

The NEPTUNE power system: design from fundamentals

Harold KIRKHAM¹⁾ Phil LANCASTER²⁾ Chen-Ching LIU³⁾ Mohamed EL-SHARKAWI³⁾
Bruce M. HOWE⁴⁾

1) Jet Propulsion Laboratory, California Institute of Technology

2) Alcatel Submarine Networks

3) Electrical Engineering Department, University of Washington

4) Applied Physics Laboratory, University of Washington

KEYWORDS: Power delivery, NEPTUNE, Mission Assurance

ABSTRACT

This paper describes the process and recounts the decisions that have been made in the design of the NEPTUNE power system. The design process has consisted of a number of top-level decisions based on trade-off studies, followed by a number of more detailed developments based on fundamental considerations of the system needs. The system will use a parallel 10-kV dc scheme to deliver power at 400 V to the user. Altogether, around 100 kW can be delivered to the load distributed along the underwater cable, and up to 10 kW can be obtained at any one science node. In the design process, several commonplace solutions of power engineering have been re-examined and abandoned.

1 INTRODUCTION

The design of the power delivery subsystem for NEPTUNE is the responsibility of a group of electrical engineers, mainly power engineers, at the University of Washington and at the Jet Propulsion Laboratory. The early designs were based on conventional practice in power engineering, and surprised no-one.

The design work began with some basic trade-offs (should the delivery system use ac or dc, should it be a parallel system or a series one), and then addressed some more detailed issues, such as the voltage and power levels to be used, and the means to start up a long string of cable sections by energizing them from shore.

After a year or so of work, the possibility of using another approach to part of the design problem was raised. With much angst, the NEPTUNE power group developed what had originally been a suggestion for powering a control system into a revised version of the power delivery approach, along the way re-examining many of the solutions they had chosen.

This paper recounts the story, and presents it as an example of the merging of cultures – in this case the power engineering culture and the ocean engineering culture.

2 BACKGROUND

The proposed NEPTUNE observatory is a linked array of undersea instruments on the Juan de Fuca plate in the northeastern Pacific Ocean 1). Fiber-optic power cable will connect land-based scientists, students, decision makers, and the public to distributed sensors above, on, and beneath the seafloor. NEPTUNE is an integrated system providing power and communications to a large number of instruments distributed over a seafloor area of about 500 by 1,000 km.

If the route eventually adopted is as shown in the many maps that have been published, approximately 3000 km of cable will be deployed, interconnecting up to around 40 science nodes. The task of the power group is to use the cable to deliver as much power as possible, as reliably as possible, to the science nodes.

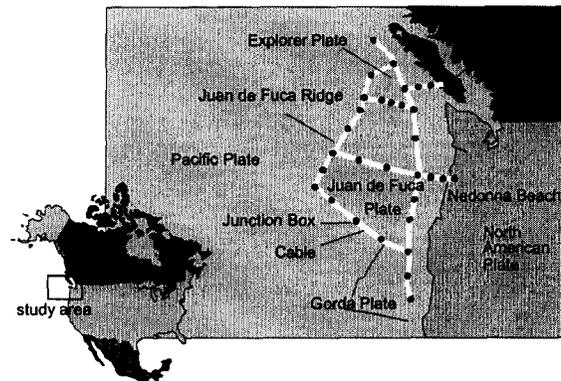


Figure 1. Proposed NEPTUNE observatory layout

Early trade-offs included consideration of ac or dc (dc was chosen because of the cost that would be incurred compensating the cable capacitance for ac use); series or parallel operation (parallel was chosen, as more power can be delivered, and the cable can be branched without the need for electronics). A voltage level of 10 kV was selected for the backbone because this is the conventional rating of the cable 2,3).

The task of designing a compact dc/dc converter was begun early in the project. The converter was specified to have an input voltage range of 5.5 kV to 10 kV (the lower limit is set by system stability considerations and the upper limit by the cable rating), and an output voltage of 400 V (this being the highest voltage that a user could conveniently convert down using a commercial off the shelf converter).

A maximum current of 10 A was imposed because the cable resistance (about 1 Ω per km) was so high. Higher currents would start to give voltage-limit problems. For example, at 50 A, the volt-drop in the cable would reach 5 kV in a distance of 100 km. This means that only 5 kV would be applied to the dc/dc converter, and this would not be enough to allow it to start up!

