

Fast Coordination of Power Electronic Converters for Energy Routing in Shipboard Power Systems

H L Ginn III**, J D Bakos*, A Benigni*

* *University of South Carolina, Columbia, SC, USA*

* Corresponding Author. Email: ginnhl@cec.sc.edu

Synopsis

A long-term goal of future naval shipboard power systems is the ability to manage energy flow with sufficient flexibility to accommodate future platform requirements such as, better survivability, continuity, and support of pulsed and other demanding loads. To attain this vision of shipboard energy management, shipboard power and energy management systems must coordinate operation of all major components in real-time. The primary components of a shipboard power system are the generators, energy storage modules, and increasingly power electronics that interface those sources and main load centers to the system. Flexible management of energy flow throughout shipboard distribution systems can be realized by automated coordination of multiple power electronic converters along with storage and generation systems.

Use of power converters in power distribution systems has continuously increased due to continued development of the power electronics building blocks (PEBB) concept which reduces cost and increasing reliability of converters. Recent developments in SiC power devices are yielding PEBBs with far greater switching frequencies than Si based devices resulting in an order of magnitude reduction of the time scales as compared to converter systems utilizing conventional IGBT based PEBBs. In addition there have also been advancements in highly modularized converter systems with hundreds of PEBBs such as the Modular Multilevel Converter. Both of those trends have resulted in the continued evolution of the Universal Controller Architecture which attempts to standardize control interfaces for modular power electronic systems. Further development of interface definitions and increasing communication and computational capabilities of new FPGA based controllers provides opportunities beyond simply supporting SiC PEBBs. Fast control coordination across the system using an appropriate communication architecture provides a degree of energy management not previously realizable in shipboard power systems. The paper will present recent research results in networked control architectures for power electronic converter coordination and control. It will demonstrate that current FPGA and gigabit speed serial communication technologies allow for a very high degree of energy flow control.

Keywords: Coordinating Control; Power Electronics; Distributed Control Architecture

1. Introduction

Present trends indicate that shipboard energy management systems will contain an increasing number of power electronic devices. A shipboard power system may have numerous multi-functional power electronic converters connecting sources, loads, and energy storage to the bus. Systems where converters are the interface between many of the main sources of energy and load centers have the ability to direct the flow of energy if the control of the converters is appropriately coordinated. This allows for reduction of systems losses by optimizing source operating points and directing load sharing and energy storage usage to meet operational requirements. Energy ramp rates at various points in the system can also be manipulated by coordinated control of energy flow through converters. An example of such a system is the simplified diagram of a notional DC distribution system [1, 2] shown in Figure 1. In this type of system all major sources and load centers are interfaced to the distribution system by appropriate converters denoted as power conversion modules (PCMs) in the diagram. The notional system has both fuel based generators and an electrochemical energy storage system (ESS). The ESS can serve the shipboard microgrid both as source and load depending on the system need and battery state of charge (SOC) condition [3, 4]. Major loads such as a pulsed load and propulsion loads are interfaced to the main bus. There are also zones of utility loads and two PCMs share the zonal load demand for each zone from two main system buses.

There has been progress in the area of modular converter systems due to continued research and development of the power electronics building blocks (PEBB) concept [5-7]. The PEBB concept has driven advancements in highly modularized converter systems with many identical subsystems such as the Modular Multilevel Converter (MMC). In addition, recent developments in SiC power devices are yielding PEBBs with far greater switching frequencies than Si based devices resulting in an order of magnitude reduction of the time scales as compared to converter systems utilizing conventional IGBT based PEBBs. Faster time scales translate to a need for more capable control systems that is usually being met through the use of FPGA based platforms. Both of those trends

have resulted in the continued evolution of the Universal Controller Architecture (UCA) which attempts to standardize control interfaces for modular power electronic systems [8, 9]. Further development of interface definitions and increasing communication and computational capabilities of new FPGA based controllers provides opportunities beyond simply supporting SiC PEBB based converters.

