To: X3T9.2 Committee (SCSI)   Document X3T9.2/88-69-R2

IBM

From   Date: August 29, 1988
Name & Phone: G. A. Marazas (407) 982-5972
U.S. Mail Address: International Business Machines Corporation
Entry Systems Division
P.O. Box 1328
Boca Raton, Florida 33432

Subject: Physical Architecture - SCSI 2 Interface, Revision 2

REVISION 2 CONTENT

The present submission adds additional functional capability in two areas: (1) A specific configuration activate command is defined using the base SCSI reset line and the config bus leads, and (2) Provision is made for a Primary Configuration master and an Alternate Configuration Master. The objective of the configuration activate command is to ensure the ability of new configurable devices to attempt normal operation on the SCSI bus even though they have not been configured. This feature accommodates failure or other nonavailability of the configuration master. It must be noted that all devices must start with unique Physical IDs if successful operation is to occur without electronic addressing having been performed. The activation command places a configurable device in a state such that it allows itself to participate in electronic assignment of its Physical ID address. Prior to activation, the device functions normally per the Base SCSI architecture and using the Physical ID pinned in its jumpers. Provision of an Alternate Configuration Master establishes a backup unit which is allowed to provide electronic address services when the primary is not able to do so.

REVISION 1 CONTENT

The present submission is a major revision and incorporates many of the comments made in the Boise Working Group meeting relative to the original version of this proposal. The most significant change is to reduce the number of active signal leads downward from the original eight to the present four. This is done to improve the implementation of the present electronic configuration feature into a single silicon chip also supporting the base SCSI architecture. The total number of active drive leads for a single ended implementation is now 18 plus 4 equal to 22. A secondary benefit is that sufficient pins are now available within the scope of a 60 pin SCSI A cable connector to explicitly support a differential implementation with push pull pairs. Additional changes are made to support electronic configuration in a multi-initiator environment and when the configuration master is at any arbitrary location on the SCSI bus rather than the previous case of requiring the Master to be at either end only. Support also is stated for a bypass function in which the configuration process steps over any SCSI device which is in a power off state while configuration is underway.

The system philosophy of electronic configuration is that operator intervention rather than automatic adjustment is the way to address the error condition in which multiple Configuration Masters are placed on the same SCSI bus. Hence, the present proposal has means to detect the presence of multiple masters but does not provide automatic means to select one of the Masters as the head Master. Resolution of this conflict occurs via Systems Operator response to an error condition noted by devices supporting this proposal.

INTRODUCTION

This proposal addresses reliability and performance tuning aspects of the Physical interface between SCSI devices and their host system interface. The intent is to define an optional extension to the present 50 lead SCSI bus supporting the first 8 bits of SCSI data path. Provisions are made to support more than 8 SCSI devices in the context of a "Wide SCSI" supporting 16 or 32 bits of data path. A significant feature of the proposed expanded interface is that new devices supporting the proposed expanded interface and associated architecture can coexist freely with devices which have not implemented the proposed interface and/or the proposed functions. No limitation shall exist relative to implementation of the proposal in a single ended or in a differential SCSI environment.

A full physical level realization of the proposal would require a 60 lead (for first 8 bits of data path) interface between the systems attachment point and the given SCSI device or enclosure box supporting the SCSI device. A likely and fully allowed physical level subset of the proposal would involve two cables between a "Systems Point" and the SCSI device. The "Systems Point" would likely be a special internal only attachment connector for those SCSI devices installed inside a Systems unit containing the host CPU. Another form of "Systems Point" would be a transition connection for those SCSI devices mounted within an enclosure box providing power and packaging support for SCSI devices installed external to the Systems unit. In both of the above "Systems Point" cases, the two cables in the physical subset implementation would be a Base SCSI 50 lead signal cable and a special four lead cable providing configuration services in a single ended environment. It is emphasized that the subset is at the physical level only; all logical functions associated with configuration reset are supported.

In both the full physical level implementation as well as the physical level subset, the first 50 leads of the interface are identical to the present SCSI standard and thus are not discussed further within this proposal. The present discussion focuses entirely on the functions and architecture supported by the additional 10 leads. Special emphasis in the discussion is given to the four lead subset providing configuration reset service in a single ended environment.

PROBLEM STATEMENT

The present proposal addresses functions which ease reliability, performance, and user friendliness problems existing in the present definition and implementation of the SCSI interface. These problems are respectively:

O RELIABILITY

The SCSI bus requires termination at both ends in order to ensure proper operation. In many cases the SCSI bus is implemented with a portion inside the systems unit and another portion external to the systems unit. Most often the external portion is exposed to relatively frequent additions/deletions of attached equipment. Too often, the proper placement of the terminator unit is neglected/ forgotten. On those occasions when the terminator is not put in place, then intermittent errors may occur and problem identification is very difficult. The present proposal contains means to detect the presence of the external terminator and in the process ensure more reliable operation.

O PERFORMANCE

Priority of data transfer on the SCSI bus is set by the value of physical ID of each connected device. During arbitration and reconnect cycles, if multiple units compete, the unit with the highest numerical value of physical ID will win. This victory is independent of the logical role in the applications environment played by the given device relative to all other devices attached to the same SCSI bus. The assumption is that the user has chosen intelligently the values of physical ID for each device. One problem is that the selection of physical ID is made by operator manual intervention via jumpers or rocker switches to be set on each device. The physical difficulty in changing the physical ID requires that a static performance optimization be made. Many application environments are dynamic and would benefit from more frequent changes of device ID than is possible today. The present proposal provides electronic means to change the device ID so that dynamic performance tuning becomes practical.

