Online Supplement to Are Call Center and Hospital Arrivals Well Modeled by Nonhomogeneous Poisson Processes?

Song-Hee Kim and Ward Whitt

Industrial Engineering and Operations Research
Columbia University
New York, NY, 10027
{sk3116, ww2040} @columbia.edu

1 Overview

We present supporting material in this online supplement to the main paper. In §2 we define the CU and Lewis KS tests of an NHPP that we consider in the main paper. In §3, we illustrate how departures from the Poisson property can have a strong impact upon performance via simulation experiments. In §**??** we give the results from addtional experiments on rounding and un-rounding. In §5 we provide an example illustrating Theorem 6, which is a theory behind the asymptotic value of the CU KS test statistic applied to an NHPP with linear arrival rate, and in §6 we present an asymptotic result paralleling Theorem 6 for a piecewise-smooth arrival rate function. More material appears in an appendix available from the authors' web pages.

2 The CU and Lewis KS Tests

In the main paper, we consider two statistical tests of an NHPP: the CU KS test and the Lewis KS test. This section describes the two tests in detail.

Conditional-Uniform (CU) Test. This test exploits the basic conditioning property of a PP. Given an arrival process over an interval $[0, t]$, we observe the number n of arrival in this interval and their arrival times T_j, $1 \leq j \leq n$. Under the null PP hypothesis, these random variables are distributed as the order statistics of i.i.d. random variables uniformly distributed over $[0, t]$. Thus, the random variables T_j/t, $1 \leq j \leq n$, are distributed as the order statistics of i.i.d. random variables uniformly distributed over $[0, 1]$. Thus the ecdf can be computed via

$$F_n(x) \equiv n^{-1} \sum_{k=1}^{n} 1_{\{T_k/t \leq x\}}, \quad 0 \leq x \leq 1,$$

and the KS statistic can be computed as in (2) of the main paper with uniform cdf $F(x) = x$, $0 \leq x \leq 1$.

Lewis Test. Lewis [1965] proposed using a different modification of the CU test, exploiting a transformation due to Durbin [1961]. Following Durbin [1961], we start with a sample U_j, $1 \leq j \leq n$, hypothesized to be uniformly distributed on $[0, 1]$. Then let $U_{(j)}$ be the j^{th} smallest of these, $1 \leq j \leq n$, so that $U_{(1)} < \cdots < U_{(n)}$. This is applied in Lewis [1965] with $U_{(j)} = T_j/t$ from the CU test. Next we look at the successive *intervals* between these ordered observations:

$$C_1 = U_{(1)},\ C_j = U_{(j)} - U_{(j-1)} \quad 2 \leq j \leq n, \quad \text{and} \quad C_{n+1} = 1 - U_{(n)}. \tag{2.1}$$

Then let $C_{(j)}$ be the j^{th} smallest of these intervals, $1 \leq j \leq n$, so that $0 < C_{(1)} < \cdots < C_{(n+1)} < 1$. Now let Z_j be scaled versions of the intervals between these new variables, i.e.,

$$Z_j = (n + 2 - j)(C_{(j)} - C_{(j-1)}), \quad 1 \leq j \leq n + 1, \quad (\text{with} \quad C_{(0)} \equiv 0). \tag{2.2}$$

Remarkably, Durbin [1961] showed in a simple direct argument (by giving explicit expressions for the joint density functions, exploiting the transformation of random vectors by a function) that, under the PP null hypothesis, the random vector $(Z_1, \ldots, Z_n)$ is distributed the same as the random vector $(C_1, \ldots, C_n)$. Hence, again under the PP null hypothesis, the vector of associated partial sums $(S_1, \ldots, S_n)$, where

$$S_k \equiv Z_1 + \cdots + Z_k, \quad 1 \leq k \leq n, \tag{2.3}$$

has the same distribution as the original random vector $(U_{(1)}, \ldots, U_{(n)})$ of ordered uniform random variables. Hence, we can apply the KS test with the ecdf

$$F_n(x) \equiv n^{-1} \sum_{k=1}^{n} 1_{\{S_k \leq x\}}, \quad 0 \leq x \leq 1,$$

for S_k in (2.3) and (2.2), comparing it to the uniform cdf $F(x) \equiv x$, $0 \leq x \leq 1$.