In a ship-wide PEBB-based power distribution system, control and measurement modules are spatially distributed. Modules that form the control system for single converters are traditionally co-located, however, modules at the application level of each converter control can be networked and furthermore with sufficient communication speed do not even have to be co-located with converter equipment. A study was performed to determine the feasibility of distributing converter application control among the modules within converters and at control layers above individual converter control, such as zonal or bus level controls. In the paper we present a new FPGA-based control architecture along with a proposed communication network topology. We then characterize the impact of the communication topology on latency as the system size is scaled. Finally, the capability of the proposed network-based control architecture for flexible energy management is demonstrated with an example. This can form the basis of a coordinating system control that allows for system wide energy management strategies [10].

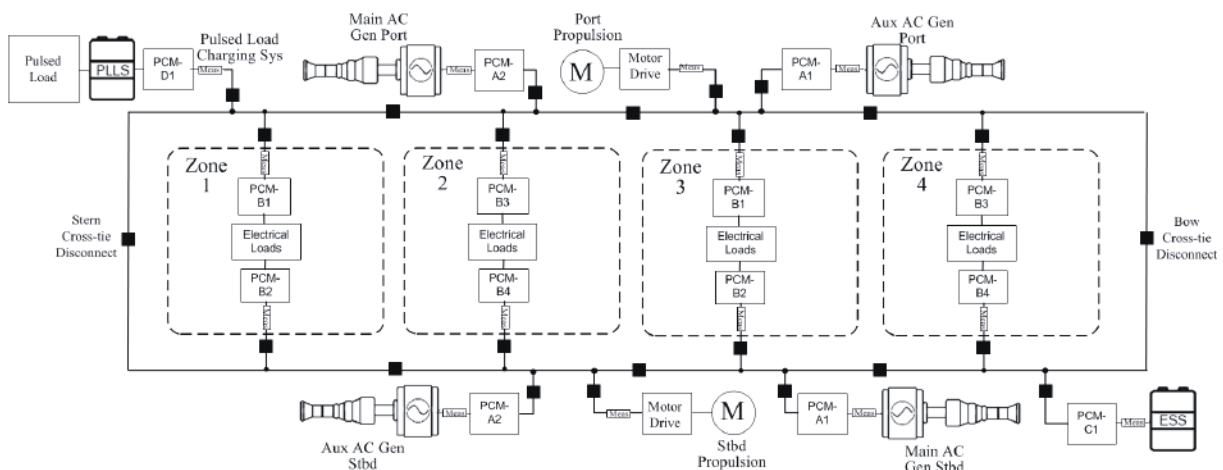


Figure 1: Simplified notional shipboard power system with power converters at major interfaces.

2. Network Based Distributed Control

The stability and performance of the system of PEBB modules is affected by the delay between when measurements are taken and when updated references are received from the controller. Since each level of the PEBB control hierarchy is connected in a local network topology, transitioning packets between control levels will also contribute to the delay. Using a multi-hop network, each control module contains a small integrated router that can both serve as a network interface and serve as an intermediate forwarding point for other messages sent among other control modules. In these types of networks, the message latency is determined by the path length between two control modules as well as the level of congestion on the channels along the path. This latency serves as a constraint for the overall control system design. As such, both the physical topology of the communication network and the routing algorithm are important considerations for the system design.

The proposed multi-hop network topology is widely used for large-scale distributed computing systems to smaller-scale networks-on-chip. However, while these networks seek to minimize average-case latency for varying dynamic traffic, a controller network must guarantee a worst-case latency for regular static traffic. Power electronic control systems consist of multiple control loops and levels or layers of control within a hierarchy. Achieving the minimum bandwidth requires that the routing algorithm equally distribute the communication data transfer load at critical bottleneck locations in the power electronics control system. Such bottlenecks occur at nodes located at control layer boundary interfaces designated as ingress/egress nodes

Several network topologies were evaluated, and some of the candidate topologies are shown in Figure 2. Figure 2(a) shows a simple 1-D bidirectional ring topology, where there is only one minimal-distance path between any two endpoints. The worst case round trip path delay is n , where n is the number of nodes (where a message must traverse $n/2$ rings in both directions). In this topology, each module requires only two bidirectional channels. Figure 2(c) shows a 2-D torus topology, which offers more than one possible minimum-length paths between any two endpoints that are not horizontally or vertically aligned. The 2-D torus has a worst-case round trip latency of \sqrt{n} and requires four bidirectional channels per node. Extending further, a 3-D torus would require six channel

per node and have a worst-case round-trip latency of $n^{1/3}$. We selected a 2-D torus as the best compromise between hardware cost and performance.