These choices are described in the trade-off documents produced by the NEPTUNE Power Group, and can be seen on the Web site at <http://neptunepower.apl.washington.edu/>

3 POWER ENGINEERING

The basic choices having been made, the task of power engineering began. Reliability of service was paramount for two reasons. First, the cost of repair was viewed as extremely high, and second, if the power system did not work, nothing in the science node would work. Furthermore, the capital investment in the observatory was so high that reputations were at stake!

Terrestrial power provided a good model of how to design a robust power delivery system. Spacecraft also teach how to make fault-tolerant systems. A few guidelines might be:

- Use a network, with interconnections. This approach provides redundant paths, and allows a good measure of fault tolerance. It is the approach used in power transmission, but (for economic reasons) not typically in power distribution. One result is that 75% of power outages are caused by distribution system problems.
- Use protective relaying judiciously, with built-in backups. Most relaying schemes have layers of protection, and overlapping zones, so that not only are failures of the power system allowed for, but failures of the protection system itself can be tolerated.
- Design for a safe mode. Most science spacecraft are designed to allow for temporary loss of communications. This loss might be caused by a spacecraft malfunction (such as when the star-tracker used for navigation mistakenly locks on to a close piece of dust instead of distant star) or by occultation (when the spacecraft passes behind a planet). Such events should not cause loss of mission.

With these ideas in mind, the power group proposed a power delivery system as follows.

3.1 Configuration

The power delivery system would be a network (as shown in Figure 1) consisting principally of one major loop with two infeeds and two cross-plate connections. At each of the locations for science nodes, the backbone cable would be interrupted, and a circuit breaker inserted so that the cable could be sectionalized to isolate a fault. The science node at this location could be powered from either side of the circuit breaker, and would control whether the breaker was open or closed.

Since the circuit was to be energized with dc, which is hard to interrupt, the circuit breaker was planned to be a combination of solid state and semiconductor switches arranged so as to commutate the current out of the main interrupter contacts. Figure 2 shows the arrangement.

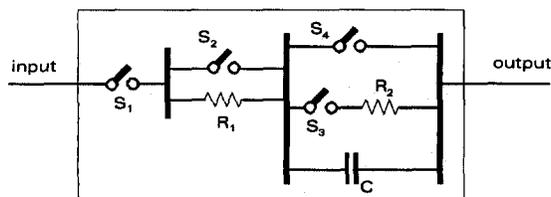


Figure 2. Circuit breaker implementation

In Figure 2, switches S_1 , S_2 and S_4 are closed when the breaker is closed. The opening sequence begins with S_4 . The line current commutates into the capacitor C , and decays exponentially to zero. After a few hundred ms, the current is small enough that S_1 can be opened, at which point the circuit has been interrupted. S_3 is closed to discharge the capacitor, ready for a repeat of the sequence. S_2 is provided to allow the resistor R_1 to act as a pre-closing resistor, charging the cable capacitance slowly, and limiting the current into a fault should there be one.

The concept of operations for the network is straightforward. The shore station energizes the first section of cable. A diode in the circuit causes the science node at the end of the cable section to power up when either of the adjacent cables is energized and communications are established with the node. Commands are sent to close the backbone circuit breaker. (All the switches in the breaker are of the normally-open type, so that power is needed to close them.) Figure 3 shows the diode/breaker configuration.

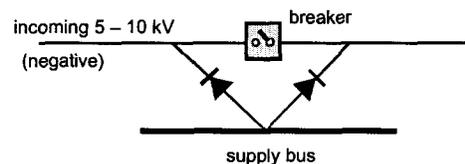


Figure 3. Diode-breaker arrangement at each cable section

With this arrangement of diodes and breakers, the supply bus is powered when one of the incoming cables is energized, even though the backbone breaker is open (the default condition). The two diodes are back-to-back as far as adjacent cables are concerned. Under the control of the node, the backbone breaker closes, and the next section of cable is energized. The sequence of events is repeated until all the backbone breakers are closed and all the science nodes energized.

Once the breakers are all closed, and the science nodes all communicating to shore, the main work of the observatory can begin.

While the observatory is in this normal condition, a number of power management activities are taking place, coordinated from shore via the communication system. For example, the currents and voltages around the network are monitored, and load-flow calculations run to examine system stability. Should the network be approaching the limit of its capability, load would be shed according to a pre-arranged priority list.

