

Network-enabled information fusion in support of underwater research

Alessandro Berni, Diego Merani, Michel Leonard

November 2007

Originally published in:

Proceedings of the 2nd International Conference and Exhibition on
Underwater Acoustic Measurements: Technologies and Results, 25-29
June 2007, IACM – FORTH, Heraklion, Greece.

About NURC

Our vision

- To conduct maritime research and develop products in support of NATO's maritime operational and transformational requirements.
- To be the first port of call for NATO's maritime research needs through our own expertise, particularly in the undersea domain, and that of our many partners in research and technology.

One of three research and technology organisations in NATO, NURC conducts maritime research in support of NATO's operational and transformation requirements. Reporting to the Supreme Allied Commander, Transformation and under the guidance of the NATO Conference of National Armaments Directors and the NATO Military Committee, our focus is on the undersea domain and on solutions to maritime security problems.

The Scientific Committee of National Representatives, membership of which is open to all NATO nations, provides scientific guidance to NURC and the Supreme Allied Commander Transformation.

NURC is funded through NATO common funds and respond explicitly to NATO's common requirements. Our plans and operations are extensively and regularly reviewed by outside bodies including peer review of the science and technology, independent national expert oversight, review of proposed deliverables by military user authorities, and independent business process certification.



Copyright © NURC 2007. NATO member nations have unlimited rights to use, modify, reproduce, release, perform, display or disclose these materials, and to authorize others to do so for government purposes. Any reproductions marked with this legend must also reproduce these markings. All other rights and uses except those permitted by copyright law are reserved by the copyright owner.

NOTE: The NURC Reprint series reprints papers and articles published by NURC authors in the open literature as an effort to widely disseminate NURC products. Users should cite the original article where possible.

NETWORK-ENABLED INFORMATION FUSION IN SUPPORT OF UNDERWATER RESEARCH

Alessandro Berni, Diego Merani, Michel Leonard

NURC, Viale S. Bartolomeo 400, 19126 La Spezia, Italy

Abstract: *Advanced communication and information processing techniques can be applied to underwater research to enable new ways of enhancing real-time experimental outcome. This paper is concerned with three main areas where network-enabled capabilities can be successfully applied: a) network infrastructure, b) information infrastructure, c) context-based data management. Network infrastructure relates to the networking infrastructure required to support sea trials, including 2-way always-on satellite links to enable distributed processing, low-cost asymmetric satellite communications systems based on DVB-S to provide unidirectional broadband services, and advanced quality of service mechanisms to allow seamless Internet Protocol operations across heterogeneous media. Information infrastructure relates to set of techniques used to collect, process, store and distribute data acquired during sea trials, including self-synchronization between web services. Context-based data management relates to interoperability aspects between heterogeneous information systems, including metadata, information identification, data reduction using genetic sequencing algorithms and the possible application of multilevel security environments, to enable information sharing in a seamless and secure way. The Dynamics of the Adriatic in Real Time (DART) sea trials, conducted between NURC and 26 international partners in 2006, will be used as a test case to describe the networking and information infrastructure implemented by NURC to fulfil scientific requirements in terms of increased collaboration and interoperability with external research partners, information sharing and guaranteed quality of service.*

Keywords: *Environmental variability, adaptation and information fusion, Network-Enabled Capabilities, Distributed Data Management and Processing*

1. INTRODUCTION

In this paper we will use the DART06 sea trials, conducted in the Adriatic Sea between NURC and 26 international partners, as a test case to present NURC networking and information infrastructures and the benefits they introduce in terms of increased collaboration and interoperability with external research partners. Such trials were performed in 2006 as part of a multi-institutional program addressing observational and modelling methodologies on small-scale instabilities in a marginal sea, producing as a secondary product a comprehensive data-model set of ocean and atmosphere properties.

From a scientific perspective the Adriatic Sea provides an interesting natural laboratory, in consideration of its wide range of environmentally-driven processes (e.g. as the result of bathymetry and coastline, wind forcing, water masses and currents, river

outflows and plumes, sea floor) and of the conspicuous accumulation of knowledge available on this area.