O FRIENDLINESS

As an additional aspect of the point above, it is difficult to set the required value of physical ID. Physical location of ID setting varies widely on devices from different manufacturers. Additionally, various system packaging concepts may make access difficult to the involved jumpers/ switches. In certain instances, tools are required to make the necessary change. As SCSI devices become more common in personal computing and data processing environments, it becomes a problem if not unreasonable burden to the user. The present proposal for electronic means to set physical ID makes SCSI substantially more user friendly.

SOLUTION ARCHITECTURE

The two basic functions discussed in the problem section are (1) detection of external terminator and (2) the provision of electronic configuration capability to define the value of physical ID for a given SCSI device. Function (1) is supported by means of two physical leads and function (2) is supported by eight physical leads -- making a grand total of 60 leads connecting the system attachment point and each SCSI device. The first 50 of these leads are exactly as defined in the SCSI 1 specification and in the A cable for SCSI 2. This section discusses the role played by only the 10 new leads required to implement the functions contained in this proposal. A following section describes flow scenarios depicting the manner in which the new functions interact with a SCSI system.

The following table provides definition for the ten new leads.

Lead	Function	Comments
51	Ground	Provides continuation to lead 50 of alternating pattern signal and ground
52	Terminator Detect	Supports subject function
53	Ground	Possible differential pair to lead 54
54	Config 1	Lead paired logically with Config 2 to form in/out bus for configuration
55	Ground	Possible differential pair to lead 56
56	Config 2	Lead paired logically with Config 1 to form in/out bus for configuration
57	Ground	Possible differential pair to lead 58
58	Config 3	Lead paired logically with Config 4 to form out/in bus for configuration
59	Ground	Possible differential pair to lead 60

60	Config 4	Lead paired logically with Config 3 to form out/in bus for configuration

Function (1) is supported by leads 51 and 52 -- or specifically lead 52. The detection of the external terminator may be performed by devices connected to the external SCSI bus. The terms internal and external are from the perspective of the system unit or box containing the means of connecting SCSI to the given system. Internal refers to that portion of the SCSI bus within the system box. External refers to that portion of the SCSI bus outside or external to the system box. In passing it is noted that the implication is that external SCSI units are in their own power envelope and are placed in an enclosure box providing power and packaging support. To ensure proper operation, the last external SCSI device must contain a terminator plug. That plug when properly in place according to this proposal will provide ground on lead 52. When the terminator is not in place, then a float signal level will be observed. In implementation, a high value of pull up resistor will be used to force the float signal level to a high level thus aiding the detection process.

Function (2) is supported by the four leads, 54, 56, 58, and 60 in the single ended implementation of the proposed architecture. In the differential case, all eight leads 53 to 60 are used. Also used to support the electronic configuration function are the Busy lead, the Reset lead and the Data BUS from the Base SCSI implementation. Figure 1 below describes the situation.

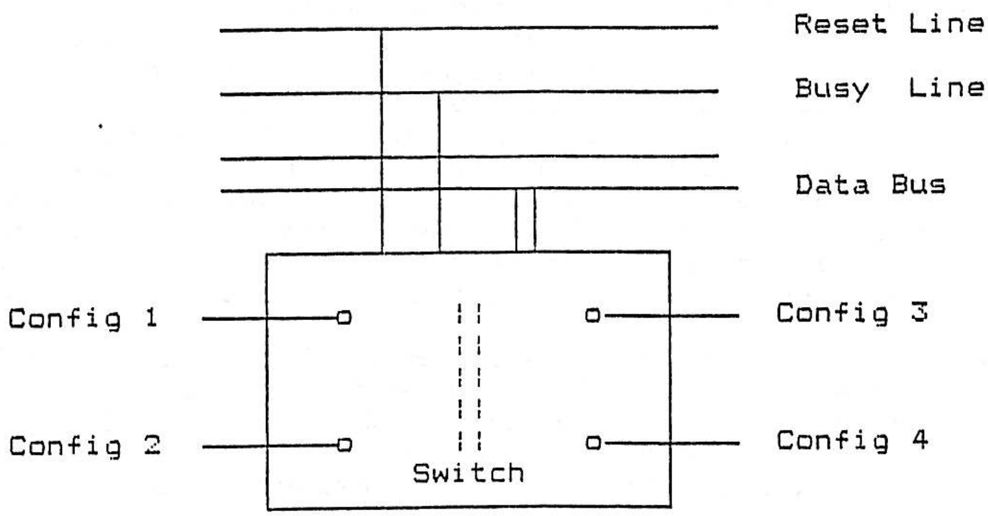

Figure 1 Electronic Configuration view of a Device

0 DEVICE DESCRIPTION

Figure 1 illustrates a SCSI device from the perspective of its participation in the process of Electronic Configuration. The device is shown in the form appropriate: (a) immediately after power on or, (b) after having processed a proposed new command sent in a normal fashion via the Base SCSI bus and telling the subject device to revert to a non configured state. Participation in the configuration process entails four special leads (54, 56, 58, 60) in the expanded function, 60 lead SCSI cable A. These leads, prior to device configuration are referred to as leads Config 1 through 4. This naming convention emphasizes the fact that the leads may serve either as input or as output for the device. During the Address Assignment phase of configuration, input and output status is be assigned to the leads. The normal position of the internal configuration switch is defined for the direction of information flow required to support Address Assignment. Transfer of signals through a given device is accomplished by the process of regeneration and new transmission on the output leads of those signals appearing on the input leads. As will be explained later, the device must also support a reverse direction flow of information to pass confirmation signals back to the Configuration Master. The reverse flow occurs when input signals appear on leads normally assigned output status. The normal direction of the configuration switch is then temporarily reversed. A bypass function, described later, provides electrical continuity between Config pairs 1 and 2 when the device is in a power off state.

