



TRENDS ON VPX

Standardization and usage

HOW NEW PROFILES CAN OVERCOME THE PROLIFERATION OF VPX PROFILES	// 4
VPX IS READY TO SUPPORT EMBEDDED SYSTEMS OF THE FUTURE	// 7
KONTRON OFFERING IS LEADING THE WAY TO THESE NEW VPX ERAS	// 8
CONCLUSIONS	// 8



The VITA65 VPX modular computer standard is evolving to better take benefit of the new high speed protocols such as Ethernet at 40 and 100 Gbits/s over the backplane. This white paper describes such on going modifications and rationalization opportunities taken.



Author:
Serge Tissot, Technical Strategy Manager
at Kontron Modular Computers, France
support.KFR@kontron.com

VPX and its supporting standards VITA 46 and VITA 65 are one of the most successful modular computing implementations available today. Modular computers are generally practical when one needs to design a dedicated parallel processing architecture to increase performance, to comply with harsh environment constraints, or to meet reliability goals through redundancy. It also offers flexible I/O thanks to associated standards such as XMC and FMC. In addition, modular computing can contribute to the longevity of the system by facilitating technology insertions, thus avoiding the redesign of the whole system. This modularity comes in the form of multiple processing, communication and I/O modules operating together. It is therefore highly desirable to ensure interoperability of modules through strict standardization rules.

In this white paper, we examine architecture trends and current evolutions of VPX technology. Currently there are a large number of VPX profiles defined, leading to legitimate questions regarding interoperability of modules from different vendors or between generations of products. Another topic that we address is how the standard sustains technology improvements over time. Finally, we present how the Kontron VPX product roadmap answers these concerns now and in the future.

HOW NEW PROFILES CAN OVERCOME THE PROLIFERATION OF VPX PROFILES

Current situation with many VPX profiles

The latest revision of OpenVPX VITA 65 specification, released as an ANSI standard in 2017, contains a large number of profiles. The profiles that define the backplane pin assignment of the slot are called slot profiles. They, for example, might define 4 lanes starting at the beginning of connector P1 to implement a Data Plane, i.e. 4 differential pairs for transmit and another 4 for receive. Likewise, the module profiles specify which protocol a module supports: for example PCI Express® or Ethernet. Three main categories of interconnect planes are defined: the Data Plane to exchange main data between modules, the Expansion Plane to provide a local mean to extend the capabilities of a module through adjacent slot(s), and the Control Plane to orchestrate and monitor the whole configuration without disturbing the main Data Plane.

This elegant and flexible method to specify an electronic module pinout and its capabilities exhibits many advantages, including the use of the same backplane routing for different communication protocols. However, this quickly leads to a large number of possible combinations which can be, and actually are, concisely listed in the standard. The drawback is that interoperability between those modules and backplanes, especially between multiple vendors if we assume a single vendor makes at least consistent choices, is not achieved. The situation might even be worse because of user defined pins which are free to be customized by vendors, and because of the number of protocols eventually supported like PCI Express®, Ethernet, Infiniband, Serial RapidIO™, etc.

40 Gbit Ethernet is disrupting the ecosystem

The introduction, and anticipated generalization, of the Ethernet protocol at 40 Gigabits/second speed in embedded computing is creating a breakthrough in modular computing applications. This protocol is now available natively in modern CPUs and peripherals, even those designed to operate in severe environments, whether their architecture is based on x86 or Arm®. Moreover, TCP offload engines and multi-core operation coming with those implementations limit the CPU load during data movements at full throughput. This makes it very attractive to use the 40G Ethernet as the main Data Plane protocol to interact between intelligent VPX modules.

