Warm-up interval

Let $W_1, W_2, \ldots, W_N$ be realizations of waiting times in a single run, and suppose we want to estimate the steady-state mean waiting time $E(W)$, defined as

$$E(W) = \lim_{j\to\infty} E(W_j)$$

by the sample mean

$$\bar{W}_N = \frac{1}{N}\sum_{j=1}^{N} W_j$$

In this estimate there are two types of errors:

- *Systematic error, or bias*
  This means that
  $$E(\bar{W}_N) \neq E(W),$$
  due to the influence of the initial conditions, which may not be "representative" for steady-state behavior;
- *Sampling (or random) error*
  The estimator $\bar{W}_N$ is of course a random variable.

To reduce the systematic error, we *delete the initial observations*, say $W_1, \ldots, W_k$, and use the remaining observations $W_{k+1}, \ldots, W_N$ to estimate $E(W)$ by the *truncated* sample mean

$$\bar{W}_{k,N} = \frac{1}{N-k} \sum_{j=k+1}^{N} W_j$$

Then one expects that $\bar{W}_{k,N}$ is less biased than $\bar{W}_N$, since the observations near the beginning of the simulation may not be representative for steady-state behavior; the parameter k is called the *warm-up interval*.

How to choose the warm-up interval k?

We like to pick k such that $E(\bar{W}_{k,N}) \approx E(W)$.

- If k is too small, then $E(\bar{W}_{k,N})$ may be significantly different from $E(W)$;
- If k is too large, then the variance of $\bar{W}_{k,N}$ (the sampling error) may be too large (its variance is proportional to $1/(N-k)$).

Waiting times (higly oscillating curve) and truncated sample means (smoother curve) for the $M/M/1$ queue with $\lambda = 0.5$ and $\mu = 1$;

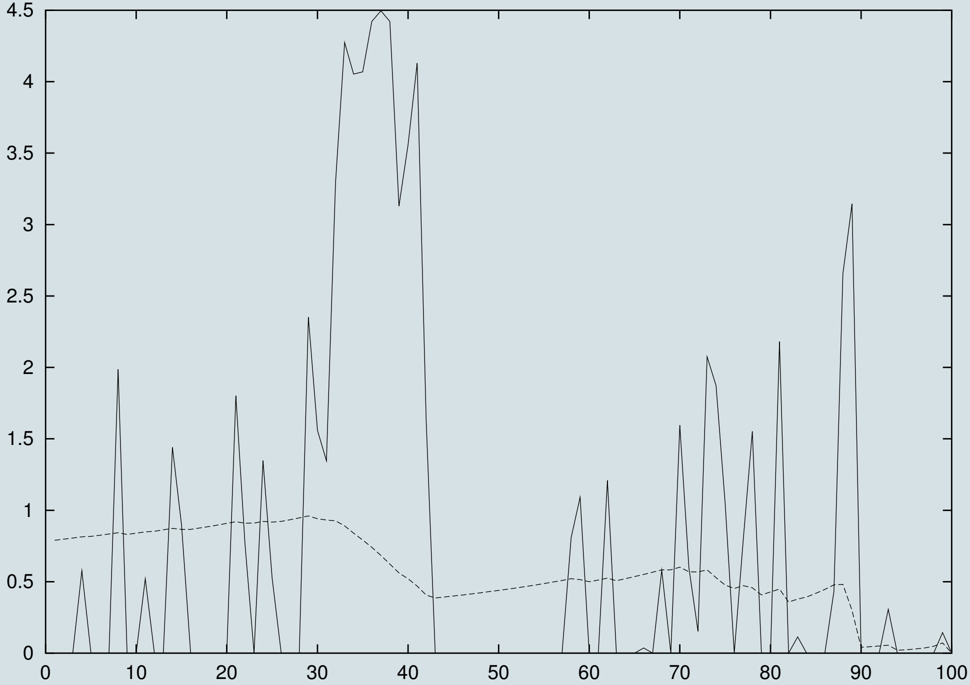

Graphical procedure to determine k

Our goal is determine a value k such that

$$E(W_j) \approx E(W)$$

for all $j > k$.

The presence of variability of the process $W_1, W_2, \ldots$ makes it hard to determine k from a single run.