The Reset lead together with the special Config leads cause the device to enable the function of setting its Physical ID address by means of the proposed electronic addressing mechanism. Prior to this activation, the device functions normally per the Base SCSI architecture using the Physical ID set in its jumpers. In this manner the device attempts normal operation on the SCSI Bus in the event the Configuration Master or its Alternate is unable to execute the electronic assignment method for establishing the address of configurable devices. It continues to be a requirement that all devices, including configurable devices, must have a unique Physical ID in order to pursue normal operation on the SCSI Bus. Note that electronic assignment of address has no requirement for unique Physical ID at the start of the address assignment process. Successful conclusion of the electronic address assignment process will leave the device with a unique Physical ID.

Also shown in figure 1 are the Busy lead and the Data Bus portion of the Base SCSI Bus. The Busy lead is required to be asserted for that portion of the process in which leads Config 1 to 4 are actively being used to convey information associated with electronic assignment of address. Asserting the Busy state insures that the Base SCSI Bus will not be used for normal data and message flow during the configuration period. The Data portion of the SCSI bus will be used to convey various Physical ID's. It is the intent of this proposed architecture that it support both the A Cable and B Cable of a SCSI 2 device. In this manner, if a future version of SCSI permits more than 8 SCSI devices on the bus, then the data portion of the B Cable may be used for extra ID bits for devices beyond the first 8.

O UNIQUE FEATURES OF THE CABLE I/F

It is noted that the expanded SCSI A cable has a split personality as defined in this proposal. The Base SCSI cable is a true Daisy Chain so that signals on it are conveyed to each and every device on the SCSI bus. In strong contrast, the special leads Config 1 to 4 function both logically and physically are a series of point to point cables between each device. It is this point to point nature of the special leads that allows SCSI devices to be selected in a simple and straight forward fashion on a one at a time basis. Thus the Configuration Master or its authorized Alternate has the ability to select and configure individual devices without confusion to or interference from other devices on the SCSI bus. It may be stated that the present proposal makes a trade-off in favor of a cable system which is more complex but a logic processing scheme which is less complex and of low implementation risk. Alternatives attempting to perform the same function may perhaps make the trade-off in favor of an approach which is less complex at the cable level but much more complex at the logical level and therefore more exposed to higher implementation risk.

Leads Config 1 to 4, are bi-directional so that they may either receive inputs or drive outputs. A key feature of their supporting logic is that leads are continuously monitored for the appearance of input. The leads operate in pairs. That pair first detecting the Address Assignment signal as an input is then designated as the input pair. The specific nature of the Address Assignment signal will be defined later. The remaining pair in the Config group is then normally intended to send output signals from the device. Among other aspects of the Configuration Process, the so-called Configuration switch is set in a manner so that authorization is given for the device to repower and send forward on the output pair the same logical signal levels observed on the input pair. In this manner, once the Configuration Switch is set, the given device will act to propagate configuration signals to the next device in a fixed direction, left to right if the left pair first saw an input or from right to left if the right pair was the one first seeing an input during the configuration process. The signal intended to be propagated in this manner is the Address Assignment signal.

O DISCUSSION OF ERROR STATES

Note well that the output pair continues to look for input signal on its leads. The Confirm Accept and Confirm Reject signals will normally flow in the reverse direction. These signals as explained later represent the responses of a device to an earlier Address Assignment signal and will appear as inputs on the output pair of devices between the subject device and its Configuration Master. It is considered a fatal error to the configuration process if an output pair should unexpectedly see an Address Assignment signal. This assessment of error follows from the fact that there is to be ONE AND ONLY ONE Configuration Master. That Master is the only entity allowed to change the Physical ID of other devices on the

SCSI Bus. Provision is indeed made for an Alternate Configuration Master. However, the Alternate is allowed to function as Configuration Master in the event that the Primary is unable to carry out the Electronic Address Assignment function. As described later, assumption of Configuration Master status is accomplished during the phase that signals are sent on the bus to command configurable devices to enable their Electronic Address Assignment function.

In a general situation, the Master may be an interior device in a linear string of devices. A given device to be configured may be either to the left or to the right of the Master. By wiring convention, Leads Config 1 and 2 will be the Input pair when the subject device is to the left of the Master. Conversely, Leads Config 3 and 4 will be the input pair when the subject device is to the right of the Master. If a mistake occurs, and more than one device believes it is the Configuration Master, then a given device may find itself with a Master on both its left and its right.

In this error case with a Master on either side, the subject device will find that Address Assignment signals occur at different points in time on both pairs of the Config leads. Recall that the first Config pair to see an Address Assignment signal is designated as the input set. With one Master in the system, that is the only pair which would ever see the Address Assignment signal. Thus, it is possible to declare the error condition of multiple Configuration Masters if the Address Assignment signal unexpectedly occurs on the Config pair designated as output.

CONFIGURATION EXAMPLES

Figure 2, on the next page, illustrates the case of a configured device. In the case shown, the Configuration Master is assumed to be to the right of the subject device. Thus, Address Assignment signals from either the Master directly or from intervening devices should only be seen on leads Config 1 and 2. In passing it is noted that it is normal and proper for the subject device to be presented the Address Assignment sequence of signals at multiple points in time during the configuration process. The first instance of this sequence occurs the device itself is having its own address redefined. Later, the device encounters the Address Assignment sequence when it passes this sequence on to other, farther devices having their addresses defined. The configuration switch within the subject device would be set in the right to left direction as shown in part (a) of the figure. Leads Config 3 and 4 would then be assigned as output so as to propagate Address Assignment signals to other devices (if any) located to the left of the subject device. Part (b) of figure 2 illustrates the location of the Config leads within the connector attached to the subject device.

