Robust time-division channel-access approach for an ad hoc underwater network

Kae Yeet Foo[a], Philip Atkins[a], Steve Pointer[b] and Carl Tiltman[c]

[a]University of Birmingham, Department of Electronic, Electrical and Computer Engineering, Edgbaston, B15 2TT Birmingham, UK
[b]Qinetiq, Winfrith Technology Centre, DT2 8XJ Dorchester, UK
[c]Dstl, Winfrith Technology Centre, DT2 8WX Dorchester, UK
kyf740@bham.ac.uk

An ad hoc underwater network is formed by a cluster of stationary nodes that can act as a source, destination, or a relay, in which data packet usually travels across multiple hops. Using a time-division scheme where the propagation delay between the nodes is used as a packet queuing buffer optimises throughput. This requires accurate knowledge of the relative ranges of all the nodes, a high level of accuracy in time-synchronisation, and restricts changes in the position of the nodes especially in the absence of a central server or master node within the network. A method that offers greater robustness is by using a time-slotted approach, where each node is offered a time-slot sufficient for one transmission to reach its maximum effective range. A passive acknowledgement scheme is implemented where a node listens for acknowledgement when the transmitted is being relayed. Nodes monitor other time-slots for the opportunity to contend for idle slots in order to enhance throughput. Results obtained from simulations demonstrate that this method is robust in supporting changes in the relative distances between the nodes, and can operate with a time-synchronisation error of up to 1 second.

1 Introduction

Acoustic modems enable underwater assets to communicate over a large distance without the constraint of physical cables. Such a point to point acoustic link plays an important role for a variety of underwater missions, such as remote command of an autonomous underwater vehicle (AUV), remote environmental monitoring, and more generally, the exchange of information between two underwater assets. Ad hoc underwater networking capability increases the role and effectiveness of point to point communications offered by existing acoustic modems. An ad hoc underwater network extends the communication range of existing acoustic modems, enables communication between underwater assets that do not have acoustic line of sight, and supports the communication between multiple users.

The underwater network discussed herein is constrained to a small group of nodes (between 3 to12) that can be deployed in an ad hoc manner without a priori knowledge of their geographical coordinates. An underwater asset that has acoustic communication capability and participates in the network is a node. Each node in the network can act as both an access point and or a relay at different instances. The nodes are autonomous in a way that they do not rely on a central server or a master node for decisions on routing or channel access [1]. Such a network can be deployed, or be formed, in various arrangements. Also, nodes may be removed, replaced or added to the network.

The challenges associated with implementing an underwater network is the slow speed of propagation, an unstable acoustic channel [2], and the finite battery life on the nodes. A network designer often needs to find the optimal point for parameters of the network such as the inter-node distance, operating frequencies and modulation methods based upon contending factors, for example ambient noise, sea-state and local noise (i.e. shipping activities). This optimisation is usually in the context of energy consumption.

This paper begins with an analytical exercise of finding the optimal cost of energy with respect to operating frequencies and inter-node distance within an operating environment with typical noise sources. This is followed by the description of a channel access scheme that is suitable for a multiple-hop underwater network. The time-division approach proposed in this paper is robust in that it allows for changes in node positioning, enabling the network to be scalable, while minimising the additional packet latencies associated with such scalability. The algorithm of this protocol is presented.

An ad hoc time-synchronisation approach for a multi-hop network without a central server is also described. Time-synchronisation is essential for channel access schemes based upon the time-division approach. In a multi-hop network, where nodes do not have acoustic line-of-sight to all the other nodes, the challenges for time-synchronisation are the autonomous selection of a reference time and the distribution of this time reference across the network.

This channel access scheme was implemented and demonstrated with the nodes of an underwater networking simulator developed in-house. It is shown that the scheme can accommodate changes to the inter-node distances and the relative positioning of nodes. The main contribution of this paper is the description and implementation of a time-division channel access scheme for a multiple hop ad hoc network. For applications where the networking of underwater assets is optimised for a single hop, readers are referred to specific research by Freitag et al. [3].

