Design of a Distributed Control Architecture for the SAMURAI Deep Submergence Manipulator

Nicholas M. Limparis, Dr. David L. Akin
Department of Aerospace Engineering
Space Systems Lab, University of Maryland
College Park, MD, USA

Abstract—**The design, development, and testing of a distributed control architecture for a dexterous deep submergence manipulator is described. Many manipulators designed for the undersea environment use a centralized control mechanism, be it a system of valves for a hydraulic manipulator or a central electronics housing for an electrical manipulator. Without the constraints of deep submergence, many modern manipulator systems have adopted distributed architectures to facilitate higher bandwidths and greater modularity for maintenance and reconfiguration. The Subsea Arctic Manipulator for Underwater Retrieval and Autonomous Interventions (SAMURAI) manipulator, developed under NASA and NSF funding, uses a distributed network of Local Processing Units (LPUs) that share the loading of the control system for the manipulator and are co-located in miniature surface-pressure housings at each of the actuator sections. The design of this system is detailed, showing the challenges of such a design, as well as the advantages and limitations of this system as compared to a centralized architecture. In addition, the paper also describes a prototype development architecture for the SAMURAI manipulator that uses commercial off the shelf motion controllers and a commercial data bus in a similar distributed architecture. This system provides reduced but still powerful functionality as compared to the custom architecture, with shorter development times and at significantly lower costs.**

***Index Terms*—Deep Sea, SeaBED, SAMURAI, Autonomous, Manipulator, Electric.**

INTRODUCTION

The Subsea Arctic Manipulator for Underwater Retrieval and Autonomous Interventions (SAMURAI) manipulator is a robotic manipulator that was designed to meet the needs of The NASA Astrobiology Science & Technology For Exploring Planets (ASTEP) program. This manipulator was finished under a National Science Foundation (NSF) grant and is designed to be able to reach full ocean depths to be able to perform biological sampling of the deep ocean for autonomous sample return to the surface [1]. Manipulators for deep submergence have been developed previously, most have been hydraulic and almost all have been controlled using a centralized control mechanism. For a hydraulic manipulator this would consist of a series of controllable valves and for electric manipulators a centralized housing to drive all of the motors of the system.

Without the constraints of deep sea applications, many robotic systems have adopted the idea of a distributed control mechanism to facilitate higher control bandwidths and to add modularity to the system for ease of maintenance and allow reconfiguration of the manipulator. The SAMURAI manipulator embraces a distributed control methodology to provide these benefits in the harsh deep sea environment [2].

THE PHYSICAL CONFIGURATION OF A DISTRIBUTED CONTROL MANIPULATOR

In order to better understand the benefits of distributed manipulator architecture, it is beneficial to grasp at a high level the physical layout of such a manipulator. The Command and Data Handling (C&DH) system is the first member of the manipulator system. The C&DH system's purpose is to interpret the commands issued to the manipulator system and distribute the corresponding manipulator commands to the local controller modules of the system. Attached to the C&DH system is a distributed controller that processes the commands issued from the C&DH system. In the nomenclature that will be used for the rest of this paper, a distributed controller will be referred to as a Local Processing Unit (LPU). Each LPU is tasked with running an internal control loop to command the motion for each of the motors attached to that LPU. The number of motors attached to each LPU is determined by the needed functionality of that particular module. Individual modules are attached together with a data interface originating from the C&DH system. A generalized representation of a distributed controller architecture can be seen in Fig. 1 demonstrating the interconnections within the system.

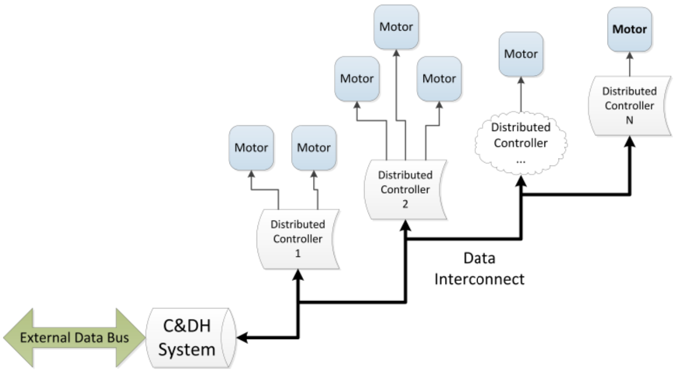

Figure 1. A conceptual diagram of a distributed control architecture for a robotic manipulator

THE LAYOUT OF THE SAMURAI ROBOTIC MANIPULATOR DISTRIBUTED CONTROL SYSTEM

The design of the SAMURAI manipulator consists of four major components. These are the C&DH system and the three joint modules. This configuration is shown in Fig. 2.

