FICON Bridge Channel Performance

Version 2.1

July 13, 2000

Wm. Todd Boyd
System Design and Architecture, Server Group
Todd Boyd/Poughkeepsie/IBM @ IBMUS
8-295-7942 or (914) 435-7942

Table of Contents

History

Version	Date	
v1.1	7/29/1999	Initial version
v2.0	4/30/2000	GA2 update
v2.1	7/13/2000	Added trademark page and pointer to configuration guidelines

.

Acknowledgment

This paper is intended to describe the performance of the FICON channel using the FICON Bridge feature in a 9032 model 5 ESCON Director. The data presented here is based upon measurements and carried out over the last several months. I would like to thank all of the reviewers for their comments but especially the following people for their continued support on this effort:

Rich Basener	Rich Basener/Poughkeepsie/IBM@IBMUS
Mario Borelli	Mario Borelli/Poughkeepsie/IBM@IBMUS
Cathy Cronin	Catherine Cronin/Poughkeepsie/IBM@IBMUS
Hsien Ko	Hsien Ko/Poughkeepsie/IBM@IBMUS
Pete McCaffrey	Peter McCaffrey/Poughkeepsie/IBM@IBMUS
Al Meritt	Al Meritt/Poughkeepsie/IBM@IBMUS
Jeff Seidell	Jeff Seidell/White Plains/IBM@IBMUS

Introduction

The new S/390 FICON channel, the latest in the long history of S/390 I/O channels, which was introduced in March of 1999, will be used with and/or in place of the present parallel and ESCON(tm) channels. The initial implementation of FICON utilizes the new FICON Bridge feature of the 9032 Model 5 ESCON Director to translate between FICON and ESCON protocols. This implementation can help allow a customer to protect his investment in ESCON control units, ESCON Directors, and other infrastructure while migrating to the new channels. Native FICON channels with FICON directors and FICON control units are the strategic direction for the I/O subsystems of future S/390 processors.

The initial FICON channel offering is depicted in the diagram at the right. It shows how a single FICON channel connects to multiple ESCON links through the ESCON Director. The number of ESCON links that can be connected to a single FICON channel is limited to 255 although the ESCON Director with 248 ports limits this to 240 (8 are used for the FICON card). Also note that only 8 of these links can be concurrently active through one FICON channel as there are only 8 protocol engines in the director bridge card.

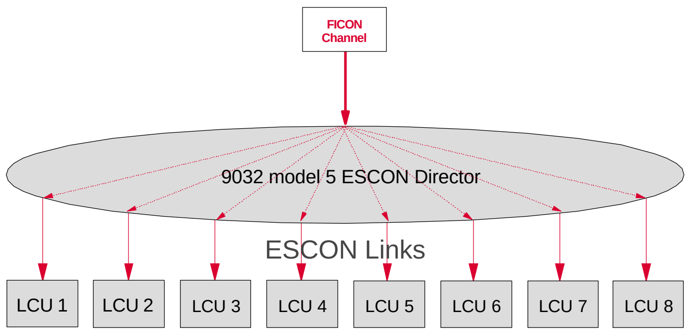

Technical Detail of FICON Bridge

The new FICON Bridge channel adds some complexity relative to the ESCON channel. The diagram below shows the switch matrix along with the ESCON ports (connected to the control units on the right) contained in the 9032 model 5 ESCON Director. On the left hand side of the pictured switch matrix are the 8 protocol engines per FICON port. Inside of the FICON channel is the microprocessor and the connection to the Fibre Channel link (labeled FC PCI). These components along with a PCI bus (not depicted) in the channel add to the infrastructure between the host and the control units.

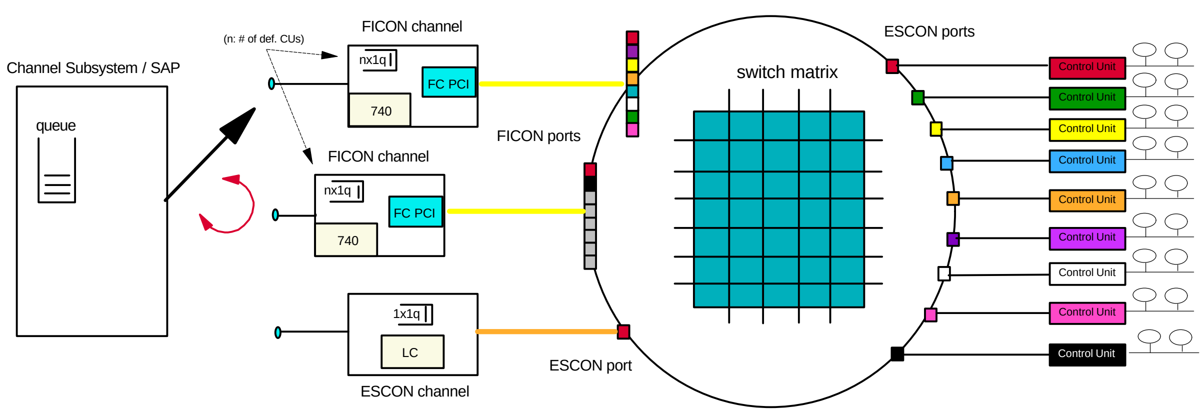

Benefits of FICON

The new FICON channel has many benefits over the traditional ESCON and Parallel channels. These benefits include:

- **Increased Distance**: With traditional parallel channels, the distance from a host processor (CEC) to the Control Units (CUs) was limited to 400 feet. This distance limitation was eliminated with the introduction of the ESCON channels in 1990 where the distance was up to 3 km without switches or repeaters and up to 43 km with switches or repeaters. The new FICON channels will now support distances up to **20 km without repeaters and up to 100 km with repeaters** without significant data rate droop.
- **Channel Aggregation**: With the initial FICON implementation that includes a bridge card in the 9032 Model 5 ESCON Director to convert to the ESCON protocol, a single FICON channel can replace multiple ESCON channels. This can effectively **provide relief for the 256 channel limitation** that is currently associated with IBM S/390 systems.
- **Reduced Infrastructure**: The infrastructure costs will be reduced by the aggregation of multiple ESCON channel fibers onto a single FICON fiber. This will be most significant when connecting host systems to remote locations either across the street, across town, or even further.
- **Increased Addressability**: The current addressing limitations for ESCON are 1024 Unit Addresses (UA) per channel. The addressability is increased to **16,384 with the FICON Channels**. This increase can help customers that are currently constrained either by the number of CUs or number of devices attached to their systems.
- **Increased Data Transfer Rate**: The data transfer rate across the link is limited to 4.5 MB/sec on a Parallel channel and 17 MB/sec on an ESCON channel. The FICON link data rate during data transfer will be 100 MB/sec; one must however keep in mind that the bandwidth of this link will be limited to the capability of the control units or the ESCON link in the case of the FICON Bridge implementation. Aggregate capacity is currently limited to 60 MB/sec.

Performance Parameters

The three main focus items from an I/O performance perspective are the I/O throughput (I/O operations per second), the bandwidth (MB/sec), and the response time (milliseconds to complete an I/O operation).

I/O Throughput

The I/O throughput of an ESCON channel today is a function of the type of control unit that is connected along with the type of channel program and number of bytes transferred. If we assume a state-of-the-art control unit, 100% read hits, and a data transfer of 4 KB (typical of many applications) an ESCON channel can execute on the order of 1000 I/O operations per second. One must keep in mind this is at 100% channel processor utilization and is without regard to any response time constraints. A more typical customer environment would execute on the order of less than 250-500 I/O operations per second (depending upon the control unit type) in order to keep the response time at a reasonable level.

The I/O throughput of a FICON Bridge channel, again assuming state-of-the-art control units and 100% read hits with 4 KB data transfers, is about 4000 I/O operations per second. Again, one must keep in mind this is at 100% utilization in an engineering-tuned environment and is without regard to response time. The recommendation to customers should be to drive these channels to a throughput rate of less than 1500-2400 operations per second (depending upon the control unit type). It must be stressed that this capability is highly dependent upon the control unit, block size, and cache hit ratios.

In order to attain this FICON throughput, more than one ESCON link must be connected through the bridge card. Be aware that the throughput of a single ESCON link connected through the bridge card will have a lower throughput than if that same ESCON link were connected directly to the control unit. The decrease in throughput is proportional to increases in response time (to be discussed later) experienced with FICON channels in a bridge environment.

As with ESCON channels, the throughput of a FICON channel is reduced in a multi-system environment with destination ports on the switch shared between systems. This is due to the fact that an operation being executed by one of the systems does not know if any other system is utilizing the shared destination port without sending the I/O operation to the switch. The additional overhead to send the operation to the switch and then handle the busy response causes the response time to increase and the throughput to be reduced. In these multi-system environments, the destination port throughput is equivalent in both the ESCON and FICON channel configurations.

Bandwidth

The bandwidth of channels are generally stressed when large blocks are transferred. Transfers from control units via ESCON channels are generally limited to about 10-13 MB/sec effective bandwidth even though the link data rate during data transfer is 17 MB/sec. This is due to the protocol, channel overheads, and control unit overheads associated with a given I/O operation.

The link data rate of a FICON channel during data transfer is 100 MB/sec but the effective data rate across the channel is approximately 60 MB/sec. This is due to a limitation of the current PCI bus implementation that is internal to the FICON channel; the utilization of this bus will be reported in RMF along with the utilization of the microprocessor in the FICON channel. In order to attain this bandwidth, operations must be concurrently active to multiple control units; operations to a single control unit will be limited to the bandwidth of a single ESCON link.

The recommendation is that the bandwidth through the FICON channel be limited to 40 MB/sec in order to keep the response time at reasonable levels. The following chart shows the response time for large block transfers across a single FICON channel compared to 1, 4, and 8 ESCON channels.

