Development of a Power System for Cabled Ocean Observatories

Bruce M. Howe[1], Harold Kirkham[2], Vatché Vorpérian[2], Tim McGinnis[1], Chen-Ching Liu[3], Mohamed El-Sharkawi[3], Kevin Schneider[3], Aditya Uphadye[3], Shalini Gupta[3]

1–University of Washington, Applied Physics Laboratory and School of Oceanography
2–Jet Propulsion Lab, California Institute of Technology
3–University of Washington, College of Electrical Engineering

Abstract

Cabled Ocean Observatories offer the potential to delivery unprecedented amounts of power to remote instruments and sensors. The availability of sufficient power will enable new instrumentation and methods. Here we describe the present NEPTUNE power system design which will be capable of delivering an average of approximately 4 kW or a maximum of 10 kW to over 40 seafloor node locations spread over approximately 100,000 sq nm of seafloor. The system will have a backbone of 3500 km of standard seafloor telecommunications cable connecting the nodes in a mesh topology. The network will have 10 kVdc parallel feed, distributed stochastic load and constant voltage output. A network of secondary extension cables will be developed that will allow the network to be extended up to 100 km from the backbone. The backbone cable has a single power conductor so a seawater ground return will be used. High availability and reliability over the 30 year life of the system is an important consideration in design and construction of the system. It is anticipated that faults will occur in the node electronics, cables, etc., so a protection system is being incorporated to allow faulted sections to be isolated and utilize the meshed topology to minimize impact on the rest of the system.

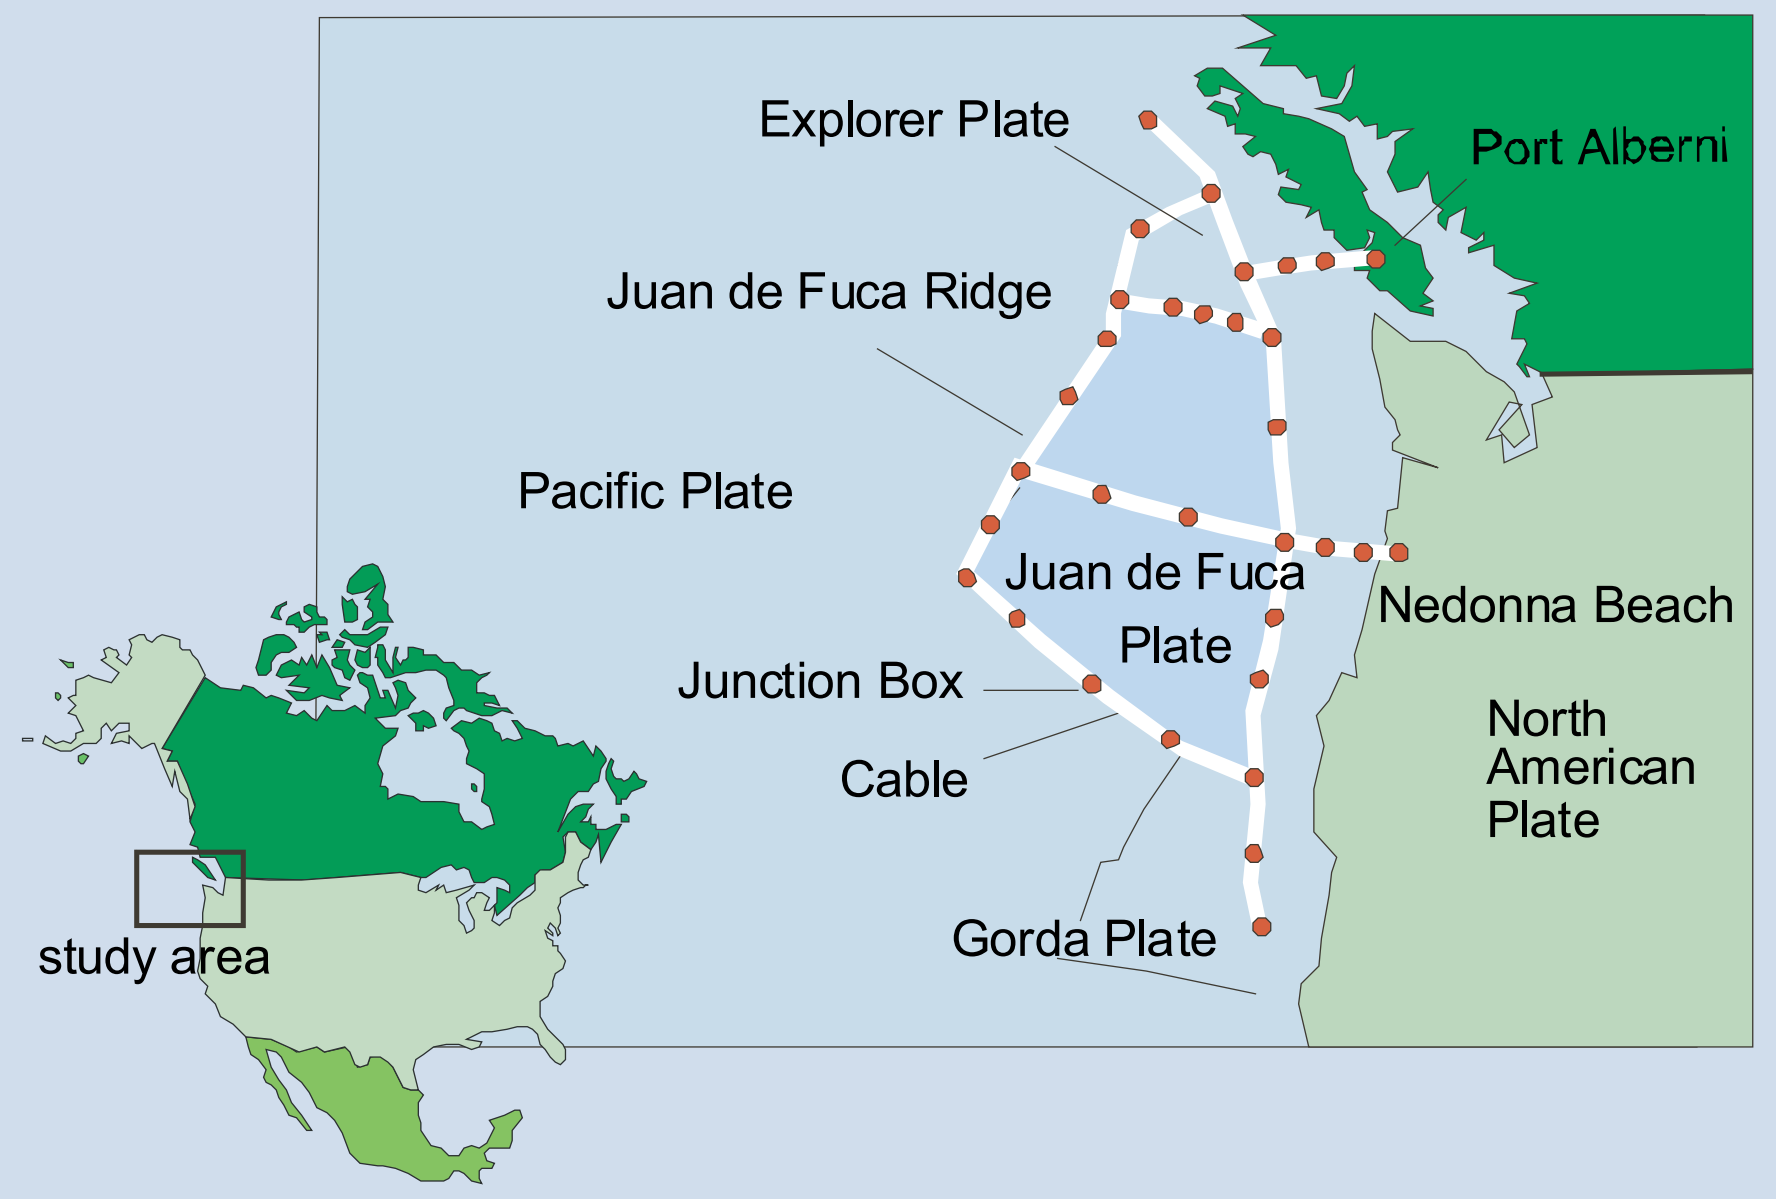

The NEPTUNE power system challenge

Design a power delivery system, using conventional submarine telecom cable, and 2 shore stations, where:

