The Design of the NEPTUNE Power System

Harold Kirkham[a] Bruce M. Howe[b] Vatché Vorpérian[a] Paul Bowerman[a]

[a]Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109

[b]Applied Physics Laboratory
University of Washington
1014 NE 40th Street
Seattle, WA 98105

***Abstract*–The proposed NEPTUNE observatory will include about 30 locations on the Juan de Fuca plate where scientific instruments can be connected for communication and power. The NEPTUNE power system is required to make available at each location the largest amount of power possible, using conventional submarine telecommunications cable. The power delivery system is based on the use of a standard cable, but it is used in an interconnected network in order to maximize both reliability and power level. The cable will be energized with medium voltage ~10 kV dc and have parallel loads, a combination that has never been built before as an interconnected network. During normal operation, it is calculated that a power level of over 5 kW can be delivered to each of the 30 nodes. Should it be needed, as much a 40 kW can be delivered to each of 3 nodes on the far west of the network, 500 km from shore, provided the power on all the other nodes is reduced to 1 kW. These power levels and distances are considerably greater than has been achieved in previous undersea observatories. The design of the sub-sea node is based on the use of grouped switching mode power supplies with series-connected inputs and parallel connected low voltage outputs. The ocean is used as a return path. Consideration of the reliability of the system plays an important role in the design of the power system. A scheme of protective relaying will enable the delivery system to continue operation even with faults in part of the network.**

I. INTRODUCTION

The proposed NEPTUNE observatory [1, 2] will include about 30 locations under the Pacific on the Juan de Fuca plate, west of the U.S. and Canadian west coasts (Fig. 1). At each location, scientific instruments can be connected to fixed nodes for communication and power.

The NEPTUNE power system is required to make available at each location the largest amount of power possible, using, for cost reasons, conventional submarine telecommunications cable.

The level of power needed is much more than for typical submarine telecommunications repeaters, which operate at a fixed current of about 1 A. It is more, too, than required for science instruments only. Lights for television, energy for experiments involving heat transfer, and energy for charging batteries are also needed.

The design that has evolved to meet these needs is novel in several respects. The power delivery system will be operated as an interconnected network in order to maximize both reliability and power level. The cable will be energized with dc. Although the concept has been discussed for many years, dc systems have never been operated before as interconnected networks.

To maximize the power capability, the network will be a parallel scheme, with the ocean providing the return path for the current. This means that at each node, the power supply has to reduce the incoming supply from around 10 kV to a more user-friendly lower voltage of 48–300 V.

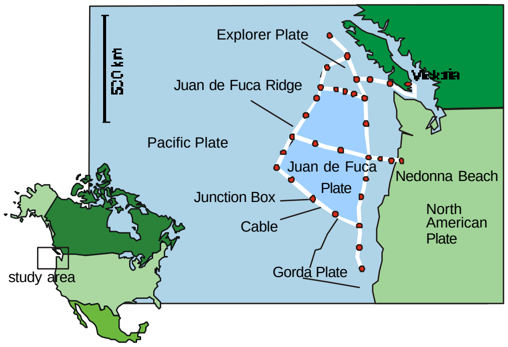

Fig. 1. The proposed NEPTUNE observing system.

In spite of the novelty of the approach, reliability will be of the same order as the conventional submarine telecommunications cable system. A lifetime of 30 years is the goal.

This paper describes the trade-offs that have been made to make the basic design choices, and the expected system performance.

II. TRADE-OFFS

Power from the shore is inserted into the network at medium voltage (MV; defined by IEEE as 2.4–72.5 kV). As a practical matter, the cable insulation on a typical telecommunication cable is rated for around 10 kV, so this value is assumed as an upper limit for the NEPTUNE backbone. The resistance of a typical cable is so high (about 1 S/km) that over the distances involved in NEPTUNE, the current is limited by the cable-volt-drop to about 10 A.

Given these constraints, the basic trade-offs are as follows:

- ac or dc on the cable?
- interconnected or radial network?
- series or parallel connected loads?