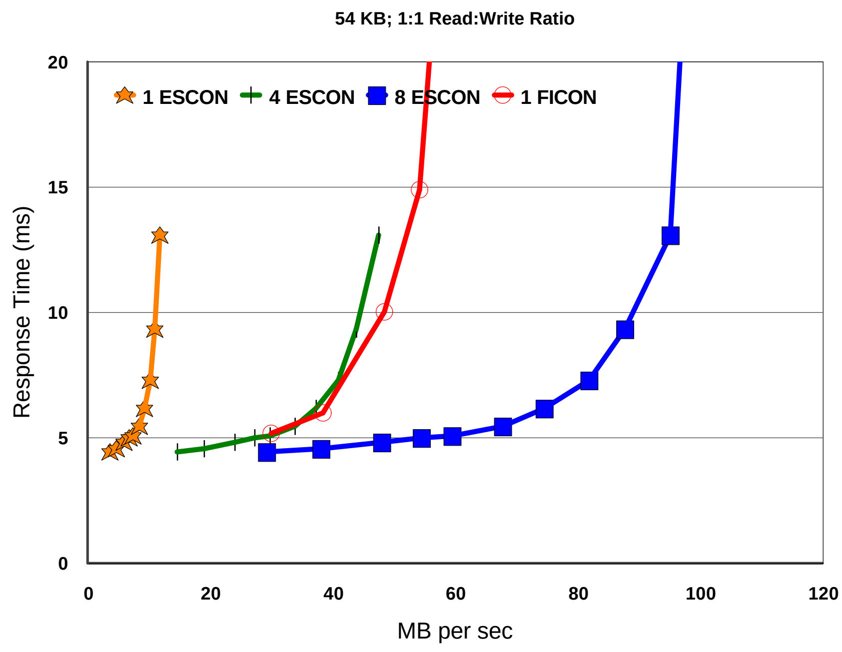

Response Time

The final performance parameter is the average response time for I/O operations. Due to the additional hardware path between the host and control unit (remember that there is still also ESCON link in the path), the time required to execute an I/O operation will always be more with the FICON Bridge case compared to native ESCON channels at short distances.

Though it is dependent upon the channel program and the control unit type, this response time elongation relative to an ESCON channel is about 400-500 microseconds in service time; service time is the single-thread response time, that is there are no queuing delays included. The following breakdown of service time was used in describing ESCON channel service times and will also be applicable to the FICON case, there are basically four components to the service time of an individual I/O operation:

- **Channel/CU Overheads**: This is the overhead associated with the channel (ESCON or FICON) and the CU in executing the channel program.
- **Data Transfer Time**: This is the actual time spent transferring data at the link data rate.
- **Propagation Delay**: This is the time a frame takes in propagating across the fiber from the channel to the control unit or vice versa.
- **Device Seek/Latency**: This time is associated only when cache misses are encountered in the control units and the actual devices must be accessed.

The following bar chart shows these components for a 4 KB read hit operation at 0km; there is no propagation delay at 0 distance and there is no device seek or latency time on cache hit operations. Since there is an ESCON link in the path between the director and the control unit, the data transfer portion is limited to the ESCON link speed and is thus equivalent for both the ESCON channel and FICON Bridge cases.

Service Time for 4KB Read (Cache Hit) at 0km

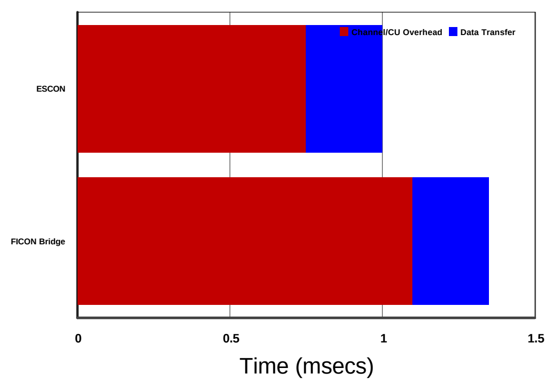

Effect of Distance on Service Time

As distance is increased in both an ESCON and a FICON configuration, the service time for the I/O operations on the FICON channel becomes less than that of the ESCON channel. An ESCON channel at distances greater than 9 km suffer from a data rate droop on the link; this reduces the effective data rate to something less than 10-13 MB/sec by elongating the data transfer portion of service time. FICON channels experience a similar data rate droop but not until much greater distances. In fact FICON channels can maintain the 60 MB/sec effective data rate up to 100 km. There are 2 main reasons for this difference between ESCON and FICON channels: increased buffer sizes in the FICON channel and improved protocols with fewer interlocks and round trips across the link.

The following two diagrams show the corresponding service time but with increased distance (10 and 20 km are depicted). One can see that the service time for the FICON Bridge channel case equals that of the ESCON channel at just 10 km.