The methodology chosen for the sea trials was to evaluate and combine different observational and modelling capabilities in a network-enabled concept of operations, comparing in near real-time operational data produced by partners connected via the Internet (meteorological, wave and ocean models, at variable resolution) with local data observations (buoys, drifters, moorings, ship and satellite-based). This required the exploitation of existing NURC data communication architectures, as well as the implementation of ad-hoc configurations and data exchange methodologies capable of meeting demanding scientific requirements within the given technological constraints.

1.1. Observational and modelling requirements

Requirements for near-real-time communications have been derived from the quantity and quality of observational and modelling tasks planned during experimental activities.

As an initial step, all activities that required network communications were categorized in seven broad categories, each with its own peculiarities in terms of data volume and time sensitivity. The following list provides a broad overview of experiment participants and their respective contribution:

- In-situ observational tasks: bottom-mounted ADCPs; SEPTR; meteo stations; wave riders; CTDs; optical surveys; drifters.
- Remote sensing observational tasks: SeaWIFS real-time ocean color; AVHRR radiometry; RADARSAT Synthetic Aperture Radar.
- Atmospheric modelling: real-time COAMPS model from NRL-SSC and NRL Monterey; COAMPS-OS run on the NURC Cluster; real-time NOGAPS model from FNMOC; real-time ALADIN meteo model from EPSHOM (France); ALADIN and LACE meteo models from Croatia Meteorological and Hydrographic Service; real-time LAMI meteo model from ARPA (Italy); ECMWF meteo model from INGV (Italy).
- Ocean modelling: real-time and hindcast NCOM and HYCOM models from NRL-SSC; HOPS model from Harvard Univ, Univ. of Ancona and INGV (Italy); real-time and hindcast ROMS model from CNR Venice and ARPA (Italy); real-time POM and OPA models from INGV; GOTM 1D vertical model from CNR Venice (Italy).
- Wave modelling: real-time SWAN model from NRL-SSC (US) , ARPA (Italy); real-time WAM model from ARPA (Italy) and Univ. of Athens (GR)
- Surface drift modelling: NCOM model prior to drifter launch from CNR-ISMAR (Italy); variational assimilation of Lagrangian data from drifters, RSMAS Univ. of Miami (US), CNR-ISMAR (Italy); hindcast Lagrangian Kalman Filter Data assimilation from Univ. of North Carolina and Univ. of California, (US)
- Tidal modelling: PC-TIDES from NRC-SSC (US).

In more general terms, three major “information domains” were identified, corresponding to logical areas where information was produced and processed, and from which information was exchanged with other logical regions using bandwidth-limited communication channels.

One domain was associated to NATO’s Research Vessel NRV Alliance, providing the seagoing laboratory for the execution of the DART06 trials, connected to NURC with a VSAT satellite connection. Conversely, another domain was associated to NURC, which

acted as a central gateway for the execution of the experiment, ensuring seamless communication with NRV Alliance, with Centre scientists involved in supporting the trials and with partners participating from remote locations. The third domain referred to experiment contributors connected over the Internet.

The following Figure 1 shows the various information domains, and the corresponding data flows. It should be noted that all traffic patterns were forced through the GEOS Data Fusion Servers installed at NURC and onboard NRV Alliance. This choice proved to be essential in providing the Quality of Service required by scientific objectives.

