NEPTUNE Gigabit Ethernet Submarine Cable System

A.R. Maffei, Woods Hole Oceanographic Institution
Woods Hole, MA. 02543

G. Massion, Monterey Bay Aquarium Research Institute
7700 Sandholdt Road
Moss Landing, CA. 95039-9644

J. Bailey, Woods Hole Oceanographic Institution
Woods Hole, MA. 02543

A. Bradley, Woods Hole Oceanographic Institution
Woods Hole, MA. 02543

H. Frazier, Dominet Systems
280 E. Gish Rd
San Jose, CA 95112

A.D. Chave, Woods Hole Oceanographic Institution
Woods Hole, MA. 02543

S. N. White, Woods Hole Oceanographic Institution
Woods Hole, MA. 02543

S. Lerner; Woods Hole Oceanographic Institution
Woods Hole, MA. 02543

D. Yoerger; Woods Hole Oceanographic Institution
Woods Hole, MA. 02543

R. Buddenberg, Naval Post Graduate School
1 University Circle
Monterey, CA 93943

Abstract- **The NEPTUNE scientific submarine cable system will wire the Juan de Fuca tectonic plate and turn it into an interactive ocean sciences laboratory. NEPTUNE will provide 30 seafloor nodes distributed over a 500 by 1000 km area to which many scientific instruments may be attached. The NEPTUNE backbone communications system will be based on commercial off-the-shelf Gigabit Ethernet components packaged as ocean bottom communications modules. This paper describes the requirements and technological alternatives that lead to this choice. In addition to the high speed communications needs, separate low-speed channels are required for out-of-band management and the distribution of high accuracy, low jitter time information. A proposed design that meets these two functions is presented. Finally, NEPTUNE instrument communications will be based on slow/fast Ethernet, and a Scientific Instrument Interface Module (SIIM) will be provided to facilitate connection.**

I. INTRODUCTION

The study of the dynamic, interactive processes that comprise the earth-ocean system requires new approaches that complement the traditional ship-based expeditionary mode that has dominated oceanography for the past century. Long-term access to the ocean is needed to characterize the diverse range of spatial and temporal scales over which natural phenomena occur. This can be facilitated using ocean observatories to provide power and communications for distributed real- time sensor networks covering large areas. Real-time networks also enable an education and public outreach capability that can dramatically impact the public attitude toward the ocean sciences.

The NEPTUNE project (http://www.neptune.washington.edu) is a joint US-Canadian effort to "wire" the Juan de Fuca tectonic plate located off northwestern North America with 3300 km of dedicated scientific fiber optic cable hosting 30 science nodes spaced a nominal 100 km apart. Each seafloor science node will provide power at the multiple kW level and two-way communications at a Gb/s

rate to many experimental packages. Aggregate backbone communications will run up to 10 Gb/s. Two shore stations link these seafloor nodes to the Internet. Fig. 1 shows the planned layout for NEPTUNE. Installation of NEPTUNE is expected to begin in the 2006 time frame.

NEPTUNE differs from a conventional submarine

Figure 1. **Bathymetric map showing the layout for NEPTUNE. The dots connected by lines denote the backbone cable and locations of the seafloor science nodes, respectively. The dashed line is the approximate location of the US and Canadian Exclusive Economic Zone.**

telecommunications system in two key respects. First, NEPTUNE requires data input and output at many seafloor sites rather than a few land terminuses. This means that transmission lasers and data switches will have to be placed underwater. Second, NEPTUNE has to distribute power at variable and fluctuating rates to many seafloor instruments in addition to energizing its own internal systems. For these and other reasons, the engineering solution for the NEPTUNE power and communications systems differ from those used in commercial telecommunications systems. However, NEPTUNE will take advantage of the submarine fiber optic cable technology used in telecommunications for its backbone, and will be installed using conventional cable laying assets.

II. CONSTRAINTS, TECHNOLOGIES, AND TRADEOFFS

A. Functional Requirements

The functional design of the NEPTUNE system must be driven by science requirements. For example, the locations of the 30 seafloor nodes in Figure 1 were determined by multi-disciplinary (e.g. geological, biological, chemical, etc.) science needs. Through an assessment of present and projected future ocean instrumentation and experiments, system parameters such as the peak and average data rate, power level, and allowed data latency and jitter (from nodes to shore) have been defined. Aggregate system capacity of up to 10 Gb/s and delivered power at the 5 kW level per site is sufficient to meet science goals. The system must also distribute accurate (1 μs) time information to seafloor instruments. These issues are further discussed in [1].

