

SPACEWIRE NETWORK TOPOLOGIES

Session: Networks and Protocols

Short Paper

Dr Barry M Cook, C Paul H Walker ...

4Links Limited, Bletchley Park, Milton Keynes MK3 6EB, England

E-mail: [Barry, Paul] @4Links.co.uk

ABSTRACT

In contrast to other ‘bus’ standards, SpaceWire offers unprecedented flexibility in the choice of network topology. Choice can be made to optimize parameters such as performance, harness mass, cost, or fault-tolerance, or to provide an optimum balance of these, for a particular mission. This paper discusses some simple topologies and how they affect these parameters. It suggests hybrid topologies that might provide an optimum balance for some missions. It points to work on graph theory and on IEEE 1355 that may give useful design insights. The subject is large enough to warrant a course lasting several days, and in a short paper we can only scratch the surface. We hope that this will inform SpaceWire users and encourage them to find out more.

1 TRADE-OFFS: WHAT IS THE TOPOLOGY TRYING TO ACHIEVE?

Each mission has its own set of goals, and one of SpaceWire’s [1] great benefits is to give the mission a wide range of topologies that provide a balance between a wide range of requirements. Inevitably, there are trade-offs between the different requirements.

Parameters that may influence, or be influenced by, the choice of topology are:


- The mass of the cable harness;
- The performance required: sometimes this will be different for different parts of the spacecraft, sometimes it is purely bandwidth, sometimes latency or delay through the network;
- The extent of fault-tolerance required, which may be different for different subsystems on the satellite.
- The power consumed by the network
- The cost of the components to build the network
- The lead time required for the satellite: can it be assembled from off-the-shelf building blocks, or is there time and budget to develop something new?

In considering the variety of topologies, we will suggest the effect that the topology might have on some of these parameters (particularly performance, harness mass and fault-tolerance). Of course topology is not the only design aspect to affect these parameters and harness mass, for example, might be reduced in networks where traffic is mostly unidirectional if a half-duplex version of SpaceWire [2] were used.


2 INTRODUCTION TO DIFFERENT NETWORK TOPOLOGIES

2.1 POINT-TO-POINT LINKS

Perhaps the simplest network is a single link between two nodes. Several spacecraft are flying with a number of such connections to make an ad-hoc network, perhaps more with the IEEE 1355-based components that have been available for some years.





Of course there is, in practice, a (hidden) network between the disjoint nodes, because in almost all cases they will need to be controlled, and will need to communicate with the TM/TC down-link. The hidden network is shown in the figure here as a cloud. At first sight, it may seem easier to build the system with the hidden network, but the same tradeoffs need to be made and a full SpaceWire network has benefits.

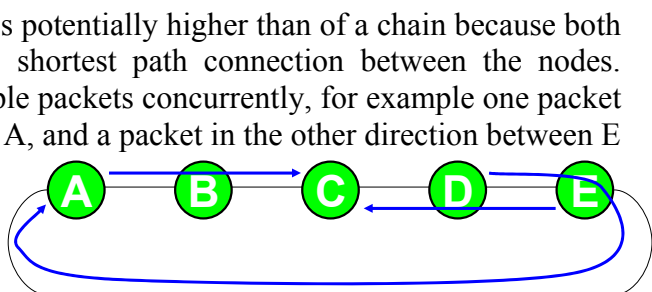



2.2 CHAINS AND RINGS

We can extend the point-to-point link by adding another node, and then another, etc., to make a chain of links, with small routing switches in each element of the chain. If connections are made between physically adjacent nodes, the chain offers the lowest harness mass of all the networks, but in other respects it is less than ideal. A packet from A to E will potentially experience contention and delay at all the intermediate nodes, limiting bandwidth and increasing latency. If a single fault occurs, such as in node C, the network becomes split into two networks, A,B and C,D which can not communicate.



The susceptibility to single points of failure can be overcome by connecting up the ends of the chain to make a ring. With this, any fault in a link or a node turns the network into a chain. Performance of the ring is potentially higher than of a chain because both directions can be used to make the shortest path connection between the nodes. Furthermore, the ring can carry multiple packets concurrently, for example one packet between A and C, one between D and A, and a packet in the other direction between E and C. Cable harness mass stays low if all the connections are made between nodes that are physically close to each other.




Strong claims are made of some ring-based networks such as Scalable Coherent Interconnect (SCI) and IEEE 1393 about 'Spatial Re-use', with several packets being transferred concurrently. It should be noted that SpaceWire provides the full benefit of Spatial Re-use while also allowing additional topologies beside the ring.


2.3 MULTI-DIMENSIONAL CHAINS AND RINGS, OR GRIDS AND TOROIDS

The chain can be extended into two or more dimensions to form a grid. This uses more ports on each node, and increases harness mass, to provide increased performance and fault-tolerance.


The multi-dimensional ring (or toroid) forms a loop on each row and column, which for a small additional cost and an increase in harness mass can provide more fault-tolerance and performance than the grid.



