Oceans '06, Singapore
(authors' manuscript – do not distribute)

ANU
THE AUSTRALIAN NATIONAL UNIVERSITY

Effective Communication in Schools of Submersibles

Felix Schill✦⅄ *& Uwe R. Zimmer*✦

✦Research School of Information Sciences and Engineering
Autonomous Underwater Robotics Research Group
The Australian National University, ACT 0200, Australia

⅄National ICT Australia Ltd., Locked Bag 8001, Canberra ACT 2601, Australia[1]

felix.schill@anu.edu.au | uwe.zimmer@ieee.org

Effective communication mechanisms in schools of submersibles are a key requirement for their meaningful deployment. Furthermore a fully distributed communication schema is preferable for reasons of reliability. The required communication form is usually many-to-many, especially omnicast [4] (or 'gossiping'). All these constraints are hard to achieve at the same time in a low-bandwidth, short range communication setup. The theoretical findings in [4] and [5] are expanded and employed for the actual underwater schools communication in the Serafina project (e.g. [2], [3]). The results are reported here.

1. Introduction

Short range, long-wave radio and optical communication links are deployed to build up and maintain a dynamical communication schedule which allows for a highly efficient exchange of all locally sampled information to all members of the school.

Short range, long-wave radio and optical communication links are deployed to form a multi-hop radio network in a school of submersibles. The nodes of this network build up and maintain a dynamical communication schedule which allows for a highly efficient exchange of all locally sampled information to all members of the school. Wireless communication under water, especially in seawater, is very limited. Commonly available high-frequency radio links such as 802.11b or BlueTooth encounter extremely high attenuation, which makes them unusable. Three methods of wireless communication are known for under water setups: acoustic (sonar modems), optical (using blue/green light, refer to [6]), and long-wave radio. In this paper we will focus on electromagnetic communication channels only. The experimental section of this paper has technical details about the transmitter and receiver modules used for the experiments. A common problem of all known wireless channels under water is that their bandwidth and range are very limited. Due to the bandwidth limitations, these links are usually only half-duplex, and have only one single channel. These constraints favour Time Division Multiple Access (TDMA) protocols, which efficiently assign time slots for sending to nodes without wasting valuable bandwidth.

Being able to efficiently distribute and share locally available data from all nodes with the entire network is a problem rarely discussed in classical communication theory. Many network applications require dedicated links between designated nodes, or have a hierarchical structure. However, in a school of identical submersible robots, there is no hierarchy. Also, local sensing only delivers sparse information about the environment, especially in an underwater setup. It is therefore crucial that submersibles share their local information with their direct neighbourhood (local control information, positions, etc.) and the whole school (measured gradients, maxima and minima of sampled data, global control parameters, etc.). The most suitable form of communication is "many-to-many", or *omnicast* (refer to [4], [5]). In some contexts this is also referred to as *gossiping*. This paper discusses the design and implementation of an efficient TDMA scheduling algorithm, and its performance regarding omnicast.

For theoretical considerations the multi-hop radio network is modelled as a graph $G = (V, E)$. This model is very common in the literature. Vertices $v \in V$ of the graph represent communication nodes. Edges $(v_1, v_2) \in E$ represent a communication link, meaning that v_1 can send messages to v_2. For simplicity, we assume that links are symmetric – $(v_1, v_2) \in E \Leftrightarrow (v_2, v_1) \in E$. A node v_i receives a message if and only if exactly one node v_k in the direct neighbourhood of v_i $(\{v_k \in V | ((v_k, v_i) \in E)\})$ sends a message in some time interval. If two or more

---

[1] *National ICT Australia is funded through the Australian Government's "Backing Australia's Ability" initiative, in part through the Australian Research Council*

nodes in the neighbourhood of v_i send at the same time, it is usually assumed in the literature that v_i will not be able to distinguish the two messages from noise. This is called a *collision* in node v_i. This paper follows this common network model for theoretical considerations, but will address the differences between model and reality in the experimental section.