The infrastructure for NEPTUNE consists of five systems: communications, power distribution, system control, time distribution, and data management and archiving. Each of these components must be designed as end-to-end systems which interface cleanly to the remainder, and must be highly fault tolerant. Physical packaging of the seafloor nodes must be accomplished in a way which facilitates science as well as maintenance. The system engineering for NEPTUNE to accomplish these goals is presently underway [4]. A description of the NEPTUNE power system design can be found in [5]. This paper will focus on the communications, system control, time distribution, and physical packaging subsystem approaches, including the requirements and technological alternatives in each area.

B. Analysis Criteria

Choosing the best backbone network technology for NEPTUNE is not a simple matter. There are three areas of analysis worth considering:

1. *Business observations:* the selected NEPTUNE technology should be close to the commercial Internet mainstream so that components can be purchased in the commercial marketplace (i.e., a COTS approach). Buying into dead-end technology or using a custom approach for a program with a 25 year life cycle like NEPTUNE can be expensive and limiting.

A sense for the wave of the future needs to be part of an analysis.

2. *Technology observations:* different networking technologies offer different capabilities and impose concomitant limitations. Internet technology is rife with solutions for problems that NEPTUNE may or may not have. There is a significant payoff from simplicity in any network design, so the technology NEPTUNE chooses to avoid is as important as that which NEPTUNE buys into.

3. *Specific NEPTUNE requirements:* the first two entries in this list don't have any seafloor-specific flavor; the conclusions would probably be similar for a terrestrial network with the same data rate and distance requirements. However, NEPTUNE clearly presents some unique requirements for a network. For the science nodes, these include packaging to fit into reasonable sized pressure cases, moderate power consumption (both because power is limited and because power represents heat that must be transferred out of a pressure case), high reliability and fault tolerance, ease and effectiveness of management, the ability to sustain upgrades as the technology evolves, and compatibility with an easily understood science interface. Many of these are difficult to define precisely, but it seems reasonable to limit the hotel load for all science node systems to about 500 W, leaving the majority of the available power at a node for science. Pressure case cost and weight rises faster than the square of the inside diameter, so keeping this contained has a real impact on cost and ease of shipboard handling. The remaining issues have to examined on a case-by-case basis.

C. Submarine Telecommunication Technologies

Originally, NEPTUNE explored how the system might be based on COTS submarine telecommunications technologies. Such systems are designed to provide voice circuits between continents, although recently they have been carrying more Internet and other data traffic than voice traffic.

Submarine telecommunication systems consist of seafloor fiber optic cables arranged in a point-to-point or, more recently, a ring topology. The seafloor plant contains only the optical amplification systems required to boost the optical signal at 60-100 km intervals. The state-of-the-art systems utilize erbium doped fiber amplifiers (EDFAs) to provide path gain. Although EDFAs are very reliable, they are also very expensive ($500K-$1M each) in submarine applications. Branching unit technology has also been developed to allow forking of seafloor cables to multiple landing sites.

Recent submarine systems employ multiple wavelengths over each fiber via dense wavelength division multiplexing (DWDM) as a way to increase capacity and Optical Add/Drop Multiplexors (OADMs) to branch wavelengths. This has been a very costly technology, as wavelength control must be precise and dispersion compensation along this system is very complex. Current DWDM seafloor systems have capacities that approach a terabit per second [2].

While the seafloor plant consists only of optical fiber and amplifiers, submarine telecommunication systems usually are based on telephone standard SONET/SDH protocols to transport voice and data information over them. This technology is optimized to carry voice traffic though it is capable of carrying data as well.

There are a number of differences between NEPTUNE requirements and a submarine telecommunications system that strongly suggest moving in a new direction. These include:

1. A huge mismatch between NEPTUNE requirements of a few Gb/s data rate and telecommunications systems operating at hundreds of Gb/s;

2. NEPTUNE requires data aggregation and switching on the seafloor, where a telecommunications system is designed as simply a fat data pipe in the ocean;

3. While telecommunications systems have stringent delay and non- interruption requirements (due to their telephone/voice traffic origins), NEPTUNE s data traffic needs are more relaxed in these areas;

4. Once a telecommunications system is installed underwater, it is rarely touched unless it needs repair. The NEPTUNE infrastructure is very likely to be extended after its initial deployment, and must be able to cope with major changes in topology and traffic.

