DATA ACQUISITION FROM HETEROGENEOUS SENSOR NETWORKS: THE CASE OF NEPTUNE CANADA THE WORLD'S LARGEST CABLED OCEAN OBSERVATORY

B. Pirenne, Ocean Networks Canada, University of Victoria, BC, Canada

Abstract

Ocean Sciences is at the crossroads: it is entering the brave new world of "Big Science". The first of a new generation of large facilities, the NEPTUNE Canada cabled ocean observatory (www.neptunecanada.ca) will be presented from the point of view of a sensor network composed of hundreds of diverse instruments. The challenges we faced will be reviewed, together with the selected network design, data management and data distribution approaches. Special emphasis will be placed on the architecture of the system and on the more recent developments and concepts used to help scientists in their exploitation of the data. Finally a number of the early discoveries made with the new facility will be briefly described.

CABLED OCEAN OBSERVATORIES

Cabled ocean observatories are remote observing systems that provide power and communication media to a host of underwater instruments and sensors. Consequently, the instruments are (almost) always on-line and sufficient power is provided to the assets to ensure uninterrupted data flow covering multiple environmental parameters at high resolution in a four dimensional space. Observatory systems considered here also provide a significant ability to remotely manage their assets (ie, provide a real-time command ability for specific instruments). As an example, NEPTUNE Canada is composed of a fully redundant 800-km cable loop and has the ability to provide 9kW of power at up to 10 different locations of scientific interest. Figure 1 shows the layout of the NEPTUNE Canada observatory as well as its currently defined 6 main locations, five of which are instrumented. They reside at depth between 20 and 2700 meters.

Each of the locations is equipped with a “node” that reduces the line voltage of 10 kVDC down to 400 VDC and offer data connection points for up to 4 Gbps. In a area covering up to a few km^2, extension cables can be run from the nodes to sites of interest, where platforms with actual instruments and sensors are installed. The platforms are typically composed of a “junction box” whose role is to be the local “power bar”, providing plugs for instrument power and communication, converting the 400 V input to 15, 24 or 48 Volts and translating the instrument serial protocol to IP where necessary.

The instrumentation measures physical and chemical parameters of the ocean (temperature, salinity, oxygen content, CO_2, currents speed and direction at different depths, ...), but also has a number of more specific devices such as underwater video cameras, electro-magnetic experiments, vertical profilers that move through the water column, small vehicles on track (crawler), ... all of which would not be possible without the availability of ample power and the ability to command them in real-time. Figure 2 illustrates the crawler, itself a device equipped with various chemical and physical sensors, cameras, etc.

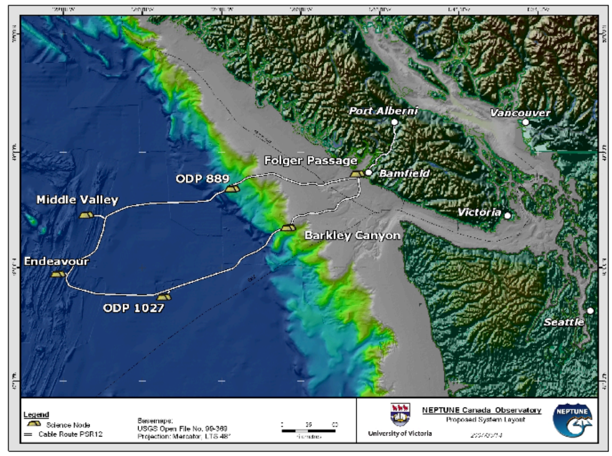

Figure 1: Map of the area covered by NEPTUNE Canada west of Vancouver Island. Please note the 800 km cable loop and the various location of scientific interest, and their "node".

The entire system represent the extension of the Internet under the Ocean, which was the vision put forward by the proponents of such a system many years ago.

Figure 2: A small tethered vehicle on track. It can roam within 50 m from its central position. It is equipped with various physical and chemical sensors and a camera.

NETWORK TOPOLOGY

The network design implements the vision of an Internet-based system, where every instrument and device is either a leaf on the tree structure or a junction point where multiple branches come together. The tree is of variable and arbitrary depth and does not impose conditions on its topology other than the fact that communication to other parts of the network will always propagate up the tree to the first common junction point between any two devices.

