Design Drivers for SpaceVNX+

A Small Form Factor Electronics Unit for Space Applications

Steve Parkes
STAR-Dundee Ltd.
166 Nethergate
Dundee, DD1 4EE UK
steve.parkes@star-dundee.com

David Givens
Industry Standards VP
Samtec Inc.
New Albany, USA
david.givens@samtec.com

***Abstract*— SpaceVNX+ [1] is a standard for small form factor equipment modules and units which is being designed specifically for space applications. Its modules are substantially smaller than a 3U Eurocard. SpaceVNX+ provides a standard platform for the implementation of the entire range of avionics applications on-board a spacecraft, from simple remote terminal units to high-performance payload data-handling units. SpaceVNX+ is intended to be complementary to larger form factor standards such as SpaceVPX [2] and ADHA [3]. This paper considers the critical design drivers for SpaceVNX+ that make it suitable for space applications, including thermal, size, modularity, connectivity, reliability and redundancy, and electrical constraints. The design of SpaceVNX+ is being driven by these considerations.**

Keywords— Avionics, Payload Data-Handling, Remote Terminal Unit, SpaceFibre, VNX+, SpaceVNX+.

1. Introduction

SpaceVNX+ (currently designated as VITA 90.5 [1]) is a development of the military/aerospace standard VNX+ [4] which is an emerging VITA standard for small form factor systems based on the earlier VNX standard. VNX is widely used in aerospace and military terrestrial applications. The ANSI/VITA standard for VNX+ is expected to be published in 2025.

SpaceVNX+ is a version of VNX+ specifically being designed for space applications. It can support single-string, dual-redundant and 1-of-M redundant architectures to suit various application reliability needs. SpaceVNX+ has a small form factor and offers high-performance and comprehensive connectivity. Further information is available for VITA members on the VITA website [1].

SpaceVNX+ is currently in its definition and prototyping stage with drafting of the formal standard expected to start late-2025. This paper aims to present SpaceVNX+ to the wider community and to encourage feedback prior to the drafting of the formal standard. The principal requirements for SpaceVNX+ are outlined and then the main architectural drivers are each considered in turn. The architectures of two SpaceVNX+ systems are presented. A single-string architecture is presented first with no redundancy and then a redundant architecture is presented. These two architectures are examples of the many architectures that can be constructed with SpaceVNX+ modules.

2. Requirements

The principal requirements for SpaceVNX+ are listed below:

- Shall be suitable for space applications in particular with respect to thermal management, radiation tolerance and reliability.
- Shall have a small size, less than a 3U Eurocard.
- Shall be modular with modules that bring together related functionality into a module and minimise the interaction with other modules (highly-cohesive and loosely-coupled [5]).
- Shall be scalable enabling payload modules to be added to a unit without an update to the unit infrastructure.
- Shall be flexible, supporting the addition of expansion modules to add capabilities to a module.
- Shall support a range of redundancy approaches including single-string, dual cold-redundancy, dual soft-step redundancy, and 1 of M and N of M redundancy.
- Shall support the use of modules designed for single-string systems in a redundant system without requiring significant additional circuitry on those modules.
- Shall support the integration of functions into one unit or the distribution of functions across multiple units.
- Shall provide high-performance (many Gbit/s) interconnect between modules.

- Shall provide fault isolation, to avoid propagation of faults between modules.
- Shall support rapid integration.
- Shall embrace simplicity as a primary means of providing reliability.

The way in which these requirements drive the SpaceVNX+ architecture is detailed in the following sections.

3. Thermal Management

Thermal design is critical to any electronics unit on-board a spacecraft. Effective conductive cooling is essential but often it is a second thought when developing a standard for an equipment unit. It is worth looking at the possible configurations of modules/boards in a unit, as illustrated in Fig. 1.