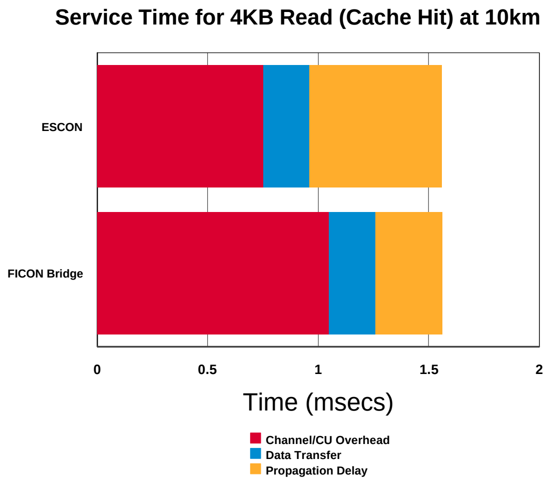

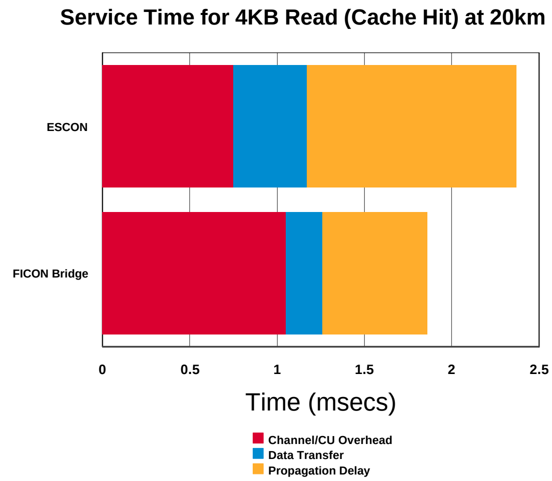

Effect of Block size on Service Time

Although the absolute elongation in service time of an I/O Operation on a FICON channel relative to an ESCON channel is independent of the amount of data transferred, the percentage elongation is less as the amount of data transferred is increased. That is for a 4 KB read operation, the service time increases from 1.0 ms for ESCON channel to about 1.4 ms for FICON (an increase of 40%). For a 32 KB read operation, the service time again increases by 0.4 ms from 2.4 ms to 2.8 ms, which is an increase of only 16%.

Since a larger portion of the overall service time is in data transfer, as distance is added into a configuration that utilizes larger block sizes the service time for I/O operations on FICON channels will be significantly less than that of ESCON channels; again this is due to the data rate droop suffered above 9 km with an ESCON channel.

The following two diagrams show the service time if a 32 KB data transfer at 0 km and at 20 km. Note that the service time for the FICON Bridge case is significantly reduced compared to ESCON channel at large distances.

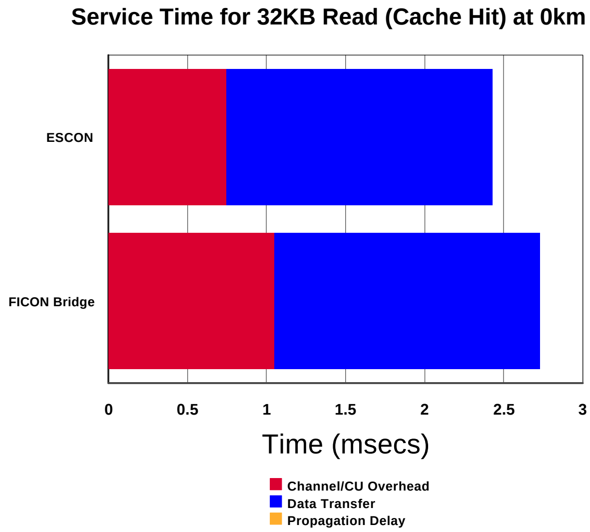

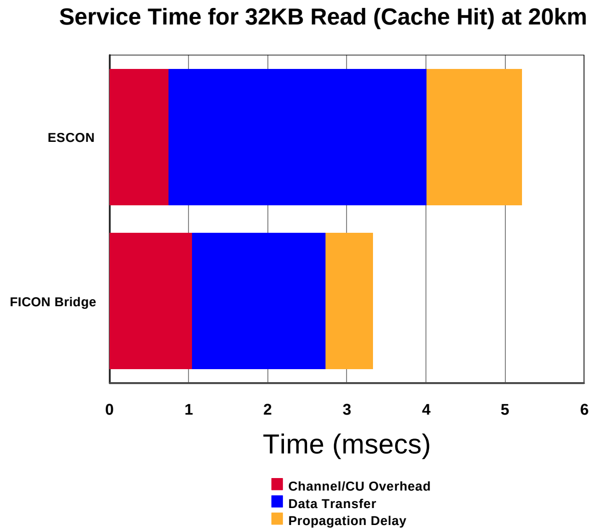

Effect of Multiplexing on Response Time

Thus far we have looked at the effects of various parameters on the service time of an individual I/O operation in order to describe the basic differences between ESCON and FICON channels. One of the most important factors in the I/O response time of a FICON channel is caused by the channel concurrently executing multiple I/O operations (up to 8 in the Bridge configuration). As the load on the channel increases, this multiplexing can increase the response time dramatically above that of an ESCON channel. This is due mainly to the fact that the channel is switching tasks among the multiple I/O operations.

The following chart compares the response time versus throughput (I/O operations per second) of a single ESCON channel, 4 ESCON channels, 8 ESCON channels, and a single FICON channel. From a system level, a FICON channel is still a single CHPID as is an ESCON channel. Physically a single FICON channel card replaces an ESCON channel card that supports 4 ESCON channels.

This chart shows two of the performance parameters. First if we look at the horizontal axis, we can see the I/O rate at which the response time increases dramatically for each of the cases. For the single ESCON channel, the recommended I/O throughput rate would be about 100-300 I/Os per second (typical for customer environments) in order to keep the response time at a reasonable level (again assuming state-of-the-art control units). For the 4 ESCON and 8 ESCON channel cases the throughput rates should be limited to about 1200 to 2400 I/O operations per second respectively. The FICON channel can support up to about 2400 operations per second in order to stay to the left of the knee of the response time curve.

