

Chapter 8

Analyzing the Sample Path

The QSIM application provides some built-in analysis capabilities. You can obtain sample means, variance estimates, and other simple statistics on several measures of performance including time spent in system components and resource utilization. You can also display simple lineplots, histograms, and boxplots to dynamically observe system behavior.

An additional feature is the ability to accumulate the sample paths in SAS data sets. See Chapter 7, “Saving and Restoring,” for more information. With these data in data sets, you can write your own programs to do more sophisticated analyses.

When calculating statistics, QSIM uses the following measures of system performance:

- number in a queue
- time in a queue
- number of multiple-servers busy
- server utilization
- time in service
- time in the system.

In addition, statistics can be calculated on samples of random variables, which are user-defined attributes.

To illustrate some of the analysis features, this chapter uses the multiple-server, single-input queuing model shown in Figure 8.1.

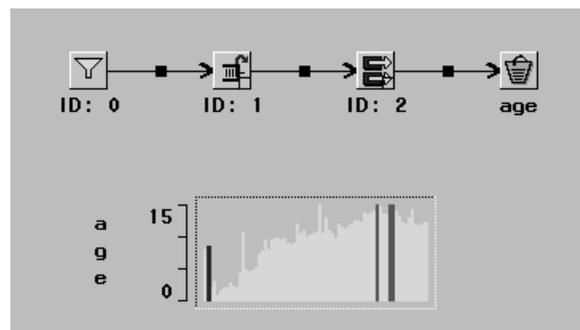


Figure 8.1. A Single Queue 3-Server Model

The inter-arrival time distribution is $\mathcal{E}(.33)$ and the service time distribution for each server is $\mathcal{E}(1)$. Transactions queue in a FIFO discipline and go to the first available server. This corresponds to the type of queuing that you encounter at check-in in many airports.

Running a Simulation

Suppose that you have built the model in Figure 8.1. You start and stop this simulation using the button on the command panel in the simulation window. The  button starts the simulation by sending a start message to each of the components in the model. Similarly, the  button stops the simulation.



Figure 8.2. The Command Panel and Buttons

Also note that the  button stops and resets the simulation by sending the reset message to each component. But first it stops the simulation, resets the simulation clock, and removes all pending actions. For example, all queues are emptied and transactions in service are discarded.

The  button toggles animation on and off. Animation is one effective way to validate a model. You simply turn animation on and observe the simulation behavior. In addition, you can add Triggers and StringHolder to display and print state information to help you with model validation.

Another way to control the simulation is with the Component Control Panel as shown in Figure 8.3.

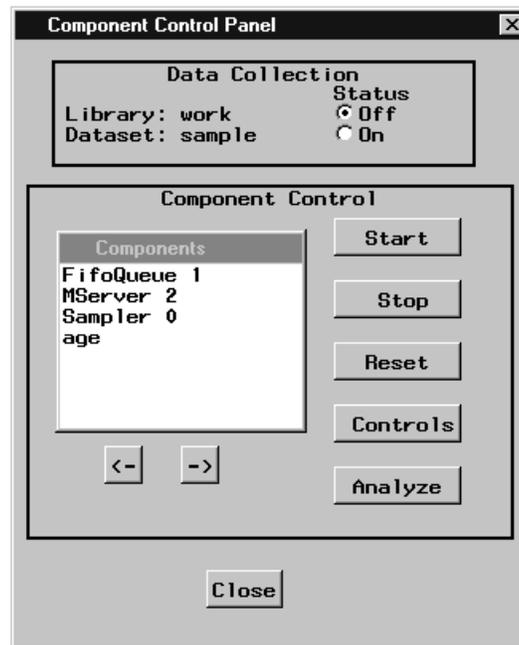


Figure 8.3. The Component Control Panel

Notice the buttons on the right side of the control panel. When any of these buttons are clicked a message is sent to each of the selected components. For example, you can select a component from the **Components** list box and start it only by clicking the **Start** button. It sends the start message to the selected component. This is typically *not* the way to start the simulation since not all components will be started, but it is the typical way to analyze component performance with the **Analyze** button.