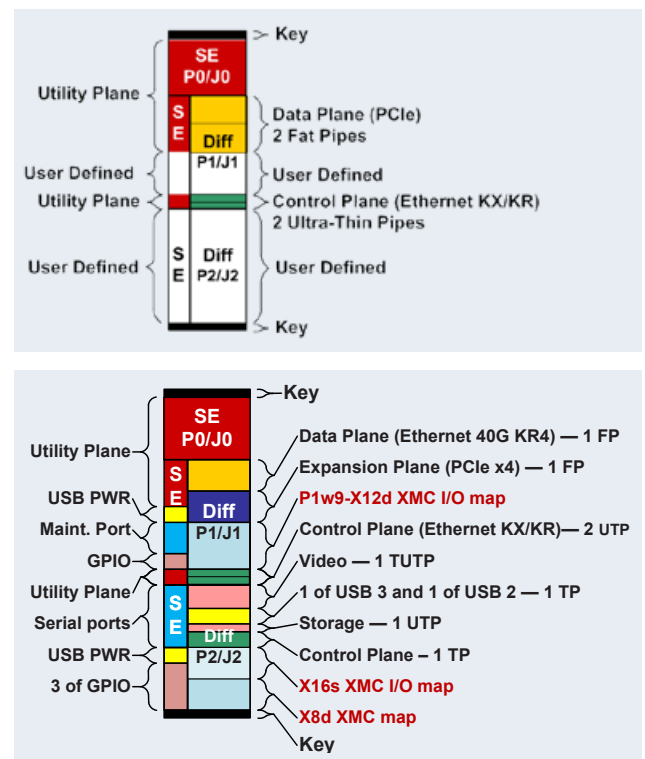
The majority of VPX modules or backplanes have not yet been designed with the goal to support 40G Ethernet as the main intra-chassis communication plane. This is creating potential pinout incompatibilities at the modules and backplane level. Many modules and backplanes

today are PCI Express® centric, using Ethernet only as a Control Plane.

The good news is that the VITA 65 working group is about to release a new version of the specification in 2019 which adds a limited set of new profiles optimized for 40G Ethernet as the main primary Data Plane. Those profiles, presented in the next section, have the potential to become the new architecture baseline for the next 5 or 10 years for modular computing. The role of PCI Express® on the backplane, once the primary high-performance data movement path, is mostly assigned to be the protocol of the Expansion Plane to connect local peripheral expansion cards.

In addition to those evolutions of the next VITA 65 standard, risks of interoperability between vendors are further reduced by defining all module rear I/O pin assignments for interfaces such as Storage, Graphics, USB, UART, etc. and removing all user-defined pins. In previous version of the standard, this kind of interfaces was implemented using these user defined pins which lead to as many implementations as vendors.

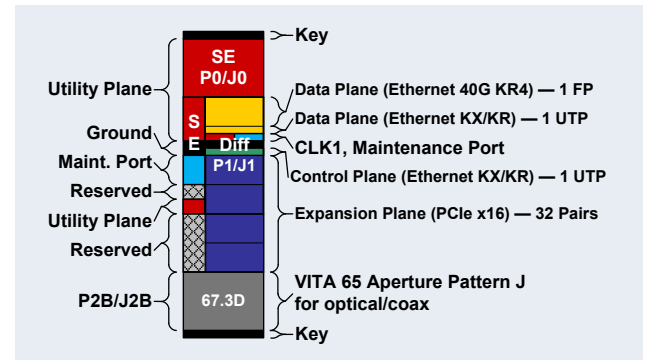
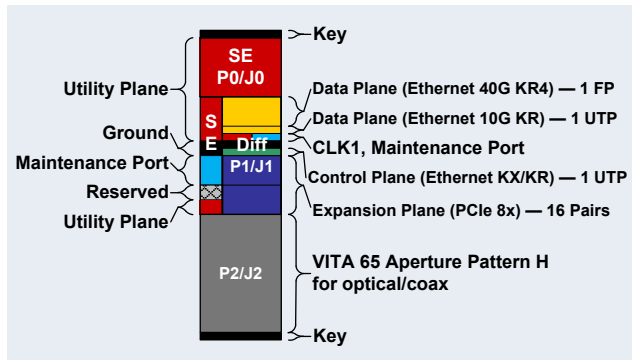
The figure below presents an example of this evolution for 3U form factor pin assignment. Starting from the popular 2F2U slot profile, a more complete, high-performance but general purpose single-board computer slot profile is defined.



// Figure 1: 3U payload profiles, current and next generation with 40G Ethernet I/O intensive

For the sake of clarity, we added in parenthesis the typical protocols expected on those slot profiles. This is actually described in the associated module profile. When looking at VPX slot profiles, one should always recognize the underlying protocols. In this case, Ethernet for the Data Plane and Control Plane, then PCI Express® for the Expansion Plane.

While the current trend is to move from PCI Express® to 40G Ethernet (or faster) for the Data Plane, PCI Express® will continue to be popular for the Expansion Plane.



