A Real-time On-Chip Network Architecture for Mixed Criticality Aerospace Systems

Majumder, Shibarchi; Nielsen, Jens Frederik Dalsgaard; la Cour-Harbo, Anders; Schiøler, Henrik; Bak, Thomas

Published in:
The Aeronautical Journal

DOI (link to publication from Publisher):
[10.1017/aer.2019.80](https://doi.org/10.1017/aer.2019.80)

Creative Commons License
Unspecified

Publication date:
2019

Document Version
Accepted author manuscript, peer reviewed version

[Link to publication from Aalborg University](#)

Citation for published version (APA):
Majumder, S., Nielsen, J. F. D., la Cour-Harbo, A., Schiøler, H., & Bak, T. (2019). A Real-time On-Chip Network Architecture for Mixed Criticality Aerospace Systems. *The Aeronautical Journal*, 123(1269), 1788-1806.
<https://doi.org/10.1017/aer.2019.80>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

A Real-time On-Chip Network Architecture for Mixed Criticality Aerospace Systems

Shibarchi Majumder, Jens F. Dalsgaard Nielsen, Anders la Cour-Harbo, Henrik Schiøler, Thomas Bak*

Department of Electronic Systems, Aalborg University
Aalborg, Denmark

ABSTRACT

Integrated Modular Avionics enables applications of different criticality-levels to share the same hardware platform with an established temporal and spatial isolation. On-chip communication systems for such platforms must support different bandwidth and latency requirements of applications while preserving time predictability. In this paper, our concern is a time-predictable on-chip network architecture for targeting applications in mixed-criticality aerospace systems. The proposed architecture introduces a mixed, priority-based and time-division-multiplexed arbitration scheme to accommodate different bandwidth and latency in the same network while preserving worst-case time predictability for end-to-end communication without packet-loss. Furthermore, as isolation of erroneous transmission by a faulty application is a key aspect of contingency management, the communication system should support isolation mechanisms to prevent interference. For this reason, a sampling port and isolated sampling buffer-based approach is proposed with a transmission authorization control mechanism, guaranteeing spatial and temporal isolation between communicating systems.

Keywords:

Mixed-Criticality System, Network On-chip, Real-time System, Embedded System, On-Chip Communication, Integrated Modular Avionics

* This research is funded by Independent Research Foundation Denmark under grant number 6111-00363B. Email: sm@es.aau.dk, jdn@es.aau.dk, alc@es.aau.dk, henrik@es.aau.dk, tba@es.aau.dk

1.0 Introduction

A modular aerospace system, where multiple applications of different criticality and certification assurance level are integrated on a shared computational resource, require analyzable, deterministic and hard-real-time end-to-end communication for certification as well as safety purposes.

Systems of different criticality levels can have different timing requirements, for example a *flight control system*, an application of Design Assurance Level A (DAL A)⁽¹⁾, i.e. the highest level of criticality in aviation standards, has hard-real-time timing requirements where the timing margins are less than 10 milliseconds⁽²⁾, whereas a multi-media entertainment system, a DAL E application, does not have any such real-time requirements at all. The on-chip communication system should have the capability of prioritizing critical applications to meet the hard-real-time requirements and eliminate the need for additional on-chip communication systems for soft-real-time requirements.

Communication between applications in an aerospace system is built upon the concept of *sampling ports*, where a fresh data packet overwrites an older data packet in a single packet buffer and the receiving application can read it single or multiple times⁽³⁾. Each subsystem may produce/consume single or multiple *channel(s)*, process variable of a specific data from IO devices or other subsystems, once per computational cycle. However, frequent communication between the applications and certain intellectual property cores (IPcores) (e.g. memory blocks or hardware-accelerators) can be expected and requires a higher communication bandwidth.

A faulty communication in the airborne system, as identified in⁽²⁾, can be detected by the consumer application with a data validation technique, except for delivery of a message to a wrong recipient. Such an error can be caused by a faulty producer application or a faulty communication system, and some protection technique is essential to identify the source of incoming messages and guarantee the authenticity of received data packets.

The recent developments in on-chip communications are primarily focused on *Network-on-Chip* (NoC) architecture and the development is driven by general purpose computation needs and focused on efficient utilization of network resources and best over-all performance⁽⁴⁾ and often neglects the time-predictability aspects.

In this work, we present an on-chip network targeting application in aerospace systems-on-chip. We propose a mixed, *priority-based* and *time-division-multiplexed* (TDM) arbitration to support different bandwidth and latency requirements of mixed-criticality systems on the same network with additional data protection and isolation mechanism for safe and time analyzable end-to-end communication.

The specific contributions of this work include-

- A real-time on-chip communication network architecture to accommodate different bandwidth and latency in the same network.
- An arbitration mechanism to support different bandwidth and latency requirements with time-analyzable end-to-end communication without packet-loss.
- A configurable isolation mechanism to prevent interference from erroneous transmission and hardware-level protection mechanism against unauthorized commu-

nication.

