

# Quality-of-Service Trade-off Analysis for Wireless Sensor Networks

Rob Hoes<sup>a,b,\*</sup>, Twan Basten<sup>a</sup>, Chen-Khong Tham<sup>b</sup>, Marc Geilen<sup>a</sup>, Henk Corporaal<sup>a</sup>

<sup>a</sup>*Department of Electrical Engineering, Eindhoven University of Technology, Eindhoven, Netherlands*

<sup>b</sup>*Department of Electrical and Computer Engineering, National University of Singapore*

---

## Abstract

Quality of Service (QoS) support for wireless sensor networks (WSN) is a fairly new topic that is gaining more and more interest. This paper introduces a method for determining the node configurations of a WSN such that application-level QoS constraints are met. This is a complex task, since the search space is typically extremely large. The method is based on a recent algebraic approach to Pareto analysis, that we use to reason about QoS trade-offs. It features an algorithm that keeps the working set of possible configurations small, by analysing parts of the network in a modular fashion, and meanwhile discarding configurations that are inferior to other configurations. Furthermore, we give WSN models for two different applications, spatial mapping and target tracking, in which QoS trade-offs are made explicit. Test results for these applications and a heterogeneous WSN combining these two applications show that the models are accurate and that the method is scalable and thus practically usable for WSN, even with large numbers of nodes. Details are given on how to efficiently implement the algorithm.

*Key words:* Wireless Sensor Networks, Multi-objective Optimisation, Pareto Analysis

---

## 1. Introduction

Wireless sensor networks (WSN) have been suggested for a diversity of applications, such as observing animals in a nature park, assistance in rescue operations, in-home entertainment systems or to monitor people's health. All these applications require different types of sensor nodes and have diverse performance demands. A lot of research has been done on topics such as the physical design of sensor nodes, self-organisation, and communication protocols, with a strong focus on energy efficiency. Recently, researchers are also considering Quality of Service (QoS) provisioning to ensure that explicit performance targets are met. Surveys suggest that there is a need for a middleware layer that negotiates between an application and a network to match QoS demands and the availability of WSN resources [1, 2]. This is challenging, because QoS requirements are often conflicting, and furthermore, it needs adequate ways to predict the behaviour and performance of a possibly heterogeneous network of nodes, under various circumstances. This paper proposes a trade-off analysis method. The current work focuses on design-time QoS provisioning, but the method is suitable for a distributed, run-time implementation as well.

A sensor node can typically be configured in many ways, as a node has several settings, such as the transmission power of its radio and its sleep/wake schedule. Moreover, configuring a whole WSN implies setting the parameters for all nodes in the network. Each such configuration has a certain effect on the quality

---

\* Corresponding author. Address: r.j.h.hoes@tue.nl

metrics of the application, to which a user may have attached (QoS) constraints. This paper describes a novel method that determines sets of parameters for nodes in a WSN that satisfy or trade off multiple application-level objectives. The method is efficient and scalable to large networks, despite the inherent exponential complexity: for an example on a 900-node network with  $27^{900}$  possible configurations, the algorithm took less than 46 seconds to complete, while never more than about a million configurations needed to be compared at the same time. Our approach uses Pareto-algebraic methods as introduced by Geilen et al. [3], in order to explore trade-offs between system properties, and to incrementally compute a set of feasible configurations. Pareto optimality is used as a central criterion to compare configurations and discard the ones that cannot be optimal, to keep their number manageable.

The current paper is an extended version of an earlier publication [4], with the following new contributions:

- The QoS analysis has been optimised in terms of memory usage and run time. Practical implementation details are presented.
- New experimental results of the optimised algorithm are given, on a larger set of test cases including larger networks. Furthermore, additional data on memory usage and the effect of node degrees on the run time are presented.
- A detailed example is given that shows that the method can also be used for heterogeneous networks with multiple node types or applications.
- An optimisation is discussed that interleaves constraint checking with the algorithm, to discard non-conforming configurations as early as possible, and further improve the algorithm's run time.

Related work is reviewed in the next section. To demonstrate that the algebraic approach is practically useful for WSN, Section 3 gives explicit models of sensor nodes and the way they work together to perform target-tracking or spatial-mapping tasks. The method is not limited to these tasks: other tasks are analysable in essentially the same way. Section 4 reviews concepts from Pareto analysis, after which Section 5 introduces the QoS-optimisation method for WSN. Algorithm optimisations and implementation details are given in Section 6; Section 7 shows test results about timing and memory usage on networks of varying sizes, and analyses the accuracy of the obtained results. A comparison with a genetic algorithm is made. Finally, Section 8 elaborates on the application of the method on heterogeneous networks, and the way constraints are used to further improve the algorithm's complexity.

