

A Delay-Tolerant Networking Framework for Mobile Underwater Acoustic Networks

Ian Katz

Department of Computer Science
University of New Hampshire, Durham, NH 03824
Email: ikatz@cs.unh.edu

Abstract—In the operation of mobile underwater acoustic networks, protocols designed for traditional wired and RF wireless networks perform poorly or not at all. Successful operation underwater - an extremely unreliable medium - requires protocols that can adapt to their environment without relying on the network to send configuration messages.

In order to support such protocols, we propose a formalized augmentation of the OSI model of networking, making use of the data that is generally discarded or not extracted from messages as they move between the individual protocol layers. Our model captures and harnesses this data in order to make sophisticated and relevant judgments about the state of the network, and to create feedback loops that can adapt each layer’s protocol to the current environment.

We further propose a Delay-Tolerant Networking framework that incorporates our augmented model, providing net-centric applications a more interactive networking API than TCP/IP network sockets. This interactivity enables the mobility of nodes to benefit rather than inhibit the performance of the network, and provides a foundation for autonomous creation of ad-hoc mobile underwater communication infrastructures to meet an application’s immediate needs.

I. INTRODUCTION

Applying Metcalfe’s Law in the domain of underwater acoustic networks, the ability to establish reliable communications between sensors, buoys, and underwater vehicles deployed in the world’s oceans would exponentially increase their value. The development of viable underwater networking technologies would enable significant advancements in oceanography, homeland and port security, assisted navigation, exploration, monitoring, and surveillance. One promising method of achieving success in this area is the use of sonic transducers to create underwater acoustic networks (UANs). However, the benefits of this technology bring new challenges, many of which remain largely unexplored.

Unlike radio waves, sound waves propagate easily through water. However, water is a poor transmission medium for a variety of reasons, including multipath distortion, shadow zones, unidirectional communication, and half-duplex operation [1]. With acoustic signal propagation speed five orders of magnitude slower than

radio waves, very low data rates, high incidence transmission errors, and high energy consumption, water is unfit to support the typical approaches employed by wired and radio-frequency wireless networking.

Furthermore, existing UAN concepts concentrate either on point-to-point communication between an operator and a single autonomous vehicle or on a network of fixed sensors. In the presence of mobile nodes such as autonomous underwater vehicles (AUVs), an efficient UAN depends on protocols that can adapt to their environment and react to topology changes without saturating the network with configuration messages.

II. UNDERWATER ACOUSTIC NETWORKING

The OSI networking model presents a transparent interface for sending messages synchronously between nodes in a network, accomplished by a series of 7 protocol layers that prune unnecessary data (protocol headers and trailers) from messages as they travel from the Physical layer to the Application layer. In its most popular implementation, TCP/IP, only 5 layers are present (Physical, Link, Network, Transport, and Application). Conceptually, there are many barriers - present on all layers of the TCP/IP model - to its adoption in the underwater domain.

On the Physical layer, water itself puts limitations on the network in the form of low bandwidth, unpredictable channel quality (loss), and a signal latency that can vary greatly between nodes. On the Link layer, no single Media Access Control (MAC) protocol is effective in all possible node configurations. (For a detailed explanation of this, see Appendix.) Network layer algorithms rely on negligible packet loss to operate effectively; the mathematical models on which they are based have no provisions for the significantly high losses present in water. Finally, Transport layer protocols require low end-to-end latencies, low round trip times, and constant connectivity in order to reliably control the delivery of data. For a more detailed discussion of the effects of a high-latency environment on TCP/IP networking, see “Delay-Tolerant Network Architecture: The Evolving Interplanetary Internet” [2].