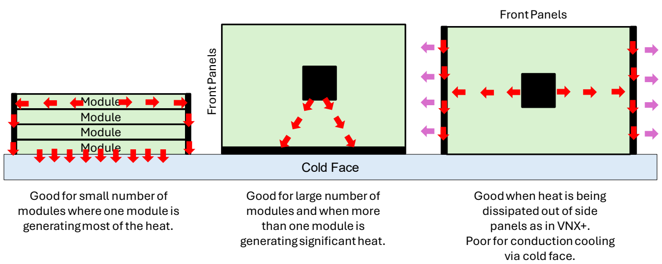

Fig. 1 Module Arrangements in a Unit

The arrangement on the left shows modules that slide in from the front and stacked one on top of the other. The cold surface which removes the heat is shown in blue on the bottom. The lower module is able to conduct heat directly to the cold surface whereas the upper modules have to conduct heat across the module, through wedgelocks, and down the sides of the modules to a relatively small contact area with the cold surface. This arrangement is effective when there are a small number of boards with one of the boards dissipating most of the heat.

The arrangement in the middle of Fig. 1 shows modules that are slid into the unit from the left with wedgelocks on the top and bottom edges. The main path to the cold surface is across the module to the lower wedgelock which provides a reasonable surface area per module to the cold surface. If the main heat generating components are moved towards the bottom of the board, the thermal efficiency is improved further. Unless heat-pipes (or similar thermal arrangements) are connected to the upper surface of the unit, the thermal path through the top wedgelock is not substantial.

The arrangement on the right of Fig. 1 has the boards sliding into the unit from the top. Heat is dissipated through wedgelocks on both sides of the module (red arrows). The area of each module in contact with the thermal surface is very small. Note a backplane is at the bottom of the modules which prevents the bottom of the board being used as an effective thermal interface.

VNX+ is often used in air cooled units where the thermal path is directly to the sides of a unit as shown on the right hand side of Fig. 1 (purple arrows - in this case there is no cold surface at the bottom, but there are convection/radiation cooling on the sides and possibly the top).

For SpaceVNX+ two module arrangements are supported. Stacked modules (left hand side of Fig. 1 are used when there are a few modules (e.g. maximum of four) and either they are all low power or only one of them dissipates substantial power. Note that this arrangement is also good from the mechanical perspective, keeping the height of the unit low compared to its footprint.

For more general applications where there are more than four modules, or more than one of them dissipates significant power the “slide in from side” arrangement is preferred (centre of Fig. 1). To maximise the thermal performance the board should be longer on the edge which is in contact with the cold surface. For a system with several boards, this is much more stable mechanically than vertically stacked modules.

4. Form Factor

The VNX+ form factor is deliberately made to fit in a very small compartment and to be easily maintainable. The length of the wedgelock in VNX+ is limited by the small size of the board (73 mm depth) and further reduced by the connector and card guide arrangement, resulting in around 50 mm available for the wedgelocks. This is small for a conductively cooled module. A VNX+ board is illustrated on the left of Fig. 2.

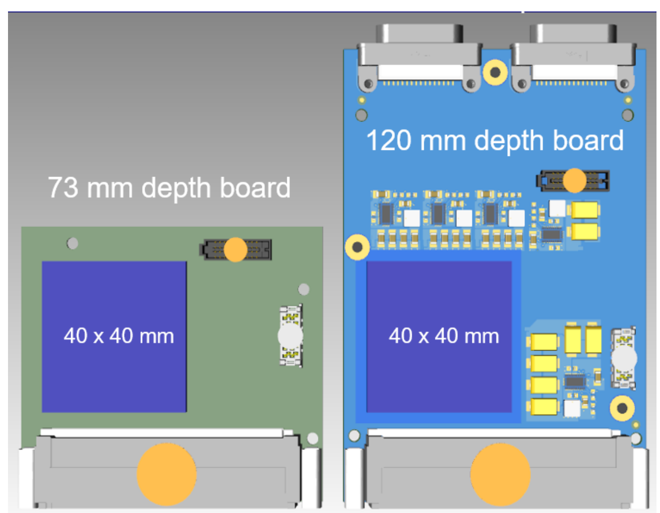

VNX+ Size Board SpaceVNX+ Extended Board

Fig. 2 Board Sizes

Spaceflight components are significantly larger than commercial/industrial/automotive components, although Space Enhanced Plastic components provide smaller size, albeit at lower reliability. An FPGA like the AMD KU060 or Microchip PolarFire is 40 x 40 mm in size. This is shown as the blue rectangle in Fig. 2. A keep out area is also required for inspection. Clearly there is not a lot of space for other components on the board when using these FPGAs, although other smaller radiation tolerant FPGAs are available and becoming available.