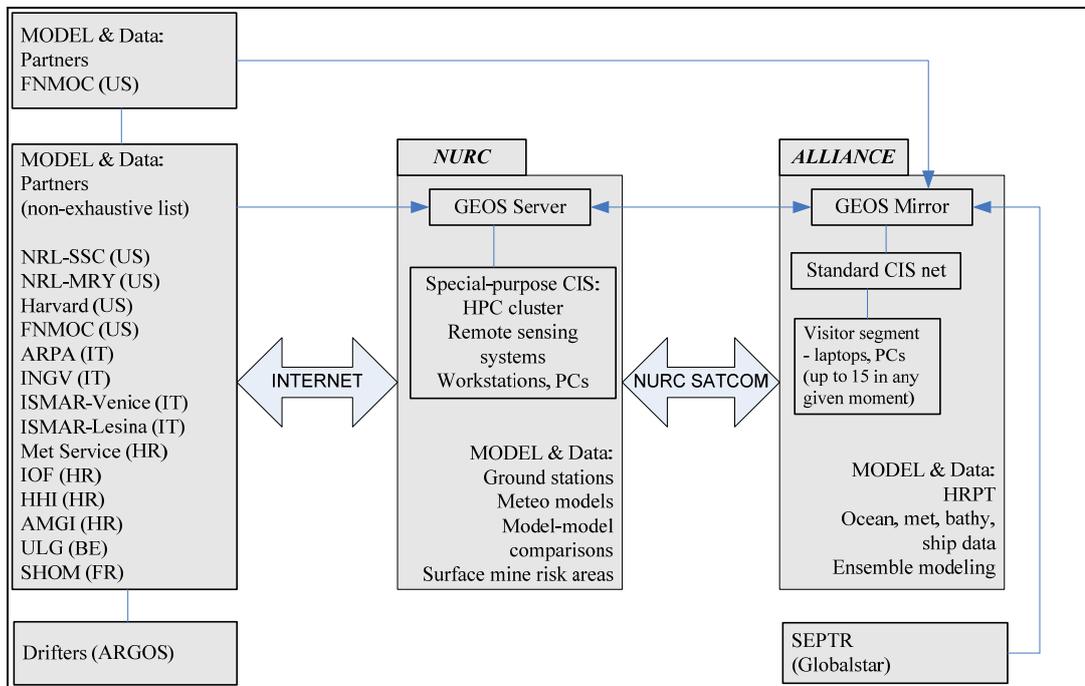


Fig.1: Summary of data exchange flows during the DART06 sea trials

A further degree of summarization of the requirement led to the estimation of daily data traffic between the different domains, as summarized in Figure 2.



Fig.2: Estimation of data volumes to be exchanged daily during DART06 sea trials

Communication capabilities normally available on maritime platforms cannot meet such a demanding requirement. Table 1 provides an optimistic estimation of the data volumes that can be supported by commonly available systems. The estimation is optimistic in the sense that it is based on the assumption of 100% efficiency and uptime, where the actual figure is normally lower because of protocol overheads and inefficiencies, rain attenuation, tracking problems with high sea states, interference and antenna shadowing from shipboard infrastructures. In addition to that, the affordability and cost-effectiveness of the various solutions may be questioned, where all services listed in Table 1 are charged by time or by volume.

<i>System</i>	<i>Data Rate [bit/s]</i>	<i>Daily data volume [MB]</i>	<i>Price per MB [Euro]</i>
Iridium	2400	24.72	75
Globalstar	9600	98.88	8
Inmarsat-B HSD	64000	659.18	18
Inmarsat Fleet 77	128000	1318.36	14

Table 1: 24-hour data transfer capability of common maritime systems

2. INFORMATION DOMAINS AND NETWORK INFRASTRUCTURE

NURC has been experimenting near-real time communications during sea trials since 1993, first using cellular phones and then Inmarsat-B HSD, at the data rate of 64 Kb/s. In 2001 a satellite communications facility was installed to provide always-on connectivity between NURC and its research vessels at data rates up to 1 Mb/s. The availability of such a flexible and cost-effective capability has driven the development of network-enabled sea trials, where experimental output is amplified by the continuous interaction between researchers operating at sea and their colleagues ashore.

In spite of this big improvement, a significant gap still exists between the capabilities provided by terrestrial networks and those made available to maritime nodes. As a consequence, choices on the placement of boundaries between the different “information domains” have been driven by the underlying network infrastructure, accounting the specific characteristics of maritime communications: lower data rate, higher latency, lower degree of reliability and availability. To ensure separation between different user communities, different Communities of Interest (COI) were identified within each information domain. As an example, the information domains at NURC and onboard NRV Alliance were actually divided into three COI, each with its own level of security and trust, to which specific communications policies could be applied.