// Figure 2: 3U payload compute intensive profiles with 40G ethernet dataplane and wider PCIe Expansion planes

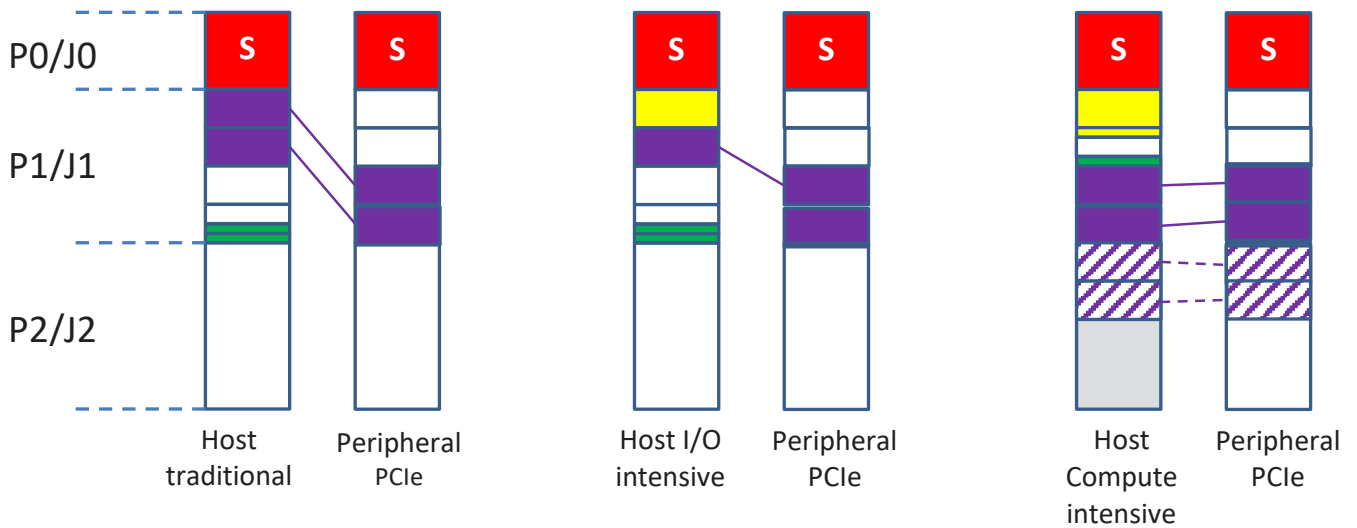
3U new profiles for next generation modular systems

Let us take a deeper look into the new profiles with 40G Ethernet Data Plane interfaces. For 3U, in addition to the second slot profile presented on Figure 1 which is generally referred to as the I/O Intensive slot profile, some computing payloads require faster and wider PCI Express® ports. These are the goals of so-called Compute Intensive slot profiles shown on Figure 2. PCIe as an Expansion Plane is allocated with a width of 8 or 16 lanes, leaving no pins for XMC mezzanine I/O, but provisioning instead Optical or RF interfaces. It can be noted that all these new profiles begin with 40G Ethernet interface at the top of VPX connector P1.

With this move to a 40G Ethernet Data Plane and the role of PCIe going to Expansion Plane, the need for PCIe switch modules has decreased. Standard 3U 1G/10G/40G Ethernet switches are still adequate to switch both the

Data Plane and the Control Plane, with typical slot profiles such as SLT3-SWH-2F24U-14.4.3 or SLT3-SWH-6F1U7U-14.4.14.

Following this tour of new 3U payload profiles for new generation systems, one could legitimately ask what profile to select for a peripheral module like a graphics card or other specialized I/O card. The two Compute Intensive profiles form the ideal platform for such devices. If the peripheral device is PCIe-centric, the Expansion Plane provides a wide, high-performance link between the peripheral module and the host (typically a single-board computer). If it is Ethernet-centric, the profiles offer both ultra-thin pipe (one lane) and fat pipe (four lane) Ethernet links, and even the provision for optical links. Lastly, should it be needed, a maintenance port is available for maintenance console access.