Figure 2: PEBB control node communication network topologies evaluated

2.1. Multi-Hop Network Topology

In a 2-D torus topology, the existence of multiple minimum-length paths between most endpoint pairs requires additional considerations in order to maximize the utilization of the network's aggregate channel capacity. In a 2-D torus of width w and height h , a message sent between nodes having addresses (x_1, y_1) and (x_2, y_2) has the following offsets in both dimensions:

$$\Delta x = \min \left(((x_1 - x_2) \bmod w), ((x_2 - x_1) \bmod w) \right) \tag{1}$$

$$\Delta y = \min \left(((y_1 - y_2) \bmod h), ((y_2 - y_1) \bmod h) \right) \tag{2}$$

The required single path routing distance is $\Delta x + \Delta y$ hops but there are

$$length_{min} = \frac{(\Delta x + \Delta y)!}{\Delta x! \Delta y!} \tag{3}$$

possible minimum-length paths. Here we assume that a single node serves as an egress/ingress point to the higher-level control layer. In this way, all nodes transmit and receive measurement and control data to/from this node, and the resulting control loop imposes real-time performance constraints on the network and the on-chip routers.

In order to estimate the minimum round trip latency for various network sizes, we developed an analytical model based on the example assumed system parameters shown in Table 1. Assuming that the chosen parameter values do not exceed the maximum bandwidth of any single channel, each packet will experience a round-trip latency

$$latency_{roundtrip} = 2 \cdot \left(\frac{latency_{Aurora} + latency_{route}}{freq_{FPGA}} + \frac{size_{packet} \cdot 8}{bw_{Aurora}} \right) \cdot \left(\frac{1}{2} w + \frac{1}{2} h \right) \tag{4}$$

Table 2 shows minimum round trip latencies for the parameter values shown in Table 1.

2.2. Routing Algorithms

As shown in Figure 3 (left), the simplest routing scheme for multi-hop networks is X-Y (also called dimension-ordered) routing, in which the network routes packets in the X dimension until the packet reaches a node that is

vertically aligned to the destination and then routes in the Y dimension [11]. X-Y routing is simple to implement and is guaranteed to follow minimal length routes.

For the traffic pattern for PEBB control networks, where all nodes periodically send one packet and receive one packet from the ingress/egress node, the north and south channels into the ingress/egress node must carry more traffic than the east and west channels. In this case, both the north and south channels will experience

$$\frac{w \cdot h - w}{2} \tag{5}$$

packet traversals while the east and west channels will experience only $w/2$ packet traversals. The east-west channels will require a maximum channel utilization equal to

$$bw_{utilization} = \frac{size_{packet} \cdot 8 \cdot \frac{w \cdot h - w}{2}}{latency_{roundtrip}} \tag{6}$$

In order to avoid this load imbalance, the routing algorithm should equally distribute the network traffic across the channels along the minimum paths, especially around the highest congested areas around the ingress/egress node. Ideally, each of the four of the ingress/egress node’s channels should experience

$$\frac{w \cdot h}{4} \tag{7}$$

packet traversals. To achieve this we propose “hub routing”, comprised of a set of pre-computed static routes between each node and the ingress/egress node, where each packet follows a path that keeps its location on the grid closest to the straight line between the node and the ingress/egress node.