- Peak power for each shore station: 100 kW
  - This is 10 or 20 times what this kind of cable normally handles

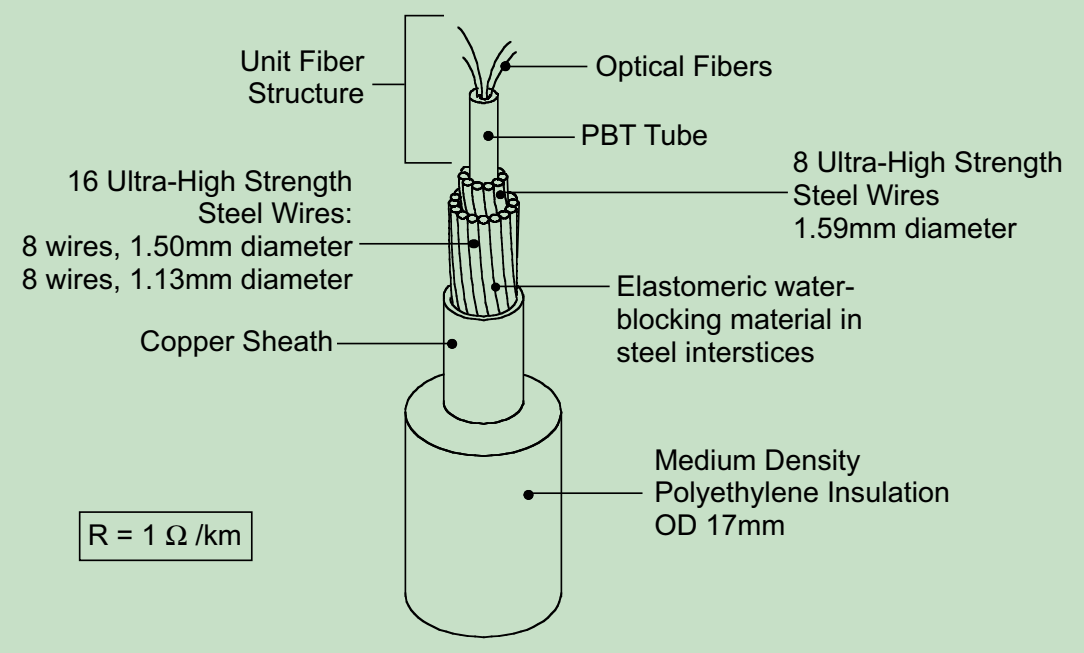

- 30 to 40 Science Nodes, peak power at each: 10 kW
  - Typical telecom cable may have repeaters, but no large loads under the water
- Max power available at ALL nodes at once: 4 kW
- Extreme reliability
  - Maximum of 4 repairs per year over 30 year lifetime
  - To achieve this, MTBF of an entire node power system must be around 1 million hours

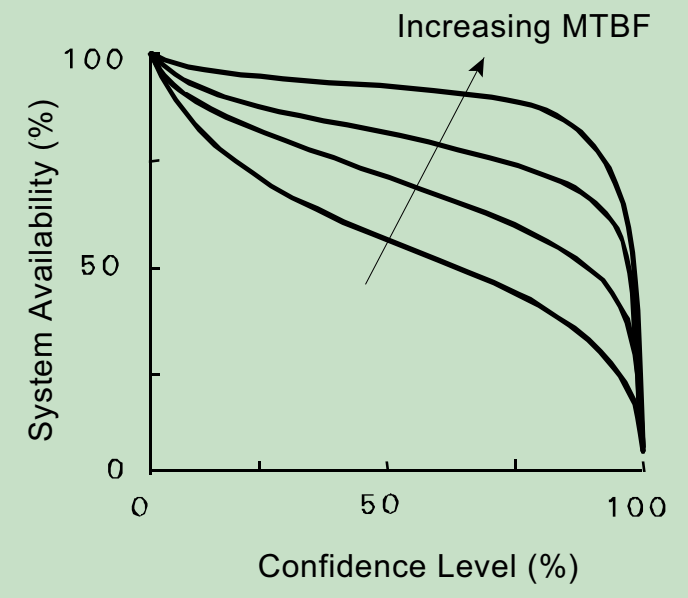

- Fault tolerant
  - Continued operation after fault
  - Minimum down time after fault
  - Cable or node faults must not cause damage to connected equipment
- Maintainable with UNOLS vessel

In other words, the problem is equivalent to reliably bringing a large amount of power to a subsea area the size of New York State using a wire barely adequate to supply a clothes dryer.

The Solution

Transmit power from shore at a relatively high voltage (10 kV) and step it down via a dc/dc converter at the Science Node. Use a parallel distribution system, with features borrowed from the utility world:

- redundancy of construction
- use of hi-reliability parts only
- a protection system for detecting and isolating faults
- a power management system

However, the design is not without difficulty

Even back in the heyday of direct current, 100 years ago, nobody ever made a dc network, and while dc is used today to transmit power over large distances, it is not done at such a low voltage as this!

- Reliability/availability is a challenge for such a novel and complex system
- Since there is a strong reliability driver, the use of a large amount of sub-sea energy storage was ruled out early on
  - This means the power system must start up from what is called a "black start, " with no power available at the node until after the node's own converter is operating
  - Fault tolerance and the protection system to ensure it are challenged by the lack of energy storage
- The reliability goals challenge the converter design, which consists of a large number of modules in series. The reliability improvement of having lower component stresses because of reduced operating voltages on each module is balanced by the fact that failure of any of the series modules could shut down the whole converter
- There are several inherent instability modes in the system
  - Voltage instability, a result of the maximum power transfer capability of the network
  - Operating stability, coordination of the multiple power sources
  - Converter stability has two aspects: the negative resistance question posed by the constant-voltage nature of the output, and the voltage/current balance question posed by the modular nature of the design
- Even the apparently simple task of interrupting the current with a circuit breaker is quite complex

Circuit Breaker

- We need to minimize arcing and transients. This is done by routing the current into a capacitor.

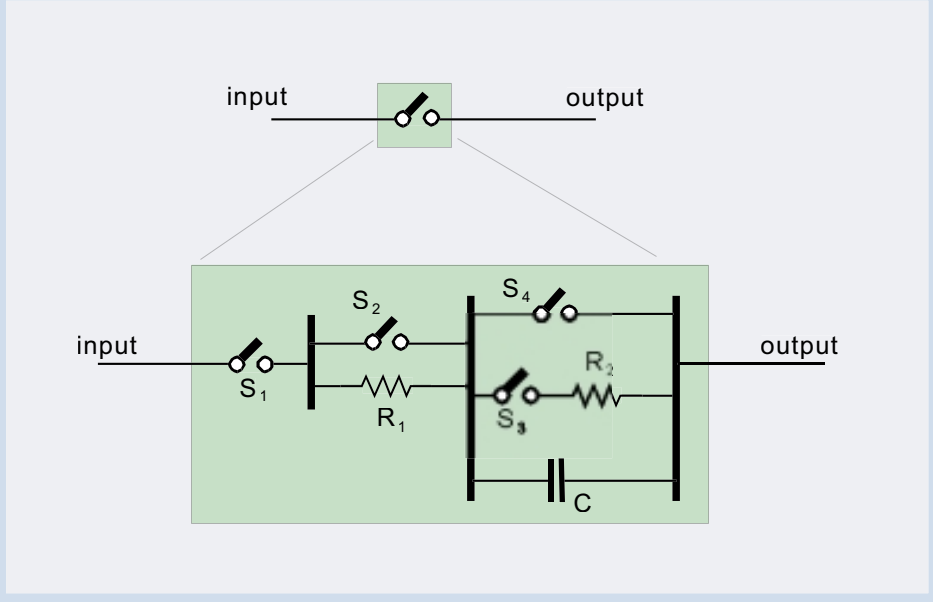

- Assume that all switches are open
- To close the breaker, switches S1 and S4 are closed first. The resistor R1 is then in series with the circuit being closed, so that the circuit can be monitored for faults before the switch S2 is closed
- To open the breaker, switch S4 is opened first. The line current then commutates into the capacitor C, and falls exponentially to zero. When the current is sufficiently low, the switch S1 can be opened. S4 is the main interruptor in this configuration
- Switch S3 is present merely to allow the capacitor C to be discharged via resistor S2