The C&DH system is housed in a titanium cylinder colloquially called the "pill" due to its capsule shape. This housing contains the two main computers, one for data handling and control and another to perform visual servoing tasks. The pill also contains the interfaces to the external data bus to receive commands and the data interconnect to the LPUs. Additionally, DC/DC converters are used to switch higher voltages, from the base AUV that SAMURAI is attached to, to lower voltages need for operation of the pill internals and for the LPUs.

Each of the joint modules in SAMURAI consists of a LPU that controls two electric motors. The shoulder module is a yaw and a pitch joint, the elbow is a pitch and a roll joint, and the wrist is a pitch and a roll joint. This is summarized in Table 1 showing that in its assembled configuration, shown in Fig. 3, the manipulator is a yaw-pitch-pitch-roll-pitch-roll (y-p-p-r-p-r) configuration.

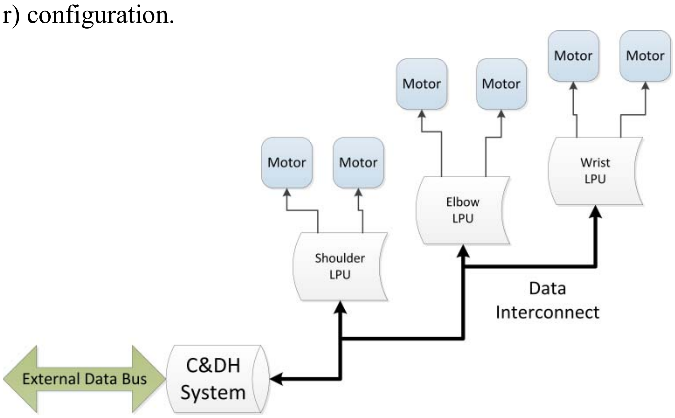

Fig. 2. The conceptual diagram for the distributed control architecture of SAMURAI

The external data bus for the SAMURAI manipulator is Ethernet to interface with the SeaBED class JAGUAR AUV developed by the Woods Hole Oceanographic Institute (WHOI) [3]. The data interconnect for the LPU is either IEEE 1394a or CAN bus depending on the LPU electronics installed in the arm.

TABLE I
SAMURAI JOINT SUMMARY

Joint Number	Location	Motion
1	Shoulder	Yaw
2	Shoulder	Pitch
3	Elbow	Pitch
4	Elbow	Roll
5	Wrist	Pitch
6	Wrist	Roll

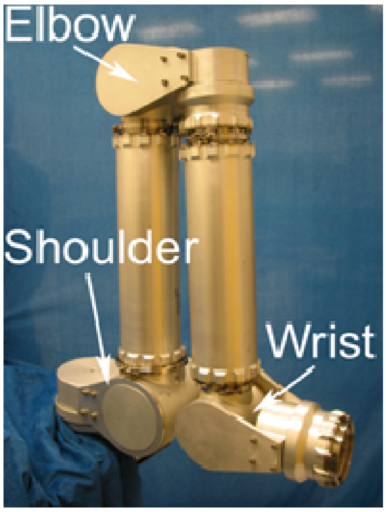

Fig. 3. The assembled SAMURAI manipulator sans end-effector

DESIGN OF THE LPUS AND THE PILL INTERNALS

The Pill Power Supply

The pill performs the vital task of acting as an interface between JAGUAR and SAMURAI. The items that had been completed from the work done under the ASTEP program were the two computers, mac minis, and the Ethernet networking infrastructure. The two computers each perform different tasks. The first one performs visual analysis of the stereographic images coming from the time synced cameras. The second performs control of manipulator and visual servoing based upon the vision data from the first computer. The processing power of the computers does not come cheap, though, in terms of electricity. Therefore, it is necessary to have a method for JAGUAR turn them off when not in use. Since JAGUAR provides the SAMURAI system with 48 volts of electricity, power conversion is a must. In this system, the pill needs to have DC/DC conversion for 12, 18.7 and 28 volts for the network switch, computers, and arm control power respectively. The pill also needs to be able to turn these voltages on and off if needed. Finally, the power conversion system needs to provide health data on the converted power. This health data includes voltage monitoring, current monitoring, and temperature monitoring. In order to provide all of this functionality a power supply unit was designed to meet these needs (Fig. 4).