By looking at the vertical axis we can see the total impact of all the factors that contribute to a response time elongation with a FICON channel. Even at low I/O rates the FICON channel response time is higher than that of the ESCON channel. This is always the case unless the configuration is operating at an extended distance.

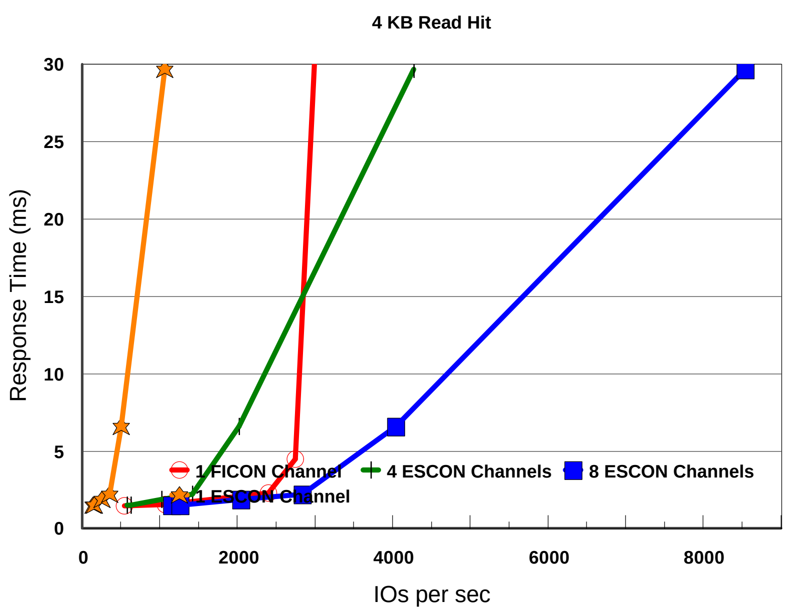

In real customer environments it is not realistic to assume that 100% cache hits can be attained. The following chart shows the response time for 4KB I/O operations versus the total I/O rate for a more typical customer environment. The configuration is 8 disk control units with 4 paths to each control unit and 80% cache hit rates. The response time for both a pure FICON (4 FICON channels) and a pure ESCON (32 ESCON channels) configuration are plotted. The 4 FICON channels can support up to 9-10 thousand I/O operations which translates to 2250-2500 I/O operations per second per FICON channel.

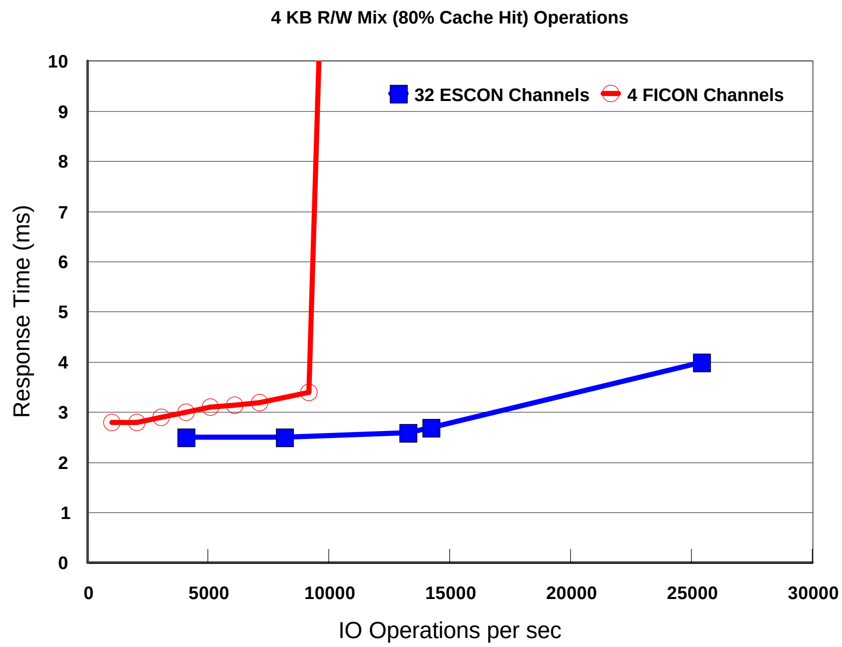

Reducing Response Time Elongation

The response time elongation can be reduced by incorporating a mix of ESCON and FICON channels to a given set of control units. By configuring the I/O subsystem in this manner a customer can get many of the benefits of FICON channels without experiencing the full impact of the increased response time.

By mixing the channels in this manner a customer can still get a significant reduction in the number of CHPIDs required. The following pictures show how in a configuration of 8 control units with 4 paths to each control unit, one could replace sets of 8 channels with a single FICON channel. In order to reduce the number of CHPIDs required. The first configuration shows 32 ESCON channels connected to the 8 control units. The second configuration would replace 8 of the ESCON channels (one to each control unit) with a single FICON channel; this effectively saves 7 CHPIDs for the channel constrained customer. The next two configurations incorporate a second and then a third FICON channel thus reducing the number of required CHPIDs down to 18 and 11 respectively compared to the original 32. The final configuration is the ultimate goal of FICON and can be used by customers that can tolerate the increase in response time associated with the FICON channel. One must keep in mind the plugging requirements associated with the FICON channel on G5 and G6 processors; physically the FICON channel card that supports 1 FICON channel takes up the same space as an ESCON card that supports 4 ESCON channels.

