SpaceFibre Virtual Channels and Flow-control

Session: Standardisation

Long Paper

Clifford Kimmery

Honeywell International Space Electronic Systems

13350 US Highway 19 N, Clearwater, FL, 33764

E-mail: clifford.kimmery@honeywell.com

Abstract

The SpaceFibre CODEC Functional Specification [1] defines a mechanism for identifying virtual channels within a SpaceFibre link and performing flow-control independently for each virtual channel. The SpaceFibre literature [2] [3] envisions the use of virtual channels for multiplexing multiple traffic flows and for providing quality of service features, but does not address the details necessary for standardized implementations.

The virtual channel concept as described in the SpaceFibre literature implements one buffer per virtual channel at each end of the SpaceFibre link. The buffer must be large enough to contain at least one SpaceFibre frame. The link receiver issues a flow-control ordered-set for any virtual channel when the buffer for that virtual channel has room for at least one frame. The link transmitter sends a waiting data frame on a virtual channel if the number of flow-control ordered-sets received for that virtual channel is greater than zero.

The ramifications of implementing virtual channels and flow-control as defined by the SpaceFibre CODEC Functional Specification and described in the SpaceFibre literature are explored. The complexity and link efficiency of various SpaceFibre virtual channel and flow-control implementation choices are evaluated and potential alternatives suggested. The topics addressed include: coupling of SpaceFibre link flow-control to the individual virtual channels, synchronization of the number of active virtual channels between the SpaceFibre link transmitter and receiver, synchronization of the maximum virtual channel frame buffer capacity between the SpaceFibre link transmitter and receiver and the effects of various Quality-of-Service factors such as bounded latency and allocated bandwidth.

Virtual Channel Background

The classical use of virtual channels is to perform bandwidth allocation on a network link as a primary mechanism for providing network QoS features. Some standard packet protocols provide virtual channel support under a different name in the form of traffic flow priority levels (RapidIO [4]) or virtual lanes (Infiniband [5]). Others (PCI Express [6] is an example) use the virtual channel term as a synonym for traffic flow priority levels. Connection-oriented protocols (Fibre Channel [7], ATM [8], SDH [9] and SONET [10] are examples) define virtual channels for statically allocating bandwidth within a link or connection. In all cases, the virtual channels gain access to the physical link based on some arbitration mechanism.

The arbitration mechanisms used by the standard protocols range from fixed time-sequencing (typically to provide guaranteed bandwidth in connection-oriented networks) and fixed packet priority (most packet protocols) to more complex schemes based on the dynamic behavior of traffic flows.

Virtual channels allow a variety of potential SpaceFibre Quality of Service (QoS) features based on allocation of link bandwidth. By establishing standard bandwidth allocation mechanisms for implementation by endpoints and routers, SpaceFibre can support the QoS needs of many different networking applications. The primary purpose of introducing virtual channels in SpaceFibre is to allow specific traffic flows to progress in the presence of congestion on other traffic flows.

SpaceFibre Virtual Channels

SpaceFibre virtual channels are part of the SpaceFibre equivalent of the SpaceWire Exchange level as shown in Table 1 – SpaceFibre/SpaceWire Network Model Relationships.

<table>
<tr><th></th><th>SpaceFibre Level</th><th>SpaceWire Level</th></tr>
<tr><td></td><td>Application</td><td>Application</td></tr>
<tr><td></td><td>Network</td><td>Network</td></tr>
<tr><td></td><td>Packet</td><td>Packet</td></tr>
<tr><td></td><td>Flow-control</td><td>Exchange</td></tr>
<tr><td></td><td>Virtual Channel</td><td>Not Used</td></tr>
<tr><td rowspan="4">SpaceFibre CODEC</td><td>Framing</td><td>Not Used</td></tr>
<tr><td>Link Control</td><td>Exchange</td></tr>
<tr><td>Encoding</td><td>Character</td></tr>
<tr><td>Serialisation</td><td>Character</td></tr>
<tr><td></td><td>Signal</td><td>Signal</td></tr>
<tr><td></td><td>Physical</td><td>Physical</td></tr>
</table>

Table 1 – SpaceFibre/SpaceWire Network Model Relationships

A primary goal of SpaceFibre is to provide transparent flow of SpaceWire packets over SpaceWire and SpaceFibre networks integrated in arbitrary configurations. Because of the significantly greater bandwidth provided by SpaceFibre links, the SpaceFibre CODEC Functional Specification defines support for a maximum of 256 virtual channels to be simultaneously active over one physical SpaceFibre link. Each virtual channel can be viewed as a virtual SpaceWire link sharing a single SpaceFibre link with up to 255 other virtual SpaceWire links.