Therefore, the idea is to make n independent replications (by using different random numbers) and employing the following steps:

1. Make n independent replications (or runs), each of length N; let $W_j^{(i)}$ denote the j-th waiting time in run i.
2. Let

$$\bar{W}_j = \frac{1}{n} \sum_{i=1}^{n} W_j^{(i)}$$

   The *averaged process* $\bar{W}_1, \bar{W}_2, \ldots$ has means and variances

$$E(\bar{W}_j) = E(W_j), \qquad \mathrm{var}(\bar{W}_j) = \mathrm{var}(W_j)/n.$$

   So its mean behavior is the same as the original process, but it has a smaller ($1/n$-th) variance.
3. Plot $\bar{W}_j$ and choose k such that beyond k the process $\bar{W}_1, \bar{W}_2, \ldots$ appears to have converged.

Waiting times for the $M/M/1$ queue with $\lambda = 0.5$ and $\mu = 1$;

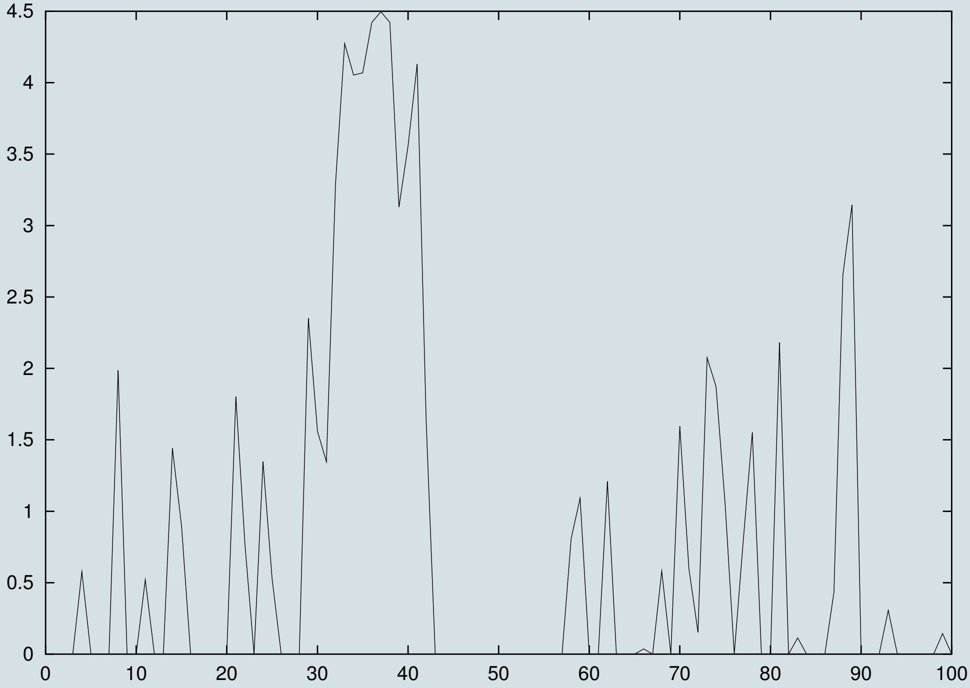

Averaged waiting times for the $M/M/1$ queue with $\lambda = 0.5$ and $\mu = 1$; number of replications is 5.

