



To: Improved SCSI Protocol Ad-Hoc Group
From: Mike Kosco (MKosco@corp.adaptec.com)
Subject: SCSI LFP - Broadcast Command Packet Proposal
Date: Tuesday, June 17, 1997

1. Background

The data transfer rates for parallel SCSI have increased significantly in the past few years. The increase promises to continue with the advent of Fast 80/100. The percentage of SCSI bus time occupied by the protocol overhead to initiate and complete a data transfer is becoming a large part of the bus utilization.

This proposal is part of a set of proposals that seek to reduce the overhead associated with command execution and data transfer.

This proposal describes a method for reducing the overhead associated with transferring commands to targets. The broadcast command packet protocol combines the nexus information with the command descriptor block and transmits it as a packet to the target device. Further improvement is gained through the use of a broadcast protocol whereby multiple packets can be sent to multiple targets.

The intent of the proposed protocol is to be compatible with the existing parallel SCSI protocol, however the maximum benefit is achieved when all devices support the improved protocol.

2. Terminology

LFP – Low-fat protocol

BCP - broadcast command packet

BCP target - target device supporting the broadcast command protocol.

BCP initiator - initiator device supporting the broadcast command protocol

BCP phase - MESSAGE OUT phase with the REQ signal negated

3. Broadcast Command Packet (BCP) Protocol

3.1 Transfer Width Exponent Code

Table 1 defines extensions to the transfer width exponent code used in the WIDE DATA TRANSFER negotiation. These additional codes allow the devices to negotiate for a narrow or wide bus with BCP.

Table 1 – Transfer Width Exponent Code

Code	BCP width	Data Width
00h	none	narrow
10h	narrow	narrow
01h	none	wide
11h	narrow	wide
21h	wide	wide
02h	none	double wide data

3.2 BCP Packet

Table 2 illustrates a broadcast command packet. The packet contains the nexus information, a packet payload and a checksum.

Table 2 - Broadcast command packet

Bit	7	6	5	4	3	2	1	0
Byte								
0	CM	UT	Reserved	Target ID				
1	Reserved	Reserved	Reserved	Initiator ID				
2	IDENTIFY Message							
3	QUEUE TAG Message							
4	Queue Tag Value							
5	Reserved							
6	Packet Payload							
21	(LSB)							
22	Checksum							
23	(LSB)							

The command or message (CM) bit indicates the contents of the packet payload is either a command descriptor block or additional messages. If the CM bit is cleared the packet payload is a command descriptor block. If the CM bit is set the contents of the packet payload is additional messages.

The untagged or tagged (UT) bit, if zero indicates that the QUEUE TAG message and queue tag value are valid. The UT bit, if one, indicates that the I/O process is untagged and the QUEUE TAG message and queue tag value shall be ignored.

The packet payload field contains a command descriptor block (up to 16 bytes in length), or additional messages.

The checksum field contains the longitudinal redundancy check (LRC) used by the target to validate the packet. The most significant byte of the checksum field shall be zero if the negotiated bus width for BCP phase is narrow.

3.3 BCP Protocol

The BCP protocol begins with a NULL selection (i.e., selection with a value of zero on data signals). This should not interfere with the SCAM protocol because it has MSG asserted during a NULL selection. Upon detecting the NULL selection, all BCP targets enter into a BCP phase and await the packet transmission.

The BCP initiator transmits the packet with consecutive ACK pulses (similar to the way an offset is done in synchronous transmission today). The BCP target accepts in the packet and decodes the target SCSI ID. If the packet is destined for it, and the checksum is good, the BCP target asserts the REQ signal within a BCP reply delay time to generate a single pulse to indicate the packet was successfully received. The pulse shall meet the assertion period requirements for the negotiated synchronous data rate.

If the BCP target detects an error in the checksum it does not assert the REQ signal following receipt of the packet.

If the BCP initiator does not detect a REQ signal asserted within a BCP time-out delay time following the transmission of a packet it ends the BCP phase by negating the ATN signal.

If the ATN signal remains asserted this indicates that the initiator is prepared to transmit another packet. The subsequent packet may have any nexus, so multiple packets can be sent during a single BCP phase to one or more BCP targets.

The BCP initiator releases the ATN signal within a BCP last packet delay time to indicate the end of BCP phase.

At the end of BCP phase the BCP target identified in the first command packet retains control of the bus and may transition to a DATA phase, a STATUS phase or a MESSAGE phase after a BCP target control delay time. All other BCP targets release the bus within a BCP target bus release time. The other BCP targets that received packets have to perform a reselection to complete their commands.

3.4 BCP Phase

Figure 1 defines the BCP phase.