Each SpaceFibre frame header contains the identity of the associated virtual channel. The data contained within the sequence of frames associated with a specific virtual channel identifier are treated by the SpaceFibre routers and endpoints as equivalent to a sequence of SpaceWire packets.

When considered in the context of a typical SpaceWire network and given the historical preference by the SpaceWire community for simple, low-complexity network implementations, the SpaceFibre virtual channel support appears to have significant capacity for future growth. A maximum-capacity SpaceFibre endpoint can bridge one SpaceWire link to one of the 256 SpaceFibre virtual channels for a maximum of 256 SpaceWire links over one SpaceFibre link. The complexity impact of virtual channel support on SpaceFibre router implementations can be significant as shown in Table 2 – SpaceFibre Router Complexity Scaling (assumes one maximum-size frame buffer per virtual channel per direction).

Ports	Virtual Channels/Port	Equivalent SpaceWire Ports	Minimum Buffer Memory (bytes)	Comment
4	4	16	32,896	Small
4	64	256	526,336	
32	4	128	263,168	
32	256	8,192	16,842,752	Maximum

Table 2 – SpaceFibre Router Complexity Scaling

SpaceFibre Flow Control

The SpaceWire protocol uses a credit-based flow-control mechanism between the link receiver and link transmitter. The link receiver issues credit tokens (flow-control characters) indicating the availability of eight characters of buffer space at the receiver. The link transmitter accumulates the credit tokens to maintain a running total of the available buffer space and transmits data characters at will if the total is greater than zero.

The SpaceFibre CODEC Functional Specification extends the SpaceWire credit-based flow-control mechanism for use with SpaceFibre virtual channels. Each credit token (flow-control ordered-set) represents the availability of buffer space at the receiver for one maximum-length SpaceFibre frame on the associated virtual channel. The SpaceFibre link transmitter accumulates the credit tokens to maintain a running total of the number of available frame buffers on the associated virtual channel and transmits frames at will if the total is greater than zero.

Virtual Channel Count Interaction with Frame Buffer Count

Since SpaceFibre flow-control uses frame granularity, the SpaceFibre link receiver cannot issue a flow-control ordered-set until a frame buffer is empty. The latency between receiver recognition of the empty frame buffer, transmission of the flow-control ordered-set, transmitter reception of the flow-control ordered-set and transmission of the next frame can be significant. Provision of sufficient frame buffers for each virtual channel to sustain the maximum throughput is ideal.

Since the Spacefibre link is a time-shared resource, the number of frame buffers per virtual channel needed to maintain a high per-channel throughput is inversely proportional to the number of virtual channels actively sharing the link. The throughput effects of flow-control latency for a virtual channel can be partially hidden by the availability of the link to that channel (if the virtual channel cannot transmit because the link is in use by another channel, the flow-control latency may be invisible).

Virtual channels with a greater share of the link bandwidth need a greater number of frame buffers since the flow-control latency is proportionally more visible. As an example, the sustained throughput capability of a minimal SpaceFibre link receiver (one virtual channel and one frame buffer) is inversely dependent on the link propagation delay.

FACTORS DRIVING SPACEFIBRE BUFFER MEMORY CAPACITY

There are three factors that drive the buffer memory capacity of a SpaceFibre link receiver: the size of each frame buffer, the number of virtual channels that can be simultaneously active on the link and the number of frame buffers used for each virtual channel.

1. In the absence of a frame-size coordination mechanism, the frame buffers must be sized for the maximum-capacity frame (255 32-bit data words).

2. The number of simultaneous virtual channels supported is an implementation decision, but greater numbers of virtual channels offer more application flexibility.