2. Omnicast communication

The required many-to-many communication form for distributed sampling and control of schools of submersibles can be formalized as:

Definition 1: (the omnicast problem) *Let* $G = (V, E)$ *be a graph describing a communication network with* $n = |V|$ *nodes. In the start state, every node* $u \in V$ *has a set* $I_u(t_0)$ *of information tokens, which contains exactly one unique token* B_u *of information. During the communication phase, a node* v *updates its set* $I_v(t+1) = (I_v(t) \cup I_w(t))$*, if and only if it successfully receives a message from* $w \in V$ *in time step* t*, and* $I_v(t+1) = I_v(t)$ *otherwise. The end state* t_f *is reached, when all nodes have the full set with all tokens,* $I_u(t_f) = \{B_v | v \in V\}$ *for all* $u \in V$.

whereas the optimal solution is specified as:

Definition 2: (***the optimal omnicast problem***) *Find a schedule* $S_G = (T_1, \ldots, T_t)$, $T_i \subset V$ *for* $i = 1, \ldots, t$, *with* T_i *being the set of sending nodes in time step* i, *such that* S_G *solves the omnicast on the network graph* G, *and* t *is minimal.*

Previous publications discussed the theoretical side of omnicast communication [5]. It could be shown that a general upper bound for optimal solutions in multi-hop wireless networks with n nodes is $2n-2$. It has also been pointed out that optimal solutions are not necessarily collision-free. An interesting point has been made, that fully connected networks actually perform worse with regard to omnicast than some connected networks with a smaller in/out-degree. Fully connected networks require n steps to achieve omnicast. There are networks for which a solution exists with less then n steps.

3. Distributed Dynamical Omnicast Routing

The theoretical analysis assumes full knowledge of the network topology. In the case of a real network, there is no entity which has this global knowledge. The information which is available to each node is only the information contained in messages which they receive. In practice this means that nodes can become aware of their 2-hop neighbourhood.

A detailed description and discussion of the DDOR method can be found in [4]. Continuous real-time simulations in a multi-tasking environment demonstrated distributed schedules which were always better than the known upper bound for optimal solutions ($2n-2$), and usually close to n and smaller than n (persistently smaller than n for heterogeneous networks with more than 40 nodes).

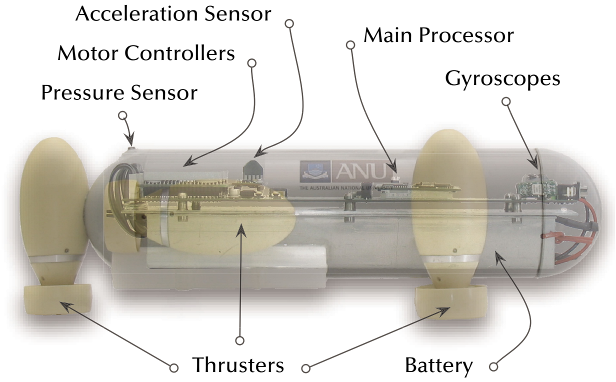

figure 1: One of the autonomous submersibles: Serafina

Distributed measurements of Omnicast performance

These definitions consider the static case in which all members synchronize before a new set of information tokens becomes available. In a practical setup this synchronization phase is omitted so that the individual omnicast schedules actually overlap and a measurement of a complete omnicast schedule (i.e. its 'roundtrip'-time) can only be done in a distributed and averaging way.

Furthermore each submersible sends its most recent set of data/measurements, so that the criterion of a static, optimal omnicast is replaced by a measurement of the maximal communication delay between nodes in a certain topological distance in the communication network.

As a formal basis this article specifies and discusses the required practical performance criteria for schools of underwater vehicles. Implications from earlier distributed scheduling works [1] and more recent schemes [4] are embedded and discussed in the context of practically achievable performances.

4. Experiments

After having shown the performance of the Distributed Dynamical Omnicast Routing algorithm, experiments have to be carried out to test the validity of the underlying network model, and to verify the performance of DDOR in a real distributed environment.