Voltage Instability

Voltage instability is the name given to the kind of instability that results when the power transfer capability of a power network is approached. Any such physical system has a power limit. In the world of electrical power, where there are feedback mechanisms to control voltage profiles across the network, the approach to the limit is characterized by the inability of the control systems to maintain voltage.

There are (at least) 3 ways to detect the onset of the condition in NEPTUNE.

- Calculate the efficiency. When instability is approached, the efficiency is approaching 50%

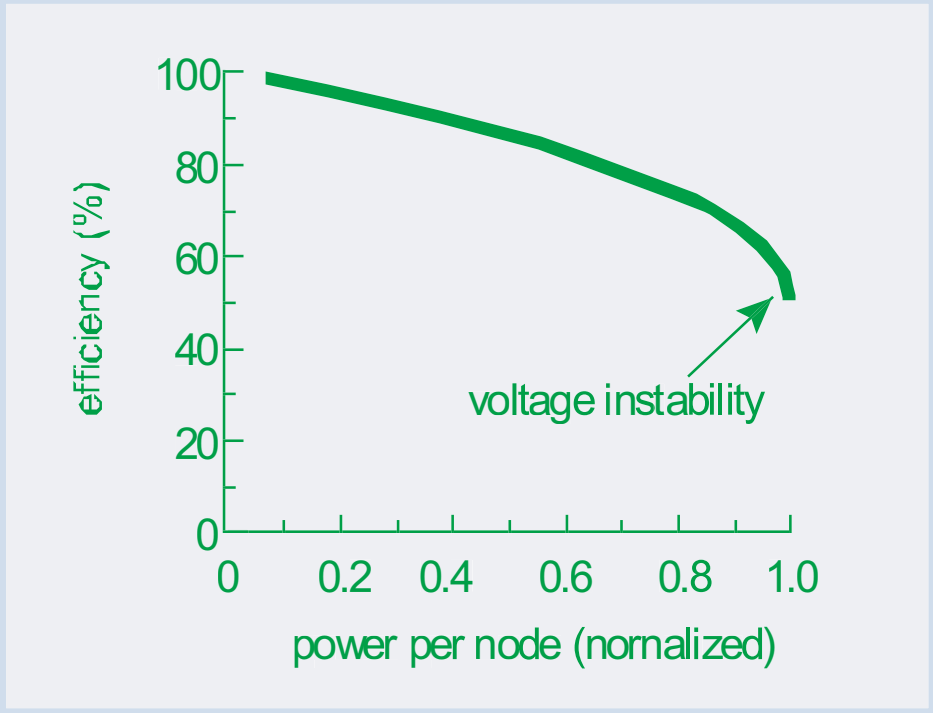

- Calculate the system Jacobian. When instability is approached, an eigenvector is approaching zero
- Look at the power flow convergence. When instability is approached, numerical convergence becomes more difficult

Interestingly, in the NEPTUNE system the constant voltage characteristic of the dc/dc converters conceals the onset of voltage collapse just as the voltage regulators would in a utility. This feature adds to the problem. However the fact that we can get a solution to the load flow problem does mean that stable operation can be achieved. The following figure shows one power flow scenario for shore stations at Port Alberni and Nedonna Beach at 10 kV (normalized to 1 per unit), and 4.7 kW load at each node.

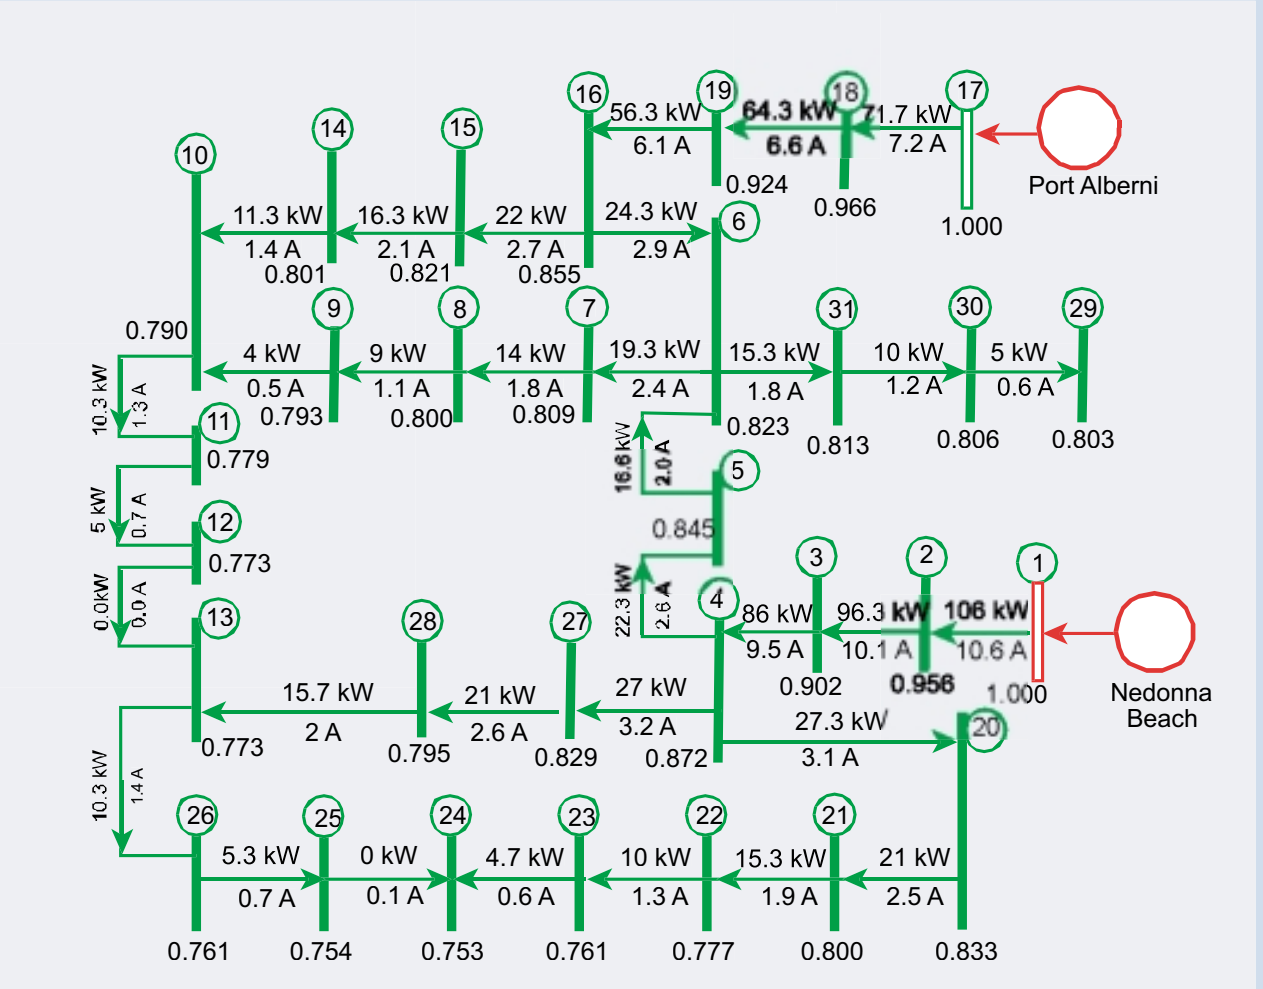

Operating Stability

Though our planned operating procedure will be radically different from utility HVdc practice, we see no particular difficulty coordinating the multiple shore station supplies.