Network Design Considerations and Choices

To minimise the cost of the system and to re-use existing off-the-shelf technologies, the use of the Internet Protocol (IP) is preferred as a transport mechanism for data packets at the user/application level. Distances and fibre technology may require another transport mechanism at the lower level. So in this instance the ISO layer 1 can be implemented using fibre optics, lasers and repeaters, on which the SONET protocol will be running. SONET packets will encapsulate layer 2 Ethernet (802.3) packets and deliver them to their end-point thanks to this standard's addressing system. At that level, a traditional network is available for implementing data communication, transport, routing, security, etc.

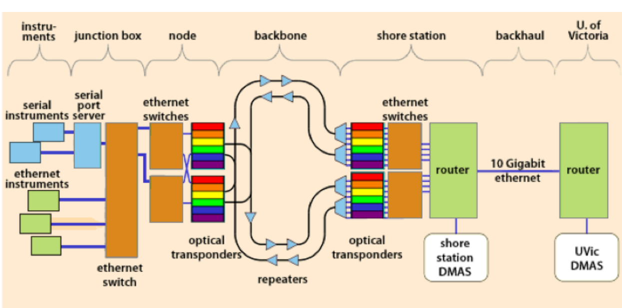

Figure 3: The example of the NEPTUNE Canada network design from a network topology point of view.

As indicated in figure 3 above, currently available oceanographic science instruments are of a legacy design, optimised for power consumption, internal recording and short stays in the water. Their typical data communication interface will be of the serial type (e.g., IEA RS-232, IEA RS-422 or IEA RS-485). To implement the vision of the observatory representing the extension of the Internet underwater, it is necessary to convert the communication protocol of the instrument to IP as close as possible to the instrument. This can be done with simple devices, typically called "terminal servers" enclosed either in the original instrument, in a can on the cable linking it to a junction box or within the junction box itself, often only metres away from the instrument.

To be complete, the structure must also accommodate multiple nodes at the same level, daisy-chained nodes; many junction boxes per node and daisy-chained junction boxes; instruments with piggy-back sensors; possibly multiple shore stations at the root of a network and finally also possibly several redundant data centres.

With a potential for thousands of individual instruments and devices attached to the network, as well as for ease of isolation of the system, it makes sense to select a non-routable set of addresses, as allowed by the IP protocol. In this case, given the complexity of the network, the familiar 10.0.0.0 address space (RFC 1918) was selected. It allows system managers and security analysts to only worry about a few selected bridges between the outside world and the private network, while allowing complete freedom of address allocation and division into VLANs etc. within the private domain.

Virtual Local Area Networks (VLAN – IEEE 802.1Q) offer service segmentation and will be the tool of choice if special categories of instruments need to be isolated from one another for security reasons. VLANs are a layer 2 feature. There are multiple examples that can be considered where VLANs use would make considerable sense in the set up of an observatory. The example of a separate management VLAN comes to mind where all non-user accessible devices will be isolated in a special management VLAN. Such devices will include all network devices on the system (on land as well as underwater) such as switches, routers, media converters, serial-IP converters; but also the facility control computers, precision clocks, etc.

Another VLAN that should be considered is one that will host all instruments that are considered of "national security concern" and would need to be especially protected or have a different management policy.

Timing and time signal

There is a requirement that all clocks on the system be synchronised with a master clock to ensure that all data have the same time baseline to ensure the ability to cross-correlate measurements from different sources. This requirement can be satisfied in a number of ways:

- convince instrument manufacturers to create smart instrument interfaces to periodically re-synchronise the internal instrument clock to the observatory's using the NTP or PTP protocol
- periodically and programmatically re-synchronise the instrument clocks through shore-based software
- time-tag all arriving measurements at the shore station.

Our current approach has been a combination of the first and third option so far, as most of the instrumentation in place is of a legacy, low-power, battery-operated type that is optimised for durability of deployment.

DATA ACQUISITION

A part of the complexity of ocean sciences stems from its plurality: an observatory such as NEPTUNE Canada is serving many different communities with different goals and relying on different types of instruments to achieve their goals: physical oceanographers and chemists will have sensors measuring directly phenomena of interest while biologists will usually rely on proxies to derive populations, species and abundances. This is reflected in the instrumentation that has to be hosted on the system.

Typical instruments will therefore usually fall into one of three categories from a data management point of view:

Table 1: Categories of data streams and instruments

Category	Instrument	Data Format
Scalar	CTDs, chemical sensors, ...	Return lists of values at regular intervals
Complex	ADCP, still cameras, ...	Return n-dimensional matrices on a regular basis
Stream	Video cameras, hydrophones	Return uninterrupted streams of bytes

For the purpose of designing a software system to manage the data flow coming from various devices connected to the infrastructure, a simple approach can be considered where all instruments are considered as sending a stream of data.