3 Performance Impact of the Arrival Process

In this section, we show the results of a simulation experiment demonstrating the performance impact of departures from the Poisson property in arrival processes. Table 1 compares the simulated performance of two $GI/M/s + M$ models where one has exponentially distributed interarrival times (hence, an NHPP with a constant arrival rate) and the other has hyper-exponentially distributed (a mixture of two exponentials, and hence more variable than exponential) interarrival times with *squared coefficient of variation* (scv, variance divided by the square of the mean) $c^2 = 2$. The cdf of H_2 is $P(X \le x) \equiv 1 - p_1 e^{-\lambda_1 x} - p_2 e^{-\lambda_2 x}$. We further assume balanced means for ($p_1 \lambda_1^{-1} = p_2 \lambda_2^{-1}$) as in (3.7) of Whitt [1982] so that $p_i = [1 \pm \sqrt{(c_X^2 - 1)/(c_X^2 + 1)}]/2$ and $\lambda_i = 2p_i$. The staffing level, s, is chosen using the square root staffing formula assuming exponentially distributed interarrival times, $s \equiv m + \beta\sqrt{m}$, where m is the offered load $\lambda/\mu = 25$. We consider three cases for the quality-of-service parameter β: 0.5, 1, and 2 (yielding s =28, 30, and 35). The results are based on 100 replications of $10^3 + 10^5$ customers (first 10^3 customers removed to get rid of the initial effect). Associated 95% confidence intervals are also shown. The results show that if we choose staffing levels assuming that interarrival times are exponentially distributed when they are actually hyper-exponentially distributed with $c^2 = 2$, then we would observe an average of 35% increase in the percentage of the customers that wait (here we consider $s = 30$, but the increase is similar in other cases). The impact on staffing is about 1 server in this example. We also note that this H_2 arrival process is not exceptionally far from a PP; other more variable processes have even greater impact on performance.

Table 1: Comparison of Simulated Performance of $GI/M/s + M$ Models with Two Different Interarrival Time CDFs.

Model	s	$E[W\|All]$	$E[W\|Served]$	$E[W\|Abandoned]$	$\%Wait$	$\%Abandon$
$M/M/s + M$	28	0.0324 ± 0.0003	0.0348 ± 0.0003	0.1004 ± 0.0006	29.96 ± 0.15	3.49 ± 0.03
	30	0.0168 ± 0.0002	0.0181 ± 0.0002	0.0878 ± 0.0006	18.23 ± 0.13	1.81 ± 0.02
	35	0.0022 ± 0.0001	0.0024 ± 0.0001	0.0647 ± 0.0012	3.41 ± 0.06	0.24 ± 0.01
$H_2/M/s + M$	28	0.0461 ± 0.0003	0.0496 ± 0.0004	0.1165 ± 0.0006	35.93 ± 0.16	4.96 ± 0.03
	30	0.0272 ± 0.0003	0.0295 ± 0.0003	0.1031 ± 0.0006	24.75 ± 0.15	2.94 ± 0.03
	35	0.0057 ± 0.0001	0.0061 ± 0.0001	0.0776 ± 0.0009	7.11 ± 0.09	0.61 ± 0.01

4 More on Un-Rounding

As indicated at the end of §2 of the main paper, this section provides additional experiment results to show that the un-rounding does not inappropriately cause the Lewis KS test to fail to reject a non-PP (in §4.1), provided that the un-rounding is not done too coarsely (in §4.2). Additional simulation results and discussion appear in the appendix.

4.1 Batch Poisson Examples

We consider two forms of batch Poisson processes, and show that as long as the rounding is not too coarse (as in the examples in the main paper, where the rounding and un-rounding are done in the units of seconds when the mean interarrival time is 3.6 seconds), the un-rounding does not inappropriately cause the Lewis

KS test to fail to reject a non-PP. The first batch Poisson process is a rate-1000 renewal process, in which the interarrival times are 0 with probability p and an exponential random variable with probability $1-p$. Table 2 shows that the un-rounding consistently detects the deviation from the PP when p is not too small (e.g., when $p \geq 0.05$). Figure 1 compares the average ecdf based on 100 replications of the renewal process where $p = 0.05$ with the cdf of the null hypothesis. We observe that the average ecdf plots for the Lewis KS test look different for raw and un-rounded, because we start with mass at 0, but un-rounding moves the mass to slightly larger values. Nevertheless, the test makes the right decision, because even though the un-rounding removes all 0 interarrival times, it still leaves too many very short interarrival times.

