NASA/TM–2010–216872

Reference Avionics Architecture for Lunar Surface Systems

Kevin M. Somervill
Langley Research Center, Hampton, Virginia

Jonathan C. Lapin, Oron L. Schmidt
Johnson Space Center, Houston, Texas

December 2010

NASA/TM–2010–216872

Reference Avionics Architecture for Lunar Surface Systems

Kevin M. Somervill
Langley Research Center, Hampton, Virginia

Jonathan C. Lapin, Oron L. Schmidt
Johnson Space Center, Houston, Texas

National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23681-2199

December 2010

The use of trademarks or names of manufacturers in this report is for accurate reporting and does not constitute an offical endorsement, either expressed or implied, of such products or manufacturers by the National Aeronautics and Space Administration.

Available from:
NASA Center for AeroSpace Information
7115 Standard Drive
Hanover, MD 21076-1320
443-757-5802

Abstract

Developing and delivering infrastructure capable of supporting long-term manned operations to the lunar surface has been a primary objective of the NASA Constellation Program. Several concepts have been developed related to development and deployment of lunar exploration systems that provide critical functionality such as transportation, habitation, and communication, to name a few. Together, these systems perform complex safety-critical functions, largely dependent on avionics for control and behavior of system functions. These functions are implemented using interchangeable, modular avionics designed for lunar transit and lunar surface deployment. There are two core concepts in this reference avionics architecture. The first concept uses distributed, smart systems to manage complexity, simplify integration, and facilitate commonality. The second core concept is to employ extensive commonality between elements and subsystems. These two concepts are used in the context of developing reference designs for many lunar surface exploration vehicles and elements. These concepts are repeated as architectural patterns in a conceptual architectural framework. This report describes the use of these architectural patterns in a reference avionics architecture for Lunar surface exploration systems.

Contents

List of Tables

List of Figures

Acronyms

ACU	Audio Converter Unit	16
AIU	Audio Interface Unit	16
ASIC	Application Specific Integrated Circuit	4
ATHLETE	All Terrain Hex-Legged Extra-Terrestrial Explorer	1
AUP	Airlock Utility Panel	21
BDD	Block Definition Diagram	2
BE	Best Effort	15
BIU	Bus Interface Unit	2
C3I	Command, Control, Communications, and Information	7
CAN	Controller Area Network	9
CAS	Common Avionics and Software	3
CCDL	Cross-Channel Data Link	8
CDK	Chassis Driving Kit	37
C&DH	Command and Data Handling	1
CEV	Crew Exploration Vehicle	5
CMC	Crew Mobility Chassis	44
CODEC	Compressor-Decompressor	8
CSA	Common Services Assembly	2
CUP	Crew Utility Panel	20
CUR	Constellation Unpressurized Rover	37
C&W	Caution and Warning	17
D&C	Display and Control	17
DDCU	DC-to-DC Converter Unit	22
DDS	Data Distribution Services	16
ECLS	Environmental Control and Life Support	1
ECS	Element Communication System	25
EVA	Extra-Vehicular Activity	4
FPGA	Field Programmable Gate Array	4
GbE	Gigabit Ethernet	9
GPC	General Purpose Computer	4
HDMI	High-Definition Motion Imagery	17
IBD	Internal Block Diagram	6
IMA	Integrated Modular Avionics	10
IMU	Inertial Measurement Unit	38
ISHM	Integrated Systems Health Management	2

MFBC	Minimum Function Backup Computer	11
LER	Lunar Electric Rover	1
LRO	Lunar Reconnaissance Orbiter	16
LSS	Lunar Surface Systems	44
MIPS	Million Instructions Per Second	4
NAS	Network Attached Storage	16
NASA	National Aeronautics and Space Administration	44
OC	Operations Computer	3
ODN	Onboard Data Network	37
OLED	Organic Light Emitting Diode	19
OMG	Object Management Group	16
ORU	Operational Replacement Unit	4
PCC	Pressurized Crew Cabin	37
PCM	Pressurized Core Module	44
PDU	Power Distribution Unit	7
PEM	Portable Excursion Module	29
PM	Pressurized Module	30
PMC	Power Management Complex	22
PUP	Portable Utility Pallet	44
PSU	Power and Support Unit	44
RC	Rate Constrained	15
RF	Radio Frequency	17
RFC	Regenerative Fuel-Cell	23
RMUX	Remote Multplexor	5
RTP	Real-Time Protocol	8
SADE	Solar Array Drive Electronics	23
SCP	Self-Checking Pair	10
SSU	Structural Support Unit	23
SWN	Surface Wireless Network	7
SysML	Systems Modeling Language	1
TDMA	Time-Division Multiple Access	15
TT	Time-triggered	15
TT-GbE	Time-triggered Gigabit Ethernet	9
VPU	Video Processing Unit	8
WAN	Wide Area Network	7

1 Introduction

Developing and delivering infrastructure capable of supporting long-term manned operations to the lunar surface has been a primary objective of the National Aeronautics and Space Administration (NASA) Constellation Program in the Exploration Systems Mission Directorate. Several concepts have been developed related to establishing an architecture of exploration vehicles and assets to support this goal [Kennedy et al., 2010]. In each of these concepts, there are a number of key systems (or elements) that provide critical functionality such as transportation, habitation, and communication, to name a few. Together, these systems and elements perform complex safety-critical functions, largely dependent on avionics for the control and behavior of system functions. The following sections describe the computing and data architecture that has been assumed in developing reference element designs for various analyses.

The term "element" keeps with the general Lunar Surface Systems (LSS) project nomenclature and refers to things that would typically be called a vehicle such as a lander or a rover. It also refers to items such as habitation elements and in-situ resource utilization elements, but the general concept is that an element is a discrete system of systems enabling a defined function.

Throughout this report, various figures are used to illustrate the topic at hand. The majority of the figures are depicted using the Systems Modeling Language (SysML) notation [SysML v1.1]. The diagrams are from a SysML model developed for Lunar Surface Systems, but are occasionally simplified with components elided for clarity. The diagrams for the reference systems listed in Section 4 are more complete.

The organization of the remainder of the document is as follows. Section 2 presents a general overview of the surface architecture. Section 3 describes the basic avionics architecture applied to each of the elements. Section 4 presents overview descriptions of reference designs for select surface elements. This report concludes in Section 5 with a brief summary, some considerations for destinations other than the lunar surface, and potential future work.

2 Surface Architecture

The lunar surface architecture is comprised of multiple cooperative systems (or elements) used for surface habitation and exploration. The aggregation of, and the interactions between, these various elements is generically referred to as the surface architecture (including interactions with other Constellation systems). The primary elements considered in developing the surface architecture support transportation, habitation, mobility, power, and communications. Each of these elements are required to provide, to varying degrees, control functions such as maintaining thermal control, navigation, and Environmental Control and Life Support (ECLS). These control functions are managed (controlled and monitored) through Command and Data Handling (C&DH) systems.

The buildup of the surface architecture is performed over several years of discrete missions. The elements are delivered to the lunar surface using lunar lander vehicles (landers). Once on the surface, the crew lives in various elements supporting habitation. Depending on the length of the mission this may be the lander itself, the Lunar Electric Rover (LER) [Harrison et al., 2008], or a habitation element [Kennedy et al., 2010]. Various notional elements have been conceived to provide surface mobility for crew and assets. These include unpressurized rovers, the LER (a pressurized rover), and the All Terrain Hex-Legged Extra-Terrestrial Explorer (ATHLETE) [Wilcox, 2008, 2009; Townsend et al., 2010]. The ATHLETE is used to transport large cargo while the rovers transport crew and smaller assets.

Communication between the various elements is supported by peer-to-peer communication over a wireless mesh-network. Communication terminals provide wireless network connectivity and routing functions for surface elements; however, each element is capable of communicating directly with other elements. Together, the surface elements, including the communication terminals and landers, comprise the lunar surface wireless network. Communications can also be relayed through orbiting lunar communications satellites or via earth-based systems extending the range of communications beyond the lunar horizon.

3 Common Avionics Architecture Framework

The avionics provide the command and data handling (processing and storage); element automation; network connectivity and routing for surface elements; and Integrated Systems Health Management (ISHM). These functions are provided by systems of interchangeable, modular avionics designed for lunar transit and lunar surface deployment. Systems are optimized towards reuse and commonality of form and interface and can be configured via software or component integration for special purpose applications (i.e. card modules used for different applications in different chassis configuration loads).

The avionics systems for the various surface elements are envisioned to be designed around a common architectural framework. These systems are fundamentally computer controlled digital systems comprised of electronic assemblies with behaviors defined through software applications. A collection of assemblies are interconnected for communications and systems control. The systems and assemblies must be maintainable and reconfigurable to accommodate multiple applications in an evolving surface architecture. To facilitate this, the electronic systems are modular and employ a standard, stackable form-factor. C&DH systems are distributed throughout a host element and are largely comprised of computing elements, data bus networks, sensors, and actuators. The main components for computing systems include computers, network interface cards, mass storage, data acquisition units, and output drivers (i.e. remote I/O). These establish a foundation for a product-line approach to developing an architectural framework for the various surface elements. The following sections describe the characteristics and aggregation of components to realize subsystems hosted in the various surface elements.

The basic high-level architecture is depicted in Figure 3-1 in a Block Definition Diagram (BDD). The BDD is the primary diagram in SysML used to depict the structure or composition of a system. The diamond-headed arrows indicate a piece-part relationship between the element and its subsystems. In this BDD, the generic element is comprised of one or more bus interface units (Section 3.1.2), a common services assembly (Section 3.1.3), some number of network switches (Section 3.1.4), and a backup computer (Section 3.2.2). The Bus Interface Units (BIUs) provide for the distribution of interfaces to sensors and actuators. The Common Services Assembly (CSA) provides core computing and communications functions such as primary computers and network routing functions (refer to Section 3.1.3). The network switch may be embedded in the CSA or separately integrated with the host element, depending on element constraints. The backup computer provides dissimilar redundancy for critical software functions and performs safety and health assessment for the element. As illustrated later, various elements can realize alternate configurations using the same basic conceptual components.

There are two core concepts in the reference avionics architecture described in this report. The first concept uses distributed, smart systems to manage complexity, simplify integration, and facilitate commonality. This trend is on-going in commercial and safety-critical applications where we imbue systems and components which were previously simple electrical or even mechanical components with computational capabilities [Hammett, 2002]. While this increases the complexity of the

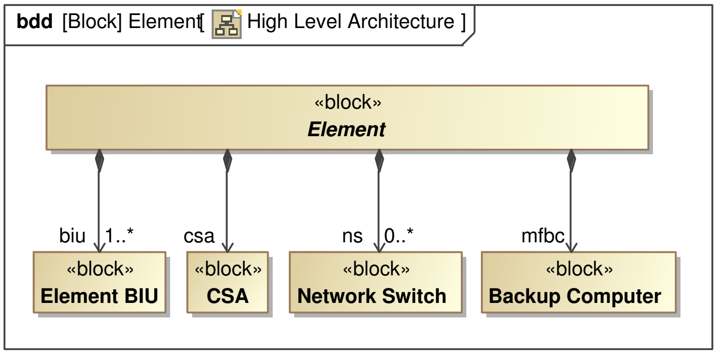

Figure 3-1. Block definition diagram of the basic computing architecture.

subsystems or component, these “smart” subsystems and components are easier to develop, integrate, use, and maintain. This is because we can partition and modularize behavior and interfaces for hardware and software systems to establish a collection of configurable, modular components that are reusable in multiple applications. This leads to the second core concept which is to employ extensive commonality between elements and subsystems. The concept of Common Avionics and Software (CAS) has developed to refer to a collection of modules that are applied to various elements and systems. This increases experience with fielded systems and reduces low-level development. It also simplifies reuse and maintenance activities. These two concepts are used in the context of developing reference designs for many lunar surface exploration vehicles and elements. These concepts appear constantly as architectural patterns repeated throughout this conceptual architectural framework. The following sections describe the use of these architectural patterns in the context of the reference avionics architecture for the LSS Project.

3.1 Computing and Data Assemblies

The primary components of the avionics systems are the computers and data systems. These are fundamental to command and data handling systems used to manage element subsystems. The computers are assumed to be high-reliability assemblies suitable for safety-critical applications. The computing system is developed around a distributed architecture. This reduces performance requirements for the control computers as well as potentially reducing power and integration complexity. Command and control are coordinated through a collection of primary computers establishing a locus of control. High-level control of the system is managed in the primary computers through goal-oriented control laws. Low-level control and application specific interfaces are distributed away from the primary computers, close to the controlled end-item (i.e. a sensor or actuator). As an example, the primary computers would support an application to manage temperature and airflow for ECLS in a habitation element. These commands would be sent to a controller in the air circulation unit where the local controller would accordingly adjust pulse width modulation signals to effect changes in fan speeds. It is also possible (for maintenance reasons among others) to “tunnel” to the low-level controller and execute precise commands; however, this is not required during nominal operations. In this way, the application (i.e. control law) can be abstracted from, although loosely coupled with, the underlying hardware. This is expected to simplify certification and integration costs as well as reduce hardware dependencies.

The computers are classified in three levels of performance: embedded controller, general purpose computer, and application specific computer. Collectively, these are referred to as Operations

Computers (OCs), OC-1, OC-2, and OC-3, respectively. Embedded controllers (OC-1) are low-level controllers, or possibly even state machines hosted in an Field Programmable Gate Array (FPGA), and they serve a specific set of functions. General Purpose Computers (GPCs) are class 2 OCs, or OC-2, and are more capable in terms of computing performance, comparable to traditional flight computers. Application specific computers (OC-3), like the embedded controllers, are intended for limited or specialized applications that require greater performance than efficiently realized using GPCs. Application specific computers can be realized using custom integrated circuits (i.e. Application Specific Integrated Circuits (ASICs)) or using high-performance FPGA devices, with the current trend favoring the latter. There is also the potential for multi-core and many-core technology to further blend the boundaries between these definitions. Because these multi- and many-core technologies are still fairly immature, they are not discussed here, so this report can focus on architecture and not on technology.

The OC-2 could be asserted as the primary computing resources among the OCs; however, all of the OCs are distributed throughout, and embedded within, the elements. The OC-2 is assumed to be capable of several hundred Million Instructions Per Second (MIPS) and consume less than 10W at 28Vdc. This includes a single board computer, network interface, and a power supply. It hosts a layered software architecture including a partitioned operating system, similar to ARINC-653, providing application partitioning and software modularity. Applications include thermal management, communications, vehicle navigation and guidance, and network services. One noted application is network file services supporting network attached storage.

The OC-3 is applied to more intensive computing. It is assumed that these utilize reconfigurable computing (such as FPGAs) for computationally intensive applications. Example applications include video and imagery processing, synthetic illumination, relative navigation, and sensor processing [Somervill, 2008]. These are assumed to be capable of the equivalent of 1000's of MIPS in processing for 15-25W at 28Vdc.

3.1.1 Form Factor

The various electronics assemblies adhere to a common, stackable form-factor. This assumption is driven by attempts to reduce mass and is supported by multiple NASA studies (Constellation Program Software and Avionics Integration Office (SAvIO) IDAC-4B Form-Factor Study and Lunar Architecture Team phase 2 studies). It must be acknowledged that considerations for serviceability have not been adequately addressed, but the primary metrics are related to reliability and maintenance. We have assumed that external assemblies are retrieved and replaced via Extra-Vehicular Activity (EVA) and that this would be on rare occasions. Assemblies serviced during EVA are maintained by box-level replacement for an Operational Replacement Unit (ORU) and lower level access (e.g. to replace modules in an assembly) is performed inside a pressurized environment. Studies to assess the ramifications of crew time related to the maintenance of stackable assemblies could indicate that the serviceability of card-in-backplane approaches outweighs the mass reduction of stackable assemblies, but that information is not available at the time of this writing.

Figure 3-2 shows one depiction of a stackable form-factor. Existing form-factors include PC/104 and mezzanine form-factors. Companies have developed space-qualified products based on the PC/104 form-factor; however, thermal management is an additional consideration as the commercial PC/104 specification does not account for conduction cooling. NASA and the Johns Hopkins University Applied Physics Laboratory (APL) have invested in developing space qualified stackable form-factors [Hodson, 2006; Seagrave, 2008; Cancro et al., 2009] but these are not standard. There are many other vendors developing stackable solutions for spaceborne computing, any of which could provide an opportunity to develop a standard suitable for lunar surface applications.

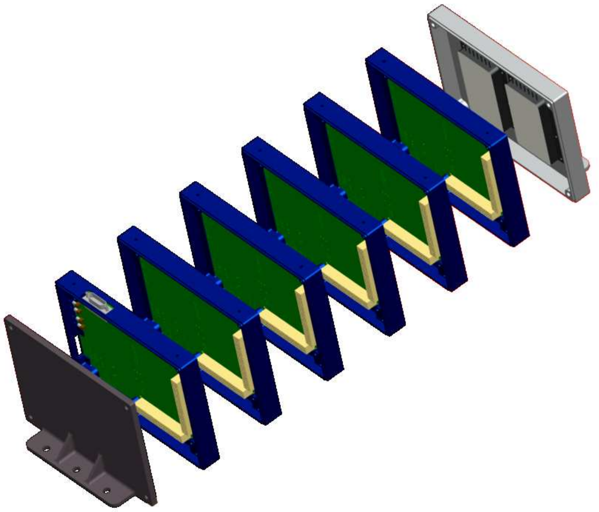

Figure 3-2. Illustration of a conceptual stackable form-factor.

3.1.2 Bus Interface Units

The BIUs are modular, network attached controllers which interface with subsystem specific functions. The BIUs provide the capability to distribute sensor and actuator interfaces in turn reducing cable harness mass by placing the digital-to-analog interface as close to the end-item as practicable. In this context, the BIU is analogous to shuttle Multiplexer/Demultiplexers (MDMs) or Orion power and data units [Anderson, 2010, pg. 11]. A similar concept, referred to as Remote Multplexor (RMUX) is employed in the conceptual lander design. The nomenclature "bus interface unit" is continued from the Crew Exploration Vehicle (CEV) smartbuyer study and is used here to clarify the distinction between the LSS reference concepts and heritage designs. Subsequent studies may result in migration towards an alternate or heritage design; however, our intent is to document the functions and current estimates for various characteristics of the BIU[1].

The primary purpose of the BIU is to facilitate communications with sensors, actuators, and subordinate systems. A typical BIU is illustrated in Figure 3-3 in the context of a host system. The element includes a control computer and an element BIU. The triangle-headed arrow from the control computer to the operations computer indicates a "specialization" relationship in that the control computer is a type of operations computer. The main constituents of a BIU are the controller, network interface, and subsystem control cards. In the figure, the controller is referred to as an embedded processor with the network interface embedded as part of the controller. The controller serves to validate commands from controlling computers and return telemetry from associated subsystems (i.e. aggregate distributed sensor and analog inputs). The BIU performs integrity assurance functions (such as voting or validation) for effectors ensuring a high-integrity communications path from the primary control system to the end-item actuator. The subsystem control cards are dedicated to various subsystems such as ECLS, power, and thermal management. An alternative segregation of interfaces could have been by interface type; however, it is believed (supported by studies performed by Orion and Altair[2]) that cards dedicated to subsystems enable greater optimization, decoupling between subsystems for power control and fault containment, and

[1] Arguably, this caveat applies to all of the systems and assemblies discussed throughout this report.

[2] The Orion and Altair studies have not been published.

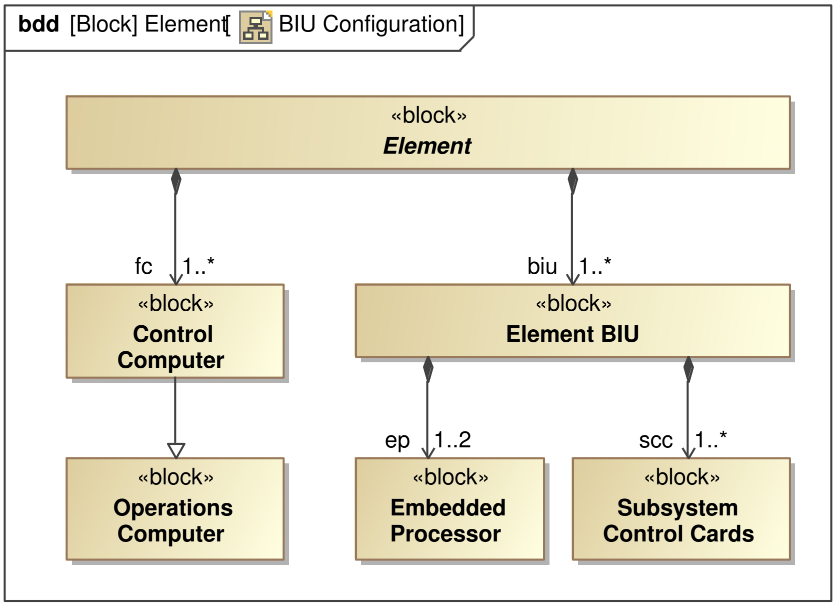

Figure 3-3. Generic vehicle Block Definition Diagram depicting BIUs

simplified integration, among other reasons [Anderson, 2010, pg. 11].

The system configuration is further illustrated in Figure 3-4 as an Internal Block Diagram (IBD). The IBD, as the name implies, depicts the internal connections of the respective system. In this figure, the control application is distributed between the control computer and the embedded processor in the BIU. High-level control is managed in the control computer which sends commands to the BIU processor over the network data bus. Status and telemetry are returned by the BIU controller to the control computer (typically part of the CSA in the host element) if required or requested.

Considering the BIU concept is intended to reduce wire mass, a few alternative concepts for wire harness reduction are also briefly discussed here. The first is that the BIU could be a wireless node, utilizing only a power connection, eliminating the data cable harness from the switch to the BIU which is estimated to be up to 10 kg per BIU. The BIU would connect to the network using the local wireless network (e.g. 802.11) using protocols to ensure the integrity of communications. Another alternative uses data over the power bus similar to that used in commercial home automation and control. The data rates available over a power bus are expected to be less than 1 Mbit/s which is likely sufficient for most BIU applications. Communications over DC power is currently being developed for industrial applications and while the concept is novel, its implications and challenges in manned space applications have not been thoroughly evaluated.