3. The number of frame buffers per virtual channel is an implementation decision that can significantly impact sustained virtual channel throughput.

DECOUPLING SPACEFIBRE LINK FLOW-CONTROL FROM VIRTUAL CHANNELS

While the flow-control mechanism defined in the SpaceFibre CODEC Functional Specification is easy to understand, it forces the SpaceFibre link receiver to allocate buffer space to each virtual channel whether that virtual channel is in use or not. As shown in Table 2 – SpaceFibre Router Complexity Scaling, SpaceFibre link receiver buffer capacity must increase linearly with the number of virtual channels on a port.

The correlation between buffer memory capacity and the number of virtual channels could be minimized by associating SpaceFibre flow-control with the framing level of the link rather than with each virtual channel. RapidIO [4] defines a transmitter-controlled flow-control mechanism that imposes responsibility for link receiver packet buffer management on the link transmitter. The link receiver implements a pool of packet buffers available to all packet priorities (the RapidIO equivalent of virtual channels) and regularly reports the number of available packet buffers to the link transmitter. The link transmitter implements the buffer allocation mechanism to guarantee that low-priority packets cannot block higher-priority packets by consuming all available receiver packet buffers.

Applying flow-control at the virtual channel level is conceptually simple and limits the decision-making complexity of flow-control management. The percentage of dedicated chip resources can be significant since the virtual channel flow-control mechanism requires correspondingly more frame buffer memory at the link receiver (because of the need for frame buffers to be dedicated to each virtual channel) and correspondingly more flow-control credit counters at the transmitter.

In contrast, applying flow-control at the framing level would allow a pool of link receiver frame buffers to be shared by all active virtual channels and a single link transmitter flow-control credit counter. The size of the frame buffer pool becomes an implementation decision that doesn’t limit the number of virtual channels that can be simultaneously active (overall link throughput would still be affected by the total number of frame buffers and the flow-control latency). The link transmitter must implement an allocation mechanism that manages the available link receiver frame buffers according to the desired QoS.

SPACEFIBRE QUALITY-OF-SERVICE OPTIONS

While the SpaceFibre literature [1][2][3] assigns Quality-of-Service support to the Virtual Channel and Flow Control levels, it leaves the goals and details undefined. Without guidance, speculation on the extent of SpaceFibre QoS support is appropriate.

Since some traditional QoS behaviors associated with reliable delivery are not associated with virtual channels and flow-control, we can assume that such behaviors are to be performed by either the SpaceFibre CODEC or at the Application level. Similarly, the Application level is appropriate for QoS behaviors designed to aid in alleviating network congestion (such as limiting packet ingress rates at the data source and providing end-to-end flow-control mechanisms).

The data-frame-based nature of SpaceFibre virtual channels suggest the feasibility of some form of guaranteed bandwidth/latency QoS similar to SDH [9] or SONET [10]. These connection-oriented network protocols use time-scheduled data frames partitioned to allocate link bandwidth. Because of their roots in telephony, the SDH/SONET time schedules are based on a 125 millisecond repeat interval (corresponding to the 8 kHz sampling rate of typical voice communications). SDH/SONET solves the more complex aspects of time-schedule-based bandwidth allocation by restricting the scheduled time increments to power-of-two divisors of the 125 millisecond interval. Each channel is assigned to a time-slot sized to provide a percentage of link bandwidth no less than the maximum needed by the channel. As a result, the connection-oriented virtual channels provide constant bandwidth regardless of the utilization profile of the application.

Typical QoS behaviors for packet networks are based on priority packet delivery. Packet networks attempt to increase the utilization efficiency of the network infrastructure by dynamically allocating bandwidth on demand rather than reserving bandwidth based on peak needs. The primary side-effect of dynamic bandwidth allocation is network congestion resulting from contention for bandwidth. By introducing packet priority into the bandwidth allocation mechanism, packet networks can allow the higher-priority packets to proceed at the expense of lower-priority packets. The network congestion still exists, but is separated into QoS classes that experience different levels of congestion.