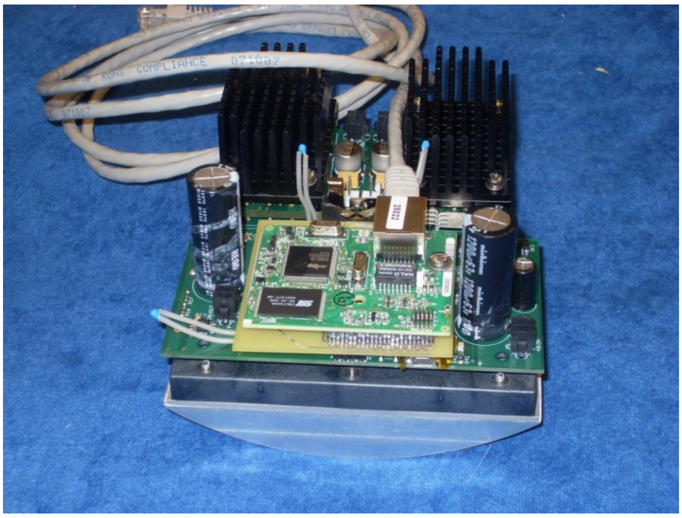

Fig. 4. The pill power supply

The Pill Power Supply has four switch mode power supplies (SMPS) on it to provide the needed power conversion from 48 volts. First, a large ½ brick, 300-watt power supply converts the 48-volt input down to 28 volts for the manipulator control power. From the 28 volts there are two ¼ brick SMPS that convert the 28 volts to 18.7 for the computers in the pill. Finally, a small SMPS converts the 48 volts to 5 volts and a linear LDO voltage regulator that converts the 5 volts to 3.3 volts for integrated circuits on the power supply board. For health monitoring of the power supply, a Digi Rabbitcore 4001 was chosen for its analog inputs and native Ethernet interface. In addition to the voltage and current monitoring that the Rabbitcore provides. The module was set up to take temperature measurements of each of the DC/DC converters in addition to the ambient air temperature of the pill. Finally, the Rabbitcore provided an interface that JAGUAR could use to boot the computers inside of the pill. With all of these items on this power supply board we quickly ran out of the room that was allotted for the power supply in the pill. We solved this issue by changing the shape of the board where the power supply extends beyond the edge for the cylindrical portion of the pill and into the hemisphere of the pill. This gives the board the unique shape that can clearly be seen in Fig. 5.

Fig. 5. The pill power supply showing its unique shape

The FireWire Network Compensation Board

The final part of the pill, completed under the NSF contract, was the FireWire camera interfaces. The machine vision cameras that were chosen for the ASTEP program are Point Grey Scorpion cameras. These cameras have a 1394a FireWire interface which provides the ability to have deterministic response of the cameras to commands. The design challenge with the camera system was not with the cameras themselves, but with how FireWire was implemented in the work done previously under the ASTEP program until its cancelation. The previous work tried to pass the 1394a signals directly through the underwater connectors and cables that were not designed for the 1394a signals. Using an oscilloscope the eye diagram for the FireWire signals were analyzed and were found to have too many reflections on the controlled impedance bus to function properly. To solve this design challenge an impedance matching network needed to be added to the system. What we selected was a device manufactured by TC Applied technologies and designed by Ecqulogics that actively compensates FireWire signals to allow them to be passed over CAT 5e or higher Ethernet cable, which is readily available as an underwater cable.

Now even with this solution for the network matching circuit a Printed Circuit Board (PCB) still needed to be manufactured. FireWire, being a controlled impedance bus, requires that all of the traces on the circuit board needed to have an impedance of 120Ω. In order to comply, we selected a PCB manufacturer that meet all of needs in terms of layers and via sizes and proceeded to calculate the transmission line impedance for a microstrip over a ground plane. From this and the information about the PCB manufactures process it was determined that our transmission line would need a width of 16.701 mils to have 120Ω impedance. With this calculation in hand, the next hurdle to jump was the need for one of the active compensation networks on both ends of the underwater cable. This put a size limitation on the final design of a 3-inch diameter circle so that it could fit in the deep submergence camera housings. Layout for these boards commenced and the finalized product is shown in Fig. 6.