2.1. Ship-shore links

With the present degree of maturity in commercial-off-the shelf (COTS) systems research vessel operators are given the possibility of increasing their efficiency in the execution of sea trials at an affordable cost. On a local scale, simple Wi-Fi networks or more advanced mesh networks can be used to provide fast and reliable wireless links with neighbouring vessels and sensor buoys at ranges greater than 10 nmi.

For wide area networking, satellite communication is a mandatory choice, although the target performance level for the choice of the system will be driven by technical and financial trade-offs. Small platforms, such as drifting buoys, do not allow for the installation of a large antenna dome, and are therefore good candidates for the installation of compact satellite telephone systems, such as Iridium or Globalstar, with data rates in the order of 10 Kb/s.

Larger platforms, instead, allow for the installation of tracking antennas capable of supporting higher data rates. Three main options are available: on-demand connection (e.g. Inmarsat-B HSD), VSAT permanent connection (e.g. private network as in the case of the NURC SATCOM Facility [1], shared service as the NSF-sponsored HiSeasNet [2]) or hybrid asymmetric connection with Digital Video Broadcasting (DVB-S) downlink, as the one provided in the Wired Ocean project supported by the European Space Agency.

Since all three technical options were already available onboard NRV Alliance, our choice for the DART06 trials was to rely on all of them, to maximise networking capabilities made available to the scientific party, both in terms of performance and of reliability, providing multiple backup options in case of failure of one system.

The main communication capability was provided through the NURC-Alliance VSAT link, with an increase of the baseline data rate from 128 Kb/s to 369 Kb/s to deal with the high volumes of data that were foreseen.

An additional communications capability was provided through a DVB-S hybrid architecture, as discussed in the following section.

2.1.1. Supplementing shipboard data links using DVB-S broadband Internet

DVB-based broadband satellite links operating in *push* or *on-demand* mode can be used in a maritime context to address bandwidth-intensive requirements, providing a cost-effective complement (in the order of 1 Euro per MB) to existing onboard communications capabilities [3]. The onboard system architecture includes two basic components: the uni-directional DVB-S channel, implemented with a satellite TV receive-only antenna with DVB/IP receiver, and the return link used to request on-demand traffic.

The return link, if required at all, can be provided using many different technical options: the baseline requirement is just to be able to access the Internet or the private network to which the land-based gateway is connected. In the NURC setup the VSAT system was used to provide the return link, although an Inmarsat-B HSD link was kept on standby to provide an alternative in case of failure of the former.

Configuration of the onboard gateway is very simple: Connection Manager Software is used to control the TV antenna and to establish a Virtual Private Network (VPN) connection with the land-based gateway through which subsequent data requests are routed. Once the request is received by the land-based gateway, the response is encoded in a DVB-MPEG2 stream and transmitted to the satellite transponder. On the receiving side, a DVB reception card installed on the onboard gateway extracts and decodes the packetized elementary stream into the original IP stream, which is then passed through the TCP stack of the operating system to the application that made the request (e.g. Internet browser, FTP client software).

2.2. Dynamic Bandwidth Allocation and Quality of Service

The NURC SATCOM Facility, based on the ND-SATCOM SkyWAN system, provides a fully meshed Multiple Frequency Time Division Multiple Access (MF-TDMA) network with bandwidth on demand and Internet Protocol (IP) local area networks connectivity.

All stations participating to the network can be connected via multiple satellite channels, and are capable of transmitting and receiving on a given channel making full use of channel bandwidth (although in different timeslots, in accordance with TDMA principles) with half-duplex communication.

Unlike in DVB-RCS systems, which are based on central hub (thus providing 2-hop communications), all stations of the NURC network incorporate all the features required for 1-hop point-to-point full-mesh networking, where the elimination of the central hub drives a 50% reduction of communication latency. All nodes in the network have the same status, with the exception of a master station broadcasting the Single Reference Burst containing the transmission plan to be used by all stations for transmission.

Large bandwidth requirements combined with limited channel capability and with high communications latency, where round trip time (RTT) perceived by users and applications is in the order of 500 milliseconds and up to 5000 milliseconds during periods of congestion, outlines a definite requirement for planning and implementation of effective Quality of Service (QoS), provided through dedicated QoS appliances.