Table 2: Test results of a rate-1000 renewal process on $[0, 6]$, in which the interarrival times are 0 with probability p and an exponential random variable with probability $1-p$. Results over 10000 iterations.

		CU		Log		Lewis	
p	Type	# P	ave[p-value]	# P	ave[p-value]	# P	ave[p-value]
0.1	Raw	9015	0.41	0	0.00	0	0.00
	Rounded	9015	0.41	0	0.00	0	0.00
	Un-rounded	9018	0.41	0	0.00	0	0.00
0.05	Raw	9306	0.46	0	0.00	0	0.00
	Rounded	9304	0.46	0	0.00	0	0.00
	Un-rounded	9308	0.46	2	0.00	0	0.00
0.01	Raw	9453	0.49	8798	0.26	7879	0.22
	Rounded	9454	0.49	0	0.00	0	0.00
	Un-rounded	9453	0.49	8877	0.37	8175	0.32

The second batch Poisson process is a modification of a PP in which every k^{th} arrival occurs in batches of size 2; the arrival rate is reduced to $1000k/(k+1)$, so that the overall arrival rate is again 1000. Table 3 shows that the un-rounding consistently detects the deviation from the PP when k is not too large (e.g., when $k \leq 9$). Figure 2 compares the average ecdf based on 100 replications of the case where $k = 6$ with the cdf of the null hypothesis. We make the same observations we did before in Figure 1.

4.2 Coarse Rounding and Un-rounding

In the main paper, all rounding was done to the nearest second, where the mean interarrival times was 3.6 seconds. It is possible to observe more coarse rounding in reality, such as to the nearest minute. In this subsection, we show that if the rounding is too coarse (to the nearest minute instead of to the nearest second in our batch-Poisson example), then the unrounding *can* hide the non-PP character of the original process. The results in Table 4 for the H_2 renewal arrival process show some loss in power when there is rounding to the nearest minute and then un-rounding, but overall the KS test still has significant power. Evidently, the

Figure 1: Comparison of the average ecdf of a rate-1000 renewal process on $[0, 6]$, in which the interarrival times are 0 with probability $p = 0.05$ and an exponential random variable with probability $1 - p$. From top to bottom: CU, Lewis test. From left to right: Raw, Rounded, and Un-rounded.

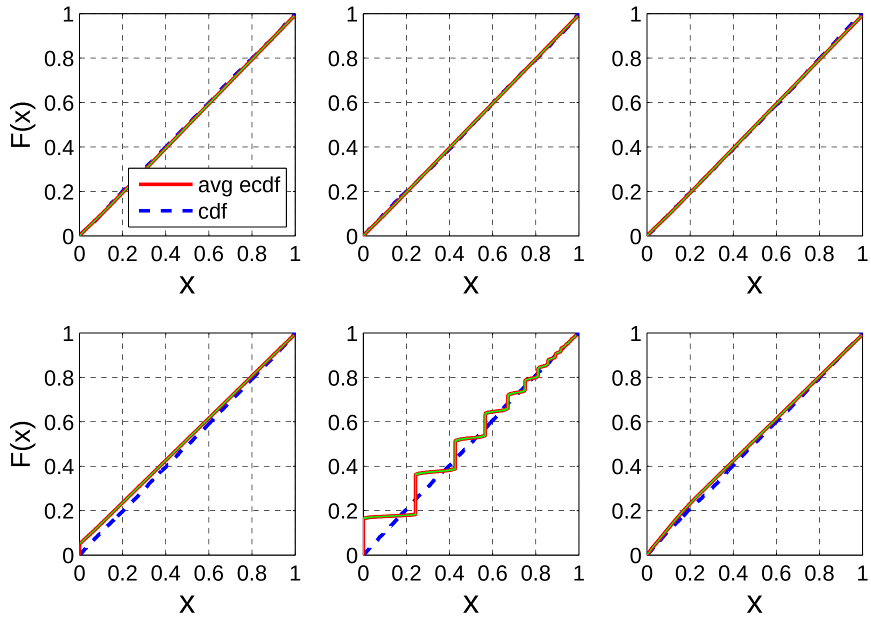

Table 3: Test results of batch Poisson processes on $[0, 6]$ in which every kth point comes in pairs; the total rate is kept the same. Results over 10000 iterations.