These questions are addressed in detail elsewhere [3]: they will be summarized here.

The question of ac or dc on the cable reduces to a relatively simple cost calculation. If ac were used (at least at 50 or 60 Hz), the charging current of the cable capacitance would be so large as to *require* compensation by shunt inductors. An order of magnitude estimate of the cost makes this alternative much more costly than the use of dc.

The use of very low frequency ac is somewhat more attractive. Supply at (say) 0.1 Hz would have the advantage of low charging current, and the further advantage of avoiding the various cable insulation problems that affect dc cables. However, at the relatively low voltages expected in this cable, these problems are solvable. Further, the use of very low frequency ac would require additional complexity in the nodes, as transformers for such frequencies are not feasible. On balance, it was felt that the complexity of ac/dc conversion could not be justified.

The decision whether the network should be operated radially or interconnected has significant impact on the reliability (availability) of the delivery system. In terrestrial power delivery systems, only the distribution system is operated radially; the remainder is interconnected. It is this interconnection that has improved the reliability to the point that the great majority of power outages are due to distribution system problems.

While it is true that operation of terrestrial power systems closer to their maximum power limit has resulted in the few outages having a more widespread effect, analysis also shows that deliberate load shedding could be employed to avoid cascading [4]. This is a possibility that will be borne in mind for NEPTUNE. Some loads are readily deferred (battery charging, for example) and others may be considered of low priority.

In order to be both series-connected and a network, a new scheme would have to be developed in order to split the power at a branch. This hypothetical new block would have to recognize sources and loads (the answer might change as loads changed over the lifetime of the project) and then deliver the same current to each of the outgoing cables, at reduced voltage. While such a device is feasible, the network that used it would still suffer the inefficiencies of a series system, and the power level that could be transmitted would be so much lower than a parallel scheme could transmit, that a parallel scheme was chosen early in the discussion of trade-offs for NEPTUNE.

III. SYSTEM DESCRIPTION

The NEPTUNE system one-line diagram is shown below (Fig. 2). The arrows on the one-line diagram show the direction of power flow, assuming all loads are equal.

A. Switching and Sectionalizing

The multiple possibilities for sectionalizing to isolate faulted nodes and faulted cable sections are also evident. In order to do this, some means of switching the faulted section must be provided. At this time, the decision has not been made whether circuit breakers or switches will be used. Circuit breakers have to interrupt fault current: in the case of a controlled system such as this, the fault current will be about the same as load current. Nevertheless, there may be some reliability advantages to switching only when the system is de-energized.

If this option is chosen, it will be necessary to implement a communication system that can operate without the nodes being active. A field-bus system operating over a king-wire is being considered. While slow (we estimate at most a few kbits/second), such a scheme would constitute a useful means of access into the power network, and possibly the communication scheme.

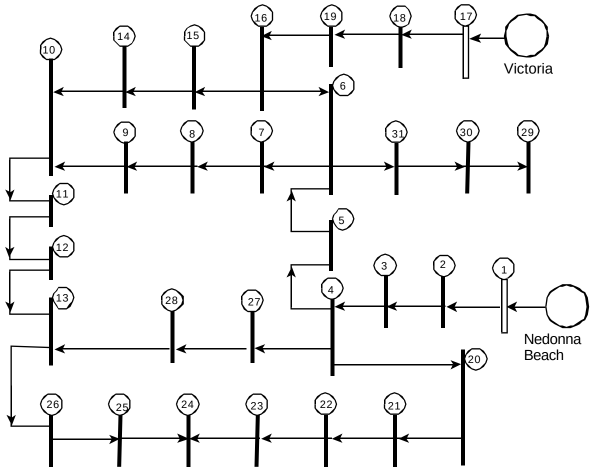

Fig. 2. One-line diagram of NEPTUNE power system

Each bus in the one-line diagram represents a switching arrangement. For the simple cable-in/cable-out locations (most nodes), a pair of diodes can be used to replace a switch (Fig. 3). The advantage is the reduced number of moving or controlled parts.