TABLE 1: DESIGN PARAMETERS

Parameter	Variable	Expected value
Maximum latency of the Aurora links	$latency_{Aurora}$	53 clock cycles
Packet size	$size_{packet}$	100 bytes
Routing latency	$latency_{route}$	1 clock cycle
Link bandwidth	bw_{Aurora}	10 Gb/s
FPGA user clock frequency	$freq_{FPGA}$	156.25 MHz
Network size	n	100 nodes
Network order, $n = o^2$	o	10 nodes

TABLE 2: MINIMUM ROUND TRIP LATENCIES.

Network size	Round trip latency
5x5	4.3 us
10x10	8.5 us
20x20	17.0 us
30x30	25.6 us
40x40	34.1 us
50x50	42.6 us

The distance between a given node at location (x_0, y_0) and a straight line is computed in the traditional way as

$$\frac{|ax_0 + by_0 + c|}{\sqrt{a^2 + b^2}} \tag{8}$$

Figure 3 (right) shows an example path computed with hub-based routing, where each packet follows a path that keeps its location on the grid closest to the straight line between the node and the ingress/egress node. Thus the maximum-loaded channels will require a maximum channel utilization equal to

$$bw_{utilization} = \frac{size_{packet} \cdot 8 \cdot \frac{w \cdot h}{4}}{latency_{roundtrip}} \tag{9}$$

Table 3 compares the minimum channel bandwidth utilization for both X-Y and Hub Routing, assuming the parameters given in Table 1. X-Y routing requires more than the available 10 Gb/s bandwidth when scaling the network to 30x30, while the Hub routing supports network sizes up to 40x40.

2.3. Related Work

Much of the recent work in FPGA-based multi-hop networks focus on networks-on-chip where a single FPGA contains all the routers comprising the network. In this case the router must be as compact as possible [12,13]. These networks typically use deflection-routing to avoid the need for incorporating buffers into the routers. In this scheme, packets may follow non-minimal routes instead of waiting in buffer when congestion blocks the minimal

path. Because deflection routing increases latency and timing uncertainty it is not appropriate for our application. On the other hand, load balancing routing algorithms developed for distributed computing [14] also generally employ non-minimal routing to maximize throughput at the cost of latency. These networks are also designed for dynamic traffic patterns, as opposed to the static patterns assumed for controller networks. Work that focuses on multi-FPGA systems often focus more on exploration of topology than specific routing algorithms, and often do not consider platform overheads [15, 16].

TABLE 3: MINIMUM LINK BANDWIDTH

Network size	$bw_{utilization}$ (Gb/s): XY	$bw_{utilization}$ (Gb/s): Hub
5x5	1.9	1.2
10x10	4.2	2.3
20x20	8.9	4.7
30x30	13.6	7.0
40x40	18.3	9.4
50x50	23.0	11.7




Figure 3: X-Y routing (left), Hub routing (right)

3. Experimental Platform

To explore the feasibility for a PEBB control network, we used an off-the-shelf KC705 FPGA board with an attached quad-SPF+ transceiver FPGA mezzanine connector (FMC) module. The KC705 has a relatively small Xilinx Kintex-7 FPGA with 203K logic slices and 2 MB of on chip RAM. The board can connect directly to PEBB hardware managers or other PEBB control level interfaces via a secondary FMC expansion connector. The FPGA boards are themselves interconnected via four optical channels to form a control network to form a closed loop control network among the boards. The boards also connect to a secondary, non-real time network through 1 Gb Ethernet for monitoring and control.

The design programmed onto the FPGA is structured as a system-on-chip (SoC), consisting of two Microblaze soft core microcontrollers, on-chip memories, and Direct Memory Access (DMA) engines connected to the four bidirectional 10 Gb/s channels using the Xilinx 66b64b Aurora link-layer protocol.

3.1. Platform Design Considerations

Each PEBB control module collects off-board measurements from the attached power electronics and encodes and transmits the measurements and control data over the multi-hop control network to other control nodes either within the same control hierarchy layer or across a layer boundary as dictated by the control loops in operation. Each node will later receive a corresponding control message from other nodes or layers. Since each operating control loop is deterministic, each control node must complete these tasks according to a fixed control period. In addition, the control system must also route messages on behalf of PEBB control modules on their path to or from other locations in the control network as needed.

