

CONGRUENT SOFTWARE, INC.
98 Colorado Avenue
Berkeley, CA 94707
(510) 527-3926
(510) 527-3856 FAX

FROM: Peter Johansson
TO: IEEE P1394.1 Working Group
DATE: October 15, 1998
RE: Virtual node IDs

This document attempts to capture the P1394.1 design effort, to date, on "virtual node IDs". A virtual node ID is a 16-bit "handle" used as a destination or source address in an asynchronous primary packet so long as *at least one bridge is on the path* between the source and the destination. Virtual node IDs supplement conventional (local) node IDs, they do not replace them.

The salient features of virtual node IDs are as follows:

- Virtual node IDs are intended to be stable across bus resets;
- A one-to-one correlation exists between a node's EUI-64 and its virtual node ID;
- So long as a sender originates asynchronous primary packets addressed to a particular virtual node ID with a minimum frequency, the sender may rely on the stability of the EUI-64 correlation it first discovered. If more than the minimum time elapses between references to the same virtual node ID, the sender shall first reread the node's EUI-64 before sending other packets to the node;
- Virtual node IDs do not vary with the path to the node. On a bus with multiple bridge portals, each portal maps the same virtual node ID to the same local node ID;
- The bridge portal immediately adjacent to the source of an asynchronous primary packet transforms the *source_ID* field from a local to a virtual node ID;
- The bridge portal immediately adjacent to the destination of an asynchronous primary packet transforms the *destination_ID* field from a virtual to a local node ID;
- One and only one bridge portal (the *alpha* portal) is responsible to synchronize virtual node ID assignments to the local nodes so that all bridge portals have identical mappings;
- Although senders are not permitted to rely on the EUI-64 identity of a particular virtual node ID unless they periodically refresh their access, the alpha portal does not invalidate virtual node ID(s) if they are not referenced;

- Virtual node IDs are invalidated when their corresponding local node ID is removed from the bus. The virtual node ID shall not be reassigned to another (newly inserted) local node until a minimum time has elapsed.

The remainder of this document describes virtual node IDs in greater detail.

Some definitions

The discussions that follow assume the following definitions:

local node ID: A 16-bit address useable as the *destination_ID* in an asynchronous primary packet so long as both the sender and the recipient(s) are on the same bus. The local node ID is the concatenation of either $3FF_{16}$ and the physical ID or else the most significant ten bits of the NODE_IDS register and the physical ID.

physical ID: A 6-bit address assigned to each node by the self-identification process that follows a bus reset (see IEEE Std 1394-1995).

virtual node ID: A 16-bit address that may be used as the *destination_ID* in an asynchronous primary packet so long as the packet passes through at least one bridge en route to its destination.

virtual ID: A 6-bit address assigned by bridge portal(s) to all of the nodes present on the portal's local bus. For any bus, all the bridge portal(s) share the same mapping from 6-bit physical ID to virtual ID.

Determination of the alpha bridge portal

In order to synchronize the assignment of virtual IDs among all bridge portals on a bus there shall be one dominant (alpha) bridge portal. The identity of the alpha portal shall be determined during the bus self-identification phase.




Figure 1 – Self-ID packet 1 format

The second self-ID packet is modified to redefine one of the reserved bits as the *bridge* bit (abbreviated as *b* in the figure above). This permits the quick identification of all bridge portals on the bus. The alpha portal is the bridge portal with the largest physical ID.


NOTE A consequence of the *bridge* bit is that bridge portals transmit a minimum of two self-ID packets even if they have fewer than four PHY ports implemented. Neither IEEE Std 1394-1995 nor draft standard P1394a mandate the manner in which implemented PHY ports are enumerated although extant PHY designs number their ports monotonically commencing with zero. Software writers should not make any assumptions about the way in which PHY ports are numbered.