3.2 Protective Relaying

The protective relaying scheme must operate in two parts of the power subsystem. In the backbone, it is responsible for clearing cable faults. In the node, it is responsible for monitoring loads and detecting overcurrents. Of the two tasks, the protection of the backbone is the more challenging, and from the system viewpoint, the more important.

Current-differential protection is commonly applied to busbars, less commonly to transformers and lines. The principle is simple: the sum of all currents into the protected element should be zero. If measurements indicate that it is not, then there is a current somewhere that is not accounted for.

Applied to the protection of the NEPTUNE cable, the comparing of the current in and out of a line section can be done via the communication system. In NEPTUNE a wideband communication

link is accessible for such a purpose. The first line of defense for cable faults was proposed to be current differential protection, operating the adjacent backbone circuit breakers.

As a backup to the current differential protection, distance relaying was proposed. In distance protection, the voltage and current at the end of a cable are measured, and their ratio is used to estimate the distance to a fault. The assumption is that the voltage at the fault is zero, and the resistance of the cable per unit distance is constant.

Because none of the parameters used in distance relaying is known precisely, there is some uncertainty about the decision to trip when the distance to the fault approaches the limit of the protected zone. (This problem does not exist with differential protection, since measurements are made precisely at the ends of the zone – indeed the location of the current transducers defines the zone.) Therefore, further protection is usually added in case the primary distance protection misoperates.

3.3 Safe Mode

It was considered prudent to have the node power system enter a “safe mode” under certain circumstances, particularly if the communication system should fail. The safe mode can be understood by examining the startup sequence in a node in detail:

- 1) The node receives the incoming power, and the supply bus is energized via a diode.
- 2) The node startup supply turns on.
- 3) The startup supply allows the main dc/dc converter to turn on.
- 4) An interlock disconnects the startup supply once the main supply is running.
- 5) All table entries go to default values. Examples:
 - outage list resets (assumes cable is good)
 - protection settings go to default
 - list of loads served goes to zero
- 6) Essential loads are turned on. Examples:
 - some (local) elements of node protection are turned on
 - communication system is turned on
- 7) A timer is started for closing the backbone breaker. The power system waits for instructions from shore
- 8) If communications are not established and the timer “times out,” the backbone breaker is closed, energizing the cable to the next node.

The Safe Mode begins at step 5. If the sequence begins from a power up, only valid communications (step 8) can move the system from the safe mode. If the safe mode is entered by command from shore, or from some internal watchdog process, the system will still wait at step 8 for further instructions.

In the event that the communication system is not working, the system will return to the safe mode with the backbone breaker closed. This means that while data cannot be obtained from or passed through the node, the power system can still use the node as a connection. Although the differential protection is not available, operation in this mode is still considered safe as the distance protection for the zone is still active.

4 THE ORIGIN OF A NEW VERSION

In the version of events just described, closing the backbone breaker required that power be available at the node. This situation was re-examined by one of us (Lancaster), who proposed an alternative powering scheme. Instead of taking power for the

breaker from the main power supply, why not provide it with its own series-type power supply based on a zener diode? This approach had been used for decades with great success in submarine cables, where the power for the optical line amplifiers (repeaters) was normally obtained this way.

In NEPTUNE, it may be necessary to add a dummy load at each breaker location to ensure some current, but there was no reason in principle that this could not be done. The result would be a supply that was likely to be more reliable than the main dc/dc converter.

In the end, this was not exactly the scheme that the NEPTUNE power group adopted. But the suggestion did trigger a re-examination of the power system approach, and a redesign from fundamentals.

The following ideas gradually emerged:

- Repair is so expensive on the backbone system that it should be avoided if at all possible. A cable-ship repair might cost as much as \$500,000. To guard against the need for such an expense, a Mean Time Between Failures (MTBF) figure of close to 1000 years should be a target for backbone components.
- Repair may be deemed acceptable in the science node, as the economics of such repair are distorted by the availability of University-National Oceanographic Laboratory System (UNOLS) ships to the university community at essentially no cost. (Unfortunately, a UNOLS ship is not capable of performing a repair to the backbone cable.)

A complicated electronics board such as might be needed in the science node probably could not be made sufficiently reliable for backbone use within the budget of NEPTUNE. The process of *qualifying* a component or subsystem for subsea (or space) use is extremely expensive.

Our only choice was to design the backbone using parts already qualified.