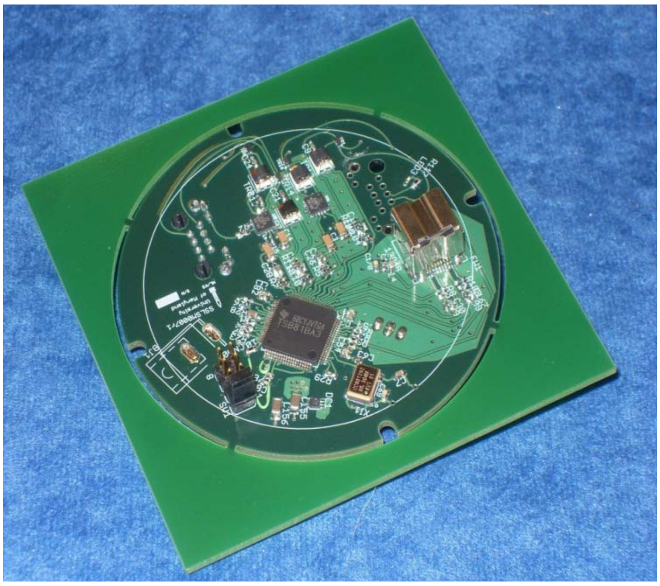

Fig. 6. FireWire Network Compensation Board

The board acts as a FireWire hub and has three nodes. Two are actively compensated and the other is a standard FireWire port allowing the computers to go from FireWire to compensated FireWire and back. In testing, the boards performed as expected providing the full S400 data rate of the 1394a standard over the underwater cables and underwater connectors.

The LPU Stack

The LPU stack consists of a computer and dual motor controllers that live at each one of the joint pairs inside of the surface pressure bubble. The design consists of a 4 circuit boards: a DC/DC board, a cpu and interface board, a communications board, and the motor control board. Figure 7 shows the finalized LPU stack and Fig. 8 shows the individual boards as listed from right to left with the CPU board having its programming board attached.

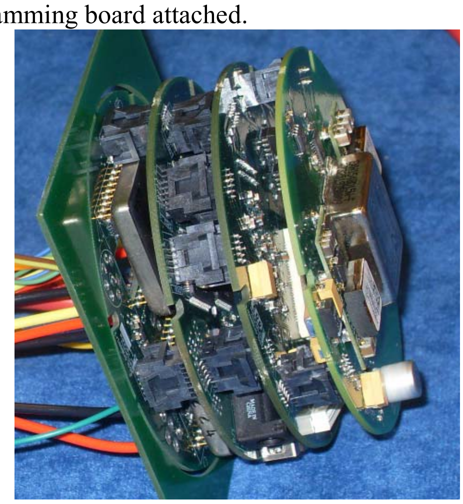

Fig. 7: The LPU Stack

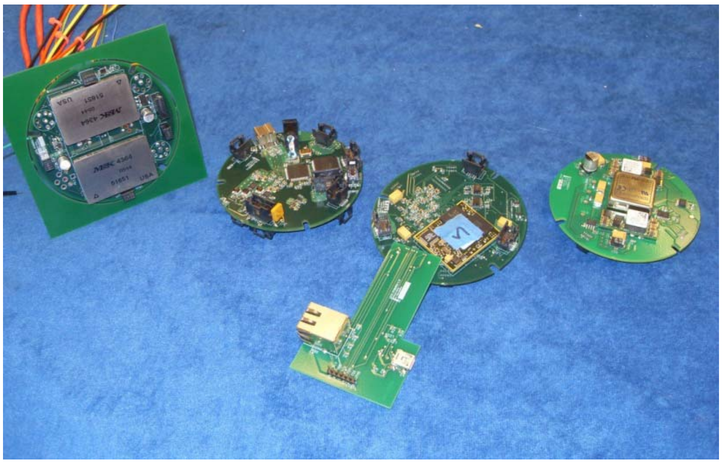

Fig. 8: The LPU Stack Disassembled

With the addition of the 4th board we still had to fit into the small 3.5 inch diameter by 2.5 inch tall cylinder that is the surface pressure bubble. It was quickly realized that this was not enough volume for the design and began to look for solutions to the issue. After many attempts to reduce the size of the electronics, it was determined that if we wanted to retain the functionality of the system that we would need to keep all of the components of the system. Therefore, we had to look for a mechanical solution. After much deliberation it was decided that a bubble extension ring could be manufactured that retained the depth rating of the system while still adding ¾ of an inch to the height. Figure 9 shows the bubble attached to the electronics penetrator plate before the addition of the extension ring.