As a starting assumption, suppose that there are 2 sources, and they are set to produce the same voltage. Each is operating as what a utility would call a "slack" generator. While utility practice is to have only 1 slack machine in any area, our simulations show that having 2 should not be a problem. Power will be drawn from each according to Ohms law. In a utility, the power would be controlled so that area generation matched area load.

It will be possible to "fine tune" the shore station voltage to change the balance of power between them. If desired, the control system could, for example, operate to have all shore stations at the same output power (rather than voltage), or to minimize cross-border power flow, or to minimize the electricity bill.

Converter Stability

Since the output of the dc/dc converter is a constant voltage (either 48 V or 400 V), the load power is constant in spite of variations in the input voltage. This means that the input current increases when the voltage decreases, i.e. the converter presents a negative resistance to the system. Normally, a negative resistance in a circuit implies instability. In the case of the dc/dc converter, the negative resistance appears because of the control action of the converter. This means that the effect is present only over the bandwidth of the control system, up to a few kHz. It is a simple matter to damp out any system oscillations by appropriate design of the converter input filter.

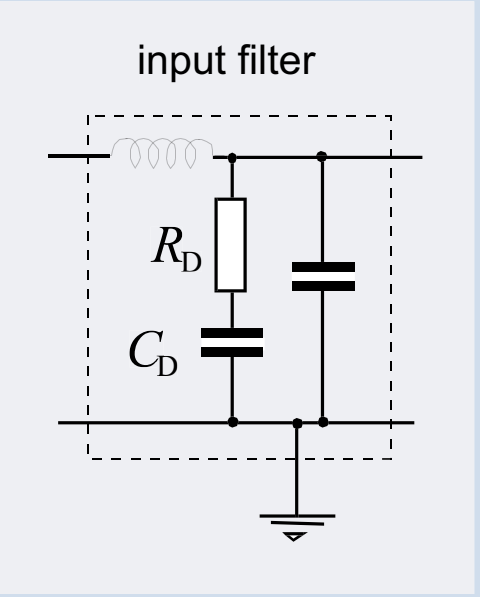

We have simulated this aspect of the system, and shown that with the filter installed, the converter is stable even for step changes in load or input voltage.

This is not the case when the filter is removed (below).

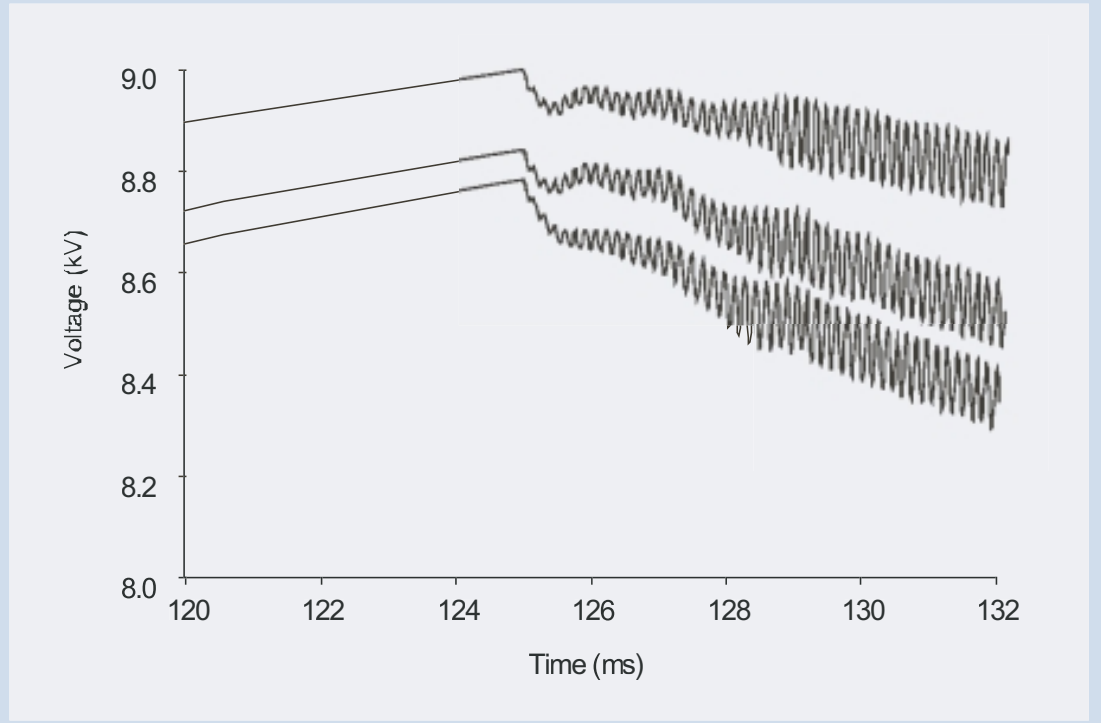

The second converter stability question arises from the modular nature of the design. We plan on a stack of converters with inputs in series and outputs n parallel.