D. Data Networking Technologies

It is clear to us that NEPTUNE requirements resemble those upon which Internet-like computer and data networks are based. This section examines network technologies available for use in NEPTUNE to determines which are applicable, which are not, and which may become available for NEPTUNE in the near future.

1) Optical IP Alternatives

Both the science user interface and shore station Internet interface will utilize Internet Protocols (IP) which have clearly emerged as the dominant internetworking protocol. There are a few non-IP solutions, such as switched ATM to the end user, but none are available in the marketplace in a practical sense, and will not be considered further. All of the pertinent technologies work over single mode optical fiber, so the physics of light transmission is much the same for all, and consequently neutral to a compare/contrast analysis. Thus, there are about four ways to build an IP backbone for NEPTUNE:

1. *IP over some switched technology* such as frame relay or asynchronous transfer mode (ATM). Assuming we have a fiber optic backbone cable, these switched technologies will in turn ride over a multiplexing technology such as synchronous optical network (SONET)/synchronous digital hierarchy (SDH). The purported attraction to this approach is the promise of switched virtual circuits connecting science users at the seafloor and on land. However, with ATM, these quality of service characteristics stop at the router, so there is limited practical payoff.

2. *IP directly over SONET*, sometimes called packet-over-SONET, eliminates the complexity of ATM/frame relay, which adds no value to networks without voice requirements, or switched virtual circuit needs. This approach was used by large ISPs a few years ago since the available highest capacity optical links were larger than router capacity. As router capacity has increased with the advent of layer 3 switches that work at line speed using routing fabrics, this approach is declining in popularity, and IP over SONET is likely to become a legacy technology.

3. *IP directly over fiber*, eliminating SONET multiplexing entirely. Since individual datagrams in an IP system are addressed by destination and application, they do not benefit from the explicit circuit-oriented multiplexing of SONET. Further, there is no need for the dim fiber represented by the backup links in SONET rings, and all fiber can be used actively with faults bypassed using dynamic routing tables. This approach is quintessential Internet.

4. *Stretched local area network (LAN) technologies* (or IP over Ethernet to preserve the colloquialism) are a rapidly growing phenomenon as the capabilities of campus area networks have been increasing. Conceptually, what started out as a single segment of LAN has undergone mitosis and gained a vertebrae. Fiber distributed data interface (FDDI) and various flavors of Ethernet (fast, Gigabit) are commonly used as backbones with a desktop Ethernet (10/100 Mb/s) fan out at the user end. While FDDI has excellent inherent fault tolerant characteristics, the FDDI market is small and diminishing to the point that chip vendors are shutting down the FDDI fabrication lines, and hence it ought not to be considered for NEPTUNE. This is largely because a fiercely competitive Ethernet market is driving technology forward, reliability up, and prices down. Further, a great deal of development effort is being focused on making these LAN technologies work over long (100 km) physical distances. Technologically, the same distance improvements could apply to Gigabit Ethernet (GbE), IP over fiber, or SONET, but business attention has largely focused on the first two.

2) Other Optical Technologies

Current GbE products on the market which interface to single mode optical fiber utilize a single optical channel rather than a Wave Division Multiplexing (WDM) approach. WDM technologies such as DWDM and CWDM (Coarse WDM) allow multiple wavelengths to operate over a single fiber. While it would be possible to design a WDM system that could implement GbE using discrete photonic components, there is little to be gained for NEPTUNE purposes. Modest fiber count commercial-off-the-shelf (COTS) submarine cables can easily accommodate NEPTUNE backbone data rate needs using standard GbE GBIC (GigaBit Ethernet Interface ConIntechnology operating

over multiple fibers. The additional engineering complexity (and thus increased cost) to implement a reliable WDM system is substantial, especially to ensure the wavelength stability of narrow-band transmission lasers at each of the seafloor nodes[1]. Additional components such as optical transponders to convert wavelengths generated by COTS components to ITU standard wavelengths may also be required. Thus, from an economic and reliability standpoint a WDM approach is not warranted for NEPTUNE.

Data networks usually utilize a mesh network architecture to provide redundant paths and ensure rapid restoration in the even of a fault. As an alternative for NEPTUNE, it would be possible to utilize a star topology in which a single pair of optical fibers (or more likely a pair of wavelengths on a WDM system) are assigned to each node and then linked directly to a shore station. This requires the same optical system complexity and requirements as a GbE WDM solution, and has a lower fault tolerance because non-overlapping paths to each node very difficult (perhaps impossible?) to build into the topology. Another limitation to this design is that the fixed number of fibers/wavelengths limits future growth. It has the advantage of eliminating the need for a router at each node in favor of a slightly simpler data switch, and probably is capable of operation at lower latency than a standard data network. These small advantages are significantly outweighed by the disadvantages. This approach has little to recommend it over a COTS GbE solution, and it was not pursued further.