2.4 CENTRALIZED SWITCHES, CONCENTRATORS AND TREES




A typical application for a SpaceWire network is to connect a set of instruments to mass memory and processing, and a convenient way to do this is with a centralized routing switch, perhaps in the mass memory or processor. In many cases, a simplified routing switch which just concentrates the traffic is adequate, particularly if the instruments do not need to communicate with each other.



A tree can be built with a hierarchy of routing switches or concentrators. By placing the intermediate devices near their nodes, this can result in a substantial reduction of harness mass. If the leaf nodes need to communicate with each other, care needs to be taken that the root switch does not become a bottleneck. If the root node and the intermediate nodes are just concentrators, with no communication between the leaf nodes, then the root node is unlikely to be such a bottleneck.

The bottleneck in the case of routing being required between leaf nodes can be overcome by building a 'fat tree' which preserves bandwidth at all levels of the hierarchy. This also provides a means of building a routing switch with twice as many ports as are available on a single chip. In the sketch here, six eight-port switches provide a single 16-port switch.




The cost of needing six eight-port switches to make a single full-bandwidth 16-port switch makes a strong argument for building switches with as many ports as possible. But, as we shall consider next, a single large switch is a large single point of failure.


3 FAULT-TOLERANCE CONSIDERATIONS FOR CENTRALIZED SWITCHES

The centralized switch or concentrator has many attractions. It is simple; each node is dependent only on the switch/concentrator and does not have to handle traffic from other nodes. If one node fails, none of the other nodes is affected, so fault-tolerance with respect to the nodes is good. On the other hand, if a single central point fails, it causes the whole mission to fail. And the cable lengths and hence harness mass is considerably higher with all the cables having to come together to one place rather than just go to nearest neighbours.

We can improve fault-tolerance by adding a cold-redundant switch or concentrator, at the cost of approximately doubling harness mass.




Many missions also need redundant instruments as well as redundant switches, and the harness mass increases again. Effectively this topology provides a 2x2 crossconnect between each node (both nominal and redundant) and each routing switch.



4 BALANCING THE TRADE-OFF, HYBRID NETWORKS

It has been seen that rings provide low harness mass at the expense of performance and centralised switches provide performance at a major cost in harness mass. Is it possible to mix the ring topology (where performance is less critical) with the centralised switch (where performance is needed)? SpaceWire does indeed provide the opportunity for such mixed or hybrid networks.



In the sketch here, the left nodes represent processors or mass memories or high-bandwidth instruments that need the dedicated connection to each routing switch. The remaining nodes do not need such high performance and so can be connected in a ring, with taps from the ring going to the switches. The number of SpaceWire links is not changed from the previous example but eight long connections to the central switches have been replaced by eight much shorter connections to make the rings.



Many variations on this theme are possible and a trade-off study is likely to show for a particular mission that an alternative is preferable. In this example, we have chosen different connection techniques to balance between performance and harness mass, another balance might be with other parameters such as fault-tolerance, where some parts of the satellite may be more critical than others.

4.1 FOCUS ON LEAD-TIME, PNPSAT


PnPSat [3] has been designed as a kit of parts that can be built and launched within a few days of the mission requirement. They construct the satellite out of standard panels, each containing a routing switch which has local connections to nodes that are bolted to the panel, together with backbone connections to other panels. The initial network chosen is a hybrid network, with a single switch per panel and a ring to make the backbone between the panels, shown notionally here.



5 FURTHER STUDY

5.1 GRAPH THEORY


There is a large volume of work on graph theory, much of it targeted on finding networks that, for a given number of ports per switch, give the maximum number of nodes for a given maximum path length between nodes. There is a ‘Moore bound’ that sets the upper limit on the possible number of nodes for a given number of ports and path length (called degree and diameter respectively in graph theory papers). Reference [4] is one of many examples of papers that consider the Moore bound and list networks that come closest to meeting the bound. Very few networks actually meet the Moore bound — one that does is the classic Petersen graph, discovered in the 19th century, which achieves ten nodes with three ports per node and with a maximum path-length of two between any nodes. Compare this with a cube, which has eight nodes, again has three ports per node, but has a path length of three between nodes on opposite corners — fewer nodes than Petersen yet a longer path length — but possibly easier to build physically.



5.2 THE NETWORK DESIGNER’S HANDBOOK

The development work on IEEE 1355 [5] produced a wide-ranging study of different network topologies and their relative performance in terms of throughput and latency. The results were published in ‘The Network Designer’s Handbook’ [6].