It is possible to stack two or three boards in a wider module which provides more board area. When this is done it should be recognised that this does not permit increased power consumption, because the wedgelock area remains unchanged. Furthermore, mechanical fixings and connectors to the daughterboards reduce the usable area on the main board. Finally, consideration needs to be given to the routing of high-speed signals (>1 GHz). Passing high-speed signals through multiple connectors and across long reaches of a PCB should be avoided. This leads to the preference from the signal integrity perspective for the main components (FPGAs and processors) to be placed on the main board, rather than on a daughterboard.

To increase the available board area and to increase the permitted power consumption the board has to be increased in length. The primary SpaceVNX+ board will have the same width as a VNX+ board (84 mm) but a depth of 120 mm, see Fig. 2. This will almost double the board area and permitted power consumption of a SpaceVNX+ board. Standard size VNX+ boards (73 mm depth) will also be permitted and the backplane connector will be in the same position on the VNX+ board. Note that a board of 84 x 120 mm will fit into a 2U or larger CubeSat, allowing SpaceVNX+ to be used on very small spacecraft as well as having the performance and capabilities for much larger ones.

5. Modularity

VNX+ is designed for readily maintainable systems, for example a unit on an aircraft being replaced in a dusty, inhospitable environment. The faulty unit needs to be replaced quickly without the use of fiddly tools to unscrew connectors and without allowing dust into the connector contacts. For this reason the front panel of a VNX+ module is normally devoid of connectors. All connections from a module to the outside world are through the backplane via connectors mounted on the chassis and wired to the backplane. These connections remain untouched when a module is replaced.

This approach is advantageous in some space missions: manned spaceflight where humans do the maintenance, or future systems that are maintained by robots. It does, however, have the disadvantage of eroding the modularity of the modules. Since their connections go through a common backplane and connectors on the chassis, modules are linked with the backplane and hence with the other modules in a unit. Putting connectors on the front panel of a module allows those interface functions directly related to the module to be on the front panel making it modular. For example, RF connections for a software defined radio would be on the front panel. Generic connections, such as module power and interconnect between modules, are via the backplane. The use of the front panels for connectors also provides more connectivity for a module.

For the vast majority of space missions, maintenance is not possible. The advantages of a modular design, which improves design reusability and provides additional connectivity, leads to SpaceVNX+ supporting connectors on the front panel. Of course, while connectors on the front panel are supported, it is also possible not to have connectors on the front panels of modules that are designed for maintainable systems.

6. Scalability

Scalability is necessarily limited by the number of modules that can be installed in the reasonable maximum size of a unit. It is also limited by the increasing infrastructure needed to support each additional module.

As we have seen, modularity reduces the number of interconnections between modules (loosely-coupled). The reduced number of interconnections helps to improve scalability as the infrastructure to support those interconnections is reduced.

Modularity and hence scalability can be enhanced by ensuring that all of the related functions are included in a module (highly-cohesive). On a small form factor board this may not always be possible. To provide more board area for a given module, daughterboards and expansion modules can be used (see section 11 for a description of an expansion module).

In SpaceVPX+, the main scalable parameter is the number of payload modules that can be supported in a unit (see section 11 for a description of a payload module). For SpaceVNX+ a maximum of eight payload modules is allowed. Beyond that an application should be separated into multiple units.

7. Flexibility

Expansion modules provide extra board space, but they can also enhance flexibility. Different expansion modules can be used with a generic payload module to customise the IO capabilities or functionality of the payload module. The common elements are implemented on the payload module and the different elements are put on expansion modules.

The expansion interface on a payload module comprises power, reference voltages, GPIO and Multi-Gigabit transceivers. If an expansion board is not required, these pins may be used for other functions – in VITA 90 these alternative pin functions are called overlays. An overlay is used when pins are used to carry different types of signal to the nominal configuration. Being able to change how pins are used adds substantial flexibility to SpaceVNX+.