		CU		Log		Lewis	
	Type	# P	ave[p-value]	# P	ave[p-value]	# P	ave[p-value]
k=1	Raw	6801	0.21	0	0.00	0	0.00
	Rounded	6796	0.21	0	0.00	0	0.00
	Un-rounded	6802	0.21	0	0.00	0	0.00
k=3	Raw	8671	0.36	0	0.00	0	0.00
	Rounded	8669	0.36	0	0.00	0	0.00
	Un-rounded	8670	0.36	0	0.00	0	0.00
k=6	Raw	9179	0.43	0	0.00	0	0.00
	Rounded	9178	0.43	0	0.00	0	0.00
	Un-rounded	9181	0.43	0	0.00	0	0.00
k=9	Raw	9196	0.45	0	0.00	0	0.00
	Rounded	9194	0.45	0	0.00	0	0.00
	Un-rounded	9195	0.45	0	0.00	0	0.00

exceptionally long interarrival times in the H_2 process can still be detected. In contrast, the results in Table 5 for the batch-Poisson example show that the Lewis test consistently fails to reject when there is rounding to the nearest minute and then un-rounding.

Figure 2: Comparison of the average ecdf of a batch Poisson process on $[0, 6]$ in which every 6th point comes in pairs. From top to bottom: CU, Lewis test. From left to right: Raw, Rounded, and Un-rounded.

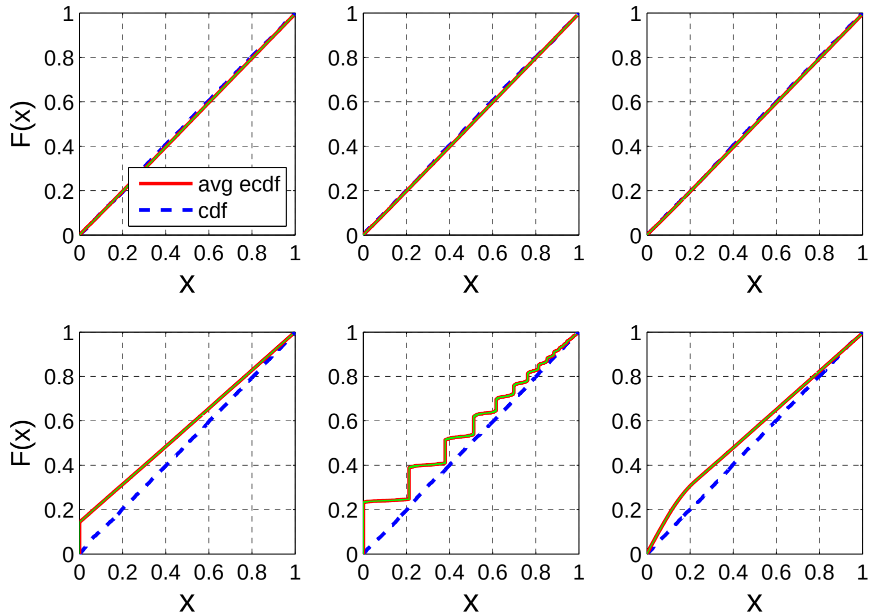

Table 4: [Compare to Tables 1 and 2 of the Main Paper] Results of the two KS tests with rounding and un-rounding

		CU			Lewis		
Interarrival Times	Type	# P	ave[p-value]	ave[% 0]	# P	ave[p-value]	ave[% 0]
M	Raw	944	0.50	0.0	955	0.50	0.0
	Rounded	927	0.40	0.1	0	0.00	94.0
	Un-rounded	949	0.50	0.0	946	0.51	0.0
H_2	Raw	705	0.21	0.0	0	0.00	0.0
	Rounded	630	0.16	0.1	0	0.00	94.0
	Un-rounded	704	0.22	0.0	42	0.01	0.0

4.3 Deciding When Un-Rounding is Needed and Will Be Effective

If the rounding is sufficiently fine (e.g., to the nearest millisecond in the examples above with a mean interarrival time of 3.6 seconds), then un-rounding is unnecessary; on the other hand, if the rounding is sufficiently coarse (e.g., to the nearest minute in the batch-Poisson renewal process example above), then it can cause a loss in power. In applications, the rounding should be judged relative to the mean interarrival time. Rounding to the nearest second was found to be important, and un-rounding was found to be effective, when the mean interarrival time was 3.6 seconds. For alternative hypotheses that differ from a PP only

Table 5: [Compare to Table 2: Rounding done to the nearest minutes] Test results of a rate-1000 renewal process on $[0, 6]$, in which the interarrival times are 0 with probability p and an exponential random variable with probability $1 - p$. Results over 10000 iterations.