Assignment of virtual node IDs

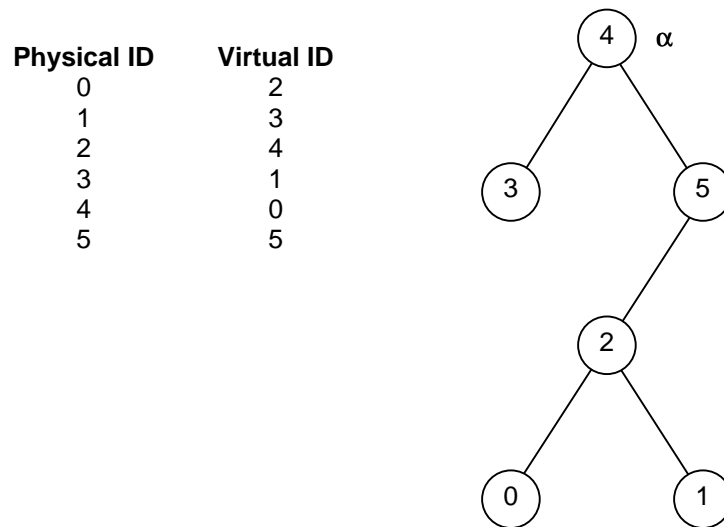
Subsequent to a bus reset, each bridge portal shall attempt to assign virtual node IDs to all nodes on the bus. The process described below is that used by a bridge portal newly

connected to a bus, *i.e.*, there are no preexisting virtual node ID assignments. The process used upon subsequent bus resets is fundamentally similar but is described later.

Virtual ID assignments shall be made from the viewpoint of the alpha portal. The topology of the bus is normalized as *if* the alpha portal is the root of the bus topology. The figure below shows the natural topology of a bus (evidenced by the natural assignment order of the physical IDs) transformed into a topology normalized to the perspective of the alpha bridge portal.



Once this transformation is accomplished, virtual ID zero is assigned to the alpha portal itself. The remaining virtual IDs are assigned during a left-wise traversal of the tree: each leaf node (or node with no remaining unvisited child links) is assigned a virtual ID in monotonically incrementing sequence. For the example already given above, the resultant assignment of virtual IDs is shown below.



Because all bridge portals use the observed self-ID packets as input and all execute the same virtual ID assignment algorithm, all portals are guaranteed to produce identical mappings

from physical ID to virtual ID. As soon as a bridge portal completes this algorithm, the virtual node IDs may be used by remote nodes to access nodes on the local bus.

The second step in the algorithm requires the alpha portal to read configuration ROM at each node to obtain the EUI-64. This information is saved by the alpha portal in a map that correlates physical ID, virtual ID and EUI-64. The information in this map is accessible to any node anywhere in the net of interconnected buses (including the other bridge portals on the local bus). Note that because the alpha portal is guaranteed to have a virtual ID of zero this map is readily available to remote nodes without an enumeration of the bus to locate bridge portal(s).

Removal of a node

If one or more nodes are removed from the bus, a reset occurs and a new set of self-ID packets are generated. The self-ID packets may be analyzed to normalized to describe the new bus topology normalized to the perspective of any node. If this is done, it is easy to determine whether or not any nodes have been added or removed.

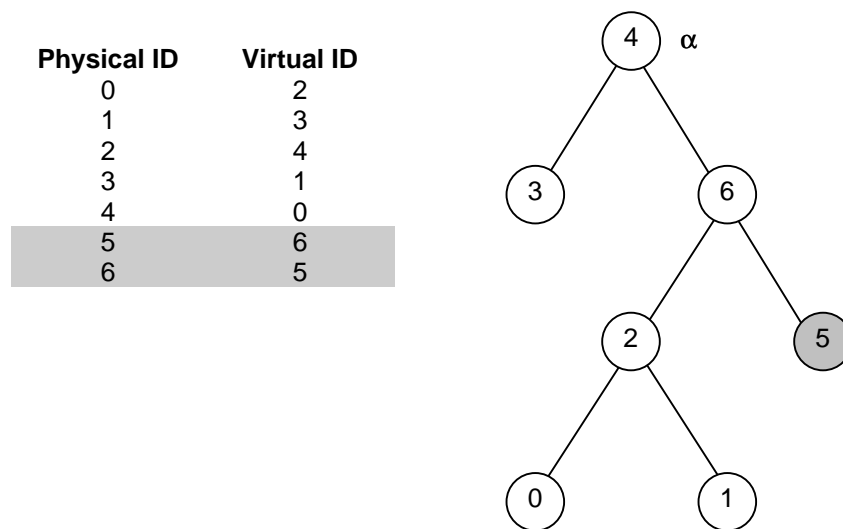
If the topology of the bus is unchanged after a bus reset, there is nothing that the bridge portals need to do: the existing physical ID to virtual ID mapping is likewise unchanged.

