AD HOC TIME SYNCHRONISATION IN A MULTIPLE HOP UNDERWATER NETWORK

KY Foo The University of Birmingham, U.K.
PR Atkins The University of Birmingham, U.K.
SA Pointer QinetiQ, U.K.
CP Tiltman DSTL, U.K.

1 INTRODUCTION

Underwater networking offers the capability to extend the communication range of existing acoustic modems, to enable communication between underwater assets that do not have acoustic line of sight, and to support the communication between multiple users. An underwater network can be deployed in various positional arrangements, and may run on different channel access and routing protocols depending on the target application. For example, inter-networking between three underwater autonomous vehicles deployed on a specific mission may not require any routing capabilities and can operate with a time-division channel access protocol. On the other hand, an underwater network designed to carry low traffic data packets over a long distance will require routing information, and may run with a contention-based channel access protocol.

In this paper, the underwater network being discussed is defined as a group of a small number of nodes (between 3 to12) that can be deployed in an ad hoc manner without a priori knowledge of geographical coordinates. Each node in the network can act as both an access point and a relay at different instances in time. The nodes are autonomous, and do not rely on a central server or a master node for decisions on routing or channel access[1].

In order to maximise the coverage and reach of the network when given a limited number of nodes, it is often desirable to deploy the nodes in a sparse distribution. As a result of such an arrangement, the nodes in the network are usually within communication range of one to three other nodes in the network. Any individual node can only communicate directly to the other nodes within the same region. An arbitrary effective communication distance in the range of 500 m to 6000 m is being considered. Practically, this depends on the modulation scheme and the environment in which the network operates. In order to send packets beyond this distance, data packets will need to travel more than one hop, over intermediary nodes that will act as a relay.

The challenges associated with implementing an underwater network is compounded by the slow speed of propagation and an unstable acoustic channel[2]. 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.

A time division protocol based upon allocating time-slots to nodes is herein proposed. Implementing a time-slotted protocol on such a network gives the advantage of predictable latency and a fair access to the channel. However, applying this protocol in the network would require two issues to be addressed: the presence of a unique clock for all the nodes to synchronise on, and the process of synchronisation over multiple hops.

This paper presents a method for performing ad hoc time-synchronisation over such a multiple hop network. This method relies on acoustic ranging between nodes, and assumes that this facility is available on the acoustic modem.

2 TYPICAL NETWORK PARAMETERS

2.1 Energy Cost Optimisation

One method of deriving the required operating parameters for an underwater network is to perform a multi-variable optimisation using the total financial cost of the network as the output metric. The major assumption made 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 transmission time per day for the following assessment (a saturated network might typically reach 7 minutes transmission time per day).

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

$$SL = NL + PL + SNR + DI_{RX} \quad (1)$$

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, as shown in Fig. 1. Three cases are considered; the quietest open-ocean conditions, high-winds coupled with heavy shipping and the worst case is that of heavy rain. It is evident that designing for the worst-case scenario would impose a 40 dB penalty – equivalent to a 10,000 increase in required energy costs. These noise figures are representative provided there are no local noise sources present (i.e. typically being further than 500 m from a noisy commercial vessel).

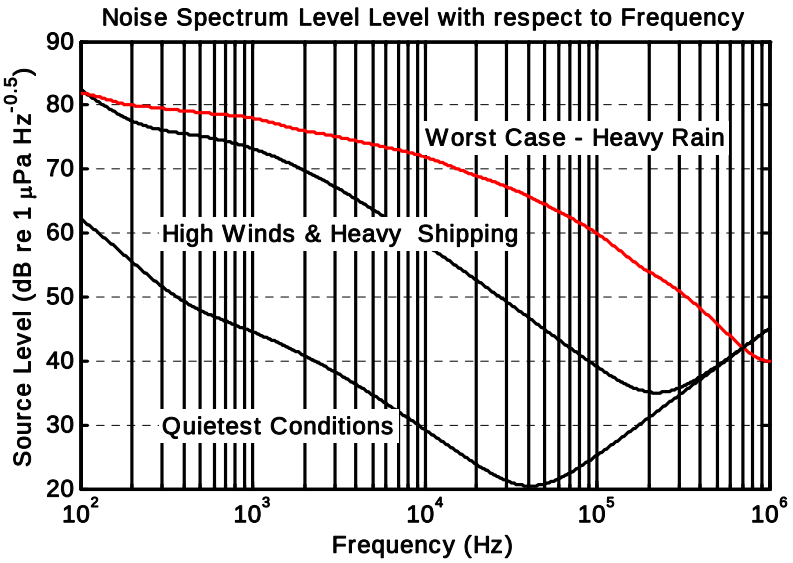

Figure 1: Typical ambient noise spectrum levels used in assessment

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, as shown in Fig. 2. 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 may migrate this operating frequency.