E.   Cross Comparisons of Viable Alternatives

Some further cross comparisons will accentuate the differences between the Optical IP alternatives vis a vis NEPTUNE requirements. SONET and ATM are solutions designed to grow an essentially voice network to increasingly large communications pipes while accommodating IP data on the side. Further, SONET is a virtual circuit, time domain multiplex approach to implementing a physical layer, while IP data is bursty, which makes SONET comparatively bandwidth inefficient. As a result of a predominantly telephone orientation, SONET is also designed to provide nearly instantaneous restoration in the event of a fault at the cost of tying up most of a fiber in standby mode. As a result, these technologies offer many features that NEPTUNE does not require, but which increase switch and multiplexer complexity, size, and cost.

Internet technology establishes connections at the transport layer (predominantly with TCP) and requires only connectionless (stateless) datagram service from the lower

layers. ATM, SONET, and X.25 establish virtual circuits at the data link layer, but only within segments and not across routers. As a result, these virtual- circuit oriented services are not as flexible as purely connectionless services such as those provided by LAN technologies or frame relay.

A SONET virtual circuit is synchronized across all the multiplexers in that circuit to a single very precise master clock. In LAN technologies and the IP over fiber solution, there are no virtual circuits and the real circuits are synchronized on a link-by-link basis. In both of these cases, the synchronization issue stops at the router, simplifying the overall system.

In the case of IP over fiber, buffering is done in router queues. In a purely switched Ethernet solution, buffering is done at the chip level. In both cases, with a constantly fully loaded network, the occasional buffer overflow may result in packets getting lost. This is used by TCP as a congestion control indicator and by the Internet as a whole as a fairness enforcement mechanism. The occasional lost packet is perfectly normal in the Internet. The global timing of SONET does not obviate packet losses, it just shifts the location of the losses to the ends of each virtual circuit (i.e., the routers). In the case of NEPTUNE, either solution gets one to the same end. Both solutions exist all day everyday in the Internet and work fine, so lost packets are really a non-issue no matter which technology is chosen.

The ability to reconstitute the network in the event of component failure or a cable break is a clear NEPTUNE requirement. This can be achieved in several ways. Since the major loop of NEPTUNE has two shore termini, each of the seafloor routers would have an adjacent router as well as any redundant routers per science node in their reachability tables. A failure in either direction would cause routers to automatically redirect all traffic in the opposite direction or simply to the second router in the node. This mechanism operates against both component failures and cable breaks. FDDI and SONET both have a ring-wrap capability, but in the sparse mesh layout of NEPTUNE, rings would have both channels in the same cable. Wrap would help against component failures, but not cable breaks. Consequently, there is little added value to the FDDI/SONET wrap feature.

The last two alternatives, IP over fiber and a stretched LAN are converging - essentially identical - solutions, although IP over fiber is an emerging technology for which COTS availability is limited at present, and is aimed at higher data rate applications than NEPTUNE. IP over fiber with DWDM salvages SONET framing, but eschews all the rest of an essentially telephone technology. The stretched Ethernet approach uses, of course, Ethernet framing. In both cases, the frames carry IP datagrams. The key stretched LAN technology is Gigabit Ethernet (GbE), which is the highest speed version of the most widely used (more than 80% of the market) data networking technology in the world. GbE routing and switching hardware is readily available from many vendors, and several are also marketing integrated Gigabit Interface Converters (GBICs) with the ability to directly drive up to 100 km of single mode optical fiber.

SONET is designed to utilize connectivity provided by telephone companies, and hence does not explicitly

---

[1] Another technically complex challenge of long haul WDM networks is maintaining proper levels and flat gain through an EDFA (Erbium Doped Fiber Amplifier). If there is any sudden change in the power level from the optical switching or wavelength translation device then the EDFAs can be thrown off kilter. Much time is spent grooming and aligning optical links matching narrow band transmission laser power with gain profiles of the EDFA across multiple hop links. Now remember that this has to be maintained 4000 meters below the surface![3]

incorporate a repeating capability. Many intercity non-telephone company SONET trunks in terrestrial WANs run alongside railroad right of ways and in decommissioned pipelines. Every 30-40 km, you can see a line hut which contains repeating equipment. The line hut looks a bit like a house trailer (but without windows), and the giveaway is a backup generator next to it. The equipment inside is receiving the incoming signal, recovering the clock and retransmitting it, which both recovers any lost photons and cleans up dispersion. SONET applied to NEPTUNE would have to provide for a repeating capability along with multiplexing distributed along the seafloor using separate systems for each.