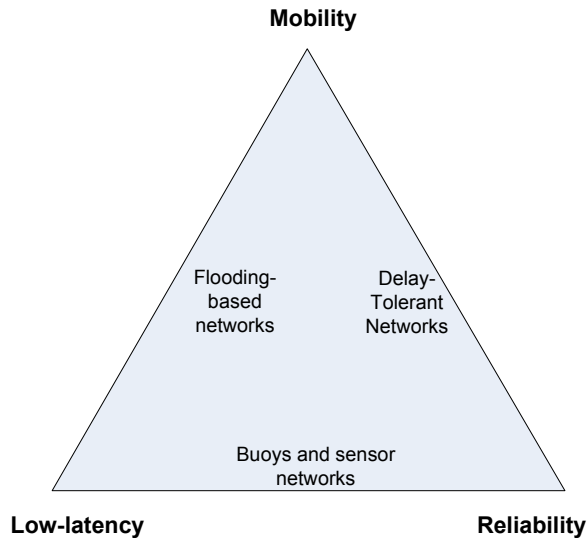


Fig. 1. Conflicting goals of mobile UAN architectures

The addition of mobility to underwater networks adds yet another complication to protocol development.

Conflicting goals of mobile UAN architectures illustrates the inherent trade-offs between the goals of mobility, low latency, and reliability are shown in Figure 1. Although many ad-hoc UAN protocols exist, none achieve all three goals. Most frequently, mobility is sacrificed or ignored in order to enable networking techniques more familiar to traditional networking. The purpose of illustrating these compromises is not to set the stage for a protocol that can achieve all of these goals, nor to choose a single set of goals to pursue. Rather, it is a demonstration of the need to be adaptable in order to participate in a network where each node may have a different (and possibly variable) set of goals.

The concept of adaptability in a mobile and high latency environment is not new; the Delay-Tolerant Network Architecture (DTN), developed for an interplanetary network of space satellites, suggests that some adaptation is necessary to balance reliability, mobility, and delay. The architecture also describes a layer above the Transport layer called the “Bundling layer” to manage disruptions in end-to-end connectivity. The Bundling layer effectively divides a network into uniquely-addressed regions of low latency, and suggests a routing mechanism to move “bundles” of data between them. This approach provides Transport layer endpoints within each region to allow low latency intra-region exchanges, while at the same time allowing inter-regional delivery despite its much higher latency [2].

One aspect of the Bundling layer prevents the adoption of this DTN Architecture in UANs: it addresses

interrupted connectivity between regions, not between nodes. Furthermore, the Bundling layer regards node mobility as a source of semi-predictable interruption (orbits and obstruction by planets), not as the ability for a node to move between regions of the network. In other words, the DTN model assumes that nodes will remain in only one region - which does not hold for mobile UANs.

Our aim is neither to develop a standalone protocol, nor a cross-layer protocol stack that will itself solve the problems of reliability, latency, and mobility. Rather, we will introduce an augmented model for networking that can adapt to its environment on all layers, a framework for Delay-Tolerant Networking between individual nodes in the network, and a more robust networking API. We will establish that protocols developed within our model will be better equipped to meet the challenges posed by the underwater environment. Finally, we will illustrate how our networking API will be better able to balance the network’s reliability, latency, and mobility based on the needs of an application.

III. THE AUGMENTED LAYER MODEL

It is more efficient for traditional networks to send all necessary configuration data in message form, since their available bandwidth rivals their available processing power and loss is negligible. In contrast, due to the low bandwidth available in UANs, rigorous analysis of the packets containing messages is more timely and reliable than waiting for additional messages to arrive. We use such analysis to form feedback loops between the inbound and outbound data on each layer, as well as between the Application layer and the other individual layers.

In our model, shown in Figure 2, we adopt the layers used by the TCP/IP simplification of the OSI model: the Physical, Link, Network, Transport, and Application layers. Beside these layers, we add four data classifications: Perception, Intuition, Management, and Adaptation.

Perception refers to the set of data on each layer, carried implicitly or explicitly, that can be used to build a more accurate view of the network’s immediate state. Perception data is made available as a shared database [3] in the Perception Plane. Intuition refers to the information that can be inferred from the combination of data in the Perception Plane with data from higher layers, computed either synchronously or asynchronously with respect to message arrivals. Adaptation is the collection of parameters for each layer’s protocol that may be adjusted or “tuned” [4], either by cues from the Management plane or automatically by the layer itself. Finally, the Management Plane is a second