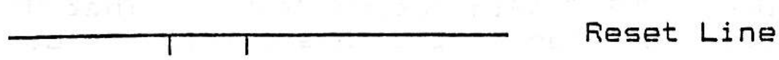

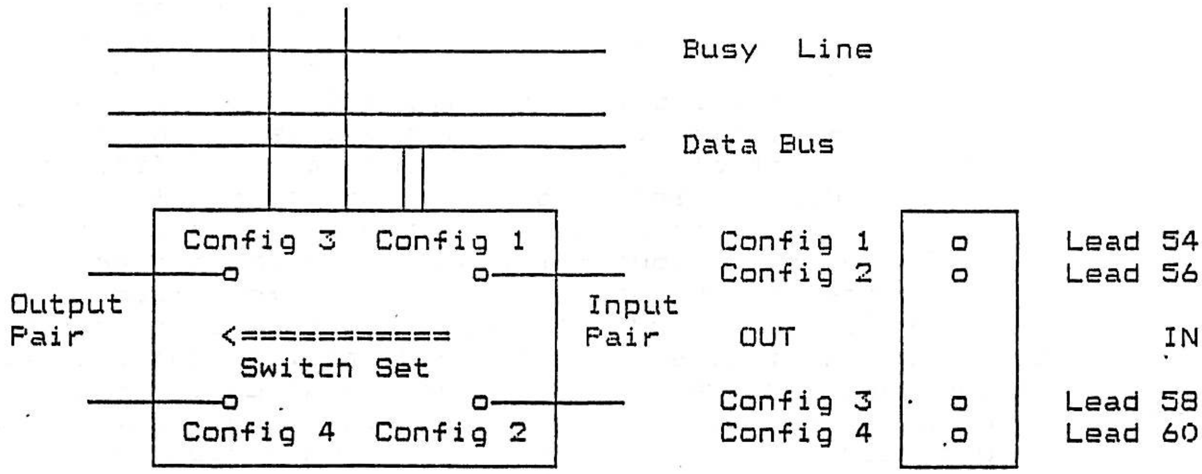

(a) Logical View (b) Connector Detail

Figure 2 Configured Device

0 CORRECTLY CONFIGURED SCSI NETWORK

Figure 3 illustrates a network of 5 SCSI devices.

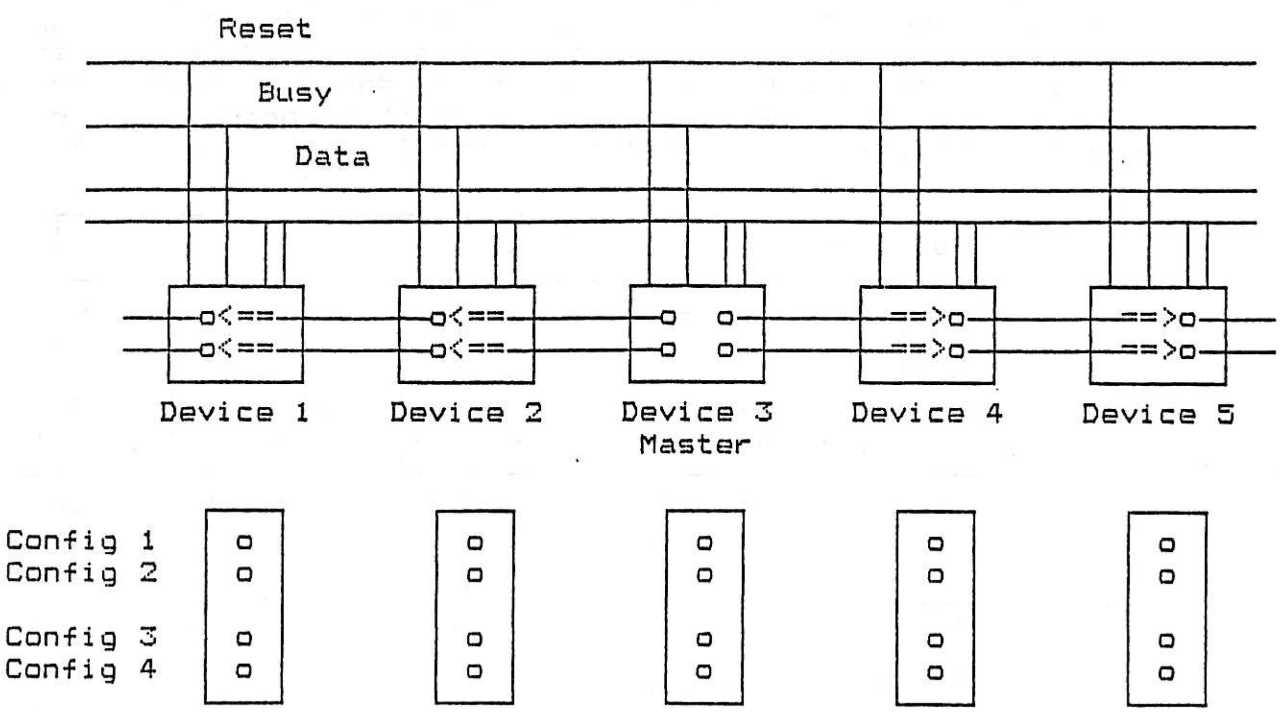

Figure 3 Network of SCSI Devices After Configuration
Case Shown is Master Sending to Other Devices