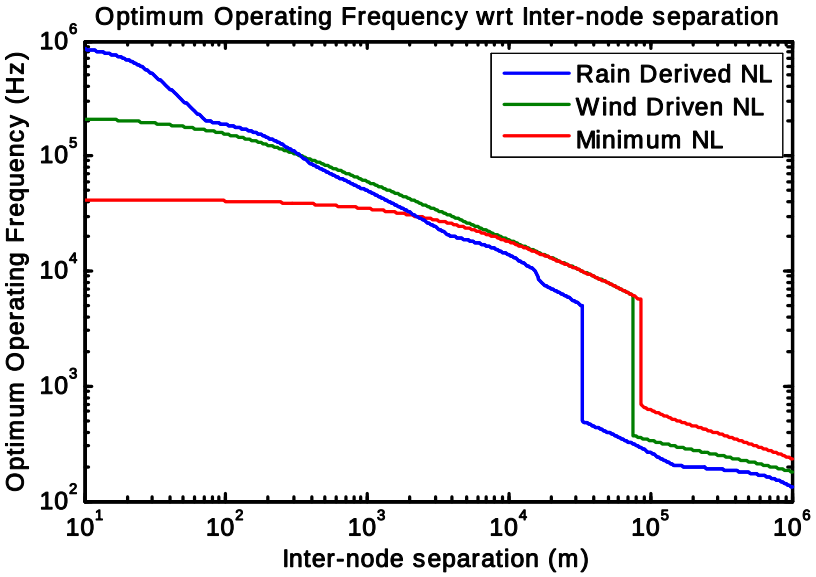

Figure 2: Optimum operating frequency as a function of distance between nodes.

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

$$N = \frac{A}{d^2} \tag{2}$$

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.

$$P_{TOT} = P_{TX} + P_Q \tag{3}$$

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. 3. 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.

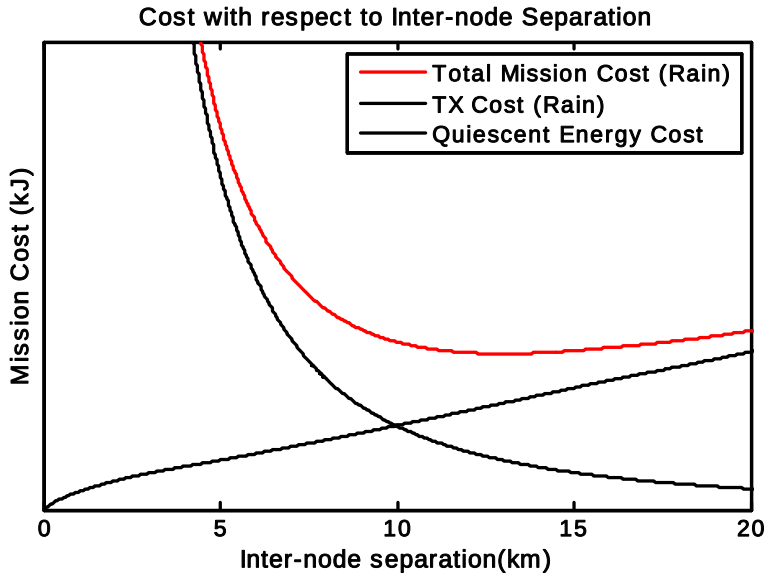

Figure 3: Illustrative plot of total mission cost as a function of distance between nodes.

It becomes apparent 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. 3, it is possible to compute the optimum inter-node separation as a function of quiescent power consumption, as shown in Fig. 4. 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.