4-1. Longwave radio modules

Longwave radio modules were designed and built for the experiments. These modules operate on a carrier frequency of 122.88 kHz, and employ differential

Medium	Conditions	Bit rate [bit/s]	Range [mm]
Air	Inside building	4096	4150
		1024	5400
Fresh water (chlorinated)	Pool (2 m deep) in 1.5 m depth	4096	4000
Salt water (Pacific)	Coastal water (4 m deep), in 3 m depth	1024	3900

figure 3: Range measurement results in various conditions

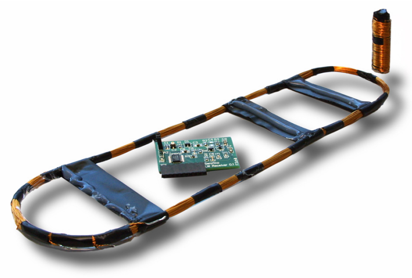

figure 2: Long wave radio transmitter inside Serafina

binary phase keying for digital modulation. The data bit rate can be chosen to be 1024, 2048, 4096 or 8192 bits per second. The modulation and generation of the carrier wave is done in software on a RISC microcontroller, using a number of hardware timers. The microcontroller runs at 9 MHz and offers 32 kilobytes of flash memory and 2 kilobytes of RAM. The generated modulated carrier is amplified and drives a tuned resonance circuit, which acts as the transmitting antenna. The receiver uses a second antenna. The signal is pre-amplified, and decoded by a dedicated longwave decoder chip, which passes on the received data to the same microcontroller that implements the transmitter. The complete modules are both transmitter and receiver.
The actual data communications is packet-based. A packet can be sent in one time slot, and consists of multiple frames, which can be up to 16 bytes long. Contained in a packet are the sender identity, the logical clock value, the schedule currently used by the sender, and the user data.

The receiver rejects frames with invalid headers. For better data integrity, the scheduling and sender information are protected by check-sums (8 bit per frame). For performance reasons, the user data block is not error-checked on the hardware layer. If needed, this can be done on higher protocol levels.

The current implementation of DDOR on the microcontroller allows a schedule length of 8 schedule slots and up to 254 nodes. Packets can be up to 256 bytes long. The duration of a time slot is calculated according to the chosen packet length and bit rate. The overhead per packet for omnicast is 14 bytes for sender information and the schedule, plus an additional 3 bytes per frame used. The binary code is less than 8 kilobytes

While not transmitting, the module consumes less than 20 mA current at 5 V. The output amplifier operates currently at 5 V output voltage amplitude, consuming 20 mA. This corresponds to 100 mW power. The transmitter output power can easily be increased by changing the drive voltage amplitude to up to 20 V, but this requires currently an external 20 V voltage supply, which complicates testing. The next design iteration will include an integrated, adjustable boost converter for dynamically changing the drive voltage to suit the conditions.

4-2. *Range measurements*

Obviously range measurements always depend on many parameters (output power, antenna efficiency and tuning, noise) and the environment. The following measurements are only examples to provide design guidelines to meet application requirements. The definition of range is the maximum distance between transmitter and receiver that provides less than 10 % frame dropout. It should be noted that for this particular implementation, the increase in frame drops is very sharp – typically an increase in distance of less than 0.3 metres corresponds to a transition from virtually error-free reception to total loss of reception. That means that within the given range, reception can be assumed to be quite reliable. The following experiments were carried out with a 0.1 meter diameter circular transmitting antenna, and a 10 mm diameter 50 mm length ferrite core receiver antenna. The drive voltage amplitude for the transmitter was 5 V. Refer to figure 3 for the results.

It can be seen that the effect of water is minimal. Even strongly conductive sea water has only a small impact on the range. The actual range might seem short, but it lies well within our expectations. For the final application, the drive voltage will be increased to 20 V, which will almost quadruple the range of the modules. A communication range of more than 10 meters is sufficient for a school of miniature submersibles, which are only 50 cm long.