3.1.3 Common Services Assembly

The Common Services Assembly (CSA) is a concept leveraged from the development of the Altair lunar lander. Originally, the CSA was conceived as the primary computing assets for the lander accommodating both manned (i.e. with an ascent module) and unmanned (i.e. in a cargo configuration) missions. During its development, the CSA has evolved from a monolithic assembly to an aggregation of assemblies that collectively realize the core functions common to each of the various missions. This concept is used in major surface elements such as the Power and Support

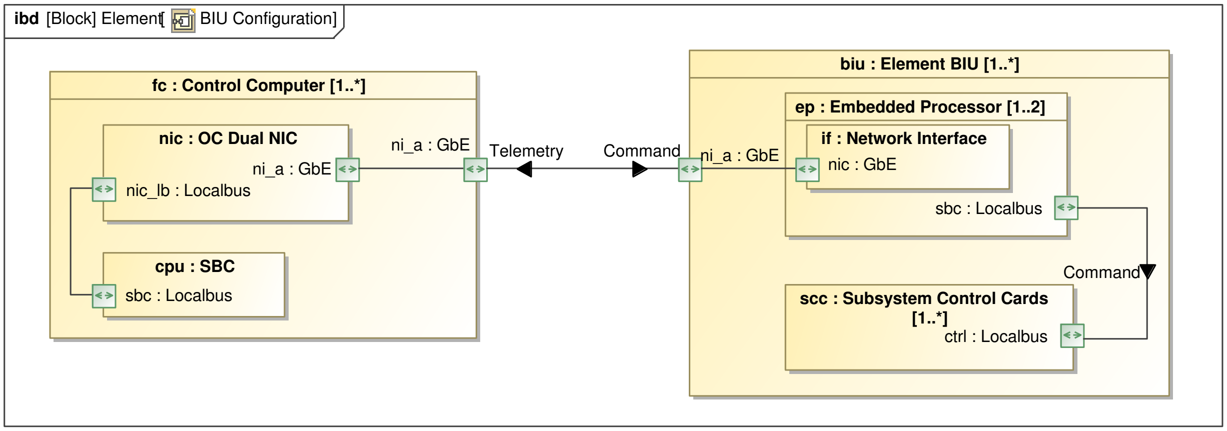

Figure 3-4. Internal block diagram depicting the interfaces between the control computer and the BIU.

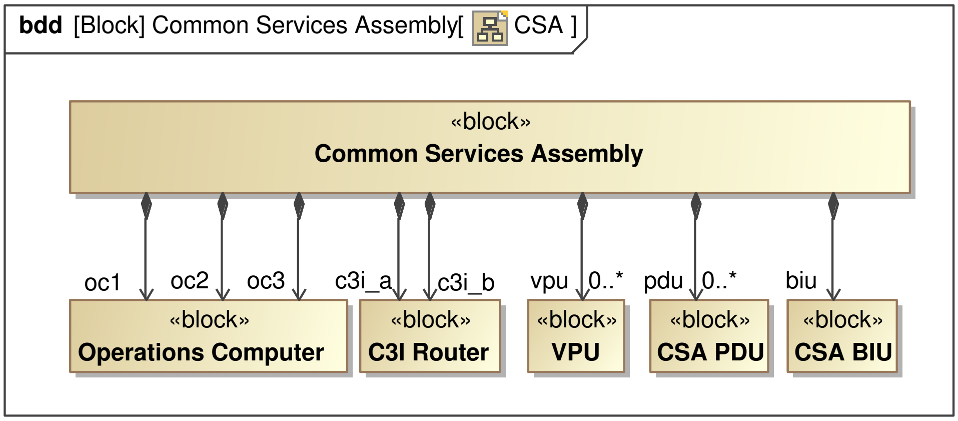

Figure 3-5. BDD for a typical Common Services Assembly.

Unit (PSU), Pressurized Core Module (PCM), or LER where the CSA is a common collection of electronics assemblies that are integrated into the element to provide core computing and control functions. The key functions of the CSA include primary computing, video processing, network and data routing, communications, thermal management, and power distribution. These functions are largely represented in the BDD in Figure 3-5. Those functions not shown are either realized via software applications or embedded as a subsystem function in the CSA BIU. For example, thermal management of the CSA is embedded in the CSA BIU as a subsystem control card as described in Section 3.1.2.

The IBD in Figure 3-6 depicts the internal connections of the CSA. The Command, Control, Communications, and Information (C3I) routers connect the assemblies of the CSA with the host element's local network and local communication systems. The local network traffic is routed through the router switches either to the computers or out the Wide Area Network (WAN) interface to the Surface Wireless Network (SWN). Power is supplied to the CSA via the CSA Power Distribution Unit (PDU) which could be integrated into the BIU or kept separate. The computers are synchronized and perform voting over the local network. An alternative configuration could

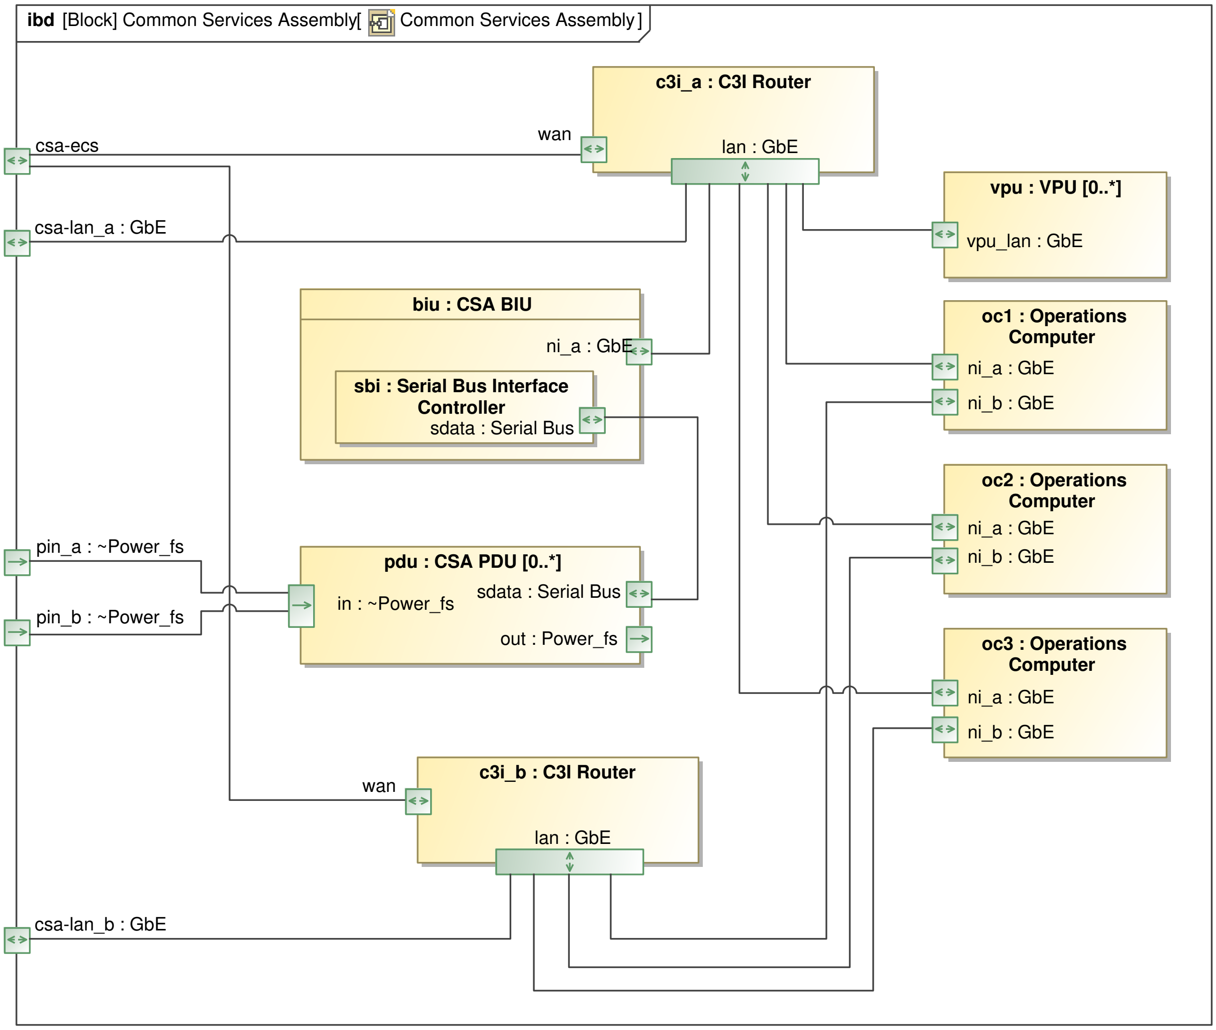

Figure 3-6. Internal block diagram for a typical Common Services Assembly.

be to implement a dedicated Cross-Channel Data Link (CCDL) between each of the processors. This alternative constrains the proximity of the CSA computers at the expense of harness mass. Assuming the local network can support fault-tolerant synchronization and agreement protocols (discussed further in Section 3.2.3), a dedicated bus is not required and using the local network allows for greater flexibility in system integration and communications.

The Video Processing Unit (VPU) is a network attached processor for Compressor-Decompressor (CODEC) processing of motion imagery. The VPU provides the ability to route and manage video streams from multiple sources to multiple destinations. Video received by the VPU can be merged with other video streams for viewing locally or remotely. The VPU can route motion imagery (merged or single channel streams) over dedicated links to crew terminals for real-time applications within the local host element. Otherwise, video is compressed and distributed over the local network using internet protocols for routing and distribution. It is assumed that the motion imagery is compressed in the H.264 format and distributed using something similar to the Real-Time Protocol (RTP) [Schulzrinne et al., 2003]. The VPU is described further in Section 3.3.4.

3.1.4 Onboard Data Networks

The elements host multiple data buses. The primary data bus for the surface elements is Gigabit Ethernet (GbE) and studies have included the use of Time-triggered Gigabit Ethernet (TT-GbE) which is discussed with fault-management in Section 3.2. Gigabit Ethernet is the specified data connection between Constellation Program vehicles; however, it has been employed internal to the lunar surface vehicles to facilitate high data rates for video and communications, as well as to accommodate science payloads and simplify integration. Internal to the elements, the GbE network is a mixed-criticality network hosting both critical and non-critical traffic. Critical traffic supporting fault-tolerant communications and element management is managed using time triggered protocols while non-critical, random data and high-bandwidth traffic (e.g. noncritical video) is handled as best-effort Ethernet traffic . The operation of these two traffic classes is transparent and assumed to be possible using TT-GbE to assure quality of service (QoS) and timing properties [AS6802 Draft]; however, other variants and protocols have the potential to meet the real-time quality of service requirements for the various elements [Rushby, 2003; Gwaltney and Briscoe, 2006]. More recent examples include FlexRay [FlexRay v2.1], and possibly even SpaceWire. This reference avionics architecture assumes that the primary onboard data network is TT-GbE and for the remainder of this report the term TT-GbE is used interchangeably with GbE.

Noting that many elements and subsystems may not require the bandwidth, or overhead of a GbE network, auxiliary serial data buses are also utilized. Examples are SpaceWire on the higher performance end of the spectrum or the Controller Area Network (CAN) bus at the lower performance end. These are used for close proximity (a few meters) low data rate connections, peripheral devices, and simple controllers. Example applications include interfaces to input controllers and subsystem interfaces such as those between a PDU and a BIU.

Considering the number of sensor and control interfaces distributed within each element and the size of some elements, such as habitation systems (e.g. core habitat, pressurized logistics modules, or pressurized crew cabin), there is an option to distribute the switch function across multiple, redundant switches. This reduces the length of cable runs to the network switches enabling subsystems and components to connect to the nearest switches. This configuration is shown in Figure 3-7 with four TT-GbE switches. Each of the switches is in turn connected to two other switches increasing network connectivity as well as resilience to switch failures. Another benefit of this distribution is that it reduces the thermal footprint of the required switching components by distributing the power load around the element. The fundamental principle here is that there are redundant paths between critical components. In this example, the CSA (see Section 3.1.3) hosts the primary computers and has redundant paths to the backup computer as well as the BIU controllers.

3.2 Redundancy and Fault Management

There are truly many considerations that come into play when considering system redundancy and fault management. During the development of the architectural concepts for LSS, a few key assumptions have been adopted. First the systems are at a minimum single-fault tolerant (1-FT) with the first fault fail-operational (i.e. no degradation in performance due to the failure). Ideally, the architecture realizes the property of composability such that when multiple elements are integrated or mated there is an increase in the fault-tolerance of the aggregate system. For example, mating two 1-FT habitation elements yields a 2-FT configuration. While this may seem to be a straight forward assumption, such configurations are non-trivial to implement and certify [Butler, 2008]. Another assumption is that the architecture can be realized using industry standard form-

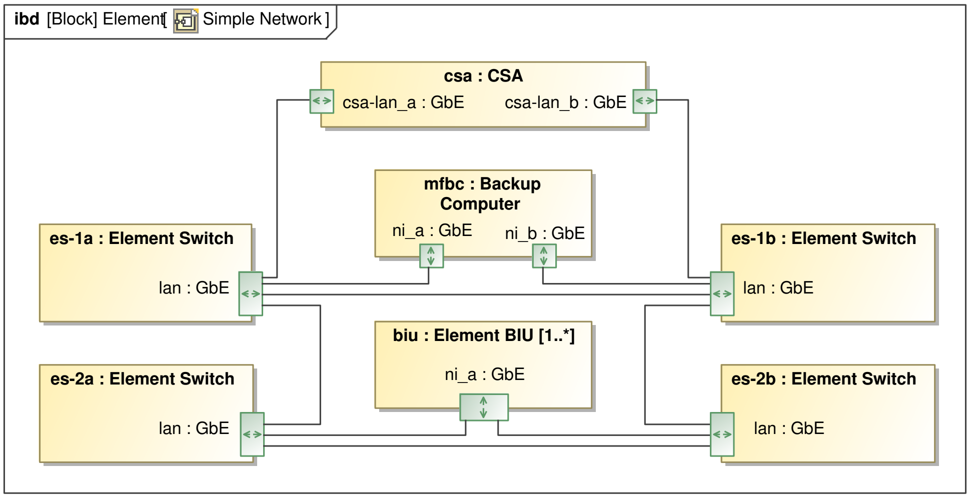

Figure 3-7. Simple network with distributed switches.

factors and components. FlexRay would be an example of this[3] [FlexRay v2.1]; it is also an example of using commercial vendor products. The final assumption is that the architecture will utilize a mixed criticality network. It is debatable as to whether this is the lower mass configuration, but intuitively, that conclusion seems reasonable[4]. This concept aligns with Air Force Research Lab development of a mixed criticality network for defense applications to enable quick integration and deployment as well as reduce system mass and power [MCAR]. When the various element designs realize a greater fidelity, this assumption can (and assuredly will) enjoy a thorough vetting.

3.2.1 Primary Computer System

The primary computing functions are realized in a complex of OC-2 class computers using an Integrated Modular Avionics (IMA) architecture [Watkins and Walter, 2007]. These could be standard single board computers (uni-processor or multi-core processors if the products are available) or Self-Checking Pairs (SCPs) similar to those used in the CEV Orion [McCabe et al., 2009]. The current architecture concepts assume the basic assembly is not an SCP, but a standard processor board that would be available as a standard product with comparable alternatives available from multiple vendors. The primary considerations in the selection of this processor configuration are application synchronization, power consumption, and product availability. Early concepts assumed SCPs due to the presumed ease of use; however, due to the expense in power and the limited vendor selection, the uni-processor is assumed. While SCPs do mask random hardware failures; they do not alleviate the need for inter-processor communication for state synchronization and agreement protocols. It is expected that either configuration will require significant software investment that can then be leveraged across the entire surface architecture. It is recommended that a detailed cost assessment including software, system reliability, and life-cycle costs drive the final selection as most of the architectural concepts discussed here are generic enough to accommodate either uni-processors or self-checking computers.

[3] TT-GbE may also be an example, when the standard is complete.

[4] An alternate approach is to utilized highly optimized, dedicated data buses for the various applications and subsystems. For example, a discrete network for video distribution and another for subsystems control separate from crew and science data networks.

Table 3-1. Property estimates for the primary and backup computers.

Assembly	Mass Estimate	Power Estimate (Nominal / Max)
Primary Computer Unit (OC-2)	3 kg	6 W / 10 W
MFBC	3 kg	3 W / 7 W
MFBC (w/ router and storage)	5 kg	9 W / 17 W

The primary computers operate in a time-synchronous fashion. Consistent state across the primary computers is managed through agreement protocols to ensure safe operation as well as detect and isolate faulty computers [Harper and Lala, 1991; Harper et al., 1988]. The voting protocols are routed over the local network. As noted above, the primary computers could utilize a dedicated CCDL; however, it is expected that a dedicated CCDL is not required. The driver for the multiplicity of computers is fault-protection from random failures, but also environmentally induced errors such as memory or processor upset. For critical operations or during periods of environmental events (solar particle events), the processors operate collectively to ensure the safety of crew and assets. However, considering that the time-to-criticality is fairly long during nominal operation of habitation elements and surface rovers, as compared to flight vehicles, systems could potentially run a single space-qualified computer monitored by the backup computer (described in the following section). This would allow for either reduced power profile or excess processing bandwidth for additional (non-critical) applications when the environment is less active. Issues of certifying the multiple configurations required to host alternate applications are complex and important, but are not germane to this discussion and not addressed here.

3.2.2 Backup Computer System

There are a number of reasons for including a backup system including common mode failures, fault-tolerance, and power modes. In the proposed architecture, a Minimum Function Backup Computer (MFBC) is added to the network as shown in Figure 3-7. This is a dissimilar, potentially lower performance, computing system responsible for tracking the operations and commands from the CSA as well as maintaining the habitation system through quiescent periods. The MFBC is connected to both TT-GbE channels where either channel can be powered down to reduce power utilization. It will utilize the TT-GbE network for time synchronization and potentially operates as a Time-Triggered entity. Alternatively, the MFBC can participate in synchronization protocols with the CSA computers across the C3I routers. The backup computer has an integrated router function such that it does not require the CSA to manage the element resources. This is to accommodate low power modes when the CSA is quiescent or during emergency situations such as partial network failure.

Table 3-1 lists the basic mass and power estimates for the primary and backup computer systems. These values represent current best estimates without margin for any uncertainty in the design. The listing for the primary computer is for a single unit which includes (in this estimate) a single board computer, a network interface and redundant power supplies. The network interface provides connections to redundant networks. The MFBC is a lower performance computer compared to the primary computer, but is envisioned to include greater functionality including a router and data storage to reduce reliance on other systems and mitigate CSA failure.

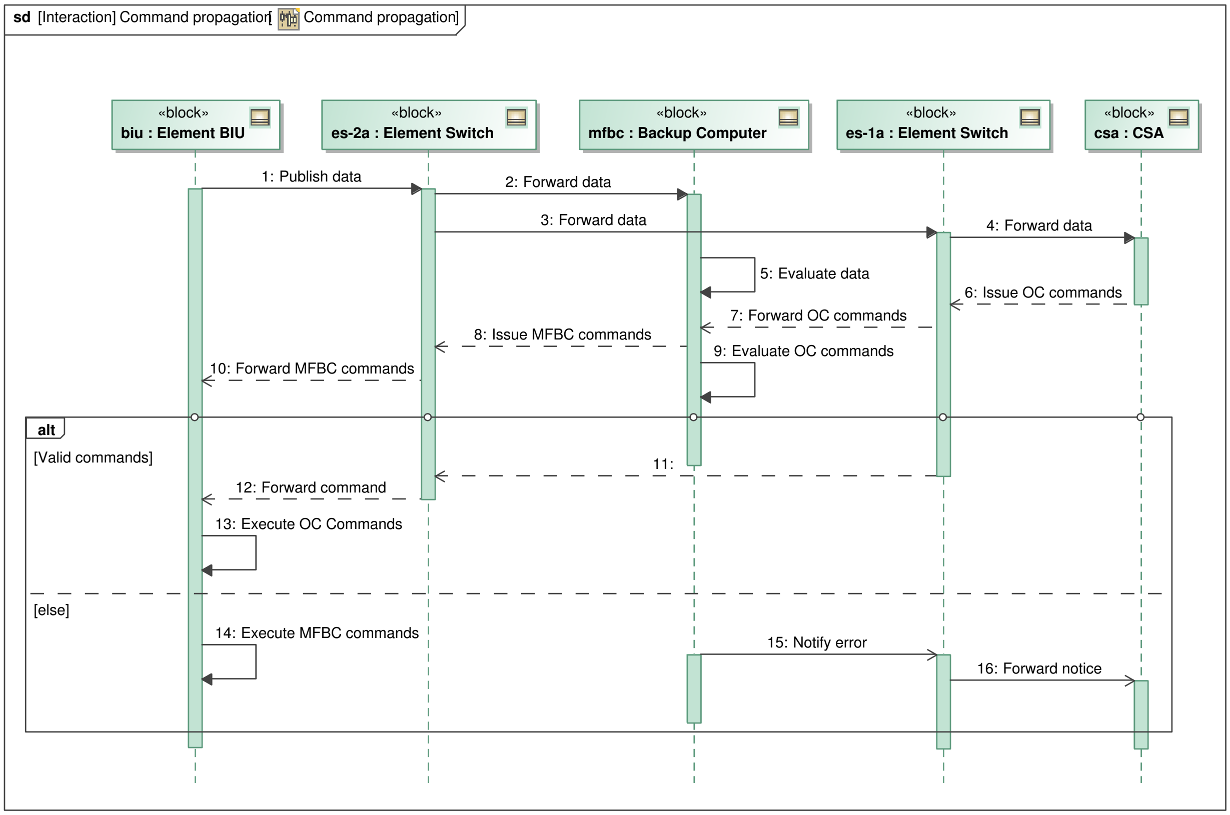

Figure 3-8. Sequence diagram of command propagation and validation.

3.2.3 Data Integrity

The computing elements in the CSA perform an agreement protocol internal to the CSA. Messages are agreed upon and signed prior to distribution to network elements. The signed messages are sent by each processor through both channels with the end-point ultimately receiving at least six command copies . The end-effectors, for example, then have potentially three discrete messages (in duplicate though). The effectors then respond to the first valid pair of commands in agreement and from discrete sources.