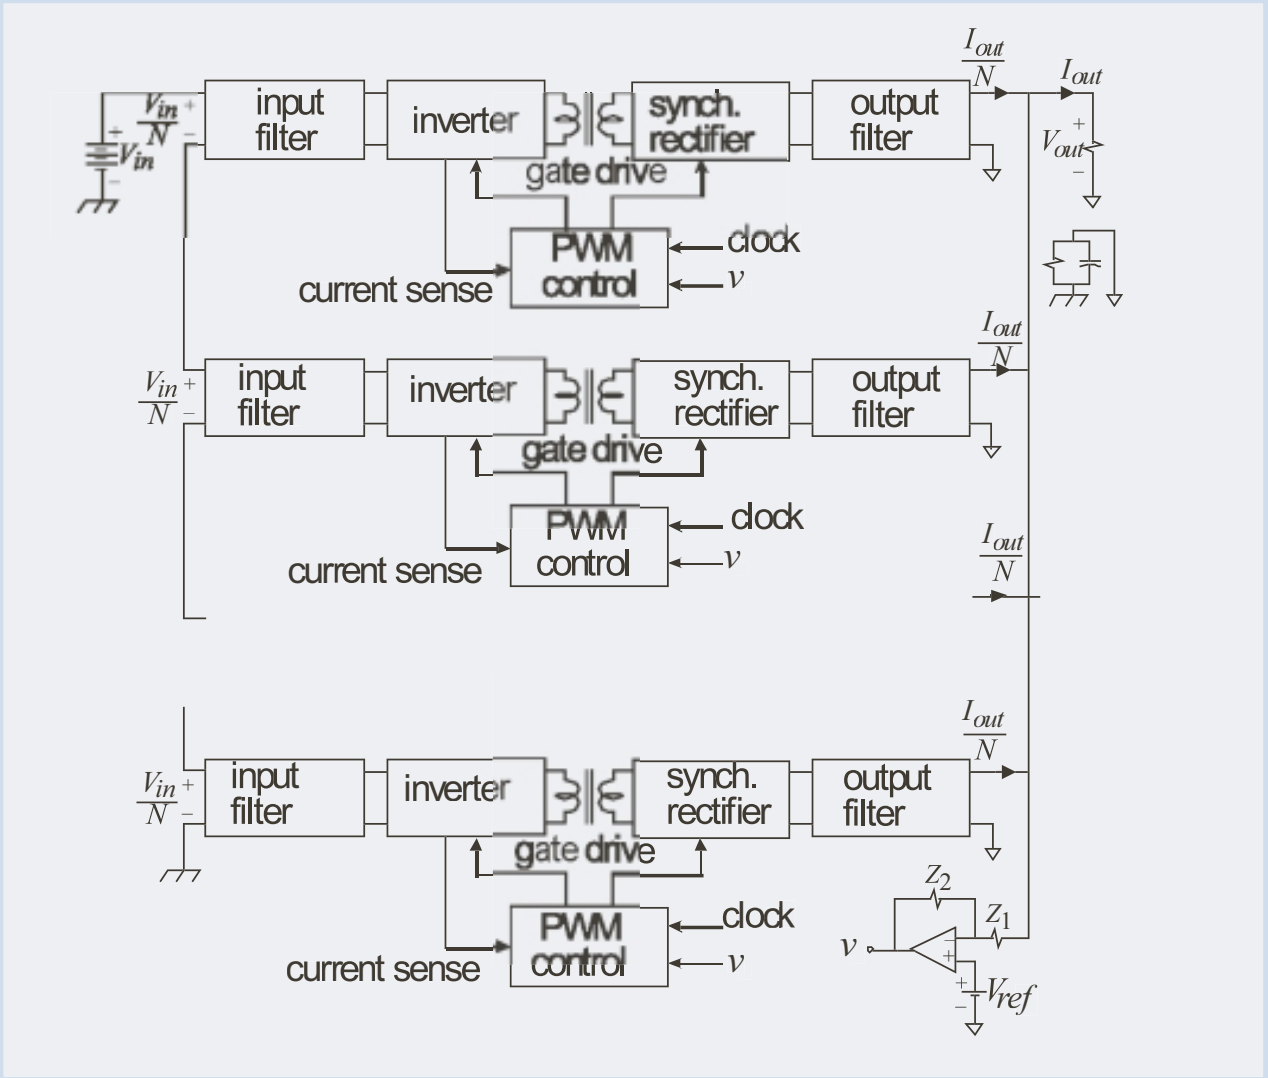

In a perfect world where all components were exactly as their labels described them, the pulse width modulation control monitoring the converter output voltage would produce identical drive pulses for the switching transistors, and each converter would drop the same input voltage.

In a practical world, where components values are subject to uncertainty, a control mechanism is needed so that the input voltage is evenly divided. This is achieved in this design by means of a current sense system that adjusts the PWM drive of each stage relative to the others.

Operations

Startup

- nodes will be powered up sequentially along parallel paths
  - first, power is applied to the node
  - a special diode arrangement allows the node to be powered from either direction, without knowing in advance which side, and without energizing the other side.

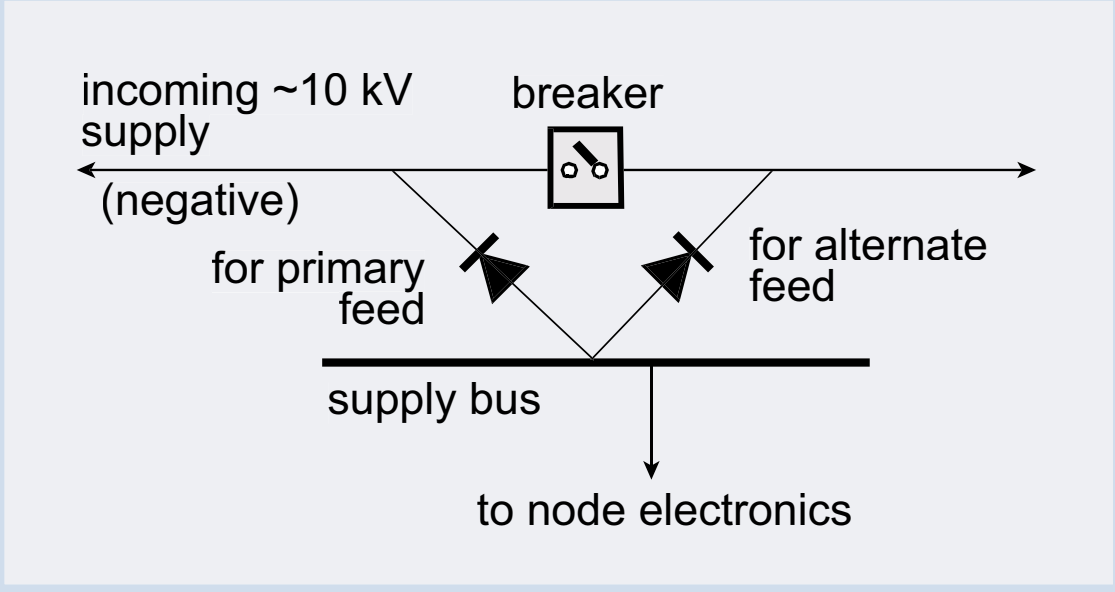

  - A startup circuit provides power for the main converters. This circuit has to work. It is therefore extremely simple in concept, and uses only a few robust parts. The trade-off is that it takes 10 seconds to start the converter
  - As soon as the main converter starts, some parts of the protection system are turned on, and the communication system is energized
  - After about a minute, the communication system is operational, and therefore the rest of the protection scheme is available
  - On command from shore, the next node in the sequence will be energized by closing the circuit breaker in the backbone (total startup time for the network is about 15 minutes)
  - As the outbound circuit is energized, it is monitored for faults. Energization takes place through a series resistor, so there is no large current spike even if the cable is faulted.
- In the event that the communication system fails for some reason, the node power system enters a safe mode. The circuit breaker is closed in this safe mode, so as to allow power to be transmitted further even in the event of a communication fault.

Normal operation is controlled by a Power Monitoring and Control System (PMACS) system of hardware and software. The Node Power Controller makes voltage and current measurements and transmits values to shore. The shore-based PMACS system performs a number of functions on a continuous basis.

