Chapter 3

Keep the messages flowing

How do you maximize the lifetime of a network of communication devices? The study group investigated approaches inside and outside the usual framework. After two days, they had an answer that gave Thales unexpected new ideas.

Imagine a group of firefighters at a big disaster site. Each firefighter has a device for communicating with the rest of the group. Every time one of them sends a message, for instance a position update, the information must be broadcasted to all the others in a wireless network. The communication devices have a limited battery capacity, and sending a message takes energy. How should the messages be sent through the network to keep the flow of information running for as long as possible?

This is an example of the more general problem that Maurits de Graaf from Thales brought to the study group. De Graaf develops algorithms that increase the lifetime of wireless ad-hoc networks of communicating nodes. In this setting the lifetime of a network is defined as the time until the first node runs out of energy. The goal is to find a broadcasting protocol that maximizes this lifetime.

De Graaf had three questions for the study group. The first one was about the current heuristic algorithm that Thales uses, how far is that from the optimal solution? The second was about a more intricate model for sending the messages and the third about variations in battery consumption. In preparing his questions, he tried to make them appealing to mathematicians from different areas. "I hoped for a diverse group, to get a new and fresh perspective."

Mathematician Nikhil Bansal from Eindhoven University of Technology joined the Thales team because the questions were theoretically most interesting: "It is a precise and technical problem with nice applications." His group started working on the first question about the quality of the heuristic algorithm from Thales. The mathematicians made two basic assumptions about the network. First they assumed it is stationary, so the connections between the nodes do not change over time. Secondly, they assumed battery use is linear and sending a message decreased the battery level by one. Receiving a message does not consume energy.

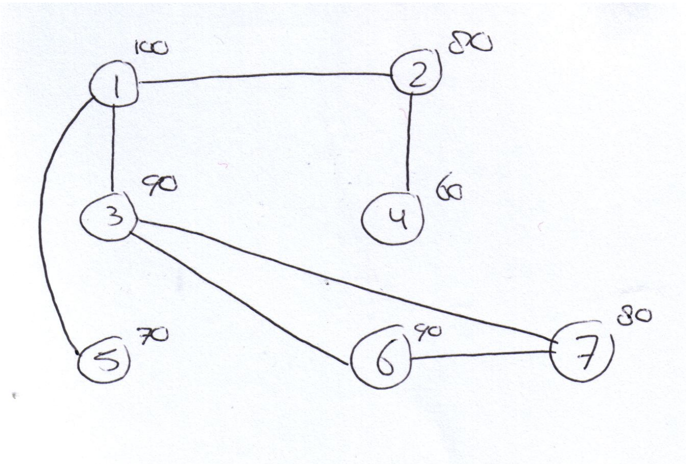

An example of a network with seven nodes. Next to each node is the current battery level (so node 3 has level 90 left). Nodes are neighbors if they are connected, so node 1 and 2 are each others neighbors, but 3 and 4 are not.

Willing and able

Thales currently uses the Maximum Willingness Heuristic (MaxWill for short). This algorithm tries to use the neighboring nodes with the highest remaining battery level for relaying messages. A key feature of MaxWill is its layered structure. A node that is broadcasting automatically sends the message to all its neighbors, the nodes that are just one hop away in the network. MaxWill selects a set of these neighbors as relays such that all the nodes that are two hops away will receive the message. The one-hop-neighbors can be seen as the first layer, the two-hop-neighors as the second. Nodes with higher battery levels get selected first as relays, because they are most willing to sacrifice their batteries.

Nikhil Bansal: "We soon found some examples that showed that MaxWill is not always optimal. We devised another algorithm that seemed to do better." Their new algorithm carefully avoids nodes with the lowest battery level. This results in a longer life-time of the network, because more messages can be sent before the first node runs out of battery. In a number of simulations the new algorithm gave on average a 60