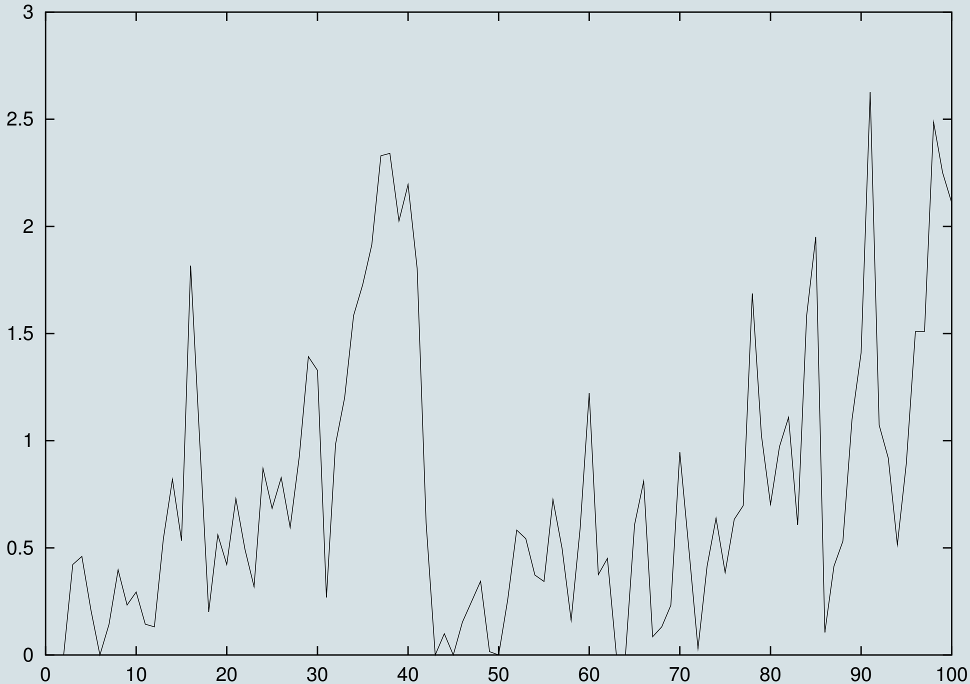

Averaged waiting times for the $M/M/1$ queue with $\lambda = 0.5$ and $\mu = 1$; number of replications is 10.

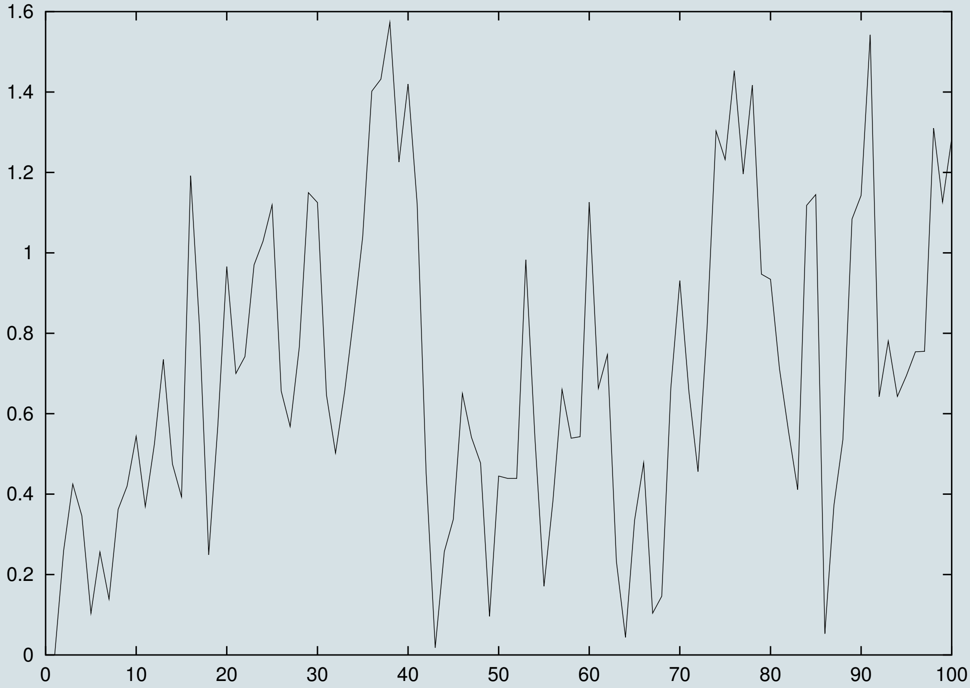

Averaged waiting times for the $M/M/1$ queue with $\lambda = 0.5$ and $\mu = 1$; number of replications is 100.