The process is illustrated in Figure 3-8. In this example, a BIU is publishing data regarding one or more subsystems. This data is forwarded to the backup computer and to the CSA via the second network switch (es-1a in the figure) in the path. Both the backup computer and the CSA OCs evaluate the data and generate commands. The CSA issues commands to the BIU via intermediate network switches in the element. These commands are also received by the backup computer which evaluates the commands to assess correct operation. The MFBC may then issue commands that are also forwarded to the BIU for execution. The BIU then receives duplicate commands from the MFBC and each of the OCs in the CSA from which the BIU selects which commands to execute. The BIU will execute the first pair of commands from separate OCs in the CSA that match. If no agreement is determined between the CSA OCs, the backup computer will be included to select a matching pair of commands. In the event that no two computers agree, the commands from the backup computer are selected for execution.

Figure 3-9 illustrates the current assumption using non-voting network interface between the

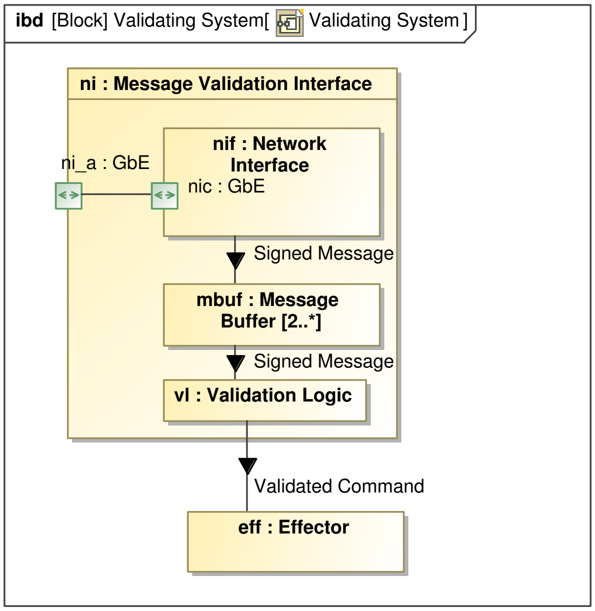

Figure 3-9. Validating network interface.

computing systems and the end-effector. This option receives and validates incoming commands for protocol compliance and valid messages are queued to the buffers. The buffered messages are then compared and only the validated, matching commands from discrete sources are forwarded to the effector. This can be implemented as either self-checking paths using an X-Y lane approach or a single interface from the network interface to the effector, depending on criticality and reliability requirements. This architecture has the benefit of a common interface that is independent of the function of the effector. It does, however, rely on signed and validated messages to ensure integrity and authenticity of messages and origins. This is similar to that described by Hanaway and Moorehead [1989] for the Orbiter main engine throttle controller.

In Figure 3-9, a single interface is shown between the validation logic and the effector. Depending on the application, there may be a need to continue redundant paths to the effector. The compare logic would serve to simplify the validation required by the effector. The effector would only respond to commands received from both links from the compare logic, but would not have to validate message signature.

The process of message validation is shown in Figure 3-10. The process is initiated by signals that a message has been received from either a primary computer in the CSA or the backup computer. The messages are buffered for validation which checks for correct signatures and errors in the packet. When sufficient valid message have been received, the valid message is issued for execution. In the event that none of the primary messages can be validated, the backup computer message is issued. Otherwise, an error is reported to the backup computer.

Figure 3-11 illustrates an alternative network interface that requires voting logic either between the network interface and the end effector (as shown) or embedded in the effector. The architecture also has a common network interface which is potentially simpler than that shown in Figure 3-9 and does not rely on message signing protocols. The additional voting logic is unique to the individual end-effector and requires two matching commands or three messages that can be averaged. The activities performed in the voting process are shown in Figure 3-12.

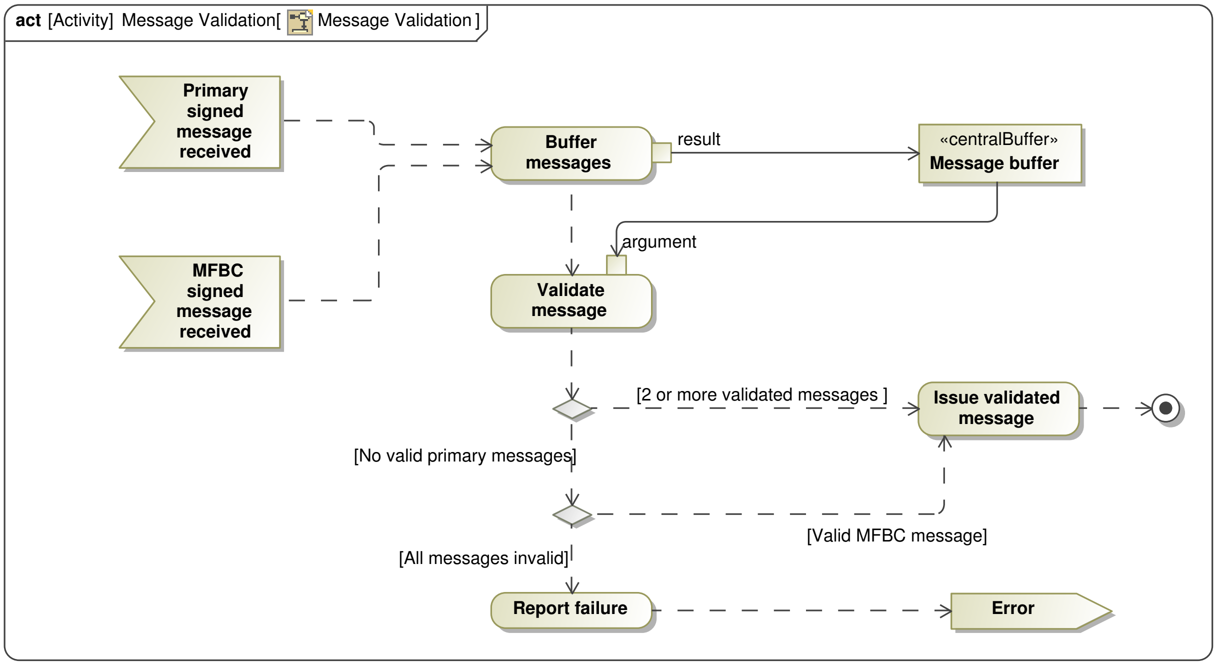
act [Activity] Message Validation[Message Validation]
Primary signed message received
MFBC signed message received
Buffer messages
result
«centralBuffer» Message buffer
argument
Validate message
[2 or more validated messages]
Issue validated message
[No valid primary messages]
[Valid MFBC message]
[All messages invalid]
Report failure
Error

Figure 3-10. Activity diagram of validation process.

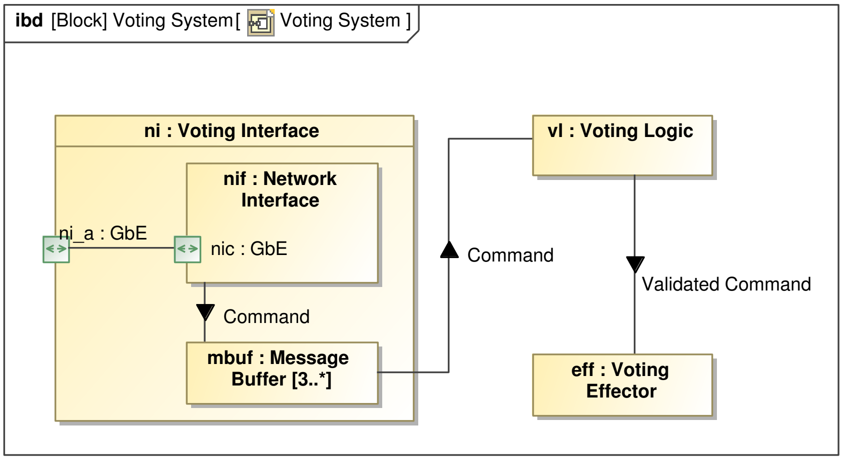

Figure 3-11. Voting network interface.

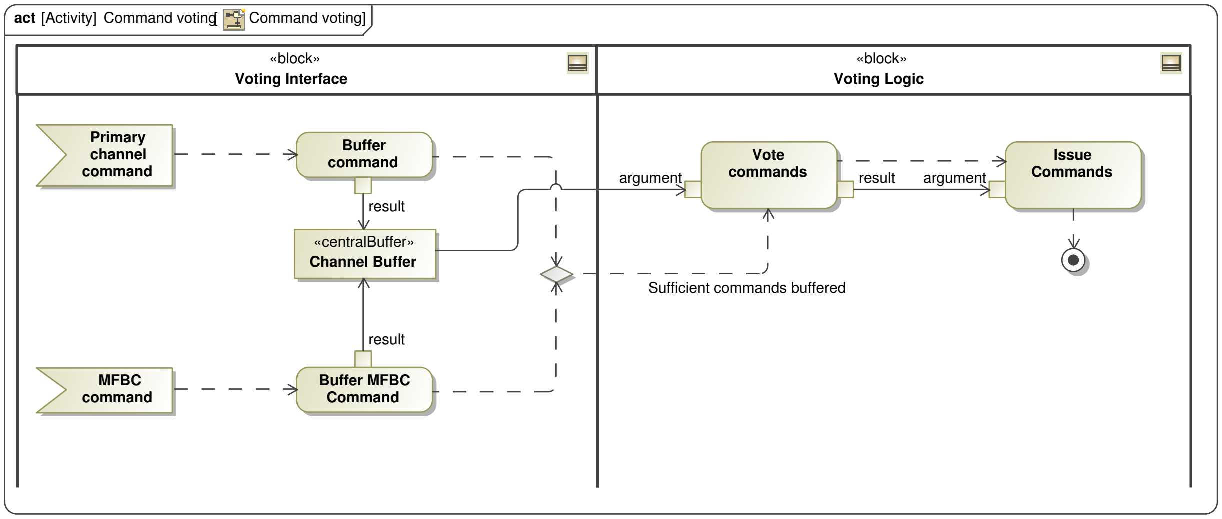

Figure 3-12. Voting a command set.

3.2.4 Clock Synchronization

Clock synchronization is fundamental to distributed, real-time applications such as those used in safety-critical systems [Cristian and Fetzer, 1995; Kopetz et al., 2004]. The computer clock synchronization can be performed using a data network in the CSA or provided by the TT-GbE network. In the former, the architecture is similar to a CCDL architecture for synchronization with the processors accessing separate channels for sensors and effectors. In the latter, network timing is provided by the TT-GbE network with the clock domain extended to incorporate the CSA.

For the option where clock synchronization is local to the computing elements, the computers perform clock synchronization and data agreement protocols over a network local to the CSA. In this configuration, the timing of the operation of the element systems is managed as software scheduled Time-Division Multiple Access (TDMA) in a master/slave configuration from the perspective of control systems. Timing is managed by the computing complex and commands are executed according to the time schedule within the CSA computers. Commands are issued to subsystems which respond in an event driven process. The TT-GbE switches forward data packets as either Rate Constrained (RC) or Best Effort (BE) traffic. For the discussion of control architecture, critical data is routed as RC traffic with guaranteed bandwidth and bounded latency. This is similar to a CCDL approach using switches for the channels and cross-strapping between the channels.

In the option with the TT-GbE timing extended to the computing elements in the CSA, the timing is provided by the TT-GbE network and resolved across the entire network [AS6802 Draft]. This however carries the potential issue of compatibility between the TT-GbE network and the CSA implementation. In this configuration, critical traffic is also routed using the Time-triggered (TT) traffic class. System operation is managed using a global schedule with network events (commands and responses) occurring as TT events.

There is also the option for a loose coupling between the two time domains. This may introduce additional latency for the synchronization across the boundary. It also adds the complexity of managing which domain is the time master in a safe and robust manner. Paulitsch and Steiner [2003] describe clock synchronization across a network of distributed clusters for fault-tolerant applications.

3.3 Data Management

A primary function of the computer network within and between surface elements is the management of data, including collection, storage, and distribution. Each of the large elements provides Network Attached Storage (NAS) services to other elements and crew.

3.3.1 Publish and Subscribe Services

Telemetry streams are broadcast to the network as services such that other elements may *subscribe* to the service to receive status and monitor the remote system. This process of broadcasting data streams is referred to as *publishing* data. One possible way to realize this functionality is to use Data Distribution Services (DDS) such as that described by Object Management Group (OMG) standard [DDS v1.2]. This enables systems to configure access to data sources as the surface infrastructure evolves and to accommodate varying mission scenarios. Access to the published data streams must be protected by security and authentication protocols.

3.3.2 Storage

As described above, large elements provide NAS services for other surface elements and redundant storage where required. This supports audio, video, and imagery storage as well as software application and science data storage. It is expected that this storage will be as large a practicable for long-term exposure in the space radiation environment. Current estimates are several hundred gigabytes (GB) to a few Terabytes (TB) in a single unit. Memory trends indicate an approximate scaling factor of two every two years. Launched in 1994, Clementine had 2.1 Gbits of storage in its solid state recorder [Garrett et al., 1995]. Applying the trend to estimate potential range of storage size yields an estimate of 2020 - 1994 = 26 years or 15 steps[5]. Doubling memory density at each of the 15 technology steps gives 2^{15}, or a factor of 32,768. Applying this factor to the 2.1 Gbits of storage on Clementine yields an estimate of 2.1 TB. Using Lunar Reconnaissance Orbiter (LRO) as a reference point, the estimate is in the range of 800 GB to 1.6 TB so it seems credible that the lunar surface architecture could have data storage units ranging from hundreds of GB to a few TB to support surface operations.

3.3.3 Audio

The crew audio from headset microphones is digitized for local intercom distribution. Audio is compressed for onboard recording and transmission to other elements. Transmission to external elements can be enabled by push-to-talk switches whereas internal audio is enabled by voice operated switches. To realize these functions, audio hardware samples and converts analog signals to digital formats such as G.711 and G.729 for distribution over the network. The two primary components are the Audio Interface Unit (AIU) and the Audio Converter Unit (ACU). The AIU provides interfaces for hardline connections to audio communication hardware. Wireless systems transmit digitally formatted audio data which is routed to the ACU for packet generation. An example of this is shown below in Section 3.4.2 describing utility panels.

[5] The observant reader deduces that $26/2 \neq 15$. The 15 steps accounts for the time delay from part selection during the design phase until the eventual launch. We've assumed a conservative four years from part selection to launch.

Table 3-2. Property estimates for the video processing unit.

Assembly	Mass Estimate	Power Estimate (Nominal / Max)
Video Processing Unit	5 kg	16 W / 25 W
Video Processing Unit (optimized)	3 kg	6 W / 8 W

3.3.4 Video

The various video streams distributed throughout LSS drive a fair amount of complexity and performance in the element data systems. As such, video becomes one of the key drivers for using a high bandwidth network such as TT-GbE. The primary resource used to manage video is the VPU which is also an example of an applied OC-3 processing system. The VPU receives uncompressed video from cameras over High-Definition Motion Imagery (HDMI) interfaces and compresses the video for distribution over the local network. The resulting video can be distributed in higher quality for local use or lower quality if it has to be transmitted over Radio Frequency (RF) networks. The primary subassemblies in the VPU are the controller board, network interface, encoder/decoder processing boards, and a power converter. The controller and network interface may be integrated as a single module. Estimates of characteristic properties are listed in Table 3-2. The first entry assumes reconfigurable computing based on FPGAs to realize the video processing. The second entry assumes an ASIC development specifically designed for the VPU to reduce power and mass. Both reflect processing capacity for four concurrent video channels. Either can be extended by adding additional video processor boards.

Uncompressed video may be distributed to displays internal to the host element if the application requires. For example, critical applications may not be able to tolerate dropouts in motion imagery packets resulting in a need for a dedicated video network between select assets . Once video streams are processed by the VPU, they are distributed over the local network. The video stream is stored in network attached storage which supports file streaming on request from any of the surface elements. Video is stored until it is transmitted to ground for storage and distribution. Once the video has been transmitted, it is deleted unless it has been marked to be saved.

3.4 Display and Control

Display and Control (D&C) systems are the primary systems for crew interactions with surface elements, allowing the crew to monitor and control both local and remote assets deployed throughout the lunar surface outpost. The D&C systems are comprised of a number of components, collectively referred to as crew workstations which provide various functions and services. The workstations are attached to the network and can access and display all stored data as well as data streams published to the network. The primary service supported by the workstation is command and control. They also support health and status services differentiating criticality and notification of fault events. Also, the workstation supports the Caution and Warning (C&W) system providing notification to the crew. The workstation display layout is managed via software configuration to suit operational modes and user preferences. The display renders text and imagery (still and motion). Data and telemetry formatting is managed partially through DDS services using self-describing data and format schemas.

A standard D&C interface and subsystem has been conceived to support a number of different LSS elements. This will provide a consistent interface for command and control and data presentation to the user simplifying training and operations and reducing risk of operator confusion. While

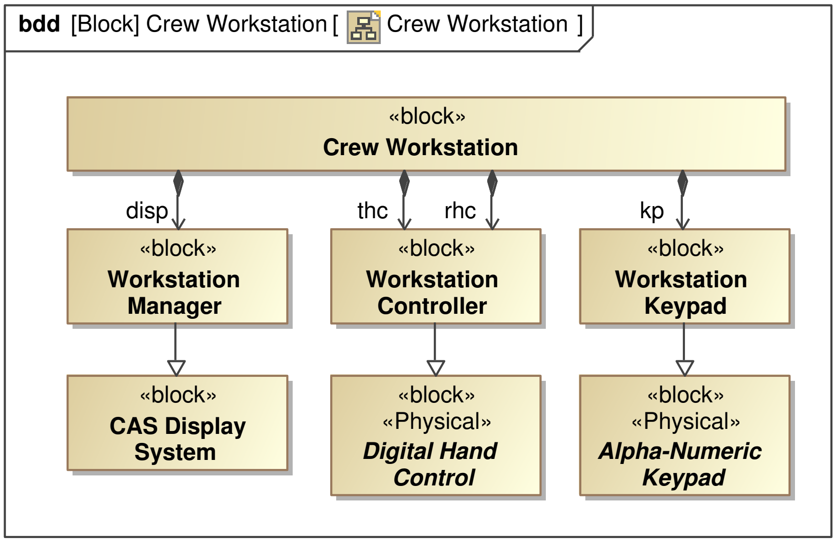

Figure 3-13. Composition of the crew workstation.

there is an intention to develop and deploy a common D&C system across LSS and the lunar lander; there are a number of considerations that must be addressed. Consideration must be given to form-factor compatibility across all of the applications. There may also be dependencies on requirements and software which are unique to the individual vehicles. The D&C systems are expected to be used primarily in a pressurized environment; however, they must also accommodate gloved-hand operation in both nominal and emergency activities in a vacuum environment. This indicates the need to understand usability issues across all of the operational modes of the individual elements.

3.4.1 Crew Workstation

The D&C system is comprised of a number of configurable components collectively referred to as the crew workstation. The high level composition of a crew workstation is shown in Figure 3-13. The primary components are the workstation manager, the workstation controllers, and keypad. The workstation manager is a CAS display system comprised of a display management computer and a display unit. The workstation manager is described in the following section. The workstation controllers and keypad are types of input controllers which are described below.

The crew workstation components connections and interfaces are shown in Figure 3-14. The workstation manager is primarily comprised of a display management computer and a display unit. The display computer is one of the classes of operations computers. It is connected to the TT-GbE network and also has an interface for direct video input for display. The display can receive compressed imagery via the TT-GbE network compressed in H.264 format. Compressed video is locally decompressed and displayed to the user. The data and imagery received from the network and video input are integrated and rendered to the display according to the layout specified by the user or source data. The current assumption for the interface between the display manager and the display unit is HDMI. The workstation manager communicates with input controllers (rhc, thc, and kp in the figure) via a local serial bus, like that used by the BIU to communicate with subordinate sensors and actuators.

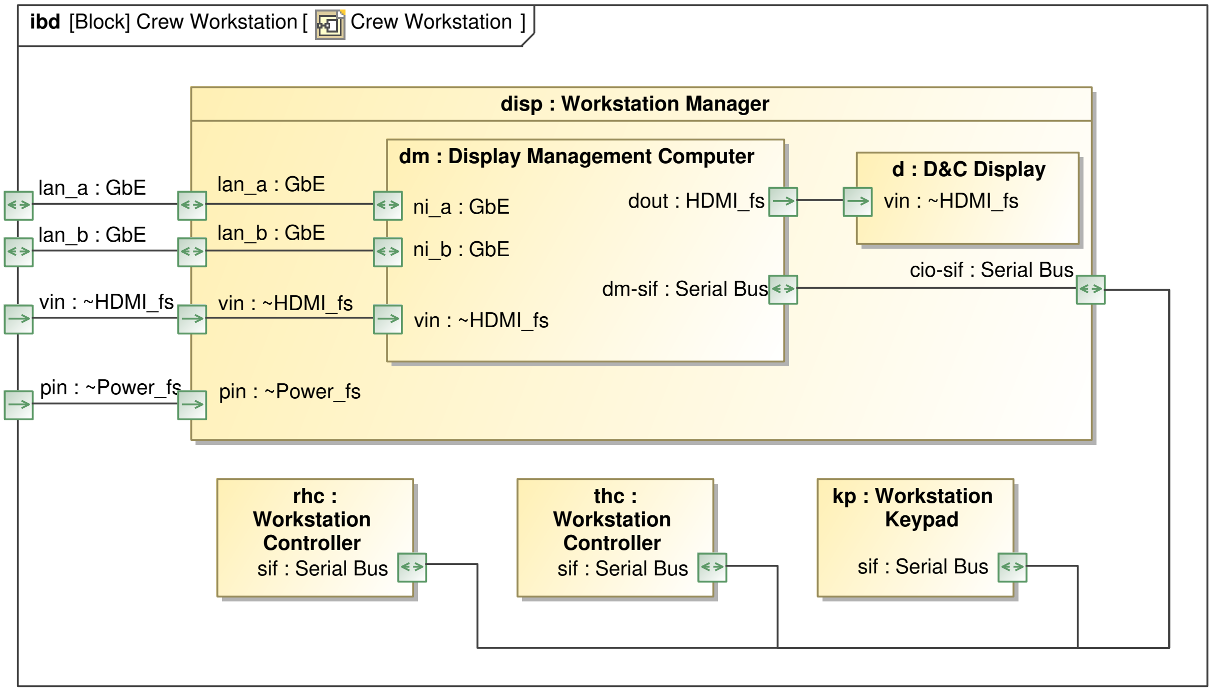

Figure 3-14. Connectivity in the crew workstation.

