SPACEWIRE NODES

Session: SpaceWire Standardisation

Long Paper

Martin Suess

ESA / European Space Research and Technology Centre, Noordwijk, The Netherlands

Albert Ferrer

Space Technology Centre, School of Computing, University of Dundee,

Dundee, Scotland, UK

E-mail: martin.suess@esa.int, aferrer@computing.dundee.ac.uk

ABSTRACT

The SpaceWire Standard ECSS-E-50-12C [1] is planned to undergo a revision. The main objective of this revision is to correct errors, remove ambiguities and to include some additional features which have been identified and agreed by the SpaceWire WG.

The title of the Standard [1] is "SpaceWire – Links, Nodes, Routers and Networks". While the Links and the Routers are well defined the experience from different implementations shows that the notion of what is comprised in a Node can differ. In conjunction with the new features that are going to be introduced during the revision of the standard to enable network discovery and PnP, the definition of a node in the SpaceWire network and the features it has to support need to be revised.

1 SPACEWIRE DEFINITIONS

As introduction some definitions of the current SpaceWire standard are reviewed.

1.1 NODE DEFINITION

The current standard [1] defines the nodes in section *10.4 SpaceWire nodes*:

"a. A SpaceWire node shall comprise one or more SpaceWire link interfaces (encoders-decoders) and an interface to the host system.

NOTE A SpaceWire node represents an interface between a SpaceWire network and an application system using the network services.

b. A SpaceWire node shall accept a stream of packets from the host system for transmission or provide a stream of packets to the host system after reception from the SpaceWire link, or do both."

SpaceWire nodes are defined here first of all and as connection point between applications and the network. Further nodes are characterised as sources and destination of packets. It is also clear that a SpaceWire node can comprise more than a single link interface.

1.2 ROUTER DEFINITION

The current standard [1] defines the nodes in chapter *10.2 SpaceWire routing switch*:

"a. A SpaceWire routing switch shall comprise a number of SpaceWire link interfaces (encoder-decoders) and a routing matrix.

NOTE The routing matrix enables the transfer of packets arriving at one link interface to another link interface on the routing switch, and the sending out from this link. Each link interface can be considered as comprising an input port (the link interface receiver) and an output port (the link interface transmitter).

b. A SpaceWire routing switch shall transfer packets from the input port of the switch where the packet arrives, to a particular output port determined by the packet destination address."

This chapter defines also that a routing switch must have an internal port with the address zero to access the configuration logic like e.g. the routing table. As safety measure this configuration port zero must only be accessible using path addressing. The maximum number of physical output ports of a router is limited to 31.

1.3 SPACEWIRE ADDRESSES

In chapter *10.2* it is also defined that packets in SpaceWire networks can be routed based on either path addresses or on logical addresses. A router always deletes the first byte of the path address (range 0 to 32) when routing a packet. When the packet reaches the destination node normally the complete path address has been deleted. A router can optionally also delete a logical addresses byte (range 32 to 254). This is done in gateways between addressing regions for regional addressing and may also be done on the final link to a node. The last point is never the case if the system is also compliant to [2] *ECSS-E-ST50-51C SpaceWire protocol identification*. There the target SpaceWire node always expects a logical address as the first byte in a packet followed by the protocol ID. There it is defined explicitly that one SpaceWire node may support several different logical addresses at the same time. These different logical addresses can for example be used to identify different virtual channels or applications in the node. The logical address 254 is the default logical address which shall be used where the target node has no other value specified or may be used when the logical address is unknown by the sending node.

1.4 TIME DISTRIBUTION

As specified in *8.12 System time distribution*, each node and router must contain one single six-bit time counter. This time counter is updated according to the Time-Codes received through any of the links or by the TICK_IN signal when the node is acting as

the time master in the network. Dependent on the value of this time counter compared to the value in the received time-code the time counter will be incremented and the TICK_OUT signal will be emitted. The consistent handling of the local time is one important aspect that is defining a node.

2 NETWORK DISCOVERY

In the frame of the discussion on Plug & Play for SpaceWire networks techniques for autonomous network discovery have been investigated. Network discovery is when a node can find out about the network topology and the functionality of the connected nodes by probing the network with interrogation packets. A lot of information on the network can be found by reading the configuration status accessible in a router through the configuration port. There the number of links of the router and their activity status should be accessible. These details can then be used to continue to discover the connected nodes and routers and in the end the complete network. One obvious problems is that when an interrogation packet is sent to an undiscovered device it is not known if this device is a node or a router. When trying to access a router with the default logical address 254 the packet will be normally spilled. On the other side, if using the path address 0 when accessing a node the same will happen. It became quickly clear that a unified method to access the configuration information of nodes and routers is needed for efficient network discovery. The currently planned revision of the SpaceWire standard gives the opportunity to introduce such modifications in the standard.