A fiber optic LAN switch performs these same dual functions all at once. Another way of looking at the Ethernet switch approach is that we are getting multiple duty out of a single piece of electronics:

1. The switch reconstitutes the signal so we get a line repeater function;
2. The switch provides a fanout to the array of science equipment to be attached to a node;
3. The switch provide alternate routing to at least three adjacent switches (upstream, downstream, and in-box redundant, presuming that more than one switch is incorporated per node).

This solution requires two boxes - a pair of Ethernet switches with GBICs at each node.

Any solutions layered upon either SONET or ATM would require more - and more complex - hardware on the ocean bottom. For instance, an IP over SONET solution requires both a pair of SONET multiplexers and a pair of switches in each node (presuming that full redundancy is desired), along with a repeating capability per fiber. Further, both ATM and IP over SONET COTS interfaces are large (typically, at least a 19 rack that is 19 high, and frequently much larger) and power hungry (typically, 600 W to several kW). This is because they are primarily used by ISPs in inter-city trunking where size is not an issue; the power consumption reflects their complexity and a high (relative to Ethernet) transistor count. Even with moderate repackaging by removing rack hardware, this would require very large and expensive pressure cases. Management complexity also increases as ATM, SONET and Internet devices tend to have different management interfaces. By contrast, GbE hardware is small (typically, a six full duplex port router fits on a 12 by 14 inch card), consumes of order 100 W, and could use uniform management interfaces for all parts of the network from backbone to user.

F.   Choice of Gigabit Ethernet for NEPTUNE

All of the above strongly suggest that the pure IP GbE solution fits the NEPTUNE requirements best. The NEPTUNE application has a diverse set of pure data applications, no voice applications, and video applications that can easily be handled as video over IP . Indeed, most video applications that one can imagine for NEPTUNE would not be highly interactive, so there is less motivation to precisely control jitter and latency which are the most common video shortcomings in an IP world that is either low capacity or highly congested. An IP-only solution is the best

mix between simplicity of the plumbing and applications available to the users. Neither ATM nor SONET provide any added value that is important to NEPTUNE, and certainly not in proportion to the increased expected life cycle costs.

On the other hand, Ethernet is now 30 years old - about twice the age of SONET and ATM - but it is still being actively developed; a 10 Gb/s Ethernet standard is projected to be out by 2002 and early 10 GbE products are now hitting the marketplace. Further, with Ethernet, backward compatibility is preserved in several ways. For instance, fast Ethernet chips can work at either 100 Mb/s or 10 Mb/s depending on the device that is attached to the port. All IEEE 802 LAN technologies can be easily bridged together. There are about 5 generations of Ethernet (802.3), two generations of 802.5 token ring, FDDI standardized by ANSI X3T9.5, and FiberChannel. Additionally, cable modems have been standardized under the DOCSIS standard rather than IEEE 802, but the technology is essentially Ethernet and it is explicitly Ethernet at the handoff. This variety is far more than NEPTUNE should allow in the interests of maintainability and supportability, but the illustration of the robustness of the marketplace and ability to handle backward compatibility are worthwhile to NEPTUNE.

G.   Future Network Technologies

The NEPTUNE network will probably not be deployed until 2006. Final design decisions are probably more than a year off. The data network industry develops new products and new technologies extremely quickly. The following technologies are on the current horizon. When commerically available they need to be evaluated for use in NEPTUNE against the criteria listed above.

The IEEE 802.3 working group is developing a 10 Gigabit Ethernet (10GbE) standard which is designed for both metropolitan area and wide area networking. Non-standard versions of 10GbE are available on the market today. It is not yet clear whether the optical fiber and components required for the eventual 10GbE standard will encourage the development of commercial switches that can meet NEPTUNE s size and power constraints.

The IEEE 802.17 working group is developing a Media Access Layer (MAC) protocol that would provide SONET-like recovery and protection features for both Gigabit Ethernet and 10 Gigabit Ethernet. 802.17 would replace the 802.3 protocol in both of these implementations and provide failure recovery times on the order of 50 milliseconds. This will be a great improvement over the several hundred milliseconds possible with current IP-layer routing protocols. Commercial implementations are expected within the next year and a half.