		CU		Log		Lewis	
p	Type	# P	ave[p-value]	# P	ave[p-value]	# P	ave[p-value]
0.1	Raw	9015	0.41	0	0.00	0	0.00
	Rounded	8580	0.32	0	0.00	0	0.00
	Un-rounded	9019	0.41	9156	0.45	8774	0.40
0.05	Raw	9306	0.46	0	0.00	0	0.00
	Rounded	8959	0.36	0	0.00	0	0.00
	Un-rounded	9307	0.46	9421	0.49	9289	0.47
0.01	Raw	9453	0.49	8798	0.26	7879	0.22
	Rounded	9178	0.40	0	0.00	0	0.00
	Un-rounded	9451	0.49	9484	0.50	9480	0.50

through their local behavior, like the two batch-Poisson process examples above, rounding to the nearest minute, which corresponds to 16.7 mean interarrival times, and then un-rounding, can virtually eliminate all power. However, the un-rounding after rounding to the nearest minute is still effective for the H_2 example above, which also has more longer interarrival times than a PP. Overall, it is reasonable to expect that rounding will not matter if it is very fine, e.g., to less than 0.01 mean service time, while there is a danger of a loss of power if it is too coarse, e.g., to more than a mean service time. The different behavior for these different examples show that it must be difficult to develop a simple criterion. Thus, we recommend using simulation to investigate in specific instances, as we have done here.

5 An Example Illustrating Theorem 6

In this section, we illustrate Theorem 6 by considering the linear arrival rate function in (12) –$\lambda(t) = a + bt$– of the main paper with $a = 1$, $b = r = 10$ and $T = 10$. We then scale by multiplying this arrival rate function by n. Thus, the expected total number of arrivals is $510n$. For each n, we divide the interval $[0, 10]$ into k_n equally spaced subintervals. Table 6 shows the performance of the CU and Lewis KS tests as a function of n for various choices of k_n. #P is the number of KS tests passed at significance level $\alpha = 0.05$ out of 1000 replications and ave[p-value] is the average p-values. We see that the conclusions of Theorem 6 are strongly supported by these experimental results.

Table 6: Performance of the CU and Lewis KS tests with different values of k_n.

			CU		Lewis	
k_n	n	k_n	# P	$ave[p - value]$	# P	$ave[p - value]$
$\lfloor n^{1/4} \rfloor$	1	1	0	0.00	0	0.00
	5	1	0	0.00	0	0.00
	10	1	0	0.00	0	0.00
	50	2	0	0.00	0	0.00
	100	3	0	0.00	0	0.00
	500	4	0	0.00	0	0.00
	1000	5	0	0.00	0	0.00
$\lfloor n^{1/2} \rfloor$	1	1	0	0.00	0	0.00
	5	2	0	0.00	27	0.01
	10	3	0	0.00	147	0.03
	50	7	0	0.00	671	0.24
	100	10	0	0.00	802	0.33
	500	22	0	0.00	913	0.46
	1000	31	0	0.00	919	0.47
$\lfloor n^{3/4} \rfloor$	1	1	0	0.00	0	0.00
	5	3	0	0.00	484	0.14
	10	5	0	0.00	825	0.35
	50	18	9	0.00	944	0.48
	100	31	36	0.01	952	0.49
	500	105	307	0.07	953	0.51
	1000	177	443	0.12	957	0.51
n	1	1	0	0.00	0	0.00
	5	5	0	0.00	872	0.38
	10	10	32	0.01	941	0.49
	50	50	612	0.21	941	0.50
	100	100	795	0.32	944	0.50
	500	500	927	0.47	946	0.50
	1000	1000	940	0.48	945	0.49

6 More on Asymptotics of the CU KS Test

In this section we present an asymptotic result paralleling Theorem 6 of the main paper for a *piecewise-smooth* arrival rate function. Table 7 shows the results for arrival rate function $\lambda(t) = 100 + 20sin(t)$ on the time interval $[0, 10]$. The results can be compared to those of Table 6. Together, they support that it is not strictly necessary that the arrival rate function be piecewise-linear in order for the asymptotic correctness of the piecewise-constant approximation.