Statistics on Component State

The queues and servers are two component types for which performance measures are typically of interest. For queues you want to know the distribution on queue length, and for servers you want to know about server utilization. Other performance measures that may be of interest are time in the queue and time in service. These are discussed in the “Statistics on Transaction State” on page 77.

There are two ways you can obtain statistics on queue length and server utilization. One way uses the chart components and the other uses the Simulation control panel shown in Figure 8.3. When the data collection status is “On” in this control panel, all changes of state are recorded to a SAS data set. These data can be summarized through the Simulation Control Panel and through the **Analyze** button.

Queue Length

The QSIM Application will generate and execute SAS code that calculates summary statistics on the length of a queue. When you want to start collecting the sample data, you select the **On** radio button in the Simulation control panel. This starts data collection. When you are ready to generate sample statistics select the component for which you want the statistics and then click the **Analyze** button. For example, you

display data on the queue length on the single-queue multiple-server model in Figure 8.1 by selecting FIFOQueue 1 in the Simulation Control panel and then by pressing **Analyze**.

```

The UNIVARIATE Procedure
Variable: number (FifoQueue 1 Length)

Weight: time

Weighted Moments

N              1228      Sum Weights      207.439375
Mean           45.9018688  Sum Observations 9521.85497
Std Deviation  1.42854557      Variance          2.04074245
Skewness       -0.9007099   Kurtosis          3.0126739
Uncorrected SS 439574.929  Corrected SS      2503.99099
Coeff Variation 3.11217301 Std Error Mean     0.09918557

Weighted Basic Statistical Measures

Location              Variability
Mean      45.90187    Std Deviation  1.42855
Median    47.00000    Variance       2.04074
Mode      49.00000    Range          14.00000
                          Interquartile Range  5.00000

Weighted Tests for Location: Mu0=0

Test      -Statistic-    -----p Value-----
Student's t    t 462.7878    Pr > |t|    <.0001

Weighted Quantiles

Quantile      Estimate
100% Max      50
99%           50
95%           50
90%           50
75% Q3        49
50% Median    47
25% Q1        44
10%          40
5%           39
1%           37
0% Min        36

```

Figure 8.4. Statistics on Queue Length

This generates the report shown in Figure 8.4 in the SAS Output Window. Figure 8.4 shows that the sample mean queue length is approximately 45.90 and that the sample variance is approximately 2.04.

The **Analyze** button also opens a graphics window with the histogram on the queue length as shown in Figure 8.5. This shows the sample distribution of queue length for the complete sample path that was collected from the time the **Collect Data** check box is selected until the time the **Analyze** button is selected.