The control system is constrained by the communication latency imposed both by the network (in terms of worst-case path length) but also the on-chip overheads of processing and forwarding packets, which may be significant since we are using relatively low-speed microcontrollers. Longer worst case delays will constrain the minimum control period for a given control layer. Likewise, as described in Section 2, the effective channel bandwidth limits the maximum size/scale of the network, since larger networks have more overlapping routing paths requiring more channel bandwidth. Like other network technologies, the effective bandwidth is dependent on the packet size. In

this model, packets comprised of fewer bytes will require a higher interrupt rate to achieve higher utilization of the 10 Gbps channel. Xilinx's DMA IP modules allow for the specification of an interrupt threshold that defines the number of received packets before the module triggers an interrupt, which effectively allows the grouping of multiple small packets between processor interrupts.

Figure 4 shows a block diagram of the design we programmed into the FPGA. The design is logically split into two subsystems mastered by a separate Microblaze microcontroller: the *controller subsystem* and the *monitor subsystem*. The two subsystems are isolated and share only one common peripheral, an on-chip BRAM that holds the controller state. Both processors have local on-chip memory from which they execute their respective program code, both processors have independent interrupt controllers, and both processors have independent timers (the monitor processor uses its timer for the TCP/IP stack).

3.2. Controller Subsystem

The controller subsystem performs the control and routing tasks on behalf of the module and is optimized for latency and determinism. To minimize the amount of unpredictable delays, we took the following steps: (1) store the microcontrollers's software and data in on-chip memory to minimize latency, (2) limit the set of interrupts to only a single timer and four channel interfaces, and (3) place the interrupt controller in fast mode, in which the interrupt controller passes the handler address directly to the processor without any software intervention.

3.3. Monitoring Subsystem

We use a non-real time 1 Gb/s Ethernet interface for monitoring and control of the module. The Ethernet subsystem runs as a fully-custom hardware IP module in the FPGA logic fabric but its TCP/IP stack runs in software. The TCP/IP stack is heavyweight and imposes unpredictable loads on the microcontroller, but when running on its own microcontroller it cannot interfere with the control subsystem.




Figure 4. FPGA based control architecture for power electronic application layer of control.

4. Experimental Results

In this section we describe characterization results of our evaluation platform.

4.1. Latency

In order to evaluate the internal latency of controller subsystem, we evaluated a single FPGA board with a loopback between channel 0 and channel 1 with software to measure the round-trip packet latency. The latency measurement includes contributions from the transmitting DMA engine, the transmitting Aurora interface, the optical transmission latency, the receiving Aurora interface, the receiving DMA engine, and the interrupt controller and are thus representative of the one hop latency.

Figure 5 shows the distribution of packet latencies over 1 million packet transmissions for a 32-byte packet and a 4 KB packet. Note that the Y-axis of the histograms is plotted on a logarithmic scale. For the 32-byte packet, 18.3% of the packets experienced 1150 to 1200 cycles of latency and 81.6% of the packets experienced 1200 to 1250 cycles of latency. On the Microblaze’s 100 MHz clock, 1200 cycles equivalent to 12 us, of which only 25.6 ns is the physical channel transmission time.

For the 4 KB packet, 99.9% of the packets experienced 1250 to 1300 cycles of latency, against a 3.2 μs transmission time. The ~100-cycle latency difference between the 32-byte and 4 KB packet size is equivalent to approximately 1 μs, caused by the higher transmission time for the larger packet.

These results indicate that the packet size has little relative effect on the end-to-end transmission latency, since a 128X increase in packet size required only a 5 to 10% increase in latency. Note that because the platform overheads are 3.9X to 468X that of the channel transmission time.




Figure 5. Observed packet transmit latency for 32 byte packets (left) and 4 Kilobyte packets (right). These results include packet transmission time over the 10 Gbps link (~3 cycles for a 32 byte packet and ~328 cycles for a 4KB packet) and all controller design overheads. Note that the Y-axis is logarithmic.

