

Decentralized Learning in Wireless Sensor Networks

Mihail Mihaylov ^a

Karl Tuyls ^b

Ann Nowé ^a

^a *Vrije Univesiteit Brussel, Brussels, Belgium*

^b *Maastricht University, Maastricht, The Netherlands*

Abstract

In the full version of this paper [1] we present a reinforcement learning algorithm with the aim to increase the autonomous lifetime of a Wireless Sensor Network (WSN) and decrease latency in a decentralized manner. WSNs are collections of sensor nodes that gather environmental data, where the main challenges are the limited power supply of nodes and the need for decentralized control. To overcome these challenges, we make each sensor node adopt an algorithm to optimize the efficiency of a small group of surrounding nodes, so that in the end the performance of the whole system is improved. We compare our approach to conventional ad-hoc networks of different sizes and show that nodes in WSNs are able to develop an energy saving behaviour on their own and significantly reduce network latency, when using our reinforcement learning algorithm.

1 Introduction

To handle the problems and limitations of contemporary Wireless Sensor Networks (WSNs), we use a reinforcement learning algorithm to optimize the energy efficiency of a WSN and reduce its latency in a decentralized manner. We achieve that by making nodes (hereby regarded as agents) develop energy-saving schemes by themselves without a central mediator. The idea behind this approach is that agents learn to reduce the negative effect of their actions on other agents in the system, based on a certain reward function. We show that when agents learn to optimize their behaviour, they can increase the energy efficiency of the system and significantly decrease its latency with minimal communication overhead.

Since communication is the most energy expensive action [2], it is clear that in order to save more energy, a node should sleep more. However, when sleeping, the node is not able to send or receive any messages, therefore it increases the latency of the network, i.e., the time it takes for messages to reach the sink. On the other hand, a node does not need to listen to the channel when no messages are being sent, since it loses energy in vain. As a result, nodes should learn on their own the number of time slots they should spend sleeping within a time window.

2 Learning Algorithm

Each agent in the WSN uses a reinforcement learning (RL) algorithm to learn an optimal schedule (i.e. sleep duration in a frame) that will maximize the energy efficiency and minimize the latency of the system in a decentralized manner. The **actions** of each agent are restricted to selecting a sleep duration for a given time window (called a frame). Agents choose their actions according to a probability distribution and use that action for a certain number of frames, in order to evaluate the effect of that action on the system. This **evaluation** is done with respect to neighbouring agents, under the belief that if each agent “cares about others” that will improve the performance of the whole system. To achieve that, we set the reward signal of each agent to be equal to the mean efficiency of its neighbours. Thus, by optimizing their own behaviour, agents will increase the performance of their surrounding nodes. After evaluating the effect of its action, each agent **updates** its action probabilities using the update rules of a classical learning automaton. This algorithm allows the learning process to be executed on-line – the algorithm adapts to the topology of the network and the traffic pattern, which typically cannot be known in advance in order to train nodes off-line.

3 Results

We evaluated our algorithm on two random topology networks of the same density, but of different sizes. We compared the performance of each setting to a network of the same size where agents do not optimize their behaviour, but rather all sleep the same pre-defined amount of time. In other words we compared the optimal “non-learning” system to the optimal one *with* learning.

In both networks we measured a moderate increase in energy efficiency and a significant decrease in latency. In the **small network** (10 nodes, 3 hops) we achieved a 10% increase in energy efficiency, while the *maximum* latency for a packet was reduced by 70%. The effect of adapting to the traffic pattern was even more pronounced in the **large network** (50 nodes, 7 hops), where agents were able to decrease the *average* latency by over 70%, resulting in three times more packets delivered to the sink, compared to using fixed schedules.

It can be concluded from our results, that agents learn to avoid “harming” other agents by adapting to the traffic pattern and therefore learning the optimal sleep duration in their neighbourhood. In other words, agents learn to sleep when their neighbours communicate (so as to avoid overhearing); stay awake enough to forward messages quickly (and thus decrease latency); and yet sleep enough (to ensure longer network lifetime).

4 Conclusion

In this paper we used a reinforcement learning algorithm to improve the performance of Wireless Sensor Networks (WSN) in a decentralized manner, in order to prolong the autonomous lifetime of the network and reduce its latency. We were able to show that when agents in a WSN use an algorithm for optimization, they can learn to reduce the negative effect of their actions on other agents in the system, without a central mediator. Our results indicate that both in a small and large network, agents can learn to optimize their behaviour in order to increase the energy efficiency of the system and significantly decrease its latency with minimal communication overhead. Our results outperformed a conventional ad-hoc network, where all agents equally listen and sleep for a pre-defined amount of time. Thus, based on our experiments we can conclude that it is more beneficial for the sensor network when nodes *learn* what actions to take, rather than follow a pre-defined schedule. In our algorithm each node seeks to improve not only its own efficiency, but also the efficiency of its neighbourhood, which ensures that the agents’ goal is aligned with the system goal of higher energy efficiency and lower latency.

References

- [1] Mihail Mihaylov, Karl Tuyls, and Ann Nowé. Decentralized learning in wireless sensor networks. In *Proceedings of the Adaptive and Learning Agents Workshop (ALA 2009)*, Budapest, Hungary, 2009.
- [2] T. van Dam and K. Langendoen. An adaptive energy-efficient mac protocol for wireless sensor networks. In *Proceedings Of The First International Conference On Embedded Networked Sensor Systems*, pages 171 – 180, Los Angeles, California, USA, 2003.