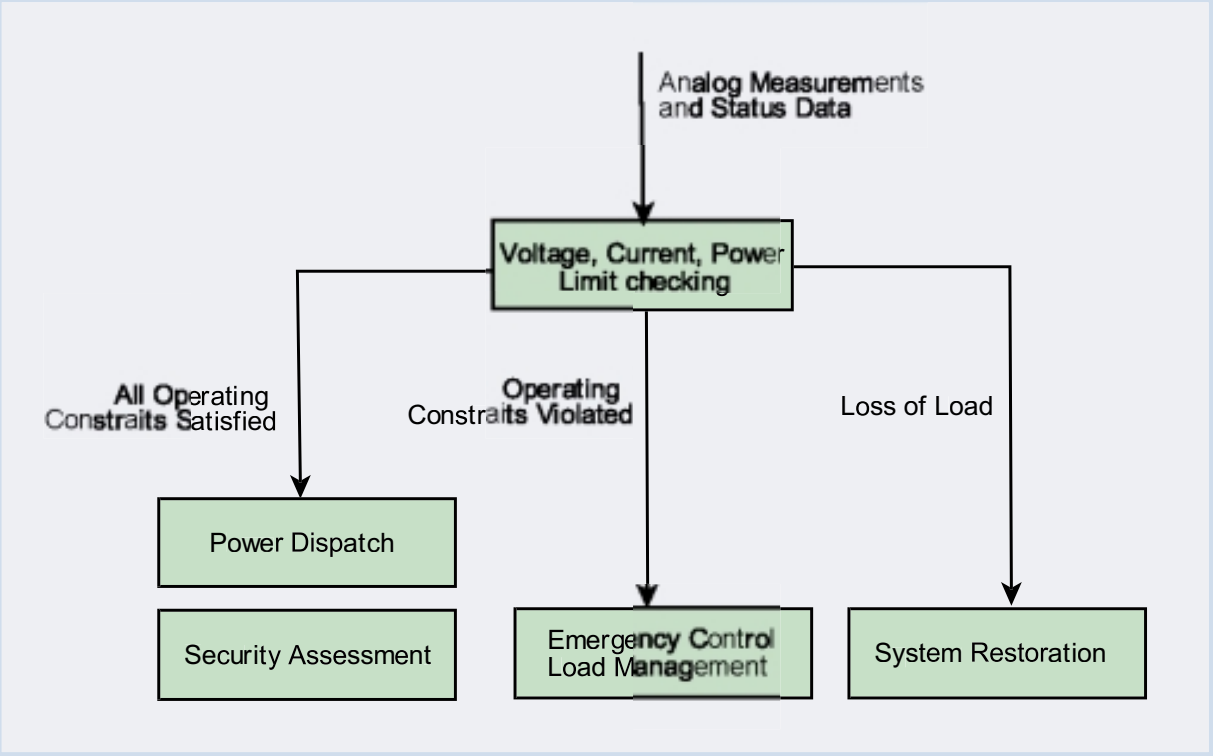

Security Assessment
- checks that all system voltage, current and power constraints are satisfied.

Dispatch
- adjusts the relative outputs of the shore stations in accordance with operating procedures decided in advance, for example, minimize the electricity bill.

Emergency Control
- adjusts voltages or loads when constraints are violated. This is where load-shedding is controlled.

Restoration
- re-energizes the system when the protection system has acted to shut all or part of it down following a fault.

Protection

Protection is more of the ambulance at the bottom of the cliff than it is the fence at the top. Nothing can prevent faults– all one can do is react appropriately

- Faults may occur in the cable and in node electronics
- Protection acts to isolate faulted parts and leave remaining parts operating
- The protection system itself may develop a fault, and so must be designed with redundancy

Cable protection

Cable faults can be caused by manufacturing or installation, defects, natural events, fishing activity or anchoring. At 10 kV under salty water, it is reasonable to assume that any cable fault rapidly becomes a short-circuit to ground. Two kinds of protection are planned for cable faults.

- Current differential compares the currents at the ends of each cable segment
- Differential protection is highly discriminatory, i.e. it does not react to faults that are not between the sensing locations

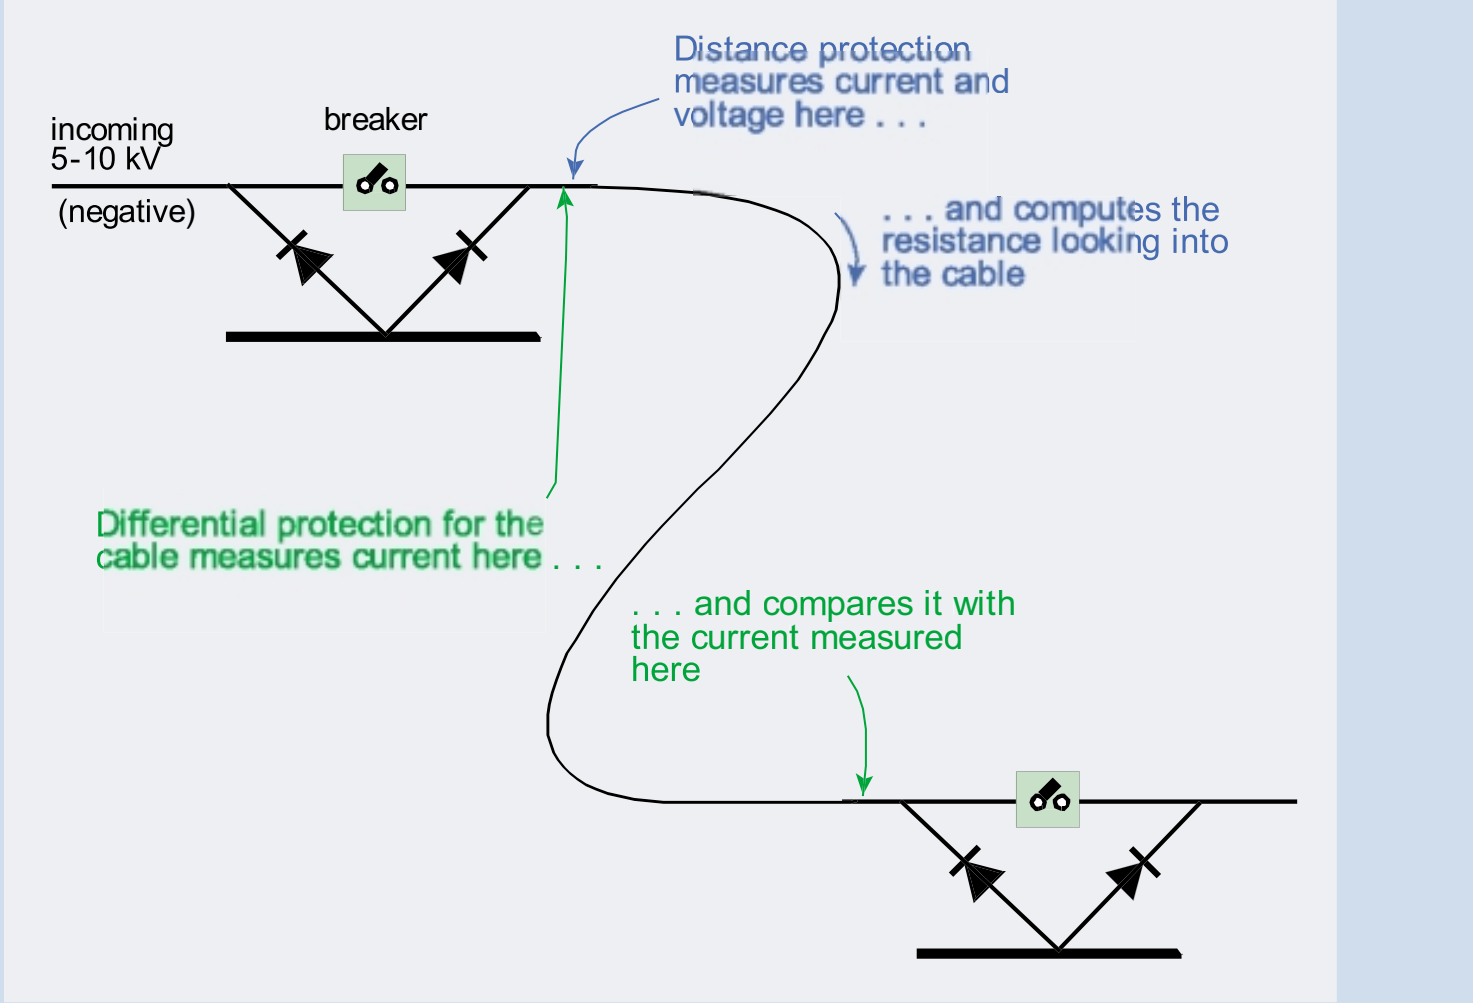

- Distance protection measures the apparent distance (based on Ohm's law) to a fault. It uses current and voltage information from one end of the cable only. It is not very discriminatory, i.e. it may overreach and react to a fault just outside its zone of coverage
- Distance protection has the advantage that it does not rely on communications

Node protection

Faults within a node are easier to detect in some ways, as the measurements needed are all relatively close together. However, depending on what the fault is, there may not be much one can do to clear it. And clearing a node fault may mean shutting the node down. In this case, the node must signal adjacent nodes to isolate it, as the individual nodes do not have breakers they can operate to isolate themselves.