Workstation Manager The workstation manager is based on the CAS display system and is comprised of three assemblies: the display management computer, the display unit, and a power supply (shown in Figure 3-15). The display management computer is an intelligent network device capable of receiving commands from a control computer, decoding video streams, metadata overlay on displays, and generating the display which is then sent to the display unit for viewing. This configuration has several benefits including decoupling the display unit from the display management computer simplifying maintenance and upgrade activities. This also reduces network traffic by only transporting data across the network which is then processed by the display computer to generate the display at the workstation. This functionality is realized using an operations computer as the display management computer and integrating some of the functionality from a VPU for CODEC processing at the workstation.

The display unit is connected to the display manager via a dedicated bus (assumed to be HDMI or something similar). It is a "simple" device for display of text and imagery to the user. The technology is currently advancing at a rapid pace, but is currently expected to be similar to or better than Organic Light Emitting Diode (OLED) technology.

Input Controllers There are a number on input controllers available to the user at the crew workstations. Cursor control is provided by devices such as a mouse, trackpoint, miniature joystick, or touch pad; although, not all of them are amenable to gloved-hand operation. An alphanumeric keypad is provided for crew input. The crew workstations provide soft keys that are both programmable and mode specific to simplify operations. Also, hand controllers, similar to flight vehicle hand controllers, provide rotational and translational input commands to operate systems, vehicles, and manipulators. Due to the variance in usability of input controllers between pressurized and unpressurized environments, multiple options are available to maximize crew efficiency. Some

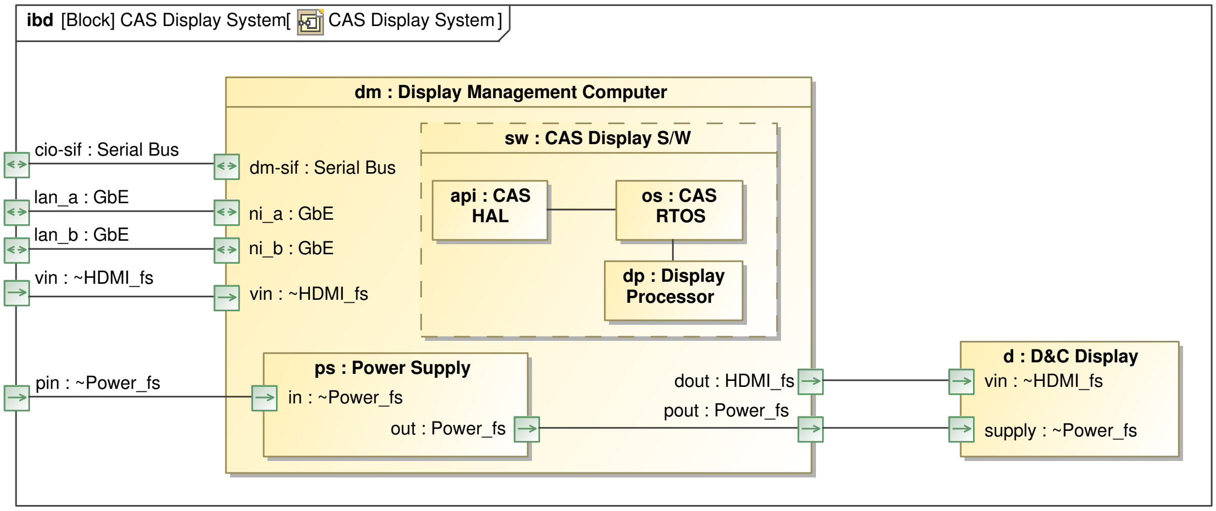

Figure 3-15. Common display manager and display system.

Table 3-3. Property estimates for the utility panels.

Assembly	Mass Estimate	Power Estimate (Nominal / Max)
Crew Utility Panel	4 kg	12.5 W / 22 W
Airlock Utility Panel	3 kg	4.5 W / 9 W

options, however, are mode dependent and may not be available in all operational environments. For example, a standard keypad may be comfortable in a pressurized environment, but not effective for a suited/gloved operator. Each of the input control devices is connected via a local serial bus.

3.4.2 Utility Panels

The habitable elements (habitation elements and large rovers) have utility panels to support crew operations. The current reference designs account for two classes of utility panels: the crew utility panel and the airlock utility panel. Both include common low power processors and hardline network interfaces. The two are differentiated by additional functions described in the following paragraphs. Estimates for the mass and power of the utility panels are listed in Table 3-3.

Crew Utility Panel The Crew Utility Panel (CUP) provides auxiliary services to simplify crew communications and operations. A CUP is intended to be installed in each habitable volume as it provides caution and warning alerts and tones for crew safety. The CUP provides system status via a display panel, but does not provide command functions beyond querying systems status; the crew would utilize a crew workstation for command and control. The CUP supports audio functions including a speaker, microphone, and wireless receiver for wireless headsets. Conceptually, the CUP could be a portable tablet with a mounting interface at a fixed location for hardline access to the element's network and recharging CUP batteries. The CUP provides network access via hardline network interfaces for portable devices (like USB or Ethernet) as well as wireless network connectivity as a wireless access point for access to the local area network. The CUP is intended to operate in a pressurized environment, but in a contingency will provide critical functions when

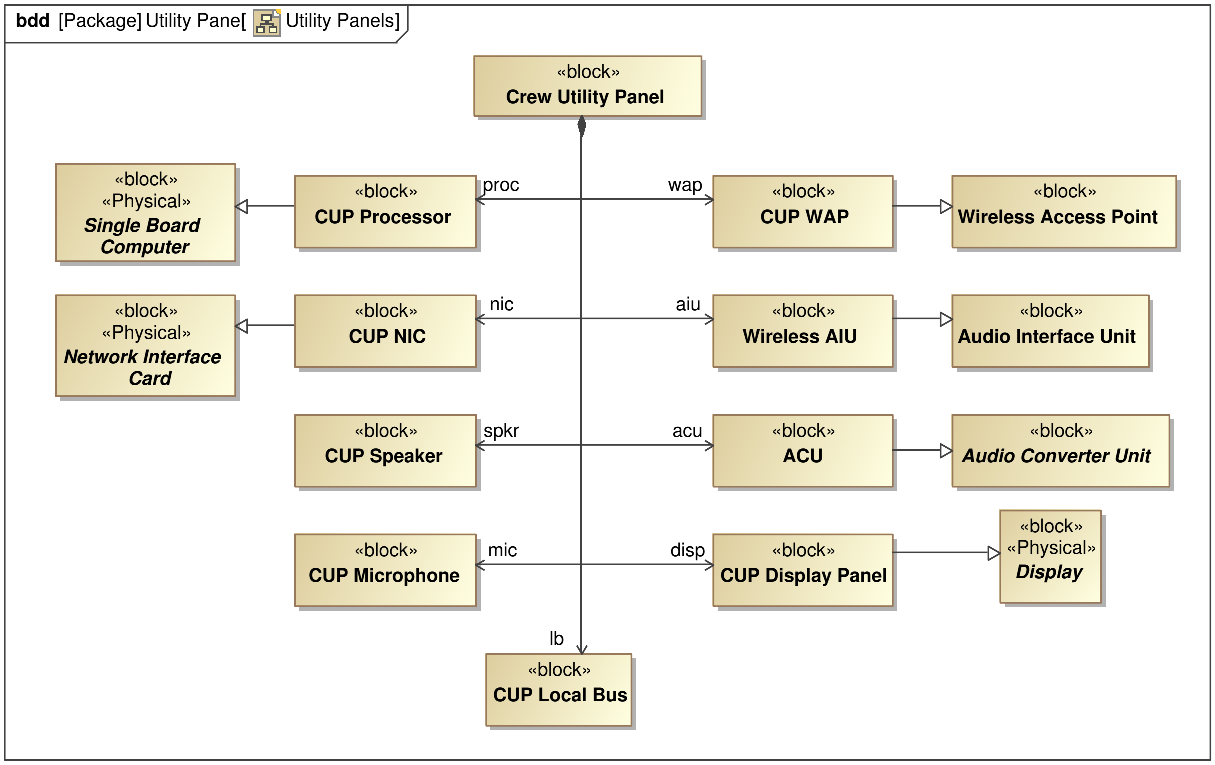

Figure 3-16. Composition of the Crew Utility Panel.

operated in a vacuum (for example, when an LER is depressurized to retrieve an injured crew member). The main components of the CUP are shown in Figure 3-16 while Figure 3-17 depicts the CUP interfaces.

Airlock Utility Panel The Airlock Utility Panel (AUP) was conceptualized to support operations by crew preparing for, or returning from, EVA operations. The AUP is simplified version of the CUP including functions to support: habitat health monitor and display screen, EVA checkout antenna (S-Band), integrated caution and warning, and suit umbilical port. The AUP provides health monitoring and status to the crew via a status panel with display and indicators as well as limited capability for commanding element subsystems. To support EVA operations, the AUP is integrated with the EVA interface panel which provides recharge and monitoring functions for the suit. The AUP has an EVA checkout antenna for communication tests prior to umbilical separation and egress and it also provides network connectivity for voice and telemetry between the EVA suit and the element. The AUP is designed for vacuum operation and does not require a cold-plate (active cooling), but does require a conduction path for heat rejection.

4 Major Elements

The previous sections present the basic concepts and functions in this reference architecture. In this section, these various architectural components and subsystems are aggregated to realize reference designs for a few of the major LSS elements. The elements presented are categorized as support, habitation, and mobility systems. The assortment of elements is not comprehensive,

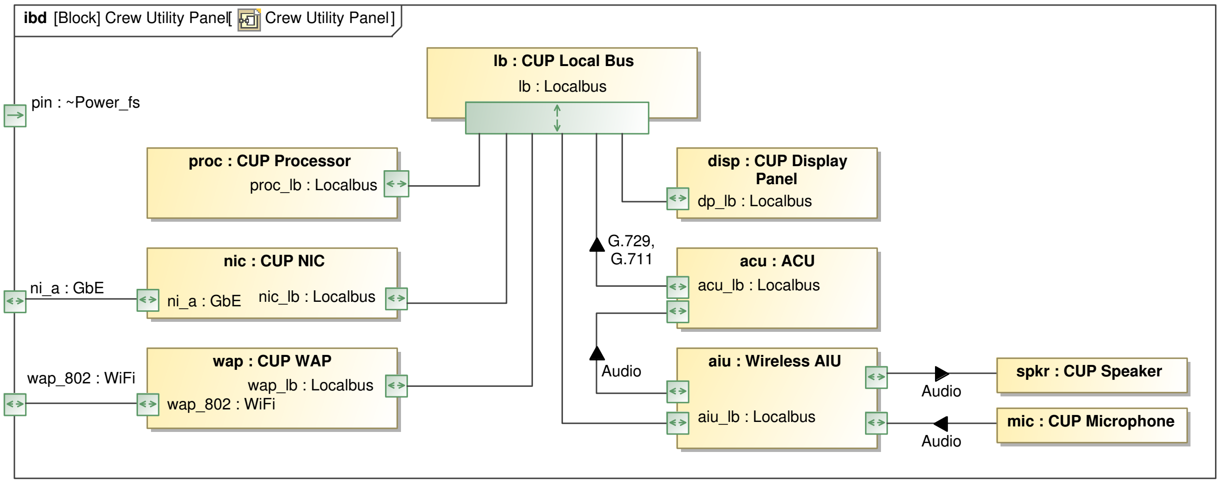

Figure 3-17. Internal block diagram of the Crew Utility Panel.

but does support discussion, illustrating the application of this reference avionics architecture to various elements.

4.1 Support Systems

The first class of major elements included in the discussion in this report is support systems. Support systems provide resources to primary systems like the habitation elements, mobility systems, and communications elements. The following sections describe the reference designs for the Portable Utility Pallet (PUP) and the PSU. Both the PUP and PSU are used to manage the storage of consumable resources, specifically power, oxygen, nitrogen, hydrogen, and water. As the name implies, the PUP is portable while the PSU is used to transport large payloads and support larger element operations such as a habitation module. When deployed, a habitation element would be permanently attached to a PSU.

A primary function of the support systems is to provide power. This functionality is managed by a collection of systems and assemblies referred to as the support systems Power Management Complex (PMC). This is shown in the block diagram in Figure 4-1. The PMC is comprised of a controller, power switching and distribution units, power converters, batteries and charge controllers, and a BIU. The connectivity of the PMC is shown in the IBD in Figure 4-2. The PMC is intended to be used as a dual redundant system using cross-strap connections to recover from failed components. Main power (assumed to be either unregulated solar array output or high voltage in the range of 300V to 400V) is received by the DC-to-DC Converter Unit (DDCU) and regulated to 120V for distribution by the PDU. Power is routed from the PDU to the power switching units for load switching. Two PMCs are cross-strapped at the PDU interface where the output of one DDCU is routed to the inputs of PDUs in both PMCs using the cross-strap interfaces (pxtrap_out and pxtrap_in Figure 4-2). Control of the PMC is managed by the controller and the BIU. Commands are received from vehicle/element level command and control via the TT-GbE network at the BIU. Command messages are routed to the controller for processing. The controller is integrated into the BIU to reduce mass and simplify integration. Also, the controller is expected to be either a class 1 or class 2 operations computer depending on what other functions (software) are allocated to it. Local control applications run on the controller issuing commands through the BIU to the various subassemblies. BIU communicates with power units (distribution, switching,

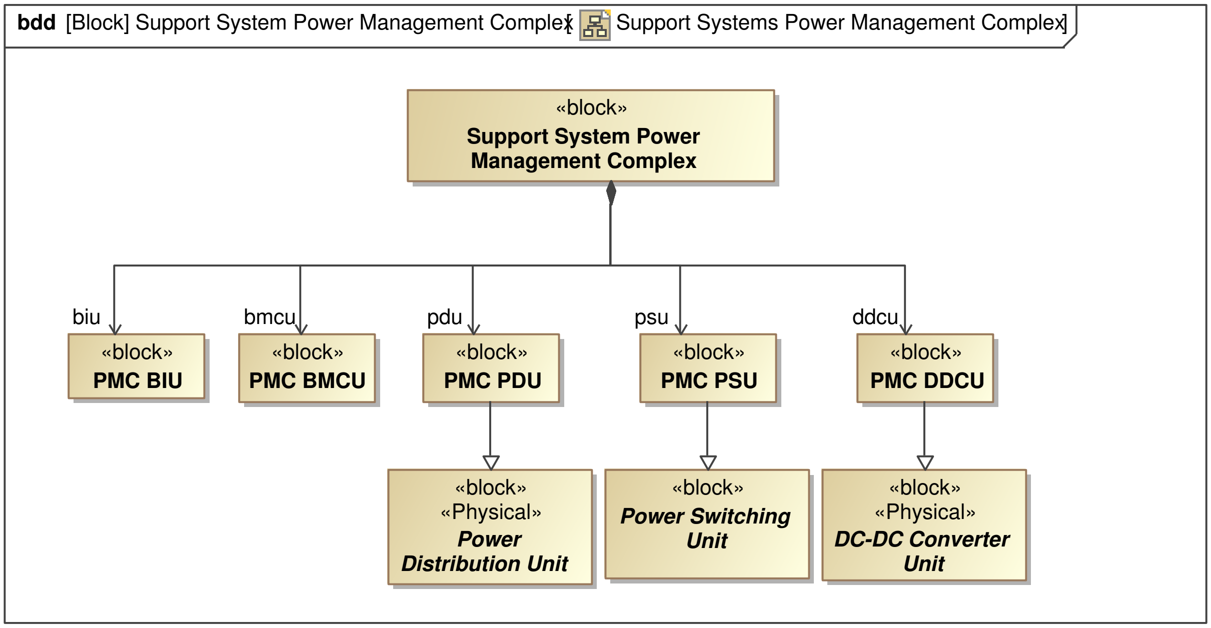

Figure 4-1. Power management complex for the support systems.

etc.) via a lower data-rate serial bus such as EIA-422 or SpaceWire. The BIU also hosts sensor and actuator interfaces such as those for thermocouples, heaters, and relay controls.

4.1.1 Power and Support Unit

The PSU combines a Structural Support Unit (SSU), used to secure and support cargo, with solar power generation, an energy storage system, a communications package, logistics storage, and resource scavenging and transfer equipment. The SSU provides the interface to the lander for habitation and pressurized logistics elements, pressurized rovers, and other cargo and payloads delivered to the lunar surface. PSUs provide the capability to sustain habitats, provide keep-alive power to surface systems and landers, provide consolidated storage of consumables, and facilitate resource scavenging from landers. The PSU can be comprised of either a battery-based or a Regenerative Fuel-Cell (RFC) based energy storage system. The structural unit is designed to work with the ATHLETE heavy-lift mobility system to facilitate cargo offloading and handling operations. The PSU supports several possible configurations for various elements, multiple power storage elements such as batteries and RFCs, and other storage units such as cryogenic tanks. This section describes one possible configuration of the PSU which includes RFCs. Figure 4-3 depicts the current PSU concept with an SSU, several tanks, and an avionics bay.

The composition of the PSU is shown in the BDD in Figure 4-4. The first component is the SSU which provides structure for the element. The PSU can support both solar arrays and RFCs. For solar array configurations, the PSU has a solar array interface and Solar Array Drive Electronics (SADE). The SADE is a special class of BIU specific to solar array control and operation. The PSU provides recharge function for swappable batteries used by portable elements such as the Crew Mobility Chassis (CMC). The PSU includes redundant power management complex to provide power to PSU subsystems, battery recharge, and mated cargo. The charging interface is a bidirectional interface to transfer power and consumables between the PSU and other ele-

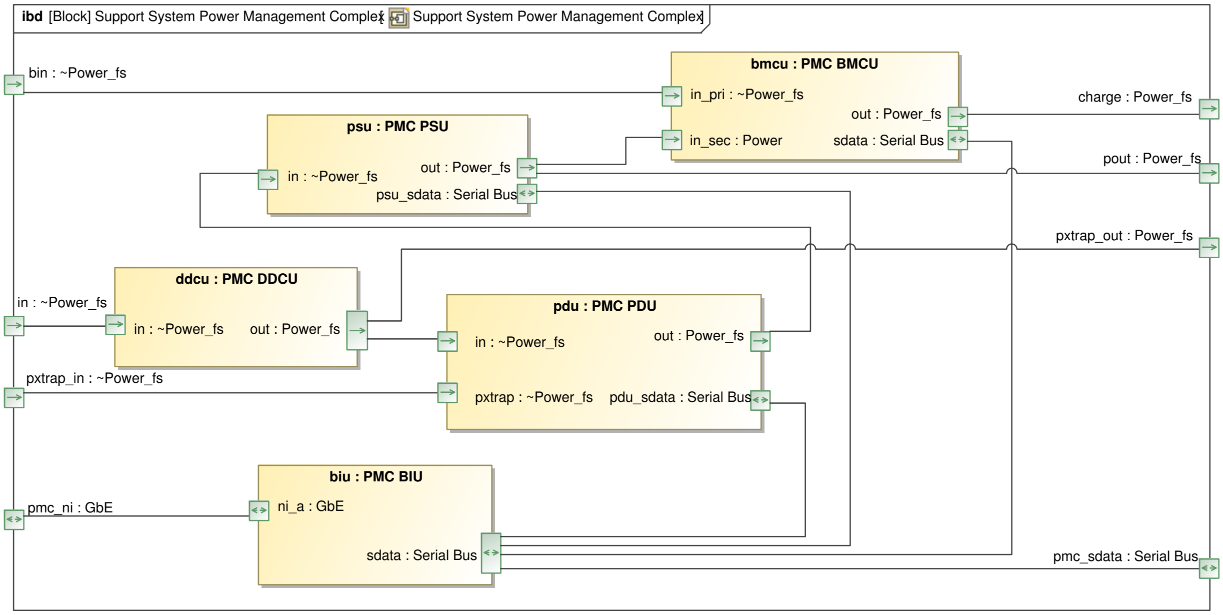

Figure 4-2. Internal view of the support systems power management complex.

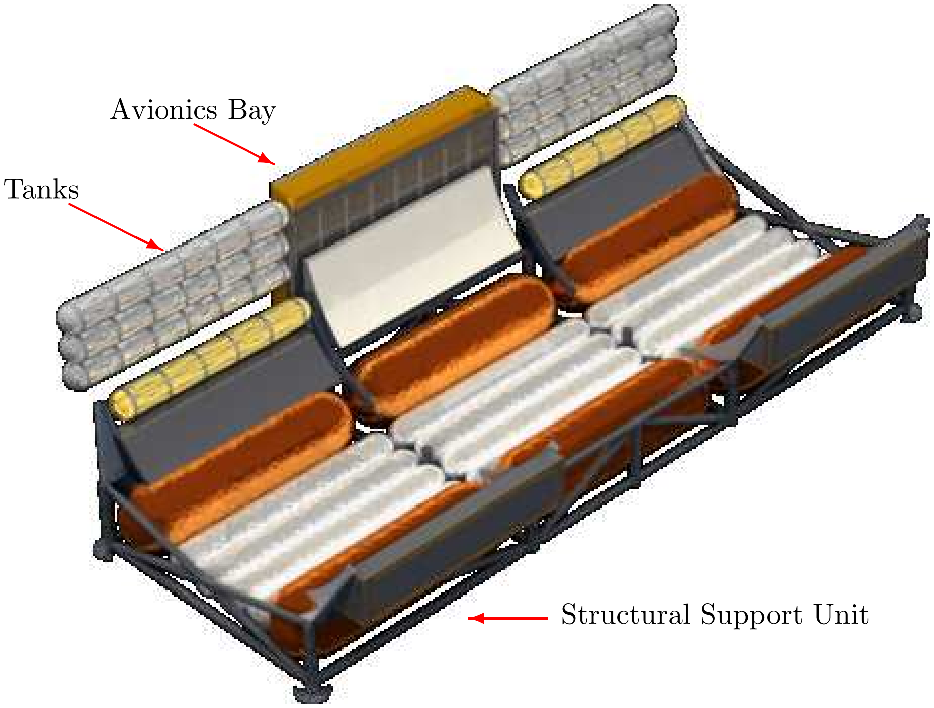

Figure 4-3. Power and Support Unit configured for a habitation element.