FifoQueue 1 Density Function Estimate

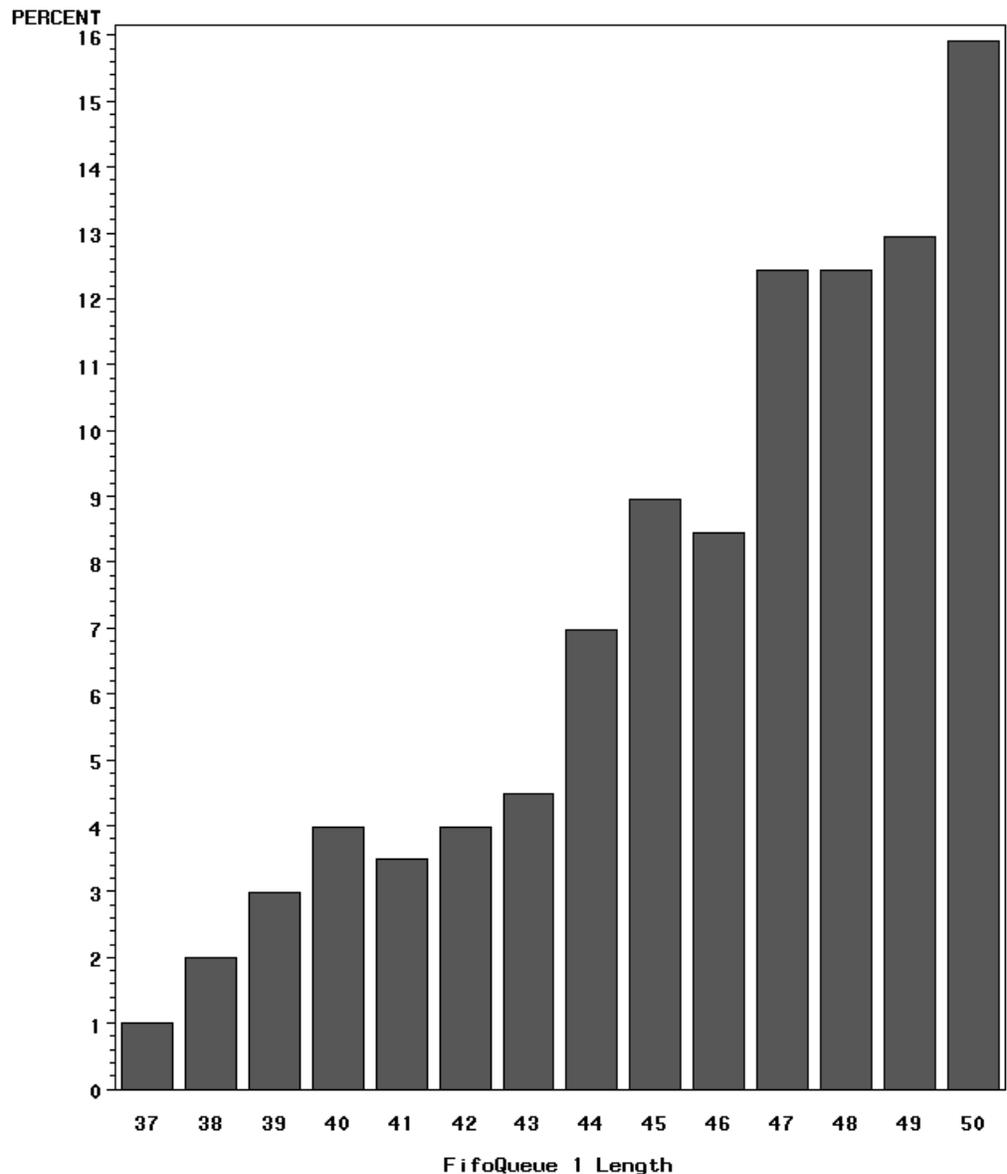


Figure 8.5. Queue Length Histogram

It is important to note that the observations in the sample path are not independent. There can be a significant amount of autocorrelation in the observations of queue length.

Another way to get Queue Length Statistics

Another way you can obtain statistics on queue length and server utilization is by using the chart components. You take the queue that you want to monitor and drag it into one of the chart components, the VHistogram for example. When you drop the queue, a control panel will open from which you can select the attribute of the component that you want sampled and the distribution to control sampling. Figure 8.6 shows an example. Here the `size` is selected to be monitored, and the component is

to be sampled with a deterministic inter-sample time of one time unit. This is the default sampling distribution. This is different than using the Simulation control panel because in that case *all* state changes are captured, when using a chart component the queue state is sampled.

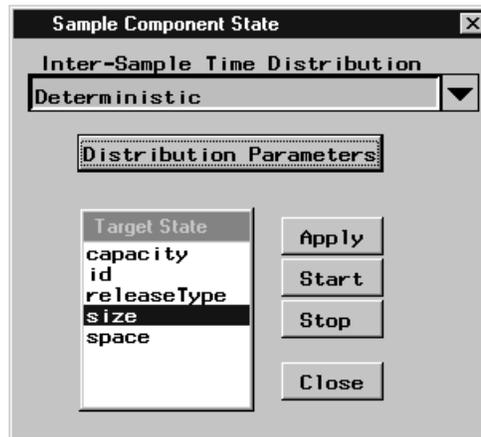


Figure 8.6. Sample Component State Control Panel

This control panel is opened when the queue labeled “ID: 1” in Figure 2.17 is dropped on the histogram.