One of the most challenging aspects of node protection is to detect small leakage currents from the backbone circuit. In general, such a leakage may signify incipient insulation failure inside the node. Since there will normally be a significant ground (ocean) current at the node, the technique depends on directly measuring the sum of the currents in the two backbone cables and the ground connection.

Restoration

Some faults will inevitably result in the entire network being shut down. This aspect of NEPTUNE's behavior is unlike what happens in the terrestrial power system, and is a result of the current-limiting nature of the shore station power supplies. (They limit fault current in order to prevent cable or node damage.)

However, while faults may occur, and the network may go down, it will only be down for a short time.

Faults will be cleared by the opening of backbone circuit breakers. They cannot act fast enough to prevent the effect of the fault reaching the shore station, but they will isolate the fault within about 20 ms. (By which time current limiting has come into effect.)

Since only the breakers closest to the fault will operate (it is planned to keep the others closed with a small amount of capacitively stored energy) when the fault is cleared, the voltage on the rest of the system will return to normal all over the system at the same time. Therefore, within a minute or so, all nodes that were not affected by the fault will be operating again.

We are studying the issue of adding more energy storage to keep the (essential) internal loads alive through a fault. If so, system operation should be resumed in less than a second. The design will hinge on the impact on the system MTBF.

Maintenance

Even though we plan on an extremely reliable system, some maintenance must be allowed for. It has been decreed that, as far as possible, maintenance will require a UNOLS vessel, not a cable ship.

This requirement means that a node has to be connected to the backbone via a single cable, as a UNOLS ship cannot handle two cables at the same time. The arrangement is shown below:

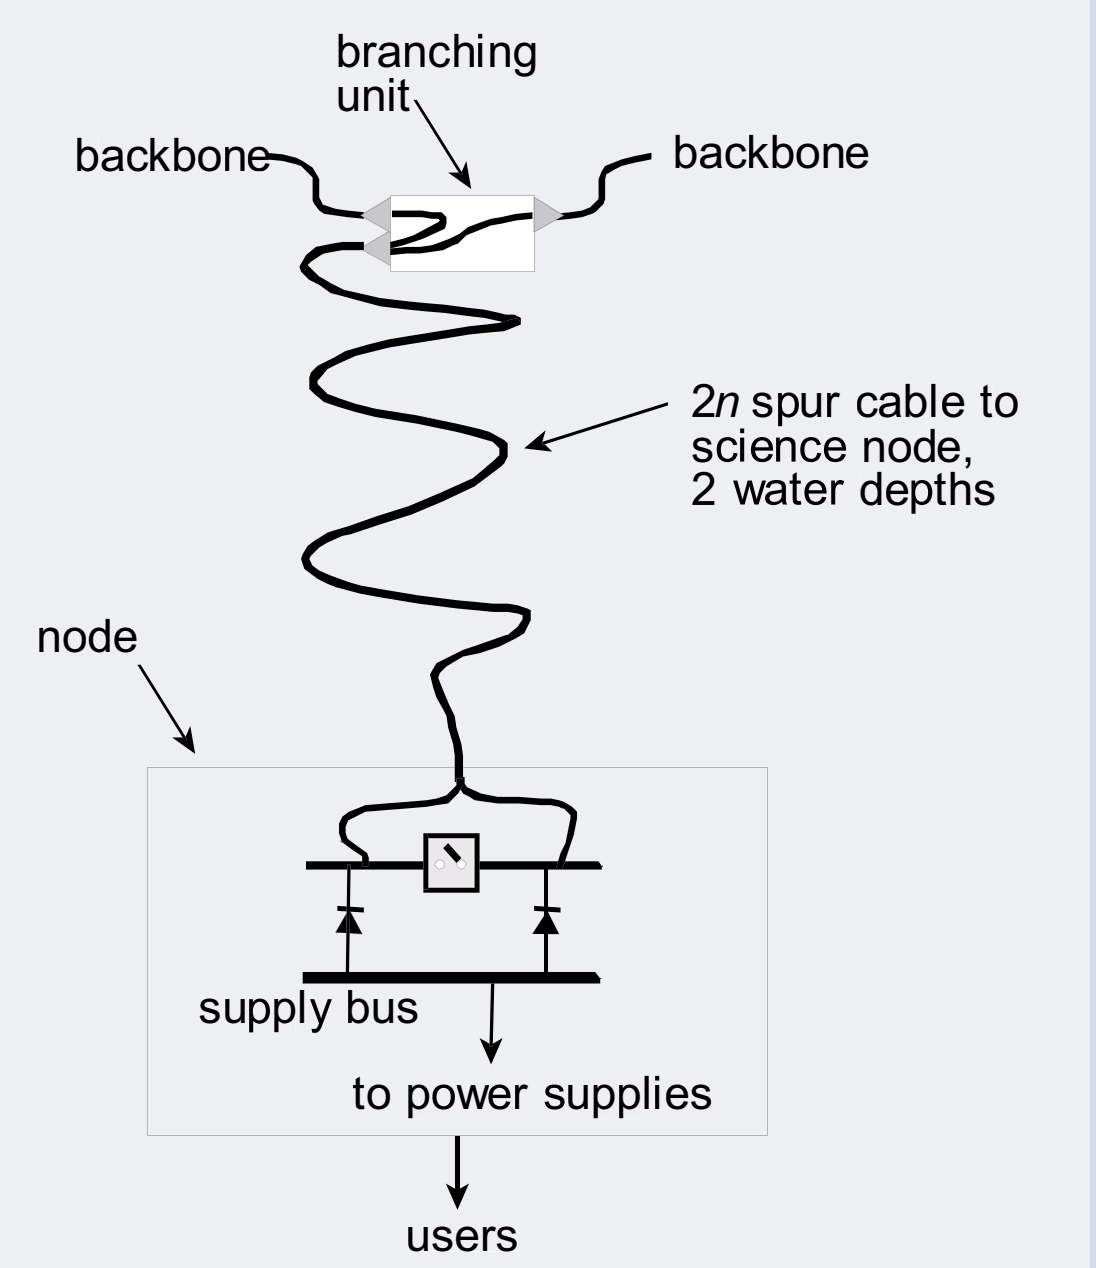

When the node is aboard the ship, it can be serviced following safety procedures that have yet to be worked out in detail. Most likely NEPTUNE will follow the common utility practice of requiring 2 separate safety mechanisms to fail before life is endangered. In the case of a node needing repair, these two things could be an open breaker at the adjacent nodes (along with the proper lockout mechanism), and a visible ground in the node. Until the visible ground is applied inside the node, the circuit would be treated as if it were live. A voltage as low as 10 kV is routinely handled in utilities, and could no doubt be handled on the ship by trained personnel.

Actual repairs would probably not be performed at sea, as most of the power system will require clean-room assembly to achieve the required MTBF figures.