We consider a *piecewise-smooth* arrival rate function λ on the interval $[0, T]$, by which we mean that there are at most $m < \infty$ points t_i with $t_0 \equiv 0 < t_1 < \cdots t_m < T \equiv t_{m+1}$, such that the arrival rate function is right continuous on $[t_{j-1}, t_j)$ with left limits at t_j for each j and differentiable with derivative $\dot{\lambda}$ on (t_{j-1}, t_j) for all j with

$$|\dot{\lambda}(t)| \le K < \infty \quad \text{for} \quad t_{j-1} < t < t_j, \quad 1 \le j \le m+1. \tag{6.1}$$

As a consequence, λ is Lipschitz continuous on each interval $[t_{j-1}, t_j)$. We will also assume that λ is *strictly positive*, by which we mean that it and its left limits at t_j, $1 \le j \le m$, are positive.

Theorem 6.1 (*asymptotic justification of the piecewise-constant approximation of piecewise-smooth functions*) *Suppose that we consider a strictly positive piecewise-smooth arrival rate function over the fixed interval* $[0, T]$ *as above scaled by* n. *Let* m *be the number of discontinuity points, so that they partition the interval into* $m+1$ *disjoint subintervals, closed on the left and open on the right (except at the right endpoint* T*) over which the arrival rate function is Lipshitz continuous. Suppose that we use the CU KS test with any specified significance level* α *based on combining data over* $(m+1)k_n$ *subintervals, where each of the* $m+1$ *initial subintervals determined by the discontinuity points is partitioned into* k_n *equally spaced subintervals. If condition (23) and (24) in the main paper hold, then the probability that the CU KS test of the hypothesis of a Poisson process will reject the NHPP converges to* α *as* $n \to \infty$.

Proof. The assumed strict positivity and Lipschitz continuity implies that the arrival rate function is bounded below by a constant $c > 0$. By these properties, there is a constant K such that the function oscillates by at most $K\delta$ over every subinterval (between discontinuity points) that is of width δ. Hence, we can apply essentially the same argument used for the proof of Theorem 3.6 of the main paper. ■

References

Durbin, J. 1961. Some methods for constructing exact tests. *Biometrika* **48**(1) 41–55.

Lewis, P. A. W. 1965. Some results on tests for Poisson processes. *Biometrika* **52**(1) 67–77.

Whitt, W. 1982. Approximating a point process by a renewal process: two basic methods. *Oper. Res.* **30** 125–147.

Table 7: Performance of the CU and Lewis KS tests with different values of k_n for a piecewise-smooth arrival rate function.

			CU		Lewis	
k_n	n	k_n	# P	$ave[p-value]$	# P	$ave[p-value]$
$\lfloor n^{1/4} \rfloor$	1	1	488	0.10	924	0.47
	5	1	0	0.00	855	0.37
	10	1	0	0.00	762	0.30
	50	2	0	0.00	527	0.16
	100	3	29	0.01	860	0.37
	500	4	0	0.00	180	0.04
	1000	5	0	0.00	104	0.02
$\lfloor n^{1/2} \rfloor$	1	1	522	0.10	926	0.47
	5	2	209	0.04	840	0.36
	10	3	875	0.30	928	0.48
	50	7	891	0.41	925	0.45
	100	10	837	0.38	940	0.46
	500	22	906	0.43	954	0.50
	1000	31	918	0.46	969	0.52
$\lfloor n^{3/4} \rfloor$	1	1	522	0.10	926	0.47
	5	3	880	0.37	947	0.50
	10	5	894	0.44	929	0.46
	50	18	944	0.51	942	0.50
	100	31	931	0.50	951	0.49
	500	105	940	0.49	957	0.50
	1000	177	943	0.49	970	0.52
n	1	1	522	0.10	926	0.47
	5	5	917	0.46	943	0.50
	10	10	947	0.48	939	0.48
	50	50	953	0.50	948	0.50
	100	100	937	0.48	949	0.49
	500	500	955	0.50	954	0.51
	1000	1000	949	0.49	965	0.52