Smaller, higher speed, less expensive, more innovative optical components have been announced regularly in the recent past. New optical switches, amplifiers, passive wavelength translators, electrical repeaters, optical cross connects, etc. may be available in future COTS network components. Although recent news is that this industry is slowing down, the NEPTUNE project must keep an eye out for products that can meet NEPTUNEs size, power, reliability, and economic constraints.

III. A Conceptual Design

A. Communications System

The NEPTUNE communications system consists of three parts: the shore station, the delivery channel (i.e., optical fibers), and the node system, and operates at two levels: the backbone system and the science instrument system.

The shore station serves as the link between the seafloor part of NEPTUNE and the data archive/ Internet. It also monitors communication system performance and makes necessary adjustment to the seafloor communications components.

The node system must aggregate data from science instruments operating at highly variable rates, switching it onto the backbone cable and ultimately to the shore stations. It must also distribute command and control information to instruments.

As justified in Section II, Gigabit Ethernet (GbE) has been selected as the backbone communications technology for NEPTUNE. GbE routing and switching hardware is readily available from many vendors. Standard converters are also available with the ability to drive 100 km or more of single mode optical fiber. This yields long haul 1 Gb/s full duplex communications using a pair of optical fibers. Higher data rates are feasible using multiple fiber pairs; the NEPTUNE goal of a 10 Gb/s backbone could be met by ten pairs of fibers. Each NEPTUNE node will contain a pair of GbE routers for redundancy.

A higher (10 Gb/s) rate version of GbE will be commercially produced in 2001, and will be evaluated when available. 10 GbE will have the same functionality as GbE, including range capability of 100 km on single mode fiber. The use of 10 GbE could substantially reduce the fiber count in the backbone cable, resulting in some cost reduction and enhanced data capacity.

The science instrument interface on NEPTUNE will be implemented using 10 and 100 Mbit/s Ethernet as the communications technology. This is further discussed in subsection E.

B. System Monitoring and Control

Monitoring and control of the NEPTUNE data networks will be implemented using the simple network monitoring protocol (SNMP) served from network supervision work stations located on shore. SNMP clients can be included in the power distribution system components to centralize and standardize all monitoring and control functions. SNMP clients can also be incorporated into the science instrument interface to simplify user level supervision. These high level system monitoring and control systems will function when the data communications backbone systems are operational.

In addition, NEPTUNE will include a low level or out-of-band control system. This will provide full control of all router functions through a standard serial interface. It can be used to configure the network and download software to the backbone routers. It also interfaces to the node power supplies. The low-level control system will be fully redundant, and operates on a pair of dedicated optical fibers.

It consists of two separate half duplex systems operating in a master-slave mode. Each link on a given node will operate in a strict slave mode with communications initiated by a master on shore. Full duplex communication can be established by the master controller using both low level control systems for special purposes like software upgrades at specific routers.

The low level control system is being implemented using the serial ASCII interchange loop (SAIL) protocol (ANSI/IEEE standard 997-1985), which is widely used in oceanographic instrument systems.

C. Time Distribution System

Accurate and precise time will be almost universally necessary for science experiments. This requires a reference standard and the distribution of a time signal across the seafloor network. In the US, the reference function is performed by the National Institute for Standards (NIST), who provides standard time, which is distributed through the global positioning system (GPS) satellites. For NEPTUNE, high accuracy time from the GPS system will be distributed over the low-level control fibers using a straightforward addition to the SAIL protocol. This allows high accuracy time to be transferred from the shore station to the nodes including correction for propagation delays on the backbone cable, and will provide synoptic time ticks everywhere on NEPTUNE. The accuracy goal of 1 µs is readily achievable.

For science users who do not require time at this accuracy, standard IP protocols are available for clock synchronization. The most widely used of these is network time protocol (NTP). Operating as a client in seafloor nodes and instruments, NTP is capable of clock synchronization with a few milliseconds accuracy across the NEPTUNE network, including correction for propagation delays. With the availability the time signal on the low-level channel described above, a node-resident NTP server might provide sub-millisecond accuracy. We are researching this area.