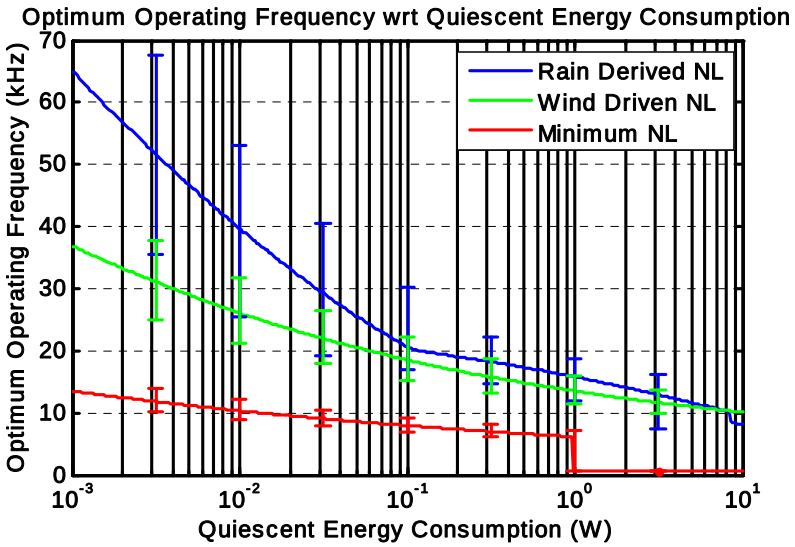

Figure 4: Optimum operating frequency as a function of quiescent power consumption.

This optimisation process also yields the financially optimum node separation as a function of quiescent power consumption, as shown in Fig. 5. The node separation is reduced under adverse weather conditions and will reduce as technology improves.

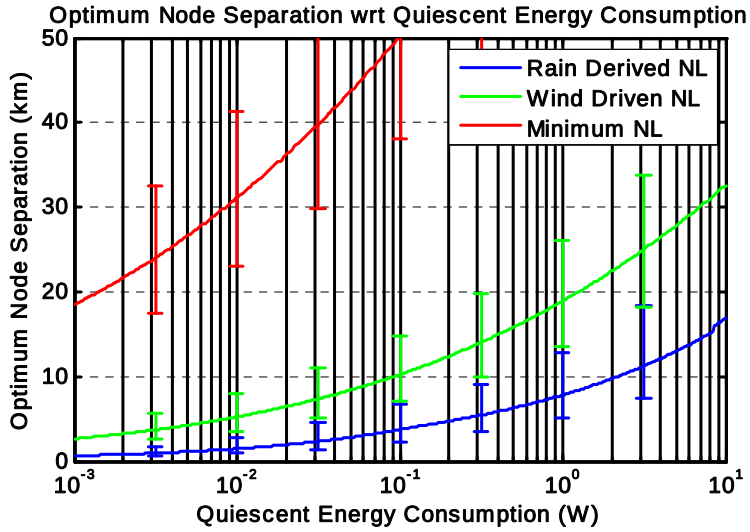

Figure 5: Optimum node separation as a function of quiescent power consumption.

Therefore, from a purely cost-driven approach, a designer would probably implement a network operating in the frequency band 8 – 16 kHz, with a desired inter-node separation of 5 – 12 km. However, at this simple level of financial optimisation, no account has been taken as to whether the required source levels are achievable given the crest factor of the modulation schemes involved and is therefore purely indicative of potential system parameters.

2.2 Asymmetric Links

Radio communication networks usually assume symmetrical links (e.g. where nodes can reciprocally communicate with one-another). In a shallow-water water environment this assumption may frequently be invalid. Consider the scenario shown in Fig. 6 where a local commercial vessel dominates the ambient noise environment.

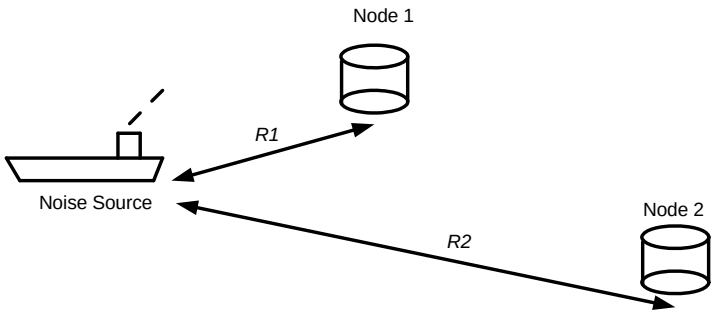

Figure 6: Typical shallow water scenario where a local noise source dominates the environment.