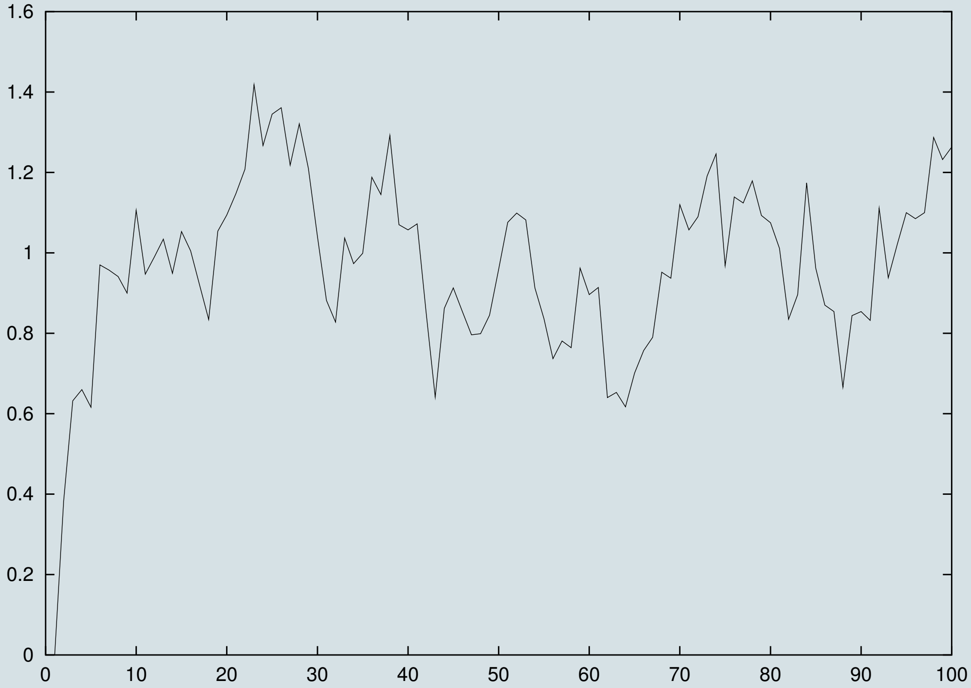

To smooth “high-frequency” oscillations in $\bar{W}_1, \bar{W}_2, \ldots$ (but leave the trend) one may consider the *moving average* $\bar{W}_j(w)$ (where w is the *window size*) defined as:

$$\bar{W}_j(w) = \frac{1}{2w+1} \sum_{i=-w}^{w} \bar{W}_{j+i}$$

for $j = w+1, w+2, \ldots, N-w$, and

$$\bar{W}_j(w) = \frac{1}{2j-1} \sum_{i=-(j-1)}^{j-1} \bar{W}_{j+i}$$

for $j = 1, \ldots, w$.

The warm-up interval k can then be determined from the plot of $\bar{W}_j(w)$ for $j = 1, \ldots, N-w$.

Moving average of averaged waiting times for the $M/M/1$ queue with $\lambda = 0.5$ and $\mu = 1$; number of replications is 10 and window size is 5.