In the case shown, Device 3 is designated as the Configuration

Master. Devices 1 and 2 are to the left of the Master. Note that for these two devices, the configuration switch is set indicating right to left forwarding of signals from the Master. Leads Config 1 and 2 are the input pair. Leads Config 3 and 4 are the output pair. A different state of affairs exists at devices 4 and 5 which are to the right of the master. Here, the configuration switch is set to indicate left to right forwarding of signals from the Master. Leads Config 3 and 4 are the input pair while leads Config 1 and 2 are the output pair. Note that the Master uses both of its Config pairs for output when sending information to devices to be configured. As noted, in this protocol, there are times when the device sends signals back to the Master. In these circumstances the Master must receive on leads otherwise used for output. Also, the sending device uses its designated input pair to send special output back to the Master. Any intermediate device between the subject device and the Master must then temporarily reverse the direction of its configuration switch in order to support the reverse direction flow from device to Master.

O SCSI NETWORK WITH CONFIGURATION ERRORS

Figure 4 below illustrates a case of multiple Configuration Masters on a single SCSI Bus. The two Configuration Masters are devices 2 and 5. In one scenario, devices 3 and 4 would declare an error state. Device 3 could be configured by device 2 in left to right mode. Device 4 could be configured by device 5 in right to left mode. The error seen by device 3 occurs when it is later attempted to be configured by device 5 using device 4 as a propagation intermediary. In similar fashion, device 4 would declare an error when it is attempted to be configured by device 2 using device 3 as a propagation intermediary. In passing it is noted that device 1 would be correctly configured by device 2 in right to left fashion.

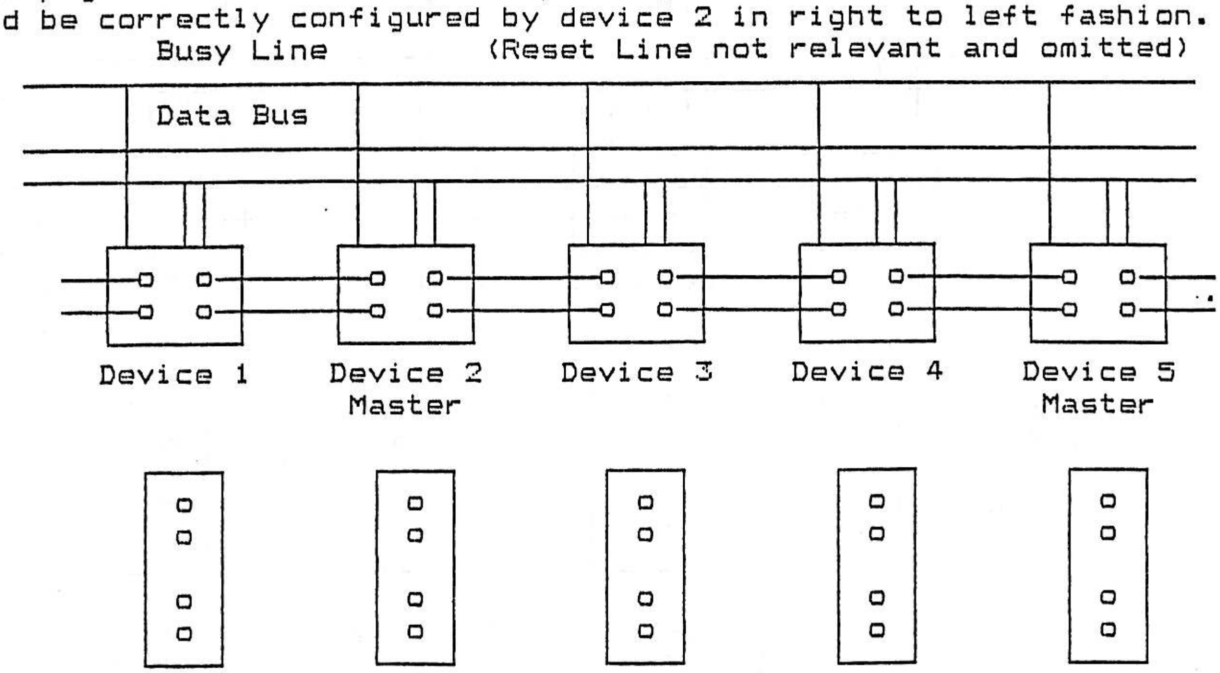

Figure 4 Example of an Incorrect Configuration

DEFINITION OF SIGNAL STATES

The table below defines the allowed states in the SCSI system as involving the reconfiguration process. Note that value = 1 is the on condition otherwise known as signal true or signal assertion.

State Name	Reset	Busy	In 1	In 2	Data Bus
Enable Config Process	1	1	1	1	ID of Master
Address Assignment	0	1	1	1	ID Sent by Master
Confirm Accept	0	1	0	1	ID Accepted
Confirm Rejected	0	1	1	0	Current ID of Device
Neutral State	0	1	0	0	No Data
Error State	0	0	1	1	Not Driven

Note the Busy condition is asserted by the Configuration Master from the start to the finish of the electronic configuration operation. This is the only time during normal operation that the special bus (Config 1 to 4) is to be in use. The only exception is within one configuration time out period after Busy is deasserted. In this period, devices may assert the special bus for the purpose of declaring a Configuration error.