Such appliances are typically conceived to support full-duplex links and their application to half-duplex links is somehow cumbersome, where in order to deliver effective quality of service, a *priori* knowledge would be required of the intentions of the remote stations participating to the network. In the absence of such a mechanism inside the network (which would in any case introduce an additional overhead for the node-to-node signalling) we chose to allocate outgoing bandwidth on each node on the basis of heuristics derived from traffic predictions and previously gathered statistics.

The appliances used during the trials (Peribit SR-50 and SR-20) allowed the definition of outbound guaranteed bandwidth and outbound maximum bandwidth both for the full link as well as for a large number of service classes within that link. In our QoS configuration policy we enforced the following stipulations:

- The sum of guaranteed bandwidth assigned to service classes cannot exceed 80% of the guaranteed bandwidth defined on the overall link. This is required to ensure that guaranteed bandwidth is always available for local system resources, such as SNMP updates and management traffic.
- Excess bandwidth is distributed between classes using Weighted Fair Queuing (WFQ) allocation, according to a priority level given to each class.
- Traffic flows (“applications”) are identified by IP source/destination addresses and TCP source/destination ports. Applications are statically assigned to service classes.

A software tool has been developed within NURC to facilitate the definition of guaranteed bandwidth, excess bandwidth and service classes on each outbound (unidirectional) link. This satellite capacity planning tool not only allows proper quality control while defining the QoS configuration for a certain trial, but greatly facilitates fine-tuning and even major reconfigurations during the execution of a trial.

As an example, we indicate in Table 2 the service classes that were defined during the DART06 trials, in order of decreasing priority in the allocation of excess bandwidth.

<i>Class name</i>	<i>NURC → Alliance [%]</i>	<i>Alliance → NURC [%]</i>
Scientific critical	40%	20%
Scientific standard	20%	10%
Network management	10%	5%
Ship management	5%	5%
DVB-S uplink	5%	40%
Default	0%	0%

Table 2: Guaranteed Bandwidth repartition during DART06 trials (sum adds to 80%)

It is important to note the difference in the allocation of guaranteed bandwidth in the two directions. This asymmetry is derived directly from the observational and modelling requirements discussed in Section 1.1 where data received onboard NRV Alliance was three times larger than data sent from Alliance to NURC and to Internet partners. Different

requirements can similarly be satisfied by changing the relevant parameters in the capacity planning software tool.

3. INFORMATION INFRASTRUCTURE

The networking infrastructure discussed in the previous sections provided the necessary building blocks to ensure the transfer of mission-critical data within the given time constraints. Additional steps were however required to ensure the seamless distribution of information between partners, enabling scientific collaboration and dissemination of interim results during the execution of the trial.

In previous work performed at NURC a central role was played by the so-called GEOS server, a geospatial information system providing the centralized platform where data fusion could be performed in near-real time during the execution of a trial. The size of the scientific effort associated to the DART06 trials, in terms of size of data collected and processed (fig. 2) as well as in the number of international partners, did not match well to this approach, suggesting the installation of a local GEOS server onboard NRV Alliance.

With two servers in place the majority of data communications between ship and shore can take place in the form of structured server-to-server transmissions, synchronized in accordance with the priority associated to a certain data type, enabling the fine-grained quality of service that increases assurance of meeting mission requirements.

The two servers were installed in two different domains of trust: onboard Alliance on the protected network segment used by the NURC team and at NURC on a DMZ segment accessible by Internet partners. This difference precluded the use of a peer-to-peer 2-way replication scheme, in favour of the use of FTP sessions initiated from the trusted side.

A coordination effort was performed prior to the execution of the trials, defining a structure of folders and subfolders where contributors could post their contents. FTP was performed using the SmartFTP [4] package, configured for near-real-time synchronization of time-sensitive datasets, while the rest of the traffic was scheduled during the night.

Network routing for the server-to-server communication was primarily across the VSAT link for both uploads and downloads, with the possibility of using the DVB-S link for download only. The rationale was to use the VSAT link for mission-critical traffic only, while keeping the DVB-S link for large, non-critical downloads.