3 SPACEWIRE NODE MODIFICATION TO SUPPORT NETWORK DISCOVERY

After the discussions within the SpaceWire working group it has been decided to leave the definition of the router unchanged. Instead the definition of the SpaceWire node shall be adapted. Most importantly a configuration port which can be accessed using the path address zero should be also introduced as mandatory feature in the nodes to support network discovery. This gives the chance to review more general the definition of SpaceWire nodes and if this definition can be made more precise and needs to be updated. When looking at various SpaceWire implementations it can be perceived that currently different notions of what is a SpaceWire node exist. Sometimes a node is associated with a single logical address or even with an individual interface but often a wider definition is used.

3.1 NODE FEATURES

With the described modification, the concept of node is tied to a single configuration port which can be accessed from all SpaceWire links which belong to this node. In this port zero configuration space, among others, information about all links belonging to the node can be found. Similar important is that the node also contains only one single 6 bit time counter with a single time interface towards the host system. Time codes arriving through different SpaceWire links are handled equally and when the node is acting as the time master the time codes are sent out through all running SpaceWire links of the node.

Both the features are the same in nodes and routers. The main feature which distinguishes nodes from routers is that the nodes are the sources and destination of

packets and that they provide an interface to the host system which is using the network services.

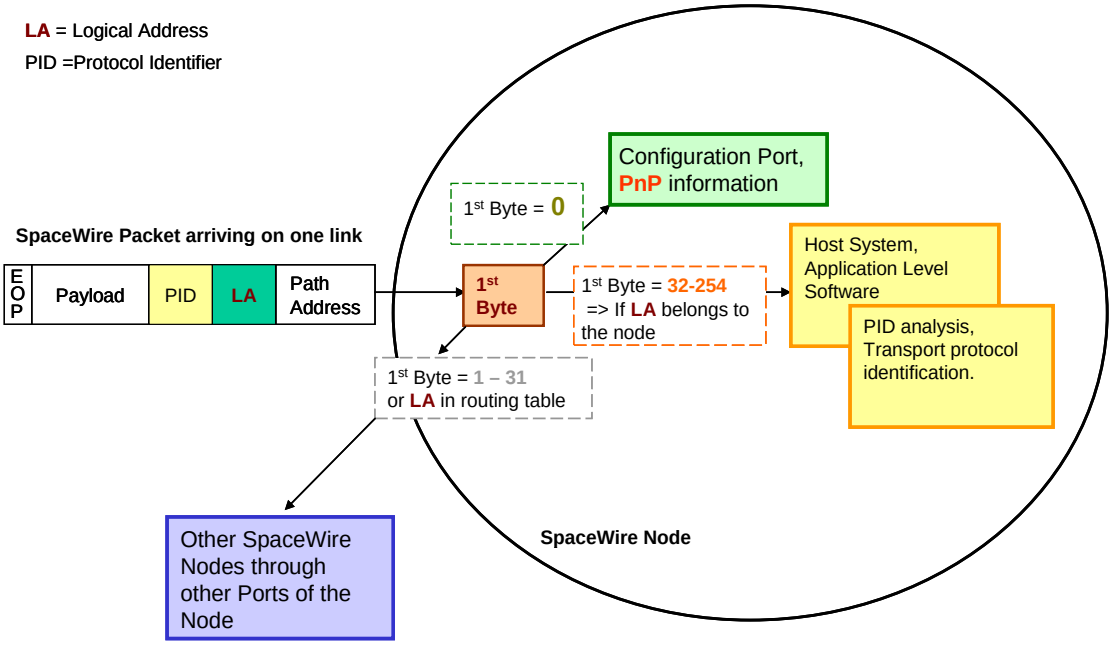

Figure 1: Handling of packet when arriving at a SpaceWire node

3.2 PACKET HANDLING BY THE NODE

The processing of a SpaceWire packet by a node following this definition is shown in Figure 1. The packet may have some leading bytes containing a path address. As specified in [2] this is followed by the logical address and the PID bytes and the payload of the packet. The node will start by analysing the first byte of the packet.

A. If the leading byte is a zero the packet will be routed to the configuration port for processing. The second byte would be expected to be one valid logical address of the node or the default logical address 254. The later is especially the case if a node is to be discovered and the logical address is not yet known by the sending node. The following handling of packet will be made in accordance with the Protocol Identifier (PID), which could for example indicate that it is a RMAP packet, a PnP packet or any other protocol supported by the node. It is important for network discovery that the node remembers the SpaceWire link through which it received the packet addressing the configuration port so that any reply to an interrogation packet is returned through the same link of the node.

B. If the leading byte corresponds to one of the Logical Addresses (LA) of the node the packet is forwarded to the host system. The PID and the rest of the packet may be analysed by hardware or software and may then be provided to the application level software for further processing.

C. One other possibility, which is not explicitly required or excluded by the current SpaceWire standard, is that the packet could be forwarded through on of the other SpaceWire links of the node. This forwarding could be based on a path address or a logical address defined in a routing table.