4.2. Bandwidth

To evaluate the effective channel bandwidth, we added a transmit command to the DMA handler that causes the software to transmit a new packet immediately after receiving a packet. We used a 2000-cycle timer interrupt to gather statistics. Figure 6 (left) plots the effective bandwidth of the channel in Megabits per second versus the packet size. The 32-byte packet size uses 38 Mbps of the channel capacity, the 512-byte packet size uses 614 Mbps, the 4 KB-packet size uses 3.2 Gbps, and the 8 KB-packet size uses 6.5 Gbps. These results are expected since there is insufficient time for the processor to process smaller packets which prevents the processor from achieving full channel utilization. In this test we lose additional performance because we only allow for up to one in-flight packet. To achieve higher bandwidth for smaller packet sizes we measured the effective bandwidth achieved by batching a set of smaller packets into a larger burst, requiring the Microblaze to interact with the DMA controller only after each burst. Figure 6 (right) plots effective bandwidth in gigabits per second achieved by batching an increasing number of packets and is consistent with the trend for increasing a one-packet transmission of increasing packet size.




Figure 6. Observed Aurora channel bandwidth versus packet size (left) and versus number of packets for a fixed packet size of 64B (right).

5. Communication-based Application Layer Control

Single hop latency in the 10µs range has been achieved for the control architecture shown in Fig. 4 which includes all necessary subsystems to implement application level control functions. This indicates that all main converter systems located at the main bus can be share a network distributed application control with total round-trip latency on the order of 100µs-200µs. This is acceptable since Application control for converter applications has a cycle time that is typically in the lower millisecond range [9].

In order to demonstrate benefits of multihop control networks at the converter application control layer a test case is presented using the system shown in Figure 7. This test system is a reduced scale version of the shipboard power system shown in Figure 1. Coordination of zone and main bus converters via communication links is depicted in the figure. Control of system energy flow above the zonal level is accomplished by the Main Bus Level Control. Within the Main Bus Level Control a Cluster is regulated as a group from the same DC bus voltage regulator. Each cluster is denoted with a shaded ellipse. The control scheme for a cluster is shown in Figure 8. The application control, which in this case regulates the voltage on the common DC bus, is external to the converters within the cluster and receives measurements and transmits control references over the multihop control network.

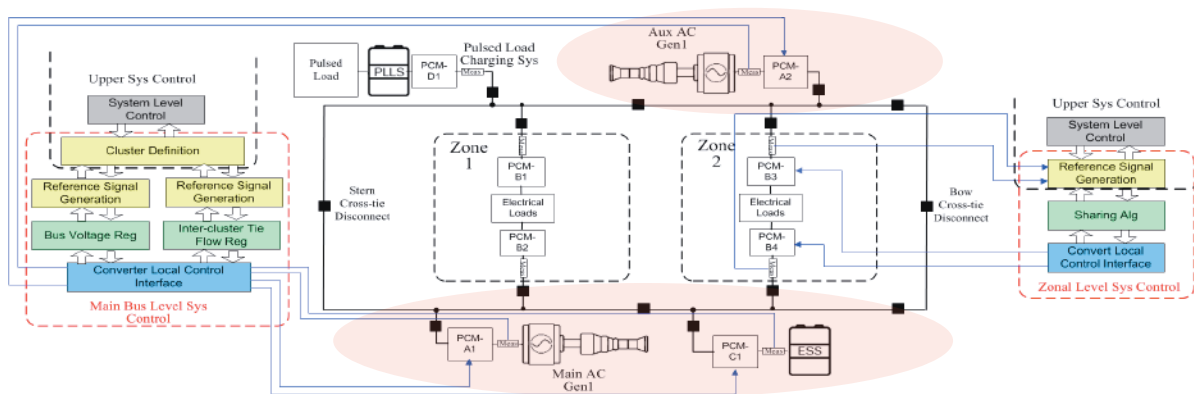


Figure 7. Reduced scale test case system with communication based coordinating control.

