Publish/Subscribe for Wireless Sensor Networks

Claude Chaudet
Télécom ParisTech
CNRS LTCI UMR 5141
Paris, France
Claude.Chaudet@telecom-paristech.fr

Isabelle Demeure
Télécom ParisTech
CNRS LTCI UMR 5141
Paris, France
Isabelle.Demeure@telecom-paristech.fr

Salma Ktari
Télécom ParisTech
CNRS LTCI UMR 5141
Paris, France
Salma.Ktari@telecom-paristech.fr

Nicola Costagliola
Télécom ParisTech
CNRS LTCI UMR 5141
Paris, France
Nicola.Costagliola@telecom-paristech.fr

Samuel Tardieu
Télécom ParisTech
CNRS LTCI UMR 5141
Paris, France
Samuel.Tardieu@telecom-paristech.fr

ABSTRACT

This extended abstract overviews some of the advantages and problems to overcome when implementing a publish/subscribe system over a wireless sensor network. Publish/subscribe usually relies on a set of intermediate nodes, brokers that gather and dispatch all the packets emissions. The presence of these brokers lightens the duty of most regular nodes, as they are the only ones who really need to know nodes identities and to maintain long-range routes. However, as computing power, memory, energy and channel bandwidth are scarce resources in WSN, defining an efficient brokering architecture is not trivial.

After presenting publish/subscribe in a wireless sensor network context, we examine various performance criteria and discuss on the brokering overlay. If a multi-brokers architecture seems necessary to distribute the load across the network, questions such as the correct number of brokers and their location are difficult when taking into account multiple criteria and that brokers need to remain synchronized. The dynamics of the system also requires attention, as brokers may appear or disappear according to the network traffic patterns or to battery levels.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Distributed Networks

Keywords

Middleware, publish/subscribe, wireless sensor networks

1. WSN AND PUBLISH/SUBSCRIBE

Wireless sensor networks (WSN) are composed by numerous low-power embedded devices that gather measurements and report them to a few collection points using a spontaneous wireless multihop ad hoc network. All the nodes are traditionally supposed to be identical, sensing and acting as communication relays, even if some more powerful nodes can sometimes endorse the role of local coordinators.

These networks are supposed to be autonomous, to require little maintenance and to last as long as possible. The nodes thus rely on a battery that should not be recharged, or at least not too often. Limiting energy consumption is a primary objective that in influences most of these networks conception choices. Alongside with production cost limitation, this explains why, at the hardware level, these nodes are generally equipped with a slow micro-controller, a low-throughput radio interface and a very limited memory space. Communication is known to be a major source of energy consumption and all strategies that aim at reducing the communication volume are expected to improve the network lifetime.

Publish/subscribe appears as an efficient communication scheme for WSN, as it introduces intermediate nodes, the brokers, whose role is to gather all the emissions of the

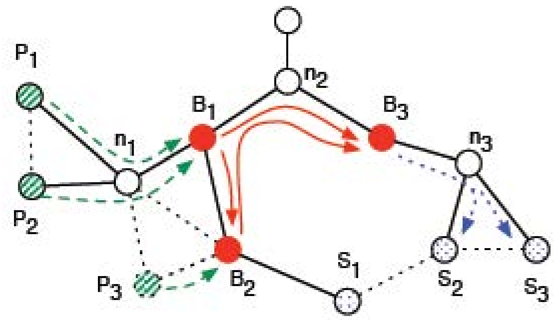

Figure 1: Example scenario

sensor nodes (the *publishers*), manipulate this information (filtering, compression aggregation, …) and ultimately transmit the resulting data to the collection nodes (the *subscribers*). The brokers are the only nodes to associate data and the receivers of this data, so emitters and receivers do not need to know their respective identities and location and do not need to be timely synchronized. As in a clustered architecture, addressing and routing require less overhead: a node only needs to maintain a route towards one broker while the brokers form an overlay to route packets towards their true destinations using any underlying routing protocol, unicast or multicast if available.