5 POWER – VERSION 2

In the redesign of the power delivery system, the following considerations assumed prominence:

- As far as possible, backbone operation and node operation should be kept segregated. Node failure should not cause backbone failure. *This was the driving consideration that emerged from the review of the breaker powering options.*
- There will be cable faults, but they will be fewer in number than for a terrestrial system of the same size, and it will be acceptable to shut down the system to clear them.
- Simplicity of design should replace the complex protection and backup arrangements of the first design.
- Component life should be maximized. This means that breakers would have to be operated when there was minimal current and voltage, and in general power dissipation inside pressure cases should be kept to a minimum.
- The control for the backbone would not depend on there being a current (hence no dummy loads).
- To reduce power consumption, all switches would be latching switches.
- The diodes of the first design would be replaced by switches. (No qualified diodes were known.)

After much discussion, it became evident that the three basic guidelines adopted earlier were due for revision. The new versions might read as follows:

- Since cable faults are relatively rare, provided the backbone can be made highly reliable, there may not be much advantage to using a network with interconnections. The redundant paths of the network are mostly wasted if there are no cuts in the cable. Monte Carlo simulations of the network bore out this conclusion.
- Protective relaying judiciously applied might mean no built-in backups. The simpler the system, the less prone to failure it may be. Utility relaying schemes can be serviced at much lower cost than a similar scheme under 4000 m of water.
- If the system can operate in a safe mode, with the loss of communications, it should be operated this way all the time. Power system capability will not then depend on the availability of the communication system.

5.1 Configuration

Without affecting the top-level topology, the considerations above led us to separate the hardware for operating the backbone from the hardware of the node. The arrangement of the building blocks for connecting the node to the backbone is shown in Figure 4, which contrasts the original design with what came to be known as Version 2.

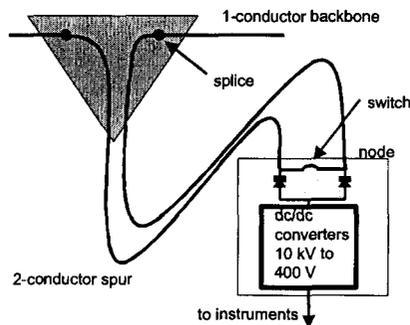


Figure 4a. Original layout of node connection

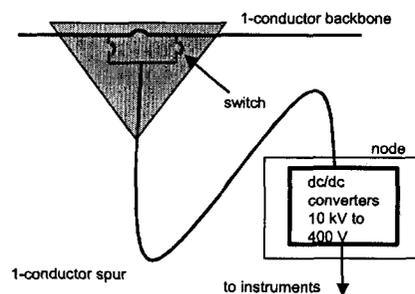


Figure 4b. Revised layout of node connection

In Version 2 the switches in the backbone are not controlled from the node, they are controlled autonomously within the branching unit. The controls there can be powered from a series (zener) supply or a simple parallel supply. There are in fact four switches, each controller having access to only two. This arrangement

overcomes the need to transfer information or commands across a possible 10-kV potential (see Figure 5).

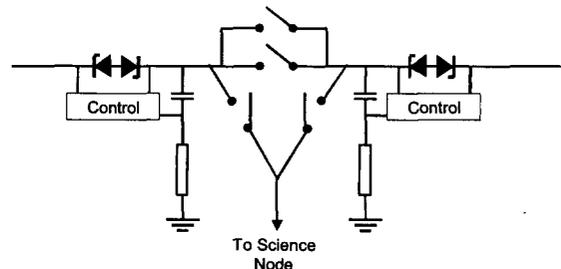


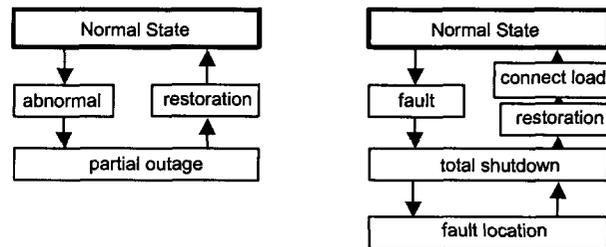
Figure 5. Control arrangement in Branching Unit

5.2 Protection: Modes of Operation

Instead of the customary power system operational philosophy, four modes of operation are identified.

- 1) Fault locating
- 2) Fault clearing
- 3) Node connecting
- 4) Normal

The differences that these modes represent compared to a terrestrial power system are evident when the way the two operate is shown diagrammatically, as in Figure 6.