Opening the single breaker in Fig. 3a has the same effect as opening one of the breakers in Fig. 3b: the two parts of the cable are isolated from one another, but the node supply bus can take power from whichever side is still connected. The double-bus double-breaker configuration in Fig. 3c has the advantage that it permits full operation of the junction even if one of the breakers fails, whether it fails open or closed. In a system designed for a long life, and with restricted access for repair, such fault tolerance may be important. The arrangement can be extended to 4 or more lines at a node.

B. Protection scheme

Complementing these various switching schemes will be a protection system based (somewhat loosely) on utility power system practice. The purpose of the protection scheme is to detect faults in the network, and minimize their effect by operating the appropriate switches.

There are several ways that faults can be detected. The simplest is based on the concept of *overcurrent*: a circuit that normally handles a current of (say) 10 A can be assumed to be faulted if the current suddenly becomes 20 A. A fuse is an example of an overcurrent protection device. (Fuses are ruled out here because they are not resettable, do not work well with direct current, and are not likely to discriminate adequately.) A problem with this approach is that it does not distinguish faults that are close from faults that are remote. More than a minimum amount of the system may thus be disconnected.

By using information from several parts of the power system, the protection scheme can do a better job of *discrimination* between faults.

For example, a *distance relay* gives better discrimination.

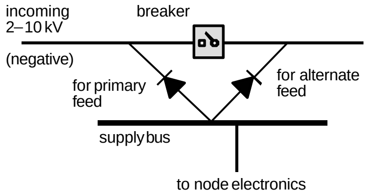

a) using back to back diodes

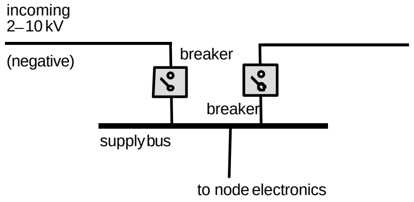

b) using 2 breakers

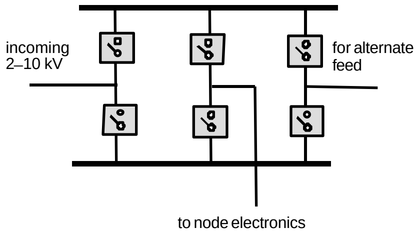

c) double-bus double-breaker

Fig. 3. Alternate ways to implement sectionalizing at a simple node

In this device, a computation based on the measurement of both voltage and current is used to estimate the fault location relative to the point of measurement.

An even better protection scheme, and one that could be used in NEPTUNE, is the *differential* scheme. In this approach, the current at one end of the section is compared with the current at the other end. If there is a difference, there must be a fault between the two points of measurement. This kind of protection relies on the existence of a fairly fast communication connection between the two ends. In NEPTUNE, this connection can be assumed to exist in the form of the NEPTUNE communication system. If the fault removes the "best" NEPTUNE connection, land-based portions of the internet could be used to complete a communications circuit. A differential protection scheme is a viable possibility.

A generic block diagram of such a protection scheme is shown in Fig. 4. Monitored parameters are used continuously to calculate the quantities on which the action decision will be based. For example, this may mean computing the location of the fault. Based on this calculation, the relay will decide whether to trip a circuit or not.

In practice, because a fault might prevent or delay communications, it is planned to adopt the utility practice of having *layers* of protection. NEPTUNE's protection scheme could include differential, distance, and overcurrent relaying. This way, the first (and best) line of defense would be the differential scheme; if the fault causes loss of communication, the distance relaying would operate; if the distance relaying failed, the overcurrent system could save the day. Because the loads are expected to be far more deterministic than those of a typical utility, overcurrent protection levels can be set quite close to the normal load values.

The protection function will be built into the junction boxes. Some aspects cannot easily be changed after installation, for example, a method to isolate faulted components must be provided for the system to function. On the other hand, some of the settings can be changed, for example, by means of updatable data tables.