0 ENABLE CONFIG PROCESS

New devices supporting the electronic address assignment procedure are not allowed to be reconfigured (change their Physical ID) unless authorized and enabled to do so. The enable function causes two changes in devices which support the present, proposed architecture. First, the device temporarily removed from functioning in a normal fashion on the Base SCSI Bus. That is device is prohibited from responding to any data and message sequence on the SCSI Bus. The device continues in the prohibited state until (a) it successfully exits the electronic address assignment operation and thus has a new electronically assigned Physical ID or (b) it receives a new command (described later) called REVERT TO UNCONFIGURED STATE. The second and complimentary change in the device is that it is authorized to enable its electronic address assignment function. The purpose of the present command is to provide some way for the device to undertake operation in as normal a fashion as possible operations on the Base SCSI Bus in the event there is no Configuration Master present --

either Primary or Alternate.

It is during the Enable Config sequence that it is decided whether the Alternate Config Master determines whether is is to assume status as Master due to lack of presence of the Primary. If an Alternate Config Master is designated, then it must have a Physical ID which is both different and numerically lower than the Physical ID of the Primary. During the subject Enable sequence, all candidate Config Masters must assert their Physical ID on the Base SCSI Data Bus. That Master with the highest numerical ID wins is to assume sole and uncontested status as Config Master. If the Primary should be unavailable, then the Alternate comes forth and assumes responsibility and authority for the Config Process.

o ADDRESS ASSIGNMENT

The Address Assignment State is issued by the Configuration Master. The intent is to define the Physical ID of a given device. Note, the device being configured may be either an Initiator or a Target. In both cases, given that they support the proposed architecture, the device is required to accept the new ID if logically able to do so. Address Assignment may be done only once in the period following a power on with an subsequent Enable Config sequence or following the command to deconfigure. The subject device is required to maintain for the duration of its configuration a record of: (a) the Value of the Physical ID accepted , and (b) the direction, left or right, of the Master. The Physical ID to be accepted is provided on the Base SCSI, Data Bus in bit position format. The configuration switch is not set until after a Confirm Accept signal is sent by the device.

The error state is to be declared and Confirm Rejection sent if two or more bits are detected for the Physical ID on the Data Bus. Error state is also to be declared if the Address Assignment Signal simultaneously occurs on both Config pairs, or if it ever occurs on the designated output pair. Both error states would occur if two or more devices on the same SCSI bus thought they were Configuration Masters.

o CONFIRM ACCEPT

The Confirm Accept State is issued by the subject device and indicates it accepts and will use the Physical ID passed to it during the Address Assignment phase. Acceptance of the new ID is mandatory if the device is logically able to do so. Refer below to part (c) Confirm Reject for conditions under which the device is allowed to reject an assigned ID.

The Confirm signal is sent as output from the device on the same Config Pair in which the Address Assignment signal was received. The confirm signal cannot be sent until after the device detects the neutral state on the newly designated input pair. Optionally, the subject device drives the Base SCSI data bus with the Physical

ID it has accepted for this configuration operation. The intent is to provide additional error checking means for the Master which expects to see the assigned Physical ID. The Master declares an error state if it sees the wrong ID. Making the retransmission of the ID optional is intended as an aid to devices implementing the proposed architecture using a combination of an existing SCSI protocol chip and additional logic specifically supporting the proposed new functions.

The subject device is to set its configuration switch immediately after sending the Confirm Accept signal. From this point on, the subject device propagates onward in the direction indicated by the configuration switch all subsequent Address Assignment signals received on the designated input leads. No action other than propagation is to be taken regarding these proper assignment signals. An error state is to be declared if the Assignment Signal is observed as input on the Config pair designated for output. The device does not propagate any Assignment Signals associated with an error condition.

o CONFIRM REJECT

The Confirm Reject State is issued by the subject device and indicates a refusal to accept the Physical ID passed to it during the Address Assignment Phase. Valid reasons to reject are: (1) The Assignment came from the wrong direction over a Config Pair assigned to function in output mode, (2) Multiple ID bits were observed on the data bus, or (3) simultaneous inputs are detected on both Config pairs.

The Reject signal is sent as output from the device on the same Config Pair(s) in which the Address Assignment signal was received. The confirm signal cannot be sent until after the device detects the Neutral State on the given Config pair. For error (1) above, the subject device optionally drives the Base SCSI data bus with the Physical ID previously accepted from the other Master in this configuration operation. (Condition 1 could not occur without a prior Config Accept.) The intent is to provide the second Configuration Master with the Physical ID to which the subject device will respond in normal operation. For errors (2) and (3) above, the Data Bus is optionally driven with the ID set in the jumpers of the subject device. Note in all cases the device sending Confirm Reject is identifying itself and providing the Master with the Physical ID to which it will respond. The subject device does not change the direction of its configuration switch if a Confirm Reject signal is sent. Again, making the transmission of the ID optional is intended as an aid to devices implementing the proposed architecture using a combination of an existing SCSI protocol chip and additional logic specifically supporting the proposed new functions.

Upon receipt of the Reject signal, the Master is to terminate its own efforts at electronic configuration. The Master ceases to assert Busy on the Base SCSI Bus. Upon deassertion of Busy from the SCSI Bus, the Master is to declare Configuration Error over both of its Config pairs.

O NEUTRAL STATE

The Neutral State is issued by a Configuration Master after ending assertion of the Address Assignment state. The Neutral State indicates that it is time to issue some form of Confirm signal from the device which is currently the object of Address Assignment. The Master waits a time out interval for a Confirm Signal. If no confirm is observed, the Master concludes that the object device does not exist. The Master then ends Address Assignment in the given direction.