2 Optimal operating frequency and inter-node separation

In deploying an underwater network, it is often desirable to estimate the optimal level of essential parameters such as the operating frequency and inter-node separation. One method is to perform a multi-variable optimisation using the total financial cost of the network as the output metric. It is worth noting that the main assumption made here is that using mass-production techniques, the cost of a node is determined by the weight of the units, which is primarily determined by the battery. For example, the battery of a 2005 vintage mobile phone contributes about one-third of the weight and cost of the unit. As the standby time of such a unit is of the order of one week, an extrapolation to a three-month deployment would imply that battery cost would comprise about 87% of the total cost. As with a mobile phone, the cost of operating an underwater network is related to the ratio of transmission time to standby time – arbitrarily chosen as one minute of transmission time per day for the following assessment (a saturated network might typically reach 7 minutes of transmission time per day).

The sonar equation may be solved in terms of required source level for a one-way communication path.

$$SL = NL + PL + SNR + DI_{RX} \tag{1}$$

where SL is the source level (including any transmit directivity), NL is the noise level (including any system bandwidth parameters), PL is the path loss, SNR is the required signal-to-noise ratio (typically +15dB for a communication system) and DI_{RX} is the directivity of the receiver (often 0dB as omni-directional transducers are used). The system is assumed to be noise limited, with the noise levels varying with frequency

The path loss has assumed a simple spherical spreading model of $PL = 20\log_{10} r + \alpha\, r$ where r is the range and α is the frequency depended absorption coefficient. Spherical spreading is assumed, as all the ray paths are generally resolved within the receiver. The lack of inclusion of surface and sea bed interaction losses and other propagation effects will lead to an overoptimistic prediction of range performance and an underestimation of the financial cost of the network. The sonar equation may be solved as a function of frequency and path length and the optimum operating frequency chosen in order to minimise the required source level. At short ranges (<1 km) the optimum operating frequency is dependent on the environmental conditions and higher frequencies are preferable (100's kHz) under severe weather conditions. As the inter-node separation increases to a few kilometres, the optimum operating frequency drops to the industry-standard band of 8 – 20 kHz frequently used by technically similar long-baseline navigation systems. At long ranges (>30 km) the effects of multiple local minima in the solution can be seen and the optimum frequency drops to a few 10's Hz. Although these figures provide the optimal operating frequency as a function of inter-node separation, it will be shown that the total cost of operating a network also influences the optimal operating frequency.

The number of nodes, N, required to fill an area, A, is given by

$$N = \frac{A}{d^2} \quad (2)$$

where d is the inter-node distance. Interestingly, if the only power used were associated with the transmission function, the optimal power-usage solution would be to use an infinite number of closely spaced nodes. However, although this would lead to an available bandwidth of many MHz, the latency in a store-and-forward network would become infinitely long. In a realistic network the total power consumption of the node, P_{TOT}, is the sum of the transmitter related functions, P_{TX}, and the quiescent power consumption associated with a continuously running receiver, P_Q.

$$P_{TOT} = P_{TX} + P_Q \quad (3)$$

Thus as the inter-node separation increases, the total quiescent power used by the network decreases but the required transmit power increases. A minimum will be encountered at some inter-node separation as illustrated for a 4 kHz bandwidth in Fig. 1. Generally, this minimum is comparatively flat and a 10% leeway in total mission cost will provide the designer with a wide range of potential inter-node separation distances.

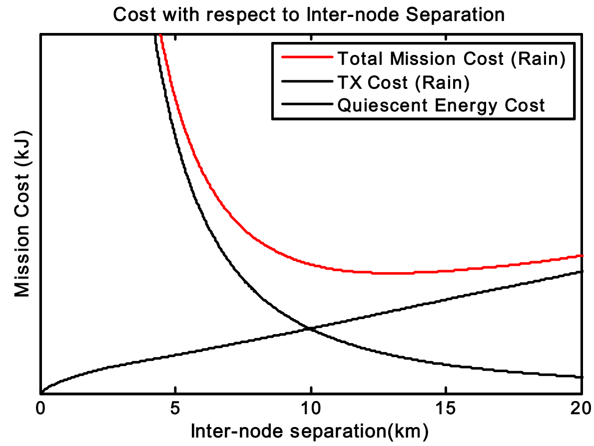

Fig. 1: Illustrative plot of total mission cost as a function of distance between nodes.