D. Physical Packaging

The seafloor nodes for NEPTUNE must be designed to facilitate the installation or removal of a large, diverse set of instrumentation. This requires the use of wet- mateable connector technology that is compatible with remotely operated vehicle (ROV) manipulators. Underwater-mateable electrical connectors for this purpose are readily available and have a good long-term reliability record. By contrast, underwater-mateable optical connectors are an emerging technology, and are less attractive from both a cost and a reliability perspective.
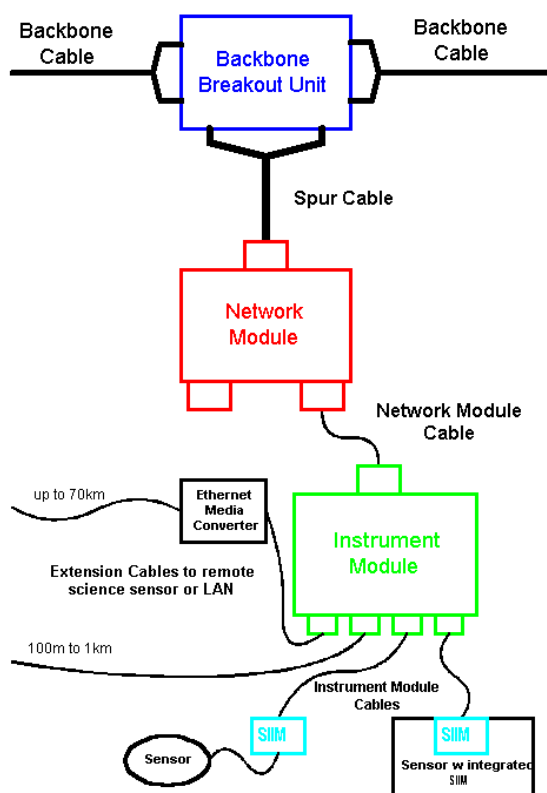
Figure 2. Sketch depicting the layout of a seafloor node.

The NEPTUNE backbone will be installed using conventional cable laying technology, and a major design goal has been minimization of the cost of maintenance of the seafloor plant. This precludes a design which requires a telecommunications-standard cable repair in the event of an electronic failure in one of the seafloor nodes. While every effort is being made to maximize NEPTUNE system reliability, it must be recognized that the seafloor installation of complex router and power system electronics will inevitably result in a lower MTBF (mean time between failures) than is achieved in submarine telecommunications systems, where the seafloor plant is extremely simple. Since extensive use of standard oceanographic ship and ROV assets in NEPTUNE is anticipated for science purposes, it makes sense to design the infrastructure so that it can be maintained using the same tools.

Figure 2 shows the layout of a science node. The backbone fiber optic cable contains an in-line backbone breakout unit (BBU) which is functionally identical to a conventional telecommunications system branching unit. The third connection to the breakout unit is a spur cable, which is 1.5 water depths long. The spur cable contains two conductors and twice as many optical fibers as the backbone cable, and serves to bring all of these connections into the network module. The BBU contains no active components, and should never require service that would necessitate use of a cable ship.

The network module (NM) contains the high (10 kV) voltage power supply and redundant backbone router equipment, along with the low level control and time distribution system. It is intended to be recoverable for maintenance or upgrading using conventional oceanographic research ship assets. This can be accomplished by disconnecting any attached instrument modules (IM), attaching a lifting cable to the unit with an ROV, and hoisting it with attached spur cable to the fantail of the research ship under dynamic positioning.

The instrument module (IM) contains the low voltage science instrument power distribution system, low (10/100 Mbit/s) speed data switches, and instrument control systems. It serves as the connection point for scientific instruments. The instrument module can either be located quite close to the network module or up to 100 km away, and more than one instrument module may be attached to a network module.

Both the network and instrument modules will be constructed like the junction box used for the Hawaii-2 Observatory (H2O) installation on the abandoned HAW-2 analog submarine telephone cable (see http://www.whoi.edu/science/GG/DSO/H2O/). Figure 3 is a cartoon depicting this ocean observatory. The junction box is about 2 x 1 x 1 m in size and is constructed entirely of titanium and plastic for corrosion protection. It contains two pressure cases to house the system electronics. An oil-filled manifold is placed about 1 m off the seafloor and houses a set of wet-mateable electrical connectors to which instruments may be attached. The H2O junction box is designed to be recoverable for servicing, and plugs into the telephone cable at a termination frame. For NEPTUNE, the IMs will be very similar to the H2O design. The NM will be permanently attached to the spur cable, and hence will have an attached gimbal and cable termination box to house the necessary connections.