At the highest level of abstraction, given the individual duty cycles of each instrument, all categories will, from time to time, return their measurements as a string of bytes. A scalar instrument may be returning the values of its sensors every second for months on end; a still camera may be programmed to take a picture every day, a video camera may be operated periodically and return a rapid succession of images.

At the same time that each instrument can be considered as a producer of a more or less continuous stream of bytes, another way to look at the problem is to see every new stream of bytes as an event that just occurred and for which some specific processing is required.

We assume here a combination of both approaches to deal with the data flow: each instrument produces data in an ad hoc, not necessarily predictable fashion. The (a)synchronous occurrence of a new sequence of data will trigger the execution of a pre-determined set of processing stages, the last of which will be the archival of said stream.

Science Data vs. Engineering Data

Clearly science data collection is the primary goal of any ocean observatory. However, sensors and instruments are attached to an infrastructure that allows them to operate. The infrastructure typically provides power and communication media to instruments and their hosted sensors. So, unless the infrastructure is “somebody else's problem” (such as is the case when all or part of the infrastructure is contracted out to an external organisation, e.g., satellite data transmission), and regulated through a service level agreement (SLA), the organisation operating the facility has to perform and support a potentially significant number of activities having to do with the oversight of the entire system.

The oversight of the system is usually a 24x7 task that involves the monitoring of a large number of subsystems dealing with power and power distribution as well as with data transmission. All of those subsystems will contain sensors that produce engineering data. The engineering data has to be acquired, converted, verified and checked against ceilings and thresholds on a permanent basis. Any value identified as going beyond pre-set bounds will generate alerts to be dealt with by observatory personnel.

In the example of NEPTUNE Canada, nodes and junction boxes, distributing power and communication facilities to science instruments, are equipped with a large amount of electrical and environmental sensors. Such sensors typically return data at the rate of one Hz. It is estimated that the nodes and junction boxes currently connected on the NEPTUNE Canada network will alone produce about 8 TB of raw scalar data per year.

The data are however essential to help predict trends, offer the ability to conduct forensic analysis to understand why an element has failed, etc. An example where trending will help observatory managers extend the lifetime of the infrastructure and establish a priority list for maintenance and recovery is the analysis of the stability of the various ground leak current sensors. Indeed, in seawater, a complete isolation of any power conductor from seawater is essential to prevent corrosion. A slowly increasing leak current (or reduced resistivity to ground) is an indication that something is amiss somewhere and could lead to accelerated corrosion of subsystems. Switching them off early will increase the lifetime of the rest of the system.

Tools have thus to be provided to engineers and “wet plant” system managers to access, examine and react to events happening underwater. The large number of individual sensors that have to be monitored calls for systems that will automatically and constantly verify that all variables remain within their pre-set boundaries. A network management system (NMS) will collect all alerts that come from any subsystem (power or communication) and draw the attention of system operators when they occur. Automating such tasks is essential to limit the operating costs of the infrastructure to a minimum and to avoid the need for a 24x7 coverage of the operation of the system, limiting the service requirement to having personnel on call.

DATA ARCHIVE

Big Science infrastructure is typically designed and built to last between 25 and 50 years: astronomical observatories, large vessels, nuclear reactors, ... after which they have to either be decommissioned or to undergo significant refurbishment, upgrades and modernisation. The case of an ocean observatory is no exception, but will likely have a life expectancy towards the lower end of the range, mostly due to the lack of experience with such system as well as the harsh and corrosive environmental conditions to which the various elements of the infrastructure are subjected.

Consequently, with funding hopefully in place to support operations during the entire period, the software systems used to acquire and store the data, monitor and control the infrastructure should be sustained and provide access to the sum total of data, information and knowledge accumulated during the complete history of the facility.

This is one of the fundamental requirements of the software system in charge of the observatory and the reason why the underwater infrastructure does not "just" extend the Internet under the Ocean.