(a) Conventional power system (b) NEPTUNE Version 2

Figure 6. system operating modes compared

In the conventional power system, normal operation is interrupted by a fault, and the system enters a state that power engineers call abnormal, and space engineers call off-nominal. The protection system operates rapidly to isolate the fault, often opening the circuit breakers *before* the fault current reaches its maximum value. The location of the fault is simultaneously reported to an operations center of some kind, and a repair crew dispatched. Transmission-system faults are routinely cleared without loss of service to any loads. Restoration is then of equipment, not load. Transmission system faults rarely result in loss of load.

NEPTUNE, by contrast, completely shuts down in response to a fault. All load is lost, and all the converters shut down. This comes about automatically because the system is *weak*, having a current limit on the shore supply in order to protect the cable from damage.

Following the shut-down, the system is put into a fault locating mode. All backbone breakers are closed (fault or no), and all load switches (to the science nodes) are opened. This action takes place

in the branching unit in response to the detection of a positive voltage on the backbone. The voltage is set low, however (around 500 V) and the current is limited to a value of perhaps 0.5 A, so the duty on the switches is not severe.

In this mode, with all the load disconnected, the resistance of the cable can be measured from shore, and an estimate made of the location of the fault.

Next, the voltage is reversed. The BU controller, sensing this reversal, acts so as to clear the fault. (Note that this is an entirely different operation from locating it.) In this mode, the voltage at the BU is used as a surrogate of the distance from the fault, and time is used to coordinate the relaying action.

When the voltage goes negative, a current will start to flow if there is a fault. A timer is started in each BU at this instant. The value on the timer is set by the voltage: the greater the voltage, the longer the time. Since the current is flowing into the fault, nodes that are closer to the fault will have lower values set in their timers, and will clear first. The fault is isolated by the closest breakers. (Almost accidentally, this approach gives a measure of backup. Should a breaker fail to operate, the next nearest will automatically operate in its place.)

When the fault is cleared, the current goes to zero, and the BU controller responds by closing the switches to the science node.

Once the switches are closed, the voltage on the backbone can be increased to its nominal value, and normal operation resumes.

6 DISCUSSION

The terrestrial power system is an example of a highly reliable system. While statistics vary from place to place, it would be an unusual first world customer who experienced more than 10 minutes down-time per year. That represents a non-availability of power of 20 parts per million. This record is achieved through the appropriate use of redundancy (from spinning reserve, and interconnected delivery networks) and the careful use of protection.

However, the terrestrial power system is also *serviced*, so this record of availability is partly the result of maintenance. It should also be noted that the cost of maintenance is not extraordinarily high: most locations on most power systems can be reached by truck, so that crews and supplies can gain easy access.

The same will not be true with the NEPTUNE observatory. Instead of borrowing from terrestrial power for the design and repair paradigm, it makes sense to borrow from the world of trans-ocean cables. When their guiding principles are adopted, the solutions are different.

We now calculate that the backbone reliability will be so high that it will not be necessary to use a network to provide redundant paths. A simple point-to-point connection will suffice. This might lead to a savings in cable costs. It certainly leads to simplification in the locating of faults.

Another difference between terrestrial power and NEPTUNE is that when a fault does occur, the entire system is shut down while it is dealt with. Since revenue is not lost, there is no economic impact. (Note that this is a question of economics, not of policy. To build a system that did not require a current-limited source on shore would require a much bigger cable than the one planned for NEPTUNE, at much higher cost.)

However, calculations show that while outage time directly due to cable faults will likely be at a slightly *lower* level than on land (there may be 10 faults in the 30-year life of NEPTUNE, each

requiring only 5 or 10 minutes to locate and recover from), total outage time will be greater. This is because the entire delivery system will have to be shut down for as many as 3 days while each of those faults is repaired. (But this is not a result of using Version 2: it would be the case however the power delivery system were operated.)

7 POWER MANAGEMENT

Both Versions of the NEPTUNE power system were designed to provide stable and secure power supply for the science nodes. To achieve this, the state of the system must be monitored, and corrective action taken if necessary. For example, as the load varies, the backbone voltages and currents will fluctuate. The voltage profile of the delivery network will have to be maintained within allowable limits. If the voltage somewhere becomes too low, the system could enter what is called voltage collapse, and shut itself down 4). Heavy load is the usual cause of voltage profile problems: the solutions can include adjustment of the voltages at the shore stations, or load shedding. These functions are performed as part of the Power Monitoring and Control System, PMACS, operating in the science node and on shore 5).