User Considerations

User "Contract"

- User will provide schedules and profiles for expected autonomous changes in power loads by the instruments
- User will provide schedules for planned power cycling or other interactions required by the instruments
- Power cycling and other interactions will be accomplished through the observatory management system

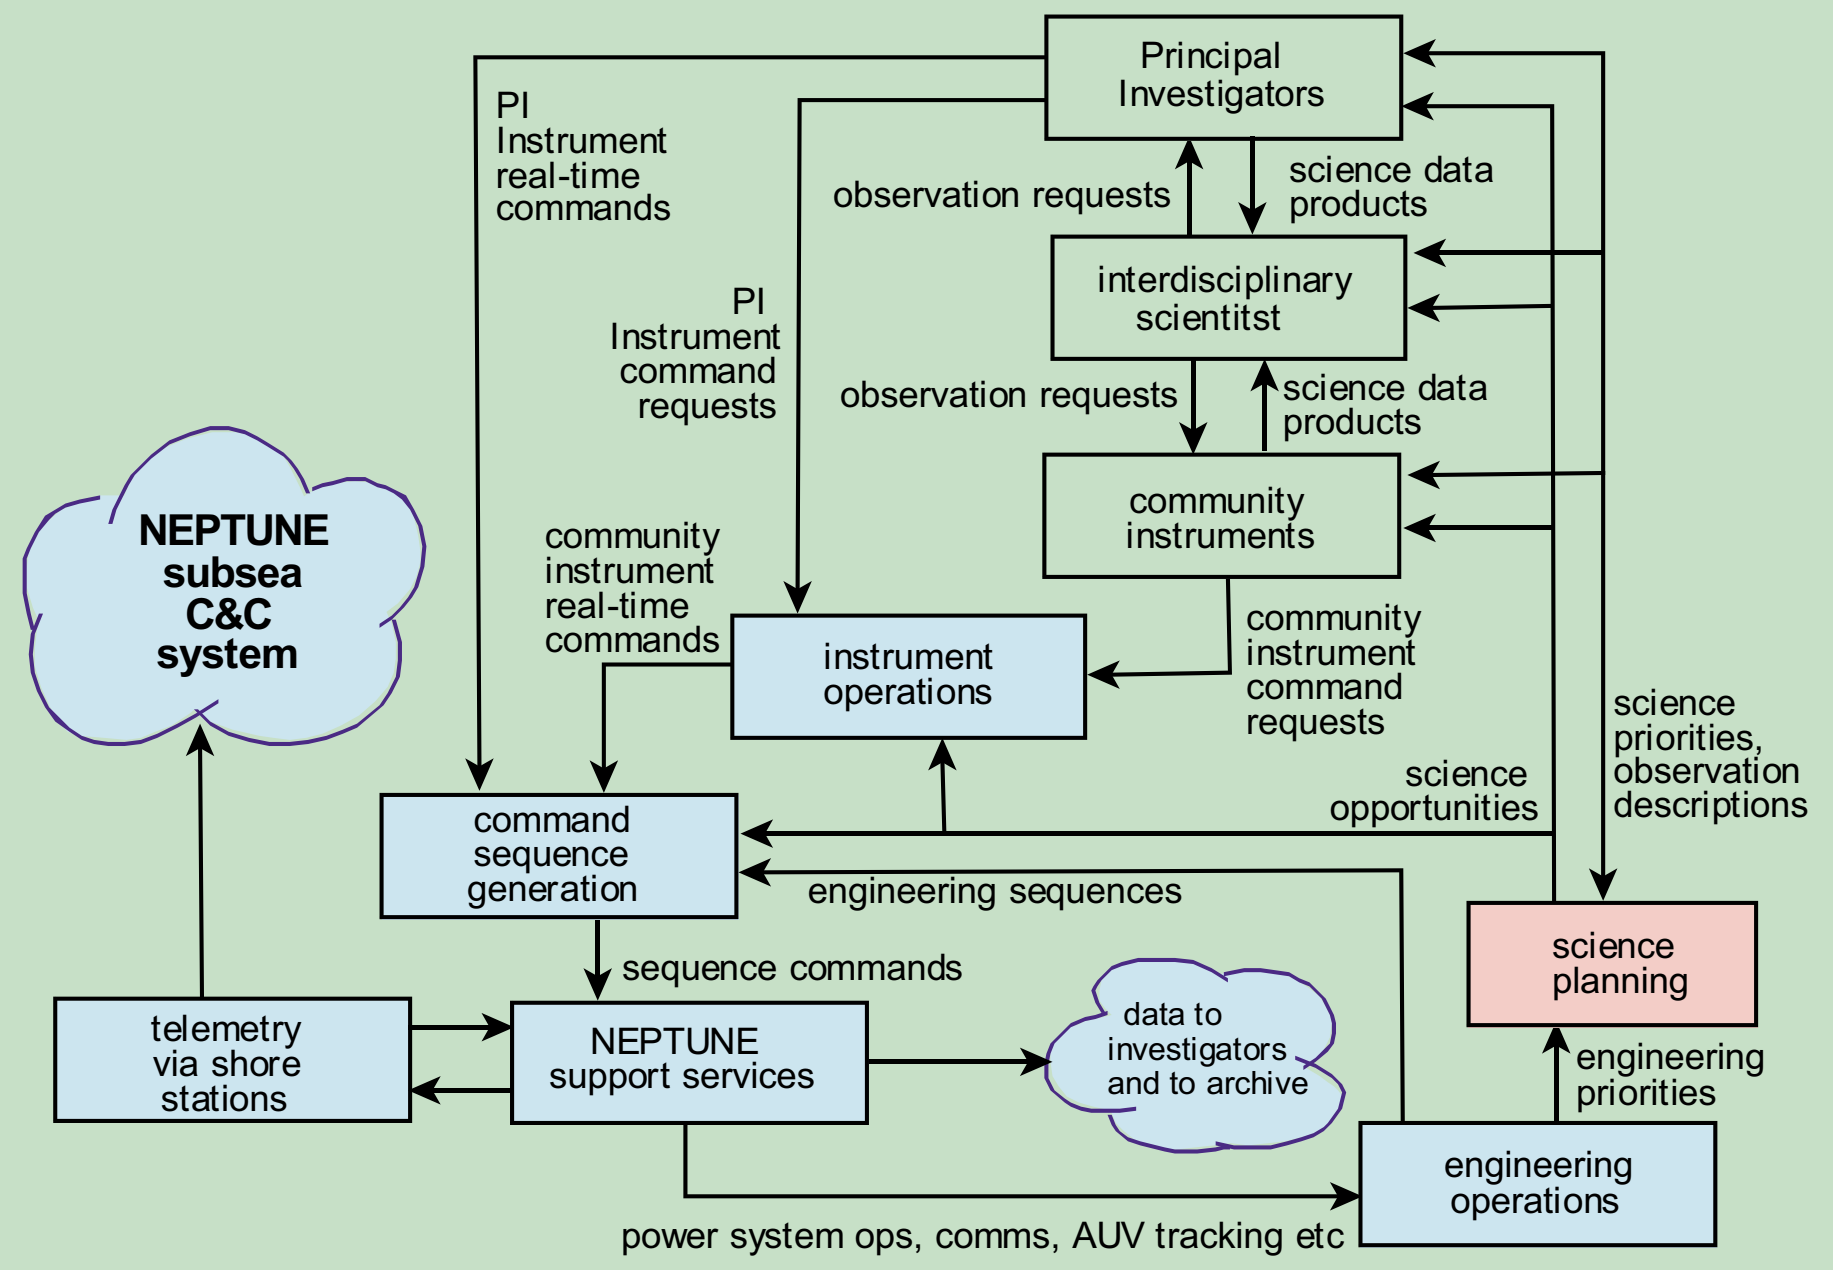

Proposed User Categories

- General – most science loads. These are energized so long as there is no particular problem or power shortage
- Deferrable – lighting, battery charging, big power users that may be disconnected if doing so will allow other load to be served in a time-critical fashion
- Essential – internal loads such as the communication system, the protection system. Powered any time the converter is operating
- High-Priority – Science loads that warrant an extra effort to keep them energized. It is anticipated that this category will include temporary members, designated in response to observatory conditions

User Interface and Sensor Requirements

- 48V and 400V available on separate pairs of pins
- Isolation of all circuitry (electronics, cables, connectors) from case/sea water is required
- Ground faults will result in instrument being disconnected
- User will need to incorporate battery, capacitor or other power storage into instrument if required for power outages
- Voltage conversion necessary if instruments require other than 48V or 400V

Concluding Remarks

While the work presented here is being done within the framework of NEPTUNE, it will be generally applicable to all cabled ocean observatories using similar cable.

Science community input is needed to review power system requirements; e.g. permissible down-time given a fault.