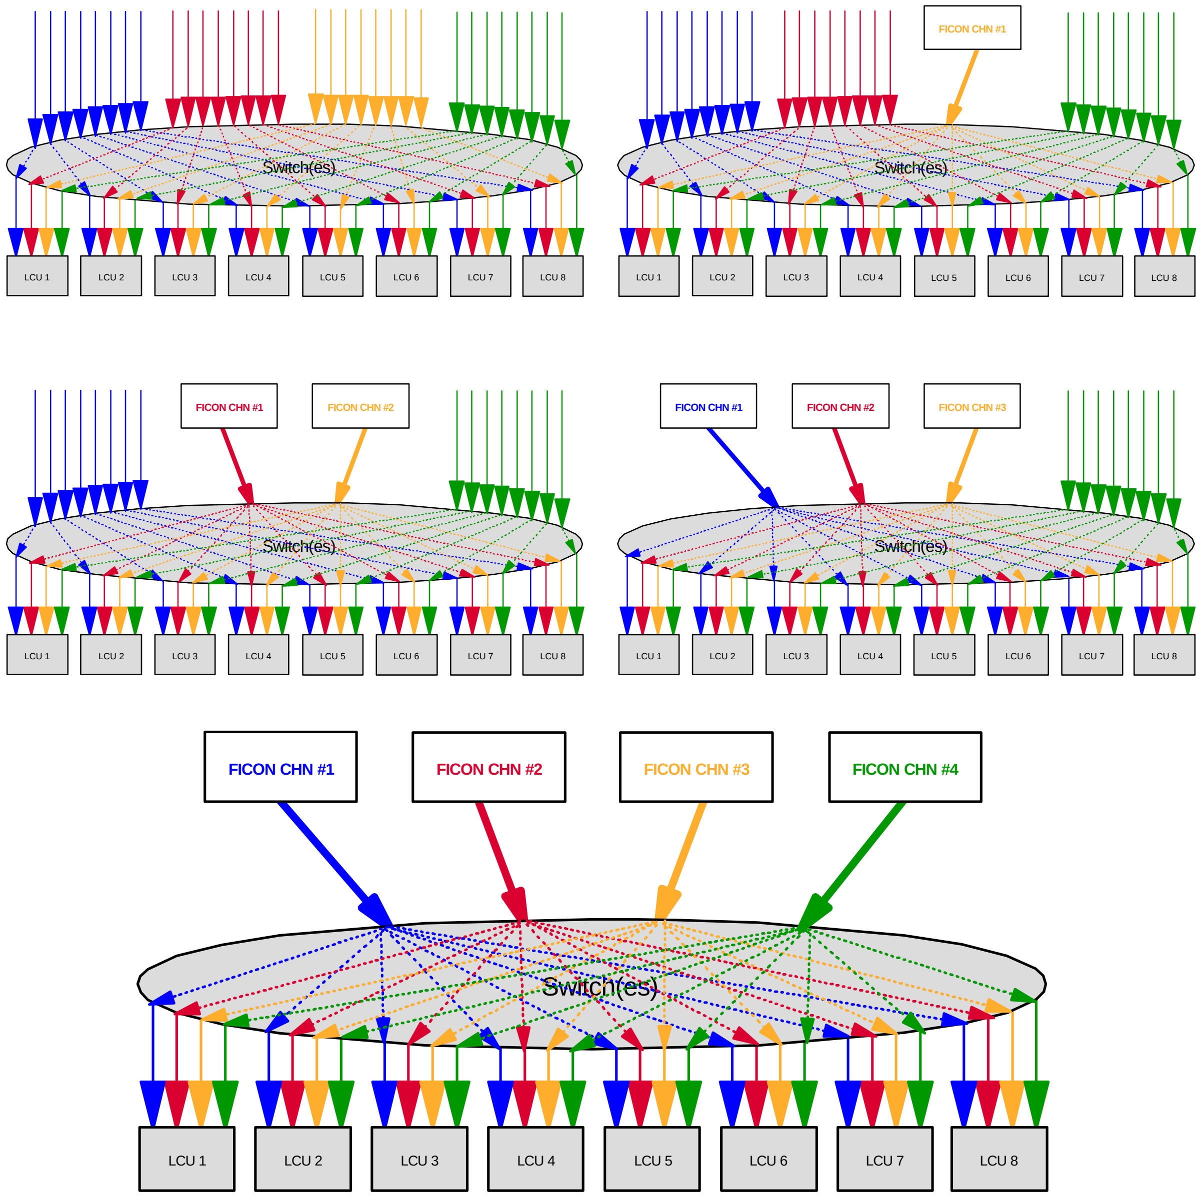

This mixing can be carried out at any level and by aggregating any number of ESCON channels onto a single FICON channel as long as the total aggregated ESCON utilization adheres to the guidelines (see Appendix A).