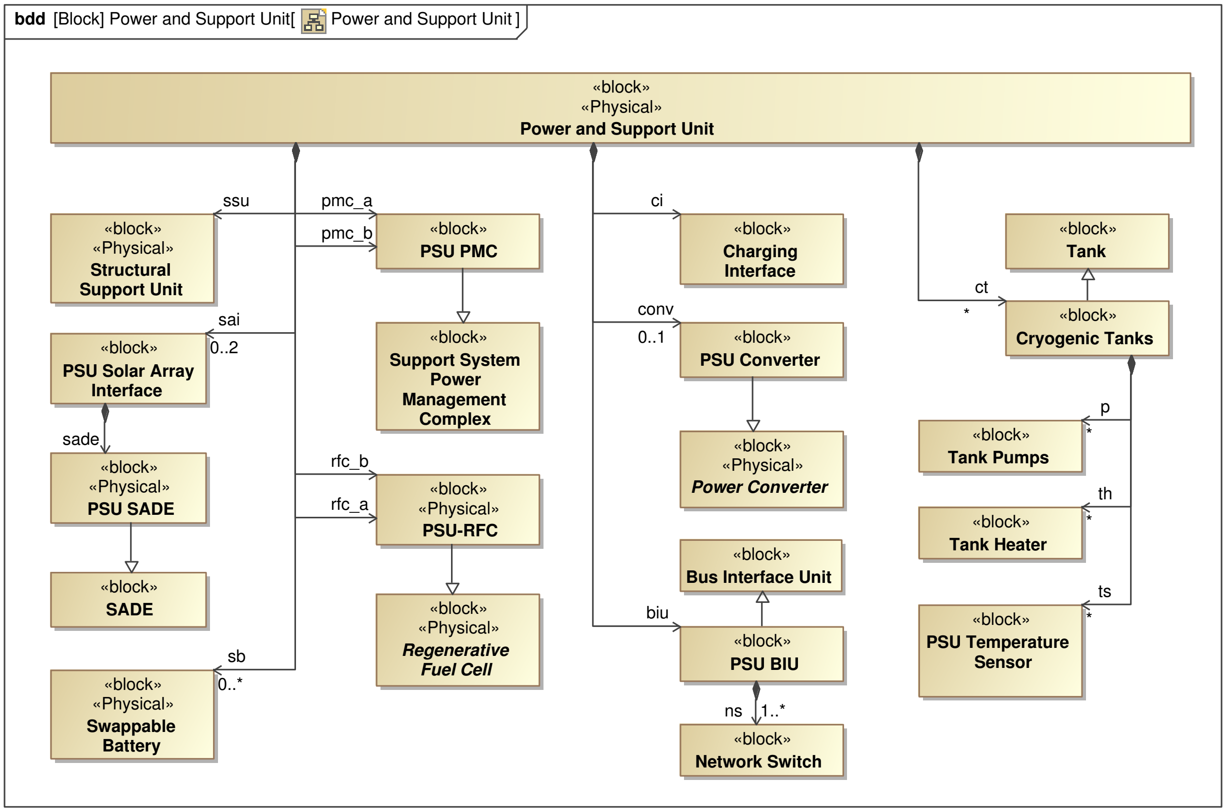

Figure 4-4. Major subsystems and assemblies in the Power and Support Unit.

ments. The final element shown is the power converter which supports options where there is a shared power supply for multiple elements (i.e. an outpost power generation element). The PSU power converter converts high-voltage, and possibly AC, power to the voltage needed by the power management complex.

The internal connectivity of the PSU is shown in Figure 4-5 and Figure 4-6. The configuration shown has solar array input with power routed through the SADE to the redundant PMCs. PMCs provide primary control for the PSU including command and control interactions with other elements (including mated payloads). The PMCs provide for battery charging and the batteries can also provide emergency power during solar eclipse and RFC failures. Communications with other elements are received through redundant TT-GbE interfaces to the PMC BIU and controller. Local communications from the PMC controller are sent via the PMC BIU over a local serial data bus as described previously. The PSU BIU includes an integrated network switch and also provides for monitoring and control of tanks, valves, and pumps. The PSU does not provide network routing functions and does not incorporate an Element Communication System (ECS). If either of these is required, the router can be integrated into the BIUand the ECS connected to a TT-GbE port on the BIU network switch. Figure 4-6 depicts the connectivity of the PSU BIU and tank management assemblies. The BIU includes the network switch connecting the network in the PSU to cargo or communications systems via the redundant TT-GbE ports. The processor in the BIU runs applications for managing PSU resources (such as cryogenic tanks) and also manages PSU functions such as command processing and health monitoring. Future reliability assessment will determine the appropriate redundancy for the BIU and its subassemblies.

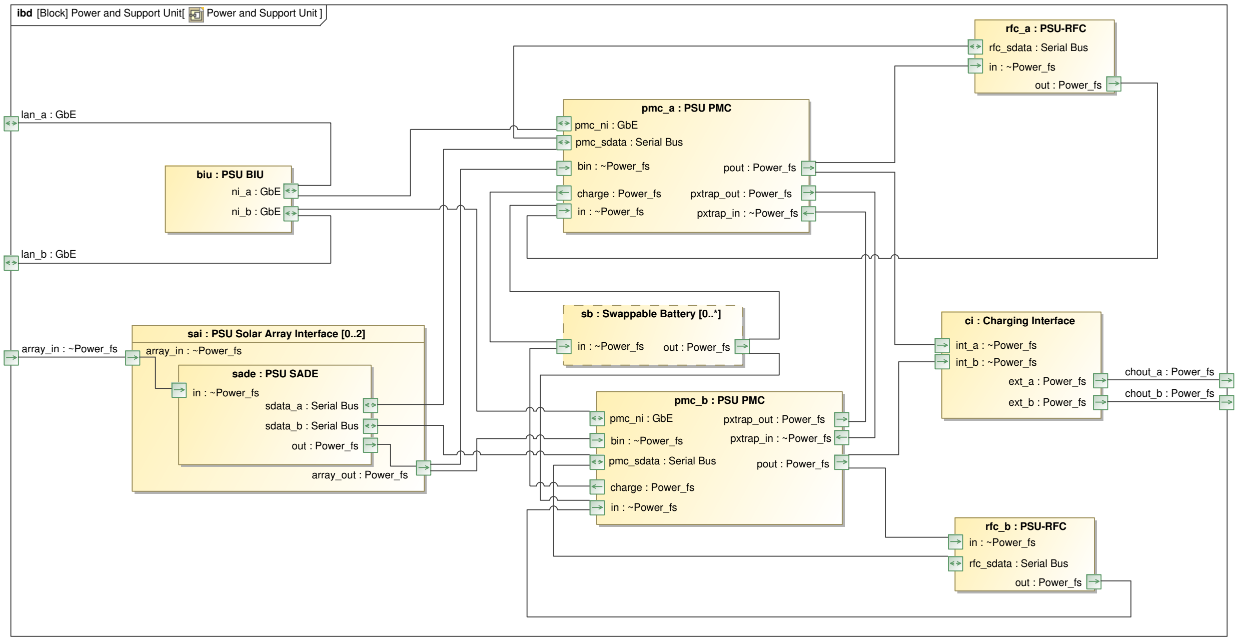

Figure 4-5. Connections in the Power and Support Unit.

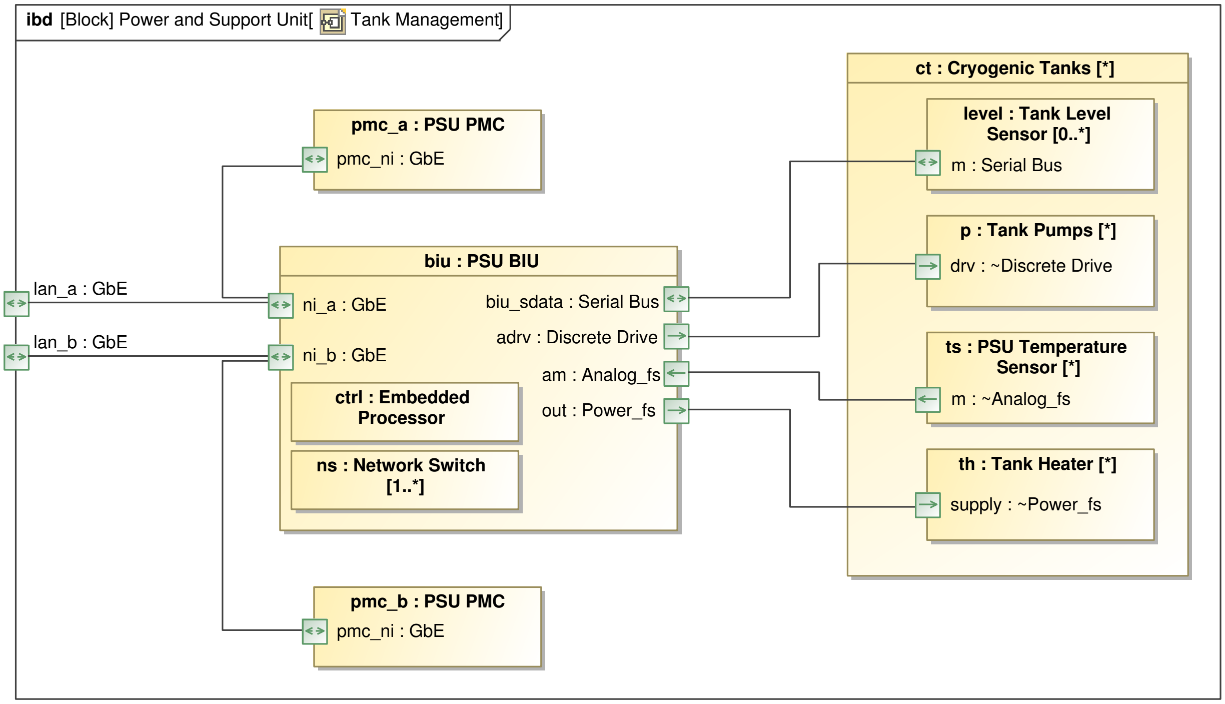

Figure 4-6. Interfaces for tank management in the Power and Support Unit.

4.1.2 Portable Utility Pallet

The PUP is a support system responsible for power generation and storage. It is also intended to enable various lunar surface excursions away from the outpost by supplementing operations and logistics. The PUP is designed to be carried by a mobility system and delivered to a remote location on the lunar surface without crew intervention. The PUP is designed to store and supply O2, water, and power to pressurized rovers, habitats, and other water storage tanks. Figure 4-7 illustrates a current concept for the PUP.

Figure 4-8 is a BDD illustrating the basic configuration of the PUP including solar arrays with control electronics, several discrete interfaces, a redundant power management complex, and removable batteries. The PUP also includes interfaces for communications using an ECS, charging, and element interface for integration as a payload. As described above, the PUP is intended to store O2 and water using cryogenic tanks. The data network for the configuration shown in Figure 4-8 is depicted in the IBD in Figure 4-9. The primary data interfaces are the element interface and the ECS interface and data is transferred to the PUP BIU. The PUP BIU is responsible for the operations of the PUP and PUP subsystems. Commands received by the PUP are processed and relayed as necessary to the SADE and the PUP power management complex. The IBD in Figure 4-10 depicts the complete data and power distribution in the PUP including control interfaces between the PUP BIU and the cryogenic tanks.

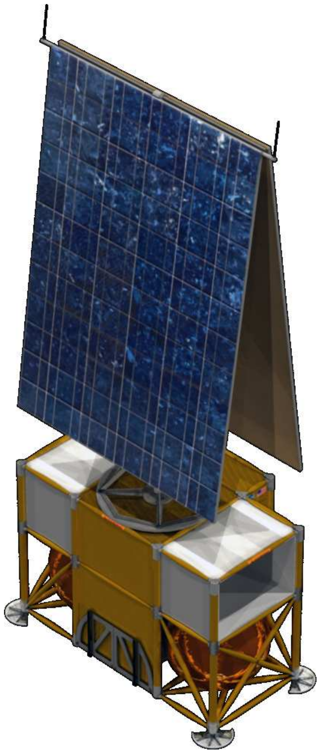

Figure 4-7. Portable Utility Pallet.

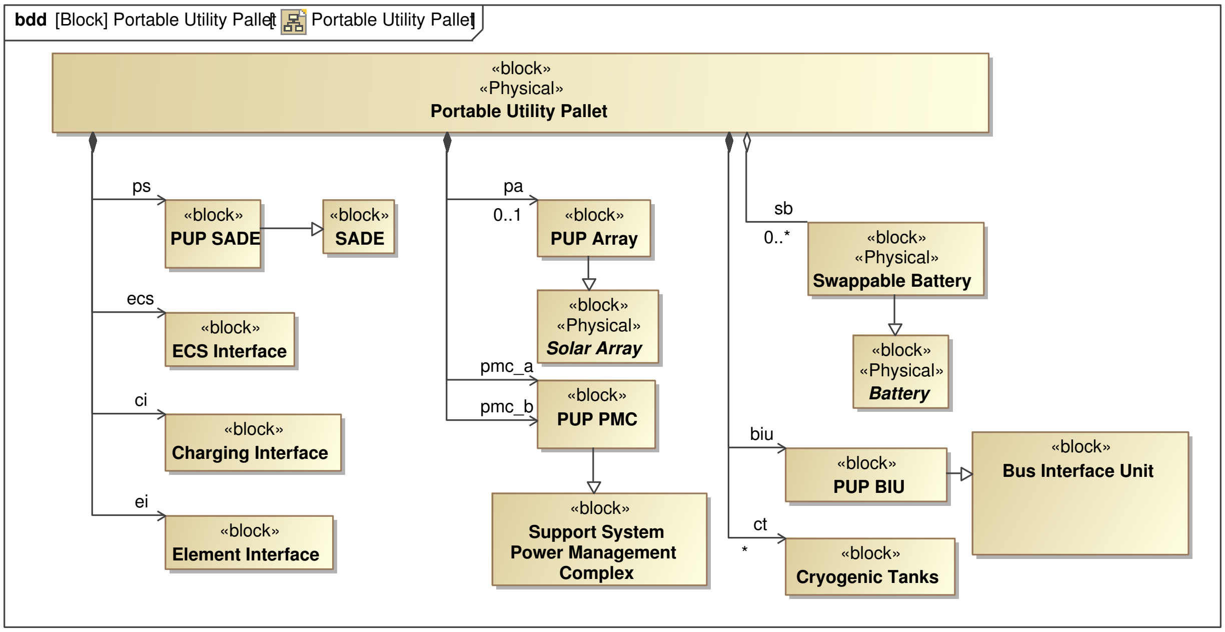

Figure 4-8. Configuration of the Portable Utility Pallet.

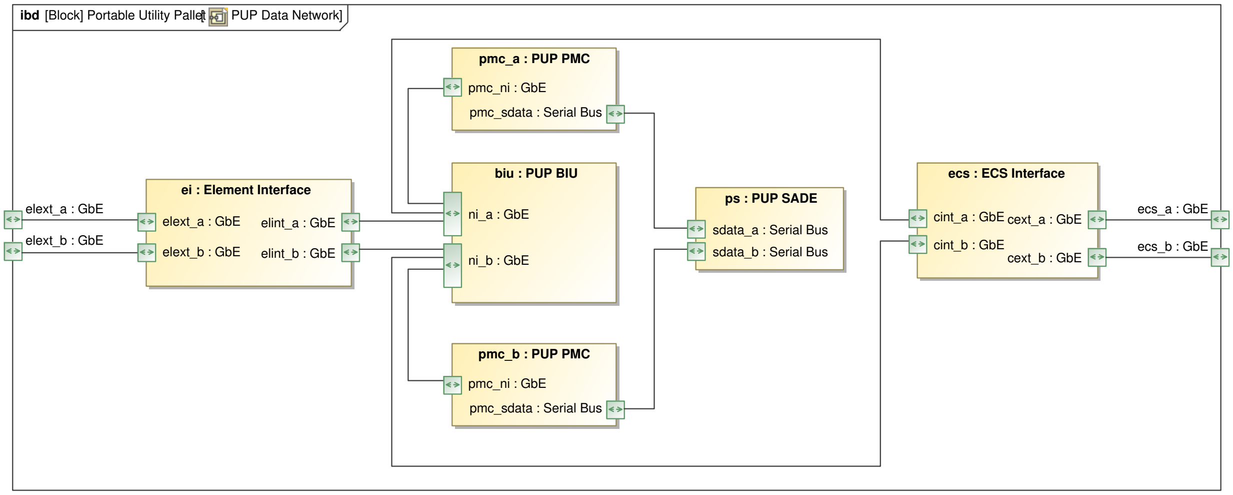

Figure 4-9. IBD for the Portable Utility Pallet data network.

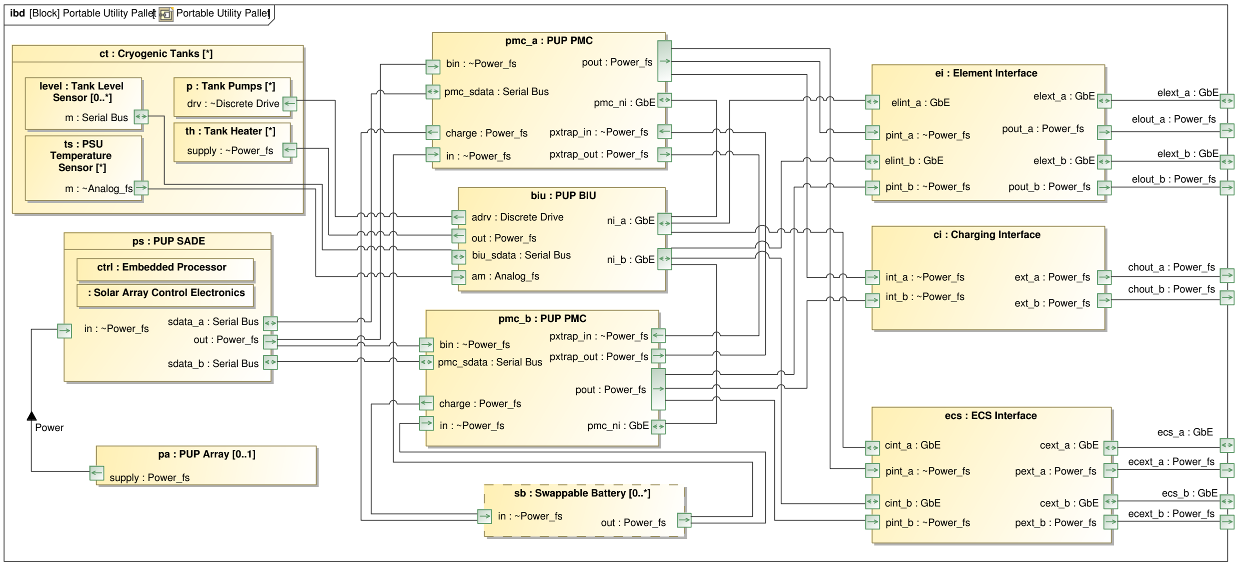

Figure 4-10. IBD for the Portable Utility Pallet.

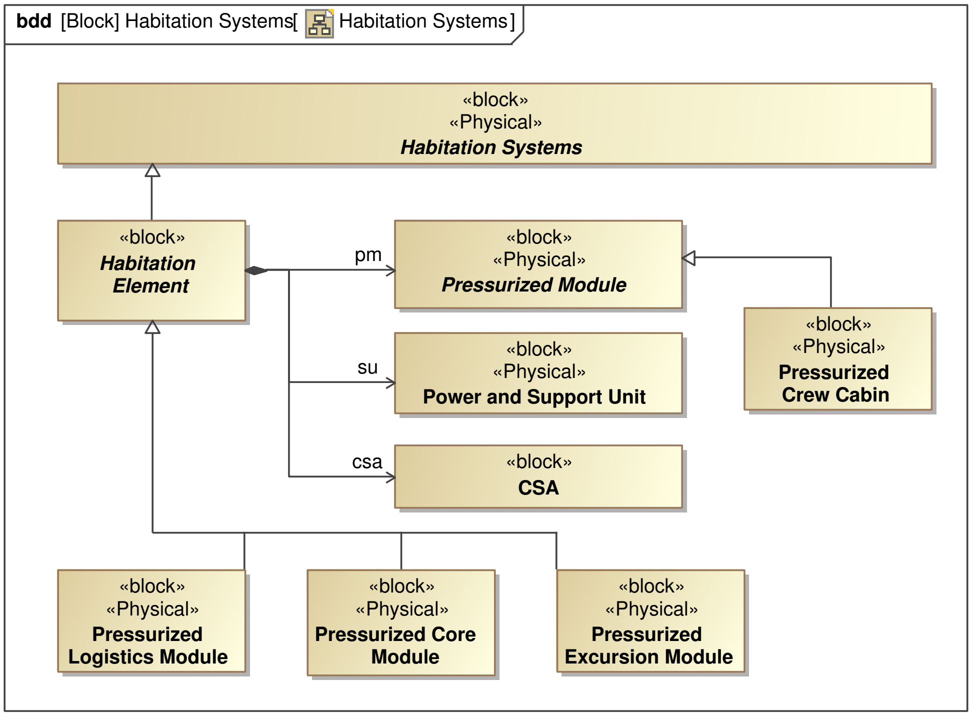

Figure 4-11. Hierarchy of habitation systems.

4.2 Habitation Systems

The habitation systems comprise a family of surface elements that provide a pressurized environment for the crew members to live and work in while performing mission tasks on the lunar surface. There are a number of architecture-level requirements that must be met by the habitat elements: reduce risk, reduce cost, achieve a basic level of crewed lunar surface stays as early as possible, and support outpost operations with a core habitat element while meeting the initial habitation functionality and volume goals. Outpost operations consist of Crew Operations, EVA Operations, Mission Operations, Science Operations, and Logistics & Supportability Operations [Kennedy et al., 2010]. Figure 4-11 is a BDD depicting the basic composition of a general habitation element as well as several specialized habitation elements. The pressurized logistics module, the PCM, and the Portable Excursion Module (PEM) are each types of habitation elements. The logistics module is used to store food, equipment, spare components, and many other resources required for long duration exploration of the lunar surface. The PCM provides primary living and working quarters for the crew and is further described in the following section. The PCM is outfitted with workstations for crew mission planning, communications, and remote operation of surface assets. The PCM also provides for non-mission activities including exercise, personal activities, meals, and rest. The excursion module is outfitted for science and exploration at remote locations away from an established outpost. As such, it is minimally configured for mission operation with workstations and equipment dedicated to science tasks. Primary mission planning and operations are managed in the LER the crew travels in. See the paper by Kennedy et al. [2010] for a more detailed description of the habitation elements.