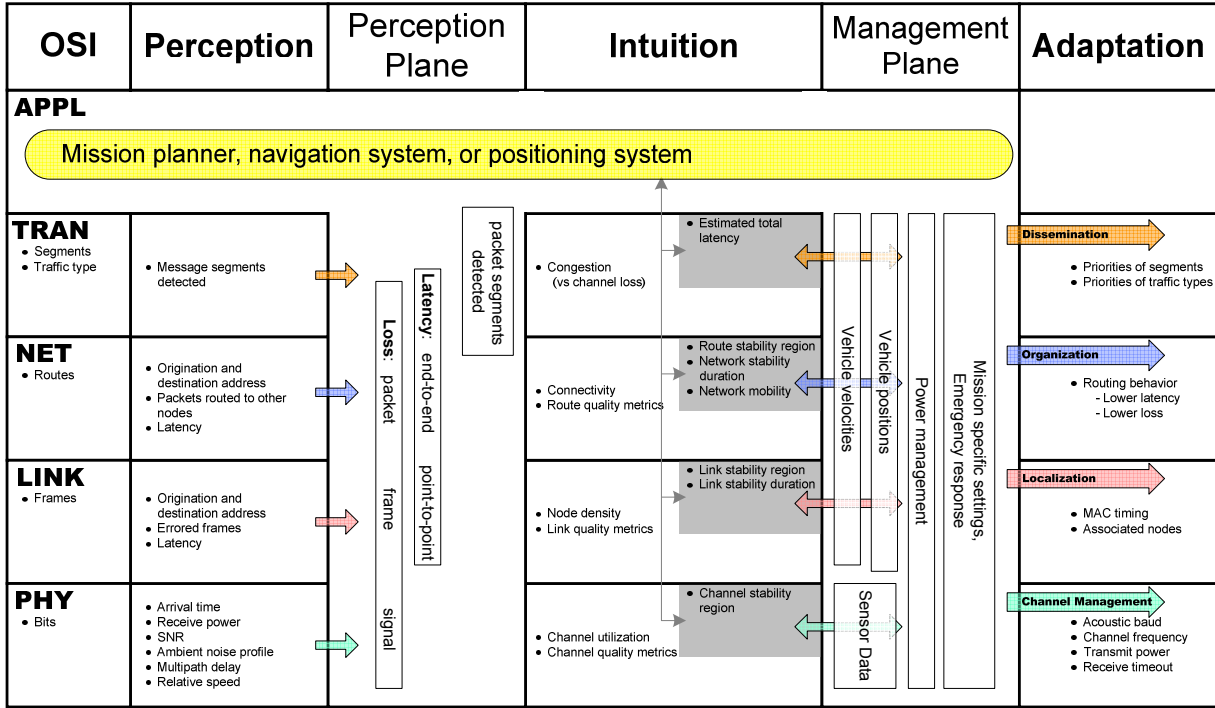


Fig. 2. An augmented layer model for adaptable protocol layers

shared database representing the relevant data from higher layers that may aid Intuition and Adaptation.

A. Flow of Information

Perception originates from incoming signals on the Physical layer, providing data left over from signal processing: the receive time, gain, distortion, and signal-to-noise ratio. The Link layer and Network layer Perception provide the source and destination addresses for all detected frames, as well as loss and latency measurements between them. The Transport layer provides a list of the message segments that have been detected.

Perception data from the Physical layer can be combined with node positions from the Management Plane, producing Intuition data: regions of exceptionally strong or weak signal quality, and overall channel utilization. Similar Intuition calculations are produced in each of the higher layers: link quality metrics and node density on the Link layer, route quality metrics and network connectivity on the Network layer. More sophisticated calculations can be made using node velocities, such as estimates of the window of time during which the current links and routes will be operational.

In the outgoing direction, Adaptation reads data from the Management Plane, when needed, to make informed judgments on the ability of a particular layer to meet the immediate needs of a message being sent onto the network. Possible examples are determining the order in which to send individual message segments on the

Transport layer, choosing a routing behavior on the Network layer, sending an emergency communication on the Link layer, or changing the communication parameters of the Physical layer device (the acoustic modem).

