

## STRUCTURE FOR NON-BLOCKING SERIAL ATTACHED SCSI INFRASTRUCTURE UTILIZING VIRTUAL PATHWAYS

### BACKGROUND

#### 1. Field of the Invention

The invention relates generally to serial attached SCSI (SAS) domains, and more specifically, relates to improving the performance of SAS domains.

#### 2. Discussion of Related Art

Serial Attached SCSI (SAS) is a popular protocol and medium for connecting the components of a networked storage/computing architecture. In general, SAS standards define a SAS domain as including one or more SAS initiators coupled with one or more SAS target devices through zero or more SAS expander devices. In general, a SAS initiator device is coupled with a SAS target device as a point-to-point connection. To enhance the flexibility in connecting a SAS initiator device to other SAS target devices, a SAS domain may include one or more SAS expanders to provide, in essence, a switched fabric for coupling any of the SAS initiators in a point-to-point connection to any of the SAS target devices of the SAS domain.

More generally SAS initiators and SAS target devices may each be referred to as SAS endpoints in that many SAS devices may act in the role of an initiator at times and may act in the role of a target at times. Regardless of the role of a device as a SAS initiator or as a SAS target, any SAS compliant endpoint may attempt to open a connection with another endpoint. It is common, for example, for a host bus adapter SAS initiator to open a connection to a SAS storage device. However, the storage device may later open a connection back to a requesting host adapter for purposes of returning requested read data. Both initiator and target devices in a connection may simply be referred to herein as a "SAS endpoint device". SAS endpoint devices and SAS expanders may be referred to herein as "SAS devices".

The SAS standards provide for a full duplex connection between two connected SAS endpoints. In principle, such a full duplex connection allows the opening endpoint to transmit frames of information (e.g., data frames) to the opened endpoint substantially in parallel with the opened endpoint transmitting information back to the opening endpoint. In accordance with these SAS specifications, such a full duplex connection allows for near full utilization of the point-to-point connection between the two SAS endpoints (some available bandwidth is always consumed by overhead exchanges associated with the protocols, and in practice it is rare to achieve full utilization due to imbalances in traffic going in either direction between the SAS endpoints). This and other related SAS specifications are well known to those of ordinary skill in the art and are also readily available at [www.t10.org](http://www.t10.org), such as SAS 2.0, revision 16, published 18 Apr. 2009 and incorporated by reference herein.

In practice, it is difficult in the SAS domain to ensure a predictable level of performance across all SAS endpoints in large topologies due to the blocking nature of open SAS connections and the variability in the duration of those connections. For example, during the SAS connection process, a SAS initiator transmits an Open Address Frame (OAF) towards the target. The OAF is propagated through the intermediate expander connections between the initiator and the target as long as the intermediate paths along the route are available. If any of the intermediate paths are not available, then the connection request is delayed while arbitration is performed at the intermediate expander con-

nections, and sometimes, an open reject message is returned to the initiator, blocking the connection attempt. This may result in unpredictable and/or chronically poor performance for certain endpoint devices in the topology. Thus, new connection requests that require inter-expander pathways that are in use may be delayed (rejected) until the established connection is closed. This is one driver of the unpredictability in typical SAS domains.

It is evident from the above discussion that an ongoing need exists to improve the performance of SAS domains.

### SUMMARY

The present invention addresses the above and other problems, thereby advancing the state of the useful arts, by providing structure for a non-blocking SAS architecture utilizing virtual connections between SAS devices. Virtual connections are implemented by encapsulating SAS frames and primitives within one or more routable virtual SAS cells. The use of virtual SAS cells allows for a connectionless virtual SAS to be implemented. This virtual SAS allows for connectionless routing, concurrent virtual pathways on a single link, etc., which improves the performance of SAS.

One embodiment is a SAS expander. The SAS expander includes a plurality of physical links (PHYs). The SAS expander further includes a Virtual Connection Manager (VCM) that is coupled with the plurality of PHYs. The VCM is adapted to exchange information over a plurality of concurrently established virtual pathways.

In another embodiment, the VCM is adapted to receive SAS frames, to segment the SAS frames into a plurality of virtual SAS cells, and to transmit the plurality of virtual SAS cells utilizing a plurality of established virtual pathways between a first PHY of the plurality of PHYs and at least two PHYs of the plurality of PHYs.

In another embodiment, the VCM is adapted to receive a plurality of virtual SAS cells comprising segmented SAS frames at a first PHY of the plurality of PHYs, to determine a virtual pathway specified by each of the plurality of virtual SAS cells, and to forward each of the plurality of virtual SAS cells from the first PHY along a specified virtual pathway to one of the plurality of PHYs.

### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of an exemplary SAS expander implementing a non-blocking SAS architecture utilizing virtual pathways in accordance with features and aspects hereof.

FIG. 2 is a block diagram of a plurality of exemplary enhanced SAS devices enabling a non-blocking SAS architecture utilizing virtual pathways in accordance with features and aspects hereof.

FIG. 3 is a block diagram of exemplary enhanced SAS devices implementing a non-blocking SAS architecture utilizing virtual pathways in accordance with features and aspects hereof.

FIG. 4 illustrates an exemplary format for a virtual SAS cell in accordance with features and aspects hereof.

FIG. 5 illustrates an exemplary format for an ADDRESS\_REQUEST virtual SAS cell in accordance with features and aspects hereof.

FIG. 6 illustrates an exemplary format for an ADDRESS\_ACK\_ACCEPT and an ADDRESS\_ACK\_REJECT virtual SAS cell in accordance with features and aspects hereof.

FIG. 7 illustrates an exemplary VCI connection state diagram in accordance with features and aspects hereof.