTo[stndSt[ssidx], s];
   AppendTo[stndL[ssidx], l];
   AppendTo[stndQt[ssidx], q];
   AppendTo[stndM[ssidx], m];
   , {sampleIdx, 1, sampleNum}
  ];
  , {ssidx, 1, ssNum}
 ];
stndM[1]
stndL[1]
stndSt[1]
stndQt[1]
stndRt[1]
```

```
{6.23085, 6.23085, 6.23085, 6.23085, 6.23085, 6.23085, 6.23085, 6.23085, 6.23085, 6.23085, 6.23085,
 6.23085, 6.23085, 6.23085, 6.23085, 6.23085, 6.23085, 6.23085, 6.23085, 6.23085, 6.23085, 6.23085,
 6.23085, 6.23085, 6.23085, 6.23085, 6.23085, 6.23085, 6.23085, 6.23085, 6.23085, 6.23085}
```

```
{0.00831995, 0.00831995, 0.00831995, 0.00831995, 0.00831995, 0.00831995, 0.00831995, 0.00831995,
 0.00831995, 0.00831995, 0.00831995, 0.00831995, 0.00831995, 0.00831995, 0.00831995, 0.00831995,
 0.00831995, 0.00831995, 0.00831995, 0.00831995, 0.00831995, 0.00831995, 0.00831995, 0.00831995,
 0.00831995, 0.00831995, 0.00831995, 0.00831995, 0.00831995, 0.00831995, 0.00831995, 0.00831995}
```

```
{459.616, 473.553, 453.723, 461.943, 479.949, 467.896, 504.242, 465.825, 473.903, 479.533, 463.854,
 463.434, 483.655, 458.153, 466.594, 482.754, 481.4, 472.362, 468.953, 480.046, 477.321, 490.003,
 459.735, 463.804, 457.969, 450.973, 472.195, 479.738, 491.073, 467.791, 471.59, 447.642}
```

```
{23.0411, 28.8926, 20.9072, 23.9363, 32.0062, 26.3718, 46.8654, 25.5002, 29.0554, 31.7948, 24.6948,
 24.5264, 33.9481, 22.4937, 25.8205, 33.4663, 32.7542, 28.3444, 26.8269, 32.0555, 30.6917,
 37.5278, 23.086, 24.6748, 22.4256, 19.9741, 28.268, 31.8989, 38.1643, 26.3272, 27.9941, 18.8936}
```

```
{482.657, 502.445, 474.63, 485.879, 511.955, 494.268, 551.107, 491.325, 502.958, 511.328, 488.548,
 487.961, 517.603, 480.647, 492.414, 516.22, 514.155, 500.706, 495.779, 512.101, 508.013, 527.531,
 482.821, 488.479, 480.394, 470.948, 500.463, 511.637, 529.238, 494.119, 499.584, 466.536}
```

## Fair: Basic Statistics

Using our standardized data, in this section I calculate the basic statistics, such as the mean and median. My objective is to ensure the data has been collected and entered correctly and also to compare the two datasets to see if they appear to be different.

```
myData = Table[
   {
    ssName[ssidx], Mean[stndL[ssidx]], Mean[stndSt[ssidx]],
    Mean[stndQt[ssidx]], Mean[stndRt[ssidx]], Length[stndL[ssidx]],
    N[StandardDeviation[stndL[ssidx]]], N[StandardDeviation[stndSt[ssidx]]],
    N[StandardDeviation[stndQt[ssidx]]], N[StandardDeviation[stndRt[ssidx]]]
   }, {ssidx, 1, ssNum}
  ];
toGrid = Prepend[myData, {"Settings", "Avg L\n(cmt/ms)", "Avg CPUt\n(ms/cmt)",
    "Avg Wt\n(ms/cmt)", "Avg Rt\n(ms/cmt)", "Samples", "Stdev L\n(cmt/ms)",
    "Stdev CPUt\n(ms/cmt)", "Stdev Wt\n(ms/cmt)", "Stdev Rt\n(ms/cmt)"}];
Grid[
 toGrid,
 Frame →
  All]
```

| Settings | Avg L (cmt/ms) | Avg CPUt (ms/cmt) | Avg Wt (ms/cmt) | Avg Rt (ms/cmt) | Samples | Stdev L (cmt/ms) | Stdev CPUt (ms/cmt) | Stdev Wt (ms/cmt) | Stdev Rt (ms/cmt) |
|---|---|---|---|---|---|---|---|---|---|
| Wait Immediate | 0.00831995 | 470.976 | 28.2259 | 499.202 | 32 | $1.76248 \times 10^{-18}$ | 12.4358 | 5.9129 | 18.3156 |
| Wait Batch | 0.00831995 | 465.452 | 24.9021 | 490.354 | 32 | $1.76248 \times 10^{-18}$ | 11.4239 | 4.49972 | 15.9122 |
| Nowait Immediate | 0.00831995 | 465.186 | 0.696187 | 465.882 | 32 | $1.76248 \times 10^{-18}$ | 12.6153 | 0.166351 | 12.7791 |
| Nowait Batch | 0.00831995 | 467.24 | 0.158056 | 467.398 | 32 | $1.76248 \times 10^{-18}$ | 10.2195 | 0.0359533 | 10.2553 |

Using our standardized data, in this section I calculate the key parameters for understanding the impact or change of the settings.

```
myData = Table[
   {
    ssName[ssidx],
    Mean[stndL[ssidx]], 100 * (Mean[stndL[ssidx]] - Mean[stndL[1]]) / Mean[stndL[1]],
    Mean[stndSt[ssidx]], 100 * (Mean[stndSt[ssidx]] - Mean[stndSt[1]]) / Mean[stndSt[1]],
    Mean[stndQt[ssidx]], 100 * (Mean[stndQt[ssidx]] - Mean[stndQt[1]]) / Mean[stndQt[1]],
    Mean[stndRt[ssidx]], 100 * (Mean[stndRt[ssidx]] - Mean[stndRt[1]]) / Mean[stndRt[1]],
    Length[stndL[ssidx]]
   }, {ssidx, 1, ssNum}
  ];
toGrid = Prepend[myData, {"Settings", "Avg L\n(cmt/ms)", "%\nChange", "Avg CPUt\n(ms/cmt)",
    "%\nChange", "Avg Wt\n(ms/cmt)", "%\nChange", "Avg Rt\n(ms/cmt)", "%\nChange", "Samples"}];
Grid[
 toGrid,
 Frame →
  All]
```

| Settings | Avg L (cmt/ms) | % Change | Avg CPUt (ms/cmt) | % Change | Avg Wt (ms/cmt) | % Change | Avg Rt (ms/cmt) | % Change | Samples |
|---|---|---|---|---|---|---|---|---|---|
| Wait Immediate | 0.00831995 | 0. | 470.976 | 0. | 28.2259 | 0. | 499.202 | 0. | 32 |
| Wait Batch | 0.00831995 | 0. | 465.452 | -1.17276 | 24.9021 | -11.7758 | 490.354 | -1.77228 | 32 |
| Nowait Immediate | 0.00831995 | 0. | 465.186 | -1.22937 | 0.696187 | -97.5335 | 465.882 | -6.6746 | 32 |
| Nowait Batch | 0.00831995 | 0. | 467.24 | -0.793276 | 0.158056 | -99.44 | 467.398 | -6.37096 | 32 |