2.0 Background and Related Work

2.1 Mixed Criticality in Aerospace System

In Integrated Modular Avionics (IMA)⁽⁵⁾, several systems and subsystems of different criticality-levels and functionalities are integrated on one hardware platform. Resource sharing and robust partitioning are the key concepts for such an implementation, where each partition is allocated a set of *spatial* resources and a mechanism in the platform that provides spatial segregation between them. The *temporal* isolation is established by allocating resources to the partitions at specific time slots and preventing access outside the time slot assigned to it. Hardware architecture of a typical IMA platform can consists of a set of computing processing modules that are grouped into clusters so that each group is connected to the same ARINC 664⁽⁶⁾ switch. Related systems and subsystems are implemented in the same group for low latency communication over the same switch⁽⁷⁾. Recent advancement in microprocessor technology provides a *many-core processor*, isolated from each other, where the IMA architecture can be implemented⁽⁸⁾ with an on-chip communication network for inter-system communication^(9,10). Such *single-core-equivalent-multicore* system, mutually isolated processors with dedicated resources, avails isolation between software running on separate processors and the requirement of isolation comes to the NoC for overall isolation in the system.

2.2 Network on Chip

The use of NoC in a real-time system imposes complex constraints in the overall design⁽¹¹⁾.

Xpipes⁽¹²⁾ is a NoC where the network is tailored to meet the bandwidth requirements at its design stage. Such a system could be hard to implement as foreseeing the exact communication load is difficult to analyze and it affects the scope of any future modification of the system. A *circuit switching* method is applied in SoCBUS⁽¹³⁾ and a concept called packet connected circuit was introduced, where a data packet is switched through a dynamic minimum route locking the circuit as it moves. This type of communication is effective where the traffic follows a fixed rule, but not effective where the data is not patterned like in avionics system where data sequence depends upon the relative state of the applications. In⁽¹⁴⁾, an alternative solution is proposed based on *backtrack probing* to avoid waiting for blocked channels to become available, seeking for alternative non-minimal routes. A synchronous circuit switching NoC is presented in⁽¹⁵⁾, a concept of spatial division multiplexing is introduced, where the lane is divided to provide physical separation between data streams.

2.2.1 Priority

A connection-less packet-switching approach is demonstrated in⁽¹⁶⁾, where the routers work independently manner and a *wormhole switching* technique is typically used. The flows are prioritized based on some fixed manner and flow with the highest priority is

given preference. The draw-back of such a design is that packets with low priorities may be dropped or stalled for a long time and has a longer latency. In⁽¹⁷⁾, the authors propose a low end-to-end latency with a *guaranteed service* traffic. In^(18,19,20), the authors address the low priority packet block problem in connection-less NoC by introducing the concept of increasing priority over waiting time. In contrast, this work offers a mixed, *best-effort* and *guaranteed-service* traffic where flow with highest priority is given preference by allocating more bandwidth while flow with lower priority is given the minimum bandwidth allocated by the system designer to maintain worst-case-time analyzable communication.

2.2.2 Time Division Multiplexing

Time Division Multiplexing (TDM) is an arbitration scheme where a resource is shared between channels in the time domain; only one channel is given access to the resources to transmit for a fixed interval of time, called *slots*.

The concept of time division multiplexing is used in⁽²¹⁾ and^(22,23,24) where the resource is allocated to channels based on time slots as an alternative to circuit switching. In⁽²⁵⁾, a *globally-asynchronous-locally-synchronous* NoC has been illustrated for real-time application with *mesosynchronous* routers. The implementation uses wormhole switching technique with TDM to prevent stall and dead-lock and provides solution in terms of no buffering, arbitration, real-time operation and no packet-loss. However, protection mechanism is not addressed in this work which focuses on Worst-Case-Execution-Time (WCET) communication.

2.2.3 Related topologies

A *star* topology, where multiple ends are connected at a single point, can furnish a single cycle end-to-end *flit* transfer with effective control and monitoring capabilities at the cost of restricted communication between ends, where only one end can transmit at any given time. Such a topology offers an efficient solution for one-to-many and low latency communication and supports easy implementation of TDM or cyclic access to each end. As multiple ends are connected to a single point (one router), the packet-routing is simple, and determinism is easy to achieve. Moreover, as all *flits* are routed through one single central router, the communication in the network can be easily monitored. Similarly, any subsystems can be isolated from the network without affecting other subsystems by restricting its access to transmit from the central router. However, the waiting time for transmission from a transmitter is linearly dependent upon the total number of transmitters in the network and can be long when a large number of communicating ends are connected.

On the other hand, in a *tree* topology, communication between the closer ends can be very fast as the *flits* need to hop through just one or a very few node(s). However, communication delay between two ends situated at two far ends of the network can have high latency as the communication gets bottle-necked at the top node.

3.0 Architecture

In this work, a hybrid of star and tree topology has been considered and this section explicitly addresses the architecture, the architectural benefits of the mixed topology approach and micro-architecture of the network components.

Figure 1: The proposed NoC architecture with 4 routers and 12 network-interfaces (NI). Communication flow between two NIs is highlighted.