Within the main bus level control, a controller regulates the total current between two clusters, in this case equal to the sum of all parallel bus-tie branches connecting the two buses. Thus, the system level control can dictate how energy flows into each zone and how energy flows across bus-ties.

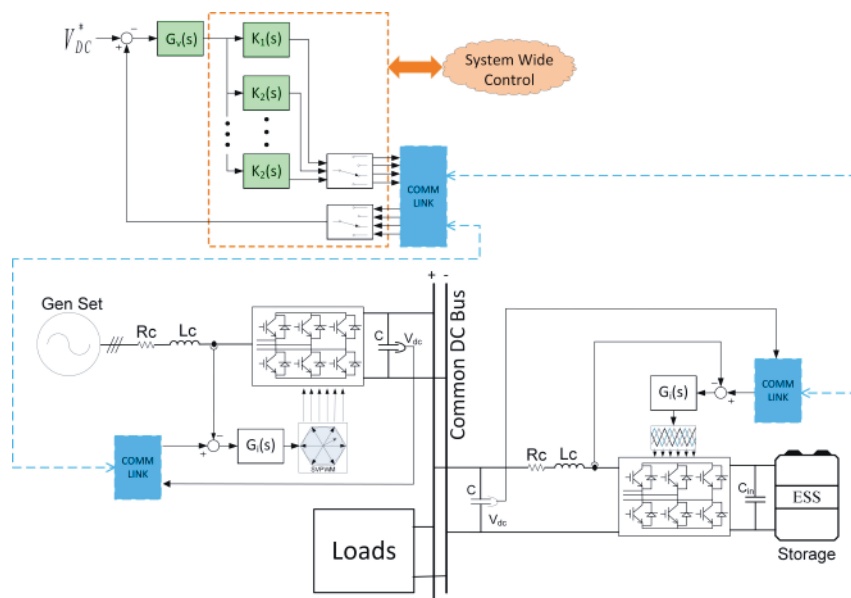


Figure 8. Zonal and main bus level control systems enable flexible routing of energy within the test system.

A test scenario was run to demonstrate the tight control of the cluster and control of energy flow across the bus tie. Initially both zones are loaded at 0.5pu with the total zonal load as the per-init base and 0.2pu is moved across the bus-tie from the Aux Generator side to the Main Generator side. Current across the bus-tie is varied from positive (Aux to Main) to negative (Main to Aux). The ESS is shut down by its control system when the SOC reaches the minimum limit of 40% which occurs just slightly after t=1 sec. Zone one has a reduction in loading to 50% and zone two has a reduction in loading to 25% at t=1.5 sec. Zone two loading then increases to 75% at t=1.75 sec.

The per-unit power of each source is shown in Figure 9. Note that the Main Generator and ESS are defined as a cluster and thus provide an equal contribution while the ESS is in operation. The contribution of the two clusters changes according to the energy commanded across the bus-tie. The bus tie current variation is shown in Figure 10. Note that limits of what can be sent across the bus-tie are determined as a function of loading on a bus. When the Aux Generator is completely unloaded, as shown in Figure 11, no more power can be moved across the tie than is drawn by the load. The effect of the control on the Main Generator current is shown in Figure 12.




Figure 9. Power per-unit of the three main bus sources.




Figure 10. Current across the bus-tie (positive is from Main to Aux)




Figure 11. Aux Generator phase A voltage and current




Figure 12. Main Generator phase A voltage and current

6. Conclusions

This paper describes a general methodology for building power electronic building blocks based converters and systems of converters, where individual PEBB modules are coupled with embedded controllers interconnected on a distributed multi-hop communication network. We evaluated two routing algorithms and used an analytical performance model to evaluate the impact of communication load balancing on system scale. Our proposed hub-based routing algorithm is capable of balancing channel load for a static traffic pattern where all modules engage in a closed control loop with a single ingress/egress point to other control layers.