The total mission cost in Fig. 1 indicates that the quiescent power consumption of the nodes plays a critical part in determining the overall mission cost. This quiescent power consumption also determines the optimum operating frequency and the inter-node distance. For example, by repeating the calculation used to determine Fig. 1, it is possible to compute the optimum operating frequency as a function of quiescent power consumption, as shown in Fig. 2. The vertical error bars indicate the range available to the designer for a 10% increase in the cost of operating the network. The right-hand extreme of this plot shows the typical quiescent power consumption achieved within the base station processing sections of a mobile phone network in the late 1990's. The power consumption of digital components has decreased by approximately a factor of ten per decade. Thus with improved energy efficiency within the receiver section of a node, the optimum operating frequency will increase. The minimum noise level plot (Sea State 0) shows a transition between local minima within the optimisation process.

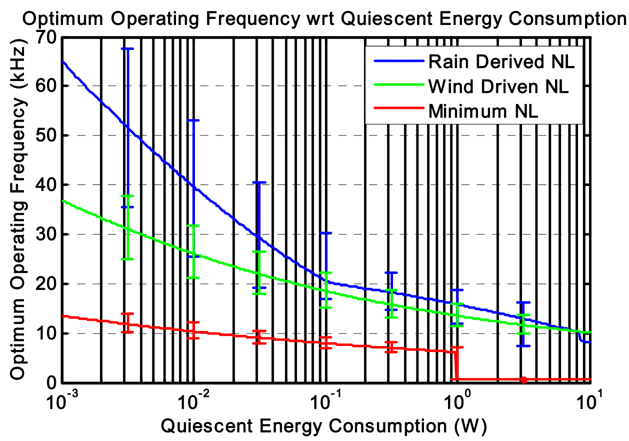

Fig. 2: Optimum operating frequency as a function of quiescent power consumption.

This optimisation process also yields the financially optimum node separation as a function of quiescent power consumption, as shown in Fig. 3. The node separation is reduced under adverse weather conditions and will reduce as technology improves.

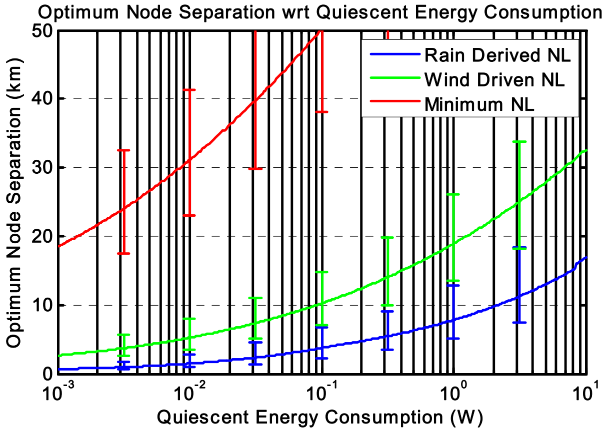

Fig. 3: Optimum node separation as a function of quiescent power consumption.

3 Robust time-division channel access

3.1 Time division approaches

Time-division channel access is a strategy where users share the access to the channel based on their allocated time. For analytical purpose, consider a cluster of three equidistant nodes with acoustic line-of-sights, where r_0 is the equidistance between the nodes. Let the propagation delay associated with r_0 be t_p. The total time required to complete a one-way transmission, t_0, is expressed as:

$$t_0 = t_h + t_d + t_p \tag{4}$$

where t_h is the length of the data packet preamble and processing time, and t_d is the length of the data payload. Each of the nodes is allocated a time-slot of equal length t_0, and let T_c be the time of a complete cycle within which all the allocated time-slots elapse once. Fig. 4 illustrates the basic time-allocation method implemented herein, with one time-slot for each node within each full cycle.

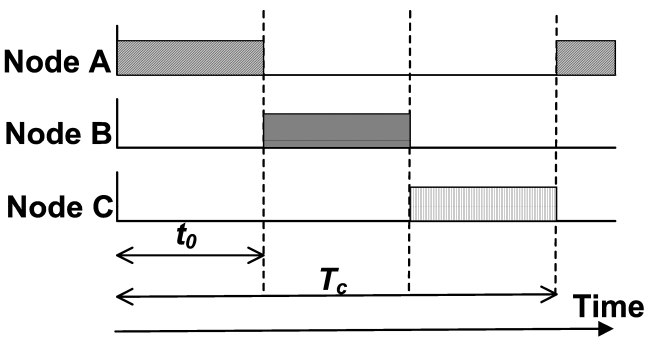

Fig. 4: Time-division channel access with a simple time-slot allocation herein implemented.