The next problem presented by the stack is how to arrange the boards within the stack. The major driving factors are thermal constraints and component heights. With the size of the bubble there is ¾ of an inch allocated for each one of the boards with a little extra for the top board in the center do to the shape of the bubble itself. Since the DC/DC board has some components that are taller than ¾ of an inch it will need to go on the top to allow for the height of the components. Additionally the motor driver board needs to be located at the bottom of the stack in order to facilitate conductive cooling to the manipulator. This leaves the communication board and the CPU board to have their positions assigned. The benefit of putting the CPU board second to the bottom is that it reduces the number of connectors that the SPI communication interface for motor controller needs to go through; however it increases the number of connectors that the FireWire signal has to go through which we need to minimize. Therefore, these requirements dictate the order of the LPU stack to be, from the bottom up, motor control, communications board, CPU board, and DC/DC conversion.

Fig. 9. SAMURAI electronics housing

With this stack up of boards, we are left with one final design challenge for the LPU stack, power distribution. By power distribution, it is meant that for the power to come into the system it must come in from the penetrator plate at 28 volts, travel up all of the four boards to the DC/DC Conversion board, and then back down to all of the individual boards below (Fig. 10). This design challenge made connector selection, and placement a difficult task since power needs to go up the connectors back down the connector while at the same time not induce electromagnetic interference on the data lines for the system. This was accomplished by using many high-density connectors to pass the power up and then back down again while carefully choosing the position of the data lines in the connectors to reduce the possibility of electromagnetic interference.

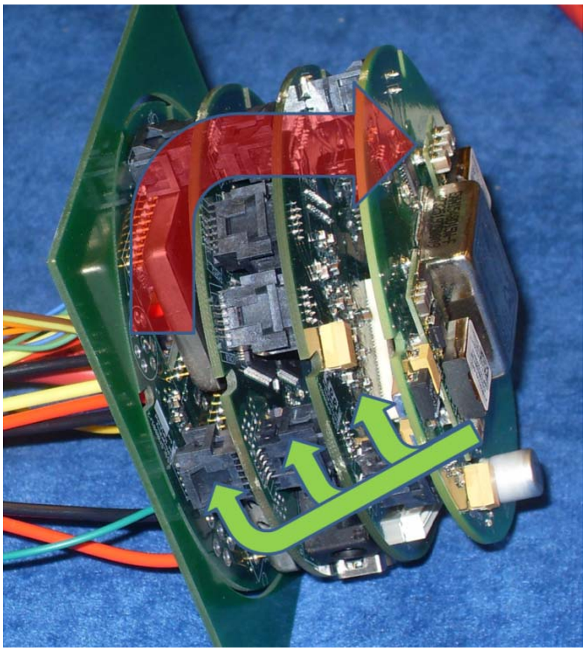

Fig. 10. Conceptual diagram of power flow within the LPU stack

The DC/DC Conversion board

The DC/DC Conversion board is a straightforward board in that it serves one task: power conversion. The complications arise from the number of voltages needed, high efficiency requirements, and height restrictions. The voltages that are needed are 5 volts, for running many of the smaller integrated circuits and optoisolators, 3.3 volts for the larger integrated circuits, and ±15 volts for motor control. Health and monitoring data is also needed to ensure the proper operation of the switch mode regulators. Therefore an analog to digital converter (ADC) is used to provide currents and voltages for each of the regulators to give real time feedback to the CPU board of the LPU stack. The finalized design for the DC/DC conversion board is shown in Fig. 11.

The CPU and Data Interface Board

The CPU board is the heart of the distributed controller system. This board runs all of the controllers that command the motors in the actuators, monitors the health of the different systems of the LPU and communicates a filtered data set back to the C&DH system in the pill. In the design for the SAMURAI manipulator the two main components of this board are the System-on-Module (SOM) and the Field Programmable Gate Array (FPGA).

The SOM consists of a microprocessor, memory, a boot ROM, an Ethernet interface and connectors to the microprocessor's address and data busses. uClinux is the operating system of choice for this system since the Analog Devices Blackfin DSP does not have a memory management unit (MMU) and therefore cannot run operating systems that need a MMU. A patch to the kernel makes the operating system "real time". The design choice for choosing a SOM over a fully integrated design on the board is twofold. First, having a SOM means that for the small number of devices that we actually will be making for this system we can ignore the need to perform the arduous task of designing, manufacturing, and testing our own design for the processor system we can use a more vetted system that can be had as commercial off the shelf (COTS). Second, the SOM allows for upgrade paths for the board without the need to redesign the entire board since the SOM can just be exchanged for a new one. Additionally, all of the controllers are stored on a micro SD card allowing for easy updates to the software of the system without the need to retest the base code stored in the ROM of the SOM.