B. The Framework and API

In this domain, it is important not to ignore mobility - both active and passive. Some nodes in the network may control their mobility, enabling network behaviors such as searching for network connectivity, or even building precisely the network that is needed at any point in time. This framework provides a method of balancing a node's mobility with its requirements for reliability and latency.

To accomplish this, we require applications to indicate the reliability and latency requirements of a message at send time. Individual layers are responsible for deciding whether they can use Adaptation to meet these requirements. We use a simple scheduler for outgoing messages, designed to deliver them when possible and queue them otherwise.

In order to support possibly delayed delivery, the API for sending messages is asynchronous. Outgoing messages are assigned a unique identifier, which is immediately returned to the application. The framework then queries each layer, in order to determine whether the message is deliverable within its specified bounds. If all the layers respond affirmatively, the message is

sent through the layers both to be wrapped in protocol information and so that Adaptation may be effected. When it reaches the Physical layer and is sent onto the network, the application is notified of its success.

Messages that require a more reliable delivery than a layer is able to provide are queued, where they will remain until either reliability improves or they exceed their limit on latency. Methods for aggregate analysis of messages on the queue are provided by the framework to the Application layer, enabling it to determine the most rewarding method of re-establishing connectivity. For example, the destinations of unsent packets can be correlated with the nodes' positions to determine what direction of travel will enable the delivery of the greatest number of packets.

As the nodes locations change, Perception will update its knowledge of the network. If acceptable delivery options are revealed, the queued messages will be automatically de-queued for delivery and the application notified of their success. Similarly, expired messages will be de-queued and the application notified of their failure.

Different requirements on reliability and latency will produce different behavior in a mobile node. For example, consider Figure 3 in which S and R are mobile nodes, B is a stationary buoy, H is a land station, and a stationary sensor network is on the sea floor. If S has a queued message to send to R, several delivery options exist. If low latency is required and reliability is not a factor, S can route the message through the sensor network. If moderate reliability is required but low latency is not, S can surface and radio the message to B to hold, anticipating that R might surface some time later (alternatively, it could send the message to H via satellite for holding). If high reliability is required or a large amount of data needs to be sent, S might achieve the best data rate by physically travelling to R.

The best course of action for message delivery depends on factors outside the scope of the networking subsystem, and it is left to the application to decide. In this way, the application finds the networking solution that most closely fits its needs.

IV. CONCLUSIONS AND FUTURE WORK

The design of our networking framework allows Delay-Tolerant Networking in an underwater environment, and is able to balance the mobility of its nodes with the communication requirements of its applications. Soliciting delivery criteria on a per-message basis gives our framework the ability to manage the trade-offs that are present in a network strained by mobility, loss, and latency.

This framework will enable the development of new protocols that can both measure and control their re-

liability and latency. In terms of traditional wired and wireless networks, UANs are attempting to solve the ultimate convergence problem in real-time: building and maintaining a reliable, efficient communications network without reliable communication between the constituent nodes. Mobility, changing environmental factors, and significant frame losses will make complete network convergence infeasible. However, net-centric applications for UANs using our framework will be able to take an active role in producing sufficient network convergence - constructing the communications infrastructure they require.

ACKNOWLEDGMENT

This research was supported in part by grant N00014-05-1-0666 from the Office of Naval Research.