The hardware may be quite complex. The protective relaying system must use independent sensors wherever possible, and be capable of dealing with failures inside the protection scheme itself (such as failed circuit breakers).

C. Power Management System

While the power system must operate autonomously (that is, without operator intervention), it must be capable of furnishing information to the system operator that could be used to modify the way operation is carried out. Some level of supervisory control, supported by an internal data acquisition system, is therefore required.

The management system will likely monitor the same parameters monitored by the metering scheme, and may share the primary transducers. Most probably it would store the data for off-line analysis. In addition, it would scan the data for the on-line generation of alarms.

A *maintenance alarm* would be generated when a condition existed that was not serious enough to warrant action by the protection system, but was serious enough to justify operator intervention. An excess power consumption by a particular load might be an example, if the power level did not result in an overload and trip.

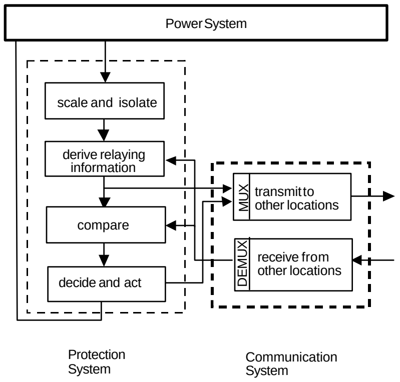

Fig. 4. Generic protection system functional block diagram

An *operation alarm* would be generated whenever a more serious condition was encountered.

The design and implementation of the power management system must be integrated with the protection scheme. Resources may be shared with this, or with the metering scheme.

D. Deliverable Power

To better understand the possible parameter space for the NEPTUNE power system design, we have simulated various configurations of the system, varying nominal shore voltage, cable resistance, and loads at each node.

The basic laws that govern the steady-state flow of power within a dc electrical network are Kirchoff's current and voltage laws and Ohm's law. Here, we have an additional constraint: the power, $P = VI$, withdrawn from each node is prescribed. To handle the nonlinearity that results in the equations, the problem is solved by iteration. The initial voltage used at all the nodes is the shore station voltage, and all the initial currents are zero.

The first set of simulations assumed the two shore stations (Nedonna Beach and Victoria) operating at 10 kV and using 1 Ω/km cable. As the load power is increased from 1 kW to 6.7 kW at each node, the efficiency falls from 97% to 63% (Fig. 5). To some extent, efficiency can be used as a surrogate for stability. As was shown in an earlier paper [3], any power system is capable of becoming unstable as its maximum power transfer capability is approached. The usual way of detecting this approaching instability is to examine the eigenvalues of the system Jacobian. An alternative is to look at the convergence of the load software. The idea of using efficiency is only workable here because it can readily be calculated during the simulation. The 6.7 kW value is close to the maximum possible power before the system (and code) go unstable. The minimum voltage occurs in the southwest leg.

If the shore voltage is increased to 15 kV, the efficiency increases from 63% to 90% for the same useful power (6.7 kW), showing that this level of power is not close to the system limit. In fact, the maximum possible power in this case is about 15.2 kW at each node.

If the cable resistance is then lowered to 0.7 Ω/km the efficiency increases from 60% to 82%. The maximum possible power per node is 21 kW, for the "best" cable 0.7 Ω/km and 15 kV. As for the 10-kV cases, the minimum voltage in the system occurs in the southwest leg. If we revert to 10 kV at the shore stations, but keep the 0.7 Ω/km value for resistance, the maximum possible power is 9.6 kW per node.

Two cases were considered where three 1000-km spur cables are added to simulate possible cables to Ocean Weather Station PAPA to the northwest, the deep Northeast Pacific (dubbed UNCLE), and south to California. In these two cases, the loads are not evenly distributed; at eight nodes that are considered by some to be scientifically more interesting, the load is higher. The eight nodes are: Juan de Fuca Strait, Endeavour Ridge, Axial Volcano, Hydrate Ridge, and at the end of each of the three long spurs. In the first of these cases, 10-kW loads are required at the selected 8 nodes, and 5 kW at the rest. This produces a stable solution with an efficiency of 74%. In the second case, 14 kW are required at the selected 8 nodes, and 2 kW at the rest. This produces a stable solution with efficiency of 71%.