D. If the first byte does not correspond to any of the options mentioned before the packet is spilled.

4 ROUTING FUNCTION IN A NODE

Whether or not to include the optional routing function described under option C as part of the definition of the SpaceWire node has been controversially discussed during previous SpaceWire working group meetings.

For example the draft SpaceWire-PnP Protocol Definition [3] states that nodes are expected to have no routing function: “packets arriving at any port on a node will be consumed by the node.”

On the other hand there exist already some devices like the SMCS332SpW (AT7911E) which include such a routing function between the SpaceWire ports of the node. Similar, the Golden Gate ASIC developed by BAE [5], which can be used to connect up to four SpaceWire interfaces through a PCI bus to the host processor, also contains a routing function between the SpaceWire ports. There have been also a number of computer boards developed which make use of the SpW-10X router (AT7910E) to interface to the SpaceWire network. The SpW-10X provides two external ports that are effectively FIFO interfaces to inject and retrieve SpaceWire packets into and form the network. These examples make clear that nodes with integrated routing function are a concept which is actually widely used.

During a discussion it was proposed that these cases could be regarded as a node being attached to a router. Conceptually this could establish again the clear distinction between the routing and the network access point function in the Space Wire network. But as this connection is part of a SpaceWire network there should be one or several SpaceWire links between the router and this node. This is certainly not the case in the examples provided above and the reason is that implementing such a very short SpaceWire link is inefficient when the connection has to be made on a board or even in a chip. In addition it would also require a duplication of the configuration port zero. Conceptually this may be even be welcome but it will result in additional implementation effort. More significantly this separation of the node and the router function would also require a duplication of the local time counter. One belonging to the router and one belonging to the node. If the router and the node are attached to each other this duplication does not make sense.

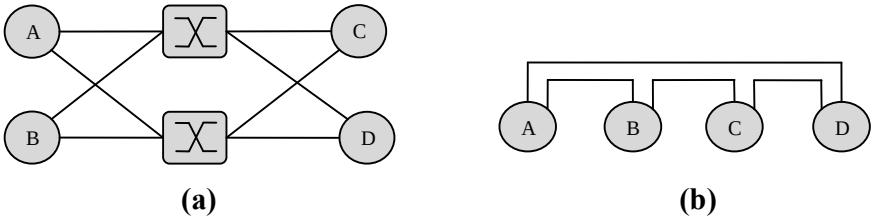

Figure 2: Example topologies to interconnect four nodes with redundancy, (a) Routing-centric topology with a redundant switch, (b) Ring topology using nodes with routing capabilities

Furthermore, the routing capability within a node allows useful network topologies as shown in Figure 2 (b). In some scenarios, a ring topology meets the requirements in

terms of bandwidth and redundancy without requiring external routing devices. The topology shown in Figure 2 (a) requires more harness, the powering of more devices and links that may not provide any advantage when for instance a simple chain of sensors is considered. On the contrary, the extra devices increase the complexity of failure cases and error recovery mechanisms. SpaceWire should not be constrained to certain network topologies and exclude other technologies which are widely used.

Finally it needs to mentioned that the presented node definition does allow the implementation of simple nodes with only a single link, nodes with several links without routing or nodes with several links with routing only between certain links. All these cases are a possible subset of the wider definition.

5 CONCLUSION

In this paper discussed the revised definition of a SpaceWire node that is needed to enable an efficient way of network discovery as part of PnP. For an interrogating packet the view of a node should be aligned with the one of a router by a mandatory introduction of a configuration port zero in every PnP enabled node as it is required in every router. The border of a node should be defined similar to the one of a router. All SpaceWire ports belonging to one node can access the same single configuration port zero and the time codes received through any link act on the same single time counter. The main distinction between nodes and routers is that a nodes provide an interface to the host system which is using the network service.

A node according to this definition may be also capable to forward a packet which it has received on one SpaceWire port through any other SpaceWire port of the node based on the physical or logical address contained in the first byte. As described this functionality is already implemented in a number of chips designed to provide SpaceWire interfaces for nodes.

It is suggested that this revised definition of SpaceWire nodes will be included in the next revision of the SpaceWire standard and that it will be reflected in the definition of the SpaceWire PnP protocol.

6 REFERENCES

1. “SpaceWire: Links, nodes, routers and networks”, ECSS-E-ST-50-12C, 31 July 2008.

2. “SpaceWire protocol identification”, ECSS-E-ST-50-51C, 5 February 2010.

3. Peter Mendham, “SpaceWire-PnP Protocol Definition”, Draft A Issue 2.1, 16 September 2009.

4. “SMCS332SpW User Manual”, SMCS_ASTD_UM_100, Issue: 1.5, 10 July 2007.

5. J. Marshall, “Evolution and Applications of System on a chip SpaceWire Components for Spaceborne Missions”, Proceedings International SpaceWire Conference 2008, 4-6 November 2008

6. SpW-10X SpaceWire Router User Manual”, UoD_SpW-10X_UserManual, Issue 3.3, 30 April 2008