4. CONTEXT-BASED DATA MANAGEMENT.

4.1. Data management considerations

Interoperability between different data formats is essential to enable distributed processing between different partners and correlation with historical data. Prior to the trial an agreement was therefore reached on the formats to be used for model results, e.g. GRIB (atmospheric models), NetCDF (ocean, atmospheric and wave models) in compliance with the COARS/CF convention [5], or ASCII compressed format, using absolute timestamps in modified Julian day.

As a general rule, the adoption of compressed formats, such as GRIB, is a wise measure towards efficient use of network capabilities. Additional improvements (ranging from 15% to 75%) were achieved during the trials by enabling the data reduction features of the Peribit appliances, based on Molecular Sequence Reduction (MSR) [6], to identify

repeated patterns inside the data stream transmitted over the VSAT link and to replace them with “shorter” labels, thus saving precious network bandwidth.

5. CONCLUSIONS

The execution of the DART06A trial (1-27 March 2006) led to the identification of some points worth being noted. First of all, a good match was observed between predicted traffic and actual traffic: NURC-Alliance traffic on the VSAT link was 32 GB, Alliance-NURC traffic was 6 GB, and DVB-S traffic was 4 GB. This translated in a daily average exceeding 1.5 GB, in line with preliminary estimations summarized in figure 2.

Second, accurate definition of data flow and replication scheduling allowed optimal use of the VSAT link, within the quality of service and guaranteed bandwidth constraints required to ensure proper support to near-real-time operations. It should be noted that savings on VSAT bandwidth thus obtained, coupled to unused bandwidth on the DVB-S link, allowed for the creation of “reserve bandwidth pool”, which could have been used to cover unplanned transmissions.

Third, by enforcing control on all phases of the scientific operations it was possible to meet ambitious scientific objectives at a very competitive cost. As an example, the cost for transmitting 1 MB of data during the DART06 trials was 6 Eurocent for the 369 Kb/s VSAT link and 40 Eurocent for the 1 Mb/s DVB-S link.

Above results were validated during the DART06B trial (14-31 August 2006), which confirmed DART06A data and enabled the fine-tuning of operational procedures.

6. ACKNOWLEDGEMENTS

Authors are grateful to Dr Michel Rixen, Chief Scientist for the DART06A and DART06B sea trials, for the support provided in the design and execution of the trial, as well as in the preparation of this paper.

REFERENCES

- [1] **Berni A.**, *The SACLANTCEN Satellite Communications Facility: system setup and operations*, SM-404, NATO SACLANT Undersea Research Centre, Italy, 2001
- [2] **Berger, J.**, HiSeasNet: a global network for oceanographic ships and observing platforms, In *Local to Global Data Interoperability-Challenges and Technologies*, pp. 4-8, 2005
- [3] **Andersen B.R. et al.**, A DVB/Inmarsat hybrid architecture for asymmetrical broadband mobile satellite services, *Int. J. Satell. Commun. Network.*, 24:119-136, 2006
- [4] <http://www.smartftp.com>
- [5] <http://www.cgd.ucar.edu/cms/eaton/cf-metadata/>
- [6] **Subathra S. et al.**, Performance Analysis of Dictionary based Data Compression Algorithms for High Speed Networks, In *IEEE INDICON 2005*, pp 361-365, 2005

Document Data Sheet

<i>Security Classification</i>		<i>Project No.</i>
<i>Document Serial No.</i> NURC-PR-2007-016	<i>Date of Issue</i> November 2007	<i>Total Pages</i> 8 pp.
<i>Author(s)</i> Berni, A., Merani, D., Leonard, M.		
<i>Title</i> Network-enabled information fusion in support of underwater research.		
<i>Abstract</i> .		
<i>Keywords</i>		
<i>Issuing Organization</i> NURC Viale San Bartolomeo 400, 19126 La Spezia, Italy [From N. America: NURC (New York) APO AE 09613-5000]		Tel: +39 0187 527 361 Fax: +39 0187 527 700 E-mail: library@nurc.nato.int