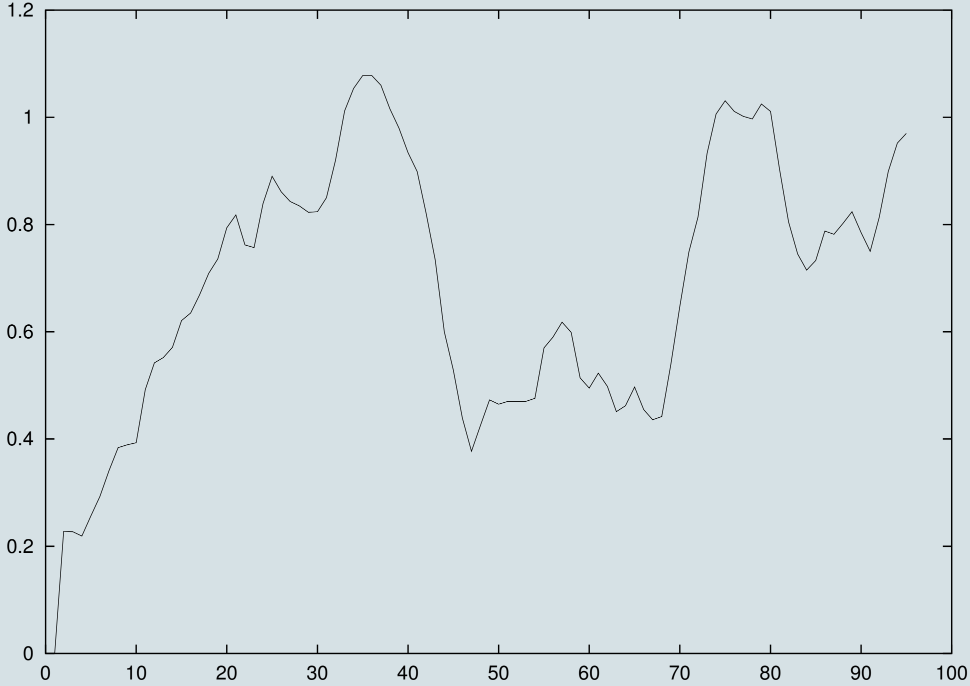

Moving average of averaged waiting times for the $M/M/1$ queue with $\lambda = 0.5$ and $\mu = 1$; number of replications is 10 and window size is 10.

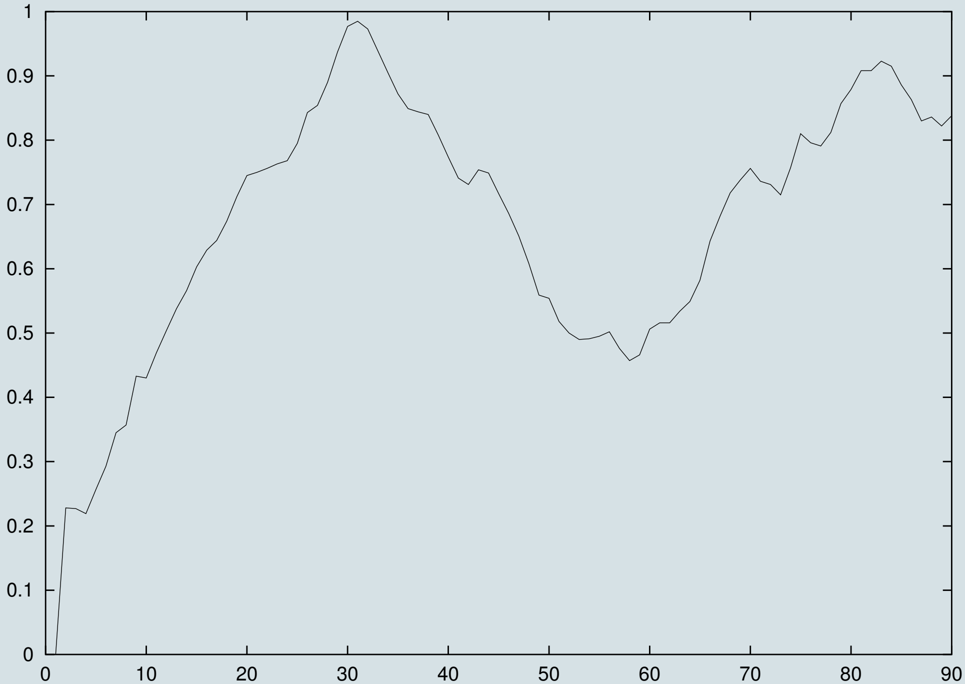

Moving average of averaged waiting times for the $M/M/1$ queue with $\lambda = 0.5$ and $\mu = 1$; number of replications is 100 and window size is 10.