REFERENCES

- [1] Cui, J., et al. Challenges: Building Scalable Mobile Underwater Wireless Sensor Networks for Aquatic Applications. *IEEE Network*, Special Issue on Wireless Sensor Networking. May/June 2006, Vol. 20, 3, pp. 12-18.
- [2] Cerf, V., et al. Delay-Tolerant Network Architecture: The Evolving Interplanetary Internet. s.l.: <http://www.ipnsig.org/reports/draft-irtf-ipnrg-arch-01.txt>, 2002. Work In Progress.
- [3] Srivastava, Vineet. Cross-Layer Design: A Survey and the Road Ahead. *IEEE Communications Magazine*. December 2005, pp. 112-119.
- [4] Raisinghani, V. and Iyer, S. Cross-Layer Feedback Architecture for Mobile Device Protocol Stacks. *IEEE Communications Magazine*. January 2006, pp. 85-92.
- [5] A Survey of Practical Issues in Underwater Networks. Partan, J., Kurose, J. and Levine, B. N. Los Angeles, CA, USA: ACM Press, 2006. Proceedings of the 1st ACM international Workshop on Underwater Networks. pp. 17-24.
- [6] Exploring Random Access and Handshaking Techniques in Large-Scale Underwater Wireless Acoustic Sensor Networks. Xie, Peng and Cui, Jun-Hong. Boston, MA, USA: MTS/IEEE, 2006. Proceedings of the MTS/IEEE Oceans 2006 Conference.
- [7] Peyravi, H. Medium Access Control Protocols Performance in Satellite Communications. *IEEE Communications Magazine*. March 1999, pp. 62-71.
- [8] A Network Layer Protocol for UANs to Address Propagation Delay Induced Performance Limitations. Xie, G. and Gibson, J. Honolulu, HI, USA: MTS, 2001. Proceedings of MTS/IEEE Oceans 2001 Conference.
- [9] Burleigh, S., Hooke, A. and Torgerson, L. Delay-Tolerant Networking: An Approach to Interplanetary Internet. *IEEE Communications Magazine*. June 2003, pp. 128-136.
- [10] ECLAIR: An Efficient Cross-layer Architecture for Wireless Protocol Stacks. Raisinghani, V. T. and Iyer, S. San Francisco, CA, USA: s.n., 2004. World Wireless Congress (WWC 2004).
- [11] Incorporating Realistic Acoustic Propagation Models in Simulation of Underwater Acoustic Networks: A Statistical Approach. Xie, G., Gibson, J. and Diaz-Gonzalez, L. Boston, MA, USA: MTS/IEEE, 2006. Proceedings of the MTS/IEEE Oceans 2006 Conference.
- [12] Location-Aware Routing Protocol for Underwater Acoustic Networks. Carlson, E.A., Beaujean, P.P. and An, E. Boston, MA, USA: MTS/IEEE, 2006. Proceedings of the MTS/IEEE Oceans 2006 Conference. pp. 1-6.
- [13] Performance Evaluation of Ad Hoc Protocols for Underwater Networks. Mupparapu, S., Bartos, R. and Haag, M. Durham, NH, USA: s.n., 2005. Proceedings of the Fourteenth International Symposium on Unmanned Untethered Submersible Technology.




Fig. 3. Multiple options for message delivery

- [14] Power and Distance Based MAC Algorithms For Underwater Acoustic Networks. Doukkali, H., Nuaymi, L. and Houcke, S. Boston, MA, USA: MTS/IEEE, 2006. Proceedings of the MTS/IEEE Oceans 2006 Conference. pp. 1-5.
- [15] Status Packet Deprecation and Store-Forward Routing in AUS-Net. Haag, Matt, et al. Los Angeles, CA, USA: ACM, 2006. WUWNet '06.
- [16] The PLUSNet Underwater Communications System: Acoustic Telemetry for Undersea Surveillance. Grund, M., et al. Boston, MA, USA: MTS/IEEE, 2006. Proceedings of the MTS/IEEE Oceans 2006 Conference. pp. 1-5.
- [17] Akyildiz, I., Pompili, D. and Melodia, T. Underwater Acoustic Sensor Networks: Research Challenges. *Ad Hoc Networks Journal*. March 2005, pp. 257-279.

APPENDIX

A. Choosing a MAC Protocol

Figure 4 shows a rough estimation of practical limits of communication with respect to node density.

TDMA protocols perform adequately up to the limits of acoustic range. However, each new time division adds another “guard band” to compensate for the variations in propagation times between the nodes. TDMA breaks down quickly for high node populations due to

the high number of guard bands and its inflexibility in allocating bandwidth to nodes that need it most.

Handshaking protocols such as RTS/CTS can provide service for a larger population of nodes, since bandwidth is allocated based on demand. However, as shown in Figure 5: The effects of distance on RTS/CTS and Random Access [6], RTS/CTS breaks down quickly as the coverage area increases [6]; latency between nodes is proportional to the number of steps in the handshaking sequence.

Random access protocols are largely unaffected by distance, at the expense of network capacity; random access has no collision avoidance mechanism. Such protocols are able to support large populations, but only when used in areas exceeding the acoustic range.