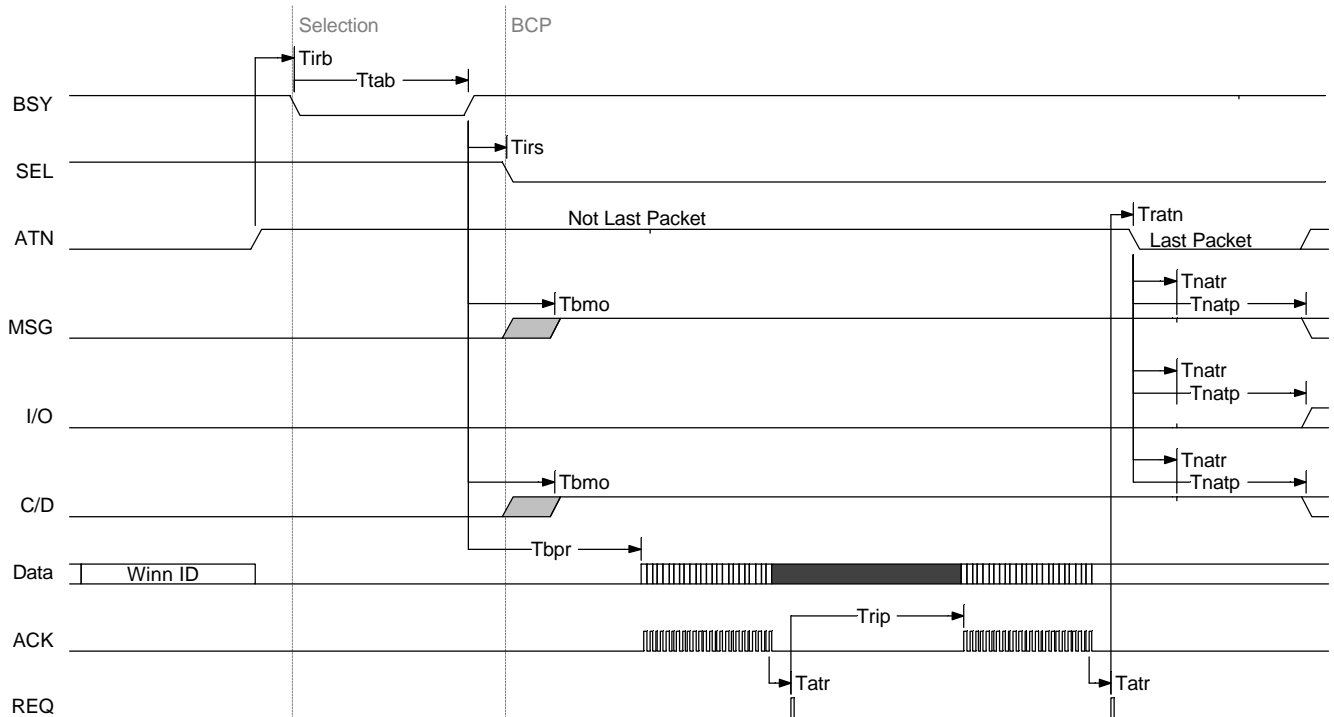


Figure 1 - BCP Phase Diagram

Table 3 defines the BCP protocol timing parameters.

Table 3- BCP Protocol Timing

Time	Description	Minimum	Maximum
T _{irb}	Initiator asserts data bus and ATN to releasing BSY (2 async system deskew delays)	90 ns	-
T _{tab}	Initiator releases BSY at start of SELECTION Phase to Target asserts BSY (bus settle delay) (selection time-out delay)	400 ns	250 ms
T _{irs}	Target asserts BSY to Initiator releases SEL (two deskew delays min)	90 ns	200 ns
T _{bmo}	First BCP Target asserts BSY to all BCP Targets driving MESSAGE OUT Phase	90 ns	200 ns
T _{bip}	First BCP Target asserts BSY to all BCP Targets ready to Initiator sends first byte of BCP packet	400 ns	-
T _{atr}	Initiator asserts last ACK to Target asserts REQ (BCP reply delay) (BCP time-out delay)	0	800 ns
T _{rip}	Target asserts REQ to Initiator sends first byte of next BCP packet	400 ns	-
T _{ratn}	Target asserts REQ to Initiator negates ATN (BCP last packet delay)	0	-
T _{bcpto}	Initiator asserts last ACK to Initiator negates ATN (BCP phase end)	1 us	-
T _{natr}	Initiator negates ATN to (non active) targets releasing all signals (BCP target bus release)	0	200 ns
T _{natp}	Initiator negates ATN to active target changes phase (BCP target control delay)	400 ns	-