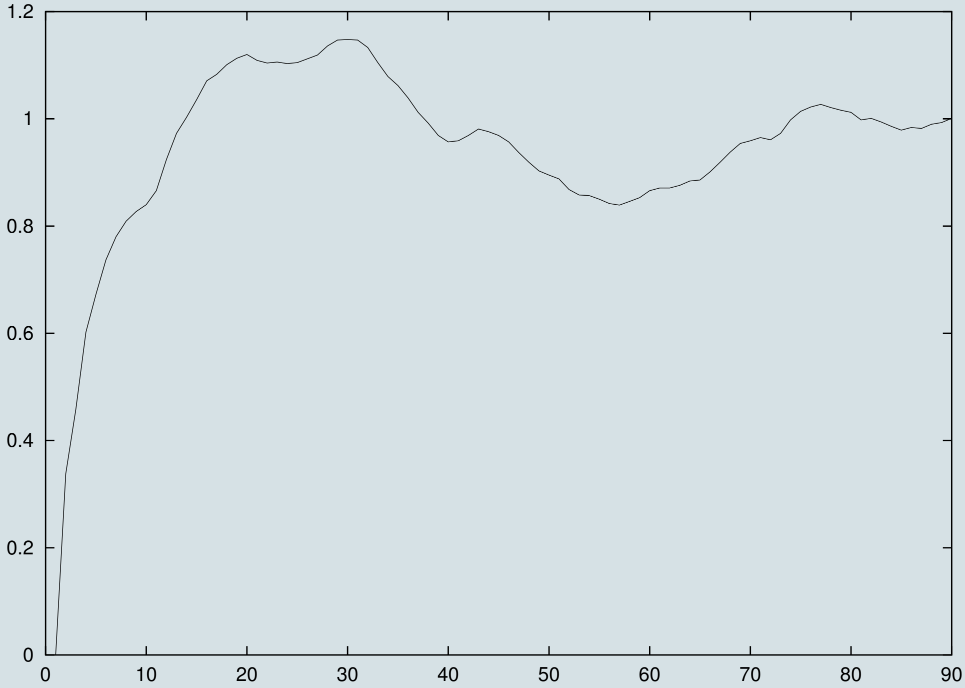

Observations:

- As k increases for fixed N, then the systematic error decreases, but the random error increases;
- As N increases for fixed k, then both systematic and random error decrease;
- Averages (or moving averages) based on n independent replications (that start in the same initial state) provide a basis for determining the warm-up interval k;
- Random fluctuations in these averages decrease when the number of replications increases;
- Systematic errors in these averages remain unaffected by increasing the number of replications.

Interval estimates

Let $X_1, X_2, \ldots, X_N$ be the output of a single run; for example, X_i is the waiting time of the i-th customer.
Suppose we want to estimate the steady-state mean

$$\mu = \lim_{i \to \infty} E(X_i)$$

by the (truncated) sample mean

$$Y = \frac{1}{N-k} \sum_{j=k+1}^{N} X_j$$

where k is the warm-up interval; k can be determined by a graphical procedure.

The sample mean Y is a *point estimate*; there are several approaches for obtaining an interval estimate for the steady-state mean μ.

Independent replications

Make n independent runs, each of N observations, where N is much larger then the warm-up interval k. Let $X_j^{(i)}$ denote the j-th realization time in run i and

$$Y_i = \frac{1}{N-k} \sum_{j=k+1}^{N} X_j^{(i)}$$

for $i = 1, \ldots, n$. So Y_i only uses the observations from run i corresponding to ‘steady-state.’

The Y_i's are i.i.d. random variables with $E(Y_i) \approx \mu$, so the sample mean $\bar{Y}(n)$ is an unbiased estimater for μ, and an *approximate* $100(1-\delta)$% confidence interval for μ is given by

$$\bar{Y}(n) \pm t_{n-1,1-\delta/2}\frac{S(n)}{\sqrt{n}}$$

where

$$\bar{Y}(n) = \frac{1}{n}\sum_{i=1}^{n} Y_i \ ;$$

$$S^2(n) = \frac{1}{n-1}\sum_{i=1}^{n}(Y_i - \bar{Y}(n))^2$$

and $t_{n-1,\beta}$ denotes the β-quantile of the Student's t distribution with $n-1$ degrees of freedom.

It is an *approximate* confidence interval for μ, because:

- $E(Y_i) \approx \mu$;
- Y_i is approximately normally distributed.

Under certain conditions (AWA; see Th. 6.3 in DES) it holds for fixed k and n, that

$$(\bar{Y}(n) - \mu)/\sqrt{S^2(n)/n} \xrightarrow{d} \tau_{n-1}$$

as $N \to \infty$, where τ_{n-1} denotes Student's t distribution with $n-1$ degrees of freedom.

Hence, the confidence interval for μ is *asymptotically valid* as $N \to \infty$.