The FPGA provides vital services for the board in that it automates the collection of data from the sensors and provides them to the SOM as if all of the different devices are a part of the memory structure for the SOM. This memory mapped I/O structure allows for large amounts of data to be transferred quickly when needed and to have it be stored in buffers when the SOM does not need that data immediately. This allows the control loops and system critical processes to occur at the rates that are needed without having to deal with the overhead of processing sensor data at fixed intervals. Additionally, this device is used to monitor health of the systems when the CPU is in a standby state and will wake the CPU if an event occurs that is out of the normal range for the LPU.

Fig. 11. The DC/DC conversion board

In order to program the SOM and the FPGA a debug and programming board was made in order to remove these components from the main stack. The reasoning behind this was to remove components that consume power but do not need to be used when the LPUs are sealed up for submergence. Additionally, moving these components to a separate board freed board space on the CPU board thereby allowing a reduction in the number of copper layers that were needed. Finally the debug board has light emitting diodes to allow error messages to be displayed when the board is connected. Figures 12 and 13 show the top and bottom of the CPU board showing the SOM and the FPGA layout respectively.

The Communications Board

The communications (COMMS) board of the LPU stack is the interface between the LPU's CPU and the C&DH system in the pill. The COMMS board performs the same task as the FireWire Network Compensation Board in that it takes the FireWire signals and matches the impedance of the cables that the signals are transmitted over. The COMMS board also has the additional task of acting as a message interface for the CPU allowing the CPU to inject massages to the FireWire bus will still allowing the C&DH system to receive and transmit message transparently on the bus. The final purpose of the COMMS board is to pass signals to and from the motor control board back to the CPU board. This places signal path constraints on the COMMS board to prevent the crosstalk of the high speed FireWire bus with the analog signals and data busses going to the motor control board and actuator on the outside of the pressure vessel. As seen in Fig. 14, the board contains a large number of connectors to be able to pass all of the data signals that are needed for the power and data interconnects. These connectors consume most of the physical board space for the COMMS board and were the driving influence for the design of this board.