The difficulty with overlays is that if the signals are different it might be possible to have potentially conflicting signals on different overlays. If a payload module with one overlay is plugged into a slot that is expecting a different overlay, it might result in damage when the modules are powered. To avoid this, chassis management has to determine what type of payload module is in a slot and what type of overlay it is using. Only if this is as expected is the payload module fully powered (12V power input enabled). Chassis management allows the use of overlays without restriction, although it is aways good practice to minimise any potential inconsistencies between overlays. SpaceVNX+ will define a set of possible overlays.

8. REDUNDANCY

Fault tolerance is the capability of a system to continue to operate in the event of a failure of one of its functional elements. A functional element is a part of a system which provides a function necessary to implement the system and can be at the component, module, unit or spacecraft level. Redundancy is used to help improve the reliability of spaceflight systems. Redundant elements are added to detect and recover from faults.

Redundancy can also be added at the chip level using Triple Mode Redundancy (TMR) or Distributed TMR (DTMR) to improve the radiation tolerance of the device against Single Event Transients (SETs) and Single Event Functional Interrupts (SEFIs).

At the module level two copies of a module are provided: the default module (also called nominal, main or prime) and the redundant module. The default module is normally the element that is in operation. If a fault is detected in the default module, it is isolated and the redundant module switched in to replace the faulty element. Operation then continues. This provides the three stages of fault recovery: fault detection, fault isolation and fault recovery referred to as FDIR. Fault isolation is important to prevent a fault propagating across a system. A fault should be detected as close to the source of the fault as possible and isolated there. The redundant modules can be cold redundant (powered down), warm redundant (powered up but idle) or hot redundant (powered up and operating, shadowing the default module and able to take over operation very quickly). Cold redundancy takes much less power and is more radiation tolerant (generally not being affected by radiation while it is powered down). Hot redundancy is much quicker to respond and takes roughly twice the power. Warm redundancy is somewhere in between.

To switch between a default and redundant module it is necessary to have switches on the inputs and outputs of a module that switch the default module out and the redundant module in.

If a module is powered down while other modules are powered, it is possible for inputs on the powered down module to be driven by active outputs of a powered module. This can damage the unpowered module. Resilience to powered inputs when an element is powered is called "cold-sparing". This is highly desirable when implementing systems which can be partially powered. In any event care must be taken to not drive any non-cold-sparing input of an unpowered module.

Redundancy can be applied at the unit level. This can be advantageous when the number of connections to be cross-strapped is small. Implementing the cross-strapping inside a unit and putting the cross-strapped connections on a backplane can be simpler, especially when the cross-strapping approach is standardised as in SpaceVNX+. It can also save significant cable mass.

Redundancy can also be provided at the spacecraft level. In this case a single-string implementation is provided on the spacecraft. Any significant fault would render the spacecraft useless. Redundancy is then achieved by launching another spacecraft. For operational missions, the time to launch another spacecraft may be a significant impediment to this approach.

Redundancy is used to enhance reliability. Avoiding single point failures and having redundant elements can significantly enhance the reliability of a system. This may also enable lower reliability, lower cost components to be used, while still achieving an overall reliability goal.

9. LOW OVERHEAD FOR SINGLE STRING

SpaceVNX+ is targeting a full range of space missions from CubeSats to large spacecraft. For lower reliability missions including in-orbit demonstrators, a single-string implementation is all that is needed. The modules in SpaceVNX+ are being designed to carry little overhead for supporting fully redundant operation when operating in a single-string application. The features required to support full redundancy will generally be outside a module. The same modules can be used in both single-string or fully redundant systems.

10. DATA CONNECTIVITY

Current FPGAs incorporate multi-gigabit transceivers (MGTs) which can support a variety of different communication protocols, including SpaceFibre-E (electrical version of SpaceFibre) [6], 10GBASE-KR4 Ethernet and PCIe. SpaceVNX+ defines high-speed interconnect suitable for these various protocols. Allowing a variety of network protocols to be used. Payload modules are typically end-points on a network, and the System Controller module incorporates an appropriate routing switch to facilitate communication between payload modules.

Besides SpaceFibre, PCIe and certain versions of Ethernet, SpaceVNX+ will also include SpaceWire [7] as an optional data network.

11. SINGLE STRING ARCHITECTURE

SpaceVNX+ can support many applications and many architectures tailored to those applications and supporting various levels of redundancy. An example architecture of a

SpaceVNX+ system without redundancy (single-string) is shown in Fig. 3.