Comments on the independent replication approach:

- It is easy to understand and implement;
- It gives reasonably good statistical performance;
- The approach applies to many output parameters (waiting times, queue lenghts, etc.);
- It can be easily used to simultaneaously estimate different parameters for the same simulation model.

Batch means

Instead of doing n independent runs, we try to obtain n independent observations by making a *single long run* and, after deleting the first k observations, dividing this run into n subruns.

The advantage is that we have to go through the warm-up period only once.

Let $X_1, X_2, \ldots, X_N$ be the output of a single run, where we have already deleted the first k observations and renumbered the remaining ones. Hence $X_1, X_2, \ldots, X_N$ will be approximately stationary. We divide the observations $X_1, X_2, \ldots, X_N$ into n batches of length b (assume $N = nb$). Thus, batch 1 consists of

$$X_1, X_2, \ldots, X_b;$$

batch 2 of

$$X_{b+1}, X_{b+2}, \ldots, X_{2b},$$

and so on. Let Y_i be the sample (or batch) mean of the b observations in batch i, so

$$Y_i = \frac{1}{b} \sum_{j=(i-1)b+1}^{ib} X_j$$

The Y_i's play the same role as the ones in the independent replication method. Unfortunately, the Y_i's will now be *dependent*.

But, under mild conditions, for *large* b the Y_i's will be approximately independent and normally distributed, each with the same mean μ and the same variance.
Hence, for b large enough, it is reasonable to treat the Y_i's as i.i.d. normal random variables with mean μ; thus

$$\bar{Y}(n) \pm t_{n-1,1-\delta/2} \frac{S(n)}{\sqrt{n}}$$

provides again a $100(1-\delta)\%$ confidence interval for μ.

Dealing with dependence

This approach 'neglects' the dependence between the Y_i's; alternatively, we may take into account (part of) the dependence when constructing a confidence interval.

The classical Central Limit Theorem states that for i.i.d. random variables $Y_1, \ldots, Y_n$,

$$\frac{\sum_{i=1}^n Y_i - E(\sum_{i=1}^n Y_i)}{\sqrt{\mathrm{var}(\sum_{i=1}^n Y_i)}}$$

is approximately standard normally distributed. This remains valid if the Y_i's are *weakly dependent*:

$$\mathrm{cov}(Y_i, Y_{i+j}) \longrightarrow 0$$

as $j \longrightarrow \infty$.

For a stationary process $Y_1, Y_2, \ldots$ we have

$$\mathrm{var}\left(\sum_{i=1}^{n} Y_i\right) = n\mathrm{var}(Y) + 2\sum_{j=1}^{n-1}(n-j)\mathrm{cov}(Y_1, Y_{1+j})$$

In the first approach we replaced the variance of the sum

$$\mathrm{var}\left(\sum_{i=1}^{n} Y_i\right)$$

by

$$n\mathrm{var}(Y)$$

As often the case in practice, the X_i's and thus the Y_i's are *positively correlated*. Then the estimator for the variance will be biased low, which gives a confidence interval that is too small.

Thus the confidence interval will cover μ with a probability that is smaller than the desired $1 - \delta$.

If we want to take into account part of the dependence, say the dependence between succesive observations, then we can replace

$$\mathrm{var}\left(\sum_{i=1}^{n} Y_i\right)$$

by

$$n\mathrm{var}(Y) + 2(n-1)\mathrm{cov}(Y_1, Y_2)$$

and then construct a confidence interval.

The unknown quantities $\text{var}(Y)$ and $\text{cov}(Y_1, Y_2)$ can be estimated by their sample estimates

$$\frac{1}{n-1}\sum_{i=1}^{n}(Y_i - \bar{Y}(n))^2$$

and

$$\frac{1}{n-1}\sum_{i=1}^{n-1}(Y_i - \bar{Y}(n))(Y_{i+1} - \bar{Y}(n))$$

In your final assignment please pay attention to:

- Input analysis
  (How does it look? Fitting? ...)
- Output analysis
  (Warm-up interval, length and number of runs, confidence intervals, ...)
- Model description
  (Modeling assumptions, simplifications, ...)
- Validation of simulation model
  (Correctness, special cases, ...)
- Presentation and discussion of results
  (What should be shown and how? ...)