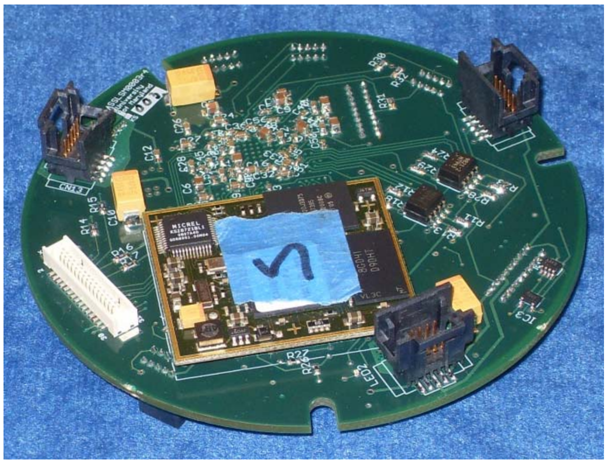

Fig. 12. Top of the CPU board showing a SOM plugged into its socket

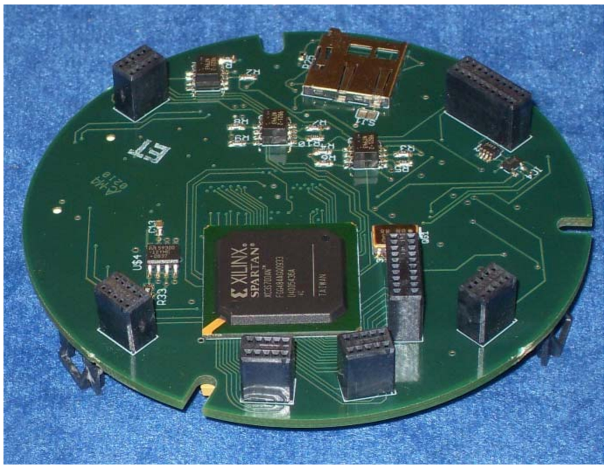

Fig. 13. Bottom side of the CPU board

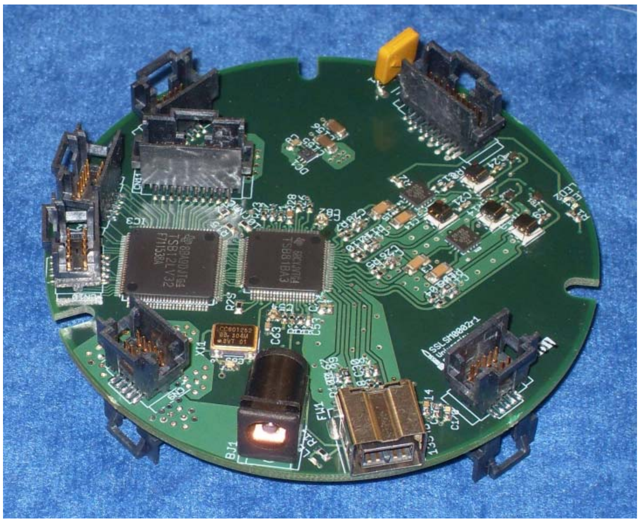

Fig. 14. Top of the communications board

The Motor Driver Board

The motor driver board is the host to the components that interface to the two motors in each one of the joint pairs. The motor drivers are two, three phase DC brushless motor drivers in hermetic packages. The support circuitry around the motor drivers monitors the currents and voltages of the motors along with the necessary motor starting capacitors. The most pressing issue that the motor board has to deal with is the design challenge of heat sinking the motor drivers to the structure of the arm. When the stack is assembled, there is only 1/8th of an inch between the top side of the motor driver and the communications board. Following some testing and analysis it was determined to not be a large enough surface to get the heat out of the motor drivers. The solution to the design challenge was to attack the issue from the other side. In order to get the heat out of the bottom of the boards large cutouts were made in the circuit boards underneath of the motor drivers. Figure 15 shows the top side of the circuit card and Fig. 16 displays the cutouts as seen from the underside of the circuit card. The final component of this board is the interface to the high pressure electrical penetrators that gives the pressure vessel's contents access to the external components of the joint module.

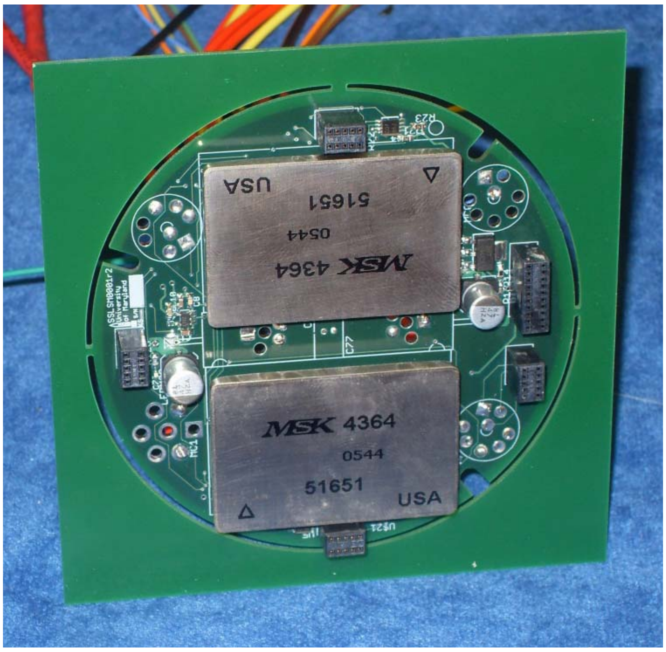

Fig. 15. Top side of the motor driver board

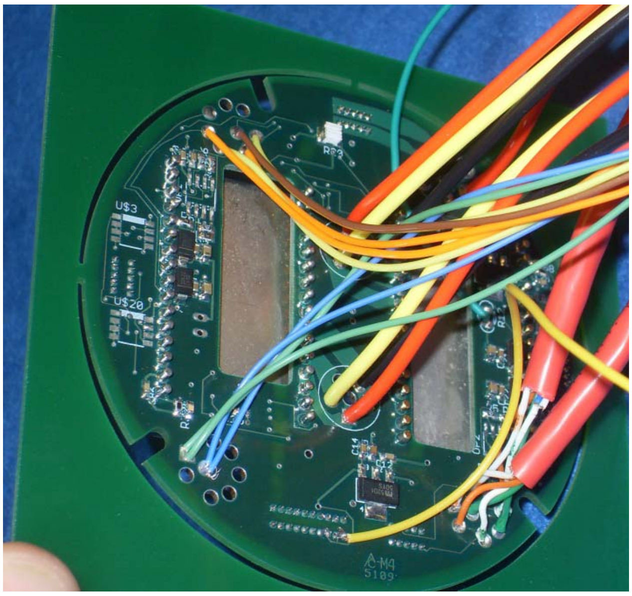

Fig. 16. Bottom side of the motor driver board showing the cutouts for the heat sinks that pass through the board.