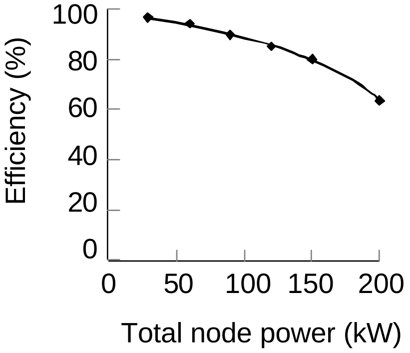

Fig. 5. Efficiency versus total useful power

Reverting to the standard 30-node configuration (Fig. 1), 0.7 Ω/km, and 10 kV, we wished to know how much power could be delivered to a small number of nodes located on the ridge. We found that it was possible to deliver 40 kW to each of the three nodes at the southern end of the Juan de Fuca Ridge (Axial Volcano and the two adjacent nodes to the south), while at the same time delivering 1 kW to all the other nodes. In this case, the efficiency is 66%, suggesting that a slightly higher power level yet is possible. This high level of power (40 kW) has implications for the node design, of course. Should all nodes be rated for this level of power? Should any?

We also explored the possibility of only one active shore station, still using 0.7 Ω/km and 10 kV. If only Nedonna Beach were active, 5.6 kW maximum can be delivered to all the nodes. If only Victoria were active, 3.9 kW maximum could be delivered to all the nodes.

In both cases, the minimum voltage is at the node farthest from the shore station (i.e., the ends of the northern and southern spurs, respectively). Repeating the scenario in the preceding paragraph, one finds that the maximum power at the distant three nodes is 30 kW if only Nedonna Beach is active, and 21 kW if only Victoria is active.

These results show that it is fair to regard each shore station power supply as a backup for the other. Even with only one shore station supplying power, a significant amount of power can be distributed to the junction boxes around the seafloor using cable with readily obtainable parameters (i.e., cable resistance and operating voltage). These power levels and distances are considerably greater than has been achieved in previous undersea observatories. Because the present topology is reasonably robust against the loss of one shore station, it may be that it is unnecessary to provide an uninterruptible power supply (UPS) at either station. This matter will receive further consideration.

While the focus of this brief discussion has been on the maximum possible delivered power, one would clearly not run an actual system near this operating point, because of the associated voltage instability. The power levels given here are, however, considered feasible.

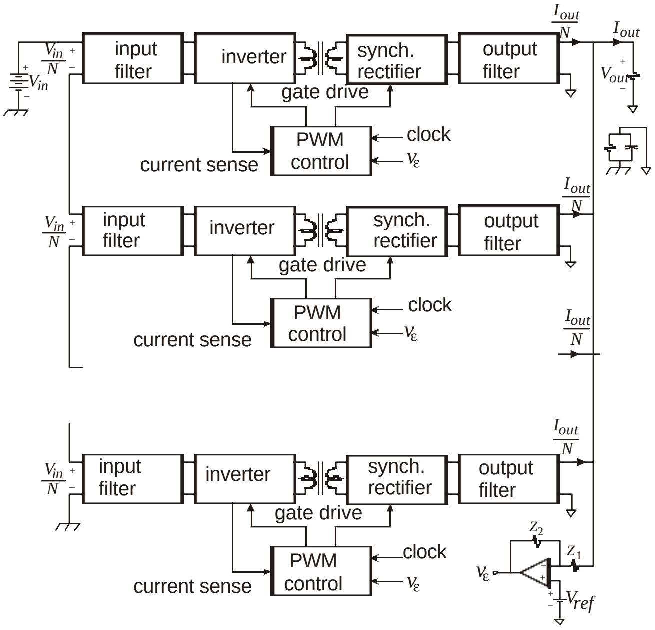

Fig. 6. Block diagram of converter

E. Converter

The dc/dc converter at each node uses high-frequency (50–100 kHz), pulse-width-modulated (PWM) technology to convert the 5-to-10 kV bus voltage to a regulated low output voltage suitable for many commercially available, low-voltage dc/dc converters. Each converter is capable of delivering up to 10 kW at an output voltage of 200 V or 48 V. A block diagram of the converter is shown in Fig. 6.