3.1 Overall Architecture

The network is built around a *hub*, interfaced with multiple routers in the network in a *star* topology and each router is attached to a single or multiple *network-interfaces* in a *reverse fat-tree* topology as shown in Figure 1. An end-to-end data packet propagation from a producer to a consumer through the network components is shown in Figure 2. Under circumstances where one or multiple router dysfunctions, such an architecture allows the operation in the rest of the network to be invariant.

Instead of conventional FIFO buffers, dedicated sampling buffers are used to provide isolation to each channel. In cases of violation of transmitting agreement i.e. maximum allowed bandwidth, only the associated sampling buffer gets overwritten (dropping of old data packets of the violating channel), the communication in the network and other data channels remain unaffected.

The *phit* size, physical channel width, is equal to *flit* size in this network, thus each *flit* can hop in a single cycle when access is provided.

Figure 2: Block diagram showing end-to-end data packet flow.

3.2 Network Interface

The *Network-Interface* (NI) in a NoC has a critical role in implementing end-to-end communication between two nodes. Figure 2 shows an example of data flow from a producer (Processor 1) to a consumer (Processor 2) via associated NIs. This section addresses the overall architecture and functionality of a NI in the proposed architecture.

Figure 3: Microarchitecture of the *network-interface*. The dashed-lines represents configuration mode operations.

Each NI has two ends, a front-end and a back-end, interfacing with the communicating end and the router respectively. A NI is connected to the router with separate transmission and reception lines for simultaneous *tx-rx* operation and interfaced with the communicating end (i.e. a producer/consumer) with standard *memory-mapping* technique. Additionally, each NI has a sampling transmission (Tx) buffer, a transmission channel index buffer and dedicated (Rx) sampling buffers for each receiving channel as shown in Figure 3.

A NI can handle a fixed number of channels and one or multiple NIs can be connected to a producer or consumer depending upon the requirement of number of channels. To send a data packet to a destination NI, a producer writes the data in the transmission buffer with the channel id of the data packet, with a standard memory writing method. Each channel has a configurable *destination address* stored in the NI, that can be configured and re-configured by the producer before starting the network by writing in the control registers. The data packets are transmitted to the associated router. Each NI has a static identification number and each NI in the network can be uniquely identified by a combination of associated router identification number and NI identification number, used as a unique destination address for transmission.

A fresh data packet written in the NI transmission buffer is sent to the connected router, concatenated with the destination address and channel id in its header. There could be application specific needs where the producer repeatedly sends exact same data packets to the consumer; To identify the reception of a fresh data packet from the NI, at the router, a single bit signal line (NI to router) is toggled by the NI on every transmission. This mechanism has additional protective benefits that are explained later.

At the beginning of a reception (data flow from a router to NI), the associated router sets a single bit state signal to *active* and the NI starts listening to the reception channel. On successful reception, the NI validates the received message by checking the source address, in the header of the incoming data packet. Like destination address, each NI has configurable expected-source-addresses (address of the producers) for each receiving channel; The data packet is saved in the sampling buffer dedicated to the channel only if the source address matches with the expected-source-address, otherwise discarded.

3.3 Router

The routers in this network operate in a fixed routing scheme without any routing algorithm. Each router has separate transmission and reception line to interface with the *hub* with two n -bit lines indicating *access-request* and *access-grant* status as shown in Figure 4, where n is the number of NIs connected to the router.

Each router has dedicated sampling buffers for each channel from each NI to guarantee isolation. The sampling buffers holds two *flits* (one data packet) with a 8-bit destination address.

Once a fresh data packet is received from a NI, the router raises a transmission request by setting the associated bit in the request line *high*. The data packet received from the NIs are stored in their associated sampling buffers unless the router gets transmission access. Once access is gained, the router transmits data packet from the sampling buffer in three *flits*; i.e. one header *flit*, followed by two payload *flits*. The

Figure 4: Microarchitecture of the router in a 3 NI configuration.

router adds a source address in the header *flit* next to the destination address.

There is no dedicated buffer for reception operation (hub to router); instead, the router packs the two payload *flits* with the source address and sends to the destination NI (refer Figure 8). Each router has a fixed and unique id, so that a router-NI id can be uniquely identified in the network.

3.4 Hub

The hub is the central and most critical component of the proposed architecture and controls arbitration. This section explains the micro-architecture of the hub.

The hub has separate transmission and reception channels for each router connected over a *cross-bar* (X-bar) as shown in Figure 5. The hub is memory-less and all the routing performed in the hub is *atomic*.

Furthermore, the hub has separate n -bit request and access lines for each router connected, where n is the number of NIs connected to each router. The hub provides transmission access to each router for a specific NI when requested by setting the associated bit *high* in the access line in a priority-TDM arbitration scheme explained in the next section. The hub enables the Rx data line only from the router with transmission access. An erroneous transmission from a faulty router outside its access time gets discarded at the hub. Once the access is provided to a router, the router starts transmitting and the hub checks for the destination router address in the header

Figure 5: Microarchitecture of the hub in 4 routers configuration. (a) transmission request lines from routers. (b) access lines to router (c) input-data lines (d) output-data lines (e) active transmission lines

flit and activates the circuit to the destination router in the X-bar. The path is locked until the last *flit* of the packet propagates through it i.e. the second payload *flit*. Each transmission line to the routers has a single bit transmission-state line that is held *high* by the hub during an active transmission to the destination router.