Since SpaceWire packet headers do not include a QoS field, SpaceFibre cannot provide QoS-based packet delivery in the same manner as other packet network protocols. The QoS level must be attached to the SpaceWire packet upon entry into the SpaceFibre network and discarded upon exit from the network. There are two straightforward ways to attach the QoS level to a SpaceWire packet: the first is by associating the QoS level with the virtual channel used and the second is to attach the packet QoS level as a field in the SpaceFibre data frame header. Associating the QoS level with the virtual channel (through network configuration) is a good fit for connection-oriented virtual channel implementations since the connection setup must be performed using network configuration in any case. Associating the QoS level with the virtual channel is unlikely to be acceptable for dynamic bandwidth allocation in any but the simplest SpaceFibre networks since the number of QoS levels supported is constrained by the number of virtual channels supported by the most limited SpaceFibre link. Attaching the QoS level as a field in the SpaceFibre data frame header allows any virtual channel supporting dynamic bandwidth allocation to convey packets of any QoS level with a minor impact on link bandwidth efficiency.

SpaceFibre could provide a combination of the guaranteed bandwidth QoS of connection-oriented networks and the priority-based QoS behaviors of packet networks in a manner similar to IEEE 1394 [11] by including time-schedule-based (also known as isochronous) bandwidth allocation as the underlying behavior for some virtual channels and allowing the remaining virtual channels to compete for access to the residue link bandwidth on a priority basis. Within a scheduling interval, each time-scheduled virtual channel has a fixed schedule position relative to the other time-scheduled virtual channels and the time-scheduled virtual channels have priority over all other virtual channels. The bandwidth allocation mechanism grants link access to each time-scheduled virtual channel in sequence and the current virtual channel transmits either a data frame or an idle frame (or nothing if some frame jitter is acceptable) based on data availability. After the sequence of time-scheduled virtual channels is completed, the other virtual channels compete for link access based on the priority of the waiting data.

A few of many possible dynamic bandwidth allocation arbitration methods are identified in Table 3 – Potential Arbitration Methods for Dynamic Bandwidth Allocation.

Arbitration Method	Description
Virtual Channel Fixed Priority	The priority of each virtual channel is fixed in hardware based on the channel number.
Virtual Channel Assigned Priority	The priority of each virtual channel is configurable by software.
Virtual Channel Rotating Priority	The priority of each virtual channel is increased by one priority level each arbitration cycle until the channel is granted access to the link. After gaining access to the link, the virtual channel is assigned the lowest priority. Each virtual channel has a different initial priority. The initial priority of the virtual channels is determined using either the Virtual Channel Fixed Priority method or the Virtual Channel Assigned Priority method.
Packet Priority	The priority of each virtual channel is established by the packet priority associated with the frame waiting to be transmitted on that virtual channel.

Table 3 – Potential Arbitration Methods for Dynamic Bandwidth Allocation

SYNCHRONIZING SPACEFIBRE LINK PARTNERS

Because different SpaceFibre router and endpoint implementations are likely to be based on conflicting goals, sufficient configurability is necessary to allow optimization of network operation. For example, a minimal SpaceFibre link receiver could be implemented using a small maximum frame size and the link transmitter must be configured to match. Another example is a link receiver implementation with the flexibility to reallocate the frame buffers to optimize either the number of virtual channels supported or the number of frame buffers per virtual channel. Similarly, a link receiver might be capable of increasing the number of frame buffers by reducing the maximum frame size.

Allowing SpaceFibre link receivers to implement a frame buffer size less than full-size (255 32-bit words) has ramifications beyond minimizing resource utilization or optimizing network operation. A SpaceFibre router receiving full-size frames from one endpoint and transmitting them to another endpoint incapable of accepting full-size frames would need to implement a form of frame segmenting to allow the data content of the larger frames to be transparently partitioned into smaller frames.

While software control of configurable link partner parameter values is consistent with the SpaceWire philosophy, a hardware-based protocol capable of communicating the link receiver maximum frame size to the link transmitter should be considered since a common understanding of maximum frame size is fundamental to link behavior. An alternative would be to require full-size frame buffers for virtual channel 0 (likely to be used during initial network configuration) after hardware reset with the ability for software adjustment.