Owing to the current feedback loop from each inverter stage, the large input voltage is shared equally by all the input inverter stages, under static and dynamic conditions (with a bandwidth of about 4 kHz), resulting in only low switching voltage stresses on the components. The current feedback loop also ensures that the load current is equally shared by all the secondary synchronous rectifier circuits, which are connected in parallel. In this design, the following are the only components that experience a 10-kV dc voltage:

1. The isolation transformers. Some of these will have a working 10-kV dc voltage between the primary and secondary windings. For testing purposes a maximum of 20 kV dc will be applied. The switching waveforms across the primary and secondary windings are no more than 200 V. These features allow for standard and simple construction methods of the isolation transformers.
2. The circuits for feedback, clock, and drive signal for the power MOSFETs. These can be achieved either by optical or magnetic (transformer) means.

All the stages are synchronized to the same clock and receive the same error feedback signal from the output voltage. The current feedback loop allows for minor adjustments for each stage in the duty cycle to ensure equal sharing of voltages and currents among all the stages, regardless of variations in component tolerances.

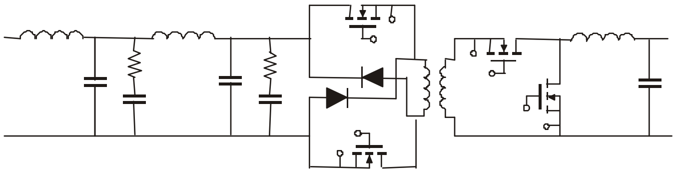

Fig. 7. One converter stage

The converter topology used for each stage is the buck-derived, two-switch forward converter. This has the lowest stresses among all isolated converter topologies and, hence, high reliability (Fig. 7). Each stage is designed to operate from an input voltage range of 100–200V. Below an input voltage of 100 V, a stage drops out. With 50 such input stages, the dc/dc converter operates from an input voltage in the range 5–10 kV and drops out below 5 kV.

The purpose of the input filter is to attenuate the high-frequency switching ripple. The filters used provide an attenuation of 100 dB resulting in an input ripple current of 30 μA.

To validate the design, a Pspice simulation was performed using four stages. On the input side, four stages were connected in series and on the output side, two series-connected stages were connected in parallel. The input voltage was set at 800 V and the output voltage was set at 100 V and 2 A. Between the voltage source and the converter, a 100-km transmission line was used with the following characteristics: 1 Ω/km, 0.2 μF/km and 1 mH/km. The component values for the inductors, capacitors, and resistors were varied within a tolerance of 10%. The waveforms obtained show that the input voltages to all the stages are nearly identical (Fig. 8). The currents in the output inductors are also identical.

IV. RELIABILITY

One of the advantages (and requirements) of a cabled observatory such as NEPTUNE is low maintenance over its planned life. A performance goal was set of no more than one node repair needed every 2 years at any of the 30 nodes over the planned 30 year life. To achieve this goal, consideration of the reliability of the system played an important role in the design of the power system. A scheme of protective relaying will enable the delivery system to continue operation even with faults in part of the network. Within each node, standby redundant converters will be used, so that with readily achievable MTBF figures, the overall system goal is met.