If the hub reads a predefined destination address (e.g. 1111 1111, which is not a valid destination address in this four router configuration), the hub broadcasts the packet to all the routers in the network.

4.0 Arbitration

In this section, we will discuss the conceptual aspects of the proposed arbitration and a generic way of implementation without concentrating on the specifics of the actual implementation in the NoC.

The goal of the arbitration is to accommodate a priority-based scheme with different bandwidth and different latency allocation to each communicating node, guarantying end-to-end time-deterministic communication without any packet loss. To accomplish, a mixed concept of TDM and priority-token-passing scheme is proposed.

Assume that, a number of producers are connected to a central node that handles the arbitration by controlling the transmission line from each producer. Each producer can have different bandwidth and different latency requirements. However, the whole concept is based on the assumption that the size of the data packet is defined and identical for all the messages. Each producer is assigned a single or multiple slots in a TDM cycle based on its bandwidth requirements and each slot in the TDM cycle has the same length as the transmission time of a data packet, to assure completion of undergoing data packet transmission i.e. if one data packet is unpacked in n *flits* and one *flit* transfer is m clock cycles long, then the slot length in the TDM cycle is $n \times m$ clock cycles. A higher bandwidth requirement of a producer is addressed by assigning a higher number of TDM slots to the producer, where a low latency requirement is addressed by assigning multiple slots at multiple intervals in the TDM cycle as shown in Figure 6.

The Figure represents a hypothetical case where access is provided to four producers by a TDM cycle of 16 slots; Assume that, producer 1 has highest priority, low bandwidth and low latency requirements, producer 4 has the lowest priority and high bandwidth requirement, where producers 2 and 3 have moderate bandwidth requirement but the priority of producer 3 is higher than producer 2. The assumption taken into consideration is not random and a relation to a practical scenario is drawn later in this section. The low latency requirement of producer 1 is accommodated by assigning multiple slots at multiple intervals to ensure that worst possible waiting time to get access is small. The higher bandwidth requirement of producer 4 is full-filled by assigning multiple slots. Producer 2 and 3 are assigned slots as per bandwidth requirements.

However, in a practical implementation, multiple slots assigned to each producer to meet its latency requirements are not always used, and such a guaranteed service traffic is not efficient in terms of resource utilization. For example, producer 1 only uses one of the multiple slots assigned to it to guarantee low latency transmission; additionally, as the TDM cycle of the network is often much shorter than the computational cycle of the communicating nodes, all of the producers do not transmit at every TDM cycle.

A priority-based token passing scheme in addition to the TDM schedule offers better resource utilization, where transmission access is given to the producers based on a concept of *dynamic priority*. The transmission access priority of a producer is determined by a prior knowledge of a priority assigned to each producer by the system designer and the slot the producer is competing for. All the fresh transmission requests are evaluated for the next slot and the ongoing transmission is never interrupted to prevent unfinished or broken data packets at the producing or consuming end. Un-served accesses requests are re-considered for the subsequent slot unless it gets transmission access. A producer drops the access request when all the associated data packets are transmitted.

Each producer has the highest dynamic priority at the slot(s) initially assigned to it in the TDM cycle and definitely gets transmission access irrespective of the access requests from other high priority producers. The producer with the highest priority gets the transmission access when competing for a free/ unused slot initially assigned to another producer.

This TDM-dynamic-priority scheme is elaborated in Figure 7 with the same four producer scenario considered earlier. The arrows show the transmission by each pro-

Figure 6: The figure shows allocated slots to different producers to meet latency and bandwidth requirements.

ducer where the thin lines represent the time in the TDM cycle when the transmission request is received from each producer. At the beginning, producer 2 and 3 competes for slot 1 (which is allocated to producer 1 in the TDM cycle) and producer 3 gets the transmission access as it has the higher priority than producer 2 and no transmission request from producer 1; however, producer 2 gets the access of slot 2 and slot 3 as the dynamic priority of producer 2 is highest as these slots are allocated to it in the TDM cycle (refer Figure 6). Producer 2 completes transmission and returns access at the end of slot 3 and producer 3 gets the access to the following slot. Transmission request from producer 1, with the highest priority is received before completion of slot 4, and access is given for slot 5. The lowest priority producer 4 with highest bandwidth requirements get the access at slot 8 when producer 3 finishes transmission. Producer 4 continues transmission unless finished at slot 15, and network is idle at slot 16.

Such an arbitration can offer deterministic worst-case latency for all the producers and guarantee transmission of packets of different priorities. This is a mixture of *best-effort* and *guaranteed-service* where best-effort is attempted when possible, but a guaranteed service is maintained under all possible conditions, even for the producers with lowest priority.