If the local noise source generates a noise spectrum level, N_0, the Noise Level (NL) received by Node 1 at range *R1* will be $NL_1 = N_0 - 20\log_{10} R1 - \alpha R1$. Similarly, the Noise Level received by Node 2 at range *R2* will be $NL_2 = N_0 - 20\log_{10} R2 - \alpha R2$. The difference in required source levels should Node 1 wish to symmetrically communicate with Node 2 is

$$SL_1 - SL_1 = 20\log_{10}\left(\frac{R2}{R1}\right) + \alpha\left(R2 - R1\right). \quad (4)$$

Taking an illustrative example of two nodes located 500 m and 1500 m from a local noise source, the required difference in source levels is 9.8 dB. The implication is that communication nodes should always incorporate an adaptive transmit power management strategy. The situation becomes increasingly severe the closer the nodes are positioned with respect to local noise sources.

3 TIME DIVISION PROTOCOLS

A time-slotted channel access protocol allocates unique time-slots to individual nodes in order to mitigate collision. The allocation of time-slots can be buffered, in which the propagation delay between two nodes, t_p, is interleaved to accommodate more than one transmission [3] [4], as shown in Fig. 7. This is effective in exploiting the propagation delay between two nodes (in the region of 0.5 to 4 seconds depending on transmission ranges) to maximise throughput. However, the caveat with this method is that the nodes need to be within pre-determined ranges from one another. As an example, where the minimum packet delay is 1 second, the propagation delay between two nodes need to be at least 2.2 seconds in order to achieve buffering for two transmissions plus the allowance of 0.1 second guard time to each transmission. This imposes a constraint on the positional arrangement of node deployment. There will also be little tolerance to geographical drifting or relocation of nodes.

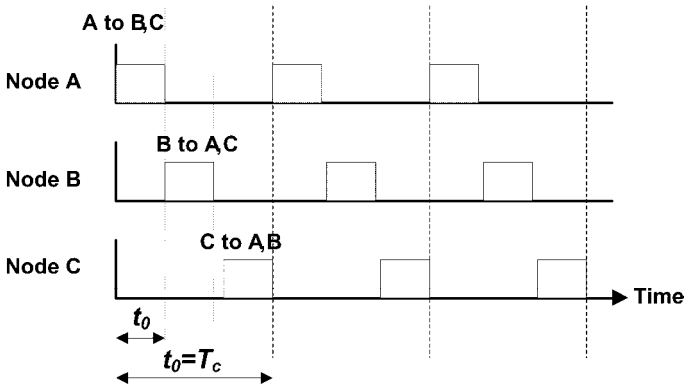

Figure 7: Buffered time-division approach by interleaving transmissions from multiple nodes into the propagation delay.

A more flexible alternative, therefore, would be to allocate time-slots that can accommodate the total duration needed to complete one data exchange session between two nodes, as illustrated in Fig. 8. The nodes are allocated one time-slot each, of duration t_0, in every full time-cycle, T_c. This is at the price of a reduced throughput, by not exploiting the propagation delay to buffer transmission from other nodes. There are two main benefits to this approach. The positional constraint is relaxed such that nodes can be deployed at any distances with respect to one another as long as their maximum propagation delays do not exceed the maximum allowed by one time-slot. Secondly, the requirement for time-synchronisation accuracy is now less stringent, enabling synchronisation between two nodes to rely only on existing modem ranging information without the necessity for added hardware complexities. The efficiency of time-slot usage can be increased by allowing nodes to passively listen during the time-slots allocated for other nodes, and contending to

use them when required – observing the fact that two nodes located at 3 or more hops away can use the same unique time-slot without causing a collision. A score table can be used to score the other time-slots in order to learn those available for contention and to adapt to changes in node positioning.

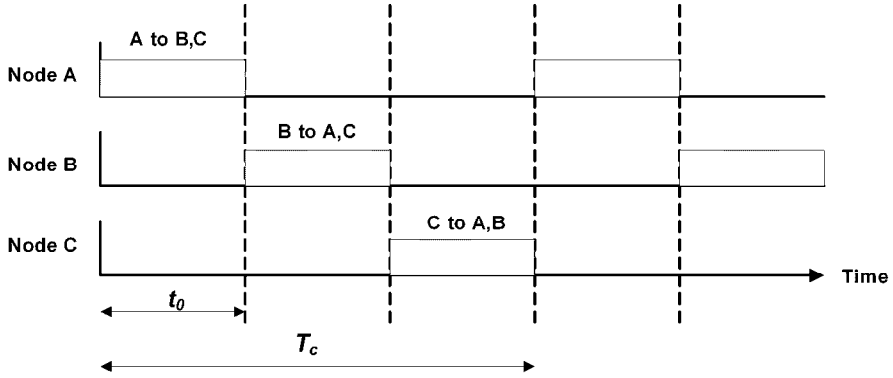

Figure 8: A simple time-slot approach, where one transmission session is completed within every time-slot.