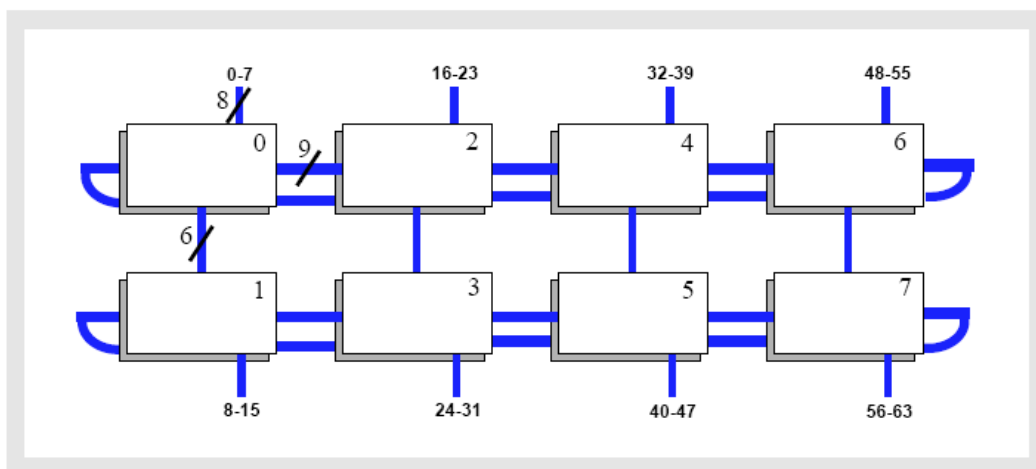
The results do not carry across to SpaceWire directly, because the work was based on the INMOS C104 routing switch which had 32 ports and used a different routing algorithm from SpaceWire’s path and logical addressing. Furthermore, the networks considered were all regular



arrays, and many SpaceWire networks are, for good reasons, not such regular arrays. Also, the traffic in the simulations was assumed to be random, or had sequences of every node communicating in turn with every other node, whereas traffic in a SpaceWire network is likely to be more predictable.

Nevertheless, study of the results yields insights that could well benefit SpaceWire users. They might also indicate where projects might benefit from network modelling by software such as OpNet (as has been done by a small number of SpaceWire users), or with hardware [7], as was done for IEEE 1355.

An example set of results from The Network Designer's Handbook is reproduced here for a two dimensional ring or torus which acts as a backbone between eight ports per switch that go to local leaf-nodes, so that it connects a total of 64 leaf nodes. The extracts shown here are taken from a draft of the book and appear on pages 78 and 79 of the published version.



6 CONCLUSIONS

SpaceWire provides unprecedented flexibility in network topology. Two extremes are (daisy) chains, where most traffic has to visit several nodes before arriving at the destination node, and large central switches, where all traffic goes via a single switch.

A number of missions have chosen some form of hybrid topology, comprising elements of rings (daisy-chains with the loop closed) and of central switches. One example is PnPSat which uses a local switch on each panel to connect to the nodes on the panel, and then uses a ring as a backbone between these switches.

By studying some of the topologies that are possible for SpaceWire, perhaps together with academic results from graph theory, the industry may gain a more complete understanding of the effect of topological choices on missions. Similarly, results obtained from simulations of IEEE 1355 networks may also give useful insights.

SpaceWire's topological flexibility brings great new opportunities for improvements in all of the parameters such as performance, mass, cost, and reliability. Increased awareness of the effects of topology on these parameters could have a significant effect on mission lead-time, cost-effectiveness, and success.

REFERENCES

In the following references, * indicates a paper that is being presented at this conference.

[1] The ECSS-E-50-12 Working Group, “**ECSS-E-50-12A 24 January 2003, SpaceWire - Links, nodes, routers and networks**”, published by the ECSS Secretariat, ESA-ESTEC, Requirements & Standards Division, Noordwijk, The Netherlands

[2] Barry M Cook, Paul Walker, “**Asymmetric SpaceWire**”, to be published, abstract available from 4Links

[3]* Donald Fronterhouse, “**PnPSAT**”

[4] Michael J Dinneen, Paul R Hafner, “**New Results for the Degree/Diameter Problem**”, accessed from http://arxiv.org/PS_cache/math/pdf/9504/9504214v1.pdf on 28 August 2007

[5] Bus Architecture Standards Committee of the IEEE Computer Society, “**IEEE Std 1355-1995, IEEE Standard for Heterogeneous InterConnect (HIC) (Low-Cost, Low-Latency Scalable Serial Interconnect for Parallel System Construction)**” IEEE, 1995

[6] P. Thompson, A.M. Jones, N.J. Davies, M.A. Firth and C.J. Wright (Eds.) “**The Network Designer's Handbook**”, Volume 51: Concurrent Systems Engineering Series, 1997, 309 pp., softcover, EOS Press, ISBN: 978-90-5199-380-6

[7] S Haas, D A Thornley, M Zhu, R W Dobinson and B Martin, “**The Macramé 1024-Node Switching Network**” Microprocessors and Microsystems, IEEE 1355 Special Issue, Ed. Paul Walker, V21, Nos 7,8, 30 March 1998, (this paper is also summarized in Chapter 8.2 of The Network Designer's Handbook)