Each habitation element is comprised of a pressurized module, a PSU, and a CSA. The interfaces and connections between these three primary components are shown in Figure 4-12. Power is

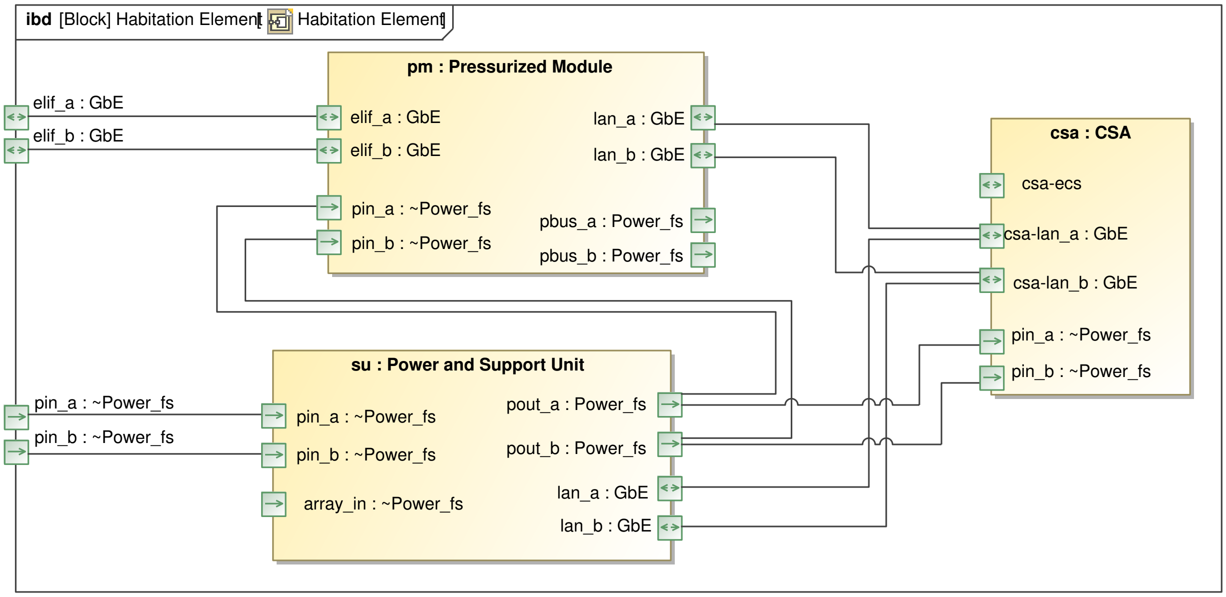

Figure 4-12. Primary modules of the generic habitation element.

received by the PSU, either from a solar array or some other power element, and distributed to the pressurized module and the CSA. Communications are received via an ECS (not shown) connected to the CSA or through the GbE interface on the pressurized module. The power outputs from the pressurized module are connected to subsystems integrated with the pressurized module and can be either internal or external to the pressure vessel. The following sections describe the reference design for the pressurized module and PCM, respectively.

4.2.1 Pressurized Module

The Pressurized Module (PM) is the basic component of a habitation element providing the pressurized environment for the crew to perform mission task objectives. As noted above, pressurized modules are used to implement several classes of surface elements including the logistics, core, and excursion modules. The basic PM is outfitted with redundant power subsystems, primary data systems complex, thermal control, and multiple hatches for ingress and egress. The power subsystem is modeled after the power management complex described in Section 4.1. The power subsystem in the PM provides only minimal power storage for emergency operations. The hatches provide interfaces for crew ingress and egress as well as to mate with other surface elements. For example, two PM can be mated to increase the habitable volume and an LER can mate with a PM allowing crew transfer without having to don EVA suits. The data systems complex in the PM provides for core data functions common to each of the pressurized elements. This includes crew utility panels, network switches, a backup computer, video processing, and cameras. The backup computer provides safety monitoring and can manage PM subsystems during quiescent and off-nominal operations. Other subsystems that would be included in a habitation element include ECLS and crew systems and have a greater dependence on the configuration and purpose of the habitation element. Figure 4-13 shows the composition of the pressurized module.

The primary data interface between the subsystems in the pressurized module is the TT-GbE network. The internal connectivity of the pressurized module is shown in Figure 4-14. The external

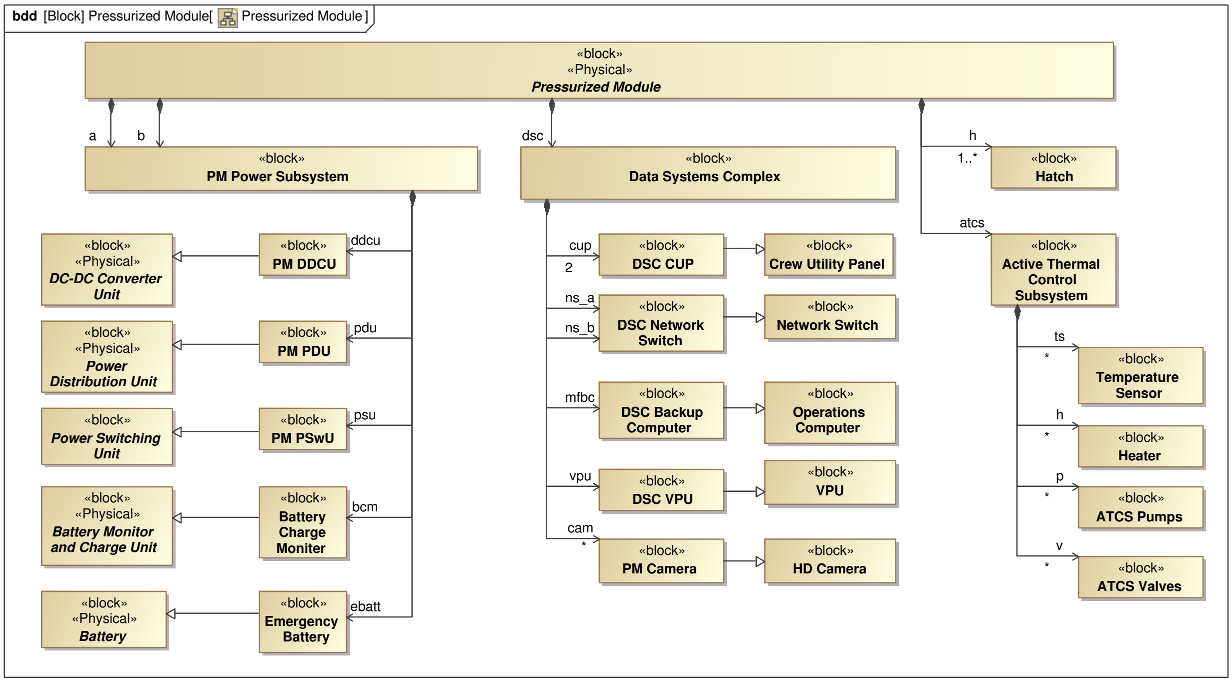

Figure 4-13. Composition of the basic pressurized module.

interfaces are through the hatches and the two network interfaces that route to the network switches in the data systems complex. The latter two network interfaces (lan_a and lan_b at the left in Figure 4-14) connect to the CSA as shown in Figure 4-12. Command and control is routed through the data systems complex which performs monitoring and distribution functions for commanding and health and status of PM subsystems. The data systems complex also includes video processing assets for internal cameras; these are not shown.

Figure 4-15 depicts a simplified IBD of the data systems complex. The switches route data between systems internal and external to the data systems complex. Command and critical telemetry is directed to the backup computer (mfbc) for systems monitoring and control. The crew utility panels provide crew interfaces for local communications (audio), caution and warning, systems status, and network access for wireless systems. The VPU encodes motion imagery and sends it over the local network to crew workstations, data storage, or communications systems for delivery to earth-based ground systems. Out-bound data traffic is sent to the C3I routers in the CSA to be routed to destinations outside the data systems complex.

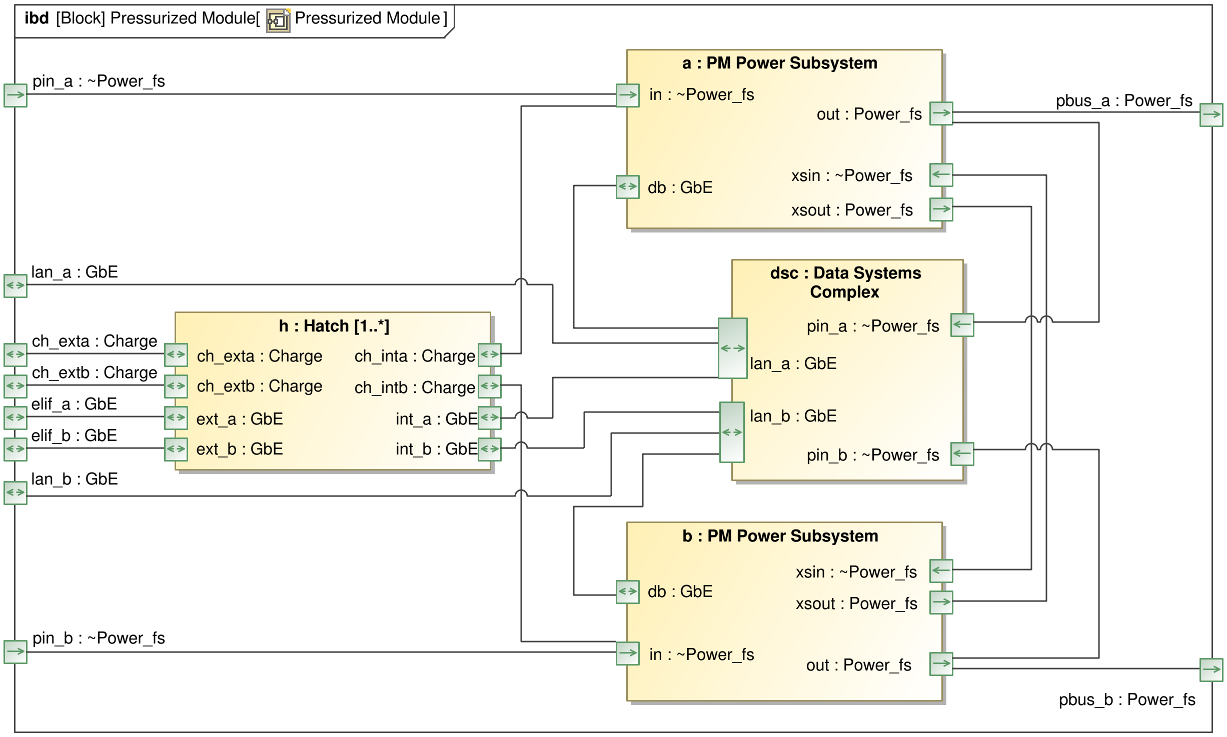

Figure 4-14. IBD for a generic pressurized module.

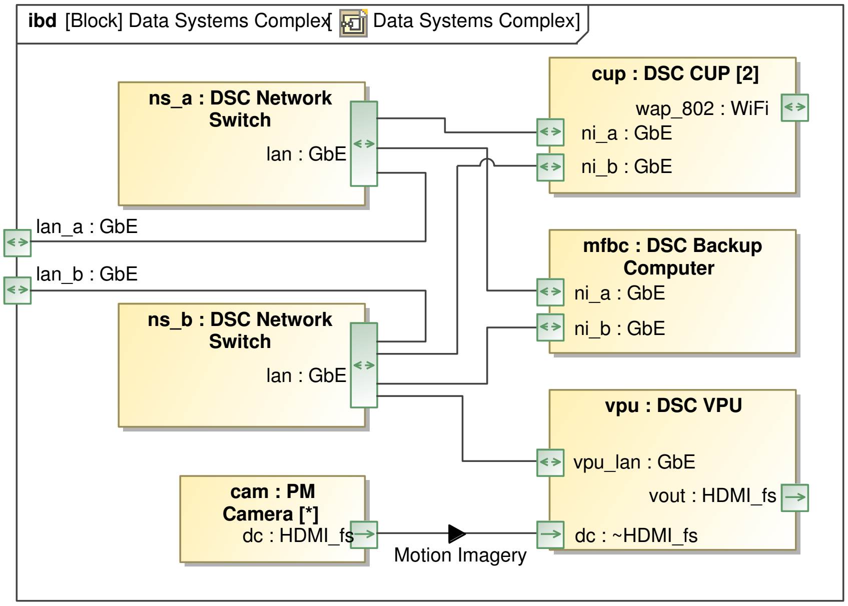

Figure 4-15. Habitation element C&DH architecture.

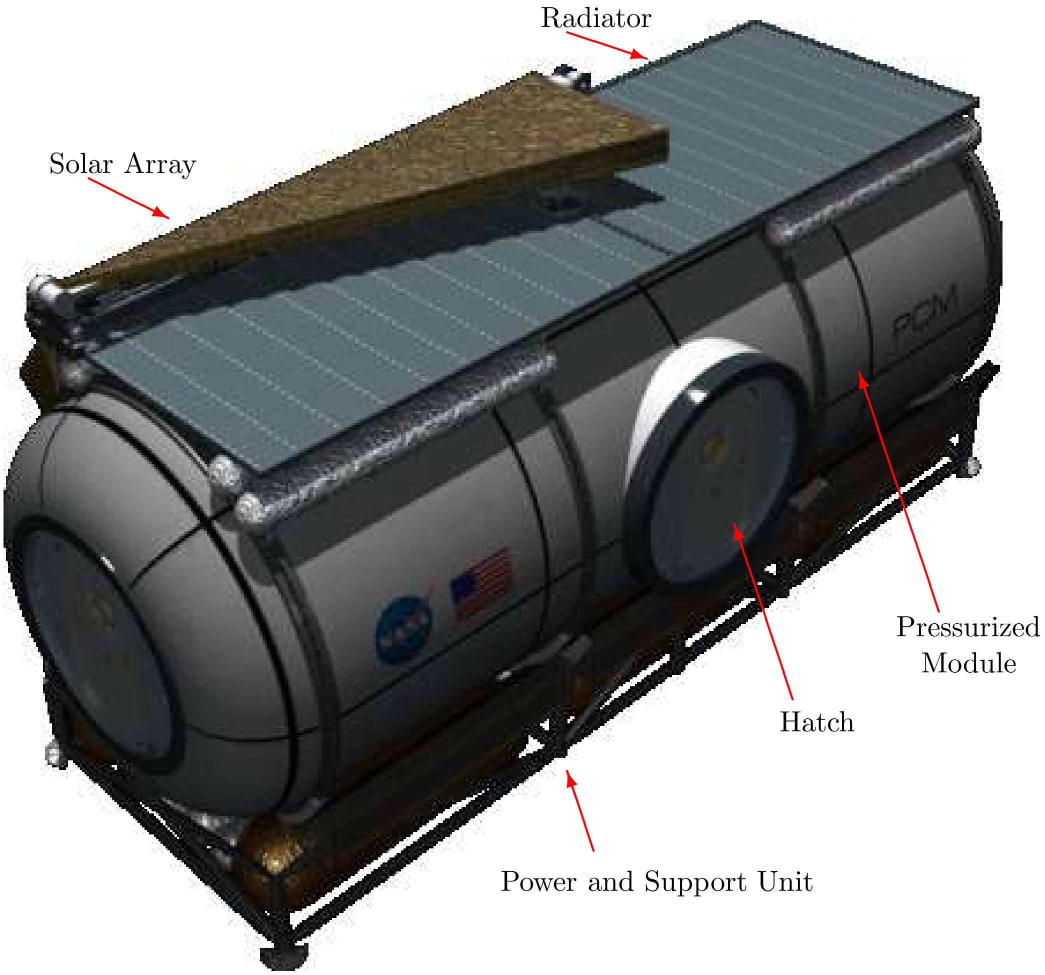

Figure 4-16. Image of the PCM habitation element.

4.2.2 Pressurized Core Module

Once delivered to the lunar surface, the PCM is the primary living and working quarters for the crew during exploration missions to the lunar outpost. An image of the current concept for the PCM is shown in Figure 4-16. In the image, a solar array is stowed top of the PCM during transport to, and across, the lunar surface. Beneath the stowed solar array, there is a radiator to reject heat from the PCM and its active subsystems. The radiator is attached to the top of the pressure module. Two hatches are shown in the figure which allow for mating to other elements and crew ingress and egress in emergency conditions. The pressure module is integrated with the PSU beneath it.

The PCM includes all of the functions and subsystems of the habitation element, including those of the pressurized module, plus the additional subsystems shown in Figure 4-17. The additional subsystems account for environmental control, communications, crew workstations, and at least one network server for file storage and network applications. The internal network connectivity between subsystems is shown in Figure 4-18. At the left, element interfaces are attached through the hatches to the network switches in the data systems complex. Data and command are then routed to a local subsystem or more likely to the C3I router in the CSA. The CSA is integrated into the

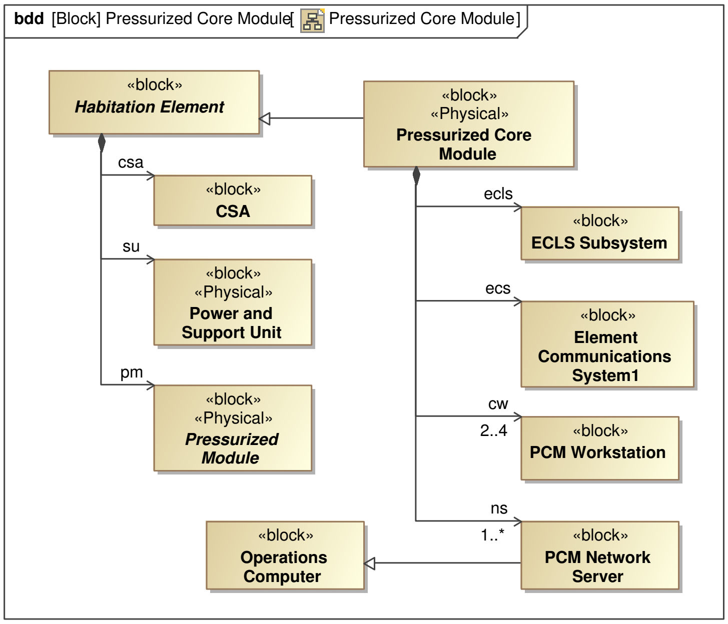

Figure 4-17. Pressurized core module configuration.

PSU external to the PCM. The CSA provides the primary computing and communications for the integrated PCM. The ECS provides RF communications over the SWN. All ECS communications are routed through the CSA. The final subsystem shown in Figure 4-18 is the PCM workstation for mission operations . The workstation is connected to the GbE network in the PM and displays video received from over the network or directly from the VPU.

ECLS is briefly described here to illustrate the integration of more complex subsystems. ECLS is responsible for environmental control and life support functions including atmospheric pressure control, air revitalization (cleaning and monitoring), water management, fire suppression, and waste management. The diagram in Figure 4-19 shows the ECLS subsystems with interfaces to the network and computing systems in the pressurized module and the CSA. The ECLS subsystem is connected to the data systems and CSA via the local GbE network. The BIUs in the ECLS subsystem perform command validation and local control. Although shown separately, later optimization may integrate these functions with BIUs that support other functions as well. Primary control applications for element-level ECLS operations are executed in the computing systems in the CSA. Low-level control applications are executed in the local BIUs managing engineering unit conversion and tight control loops where necessary. For example, the control application in the CSA does not require calibration data for the analog signal conditioning or data acquisition performed by the pressure control system. Figure 4-20 provides further detail for the water management subsystem of ECLS. The WM BIU receives and generates commands that are translated to discrete commands sent over a local serial bus or as power to a water management component. In this case, the serial bus is routed to the level sensor in the water storage tank, waste water recovery system, and the waste water recovery storage tank.

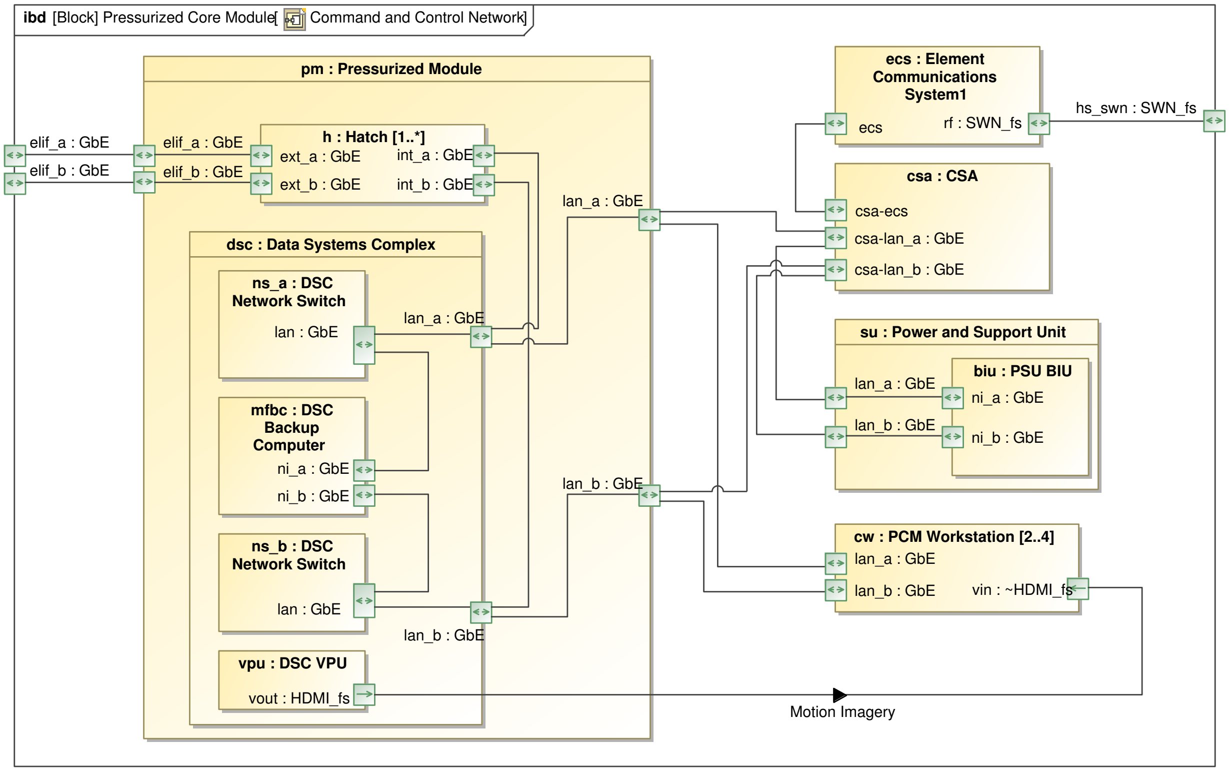

Figure 4-18. Command and control systems in the pressurized core module.