A separate, but integrally related, subject is the availability of the overall NEPTUNE system. Repair missions will require some lead time to acquire the resources, and a repair mission will only be possible during an abbreviated portion of the year (from May through September). Hence, even when failures occur and repairs are required, it may take more than 6 months to accomplish. The NEPTUNE system must be robust to these failures. Consequently, it is designed so that most failures that occur will only result in loss of redundancy. Most single points of failure that are in the system can be isolated and result in the

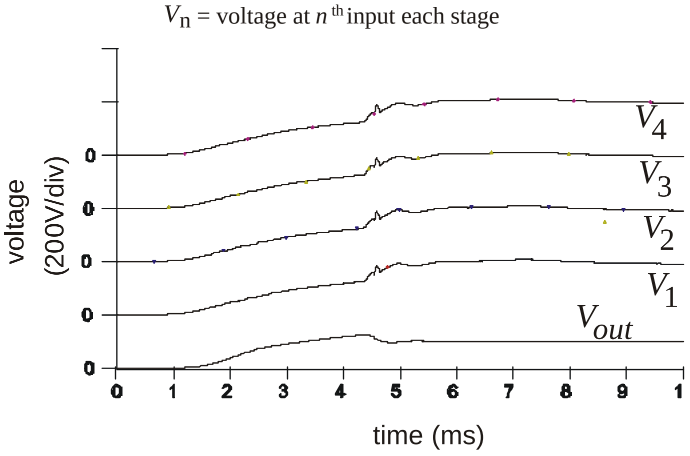

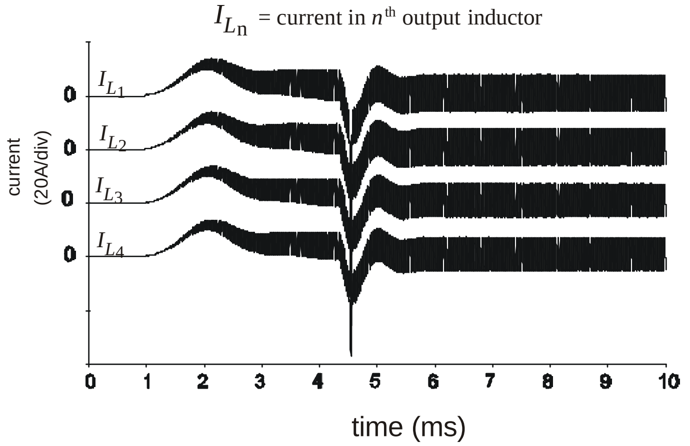

Fig. 8. Pspice simulation

loss of power to one node only so that the scientific objectives are not compromised.

The scheme in Fig. 3 provides for breakers at either end of a cable to allow for isolation of the cable in the event of a break in the cable and consequent short to the sea return. A single break in the cable that does not short to the ocean or an open failure of a breaker will not cause any science loss since the power supplied from a single shore station is sufficient to provide power to the entire undersea configuration. Likewise, a short to sea within a node power system can be isolated by opening the breakers at the nodes on either side.

There will likely be two power converters in parallel in a node for redundancy. They can be isolated by opening two breakers if a converter happens to fail. If, upon analysis, the two-converter design proves not sufficiently reliable, then a third converter can be added in parallel to provide the reliability needed. The design would contain an additional on-line spare. The startup supply element should be simple enough to be highly reliable, but could also be redundant, one in series with each converter. In the current design, the only single string elements of a power node are the primary and secondary supply busses.

Given the high cost of repairing the system, both in terms of budget and time, there is a strong impetus to make the system robust and reliable. As noted above, additional redundancy may be needed to meet the reliability goal or may be determined to be cost effective when traded against the cost of repair. Other options are also available to provide assurance. These include:

- Assuring that the electronics designs are not using parts close to or above their maximum ratings
- Using higher quality (e.g., military or space qualified) parts
- Maintaining the temperature at low levels (which should be fairly easy with a convenient heat sink via the ocean)
- Performing worst case analysis on the circuits to assure that parametric shifts over the life of parts does not result in out-of-specification performance.

From the network point of view, reliability is gained by operating the power delivery scheme as a network. A scheme of protective relaying will enable the delivery system to continue operation even with faults in part of the network.