There are five types of module:

- Power Supply (PSU);
- Power Switch (PSX);
- System Controller (SC);
- Payload Module (PM); and
- Expansion Module (EXP).

The **Power Supply (PSU)** provides the two voltage rails used by SpaceVNX+: 3V3AUX and 12V. 3V3AUX is used to power auxiliary circuitry related to chassis management. 12V is used for the principal power rail to the System Controller (SC), Payload (PM) and Expansion (EXP) modules. The power supply takes input power either from a front panel connector or from the backplane. The input voltage is typically 28V nominal, but other input voltages are permitted to suit the application. When the PSU is provided with input power, a microcontroller on the PSU is enabled first. This microcontroller provides first-stage chassis management (first-stage chassis manager) as well as controlling the power supply and related power switches. The microcontroller confirms that the PSU module is in the correct slot of the SpaceVPX+ rack using the slot's geographic address, that the input power is at an acceptable voltage, and that the 3V3AUX and 12V regulators are operating correctly. It then provides 3V3AUX and 12V to the power switch module (see the bottom arrow in Fig. 3.). The first-stage chassis manager then switches on 3V3AUX to the System Controller. It talks to the SC over an I2C link to confirm that it is a SC module. The first-stage chassis manager then switches on 12V power to the System Controller.

The **Power Switch (PSX)** switches 3V3AUX and 12V power to the System Controller, Payload and Expansion modules. The 3V3AUX and 12V power switches to the SC are operated by the first-stage chassis manager on the PSU using discrete control lines. The 3V3AUX and 12V power switches to the PMs and any EXP modules are operated by the second-stage chassis manager on the SC module again using discrete control lines. Discrete control lines are used rather than I2C, to speed up the response to a fault. The PSX module can be thought of as an extension to the PSU board. If only a few modules are being implemented in a system, both power supply regulation and power switch may be integrated on the PSU module.

The **System Controller (SC)** provides the second level of chassis management controlling the 3V3AUX and 12V power switches that switch power to the PMs on the PSX module. The SC controls all the PM modules (and associated EXP modules) providing an I2C connection to each PM along with reset, and status signals. The SC distributes a 100 MHz reference clock (REFCLK) and 1 pulse per second (PPS) auxiliary clock (AUXCLK) to each of the PMs. The SC also provides the data routing-switch when required, with a data connection to each PM.

The **Payload Module (PM)** provides some required function. Together the one or more PMs in a unit provide the capabilities required by the application. The PSU, PSX and SC all provide support services to the PMs. Should a PM need to communicate with one or more other PMs, it will do so via data links that are typically connected to a data switch in the system controller. It is also possible for the data links available to be directly connected to other PMs. Example functions of a PM include sensor interfacing and control, data compression, data storage, data encryption, data fusion and down link encoder.

The **Expansion Module (EXP)** is a companion module to a payload module. In essence it is like a daughterboard on a PM, except that the connections between the PM and the EXP modules are via the backplane connector. This approach avoids the need for daughterboard connectors which would take up valuable space on the PM board and daughterboard. It also provides another wedgelock to take heat from the EXP module, whereas a daughterboard would share a wedgelock with the PM board. In a unit, the EXP board is adjacent to its PM board. The backplane connections between the PM board and the EXP board include regulated power, GPIO, Multi-Gbit/s transceivers (MGTs) and reference voltages for the GPIO. The EXP board interface contains signals that are very much like those of an FMC daughterboard. From the system perspective the EXP board is an expansion of the PM and it is connected only to the signals from the PM, except for 3V3AUX and 12V power which, if required, can be provided from the same power connections as for the PM. It is also possible for the EXP board to be a COTs processor under control of and reporting to the PM it is attached to.

The chassis management function (chassis manager) is split between the PSU and the System Controller because a typical space-grade microcontroller has a limited number of I2C ports – not enough to have point-to-point I2C links to each SC and PM module. For this reason, chassis management is separated into two stages. The first stage on the PSU microcontroller is responsible for checking the PSU and SC. The second stage is on the SC, which is likely to include an FPGA. It is responsible for checking all the PMs and their expansion boards. For each slot, the relevant chassis manager checks that each slot contains the correct module, before it permits that module to power up fully. Once all the modules have been checked the SC can power up the PMs and associated EXP modules as required.