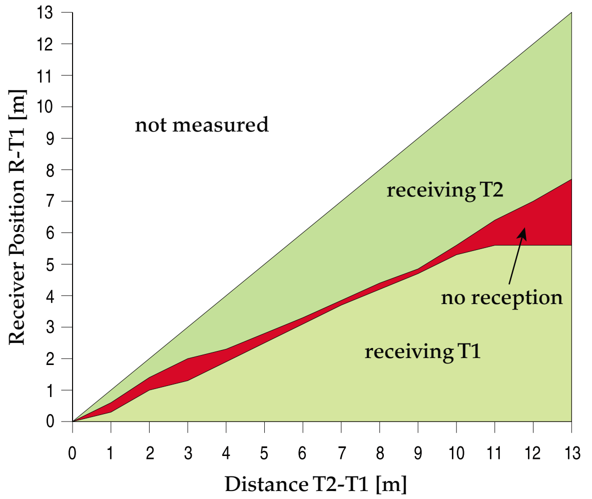

figure 4: Reception zones and collision zone

4-3. Interference measurements

An important question is how transmitters interfere with each other, and how receivers experience message collisions. For this experiment, two almost identical transmitters T1 and T2 were placed with varying distance from each other in one meter increments. Both transmitters have approximately 5 m range, and are continuously sending out frames containing their own identifier. A receiver was placed on the straight line connecting the two transmitters. The receiver was then moved on this line to determine the boundaries of the region were the receiver reliably receives T1 or T2, or has no reliable reception. The results are shown in figure 4.

It may be surprising to see that the zone in which collisions occur is actually quite narrow. The theoretical network model which is commonly used in the literature would have predicted collisions for the entire regions where the ranges overlap. In reality, the receiver can easily decode the respective stronger signal of the closer transmitter. Only in the middle between the transmitters, there is a narrow zone where reception is distorted. This obviously means that the assumed theoretical network model is wrong in the assumption that a node only receives a message if exactly one of its neighbours is transmitting within some time interval. A better model would be that it receives the message from the closest transmitter which is currently sending, if no equally close transmitter is sending at the same time. However, since there is no direct way for measuring these distances, nodes are generally unaware of the geometric arrangements of the network. Only the topological information can be reliably determined, which leads again to the common graph model.

The remaining question is now how this affects the performance of omnicast scheduling algorithms. This question is easy to answer for the case of already established, collision-free schedules, such as global solutions or the schedules that DDOR will eventually produce after a stabilization period. If a schedule is collision free in the graph model, then it will still be collision free in reality. The performance is not affected at all.

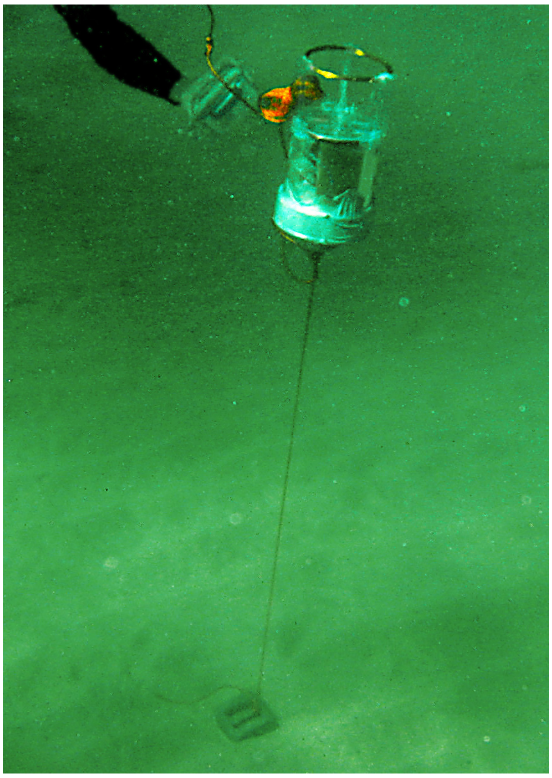

figure 5: Long wave radio experiments in the Pacific Ocean

The second case are established schedules containing collisions in the graph model. These schedules will contain less collisions in reality - only those collisions remain, where a receiver is exposed to two equally distant sending transmitters. With regard to omnicast performance, this means that the performance will not be worse, but might be better, since more information that theoretically assumed is in fact exchanged.