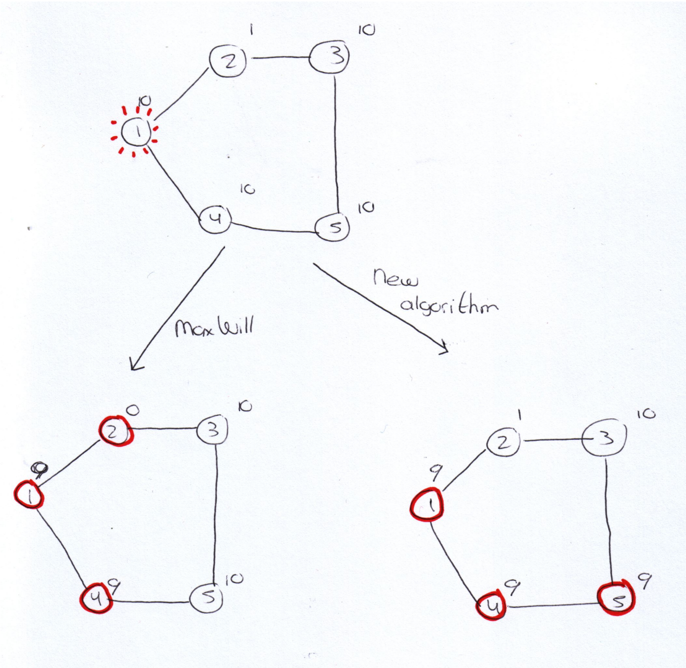

With MaxWill this network is down after sending one message from node 1, because its neighbors 2 and 4 have to relay the message to reach all two-hop neighbors. So node 2 will immediately run out of battery. In the new algorithm, nodes 4 and 5 relay the message and it is possible to send at least one more message.

A hard problem

The new method does not always find the optimal solution. In fact, it is impossible to find a fast algorithm that does this. The mathematicians proved that the problem of maximizing the network life-time is NP-hard. This is a class of problems for which a solution can be checked efficiently, but for which there is no fast algorithm for finding such a solution. You may think of a sudoku puzzle. It takes quite some time to finish one, but it is really easy to quickly verify that a given solution is correct.

So there is no algorithm that rapidly gives the optimal way of sending messages. Even more depressingly, the study group showed that it is also NP-hard to approximate the best solution. So every fast algorithm will find a solution that is at least a factor worse than the optimal lifetime.

The study group emailed their results to Thales on Tuesday night, feeling that they had solved the first problem and were ready to move on to the other two questions. Bansal: “But on Wednesday Maurits de Graaf came to Eindhoven explained that their real problem was slightly different. The demand that all the second-hop-nodes were covered after the first relay was not a choice from their algorithm, it was a constraint.

This layered approach was obliged in the way all their systems were implemented. Our algorithm was impossible to use for Thales, we picked the wrong relayers."

It took De Graaf some time to appreciate the new algorithm: "My first reaction was that this was all wrong. However, the more I thought about it, the more I liked this new simple mechanism. We would not have come up with something like this at Thales, because the idea is different from the framework we are currently using. However, the algorithm from the study group performs so much better that we now consider adapting or even leaving the framework."

During a forest fire the firefighters update each other on the fire details at different locations.

Layered approach

After De Graaf's visit the study group went back to the first question, but now with the layers as in the preconditions. To compare MaxWill with the optimal solution, the problem was reformulated as a linear program. In this setting there is a target number of rounds of messages. A round is series of broadcasts where each node occurs exactly once as the source. It is a reasonable assumption that the nodes all have to send regular updates, because they have to update the information from their location. The problem is now to decide if there is a way to send the target number of rounds through the network. When the number of nodes is not too large, this decision problem can be solved rather fast. Repeating this procedure for higher and higher numbers of rounds until the problem becomes infeasible, yields the optimal number of rounds.

To compare the results from MaxWill with the optimal solution a number of networks were randomly generated. Two nodes had a probability of 50Simulations with other types of networks showed that MaxWill performs best in sparse networks.

In their final presentation the mathematicians apologized for not answering the other questions: "We had too much fun with problem one to start on the rest, but if the study group had been one week longer we would have done more." Their paper gives some first ideas for the other questions, but these topics would deserve a separate article.

Two directions

For the first question the final conclusion of the study group is twofold. If Thales decides to switch to a protocol that does not require the layered approach, their new algorithm performs much better than MaxWill. If they want to stay within the current framework, the linear problem indicates in which types of networks MaxWill performs well and not so well. For the latter cases another approach would be useful. For small networks the linear program can serve as a better heuristic.

Thales decided to further investigate both directions. Maurits de Graaf: "A student is coming to do an internship at Thales, he will find the report from the study group on his desk. Part of his task will be to look at the impact of dispensing our framework. Can we retain the good properties of the current setup and combine them with the new algorithm? And if we remain within the framework, how can we use the ideas from the study group to improve the MaxWill heuristic?" He concludes that the study group is a great initiative: "Even if you can not directly implement the results, you gain a better insight and departure points for further research."

Team Thales

Nikhil Bansal, David Bourne, Murat Firat, Maurits de Graaf, Stella Kapodistria, Kundan Kumar, Corine Meerman, Mihaela Mitici, Francesca R. Nardi, Björn de Rijk, Suneel Sarswat and Lucia Scardia.