O ERROR STATE

The Error State, if necessary, is entered immediately after Busy is deasserted from the Base SCSI Bus. The error signal is initially issued by those devices, including Masters which have detected an error during the Electronic Configuration operation. Any Devices seeing this signal on one of its Config Pairs takes note of the error condition but does not attempt to propagate it further. After another time out interval, the error state is cleared by all devices from the special Config Bus.

The intent of the error state is to call attention to all devices on the SCSI Bus of the fact of a fatal error having occurred during the configuration operation. It is expected but not required that operations cease until human intervention acts to clear factors responsible for the error state. In most instances, it is expected the errors if any will be associated with multiple Configuration Masters on the SCSI Bus. By design philosophy, it is believed that human intervention rather than automatic procedures be used to identify the problem and to take appropriate action.

CONFIGURATION SCENARIO

The reconfiguration process must begin with the MASTER forming a map of SCSI devices and their Physical ID's. The mapping process must proceed first with a check of the SCSI bus for all old devices not supporting the proposed architecture for electronic configuration. A reset must have been issued to break the setup which previously may have been put in place. The reset applies only to devices supporting the Electronic Address Assignment architecture. The reset for these devices is accomplished by the Enable Config Process sequence. After configuration reset there is the guarantee that no new device will respond on the base SCSI bus. This being the case, the initiator uses the base SCSI bus to make an Inquiry against each allowed Physical ID, 0 to 7 in the normal case. Note, that two old devices must not have the same Physical ID.

Upon completion of the mapping process, the knowledge is obtained as to those Physical ID's available for reassignment and those ID's in use by given old devices and thus unavailable for reassignment. The configuration then proceeds relative to new SCSI devices supporting this architecture. It is noted as an emphasis point, that multiple of the new devices may have duplicate Physical ID's prior to reconfiguration. In fact, a new device may start with an ID which is the duplicate of the ID of an old device. After configuration as described above, all devices must have unique Physical ID's.

Please refer back to Figure 3 to observe one possible instance in which a single Configuration Master acts to configure a SCSI Bus. As shown the Master is Device 3. Assume the Master begins configuration after a power down, power up cycle. Config Reset has been accomplished via the Enable Config Process sequence. If present, the Primary and Alternate Config Masters have resolved the question as to who proceeds with Electron Address Assignment as Master. The Master starts by issuing an Inquiry command to each device on the bus. At this point only old devices not supporting the Electronic Configuration are able to respond. By design, new devices cannot participate on the Base SCSI Bus until they have been electronically configured. The old devices respond with their Physical ID as set via jumpers. The Master builds a logical map with those old devices and from this produces a list of ID's not in use by old devices. The Master then arbitrates for and upon winning an arbitration on the Base SCSI Bus, asserts Busy so as to prevent operation at this time by the old SCSI devices.

Assume next the configuration operation first proceeds in a left to right fashion. The sequence of signals are:

<> Cycle 1 --
- o Master Asserts Address Assign with New ID for Device 4
- o Master takes special bus to the Neutral State
- o Device 4 returns Confirm Accept and sets Config Switch
- o Device 4 takes Special Bus to Neutral State

<> Cycle 2 --
- o Master Asserts Address Assign with New ID for Device 5
- o Device 4 receives and forwards signal to Device 5
- o Master takes Special Bus to the Neutral State
- o Device 4 receives and propagates signal to Device 5
- o Device 5 returns Confirm Accept and sets Config Switch
- o Device 4 receives and propagates signal to Device 3
- o Device 5 takes Special Bus to Neutral State
- o Device 4 forwards Neutral State to Device 3

<> Cycle 3
This is a repeat of Cycle 1 using device 2 instead of 4

<> Cycle 4
This is a repeat of Cycle 2 using device 1 instead of 5

TERMINATOR DETECT

The mandatory portion of this proposal is that if terminator detect is implemented, then it proceeds in relation to lead 52. Differences exist as to whether the detect operation is restricted to the digital case of present/ not present or to an analog case with capability to detect the number of terminators and perhaps even their general location on the Bus. Consequently, it is left as an implementation option in the digital or analog level of function is provided. The present submission describes the digital case of none versus one or more. If an initiator fails to detect ground on lead 52, then a declaration is made that the required external terminator is missing. Initiator code will make the decision to declare an error condition based on policy guidelines previously established by Host software. As agreed, the initiator can return some level of error condition in response to a service request. In the alternative, the initiator can retain the knowledge of error condition to relay to the HOST when some appropriate inquiry is made of SCSI status.

BYPASS FUNCTION

It is desired that configuration services be impacted to the least extent possible by a unit at an arbitrary location on the bus and existing in the power down state while configuration is underway. The concern is that the power down unit might block provision of configuration services to devices "down stream" and which are in the power up state and otherwise able to participate in electronic configuration. The resolution of this concern is to provide a bypass function accessible to each device. The notion is that means be provided which operate in a power down state at individual devices such that Config Pair 1 is directly connected to Config Pair 2 in a bi-directional sense.

In essence, the bypass function acts in a manner equivalent to a switch which is normally closed under power down and normally open in the power on state. Many different physical implementations exist as alternatives to provide the desired function. At one end of the solution spectrum there is the mechanical interlock switch which opens or closes the desired contacts when the power switch is on or off. Here, the function is performed with a reliable item of equipment and in a manner that does not consume power from the SCSI Bus. At a different point in the spectrum is a solid state switch (any one of several varieties possible) which performs the desired function but which also consume modest power from the SCSI Bus. At a still different point in the spectrum is a relay device which performs the function while potentially consuming relatively high levels of power from the bus.