There are variants of time-division schemes where the propagation delay may be used as a single-transmission duplex buffer [4] or a multiple transmission buffer [5]. Buffering the propagation delay is efficient, and seeks to maximise throughput while reducing latency. However, they are based on the assumption that the distance between all the nodes are known, and remains the same. If not, they need to be learned and their changes need to be tracked. The use of the propagation delay as a buffer also constrains the flexibility of node arrangement within the network. Hence, in order for the protocol to operate efficiently, it is necessary to pre-determine the arrangement of nodes within the network prior to deployment.

In a network where the deployment of nodes was carried out in a completely ad hoc manner, the arrangement of the nodes is usually not known a priori. Also, since each node has line-of-sight to a relatively small number of nodes in the network, it is not straightforward to learn and track changes in the node arrangement before autonomously allocating time-slots to each of the nodes.

3.2 Robust time division scheme

The typical underwater network under consideration consists of 5 to 20 nodes, and the nodes are usually deployed with a distance of 1000 to 4000 m from one another, in a sparse and multiple-hop arrangement. For this network a suboptimal, yet robust, scheme is to allocate time-slots for each node, without using the propagation delay as a buffer, as shown in Fig. 4, but nodes contend for other time-slots that are perceived as idle. This trades channel efficiency for increased robustness for network deployment and maintenance.

With this scheme, a maximum of N_{max} slots are made available for allocation, each of length t_0, forming a complete cycle of length T_c. Each node in the network would have an individual time-slot. A network of 15 nodes would have 15 unique, equally long, time-slots within a full time-cycle. In order to increase throughput and to exploit the fact that nodes more than 2 hops away do not mutually interfere, nodes contend for other slots that are perceived to be idle. The number of slots that a node contends for within T_c, $N_{contend}$, is dependent upon N_{max} and the expected number of neighbours in the network, K, where

$$N_{contend} \le \left| \frac{N_{\max} - 1}{K^2} \right| \tag{5}$$

and $|N|$ denotes the rounding of N to the nearest integer. This expression states that the number of slots that a node can contend for within each full-cycle of T_c is equal to the total number of slots minus the slot belonging to the node, and divided among the number of nodes that are likely to cause collisions when contending on the same time-slot. This is taken as K^2 to account for the hidden node effect, where two nodes that do not have acoustic line-of-sight with each other would still cause a collision at a third node located between them. The length of the time-slot, t_d, is as described in Eq. (4), where t_d and t_p are set to be the maximum length of the data payload and the propagation delay associated with the maximum inter-node distance (e.g. 4000 m) respectively.

The synchronisation of the time-epochs on all the nodes is conducted by an ad hoc time synchronisation protocol [6]. This protocol relies on the availability of two-way ranging facility on the acoustic modem, and is as follows:

(a) A session is delegated for time synchronisation, during which no data packet (i.e. from the data layer) will be transported. This can be set to a specific time of the day (e.g. between 0000 and 0030 hours each day) or manually invoked by an operator. If invoked manually, the synchronisation session needs to be extended in order to take account of existing data queues within the network.

(b) Within the session, a time-reference node is autonomously elected (e.g. by using the node with the largest or smallest ID) – this must be recognised by all the nodes in the network, both existing and recent. An alternative would be a pre-elected node.

(c) The time reference node broadcasts its time-reference. Where the largest ID approach is implemented, most of the nodes would broadcast their time-reference on the assumption that they could be the node with the largest node ID, with the exception of the node with ID 1, and nodes that had received broadcasts from another node with a larger ID before their own attempt to broadcast.

(d) Nodes that receive the broadcast then record the reference node and time, and derive their range from the broadcaster by performing a two-way ranging operation to calculate the associated propagation time-delay. By adapting the broadcaster's time-reference at the time of reception, and then removing the element of the propagation delay, the time-epochs of the nodes are synchronised.

(e) If the new time-reference is fresher than the one held locally, it is updated and then rebroadcasted. The freshness parameter is decided by the rule for autonomous clock selection. For example, with the largest ID rule, a node with ID 3 that had just synchronised to node ID 4 would resynchronise and rebroadcast its time when a time-reference from node 5 is later received, but not vice versa.

(f) The time-slot in which nodes initially broadcast their time-references may not be synchronised and thus result in the collision of broadcasts. Therefore, a passive acknowledgment mechanism is necessary to invoke a broadcast retry. The passive acknowledgement mechanism operates by a node listening for ranging attempts or rebroadcasts of its time-reference by another node with a smaller ID. If these are not heard after a number of time-cycles, the node would retry to broadcast its time-reference using a different time-slot, up to a maximum number of retries.