It is much harder to answer the question how this difference between model and reality will affect the stabilization of a schedule in a distributed environment. Nodes will receive messages which they would not have received in the theoretical model. In the case of DDOR, this does not affect the calculation of schedules, which is based only on the information which nodes lie in the 2-hop neighbourhood. The information which can be gained by *not* receiving a message which should have been received is not exploited by DDOR. The calculated schedules will eventually be collision free. Therefore the different conditions in reality will not affect the performance of DDOR - however, the start-up phase in reality might differ from the simulations. Since more messages will be received, which are lost in the simulation, the stabilization is expected to be quicker.

5. DDOR in the physical world

The complete DDOR protocol has been implemented in four water proven, and battery powered radio transceivers. All possible combinations of network topology and joining/leaving sequences of individual nodes have been tested. A compact, collision-free schedule has been created in all configurations and in all nodes.

More interesting are the configuration, and re-configuration times which could be achieved. The joining time depends on the lengths of the established schedules in this neighbourhood. In the four node system, this means that after a maximum of 15 schedule-slots each nodes can be joined into the network. Nodes are deleted after five missed schedule-slots. Even considering extremely low bit-rates, all those reconfiguration procedures are completed after a few seconds of real-time, which is perfectly acceptable for dynamic underwater school-communications. The achieved schedule performance is identical with the simulated measurements (as reported in [4]): always less than $2n-2$, usually close to, or smaller than n.

6. Conclusions

A distributed scheduling algorithm suitable for schools of underwater vehicles has been presented and implemented. It could successfully be shown that the Dynamical Distributed Omnicast Routing (DDOR) algorithm performs well with regard to information dispersal both in simulations and in reality. The algorithm requires only a small amount of memory, computations and communication overhead, which makes it suitable to be implemented on cost-efficient, low-power, low-bandwidth longwave radio modules. Experiments could show that digital longwave radio is a suitable method for wireless communication in sea water, and that the range is only minimally affected compared to transmission in air.

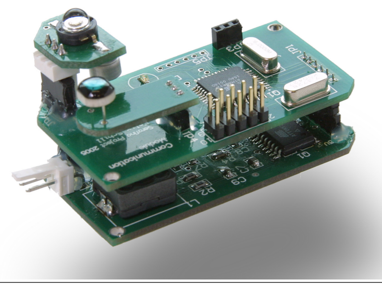

figure 6: LED communication transmitter/controller

Further experiments regarding interference of two transmitters showed that the commonly assumed method to model a wireless network using a graph is imprecise. The correct modelling of collisions requires geometrical information. However, the traditional graph network model is still valuable for computing schedules without collisions - these schedules will remain collision free in reality as well.

References

[1] Ted Herman and Sébastien Tixeuil; *A Distributed TDMA Slot Assignment Algorithm for Wireless Sensor Networks*; In ALGOSENSORS, pages 45–58, 2004.

[2] Shahab Kalantar, Uwe R. Zimmer; *Distributed Shape Control of Homogeneous Swarms of Autonomous Underwater Vehicles*; to appear in Autonomous Robots (journal) 2006

[3] Shahab Kalantar, Uwe R. Zimmer; *Control of Open Contour Formations of Autonomous Underwater Vehicles*; International Journal of Advanced Robotic Systems, Dec.'05

[4] Felix Schill, Uwe R. Zimmer; *Distributed Dynamical Omnicast Routing*; to appear in Intl. Journal of Complex Systems '06

[5] Felix Schill, Uwe R. Zimmer & Jochen Trumpf; *Towards Optimal TDMA Scheduling for Robotic Swarm Communication*; Proceedings of the TAROS (Towards Autonomous Robotic Systems) intl. conference, September 12-14, 2005, London

[6] Felix Schill, Uwe R. Zimmer & Jochen Trumpf; *Visible Spectrum Optical Communication and Distance Sensing for Underwater Applications*; Proceedings of the Australasian Conference on Robotics and Automation, December 6-8, 2004, Canberra, Australia