Table 2: Life expectancy of different elements of the System

Element	Longevity
High-level design, topology, external environment	Lifetime
Hardware Architecture	10-15 years
Programming language	10+ years
Operating Systems	10 years
Storage Technology	8-10 years
Design of the main software elements	7 years
Operational computers	4-5 years
Storage system	3-5 years

The numbers in Table 2 above indicate the expected life expectancy of the various elements of any large system and illustrate that throughout its lifetime, constant changes and update will have to take place to keep it operating efficiently and economically as, as is often the case, running an ageing infrastructure is more expensive than a timely adoption of new technologies:

- Old hardware will cost more and more to keep running (e.g., keeping lots of small disk drives in operation rather than a few large ones)
- Old software implementation (legacy software) may make it more difficult to find suitable developers who know about the language, OS, etc.
- Novel instrumentation design or radically different ways of using the underwater infrastructure might lead to the impossibility to continue operating with the assumptions that led to the elaboration of the system to that date. (Disruptive technologies).

OPERATION SUPPORT

A large underwater observatory has many physical components. It also represents a facility that has to have a long life time and will therefore host several generations of caretakers. The complexity is so large that it is impossible for a single person or small group of people to remember everything about the system. Examples of essential information abound: installation date and position, date of recalibration of an instrument and the formulae that have to be used for each of its sensors; when the instrument was turned on and off and by whom, ... This information is absolutely critical to understand the data that any instrument produces. Moreover, when dealing with a multiyear archive of data from instruments with a complicated history, understanding that history is necessary for data users to have some trust in the data quality.

The considerations above imply that the amount of information to be recorded, maintained and presented to users about any component of the observatory is tremendous and usually much more considerable than what casual observers would imagine.

DATA ACCESS

Traditionally, data access consists in providing search screens and a result download facility to users. A number of files are downloaded and have to be individually processed by the user, usually in isolation, with local resources and locally developed or installed software. This model no longer works for disciplines where the amount of data is multiplied by a large number of orders of magnitude while the amount of users remains constant. The model that is currently emerging involves a shift away from the search-download-process approach. The concept of Web 2.0 with its participatory approach is calling for something quite different where users use their web browsers to perform all activities related to the scientific process. Some of the differences are as follows:

- On-line collaborations with remote colleagues and students are the norm. Data volumes are so large and so multi-disciplinary that it is often necessary to seek out the support and advice of colleagues in different disciplines to support a particular project execution. The new collaborators may not be co-located and may work at different times but a "work space" is available for all members of a work group to perform all tasks from data search and examination all the way to the redaction of the final paper.
- Searching and sifting through data is done using other criteria and sources of information than previously available such as annotations provided by "crowdsourcing" activities and data from other observatories using interoperability concepts.
- There is little need to download data: data processing facilities on the Grid or in a computer

Cloud are available through privileged links with the archive. Data processing software libraries and templates are available to run against the data. Instead of downloading data, the new concept encourages the upload of new code to run on the server. New code can first be tested, refined and maybe later made available for all to use.

- With compute facilities becoming utilities, with storage capacity available on the network, there is no need to spend money and time maintaining one's own infrastructure. Shared infrastructures are always available at the other end of the high-capacity network.

SOME OF THE FIRST RESULTS

There is no space on such a summary paper to list, explain and illustrate the findings, discoveries and new knowledge acquired through a novel facility such as NEPTUNE Canada. So the author will refer the news posted on the observatory's home page for up-to-date information. The prospects for new findings are very important as such a system has never been built before, as the spacial, time resolution and accuracy of the measurements are increased by several orders of magnitude and that NEPTUNE Canada is supporting no less than five distinct science disciplines (ocean physics, chemistry, biology, plate tectonics and computer science and engineering). Moreover, it is opening the prospect of multi-disciplinary science discoveries.

REFERENCES

[1] C.R. Barnes, M.M.R. Best, and B. Pirenne. NEPTUNE Canada cabled ocean observatory: Final installation and initial results. APEGBC Innovation, March/April issue, pp. 30-33, in press.

[2] M.M.R. Best, C.R. Barnes, B.D. Bornhold, F.R. Johnson, P. Phibbs and B. Pirenne. Live data from the coast to the deep sea: NEPTUNE Canada. AGU Ocean Sciences Conference, Portland, February 2010.

[3] C. R. Barnes, M. M. R. Best, F. R. Johnson, P. Phibbs and B. Pirenne. The NEPTUNE Canada Project: installing the world’s first regional cabled ocean observatory. Chapter in book on Ocean observatories, P. Favali et al. (Eds.), Praxis Publishers, in press.

[4] B. Pirenne. Data Management and Archiving System for seafloor observatories: acquisition, archival, analysis, interoperability. Chapter in book on Ocean observatories, P. Favali et al. (Eds.), Praxis Publishers, in press.