(g) It is essential that the record of the master time-reference is reset to the local node ID at the end of the time-synchronisation session. This enables resynchronisation to the node with the largest ID in any future time-synchronisation session.

After time-synchronisation is completed, all the nodes would observe the same time-slot at any point in time. Each node transmits within its allocated time-slot, as well as time-slots belonging to the other nodes that are perceived as idle, up to a total of $N_{contend}$. Nodes observe other time-slots and maintain a time-stamped score table on the level of activities during the time-slots allocated for other nodes. Nodes choose to contend for the slots perceived to be most idle. The scoring also takes into account failed and successful contentions. The observation window of the scoring table progresses with time in order to adapt to node repositioning or replacement.

The acknowledgement mechanism is based upon the sender overhearing the transmission when the packet is being forwarded onto the next hop down the route. If the receiving node is the final destination, the receiver would just send a short acknowledgement to the packet forwarder.

This scheme is robust because node positioning need not be known a priori, and there is no need for a master or central server in the network to organise the allocation of time-slots. Nodes can be deployed in a multihop arrangement. Also a node maybe replaced or repositioned in the lifetime of the network. The data packet length plus processing time can be longer than the propagation delay. This is quite common where the propagation delay is only in the region of 1 to 3 seconds, in which case buffering the propagation delay may not always be feasible. In addition, the accuracy required of time-synchronisation is relatively less demanding compared to schemes that buffer the propagation delay.

3.3 Simulation results

A network of six nodes is deployed in a simulator developed in-house using the C++ platform. The nodes are deployed in a multiple-hop arrangement, as shown in Fig. 5. The dotted ring around each node denotes the effective communication range. The average inter-node distance was 4000 m, giving a t_p of 3 s. The maximum data packet length is 4 s and t_h is arbitrarily set to 2 s. A guard time of 1 s is added, giving a time-slot length, t_0, of 10 s. The complete cycle, T_c is therefore 60 s. Nodes are allowed to contend for one other time-slot (Eq. (5)) within each complete cycle.

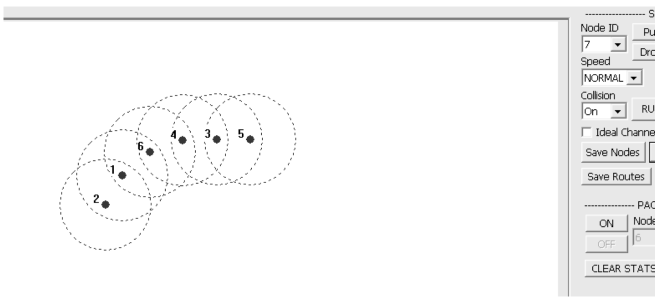

Fig. 5: Scene 1 – a network of six nodes deployed in a multiple hop arrangement in the simulator.

After deployment, the time-slots on the nodes were not synchronised. Fig. 6 and Table 1 show the time-epoch of the nodes before and after synchronisation, where an accuracy of 0.46 second standard deviation was achieved.

(a)

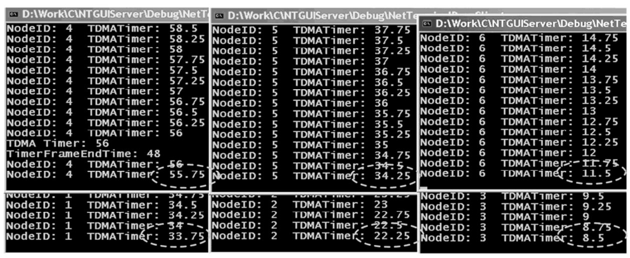

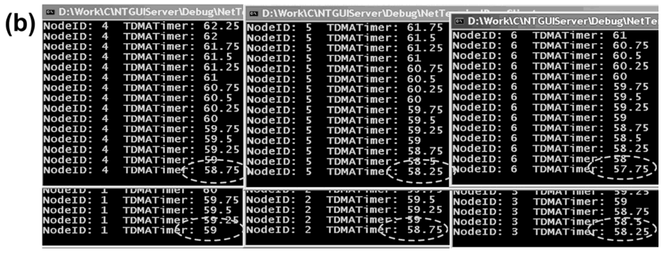

Fig. 6: Screen snapshot of the time-slots for all the nodes (a) prior to time-synchronisation and (b) after time synchronisation.