The effect on the response time of this mixing is represented in the chart below. The graph shows the response time verus the total I/O rate as the channel configuration is changed from 32 ESCON channels to 16 ESCON/2 FICON channels to 4 FICON channels; these channels are again connected to 8 disk control units with 4 paths to each control unit.

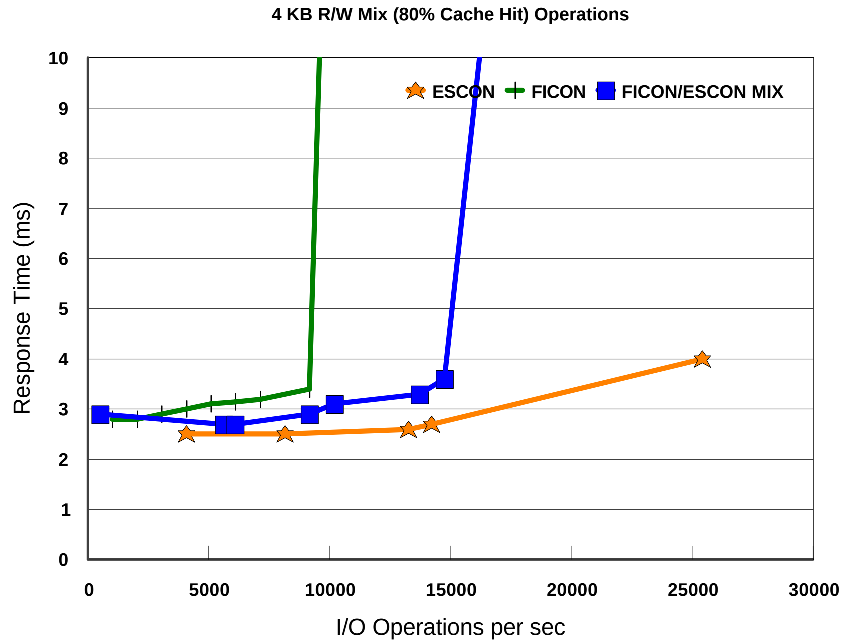

This mixing of channels is quite natural since many customer environments utilize multiple ESCON channels to a control unit in order to handle the peak throughput experienced at some time during day. This spike in I/O activity generally moves across the various control units throughout the day and installing FICON channels in this manner will allow a customer to balance the work load while using less CHPIDs than would be required with ESCON channels.

Impacts to Customer Environments

Thus far we have described the impact that FICON will have on individual I/O operations but the real question that must be addressed is what impact using FICON channels with bridge feature in place of ESCON channels will have at the system level. In order to analyze this, measurements on various customer representative workloads were taken, first with all ESCON channels and then with FICON channels replacing some of the ESCON channels .

IMS

The first environment measured was an IMS environment which is representative of an On-Line Transaction Processing (OLTP). Various channel configuration mixes were measured including:

1. All ESCON (64 ESCON channels to 16 control units),
2. FICON/ESCON mix to 2/3 of the database (2 FICON channels & 16 ESCON channels to 8 control units),
3. FICON/ESCON mix to all of the database (4 FICON channels & 32 ESCON channels to 16 control units).

In this environment, an average of 125% ESCON channel utilization (see Appendix C for guidelines) was aggregated per FICON channel resulting in a FICON channel utilization of 20%. Relative to the all ESCON configuration, when FICON channels were added the I/O response time increased 0.4-1.0ms (7-17%). This I/O response time increase translated into an increase in the transaction response time of 9-30%; this response time is on the order of 0.1-0.3 seconds and thus this increase will be imperceptible to a user. The throughput of the system was reduced by only about 1%-3%.

It should be noted that the changes in response time and in ITR will vary from customer to customer depending upon the load on the system.

TSO

The second environment measured was TSO which is representative of the work done by a typical OS/390 end-user community interactively developing and testing programs. In this environment, the channel configurations measured were:

1. All ESCON (64 ESCON channels to 16 control units)
2. FICON/ESCON mix to all of the database (4 FICON channels & 32 ESCON channels to 16 control units).

In this environment, an average of 175% ESCON channel utilization (see Appendix C for guidelines) was aggregated per FICON channel resulting in a FICON channel utilization of 32%. Relative to the all ESCON configuration, when FICON channels were added the I/O response time increased 0.3ms (12%). This I/O response time increase translated into an increase in the transaction response time of only 3ms (6% on a base response time of 47 ms) and the throughput of the system actually increased by 0.4%.

Batch Environment (CBW2/CB84)

Two of the LSPR workloads that are representative of batch environments were also measured; these two workloads are CBW2 and CB84. The workloads were first run on systems with only ESCON channels in the configuration and then various amounts of FICON channels incorporated into the configuration; with up to 150% (for CBW2) and 200% (for CB84) ESCON channel utilization aggregated per FICON channel resulting in a 20% FICON channel utilization. Relative to the ESCON only case, the elapsed time to complete the equivalent jobs with FICON channels in the configuration varied from an improvement of 1.7% to a degradation of 3.0%. Based upon this customers should expect to see a batch window change of +/- 2-3% depending upon their environment.