A Monte Carlo simulation of the 32-node network (Fig. 1) was performed to obtain initial estimates of expected system availability, given estimates of mean-time-between-failures (MTBFs) of cable sections and node breakers. We assume the reliability of a breaker is an exponential based on 1,000,000 hours MTBF and each connecting cable is exponential based on 10,000,000 hours MTBF, and each repair takes 3 months. The two shore station nodes are assumed perfect, and all other aspects of the nodes are ignored. A system lifetime of 30 years is used. One thousand realizations were performed with results shown in Fig. 9.

For approximately 85% of the realizations of the possible future states of the system, the availability is greater than 0.9, i.e., during the 30-year lifetime, at most 3 years are spent in repair. With the system under repair due to component failure, availability is counted as zero even though most of the nodes are functioning. Most of the time, at least 31 nodes are available and

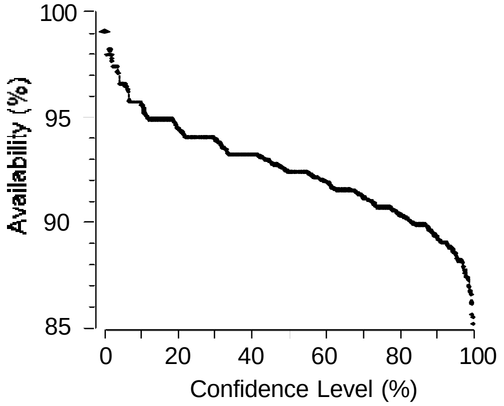

Fig. 9. NEPTUNE power system availability versus one-sided confidence limit for given availability

functioning, occasionally 30, 29, or 28, and rarely fewer than 28. For example, in 95% of the cases, the availability was greater than 0.967, i.e., only one node down.

V. CONCLUDING REMARKS

Some important decisions have been made – the NEPTUNE power system will be dc, will operate as a network, will have parallel loads – and are documented elsewhere. Some crucial decisions remain before the design work can continue very far. For example, what voltage levels must be available on the output of the dc/dc converter? Will there be a king-wire communication system? Will the double-bus double-breaker approach be used even at a simple node? Some of these questions have been discussed in this paper – their resolution will be documented in the growing library of documentation that NEPTUNE is accumulating.

A good deal of design work remains. Some of this work is at the system level, and some at the subsystem and even the component level. The recent addition of Drs Chen-Ching Liu and Mohamed El-Sharkawi at UW, Tim McGinnis at APL, and of George Fox at JPL increases the power group's strengths across this spectrum. We feel we have assembled a team with wide-ranging expertise, and we are moving forward with growing momentum.

For the NEPTUNE project, the power system is in many respects the pacing item. The task of designing the dc/dc converter is in hand and we plan to have a prototype within a year. Other elements of the work, for example, the operations software and some elements of the protection system, may not be ready for a year after that. The challenge of meeting the goals will, we anticipate, be rewarding. That we are contributing to an important new facility is even more so.

Acknowledgments

The National Ocean Partnership Program (NOPP) supported much of this work. Part of the research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Reference herein to any specific commercial product, process or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

References

1. NEPTUNE Phase 1 Partners (University of Washington, Woods Hole Oceanographic Institution, Jet Propulsion Laboratory, Pacific Marine Environmental Laboratory), *Real-time, Long-term Ocean and Earth Studies at the Scale of a Tectonic Plate: NEPTUNE Feasibility Study* (prepared for the National Oceanographic Partnership Program), University of Washington, Seattle, 2000.
2. Delaney, J. R., G. R. Heath, A. D. Chave, B. M. Howe, and H. Kirkham, NEPTUNE: Real-time ocean and earth sciences at the scale of a tectonic plate, *Oceanography*, vol. 13, pp. 71–83, 2000. Also available at http://www.neptune.washington.edu
3. Howe, B. M., H. Kirkham, and V. Vorpérian, Power system considerations for undersea observatories, *IEEE J. Ocean. Eng.*, submitted, 2000.
4. Amin, M., Toward self-healing energy infrastructure systems, *IEEE Comput. Appl. Power*, vol. 14, pp. 20–28, 2001.