E. Science Instrument Interface Module

A specification defining communications, power, timing, instrument control, and metadata characteristics for NEPTUNE-attached instrumentation is being developed. The NEPTUNE systems engineering team recognizes that support
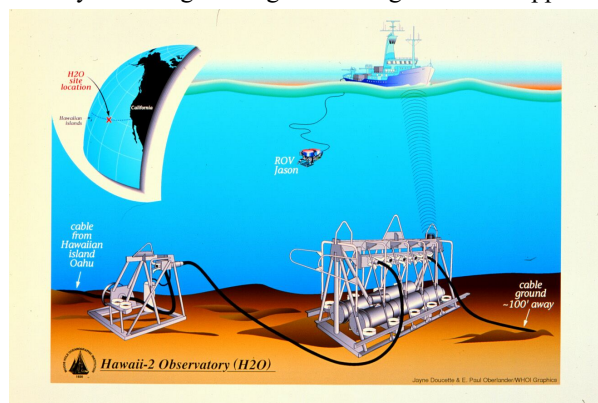
Figure 3 Cartoon showing the physical design of the Hawaii-2 Observatory. See text for discussion.

for existing instruments, which do not already conform to this specification, will have to be provided. Although it might be possible to connect instruments directly to NEPTUNE via ROV-mateable connectors on the IM Ethernet interface, the NEPTUNE project is developing Scientific Instrument Interface Modules (SIIMs) to meet this need.

SIIMs are based on work initiated at the Monterey Bay Research Aquarium Institute (MBARI) on a concept they call a Puck . As can be seen in Figure 2, SIIMs are separate units either inserted between the instrument and the IM or built into the instrument itself. There are two sides: the instrument side and the IM side.

One example of a SIIM would accept a serial connection from the instrument side and convert it to an Internet-compatible data stream (perhaps using the TCP/ IP telnet protocol) on the Ethernet side. Such a SIIM would contain an SNMP agent, which communicates with the shore-based Network Management System (NMS) to identify the instrument by serial number and report on instrument health. Equipment similar to these simple versions of SIIMs are commercially available today and can be adapted for subsea applications.

SIIMs are one-to-one paired with individual instruments and are capable of identifying the instrument to the NEPTUNE infrastructure via SNMP. Up-to-date metadata (calibration values, etc.) and software drivers, residing in the NEPTUNE data management system, are referenced via this identification number. In some cases the identical information will reside in SIIM non-volatile memory. In these cases NEPTUNE control software will be capable of reading and updating SIIM-resident information.

Different SIIMs will be developed for other instrument needs. SIIM features to be added might include voltage/power conversion, uninterruptible power supplies, different size serial data buffers, copper-to-fiber Ethernet conversion, etc. The specific list of features to be included in the first SIIMs will depend on a survey of the type, number, and needs of the instruments that will be initially deployed on NEPTUNE.

IV.   Next Steps

The technical requirements, constraints, and available technologies for building a plate-scale subsea, high-speed, data network to be used in the support of oceanographic research have been outlined. This network is modeled after existing computer networks and uses components similar to those used to build the Internet.

Further work using a software package called Opnet Modeler is underway to model network topology and test routing protocol configurations and other parameters to see how the network will respond to cable breaks, high bandwidth applications such as HDTV video transmissions, and increased sensor load over its year lifetime. A seven-node testbed network has also been built to compare the model results to reality.

Another effort is underway to run system-level reliability studies of the network design using available reliability data for COTS components that might be incorporated into the NEPTUNE systems. This should provide an indication of the level of redundancy necessary for the components.

Future studies are also planned to test the ability of ROV-mateable connectors to carry 10BaseT, 100BaseT, 1000BaseT, and 10GbE over both fiber and electrical conductors to be used in NEPTUNE.

REFERENCES

[1] Neptune Partners, 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 (NOPP) , University of Washington, Seattle, WA. http://www.neptune.washington.edu/pub/documents/feas_study_2000.html, June 2000.

[2] D. Williams, An Oversimplified Overview of Undersea Cable Systems , http://nicewww.cern.ch/~davidw/public/SubCables.html, 1997.

[3] T.W. Chung, J. Coulter, J. Fitchett, S. Mokbel, and Bill Bt. Arnaud, Architectural and Engineering Issues for Building an Optical Internet , http://www.canarie.ca/~bstarn/optical-internet.html, 1998.

[4] D.H. Rodgers et. al., NEPTUNE Regional Observatory System Design , Proceedings of Oceans 2001 Conference, Honolulu, HI., Nov. 2001.

[5] H. Kirkham, B. Howe, V. Voperian, P. Bowerman, The Design of the NEPTUNE Power System , Proceedings of Oceans 2001 Conference, Honolulu, HI., Nov. 2001.