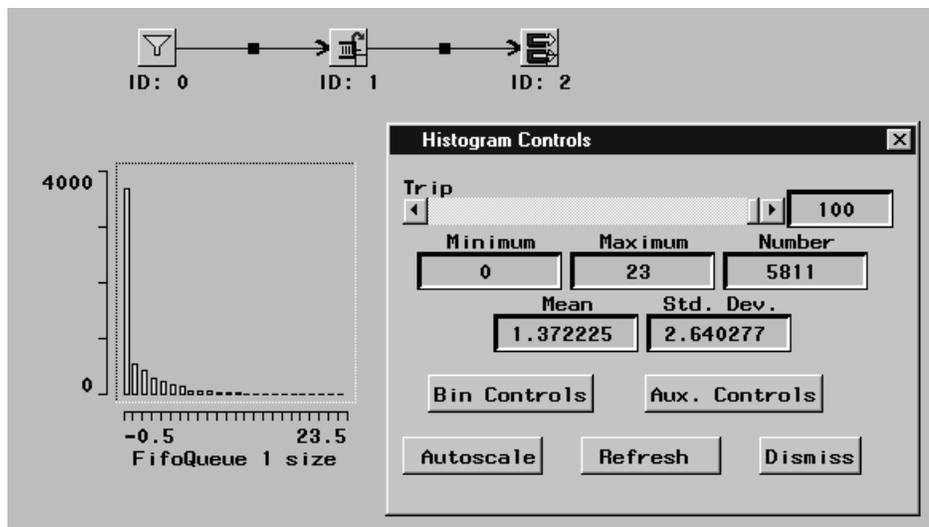


Figure 8.7. Vertical Histogram of the FIFO Queue length

Figure 2.17 also shows the histogram control panel opened. Notice that the control panel shows the simple statistics: minimum, maximum, mean, and standard deviation. Also note that when a component is dropped on a chart, a bucket is automatically associated with the chart component. The bucket is where the sample is collected. The controls for this implicit bucket can be viewed by clicking the **Bin Controls** button in the Histogram Controls window. It opens the control panel for the bucket as shown in Figure 2.15.

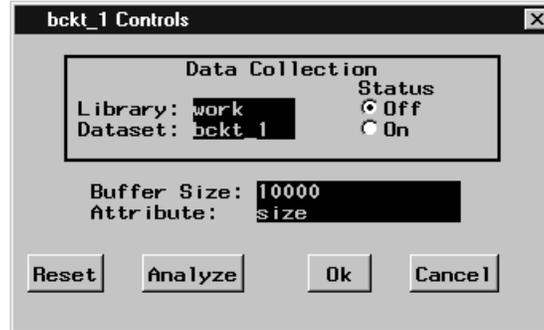


Figure 8.8. Bucket Control Panel

You can display more detailed statistics on queue length using SAS data sets. The “Data Collection” section in the Bucket control panel shows a data collection **Status** check box. If this is selected, then data on the component state changes are routed to the SAS data set named WORK.BCKT_1. If sample data have been collected for a time, then clicking the **Analyze** button will generate a univariate analysis performed by the UNIVARIATE procedure and a histogram.

Server Utilization

The utilization of a server is another measure of system performance that is often of interest. Estimates on the probability that a server is busy is one statistic for measuring utilization. If after collecting data using the **Collect Data** check box on the Simulation Control panel, you select a server from the components list box and click the **Analyze** button, you get a printout similar to that shown in Figure 8.9.

```

The UNIVARIATE Procedure
Variable: number (Utilization of Server 2)

Weight: time

Weighted Moments
N          4466      Sum Weights    2936.6774
Mean       0.77599376 Sum Observations 2278.84334
Std Deviation 0.33812444 Variance        0.11432813
Skewness   -2.6100791 Kurtosis         9.00063288
Uncorrected SS 2278.84334 Corrected SS    510.475122
Coeff Variation 43.5730869 Std Error Mean    0.00623948

Weighted Basic Statistical Measures

Location          Variability
Mean             0.775994      Std Deviation    0.33812
Median          1.000000      Variance         0.11433
Mode            0.000000      Range            1.00000
                                   Interquartile Range 0

NOTE: The mode displayed is the smallest of 2 modes with a count of 2233.

Weighted Tests for Location: Mu0=0

Test          -Statistic-      ----p Value-----
Student's t   t 124.3683      Pr > |t| < .0001

Weighted Quantiles
Quantile      Estimate
100% Max      1
99%           1
95%           1
90%           1
75% Q3        1
50% Median    1
25% Q1        1
10%           0
5%            0
1%            0
0% Min        0

```

Figure 8.9. Statistics on Server Utilization

The mean shown is approximately .77. This is an estimate of the probability that the server is busy.