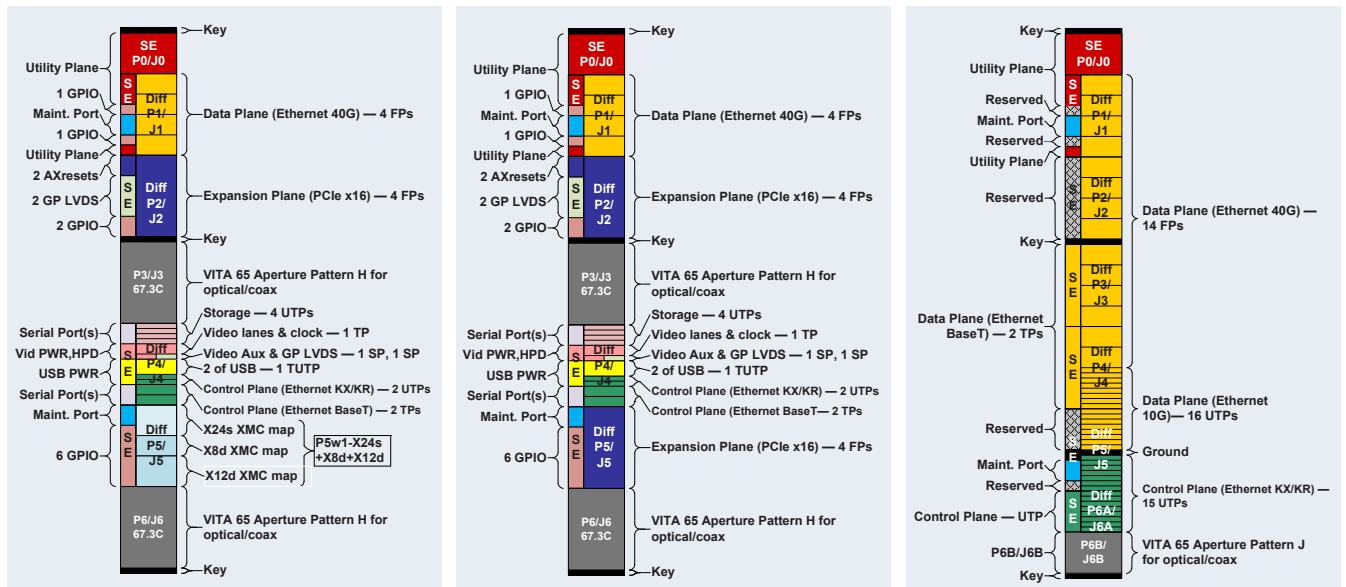


// Figure 3: PCIe peripheral pin assignment optimization

6U profiles for next generation modular systems

6U VPX systems exhibit fewer pinout constraints because of the higher number of connectors available, P0 to P6. However, the high number of user defined pins traditionally found in VITA 65 slot profiles creates at least a risk of interoperability issues between vendors.

A new set of modern 6U profiles has therefore been defined in the new version of the OpenVPX standard targeted to be released in 2019. Those slot profiles have been summarized in Figure 4 below, including the expected popular module profiles defining the underlying protocols shown in parenthesis.



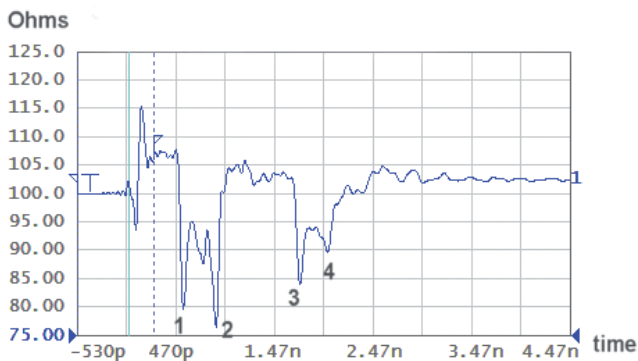
// Figure 4: 6U new VPX profiles for Payloads and Switch

VPX IS READY TO SUPPORT EMBEDDED SYSTEMS OF THE FUTURE

Does it still make sense to introduce into the VITA 65 OpenVPX standard new slot profiles close to 10 years after the initial release of the standard? The answer from our point of view is definitely yes. VPX technology is used in applications for the long term, and evolutions are needed for higher interconnect speed such as 40G Ethernet now, and 100G Ethernet in the near future. In addition, these new profiles incorporate lessons learned over the last 10 years of OpenVPX implementations.

Higher speeds now feasible

The interconnect speed on VPX copper backplanes has been so far limited to around 10 Gbits/s per channel, allowing 40 Gbits/s Ethernet KR4 (4 transmit+receive copper pairs on backplane). However, applications are reaching for 100 Gb/s Ethernet KR4 where each channel needs to run at 25 Gb/s. In a previous white paper "High data rates over the VPX infrastructure, revision 2 – 2017", we described the main reason for this limitation, highlighted in the impedance diagram below.



- 1: Tx module connector via
- 2: Tx backplane connector via
- 3: Rx backplane connector via
- 4: Rx module connector via

One can clearly see the impedance drop to less than 80 Ohms happening inside the press fit hole of the backplane connector and the module connector. This creates signal reflections back to the transmitter, altering the propagation of the bit and potentially mixing with the bits after.