Figure 1 illustrates such a scenario in which three publishers, *P1*, *P2* and *P3*, send measurement results (dashed lines) to their closest brokers (*B1* and *B2*). The brokering system is composed here of three nodes among which B2 and B3 registered some subscriptions from nodes S1 to S3. The different brokers then exchange packets together (plain lines), possibly using publish/subscribe between them, in order to conciliate the required data so that B2 can send the data flow to *S1* and *B3* to *S2* and *S3* (dotted lines). Inside the brokering system, any data manipulation algorithm can be applied (aggregation, fusion, filtering, encryption, etc.), therefore the flow conveyed to the subscribers is not necessarily the sum of the incoming flows.

1.1 Brokers overlay and performance

All communications pass through the brokers overlay and the shape of this overlay - i.e. the number of brokers and their location - and its maintenance shall have a significant impact on all the network performance criteria. To a certain extent, brokers share the load and increasing their number should lead to a fairer distribution of the resources, and to a better balance of energy consumption among these nodes.

Communication is the dominant source of energy consumption for sensor nodes. Data packets should thus take the shortest paths from publishers to the brokers and from the brokers to the subscribers in order to reduce the number of re-emissions and minimize the load on the network. This would militate for bringing brokers close to publishers and subscribers, leading to a dense overlay. However, over a wireless channel, all frames are broadcasted and filtered by the receivers based on the MAC addresses. The energy necessary to emit a frame is comparable to the energy necessary to receive it due to decoding and noise filtering. Therefore, the consumption of a given node does not only depend on the number of frames it emits and forwards, but also on the number of frames emitted in its neighborhood. Brokers form the core of the network, they therefore attract a lot of traffic and all nodes close to a broker will see their energy consumption increase with channel activity, which militates for a sparse overlay of brokers.

As long as energy only is concerned, a compromise exists between routes lengths and the number of collateral receptions. Finding this optimal point requires solving an optimization problem roughly described as finding the overlay that minimizes the average energy consumption of sensor nodes, or the maximum energy consumption, or the consumption fairness, under a set of constraints defined by the forwarding / routing scheme and by the network topology. This problem is complex to solve and may require to examine all the possible brokers configurations and heuristics shall be necessary for implementation in an online scenario. Considering the additional performance criteria does not ease this task. Memory, for instance, is also a scarce resource, which is consumed by packets queues, but also by routing and publishers and subscribers tables. The amount of each node's memory that is dedicated to publish/subscribe limits the maximum load that a broker can support. However, a certain level of synchronization is necessary between the brokers and increasing the number of brokers does not lead to a proportional decrease in memory space, as some table entries need to be shared and thus replicated. A compromise, similar to the one that exists for energy, is applicable to memory independently.

Finally, classical network metrics such as delivery time, channel congestion, etc. also impose some constraints that shall shift the equilibrium point towards more or less brokers. These brokers should be well distributed, located in central areas to be close to most publishers and subscribers, but not in too dense areas in order to minimize collateral receptions. Various algorithms and local metrics, based for instance on centrality measurements, can be tuned to achieve the desired goal, but no precise objective can be defined without looking at the cost of adding or removing a broker.

2. BROKERS OVERLAY DYNAMICS

The brokers overlay is eager to evolve during the network lifetime, adapting to the network dynamics. A node should stop serving in the brokers overlay when its battery gets low, but the overlay shall also adapt to nodes mobility and to traffic mobility, for instance when publishers and subscribers appear or disappear.

Brokers then need to regularly monitor the network status and auto-evaluate their own suitability based on various criteria. If the performance criteria mentioned in the previous paragraph may constitute the base of the evaluation metric, brokers selection shall be based on a distributed election algorithm that bases decisions on a combination of topological and traffic-related metrics and in which some nodes may not be candidates to reflect the low-battery situation for example. Local metrics such as the egocentric vision of centrality, correctly weighted by the traffic allow to rank and compare candidates, but actually creating new brokers or deciding to pass the brokering duty to another node requires first to exchange publishers and subscribers tables, and second to make the network, or at least the concerned publishers and subscribers, aware of this change, which has a cost in terms of communication and hence energy.

3. CONCLUSION

In this short article, we exposed a few advantages, requirements and guidelines to implement a publish/subscribe system over a wireless sensor network composed of several low-end devices. Taking into account multiple nodes criteria such as energy, memory and network-related metrics make the optimal difficult to define and thus to find. Taking into account the system online operation and the maintenance cost of the associations tables influences us to favor stability at the cost of optimality.

4. ACKNOWLEDGEMENTS

This work is funded by the Diaforus project, granted by the French National Research Agency (ANR).