An Alternate LPU

A design study was performed to see if we could reduce the complexity of the LPU stack and yet retain most of the core functionality of the system. Many systems were considered and the Whistle motion controller was chosen from Elmo Motion Control. It is an industrial motion controller that was small enough to fit into the bubble with the extension ring. Heat sinking for these devices is an issue since the do not have a good conduction path to the outside of the bubble. An image of a prototype Whistle based LPU is shown in Fig. 17.

One of the features of the LPU stack that was lost by going with the Whistle modules is that they cannot self-identify their relative location in the chain, e.g. the third module in the chain knowing that it is the third, instead it only knows that it is address X. Instead the modules must be programmed and the C&DH system notified of the organization of the manipulator in order to allow the manipulator to move properly. Without the capability to self-identify the LPU node location in the chain, the ability to swap modules is not possible without reconfiguration of the C&DH system. Regardless of the features lost from the original LPU stack, the Whistle based LPUs still proved that for static configurations of joint modules; a LPU can consist of a set of simplified controllers that only run the control loops for each of the motors in the manipulator.

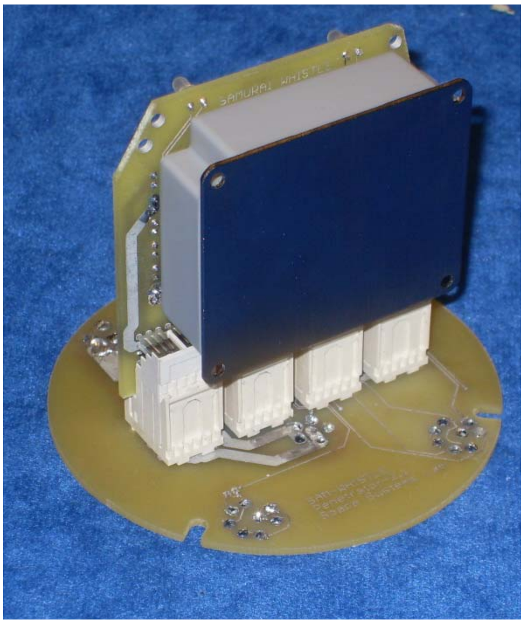

Fig. 17. Prototype of a Whistle module based LPU

Conclusions

The use of a distributed control architecture for the manipulator presents a unique challenge for the system design. The many distributed LPUs require the manipulator system to have many surface pressure vessels. This adds mechanical and electrical complexity. However, distributed control architectures add capabilities, such as modularity and the ability to exchange and reconfigure joint modules for changing workspace needs. Although, this architecture may not be necessary for a particular mission, the ability to make changes to the system as needed to meet mission objectives makes a strong case for considering a distributed control architecture for a deep submergence manipulator.

The designs presented in this paper for the SAMURAI system are one method of implementing a distributed control architecture. These designs meet the requirements for the mission that the SAMURAI system was tasked with. It was designed to allow future expansion of the system and to be capable of meeting the future performance needs of the system.

Acknowledgment

The authors would like to thank our sponsors throughout our development of these systems: David Lavery and the National Aeronautics and Space Administration, work conducted under grant NNG04GB31G; Kandice Brinkley and the National Science Foundation, for work conducted under grant OCE0752347; and the Maryland Industrial Partnerships program, under grant 4211.

The authors would also like to thank Hanumant Singh and Robert Sohn of the Woods Hole Oceanographic Institute for their collaboration, help and use of the Jaguar AUV.

References

Lewandowski, C.; Akin, D.; Dillow, B.; Limparis, N.; Carignan, C.; Singh, H.; Sohn, R.; , "Development of a deep-sea robotic manipulator for autonomous sampling and retrieval," *Autonomous Underwater Vehicles, 2008. AUV 2008. IEEE/OES* , vol., no., pp.1-6, 13-14 Oct. 2008 doi: 10.1109/AUV.2008.5290533

Barrett E. Dillow, David L. Akin, and Craig R. Carignan, "Development and Testing of a Dexterous Manipulation Capability for Autonomous Undersea Vehicles" AIAA 2009-2003, *AIAA Infotech@Aerospace Conference and AIAA Unmanned/Unlimited Conference*, Seattle, WA, April 2009

Singh H., Can A., Eustice R., Lerner S., McPhee N., Pizarro O., and C. Roman. "SeaBED AUV Offers New Platform for High-Resolution Imaging," EOS, *Trans. of the AGU*, 2004; 85(31): 289, 294-295.