To address this, main VPX connector manufacturers are introducing parts with smaller diameter press fit pins, both on the backplane side and the module side, so that the impedance drop issue is mitigated. The smaller via diameter helps to maximize the distance of the signals to reference planes, thereby decreasing the capacitance and increasing the inductance, thus keeping the impedance L/C close to the target value in the hole.

Those new higher speed VPX connectors are specified in the upcoming VITA 46.30 standard and are interoperable with existing connectors when speeds above 10 Gbits/s is not required. The implementation on VPX of the PCI Express® generation 4 at 16 Gbits/s and the subsequent versions will also benefit from those improvements.

Other VPX features ready to support the future


Going to even higher interconnect speeds may require using optical links, whether those links stay internally inside the chassis or go outside of the enclosure. Those optical links going from the VPX module, through the backplane via blind-mate connectors to the rear of the chassis, are well defined in the VITA 66 standard documents. Most of the new 3U and 6U module profiles presented in the previous chapters have a defined connector area for hosting optical links.

The VPX standard is also well designed to accommodate high power modules with appropriate cooling techniques. The last version of the standard recommends designing new modules based only on +12 V power supply. With a current capacity per module of 22 amperes for 3U and 28 Amperes for 6U, this equals to a possible power of respectively 264 Watts for a 3U module and 336 Watts for a 6U module. This is generally much more than needed, with true limitations coming from the cooling subsystem. The different cooling methodologies for VPX are normalized in the VITA 48 standard documents, and include air cooling, conduction cooling, and liquid cooling in different variants.

Of course, not all applications need the highest speeds and power dissipations. The VPX standard is also well suited for other modular computing use cases such as low power application, safety applications, etc. However, it is important that the supporting infrastructure and ecosystem, defined for the long term, remain the same in all cases.

KONTRON OFFERING IS LEADING THE WAY TO THESE NEW VPX ERAS

VPX products from Kontron fully embrace these new architectures and profiles. In 3U form factor, CPU boards based on the Intel® Xeon®-D processor and featuring 40G Ethernet are available in both I/O Intensive and Compute Intensive profiles. Those released products are shown in the figure below.




// Figure 5: Kontron VX305C-40G and VX305H-40G, up to 12 cores Xeon-D1500 CPU

For Ethernet switching in 3U, the VX3920 L3 switch in the next figure is recommended. It features up to 6x 40G Ethernet ports on the rear (each 40G port can also be configured as 4x 10G or 1G ports).



// Figure 6 : Kontron VX3920 40G/20G/10G Ethernet switch level L3

For 6U form factor, Kontron is releasing an L3 Ethernet switch with 72 lanes allowing 100G, 40G, 10G and 1G operations. This VX6940 product adopts the new VITA 6VITA 65 profile SLT6-SWH-14F16U1U15U1J-10.8.1 shown on the rightmost diagram of figure 4. 6U high performance processing boards will be available from Kontron using the two other profiles of figure 4. Figure 7 shows the current generation of 6U processing board from Kontron, the VX6090, featuring two high-end Xeon®-D1500 CPUs.



// Figure 7 : Kontron VX6090 dual Xeon-D1500 6U CPU module

CONCLUSIONS

The Ethernet 40G architecture evolution requires new VPX profiles, and the VITA65 OpenVPX standard revision 2019, under ANSI ratification process, paves the way for that. It takes the opportunity to narrow the number of profiles supporting this protocol and to improve vendor compatibility by removing user defined pins. It is already possible to evolve to 100G Ethernet for copper backplane with the use of higher performance VPX connectors, or to take advantage of optical and RF connectivity through the backplane.

This forms a rock solid baseline to implement embedded data processing machines of the future.

About Kontron – Member of the S&T Group

Kontron is a global leader in IoT/Embedded Computing Technology (ECT). As a part of technology group S&T, Kontron, together with its sister company S&T Technologies, offers a combined portfolio of secure hardware, middleware and services for Internet of Things (IoT) and Industry 4.0 applications. With its standard products and tailor-made solutions based on highly reliable state-of-the-art embedded technologies, Kontron provides secure and innovative applications for a variety of industries. As a result, customers benefit from accelerated time-to-market, reduced total cost of ownership, product longevity and the best fully integrated applications overall.

For more information, please visit: www.kontron.com



GLOBAL HEADQUARTERS

KONTRON S&T AG

Lise-Meitner-Str. 3-5
86156 Augsburg, Germany
Tel.: +49 821 4086-0
Fax: +49 821 4086-111
info@kontron.com

www.kontron.com