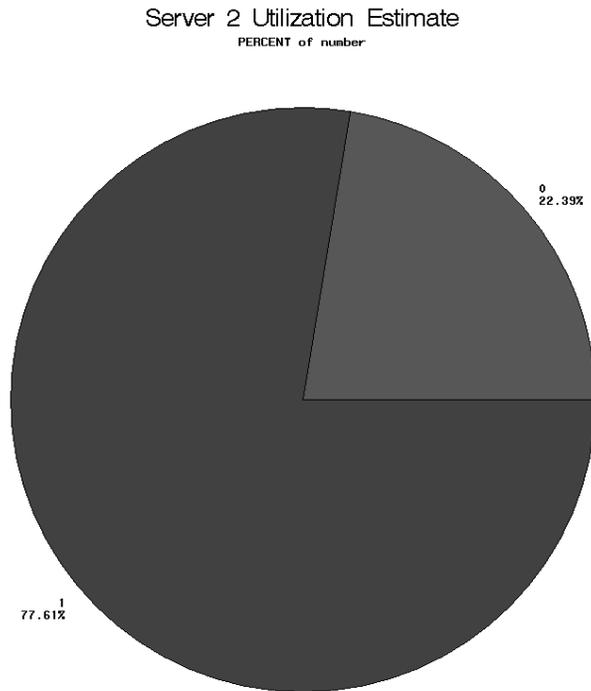


Figure 8.10. Server Utilization

The pie chart in Figure 8.10 shows the percent of time that the server is busy.

Another way to get Server Utilization Statistics

Another way to view server utilization statistics is with a chart component. One component that can be used is the VHistogram. Figure 8.12 shows a VHistogram and a Histogram control panel. This histogram had an MServer dropped on it and is monitoring the number of busy servers.

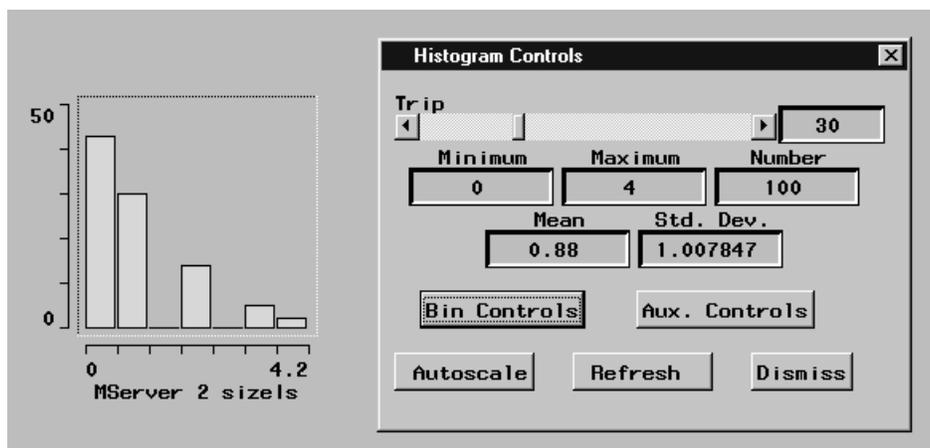


Figure 8.11. Histogram of Server Utilization

DATA Step Program for Extracting Information

When you click the **Analyze** button in the Simulation control panel (see Figure 8.3), the QSIM Application executes SAS code to extract and summarize the sample path saved in the WORK.SAMPLE data set. You, too, can write a SAS program to extract and analyze these data. For example, suppose you want to subset the data on the time the queue is in each state. The following DATA step subsets WORK.SAMPLE with those observations that have state information on queue with ID 5:

```
data subset;
  set sample;
  if id=5;
  keep timenow number time;
  label number="Queue Length";
  time = ( timenow - lag(timenow) );
  number = lag(state);
run;
```

The resulting data set has three variables: TIMENOW for the time that the state changes; NUMBER for the number in the queue; and TIME for the length of time in that state. The following SAS code executes the UNIVARIATE procedure to produce summary statistics and the GCHART procedure to produce a histogram as shown in Figure 8.5.

```
proc univariate data=subset;
  weight time;
  var number;
run;

proc gchart;
  label = ``Time in Queue``;
  vbar number / subvar=time discrete;
run;
```