It is against the background of the aforementioned time-slotted channel access scheme that the requirement to achieve time-synchronisation for all the nodes arises. In order to achieve this, there needs to be a rule for autonomous selection of a unique clock amongst the nodes. This unique clock shall be the time-reference for synchronising the time-epochs on all the nodes. Following the selection of a unique clock, a protocol to disseminate this time-reference accurately and in an orderly manner across multiple hops is required.

4 AD HOC TIME SYNCHRONISATION

4.1 Unique Clock Selection

The first task is to identify and choose a unique clock to which all the nodes would synchronise. Immediately upon deployment, the only information that distinguishes the nodes is their unique address or identification (ID) number. Hence, a simple rule for unique clock selection is to synchronise to the clock on the node with the largest ID number. This rule forms the basis for an autonomous master-to-slave relationship for time-synchronisation, where a node synchronises to the time reference received from another node that has a larger ID.

An alternative is to pre-select a node that will act as a master clock. This is at the expense of introducing a thin layer of hierarchy for the time-synchronisation process, and therefore this node, or node ID, must always be active within the network.

4.2 Time Reference Dissemination

With a rule ensuring that all the nodes will only synchronise to one unique clock within the network, the next step is to implement a protocol for the dissemination of this time-reference across multiple hops. This is described by the following:

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 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. It is worth noting that 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 ID 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.

Both time-reference broadcasts and ranging operations should be performed within allocated time-slots. The allocation of time-slots is usually based upon the ID of the nodes. If the node density is low, nodes may contend for additional time-slots within each time-cycle as transmissions only affect nodes up to 2 hops away (inclusive of the hidden node). The recommended number of additional slots to contend for is:

$$N_C < 0.5\ ((N_T - 1) - (2K - 1)) \qquad (5)$$

where N_C is the number of additional slots to contend for, N_T is the total number of time-slots in one full-cycle and K is the average number of neighbours per node.

Residual synchronisation errors between adjacent nodes (i.e. from uncompensated system delay) should not exceed t_g/H, where t_g is the guard time in each time-slot and H is the maximum number of hops travelled by any data packet in the network.

5 SIMULATION RESULTS

This time-synchronisation method was tested in a network simulator, as shown in Fig. 9. A network of 6 nodes was arranged linearly, with an arbitrary inter-node distance of 3000 m. The channel was assumed ideal, such that no packets were lost other than that due to collisions. Nodes only have acoustic links with adjacent neighbours.

The time-slotted protocol was manually initiated on each node, thus the nodes' time-epoch were not synchronised. For example, at a specific instance in time, node 1 might be observing time-slot 4 while node 2 was observing time-slot 5. If the nodes were synchronised, all the nodes would be observing the same time-slot (and its associated progression) at any instance in time. A screen capture in Fig. 10 shows the time-epoch of all the six nodes in the simulator prior to the time-synchronisation, where the last line in each window shows the node ID and the corresponding time-epoch at the time of the screen capture.

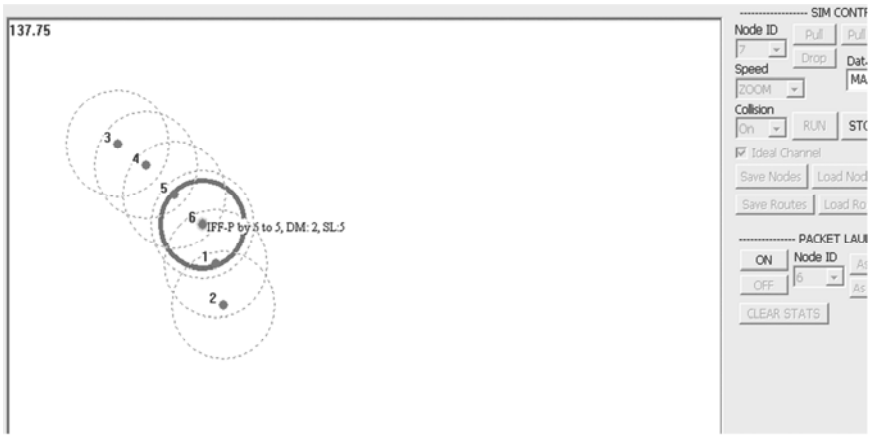

Figure 9: An arbitrary network of 6 nodes in the simulator.