The effectiveness of the arbitration can be better understood by analyzing in the context of a flight control implementation, where signals with different functionalities and requirements can be categorized as - discrete, sampling and streaming signals. Discrete signals are triggered on the occurrence of some event that are not frequent

Figure 7: The figure shows transmission request and transmission access in the proposed priority-TDM cycle.

but needs low latency end-to-end transmission to meet hard-real-time constraints. Sampling data are regular, between subsystems or IO devices that are transmitted limited times (mostly once) per computational cycle. Streaming data, like log/data recorder or multimedia has high bandwidth but lenient latency requirements.

The hypothetical example we considered earlier in fact represents the same framework, where producer 1 represents discrete, producer 2 and 3 represent sampling and producer 4 represents streaming data transmission. When all the producers obey the transmission agreement, the operation takes place as explained. In case of a dysfunction in a high priority transmission, where the producer transmits more than the agreement, the arbitration guarantees transmission from low-priority producers, that is not achievable by a best effort only traffic.

5.0 Implementation

Microarchitectures of the network components and the arbitration have been explained in previous sections. This section explains the operation of the network.

5.1 Operation

The network must be configured before operation by configuring the destination and sources addresses in the NIs for each channel. This is done by writing in the destination address registers and expected-source-address registers with a standard memory writing technique. There is no memory access mechanism like DMA in this network

and at the beginning of a transmission, the producer pushes a data packet to the associated NI by a standard writing method for sending it to a pre-configured destination. Recall, the NIs interface with the communicating ends via a memory-mapped interface. As size of the data packets and the number of flits per data packet are predefined and fixed to avoid any skew, the need of a tail *flit* is obsolete in this architecture. The payload size is set to 8 *words*, that should accommodate all data-types used in control applications. The producer is responsible for evaluating the data size before transmission; if the data-type is greater than the payload size, the data should be segmented and each segment is sent separately, however, a data-type less than the payload size does not need any special treatment.

Once the writing process by the producer is complete, the data packet is transferred to the router in the next clock cycle. data packet received from the NI at the router contains 8-bit destination address, followed by 2-bit channel id, followed by the payload as shown in Figure 8. The *flitization* and *de-flitization* is done at the router on the received data from NIs and the hub respectively. The data packet from a NI is unpacked into *flits*, with a single header flit followed by the payload flits for transmission. The router adds a source address i.e. a concatenation of the router address, NI address and channel id in the header, which is later used for authentication. We have used a concept of the dynamic header where the size and information in the header changes as the packets flows through the network to reduce the amount of data flow as shown in Figure 8.

The router sets the associated transmission request line *high* and the line is held high until all the flits from that specific NI are transmitted. The transmission access given to the routers are NI specific and the router only transmits packets from the associated NI. This is how the notion of prioritized arbitration implemented in the hub is carried to the NIs.

Figure 8: Packing and un-packing of data packet and data packet header at different stages of flow. *dst* : 8-bit destination address; *src* : 8-bit source address; *NI* : 2-bit NI address; *ch* : 2-bits channel id

The hub consumes a single flit from the Xbar for the router with the transmission access every clock cycle. A link is established between an input *phit* and an output *phit* based on the destination addressed carried by the first *flit* and associated transmission state line to the router is set to *high* in the same clock cycle. This path is maintained for at least next two *flits* (two cycles) and until the hub revokes the access.

The input line from the hub to the router can consume a *phit* at every cycle. Three successive incoming *flits* are pipe-lined and re-structured for transmission through router-to-NI line that has a different *phit* size. The destination router de-*flit*ize the *flits* with only destination channel id and the source address in the header before sending it to the destination NI.

At the consumer end, the channel of the received data packet is identified from the channel id in the header. Further the source id in the header is evaluated to check the authenticity of the producer. The packet is stored in the receiving channel buffer if the source address matches the expected source address, pre-configured in the NI.

5.2 Scheduling and Latency Analysis

The arbitration mechanism needs a static schedule before the network can operate, where slots for each channel should be configured as per latency and bandwidth requirements. This scheduling is done by the user and a separate process. A low latency requirement is fulfilled by assigning multiple distributed slots in the TDM cycle. This could be a complex process to strategically accommodate multiple slots in the TDM cycle as adding a new slot changes the TDM cycle time and affects the latency of other schedules. Moreover, the maximum number of slots in the TDM cycle is also limited due to physical limitation of resources. In this work, the maximum number of slots is fixed to 96.

A tool is developed that computes the schedule with an iterative method. The user needs to input the number of producers and bandwidth and latency requirements for each producer. The tool initiates by assigning number of slots based on bandwidth requirements only, where higher requirement of bandwidth is accommodated by assigning more than one slots to the channel. Next, the tool sequentially picks the channel and inserts additional slots or removes slots assigned to the selected channel to meet latency requirements. Asserting or removing slots for one channel affects the schedule of other channels, and the tool iterates the process until the latency requirements are met for all the channels. The tool outputs the schedule and the total number of slots in the TDM-cycle. If the number of required cycles computed by the tool exceeds the physical limitation of the network, either the network needs to be reconfigured or the latency requirements should be more lenient, or number of channels can be reduced.