Volume Dumps

The final system level measurement is the elongation of time to execute volume dumps from disk to tape. This experiment was carried out by varying the number of tape controllers and disk controllers connected per FICON channel; there was no mixing of tape and disk controllers on a single FICON channel. With 1 to 3 controllers connected per FICON channel, the elapsed time to complete the volume dumps was increased by 2-8%. This elongation grows to 15-30% as the number of controllers is increased above 3.

This elongation in carrying out volume dumps will be seen by all customers that use FICON and should be managed with caution especially with those customers that have a given time window during which the backups must be completed. By managing the configuration, the elongation that is experienced can be controlled to less than 10%.

Summary

Several of the benefits of the new FICON channel such as the channel aggregation and increased addressability are presently required by many customers. There is, however, a tradeoff that must be made in order to get these and the other benefits of the FICON channel. This tradeoff comes in the form of an elongated response time for individual I/O operations at data center distances; this elongation can be minimized by carefully managing the incorporation of FICON channels into existing configurations.

Even though the response time elongation can be quite significant in some cases (up to 40-50% for some cache hit operations), the effect is not nearly so dramatic at the system level due to the fact that the I/O operations are only a fraction of the work being carried out by the system. Care should be taken when using FICON channels to connect to control units that are heavily utilized during batch windows. In this environment, the total elapsed time to execute batch workloads, including disk volume dumps, could potentially be elongated.

The following table summarizes many of the key performance characteristics for FICON Bridge channels.

	FICON	ESCON
Link data rate	100 MB/sec	17 MB/sec
Effective date rate	55-60 MB/sec	10-13 MB/sec
Maximum Throughput (CU Dependent)	3600 per sec	1000 per sec
Throughput (4KB I/O operations) at reasonable response time (CU Dependent)	1500-2400 per sec	250-500 per sec
Distance (no repeaters)	10km;20km RPQ	3 km
Distance w/o significant data rate degradation	100 km	9 km

This chart summarizes the key decision criteria to consider during an ESCON to FICON migration.

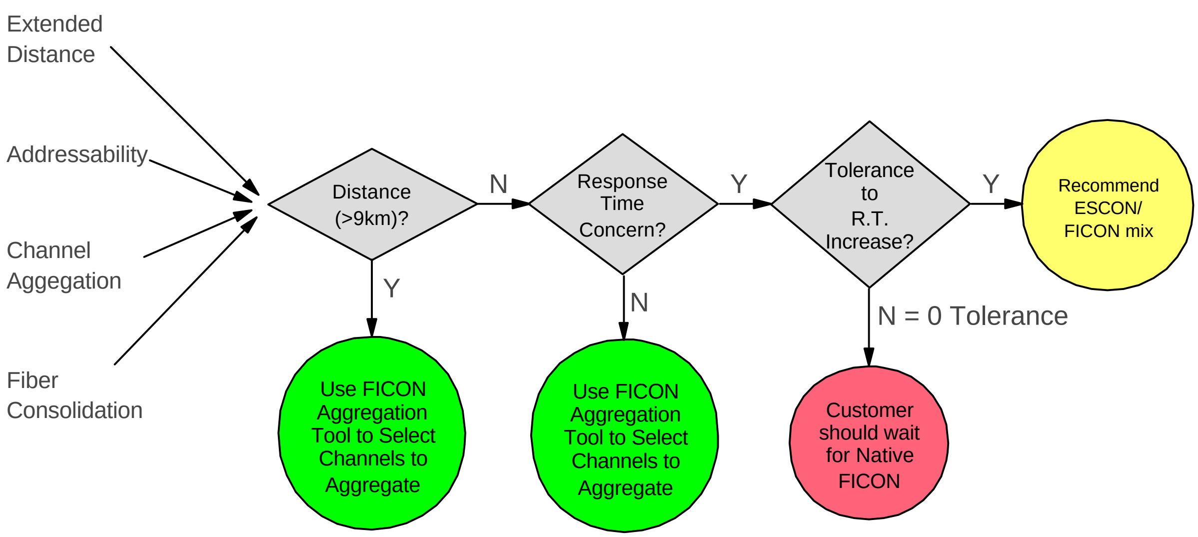

Appendix A: ESCON-to-FICON Bridge Channel Aggregation Guidelines

For configuration gidelines please use the 'Hints and Tips' section under the following URL:

http://www.ibm.com/support/techdocs

or refer to the

FICON Planning Redbook (SG24-5456).

Appendix B: Trademark Information

IBM Corporation
Marketing Communications, Enterprise Systems Group
Route 100
Somers, NY 10589

Printed in the United States of America, 04-00

This publication was produced in the United States.
IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the products or services available in your area.

You can find additional information via IBM's World Wide Web server at .ibm.com.

IBM Hardware products are manufactured from new parts or new and serviceable used parts. Regardless, our warranty terms apply.

Actual performance and environmental costs will vary depending on individual customer configurations and conditions.

By "Year 2000 ready", IBM means the capability of a product, when used in accordance with its associated documentation, to correctly process, provide and/or receive date data within and between the 20th and 21st centuries, provided that all products (for example, hardware, software and firmware) used with the product properly exchange date data with it. IBM's statements concerning the Year 2000 readiness of its products are Year 2000 Readiness Disclosures under the Year 2000 Readiness Information and Disclosure Act of 1998.

® ESCON and FICON are registered trademarks of the International Business Machines Corporation.