It is neither necessary nor desired to specify implementation details regarding implementation of the bypass function. What is done here is to recommend strongly that devices implementing the proposed architecture for electronic configuration also implement the bypass function. It should be noted that SCSI devices are quite likely to be in a support box of some type providing power and packaging. In instances such as magnetic media DASD and optic media DASD, these units are likely to be inside the systems box or in some support box separate separate from the system box. In either case -- System box or Support box -- the DASD units, etc. receive the benefits of a support structure considerably larger and more rich than that provided within the constraints of the 3 1/2 inch or 5 1/4 inch form factor. These support boxes/structures may well contain more than one DASD unit. Thus, support can be provided more economically on a shared basis over multiple devices when the bypass function and/or cable connection means are done at the support box level rather than at the individual device level.

SYSTEM INTERACTIONS

To complete the statement of the subject proposal it is necessary to address various miscellaneous factors governing operation of the device and the several others parts of the system. These miscellaneous factors include the means for deconfiguring a device, detecting transitions of signal state on the Config bus, and trade-offs regarding drive current versus timing allowances.

O NEW COMMANDS AFFECTING CONFIGURATION STATE

As previously described, devices supporting this proposed architecture go into a non configured state after undergoing a power down, power up cycle. That is to say, devices upon power up revert to the Physical ID set in their jumpers. Also, devices of the subject type are prohibited from full participation on the SCSI bus once they have accepted the Config Enable sequence. The limited participation allowed unconfigured, new devices is that they may listen for commands sent them over the base SCSI bus -- they may not originate traffic on the bus until configured. The listening process occurs only when Busy is not being asserted. Unconfigured devices listen only after end of the configuration operation is completed.

The first new command is titled -- REVERT TO UNCONFIGURED STATE. The function performed is to permit the Configuration Master to reconfigure one or more devices at any arbitrary point in time after initial power up. The command is limited to a new device supporting the subject architecture and requires that the subject presently be in a configured state. Thus, the subject device is responding as a device with Physical ID given it via the electronic configuration operation.

To the "Electronic" ID is sent the present REVERT command which causes the subject to respond in the same manner as it did previously after entering the Config Enable state. This is to say, the device:

(a) Replaces the "Electronic" ID with its Jumper ID
(b) Sets to neutral its configuration switch

(c) Begins once more to monitor and respond to the Config Bus.
(d) Goes into limited participation mode on the base SCSI Bus.
In this manner the subject device is once more capable of being reconfigured as described by the proposed architecture. The command itself is of the six byte variety and with a command code to be designated. There are no parameters modifying the command other than the obvious specification of the Physical ID of the subject device.

The second and final new command is titled RECOVER FROM CONFIG ERROR STATE. By virtue of the proposed architecture, the subject device must respond to the Address Assignment state with one of the two defined types of Confirm response. In both instances, the subject device optionally states the Physical ID to which it will respond. Thus, for those devices having responded with their ID there is the possibility of sending them the RECOVER command. The object of the RECOVER command is to allow a SCSI device to participate in a normal fashion on the SCSI bus in recovery from an error in the configuration operation. The command itself is of the six byte variety and with a command code to be designated. There are no parameters modifying the command other than the obvious specification of the Physical ID of the subject device.

O DETECTING TRANSITIONS ON THE CONFIG BUS

The intent is to permit operation of the Config bus at a very low data rate so that implementations may adopt low drive current Driver/Receivers. Reduction in the drive current level is useful in that it minimizes the difficulty of implementing a new SCSI protocol chip supporting the proposed architecture in combination with the Base SCSI architecture. The reduction in data rate acts to simplify operation of microprocessor supporting the hard wired portion of the implementation.

The intent is monitor the four leads of the Config Bus via a polling mechanism rather than an interrupt mechanism. During electronic configuration the subject device can at best engage in limited participation on the SCSI bus until configuration has been completed. During the active phase of configuration, the Busy condition is maintained by the Master and the subject device is to have no operation other than in terms of the Config Bus. Thus, the support microprocessor has ample time to devote to monitoring the two Config pairs for some input signal.

One technique of operation is for the support microprocessor to monitor each of the four Config leads on a given time cycle. A record is maintained of the last valid signal level (True or False) validated on the given line. At some point, the poll operation will reveal a new sample with value different than the previously accepted value. A qualification procedure might then begin so as to determine the validity of the new value. Possibly the qualification proceeds by the mechanism or requiring some number of successive poll samples to each be the same and all different from the previously accepted value. Upon completion of the qualification process, the proposed architecture describes what is to be done with the newly determined signal value. As indicated in the next section, trade-off analysis is currently underway to fix values for drive current used by the Config bus and the associated data rate and transition rate for Config cycles. It is anticipated the resulting operation point would involve state transitions on the Config bus completing on the order of milliseconds.

O DESIGN TRADE-OFFS INVOLVING DRIVE CURRENT

Evaluation work in this subject area is still underway with no recommendations yet available. The general nature of the trade-off is as follows. Given that a low data rate is acceptable, then in a single ended implementation it is possible to balance reductions in drive current with higher values for resistors in the special terminator network supporting the four Config leads. As the resistance in the terminator network increases, then state transitions on the given lines (Config 1 through 4) become more sine wave in shape. Thus, a sharp leading edge square wave with associated short time duration transition zones is replaced by more rounded wave forms with longer time duration required for the transition in state.

The present approach for electronic configuration is fully capable of operation at a transition rate far slower than the four megabytes per second associated with normal data transfer operations in synchronous mode SCSI. In consequence of the tolerance for relative low rate of information flow rate, the trade-off is possible as indicated above.

Gerald Marazas

G. A. Marazas