It is possible in a single-string unit for 3V3AUX to be bused from the PSU to the other modules avoiding the need for 3V3AUX switches. However, this removes the fault isolation provided by those switches, so that a fault on any module affecting 3V3AUX would propagate to all modules.

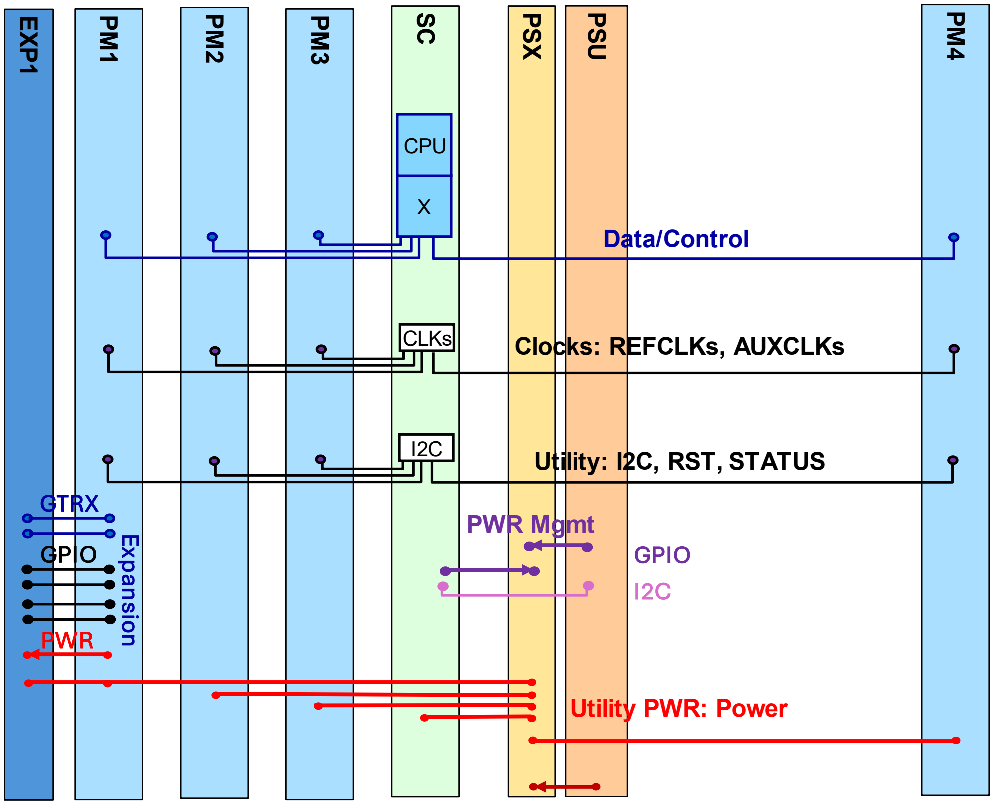

Fig. 3 SpaceVNX+ Single-String Architecture

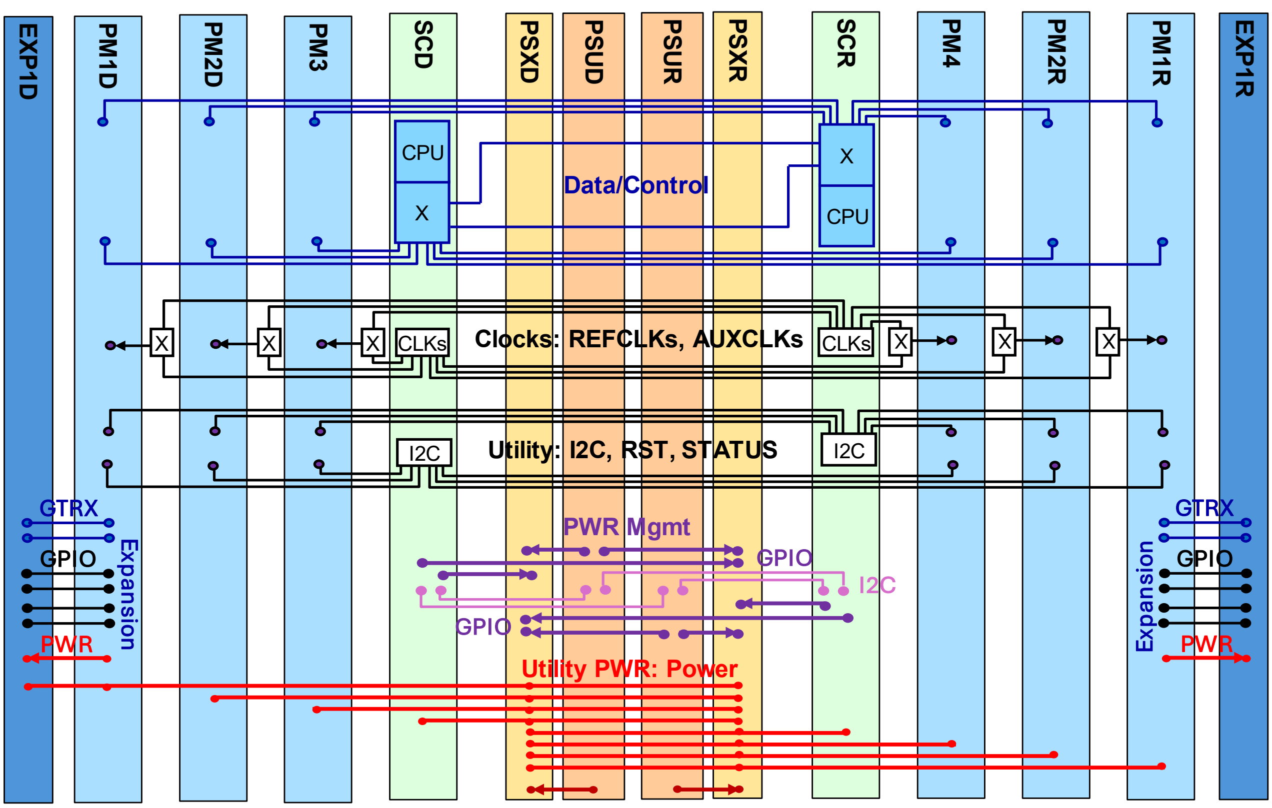

Fig. 4 SpaceVNX+ Architecture with Redundancy

12. ARCHITECTURE WITH REDUNDANCY

In Fig. 4, the single-string architecture of Fig. 3 is expanded into an architecture with redundancy. The default modules have a suffix ‘D’ and the redundant ones ‘R’. Modules without such a suffix are single-string modules. Redundant Power Supply (PSUR), Power Switch (PSXR) and System Controller (SCR) have been added providing full redundancy for the default modules (PSUD, PSXD and SCD). The default payload modules (PM1D and PM2D) have been provided with redundant copies (PM1R and PM2R). Since PM1D has an Expansion module (EXP1D), PM1R is provided with a redundant copy (EXP1R). The other two payload modules, PM3 and PM4, are not provided with redundant counterparts – either their reliability is sufficiently high or they are not critical elements and their failure would not compromise the principal operation of the unit.

The redundant PSU, PSX and SC are identical to the default PSU, PSX and SC, whether they are default or redundant modules depends only on the geographic address of the slot they are plugged into.

Power is passed from the PSUD to PSXD and from PSUR to PSXR. Normally the default PSU would be provided with input power and would then pass power to PSXD. The redundant PSU would not be powered. If PSUR is powered, it first checks that PSUD is not powered or is powered but showing a fault condition, and would then power up and pass power to PSXR. If both PSUs are receiving power and PSUD is working properly, it will be PSUD that is enabled and power will be passed to PSXD.

The connection from a power switch to a PM or SC module is called a power channel. There is a separate power channel for each PM and each SC module, and for each voltage 3V3AUX and 12V. Each power switch has a connection to each of those channels, so the 3V3AUX power channel to PM1 is connected to a 3V3AUX power switch in both PSXD and PSXR. This can be seen in Fig. 4. Power will be provided from PSXD or PSXR, depending on which one is active, A power channel delivers power when the appropriate power switch on the active PSX is turned on.