The end-to-end latency is dynamic and depends upon the number of communicating ends and load on the network. However, the worst-case-latency only depends on the number of communicating nodes used in the network and fixed unless the configuration is modified. The worst-case-latency can be computed by reversing the concept of scheduling, as –

$$L_{channel} = \left(\left\lceil \frac{S_{total} - 1}{S_{channel}} \right\rceil + 1 \right) \times t_{slot} + 1 \quad (1)$$

where, latency of a channel in clock-cycles is $L_{channel}$, $S_{channel}$ is the number of slots assigned to the channel, S_{total} is the total number of slots in the TDM-cycle and t_{slot} is the clock cycle per TDM-slot.

5.3 Protection and Isolation

Data protection and established isolation are one of the primary concerns for application in mixed criticality-systems and a key contribution of this work. This section elaborates the isolation and protection aspect of the architecture in end-to-end packet flow.

An arbitrary transmission starts with a producer writing a data packet in the transmission buffer of the associated NI. There is no channel specific buffer for packet under transmission in the NI, however, the packet is transferred to the connected router atomically, establishing a temporal isolation between two successive packets from the same producer. Routers have dedicated sampling buffer to hold the packets under transmission, unless the transmission access is gained. The sampling buffers are isolated registers in the physical hardware, offering spatial isolation between each data packet. In a dysfunction condition where producer violates the transmission agreement and a new data packet is received at the router before the previous packet is transmitted, the old-packet gets over-written by the new packet, but data packets in other buffers remain un-affected.

The arbitration is implemented in the hub and transmission access is provided in a deterministic schedule, guaranteeing access to each producer. The hub controls the transmission lines with circuit switching mechanism and only the router with transmission access is connected to the X-bar at any point of time, ensuring no-packet collision. The memory-less hub operations are atomic, establishing a temporal isolation. On the receiving end of the router, *flits* are packed and forwarded to the destination NIs. At the NI, each channel has its dedicated sampling buffer where the fresh data packet is saved for the consuming end to read. A dedicated sampling buffer provides a spatial isolation that prevents each feature from getting overwritten by data packet from another channel before consumed by the consumer application when transmission-agreement is violated by a faulty application.

Figure 9: Flow diagram showing *temporal* and *spatial* isolation in different stages.

The arrows marked in red shows temporal isolation and the blocks and arrows marked in green show spatial isolation.

Communication between multiple systems is prone to erroneous transmission from a faulty application to a wrong recipient. Such a fault is hard to detect in the software if the faulty data is ranged within the expected data range at the consumer end.

In this network, the destination address is configured in the NI for each channel and a dysfunction in the producing application cannot tamper with the destination. Additionally, the receiving NI has a pre-configured authorized source address for each channel. On reception of a data packet, the consumer NI checks for the source address before registering the message in the reading buffer. The source address is added by the router in the header during propagation and the application has no control over it. Such a mechanism provides a two-step protection to prevent transmission to wrong

recipient.

Additionally, faults like frozen data are hard to detect where the producer system may transmit the same data to the consumer if there is no change in physical state, for e.g. in cruise or hover condition, a flight control application can correctly send the same attitude data to a display application. A time-stamping in DMA based solutions significantly increase data-flow and need additional software feature to handle the timing data. In this work, the data packet transmission from the NI to the router is accompanied by a toggling signal that changes on every fresh transmission from the producer in spite of the content of the data packet; the router only registers the data from the NI when the toggling signal changes state, ensuring the transmission of only fresh data packets.

6.0 Results and Discussion

6.1 Experimental Setup

All the hardware is defined in *Verilog* HDL and synthesized on FPGA threads. In this work we have used *Xilinx ARTIX 7* and *Intel Cyclone V* SoC chip, although the hard embedded processor on the SoC was kept untouched. The board has a default 50 MHz oscillator and two external oscillators of 80 MHz and 100 MHz has been used for experimentation.

Each network component (NI, router and hub) are separate modules and defined as a Quartus custom/ external IP Cores, written in Verilog. Intel NIOS II soft-processors are used as producers and consumers and connected with the NIs with avalon-memory mapped interface. All the network components and the communication ends share the same global clock and reset signals. The components are inter-connected with Intel's Quartus Platform designer tool. The connections between network components (NI-router and router-hub) are not visible to the platform designer tool and should be externally connected by editing the top-module before synthesization. Quartus Prime lite edition tool has been used for synthesis.

6.2 Performance

To evaluate the performance of the proposed network architecture, an example network has been configured with four routers and twelve network-interfaces as shown in Figure 1.

Table 1 shows worst case latency analysis in different network configurations. Note that with increase in number of channels the latency of each channel increases.

Table 1 shows worst case latency analysis without any priority. However, if a channel has a lower latency requirement, meeting that requirement increases the worst case latency of other channels in the network. Figure 10 shows the effect of lowering the latency of one channel in the rest of the channels in a 36 channels configuration.