NODE ID	BEFORE SYNC	AFTER SYNC
1	**33.75**	**59.00**
2	**22.25**	**58.75**
3	**8.50**	**58.25**
4	**55.75**	**58.75**
5	**34.25**	**58.25**
6	**11.50**	**57.75**
Std Deviation	**17.48**	**0.46**

Table 1: Time-slots of nodes prior to and after time-synchronisation.

A total of 50 data packets with the maximum length of 4 s were sent from node 2 destined for node 5, and vice versa. Each pair of data packets, one in each direction, were introduced at an interval of 180 s apart, so as not to saturate the network and cause a growing queue of data packets at the source nodes. Routes were already set up before introducing the data packets. This was then repeated after rearranging the nodes into having different inter-node spacing as shown in Fig. 7. The results for both scenarios are shown in Table 2. Latencies for the first 15 pair of packets were excluded from the calculations to exclude the effect of contentions required for learning the best idle slot. The latency results are consistent with the protocol, which allowed nodes to transmit twice in every complete cycle of 60 seconds. The results also indicate that a change in the inter-node spacing had no significant impact on the normal operation of the channel access scheme.

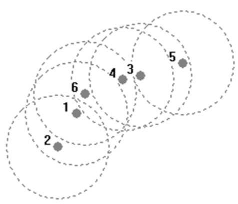

Fig. 7: Scene 2 – nodes are rearranged to have non-uniform inter-node spacing.

ROUTE	SCENE-1	SCENE-2
2 to 5	**29.2 sec/hop**	**39.2 sec/hop**
5 to 2	**36.4 sec/hop**	**33.9 sec/hop**
Average Latency	**32.8 sec/hop**	**36.5 sec/hop**

Table 2: Packet latencies for both scenarios.

5 Conclusion

A simple time-division channel access scheme suitable for an ad hoc, multiple-hop underwater network has been described. The scheme operates along with ad hoc time synchronisation, passive acknowledgement and contention for idle slots. In comparison with techniques that exploit the propagation delay, this scheme is suboptimal in its achievable throughput. However, this scheme offers better robustness in terms of network scalability and positioning. Also, the requirement on the accuracies of time-synchronisation is relatively small.

Simulations were carried out with six nodes in order to test the viability of this scheme, under the assumption of ideal channel conditions and the availability of routing information. The results demonstrate that time synchronisation for all the nodes can be achieved in a multiple-hop network in the absence of a pre-determined master, up to the accuracy of 0.5 s standard deviations. Packet latency results are consistent with the protocol of the channel access scheme, and changes in inter-node spacing across the network had no significant impact on the normal operation of the network.

References

[1] K.Y. Foo, P.R. Atkins, T. Collins, C. Morley, and J. Davies, "A routing and channel-access approach for an ad hoc underwater acoustic network," *MTS/IEEE TECHN-OCEANS 2004. (MTS/IEEE Techno-Ocean)*, vol. 2, Nov. 2004, pp.789-79

[2] E.M. Sozer, M. Stojanovic, and J.G. Proakis, "Underwater acoustic networks," *IEEE Journal of Oceanic Engineering*, vol. 25, no. 1, Jan. 2000, pp.72-82.

[3] Stojanovic, M.; Freitag, L.; Leonard, J. and Newman, P.; "A network protocol for multiple AUV localization," *MTS/IEEE Oceans '02*, vol. 1, 29-31 Oct. 2002, pp. 604-611.

[4] B. Hou, O.R. Hinton, A.E. Adams and B.S. Sharif, "A time-domain-oriented multiple access protocol for underwater acoustic network communications," *MTS/IEEE OCEANS 1999*, vol. 2, Sept. 1999, pp. 585-589.

[5] I.P. Morns, O.R. Hinton, A.E. Adams, B.S. Sharif, "Protocols for sub-sea communication networks," *MTS/IEEE OCEANS 2001*, vol. 4, pp. 2076-2082.

[6] K.Y. Foo, P.R. Atkins, S.A.Pointer and C.P.Tiltman, "Ad hoc time synchronisation in a multiple hop underwater network," *Proceedings of the Institute of Acoustics Spring Conference 2008,* Widening Horizons in Acoustics, April 2008.