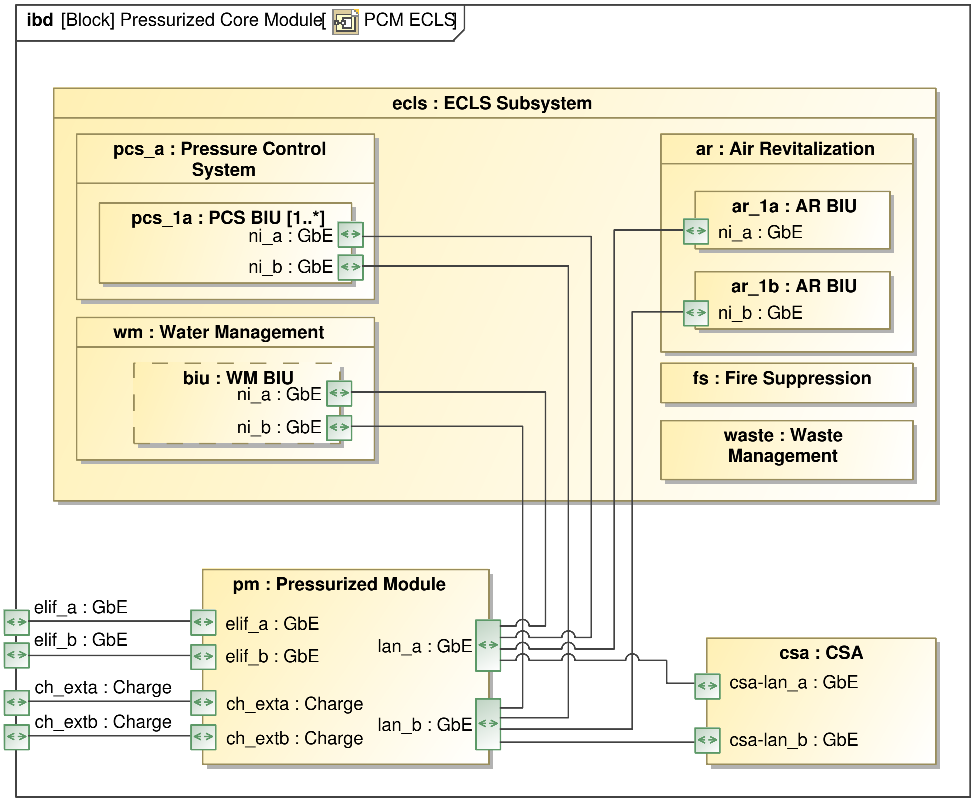

Figure 4-19. Pressurized Core Module (PCM) ECLS subsystems.

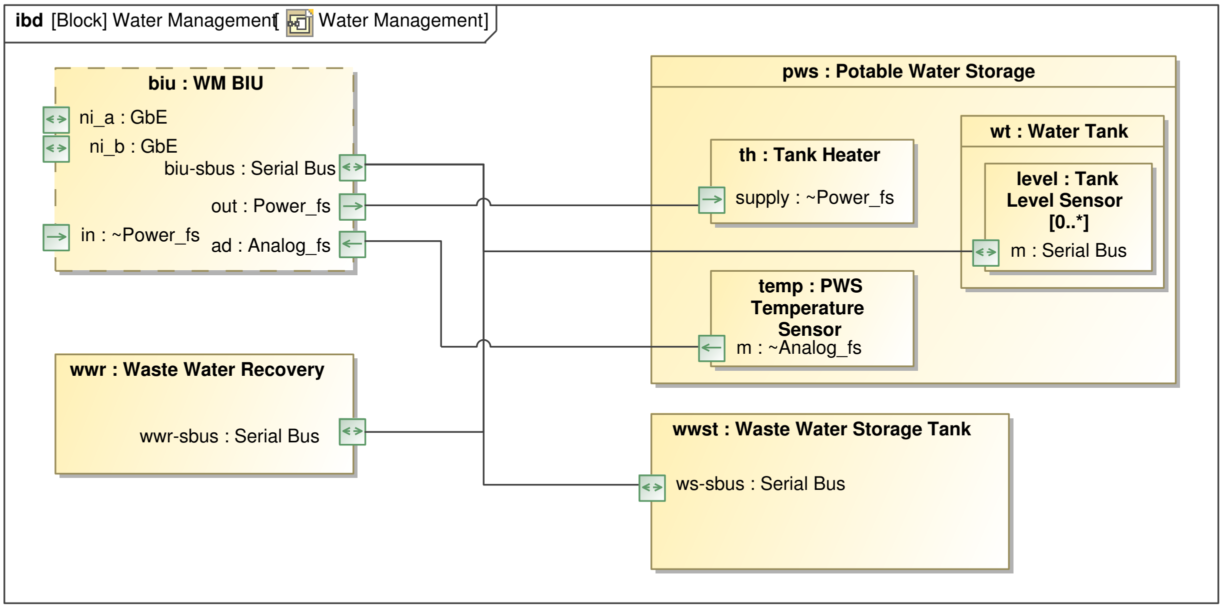

Figure 4-20. Water management subsystem.

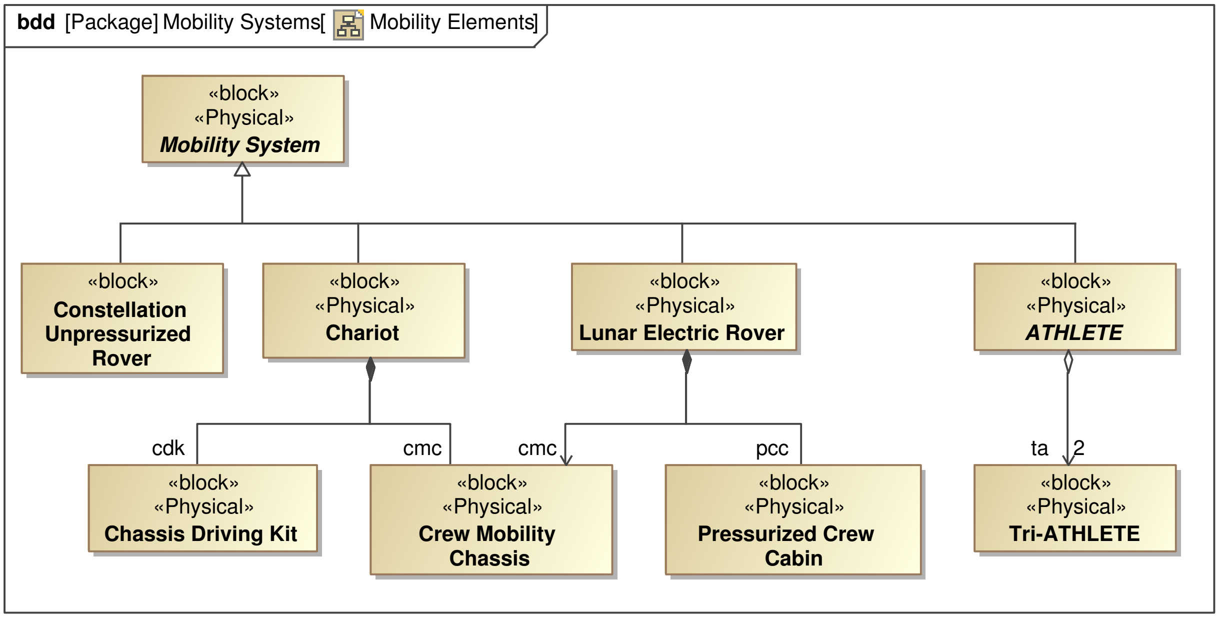

Figure 4-21. Primary mobility systems in the Lunar Surface Systems architecture.

4.3 Mobility Systems

The mobility systems (Chariot, LER, and ATHLETE for example) utilize the same architectural concepts and patterns as those described above. For discussion purposes, the following sections focus on systems derived from the CMC, but similar arguments are readily made with the ATHLETE. In either case, the mobility systems provide mobility to payloads and receive some level of command and control from external sources. These commands are processed by computing resources in a CSA where local control commands are generated and distributed over the CMC Onboard Data Network (ODN). This architecture utilizes the same primary components described in the previous sections in a smaller scale network. The primary mobility systems are shown in Figure 4-21. The Constellation Unpressurized Rover (CUR) is a small rover similar to those used during the Apollo missions. The Chariot is a configuration of the CMC with the Chassis Driving Kit (CDK). The Chariot also provides unpressurized mobility for the EVA crew. The Lunar Electric Rover (LER) includes a Pressurized Crew Cabin (PCC) integrated onto a CMC for a mobile, pressurized environment for the crew during lunar surface exploration. The ATHLETE is a heavy-lift vehicle enabling the transportation of large assets such as habitation systems, communication terminals, and possibly even landers. The ATHLETE is comprised of two three-legged tri-ATHLETEs that mate with a PSU to carry cargo. The following sections illustrate the application of the architectural concepts as applied to the CMC, Chariot, and LER.

4.3.1 Crew Mobility Chassis

The CMC is a six-wheeled vehicle able to function as an unpressurized or pressurized rover or without any crew at all [Harrison et al., 2008; Bluethmann et al., 2010]. Various payloads and tools can be attached for conducting science, exploration, or outpost construction tasks. In its basic configuration, the CMC can be remotely operated to carry unpressurized cargo. A CDK can be added to allow a suited EVA crew member to ride on and command the rover while carrying unpressurized cargo. A PCC can be attached to the CMC to form an LER providing a pressurized environment for “shirtsleeve” crew members to ride inside, donning suits to egress the habitable volume and conduct EVAs. The CMC can automatically dock with a charging station (such as the PUP) allowing it to recharge its internal batteries. It can also carry fuel cells or additional batteries for extended range excursions. The CMC is comprised of several subsystems to support vehicle operations and autonomy. The primary computing is provided by a CSA integrated into the CMC chassis. A redundant power management complex provides power storage, conditioning, and distribution to CMC subsystems and payloads. Several driving cameras provide for autonomous relative navigation, surface video and inspection, and video for an operator. An ECS provides communication over wireless network to local lunar surface assets via the surface wireless network, via communications resources in lunar orbit, or direct-with-earth communications. The CMC has six wheel assemblies with active suspension and wheel motor controllers which function similar to BIUs. The six wheel controllers are responsible for controlling the velocity of the vehicle as well as control of the active suspension. This local control implements a high-rate control loop (inner loop) which is managed by the vehicle lower rate control loop (outer loop). Navigation is supported by star trackers and Inertial Measurement Units (IMUs) as well as RF transponders and cameras for terrain relative navigation. Removable batteries can be interchanged with a PUP for charging and additional power converters provide high-voltage power to the wheel assemblies. For use in crew-critical operations, the controller (an operations computer) in the power management complex can serve as a backup computer if required. However, it is expected that the CMC is usually configured with a payload such as the CDK or PCC, both of which have onboard processing sufficient for backup computer and safety monitoring functions. The aggregation of the components used to realize a CMC is shown in the BDD Figure 4-22.

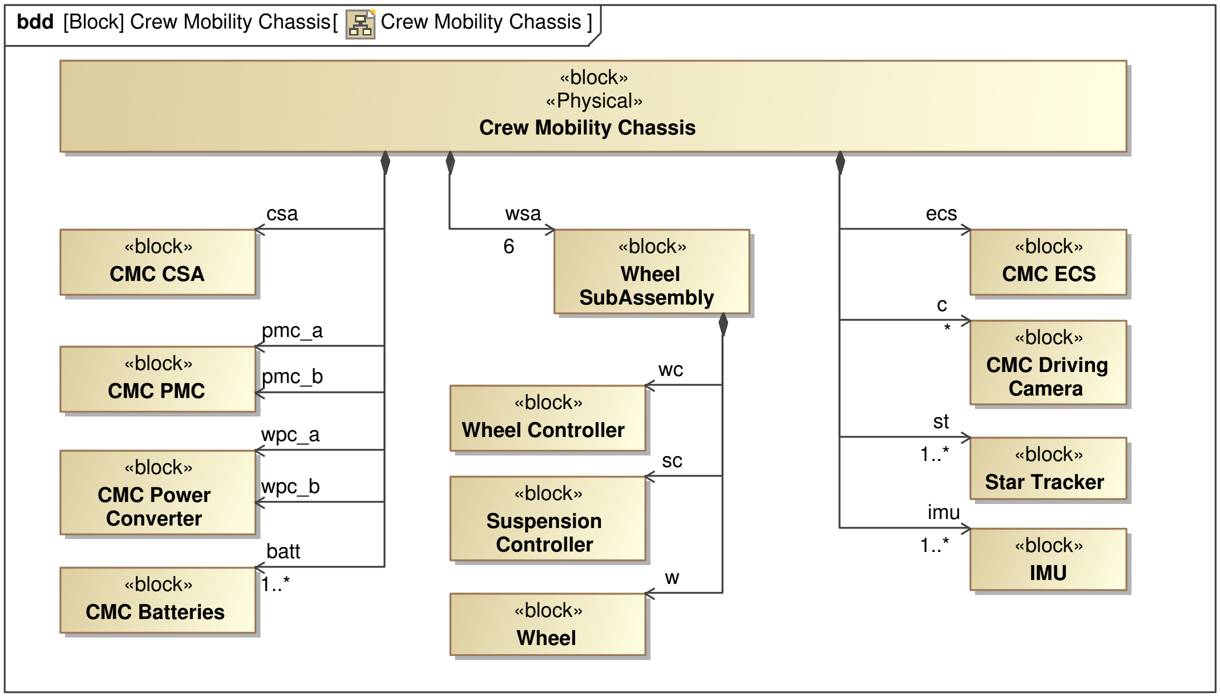

Figure 4-22. Composition of the crew mobility chassis.

The internal connectivity of the primary components of the CMC is shown in Figure 4-23. Two power interfaces are connected to the power management complexes for charging the CMC power storage. GbE interfaces are routed to the CSA for C&DH and command and control processing. Access to available wireless networks is provided by the ECS which routes communications to the C3I router in the CSA. The star tracker and IMU are shown connected to the serial bus from the BIUs in the PMCs, but they could also be connected via the local GbE network. Video from the driving cameras is sent to the VPU in the CSA for processing and encoding. Finally, the power converters are powered by the PMCs and provide high-voltage to the wheel assemblies. Control of the CMC power converters is provided by switching the load in the PMC.

4.3.2 Chariot

The Chariot configuration is comprised of a CMC and a CDK that hosts the control interface (display, input panel, and hand controls) and suit interfaces [Harrison et al., 2008]. A recent concept for the Chariot is shown in Figure 4-24. The image depicts an unmanned Chariot able to transport four crew members, two at control turrets and two on a rear platform. The Chariot can also carry payload (such as the PUP or science samples) on the rear platform. Figure 4-25 shows an IBD for the Chariot configuration illustrating data interfaces. The CDK augments the functions of the CMC by providing command and control interfaces for onboard crew with the CDK control station. The CDK control station is a class of crew workstation (see Section 3.4.1) configured specifically for operation by suited crew in a vacuum environment. Commands and data are transferred over the GbE interfaces and video can be displayed for the crew either real-time from the VPU over the HDMI interface or encoded over the GbE network.

4.3.3 Lunar Electric Rover

The LER provides a mobile, pressurized environment for the crew to perform exploration, mission, and science tasks. The LER is comprised of a PCC integrated onto a CMC (refer to Figure 4-21). The interfaces of the LER are shown in Figure 4-26. As described above, the CMC has a SWN interface and charging interfaces. Between the CMC and PCC, power is supplied to the PCC and there is a bi-directional GbE interface for data, command, and control transmission. Not shown are the interfaces for mating to other pressurized modules such as the PCM. The PCC has two hatches with the interfaces described for the pressurized module in Section 4.2.1.

The PCC is a pressurized module designed with a smaller volume for mobile applications. Some of main subsystems in the PCC are shown in the BDD in Figure 4-27. The PCC includes power distribution and a data systems complex as described in Section 4.2.1. The PCC has two workstations for controlling the LER (as well as other remote surface assets) and performing mission tasks. It also provides ECLS to sustain the crew. The thermal control subsystem (not shown) incorporates active and passive techniques for thermal management [Stephan, 2010]. The internal connections for these systems are shown in the IBD in Figure 4-28.

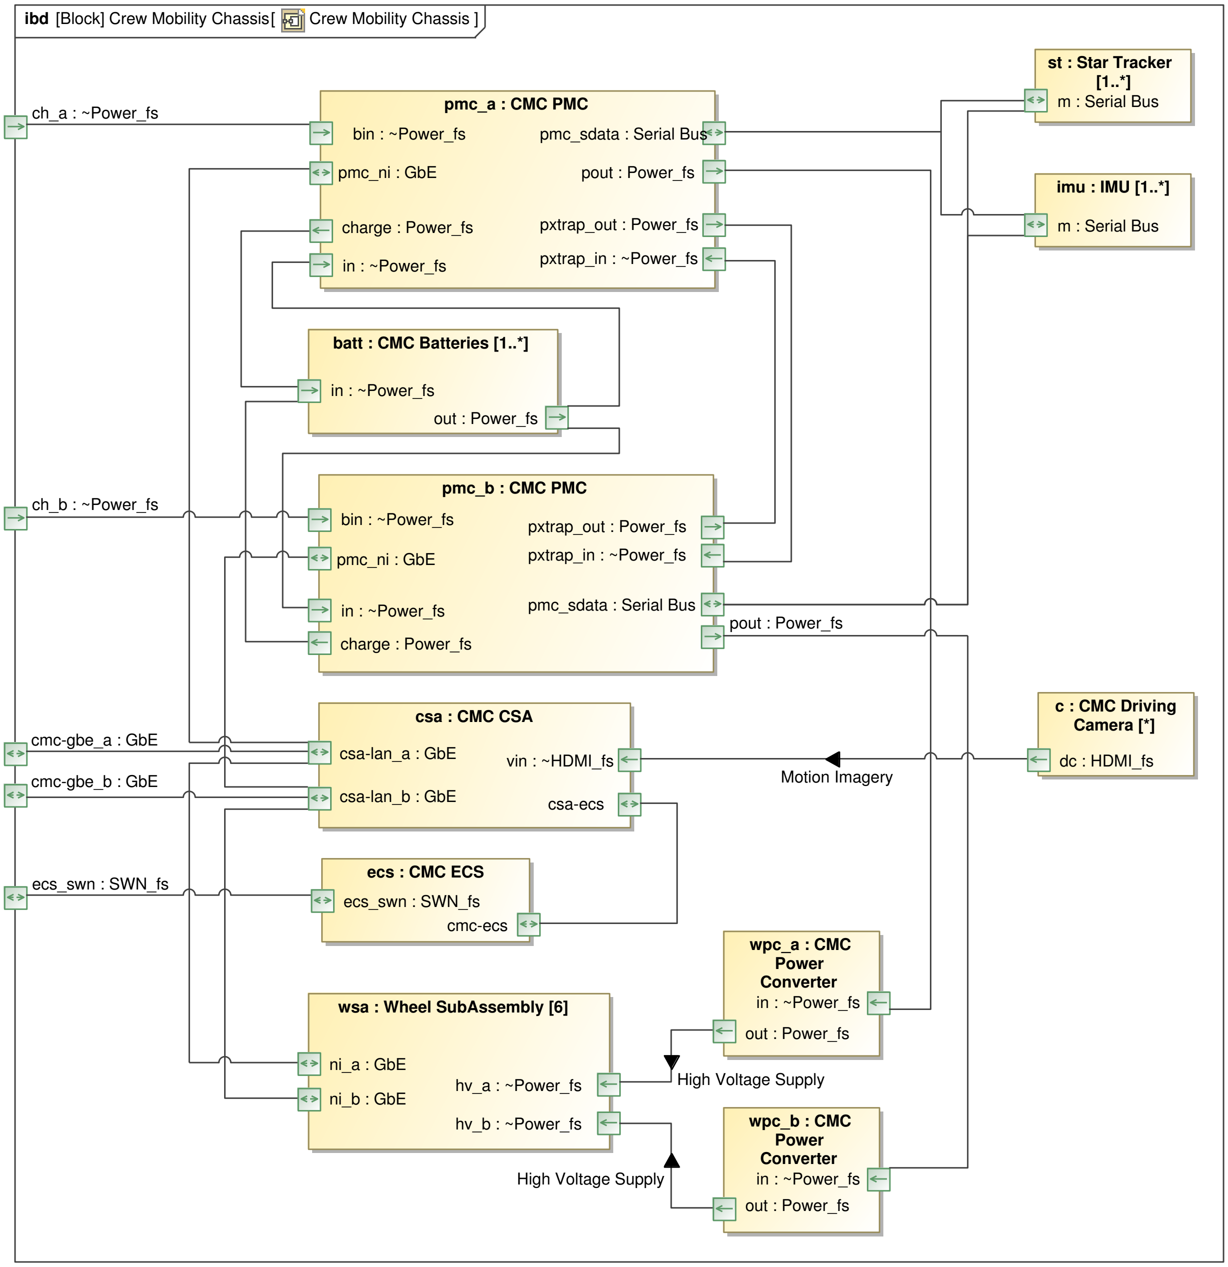

Figure 4-23. IBD for the CMC.

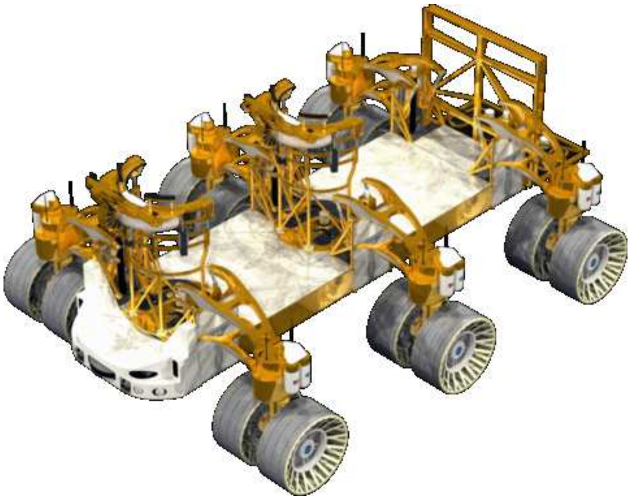

Figure 4-24. Chariot configuration with Crew Mobility Chassis (CMC) and Chassis Driving Kit (CDK).