The bandwidth of the network depends upon the network clock frequency. The user can avail different oscillator depending on the bandwidth requirement. Table 2 shows the minimum bandwidth for a 4 router, 36 channels configuration with different clock frequencies.

Table 1: WORST CASE LATENCY ANALYSIS FOR DIFFERENT CONFIGURATIONS WITH 8 BYTES PAYLOAD AND 50 MHz OSCILLATOR. ALL CHANNELS HAVE EQUAL PRIORITY.

Number of Components			Latency	
Router	NIs	Channels	in cycles	in msec
2	2	6	19	0.00038
	4	12	37	0.00074
	6	18	55	0.00110
3	3	9	28	0.00056
	6	18	55	0.00110
	9	27	83	0.00166
4	4	12	37	0.00074
	8	24	73	0.00146
	12	36	109	0.00218

Figure 10: The bars in red represents the latency of a channel with low latency requirements. The bars in blue is the latency of other channels. The horizontal axis shows the number of slots assigned to the low-latency channel, vertical axis shows the latency in cycles.

Table 2: WORST-CASE BANDWIDTH OF A CHANNEL IN A 4 ROUTERS 36 CHANNELS CONFIGURATION WITH DIFFERENT NETWORK CLOCK FREQUENCIES. ALL CHANNELS HAVE EQUAL PRIORITY.

Clk (MHz)	Mega-bits-per-sec	Packets-per-sec
50	29.357	458714
80	46.972	733944
100	58.715	917430

7.0 Conclusion

We have proposed a network-on-chip architecture for the intended application in real-time mixed-criticality systems like integrated modular avionics platforms, that has some unique benefits- real-time end-to-end communication with isolation between data packets under transmission, different latency and bandwidth allocation in the same network and protection mechanism for authentic transmission that plays critical role in safety-critical application. Additionally, the concept of combined priority and time division multiplexing arbitration has been extended for better utilization of the network resources to allow more low priority applications to utilize the network

while maintaining determinism worst-case-latency for all the applications. However, the topology is subject to a linear extension in worst-case latency with expansion.

IMA is new technology and the guidance and requirements are evolving. The use of multi-core processors is new in today's avionics and the exact requirements of inter-core and inter-application communication still under investigation⁽²⁾. The performance of the proposed architecture in terms of bandwidth and latency is more than adequate to meet the requirements of conventional on-board applications. The fixed resources for the network components set a limit to the performance capabilities of the proposed architecture and increased bandwidth or low-latency demand in one channel affects the other channels. However, the worst-case performance is deterministic and analyzable for the system designer and no anomaly occurs during run-time. The hub in this architecture is the most critical component and could be a single point of failure. However, the applications do not have any effect on the hub, and the hub is only susceptible to hardware failures. A redundant implementation of the hub or the entire network can be considered for enhancing reliability measures.

The work was mainly focused on meeting the requirements of safety-critical aerospace applications and the scope of efficient resource utilization was not considered. Furthermore, the isolation mechanism degrades the efficiency of resource utilization as compared to general purpose communication. Further extension of this research to support inter-chip communication and scalability can be addressed in future research. We have implemented a lighter version of the proposed network in an asymmetric multiprocessor architecture to demonstrate improvements in the reliability of on-board computations in small airborne platforms⁽¹⁰⁾.

8.0 Acknowledgement

The authors would like to thank John-Josef Leth from Aalborg University and Martin Schoeberl and Jens Sparsø from Technical University of Denmark for their insightful comments and helpful discussion.

REFERENCES

1. FAA, Software Consideration in Airborne Systems and Equipment Certification (December 1992).
2. FAA, Assurance of Multicore Processors in Airborne Systems, DOT/FAA/TC-16/51 (July 2017).
3. R. L. Alena, J. P. Ossenfort, K. I. Laws, A. Goforth, F. Figueroa, Communications for integrated modular avionics, in: 2007 IEEE Aerospace Conference, 2007, pp. 1–18. doi:10.1109/AERO.2007.352639.
4. S. Hesham, J. Rettkowski, D. Göhringer, M. A. Abd El Ghany, Survey on real-time network-on-chip architectures, International Symposium on Applied Reconfigurable Computing (2015) 191–202.
5. I. Radio Technical Commission for Aeronautics, RTCA: DO-297: Integrated Modular Avionics (IMA) Development Guidance and Certification Considerations (2005).