Although knowledge of the number of virtual channels supported between link partners is not necessary for proper link operation, there are potential benefits when frame buffers are reassigned from unused virtual channels to increase the number of frame buffers available to in-use virtual channels. In the absence of that knowledge, the link receiver must issue flow-control ordered sets for each virtual channel it is capable of supporting. The link transmitter must ignore any flow-control ordered-sets associated with a virtual channel it doesn't support. Although the link transmitter will never transmit frames on an unsupported virtual channel, the link receiver can only determine that a virtual channel is supported/in-use when a frame is received.

SUMMARY

There are a number of concerns that the SpaceFibre community needs to address as the details of the Virtual Channel and Flow Control protocol levels are defined. Some derive from the traditional SpaceWire preference for low complexity and simple behavior. Others offer opportunities to impact the flexibility of SpaceFibre and the ability to optimize the performance of SpaceFibre networks.

Because SpaceFibre links operate at much greater data rates than SpaceWire links, the number of virtual channels necessary to fill the bandwidth of a SpaceFibre link with typical SpaceWire traffic is substantial (for example, a 3 Gbps SpaceFibre link can carry roughly 30 SpaceWire links each operating at 100 Mbps). The combination of high virtual channel count and the full-size data frame defined by the SpaceFibre creates a need for large frame buffer memory capacity in SpaceFibre link receivers. SpaceFibre routers are particularly affected because of the need to support multiple ports.

The SpaceFibre definition of flow-control on each virtual channel exacerbates the frame buffer memory capacity issue by forcing link receiver frame buffers to be reserved for use by inactive virtual channels. A flow-control method that is independent of the virtual channel implementation would allow the link receiver to use its complement of frame buffers for any virtual channel. By considering a transmitter-controlled flow-control mechanism similar to that defined by RapidIO, the community can make the link receiver frame buffer capacity an implementation decision.

SpaceFibre has the facilities to support a variety of QoS behaviors ranging from guaranteed bandwidth based on time-scheduled frame transmission to various forms of dynamic arbitration. The SpaceFibre community has the opportunity to define standard QoS behavior or to follow the SpaceWire precedent and leave QoS capabilities as implementation decisions.

The full-size SpaceFibre data frame can be an issue for resource-constrained implementations. Allowing the maximum size of data frame buffers to be less than full-size would be beneficial in such circumstances, but introduces the need for link transmitters to segment large frames to fit within the size established for the link.

[1] Steve Parkes, et. al., “SpaceFibre CODEC Functional Specification Draft A” , SpaceFibre Outline Specification.pdf, 31-Oct-2007, University of Dundee

[2] Martin Suess, et. al., “Future Focus: SpaceFibre”, 2006 MAPLD International Conference, 20-Sep-2006, http://klabs.org/mapld06/seminars/spacewire/presento/21.ppt

[3] SpaceFibre, Steve Parkes, et. al., 23-Feb- 2008, UoD-SpaceFibre.pdf

[4] RapidIO Trade Association, “RapidIO Rev. 2.0”, 03/2008, http://www.rapidio.org/specs/current/Rev2.0_stack2.zip

[5] InfiniBand Trade Association, “InfiniBand Architecture Release 1.0.a”, June 19, 2001, http://www.infinibandta.org/specs/

[6] PCI-SIG, “PCI Express Base Specification Revision 2.0”, December 20, 2006, http://www.pcisig.com/specifications/pciexpress/specifications

[7] T11/Project 1822-D/Rev 8.1, “Fibre Channel Switch Fabric - 5 (FC-SW-5)”, 07/21/2008, http://www.t11.org/ftp/t11/pub/fc/sw-5/08-412v0.pdf

[8] ITU-T I.150, “B-ISDN Asynchronous Transfer Mode Functional Characteristics”, 02/99, http://www.itu.int/rec/T-REC-I.150-199902-I/en

[9] ITU-T G.707, “Network Node Interface for the Synchronous Digital Hierarchy (SDH)”, 01/07, http://www.itu.int/rec/T-REC-G.707-200701-I/en

[10] ANSI T1.105-2001, “Synchronous Optical Network (SONET) - Basic Description including Multiplex Structure, Rates, and Formats”, May 2001

[11] IEEE Std 1394-1995, “IEEE Standard for a High Performance Serial Bus”, 1995, ISBN 1-5593-7583-3