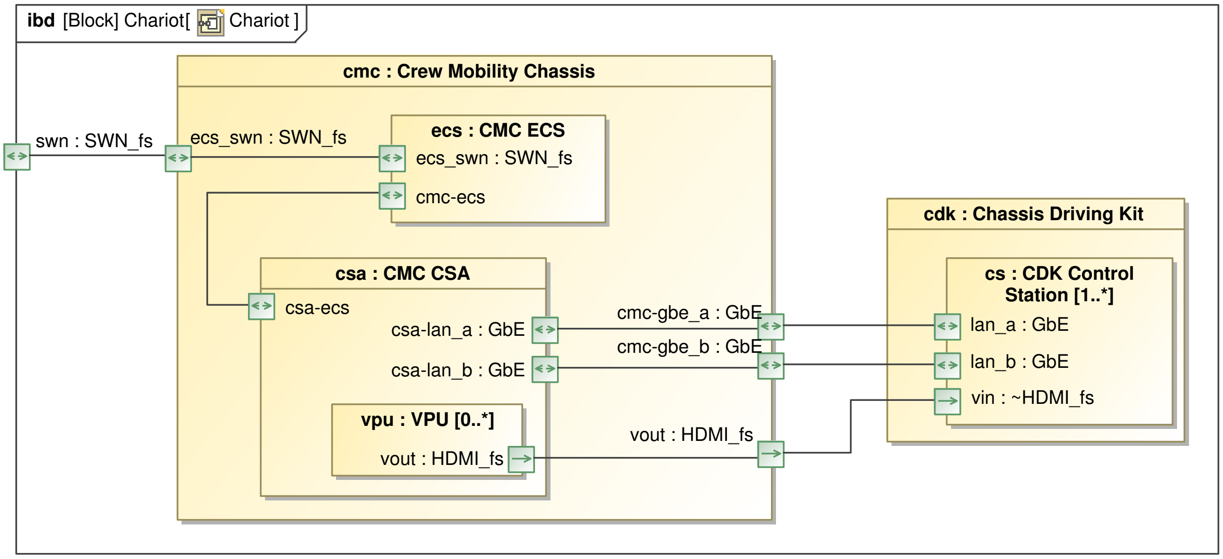

Figure 4-25. IBD for the Chariot configuration.

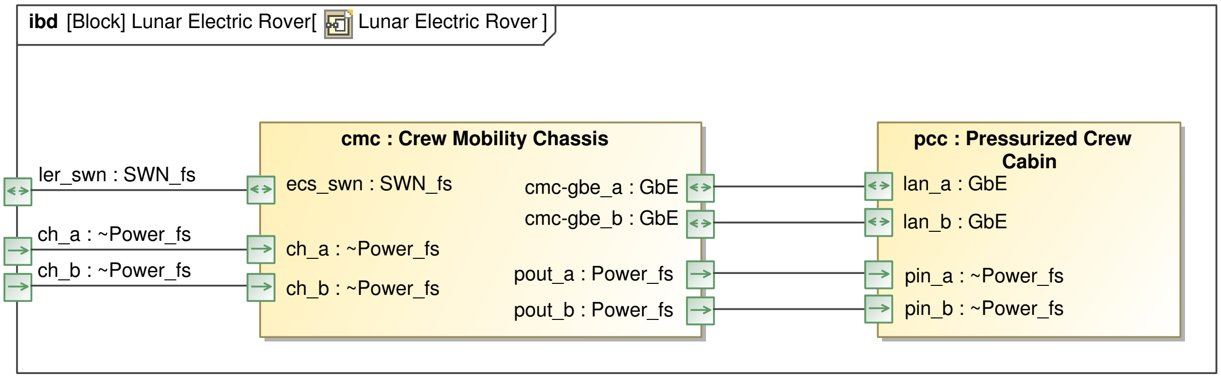

Figure 4-26. IBD for the Lunar Electric Rover (LER).

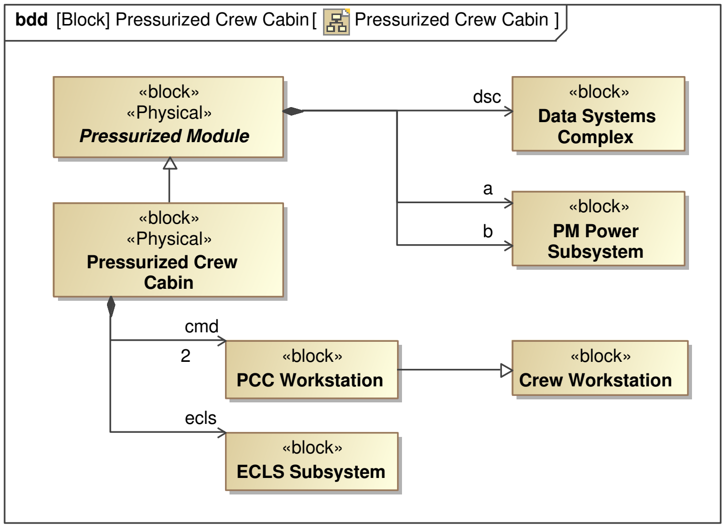

Figure 4-27. Composition of the Pressurized Crew Cabin (PCC).

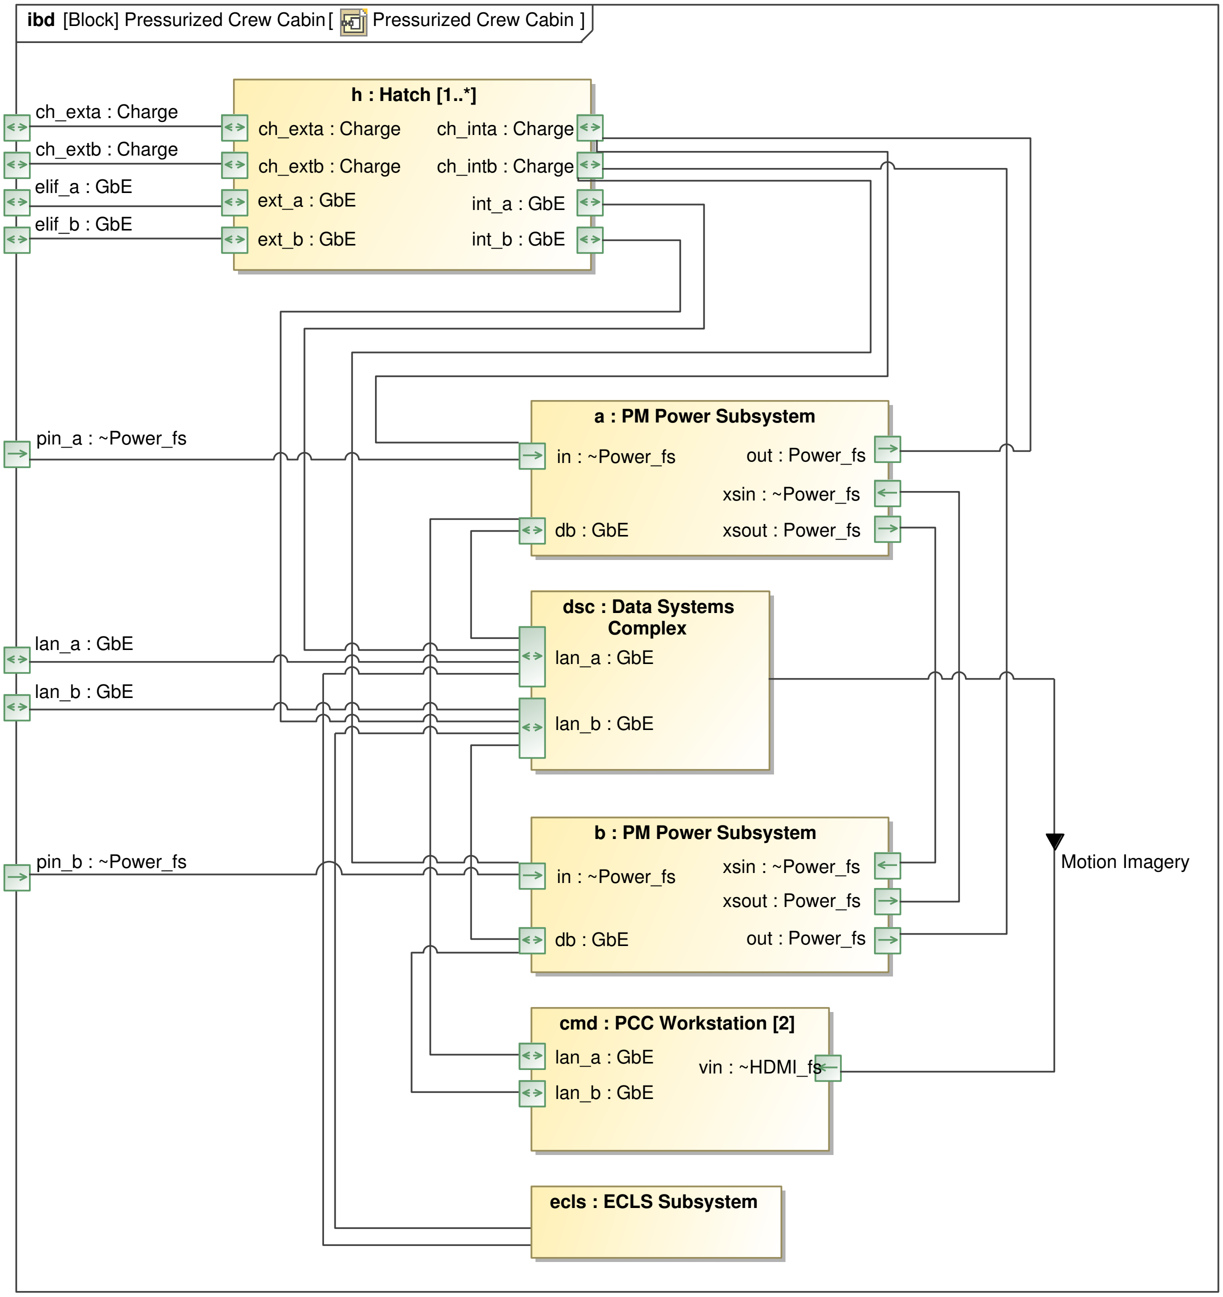

Figure 4-28. IBD for the Pressurized Crew Cabin (PCC).

5 Summary

This report describes the current concepts for the reference avionics architecture used in Lunar Surface Systems (LSS) elements. The progression of the discussion starts in Section 2 with a brief overview of the surface elements that rely on the avionics and data systems described in this report. Section 3 describes assemblies and functions that are collected to develop subsystems used in multiple surface elements. These establish patterns of reuse and integration that are then aggregated to develop reference architectures for complete elements and vehicles. The integrated reference designs for several major elements are described in Section 4. Among these are the Power and Support Unit (PSU) followed by descriptions of the Pressurized Core Module (PCM) as a representative habitation element and the Chariot configuration of the Crew Mobility Chassis (CMC). The description of the element reference designs serves to both detail the vehicle design as well as illustrate the use of subsystems repeated in architectural patterns. For example, the support system power management complex described in the introduction of Section 4.1 is reused in the PSU, Portable Utility Pallet (PUP), habitation elements, and mobility systems. Also, the general reference architecture incorporates both string-level redundancy as well as dissimilar redundancy as part of initial considerations for common cause failure mitigation.

5.1 Alternate Destinations

Recent governmental and programmatic activities have altered the primary objectives of the National Aeronautics and Space Administration (NASA) Exploration Systems Mission Directorate related to manned spaceflight. The concepts presented here reflect systems associated with lunar surface deployment. The concepts and assumptions are grossly applicable to many space destinations as the avionics architecture conceived of to date is intended to be flexible and evolvable. There are, however, a number of assumptions that would have to be revisited. Among these are assumptions about timing for critical operations and phases of operations, the interoperability of components and systems, and maintenance considerations. In-flight vehicles have dynamic intervals different from those experienced by mobile and stationary surface assets. The radiation environment, although not addressed in this report, has driven many of the decisions related to reliability, redundancy, and fault-management. Some vehicles and the quantity of various elements are the result of targeting exploration on the lunar surface. The exploration of many in-space destinations will likely not require wheeled vehicles nor will they merit the interoperation of a large collection of nearby assets. This potentially increases the need for autonomy in the various elements. Finally, a major objective in developing the described lunar reference avionics architecture was to assess and establish commonality across the various elements. The collection of relevant elements is dependent on the destination which in turn affects the various parameters and amount of commonality across the architecture. The overall architecture will probably still be applicable; however, subtle details of implementation and subsystem partitions will likely shift.

5.2 Future Work

The architectural concepts presented in this report are neither exhaustive nor complete. The intent is to both capture the concepts that have developed so far and establish a basis for future analysis. In the continued development of manned space exploration, there are a number of objectives that are immediately evident that can benefit and build upon this work. Current programmatic changes are evaluating destinations other than the lunar surface. The first task is to revise the architecture to account for new mission objectives and reference missions. A second task is to assess mission safety to develop reliability metrics to evaluate and mature architecture. This is similar to, and would extend, the work described by Borer et al. [2010]. The third task is to develop standards for interoperability including form-factor, power, communications, and software. These are required to efficiently realize the interoperability described in this report. The fourth immediate task is to mature the design relative to instrumentation and subsystem integration. There are significant opportunities for subsystem integration that will reduce system mass and power. Also, detailed instrumentation planning and designs are incomplete and will affect subsystem partitioning and integration.

References

P. Anderson. Orion Electrical Power System An Achievement in Space Power Technology. In *IECEC 2010, 8th Annual International Energy Conversion Engineering Conference*, Nashville, TN USA, 25-28 July 2010.

AS6802 Draft. Aerospace Standard Time-Triggered Ethernet. Aerospace Standard AS 6802, SAE Aerospace, 2010. Draft.

B. Bluethmann, E. Herrera, A. Hulse, J. Figuered, L. Junkin, M. Markee, and R. O. Ambrose. An active suspension system for lunar crew mobility. *Aerospace Conference, 2010 IEEE*, pages 1–9, 6-13 Mar. 2010.

N. Borer, I. Claypool, D. Clark, J. West, R. Odegard, K. Somervill, and N. Suzuki. Model-driven development of reliable avionics architectures for lunar surface systems. *Aerospace Conference, 2010 IEEE*, pages 1–21, 6-13 Mar. 2010.

R. W. Butler. A Primer on Architectural Level Fault Tolerance. Technical Memorandum TM-2008-215108, NASA Langley Research Center, Hampton, Virginia, Feb. 2008.

G. Cancro, P. Eisenreich, G. Oxton, S. Ling, and K. Balon. NASA SmallSat modular hardware and software standardization. *Aerospace conference, 2009 IEEE*, pages 1–19, 7-14 Mar. 2009.

F. Cristian and C. Fetzer. Fault-tolerant external clock synchronization. In *Distributed Computing Systems, 1995., Proceedings of the 15th International Conference on*, pages 70–77, 30 May -2 June 1995.

DDS v1.2. Data Distribution Service for Real-Time Systems Specification. Specification formal/2007-01-01, Object Managment Group, Jan. 2007.

FlexRay v2.1. FlexRay Communications System Protocol Specification. Standard Version 2.1, Revision A, FlexRay Consortium, 22 Dec. 2005.

H. Garrett, M. Johnson, J. Ratliff, A. Johnston, S. Anderson, and W. Stapor. Single event upset effects on the clementine solid state data recorder. *AIAA Journal of Spacecraft and Rockets, Special Issue on Clementine*, Oct. 1995.

D. Gwaltney and J. Briscoe. Comparison of Communication Architectures for Spacecraft Modular Avionics Systems. Technical Memorandum TM-2006-214431, NASA Marshall Space Flight Center, Huntsville, Alabama, June 2006.

R. Hammett. Flight-critical distributed systems - design considerations. In *Digital Avionics Systems Conference, 2002. Proceedings. The 21st*, volume 2, pages 13.B.3–1–13.B.3–8, 2002.

J. Hanaway and R. F. Moorehead. Space Shuttle Avionics System. Special Publication SP-504, National Aeronautics and Space Administration, 1989.

R. E. Harper and J. H. Lala. Fault tolerant parallel processor. *Journal of Guidance, Control, and Dynamics*, 14(3):554–563, May - June 1991.

R. E. Harper, J. H. Lala, and J. J. Deyst. Fault tolerant parallel processor architecture overview. *International Symposium on Fault-Tolerant Computing (FTCS '88)*, pages 252–259, June 1988.

D. Harrison, R. Ambrose, B. Bluethmann, and L. Junkin. Next Generation Rover for Lunar Exploration. *Aerospace Conference, 2008 IEEE*, pages 1–14, 1-8 Mar. 2008.

R. Hodson. SPACE-104: A stackable solution for space electronics. *PC/104 and Small Form Factors*, pages 32–35, 2006.

K. J. Kennedy, L. D. Toups, and M. Rudisill. Constellation Architecture Team-Lunar Scenario 12.0 Habitation Overview. Technical Report 20100003415, NASA Johnson Space Center, Houston, Texas, 2010.

H. Kopetz, A. Ademaj, and A. Hanzlikm. Integration of internal and external clock synchronization by the combination of clock-state and clock-rate correction in fault-tolerant distributed systems. In *Real-Time Systems Symposium, 2004. Proceedings. 25th IEEE International*, pages 415–425, 5-8 Dec. 2004.

MCAR. A Research Agenda for Mixed Criticality Systems. In *Cyber Physical Systems (CPS) Week 2009 Workshop on Mixed Criticality*. Air Force Research Laboratory, 16 Apr. 2009.

M. McCabe, C. Baggerman, and D. Verma. Avionics architecture interface considerations between constellation vehicles. In *Digital Avionics Systems Conference, 2009. DASC '09. IEEE/AIAA 28th*, pages 1.E.2–1–1.E.2–10, 23-29 Oct. 2009.

M. Paulitsch and W. Steiner. Fault-Tolerant Clock Synchronization for Embedded Distributed Multi-Cluster Systems. In *Real-Time Systems, 2003. Proceedings. 15th Euromicro Conference on*. IEEE Computer Society, 2003.

J. Rushby. A comparison of bus architectures for safety-critical embedded systems. Contractor Report CR-2003-212161, Computer Science Laboratory, SRI International, Menlo Park, CA, Mar. 2003.

H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson. RTP: A Transport Protocol for Real-Time Applications. Request for Comment RFC3550, The Internet Society, July 2003.

G. Seagrave. SpaceCube: A Reconfigurable Processing Platform for Space. *Military / Aerospace Programmable Logic Devices (MAPLD) 2008*, 15-18 Sept. 2008.

K. Somervill. Lunar Applications in Reconfigurable Computing. *Military / Aerospace Programmable Logic Devices (MAPLD) 2008*, 15-18 Sept. 2008.

R. A. Stephan. Overview of NASA's Thermal Control System Development for Exploration Project. *40th International Conference on Environmental Systems*, 11-15 July 2010.

SysML v1.1. OMG Systems Modeling Language Version 1.1. [online], Nov. 2008.

J. Townsend, J. Biesiadecki, and C. Collins. ATHLETE mobility performance with active terrain compliance. *Aerospace Conference, 2010 IEEE*, pages 1–7, 6-13 Mar. 2010.

C. Watkins and R. Walter. Transitioning from federated avionics architectures to integrated modular avionics. In *Digital Avionics Systems Conference, 2007. DASC '07. IEEE/AIAA 26th*, pages 2.A.1–1 – 2.A.1–10, Oct. 2007.

B. H. Wilcox. ATHLETE: An Option for Mobile Lunar Landers. *Aerospace Conference, 2008 IEEE*, pages 1–8, 1-8 Mar. 2008.

B. H. Wilcox. ATHLETE: A cargo and habitat transporter for the moon. *Aerospace conference, 2009 IEEE*, pages 1–7, 7-14 Mar. 2009.

REPORT DOCUMENTATION PAGE	*Form Approved* *OMB No. 0704–0188*

1. REPORT DATE *(DD-MM-YYYY)* 01-12-2010	**2. REPORT TYPE** Technical Memorandum	**3. DATES COVERED** *(From - To)* 3/2007 – 9/2010

4. TITLE AND SUBTITLE Reference Avionics Architecture for Lunar Surface Systems	**5a. CONTRACT NUMBER**
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S) Kevin M. Somervill, Jonathan C. Lapin, Oron L. Schmidt	**5d. PROJECT NUMBER**
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER 075585.01.01.01.01.04
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) NASA Langley Research Center Hampton, Virginia 23681-2199	**8. PERFORMING ORGANIZATION REPORT NUMBER** L–19953
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) National Aeronautics and Space Administration Washington, DC 20546-0001	**10. SPONSOR/MONITOR'S ACRONYM(S)** NASA
	11. SPONSOR/MONITOR'S REPORT NUMBER(S) NASA/TM–2010–216872

12. DISTRIBUTION/AVAILABILITY STATEMENT

Unclassified-Unlimited
Subject Category 62
Availability: NASA CASI (443) 757-5802

13. SUPPLEMENTARY NOTES

An electronic version can be found at http://ntrs.nasa.gov.

14. ABSTRACT

Developing and delivering infrastructure capable of supporting long-term manned operations to the lunar surface has been a primary objective of the NASA Constellation Program. Several concepts have been developed related to development and deployment of lunar exploration systems that provide critical functionality such as transportation, habitation, and communication, to name a few. Together, these systems perform complex safety-critical functions, largely dependent on avionics for control and behavior of system functions. These functions are implemented using interchangeable, modular avionics designed for lunar transit and lunar surface deployment. There are two core concepts in this reference avionics architecture. The first concept uses distributed, smart systems to manage complexity, simplify integration, and facilitate commonality. The second core concept is to employ extensive commonality between elements and subsystems. These two concepts are used in the context of developing reference designs for many lunar surface exploration vehicles and elements. These concepts are repeated as architectural patterns in a conceptual architectural framework. This report describes the use of these architectural patterns in a reference avionics architecture for Lunar surface exploration systems.

15. SUBJECT TERMS

avionics, architecture, elements, exploration, lunar

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			STI Help Desk (email: `help@sti.nasa.gov`)
U	U	U	UU	62	**19b. TELEPHONE NUMBER** ***(Include area code)*** (443) 757-5802