The power switch control signals for power channels to the SCD and SCR modules come from the PSU associated with the PSX. The PSX control signals for the switches providing 3V3AUX or 12V power to the PMs, and any associated EXP modules, are provided by both the nominal and redundant System Controllers, whichever is active.

When 3V3AUX is switched on for a PM, the active System Controller can use its I2C link to that PM to confirm that it is the expected PM in the slot and that if an EXP module is connected to that PM, it is the correct EXP module. Only if all is well and all PMs have been checked, will the active SC turn on 12V power to the PM. It is necessary to check all PMs before powering up the PMs, in case one of the other slots contains an incorrect PM.

VNX+ and SpaceVNX+ use I2C for chassis management. It is difficult to use I2C in a redundant configuration owing to the bidirectional nature of the I2C SDA signal. Most small space-grade microcontrollers have at least three I2C controllers. For these reasons it has been decided to use I2C and to have two I2C interfaces on each PM, one I2C interface from SCD and the other from SCR. This is thought to be an acceptable overhead in return for being able to keep I2C as the chassis management communication protocol.

The active SC provides reference and auxiliary clocks (REFCLK and AUXCLK) to each of the PMs. These clocks are distributed using LVDS with matched length tracks so that the clocks are synchronous to all the PMs. It is necessary to use an LVDS switch to switch between the clock from SCD and that from SCR, even though one of them is not running. The clock switches could be placed on the PSX modules, on the payload modules or on the backplane. If they are placed on the PSX module the PSX modules would no longer be identical; PSXD would handle four of the clock switches and PSXR would handle the remaining four (of the possible maximum of eight – one for each possible PM). If the clock switches are on the PMs, they would represent an unnecessary overhead on the PM for single-string applications. The clock switches are therefore placed on the backplane, adjacent to the PM they are serving and using the power from that PM.

There are two data networks, the one with its routing switch on SCD and the other with its routing switch on SCR. Each routing switch is connected to each PM. For a redundant system, that means that the PM has twice as many data network interfaces as required when a PM is used in a single-string system. However, these spare data interfaces can be used for other purposes in a single-string system, or they can be disabled. The overhead is small.

Using SpaceFibre-E, for example, the data network interface to a PM is capable of a data rate of 10 Gbit/s in each direction, giving a bisection bandwidth for the switch of 80 Gbit/s.

In a system with two System Controllers, it is possible to have both of them operating. That provides twice the data network bandwidth. It is also possible to have different data network protocols running on each data network, provided that this is supported by all the PMs.

13. CONCLUSIONS

SpaceVNX+ is a standard for high-performance, capable, small form factor electronic systems for space applications. Standard modules can be configured in a single-string system or those same modules can be used in a fully redundant or partially redundant configuration.

The authors of this paper would welcome comments on the proposed SpaceVNX+ standard prior to the drafting of the formal standard document.

ACKNOWLEDGEMENTS

The discussions with the various members of the VITA 90.5 working group have been very helpful in focusing the work on SpaceVNX+. The VITA 90.5 working group is open to VITA members.

REFERENCES

[1] VITA, https://vita.com, last visited October 2025.

[2] ANSI/VITA, “SpaceVPX System Standard”, ANSI/VITA 78.0-2022, August 2022.

[3] ESA, “Advanced Data Handling Architecture (ADHA)”, ADHA-2 Workshop, ESTEC, November 2022.

[4] VITA, “VNX+ Base Standard, DRAFT – Not Approved”, VITA 90.0, Latest revision available at https://vita.com, last visited October 2025.

[5] Stevens, W. P.; Myers, G. J.; Constantine, L. L. (June 1974). "Structured design". IBM Systems Journal. 13 (2): 115–139. doi:10.1147/sj.132.0115.

[6] ECSS Standard ECSS-E-ST-50-11C, “SpaceFibre – Very High-Speed Serial Link”, Issue 1, European Cooperation for Space Data Standardization, May 2019, available from http://www.ecss.nl.

[7] ECSS Standard ECSS-E-ST-50-12C Rev.1, “SpaceWire – Links, nodes, routers and networks”, European Cooperation for Space Data Standardization, May 2019, available from http://www.ecss.nl.