A FPGA design is developed that is decomposed into two mostly isolated subsystems. One of these systems is designed for real-time control and control network routing and the other for non-real time instrumentation and monitoring. The network performance of a 10 Gbps communication infrastructure was characterized including all of the overhead of sub-systems that provide a flexible platform for application control. The test numbers obtained using the 2-D Torus network configuration along with the developed Hub Routing method have shown that the application layer of control can function as the most fundamental system layer within a distribution system comprised of many power electronic converters.

Using the control architecture test results and a notional shipboard system it was demonstrated that current FPGA and gigabit speed serial communication technologies allow for a very high degree of energy flow control in power electronics based distribution systems.

7. Acknowledgements

This work was supported by the Office of Naval Research under award N00014-16-1-2956.

8. References

- [1] T. J. McCoy, “Integrated Power Systems- An Outline of Requirements and Functionalities for Ships”, Proceedings of the IEEE, Vol. 103, No. 12, Dec 2015, pp. 2276-2284.
- [2] N. Doerry, “Naval Power Systems”, IEEE Electrification Magazine, Vol. 3, No. 2, June 2015, pp. 12-21.
- [3] Yann Riffonneau, Seddik Bacha, Franck Barruel, Stephane Ploix, “Optimal power flow management for grid connected PV systems with batteries” IEEE Transactions on Sustainable Energy, Vol.2, No. 3, July 2011.
- [4] Shuo Pang, Jay Farrell, Jie Du, and Matthew Barth, “Battery state-of-charge estimation” Proceedings of the American Control Conference, Arlington, VA June 25-27, 2001.
- [5] T. Ericson, “Power Electronics Building Blocks – a Systematic Approach to Power Electronics,” Proceeding of the 2000 IEEE Power Engineering Society Summer Meeting, Volume: 2, 16-20, July 2000, pp. 1216-1218.
- [6] F. Wang, S. Rosado, D. Boroyevich, “Open Modular Power Electronics Building Blocks for Utility Power System Controller Applications,” Proceedings of IEEE 34th Power Electronics Specialist Conference, Vol. 4, 15-19, June 2003, pp. 1792-1797.
- [7] T. Ericson, “SiC-PEBB based zonal distribution system architectue,” Ericson Innovations, May 2013.
- [8] IEEE Std. 1676-2010, “Guide for Control Architecture for High Power Electronics (1 MW and Greater) Used in Electric Power Transmission and Distribution Systems”, 11 Feb. 2011.
- [9] Ginn, H.L., Hingorani, N., Sullivan, J.R., Wachal, R., “Control Architecture for High Power Electronics Converters”, Proceedings of the IEEE, Vol. 103, No. 12, Dec. 2015, pp. 2312-2319.
- [10] M.R. Hossain and H.L. Ginn III, “Real-Time Distriubed Coordination of Power Electronic Converters in a DC Shipboard Distribution System,” IEEE Trans. On Energy Conversion, Vol. 32, No. 2, June 2017, pp. 770-778.
- [11] W. Dally and B. Towles. Principles and Practices of Interconnection Networks. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2003.
- [12] N. Kapre and J. Gray, "Hoplite: Building austere overlay NoCs for FPGAs," Proc. 25th International Conference on Field Programmable Logic and Applications (FPL).
- [13] Nachiket Kapre, "Implementing FPGA Overlay NoCs Using the Xilinx UltraScale Memory Cascades," Proc. 25th IEEE International Symposium on Field-Programmable Custom Computing Machines, 2017.
- [14] Arjun Singh, William J. Dally, Amit K. Gupta, Brian Towles, "GOAL: a load-balanced adaptive routing algorithm for torus networks," Proceedings of the 30th annual international symposium on Computer architecture 2003.
- [15] Trevor Bunker, Steven Swanson, "Latency-Optimized Networks for Clustering FPGAs," Proc. 21st Annual International IEEE Symposium on Field-Programmable Custom Computing Machines, 2013.
- [16] Andrew G. Schmidt, William V. Kritikos, Rahul R. Sharma, Ron Sass, "AIREN: A Novel Integration of On-Chip and Off-Chip FPGA Networks," Proc. 17th IEEE Symposium on Field Programmable Custom Computing Machines, 2009