PMACS is shown functionally in Figure 7. Through the communication system, breaker status data and analog measurements (voltage, current and power measurements) are acquired and archived. Operating constraints (voltage, current and power limits) are checked, and adjustments (usually load shedding) made if necessary.

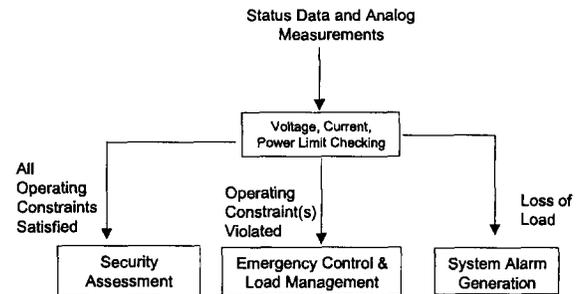


Figure 7. PMACS functions

One of the principal requirements of the PMACS system is that it must be robust against missing data. This requirement is normally met by an operation called *state estimation*, in which the available measured data are fed into a model of the system, and the system equations solved so as to minimize some cost function, usually the total least-squared difference between the measurements and the model. The state estimation problem is different between Version 1 and Version 2 because in Version 2 a good deal of the system is not observable (or controllable, for that matter) in the controls sense.

For example, since the Tee point on the backbone is located inside the autonomously-controlled Branching Unit, neither the voltage at the Tee nor the status of the breaker is known to PMACS. The voltage can, of course, be estimated, and from such estimates, breaker status can be inferred. However, measurement errors accumulate in the estimation, and the problem is much more challenging in Version 2 than in Version 1, where practically any parameter could be directly measured. The state estimation problem is addressed in a paper to be presented at the IEEE PES T&D Conference later this year 6).

8 CONCLUSIONS

In a re-examination of the design of the power delivery system for the planned NEPTUNE observatory, the designers were driven to consider reliability of supply as paramount. They borrowed from common practice in power systems, but even more from the principles and practices of subsea telecommunications. Two cultures, each with a commendable record of reliability, were merged, and the best of each used for NEPTUNE.

As a result, the NEPTUNE power system will combine several concepts novel to power systems in its design and its operation. Overall, it is expected to provide a highly reliable infrastructure for the science it supports.

ACKNOWLEDGMENTS

The authors would like to acknowledge the contributions made by our colleagues in the NEPTUNE Power Group in many useful discussions. Support from the National Ocean Partnership Program (Grant # N00014-99-10129), the National Science Foundation (Grant OCE 0116750), and the authors' institutions is gratefully acknowledged.

REFERENCES

- 1) J. Delaney, G. R. Heath, A. Chave, H. Kirkham, B. Howe, W. Wilcock, P. Beauchamp, and A. Maffei,: NEPTUNE Real-Time, Long-Term Ocean and Earth Studies at the Scale of a Tectonic Plate, *Proc. Oceans 2001, MTS/IEEE Conference and Exhibition*, Vol. 3, 2001, pp. 1366-1373.
- 2) B. Howe, H. Kirkham, and V. Vorperian,: Power System Considerations for Undersea Observatories, *IEEE J. Oceans Eng.*, Vol. 27, No. 2, April 2002, pp. 267-274.
- 3) B. Howe, H. Kirkham, V. Vorperian, and P. Bowerman,: The Design of the NEPTUNE Power System, *Proc. Oceans, 2001, MTS/IEEE Conference and Exhibition*, Vol. 3 , 2001, pp. 1374-1380.
- 4) K. T. Vu, C. C. Liu, C. W. Taylor, and K. M. Jimma,: Voltage Instability: Mechanisms and Control Strategies, *Proceedings of the IEEE*, Volume 83, Issue 11, Nov. 1995, pp. 1442-1455.
- 5) K. Schneider, C-C. Liu, T. McGinnis, B. Howe, H. Kirkham,: Real-Time Control and Protection of the NEPTUNE Power System, *Proc. Oceans, 2002, MTS/IEEE Conference and Exhibition*, Vol. 2, 2002, pp. 1799-1805.
- 6) Liu, C-C., Schneider, K., Kirkham, H., Howe, B.M.: State Estimation for the NEPTUNE Power System *To be presented at the IEEE Power Engineering Society Transmission and Distribution Conference*, Dallas TX, 7-12 September 2003.