In summary, there is no “best” MAC protocol. Each has a unique and feasible combination of node population and coverage area for which it will perform better than other MAC protocols - and none of the protocols will work in all possible cases. The ability to recognize the boundary conditions for a given MAC protocol and

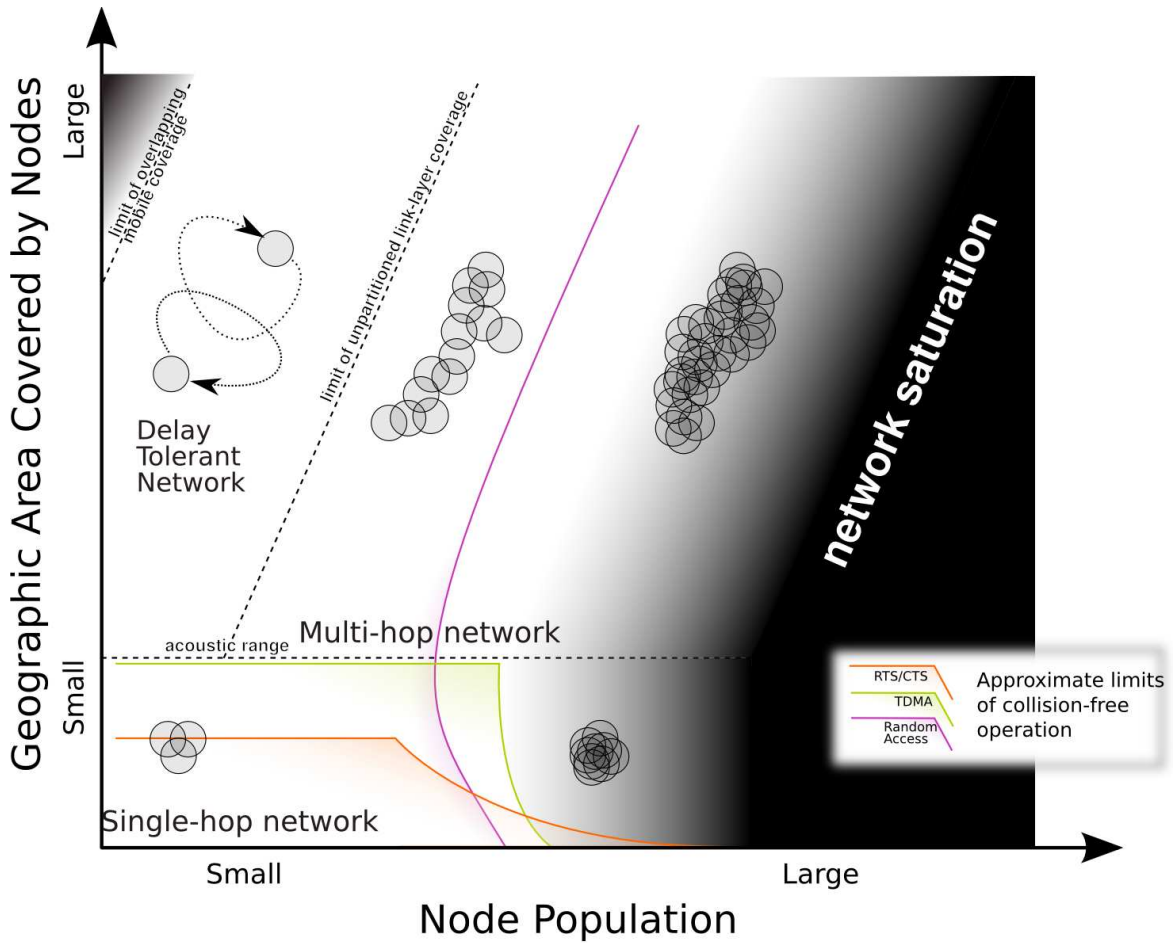


Fig. 4. Choosing the MAC protocol for node density

the ability to automatically change a network to a more appropriate protocol remain open problems.




Fig. 5. The effects of distance on RTS/CTS and Random Access [6]