If a node has been removed, its virtual node ID is no longer valid. Bridge portals that receive request subactions addressed to an invalidated virtual node ID shall return a response code that indicates the error. Bridge portals that receive response subactions addressed to an invalidated virtual node ID shall discard the response packet.

The bridge portal shall not reassign the virtual node ID to another node until at least $2 * \text{VIRTUAL_ID_REFRESH}$ seconds have elapsed without the receipt of any subaction addressed to the invalidated virtual node ID. This is necessary to provide a guard band for the originators of subactions whose *destination_ID* is equal to the invalidated virtual node ID.

Insertion of a node

Just as for the removal of a node, the insertion of a new node causes a bus reset. If analysis of self-ID packets detects the insertion of one or more node(s), the new node(s) shall each be assigned a virtual ID from the pool of available virtual IDs. The assignment shall use the same algorithm described above; the bus topology shall be normalized to the perspective of the alpha bridge portal and nodes without virtual IDs shall be assigned the smallest available virtual ID according to the traversal order of the tree.



Because invalidated virtual IDs shall not be returned to the pool of available virtual IDs until they have remained unreferenced for $2 * \text{VIRTUAL_ID_REFRESH}$ seconds, it is possible to exhaust the supply of available virtual IDs if many new nodes are inserted on a bus. In this case, some newly inserted nodes may not be accessible from a remote bus until a virtual ID becomes available. The bridge portal indicates this condition by mapping the physical ID of the node to an illegal virtual ID, $3F_{16}$. As virtual IDs are returned to the available pool, the bridge portal shall assign them to any nodes awaiting a virtual ID. Care must be taken that the algorithm executed by the bridge portal(s) assigns the newly freed virtual ID to the same node.

Virtual node ID refresh

A remote node that originates request subactions addressed to a virtual *destination_ID* cannot rely on the identity of the addressed node unless it has periodically refreshed its reference to the virtual node ID. If $\text{VIRTUAL_ID_REFRESH} + \text{REMOTE_TIMEOUT}$ seconds elapse without either a) the successful completion of a transaction addressed to the node or b) the successful completion of a refresh command addressed to the alpha portal on the node's local bus, the subaction originator shall not use virtual node ID as *destination_ID* in a subaction without first verifying that the virtual node ID still addresses the same node.

The correlation between virtual node ID and EUI-64 may be performed in either of two ways. Two quadlet read transactions may be used to obtain the EUI-64 from the addressed node's configuration ROM. Or a query may be sent to the alpha portal for the bus on which the addressed node

Obtaining a physical ID from a virtual node ID

There may be circumstances for which it is necessary to obtain the physical node ID of a remote node. For example, if SBP-2 is in use and the initiator is on one bus and both the target and the data buffer are on another bus, the initiator requires the physical node ID of the data buffer in order to communicate it to the target in data transfer ORBs.

The simplest way to accomplish this is to make the mapping between physical ID, virtual ID and EUI-64 (already maintained by bridge portals for their own purposes) visible to remote nodes.

Other issues

Additional concerns remain to be addressed in a subsequent revision of this document. They are not discussed at present because the author is not certain that consensus has been achieved by the working group. These include:

- * The C language definition of a standard algorithm for the analysis of self-ID packets normalized to the perspective of the alpha portal and the subsequent assignment of virtual IDs to newly inserted nodes;
- * Procedures for the assignment of virtual node IDs when two formerly separate buses are joined by a connection between two nodes that are not bridges;
- * The effects of virtual node IDs on the definition of global broadcast previously agreed by the working group. We may wish to reexamine the need for any global broadcast mechanism whatsoever; and
- * Whether or not a standard format for asynchronous stream packets should be defined that includes *source_ID* in a fixed location. This would permit bridge portals to map this address to a virtual node ID just as is done for asynchronous primary packets. Note that the broadcast nature of stream packets eliminates the need for *destination_ID* in any fixed location.