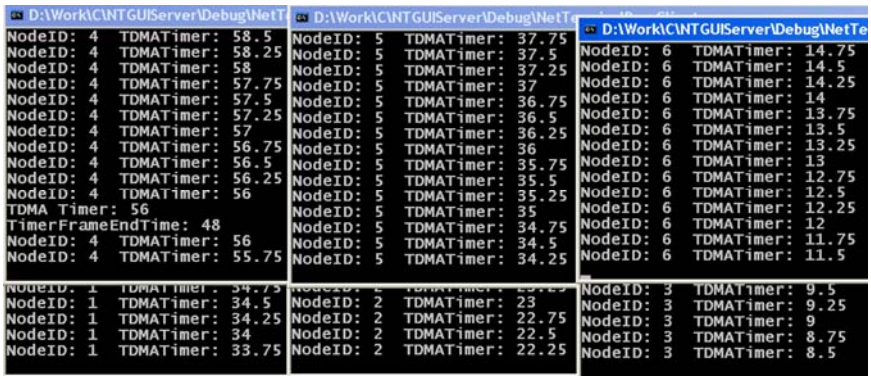

Figure 10: Nodes and their corresponding observed time-epochs prior to time-synchronisation.

Time-synchronisation was then initiated by manually sending a command to node 6. The following chain of events occurred:

a) Node 6 transmitted a time-broadcast.
b) Nodes 1 and 5 both received the broadcast, and waited for their respective time-slot to perform a ranging operation with node 6.
c) After a successfully ranging operation, nodes 1 and 5 each rebroadcast their new time, tagged with the ID of node 6.
d) Node 2 then perform ranging with node 1, while node 4 with node 5, after each receiving the respective time-broadcasts.
e) This process repeated until the boundaries of the network was reached at nodes 2 and 3.

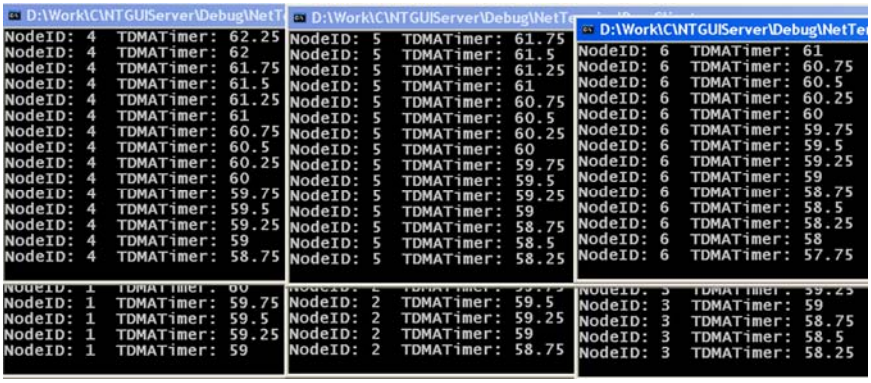

Figure 11: Nodes and their corresponding observed time-epochs after time-synchronisation.

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-epochs of nodes prior to and after time-synchronisation.

The results are shown in Fig. 11 and listed in Table 1. The simulation was repeated for a further ten times with random initial time-epoch for each node. In all ten cases, the calculated standard deviations of the observed time-epochs after time-synchronisation are within the range of 0.41 to 0.68 seconds.

6 CONCLUSION

A typical ad hoc underwater network would consist of a relatively small number of nodes deployed sparsely in order to fully utilise the communication range of the acoustic modems. Therefore, data packets are likely to travel across multiple hops. The inter-node distance is bounded by factors such as sea-state, noise level, and the cost per transmission in terms of energy consumption.

A simple time-division protocol is described, where time-slots allocated to the nodes allow the completion of one transmission session by taking into account the maximum propagation delay between adjacent nodes. This method increases scalability, but requires time-synchronisation across multiple hops to be performed in the absence of a central server. Therefore, an ad hoc time-synchronisation method is presented. This enables all the nodes within the network to be synchronised even if they have no direct acoustic links to one another. Simulations with a network of 6 nodes were carried out to demonstrate the accuracy of the method under the ideal channel assumption. The results demonstrated that the standard deviation of the time-epochs after time-synchronisation is well under 1 second, and is only a small fraction of the length of a typical time-slot. This suggests that the method is a viable solution to the problem of achieving time-synchronisation in an ad hoc network.

7 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. 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.
4. 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