6. C. M. Fuchs, A. S. Schneelee, E. Klein, The evolution of avionics networks from arinc 429 to afdx, in: In Proceedings of the Seminars Future Internet (FI), Innovative Internet Technologies and Mobile Communication (IITM) and Aerospace Networks (AN), Summer Semester 2012, p.65-76, Technische University of Munich.
7. P. Bieber, F. Boniol, M. Boyer, E. Noulard, C. Pagetti, P. Bieber, F. Boniol, M. Boyer, E. Noulard, C. Pagetti, N. Challenges, New Challenges for Future Avionic Architectures . To cite this version : New Challenges for Future Avionic Architectures (2015) 1–10.
8. Q. Perret, P. Maurere, E. Noulard, C. Pagetti, P. Sainrat, B. Triquet, Temporal isolation of hard real-time applications on many-core processors, in: 2016 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS), 2016, pp. 1–11. doi:10.1109/RTAS.2016.7461363.
9. F. M. P. R. Moustapha Lo, Nicolas Valot, IMPLEMENTING A REAL-TIME AVIONIC APPLICATION ON A MANY-CORE PROCESSOR, in: 42nd European Rotorcraft Forum (ERF), Lille, France, 2016, pp. 1–10.
10. S. Majumder, J. Nielsen, T. Bak, A. La Cour-Harbo, Reliable flight control system architecture for agile airborne platforms: an asymmetric multiprocessing approach, The Aeronautical Journaldoi:10.1017/aer.2019.30.
11. K. Sano, D. Soudris, M. Hübner, P. C. Diniz, Applied reconfigurable computing 11th International symposium, ARC 2015 Bochum, Germany, april 13-17, 2015 proceedings, Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) 9040 (2015) 191–201. doi:10.1007/978-3-319-16214-0.
12. D. Bertozzi, L. Benini, Xpipes: a network-on-chip architecture for gigascale systems-on-chip, IEEE Circuits and Systems Magazine 4 (2004) 18–31.
13. D. Wiklund, D. Liu, Socbus: switched network on chip for hard real time embedded systems, in: Proceedings International Parallel and Distributed Processing Symposium, 2003, pp. 8 pp.–. doi:10.1109/IPDPS.2003.1213180.
14. P. H. Pham, J. Park, P. Mau, C. Kim, Design and implementation of backtrack-ing wave-pipeline switch to support guaranteed throughput in network-on-chip, IEEE Transactions on Very Large Scale Integration (VLSI) Systems 20 (2) (2012) 270–283. doi:10.1109/TVLSI.2010.2096520.
15. P. T. Wolkotte, G. J. M. Smit, G. K. Rauwerda, L. T. Smit, An energy-efficient reconfigurable circuit-switched network-on-chip, in: 19th IEEE International Parallel and Distributed Processing Symposium, 2005, pp. 155a–155a. doi:10.1109/IPDPS.2005.95.
16. E. Bolotin, I. Cidon, R. Ginosar, A. Kolodny, Qnoc: Qos architecture and design process for network on chip, Journal of Systems Architecture 50 (2) (2004) 105 – 128, special issue on networks on chip. doi:https://doi.org/10.1016/j.sysarc.2003.07.004.
17. S. H. Lo, Y. C. Lan, H. H. Yeh, W. C. Tsai, Y. H. Hu, S. J. Chen, Qos aware binoc architecture, in: 2010 IEEE International Symposium on Parallel Distributed Processing (IPDPS), 2010, pp. 1–10. doi:10.1109/IPDPS.2010.5470359.

18. E. d. F. Corrêa, L. A. d. P. e. Silva, F. R. Wagner, L. Carro, Fitting the router characteristics in nocs to meet qos requirements, in: Proceedings of the 20th Annual Conference on Integrated Circuits and Systems Design, SBCCI '07, ACM, New York, NY, USA, 2007, pp. 105–110. doi:10.1145/1284480.1284514.
19. C. H. Lu, K. C. Chiang, P. A. Hsiung, Round-based priority arbitration for predictable and reconfigurable network-on-chip, in: 2009 International Conference on Field-Programmable Technology, 2009, pp. 403–406. doi:10.1109/FPT.2009.5377690.
20. J. Diemer, R. Ernst, M. Kauschke, Efficient throughput-guarantees for latency-sensitive networks-on-chip, in: 2010 15th Asia and South Pacific Design Automation Conference (ASP-DAC), 2010, pp. 529–534. doi:10.1109/ASPDAC.2010.5419828.
21. M. Millberg, R. J. A. Nilsson, E. and Thid, Guaranteed bandwidth using looped containers in temporally disjoint networks within the nostrum network on chip, no. 6, 2004. doi:10.1080/00207210600562645.
22. K. Goossens, J. Dielissen, A. Radulescu, Aethereal network on chip: concepts, architectures, and implementations, IEEE Design Test of Computers 22 (5) (2005) 414–421. doi:10.1109/MDT.2005.99.
23. K. Goossens, A. Hansson, The aethereal network on chip after ten years: Goals, evolution, lessons, and future, in: Design Automation Conference, 2010, pp. 306–311. doi:10.1145/1837274.1837353.
24. R. A. Stefan, A. Molnos, K. Goossens, daelite: A tdm noc supporting qos, multi-cast, and fast connection set-up, IEEE Transactions on Computers 63 (3) (2014) 583–594. doi:10.1109/TC.2012.117.
25. E. Kasapaki, M. Schoeberl, R. B. Sorensen, C. Muller, K. Goossens, J. Sparso, Argo: A Real-Time Network-on-Chip Architecture with an Efficient GALS Implementation, IEEE Transactions on Very Large Scale Integration (VLSI) Systems 24 (2) (2016) 479–492. doi:10.1109/TVLSI.2015.2405614.