A similar SAS program subsets the data on server utilization and produces the output in Figure 8.10 and Figure 8.9.

```
data subset;
  set sample;
  if id=4;
  keep timenow number time;
  label number="Utilization";
  time = ( timenow - lag(timenow) );
  number = lag(state);
run;

proc univariate data=subset;
  weight time;
  var busy;
run;
```

```
proc gchart;
  pie busy / sumvar=time discrete percent=outside;
run;
```

Statistics on Transaction State

As a transaction goes through the simulation, it maintains a history of the time spent at various components. When a transaction arrives at a Bucket, it can save any of its attributes. The name of the data set is derived from the number of buckets in the model and is displayed in the Bucket control panel as shown in Figure 2.15.

Time in the System

By default, the Bucket component accumulates the age or system time of the transactions arriving to it. However, it can accumulate any attribute of the transaction including user-defined attributes. The Bucket control panel has a check box which when set, records the attribute values in the SAS data set. Figure 2.15 shows such a control panel. If a simulation has been executing and the **Collect Data** check box has been set, then you can click the **Analyze** button to submit SAS code that calculates statistics on the sample.

Univariate statistics are calculated and histograms are printed for the attribute named in the Bucket control panel.

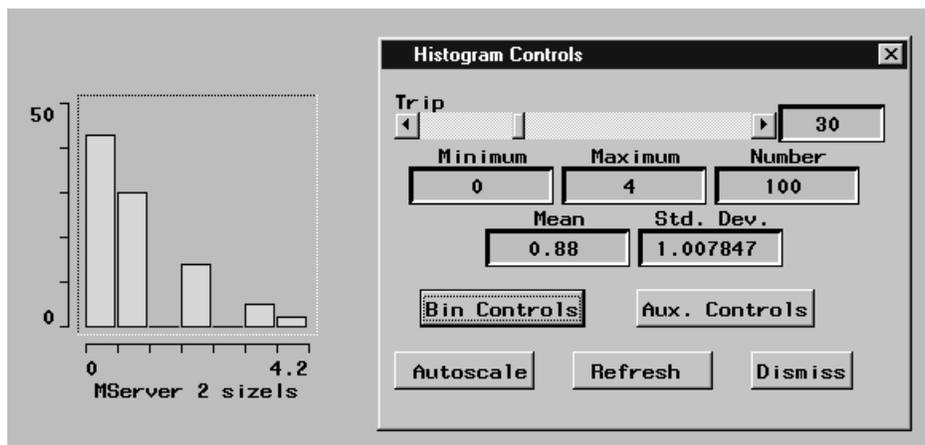


Figure 8.12. A Histogram Showing Time in the Queue

The control panel for the VHistogram maintains statistics on the sample in the Bucket and displays them as shown in Figure 8.12.

The correct bibliographic citation for this manual is as follows: SAS Institute Inc., *SAS/OR[®] User's Guide: QSIM Application, Version 8*, Cary, NC: SAS Institute Inc., 1999. 110 pp.

SAS/OR[®] User's Guide: QSIM Application, Version 8

Copyright © 1999 SAS Institute Inc., Cary, NC, USA.

ISBN 1-58025-488-8

All rights reserved. Printed in the United States of America. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, by any form or by any means, electronic, mechanical, photocopying, or otherwise, without the prior written permission of the publisher, SAS Institute, Inc.

U.S. Government Restricted Rights Notice. Use, duplication, or disclosure of the software by the government is subject to restrictions as set forth in FAR 52.227-19 Commercial Computer Software-Restricted Rights (June 1987).

SAS Institute Inc., SAS Campus Drive, Cary, North Carolina 27513.

1st printing, October 1999

SAS[®] and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. [®] indicates USA registration.

IBM[®], ACF/VTAM[®], AIX[®], APPN[®], MVS/ESA[®], OS/2[®], OS/390[®], VM/ESA[®], and VTAM[®] are registered trademarks or trademarks of International Business Machines Corporation. [®] indicates USA registration.

Other brand and product names are registered trademarks or trademarks of their respective companies.

The Institute is a private company devoted to the support and further development of its software and related services.