

Performance of STAR-System

STAR-Dundee's new software stack, STAR-System, provides a high performance software stack for accessing STAR-Dundee interface and router devices, including the STAR-Dundee PCI Mk2, cPCI Mk2 and PMC Mk2 devices and the new STAR-Dundee PCIe device.

One of the primary goals of the STAR-System APIs and drivers is to provide high performance while keeping CPU usage to a minimum. This application note describes how performance of the system is measured, and how you can repeat these tests. Most importantly, it provides details of the performance that can be achieved.

STAR-System Performance Tester

The STAR-System Performance Tester is an application provided with STAR-System. It was written to allow tests to be run to test the performance and CPU usage of both the software stack and the devices being used. It is also provided as source code, so is a useful example of how to write code to get the best performance from the API.

The Performance Tester provides three different types of test, each of which is described below. These tests can be used to transmit packets, receive packets, or both. Transmit only tests can be used to check the performance of the transmitter. The packets can be sent to an address where the packets will be dropped (e.g. an invalid port on a router) or to a receiver which is capable of receiving at the maximum link rate being used. Transmit only tests can also be used to test communication with another instance of the Performance Tester, which may be running on the same PC, or on a different PC. In such cases, the other instance of the Performance Tester should be used to receive packets. The most common way to use the software, however, is to both transmit and receive packets.

Maximum Data Rate Test

The maximum data rate test transmits and receives a large number of packets and times how long it takes from the first packet being transmitted to the last packet being received. It then uses this information to calculate the data and packet rates, and the average time per packet. The total CPU usage is also recorded, as is the CPU usage of the Performance Tester process. Assuming there is no other activity on the PC while the tests are running (which is unlikely due to virus scanners, etc.) the difference between these two figures will probably indicate work being done by the device driver, outside the threads of the Performance Tester process.

This test is repeated for each of the packet sizes specified. A minimum and maximum packet size is specified, as is the step to get from the minimum to the maximum. In each run of this test, the packet sizes remain fixed.

Random Packet Size Data Rate Test

The random packet size data rate test is very similar to the maximum data rate test – it transmits and receives packets as quickly as possible, recording data rate, packet rate, average packet time and CPU usage. The difference from the maximum data rate test is that the size of the packets being sent is random. Each test transmits a number of different packet sizes, each of which is between the current minimum size and the maximum size specified. The minimum size increases after each test by the step value specified.

Although this test does not provide an accurate measurement of the system (each run of the test is likely to produce different results), it is a useful indication of how the system can cope with packets of various sizes.

Latency Test

The latency test transmits a single packet then receives a single packet. It transmits the next packet once the last packet has been received. The test is repeated a number of times for a single packet size, before displaying the data rate, packet rate, average packet time and CPU usage. The test is then repeated for each packet size specified.

The important information from the latency test is the average packet time. This indicates the average time taken to transmit a packet and then receive that packet – the combined transmit and receive latency.

Performance Figures

The results obtained by the Performance Tester can be written to file, generating a text file containing tab separated values for the figures recorded in each test. This text file can be loaded in to a spreadsheet such as Microsoft Excel and graphed.

The graphs below show the results of running the Performance Tester on a reasonably low spec PC (using an Intel Core 2 Duo E7500 running at 2.93 GHz, with 2 GBytes of memory, and bought for around £300 in October 2011), with a SpaceWire PCI Mk2 device connected, and using various operating systems. The tests involve transmitting packets out of port 1 of the device, with a cable connected between ports 1 and 3, and receiving the packets on port 3.


Each test was performed with the links running at the PCI Mk2's maximum link speed of 200 Mbits/s. Note that SpaceWire encodes an 8-bit character as 10 bits, which means the maximum theoretical data rate achievable on a 200 Mbits/s link is 160 Mbit/s. Each packet is also terminated with a 4-bit end of packet marker which affects the data rate that can be achieved, particularly for smaller packets, e.g. a 1 byte packet will be represented using 14 bits.

The performance test results graphs show the packet size in bytes along the x-axis, with the data rate in Mbits/s and the processor usage as a percentage along the y-axis. The latency tests do not include processor usage figures, as these tests aren't designed to push the system to its limits. As with the performance test results, the latency test results show the packet size on the x-axis. The y-axis shows the average packet time, and hence the average latency, in microseconds.


Windows 7 32-bit Tests

The following tests were run on the 32-bit version of Windows 7.

Data Rate Test Results




Latency Test Results




Windows 7 64-bit Tests

The following tests were run on the 64-bit version of Windows 7.

Data Rate Test Results




Latency Test Results




Ubuntu Linux 32-bit Tests

The following tests were run on a 32-bit 2.6.38 Linux kernel using the Ubuntu 11.04 distribution.

Data Rate Test Results




Latency Test Results




Ubuntu Linux 64-bit Tests

The following tests were run on a 64-bit 2.6.38 Linux kernel using the Ubuntu 11.04 distribution.

Data Rate Test Results



Latency Test Results



Assessing the Results of the Tests

These graphs show that performance of STAR-System is relatively consistent across all platforms tested. Data rate throughput performance is very good, with the rates being very close to the maximum possible rates that can be achieved, even with packets as small as 60 bytes. The latency of the packets is reasonable, considering these tests were not being performed on real-time operating systems, with a worst case of less than 100 microsecond latency in addition to the latency on the link.

Latency is better on Windows than on Linux. There is very little difference between the 32-bit and 64-bit versions of Windows, although latency is slightly better on 64-bit. The granularity of the latency timing is not as good as on Linux, which is why a stepped graph is seen.

On Linux, throughput performance is better than Windows with smaller packets, and CPU usage is lower. Again, there isn't much difference between the 32-bit and 64-bit figures, although latency on the 32-bit kernel seems to be more variable.

For comparison, the graphs on the following page show the maximum theoretical data rate that can be achieved on a 200 Mbits/s link, and the minimum latency time to transmit a packet on a SpaceWire link at this rate. Note that these data rate figures are when packets are only going in one direction. If packets were being transmitted in both directions, the maximum rate would be between 4 and 8 Mbits/s lower due to the 4-bit FCTs which would also be transmitted on the link in each direction.

Although these results show that STAR-System and the SpaceWire PCI Mk2 provide excellent throughput and latency, STAR-Dundee's engineers are always endeavouring to make improvements. It is hoped that the CPU usage of STAR-System on all platforms can be reduced, and attempts will also be made to improve the latency on Linux in the near future.

Getting the Best Performance from Your Applications


An important point to note with any tests performed using the Performance Tester is that all tests indicate the maximum possible rates that can be achieved. These figures might not be achievable when other tasks are being performed with the packets and data being transmitted and received. For example, writing data to a file can be very time consuming, so it's unlikely that the same data rates can be achieved when writing each received packet to a file.

The code in the Performance Tester is also written to achieve the best possible performance. Multiple packets are transmitted at any one time when running the maximum and random packet size data rate tests, and it's unlikely that the same levels of performance will be achieved if each packet is transmitted individually.

When writing an application it's important to take in to consideration the characteristics of the packets to be transmitted and received. If packets should be transmitted with the minimum possible latency, then it might be necessary to transmit each packet as soon as it is ready. If throughput is most important, it might be better to transmit packets in groups.

STAR-Dundee prides itself on its support and can provide suggestions on the most effective implementation method to achieve the performance required for a given situation.

Theoretical Maximum for Performance Tests



Theoretical Maximum for Latency Tests