## 2. Related Work

Chen et al. [1] give an overview of recent approaches and challenges related to QoS support in WSN. There are some network protocols that offer QoS support [5, 6], often based on delay constraints, but only a few approaches look at application level demands. One example of the latter type attempts to guarantee a minimum reliability level while maximising network lifetime [7]. There is a clear need for a middleware system for WSN that supports application level QoS provisioning [2, 8]. However, this is still a mostly open research problem. Our algebraic framework for handling multiple QoS requirements at the same time is an important step in this direction. This paper focuses on design time trade-off analysis, but the compositional analysis is well suited for distribution, and hence run-time application. To the best of our knowledge, this approach is unique in the area of WSN.

The Pareto-optimality criterion is a general concept that originally comes from economics. The Pareto points of a system precisely capture all the trade-offs in a multi-dimensional optimisation space. In engineering, it is used, for example, in design-space exploration for embedded systems [9, 10]. The recent development of Pareto algebra [3] offers a very structured way of analysing the design space. More traditional ways to find Pareto optimal solutions include genetic algorithms [11] or related algorithms like tabu search. The disadvantages of these approaches are their random nature and the use of a flat search space: they may find some Pareto solutions fairly quickly, but may find only a limited subset of all points. Finding all solutions requires an exhaustive search of the whole space. Our algebraic approach intelligently searches the whole space in a hierarchical manner; it is able to quickly find all locally Pareto-optimal solutions at different levels and incrementally combine them into a complete set of solutions, thus providing a better basis for configuring WSN under often conflicting QoS constraints. Q-RAM [12] is another framework that uses the

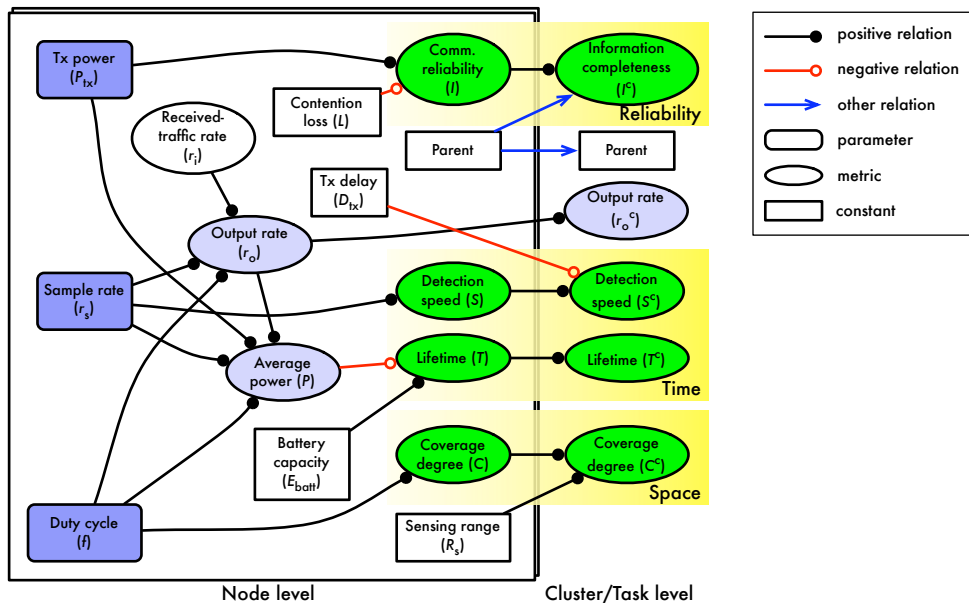


Fig. 1. Hierarchical trade-off model: relations between parameters (left), node-level quality metrics and task-level quality metrics (in the shaded boxes).

Pareto criterion to find QoS trade-offs, but it does not use algebraic trade-off computation and it focuses on resource allocation for multiple tasks sharing a single resource, which does not directly apply to WSN configuration.

### 3. WSN Models

Before introducing the Pareto method, we define basic models of a WSN for two different tasks (applications), and reveal the inherent trade-offs. These models serve as examples that can be extended or adapted as needed.

#### 3.1. Architectural Overview

Consider a network that consists of a collection  $\mathcal{N}$  of identical sensor nodes. The nodes are randomly scattered in an area, and do not move once deployed. The end-user employs the sensor network for a specific task; we look at two different tasks in this example. Firstly, *spatial mapping* (SM), in which all nodes periodically take samples that are sent to the user, for instance to determine the temperature profile over the area. The second task is *target tracking* (TT), in which the objective is to detect and follow target objects. The main difference is that nodes in SM continuously transmit data, while a node in TT only sends a report if it detects a target. We assume that the node locations are known and the network automatically creates a routing tree. The user is in direct contact with the root node of the network.

We use a hierarchical system of requirements and hardware parameters, where the task that runs on the network forms the highest level (the *task level*) and each node is an entity at the lower *node level* (see Figure 1, explained in more detail below, for the TT task). The node level contains a node's hardware settings, such as sample rate and transmission power. From these low-level *parameters* we derive *quality metrics*, like reliability and lifetime. We may use additional metrics to store intermediate results. Quality metrics form the interface between the node and task levels. From the metrics of all nodes together, we derive task-level quality metrics, such as coverage degree and network lifetime. Between the node and task level, we introduce additional *cluster levels*. A cluster is defined as a collection of nodes with the restriction that these nodes form a sub-tree of the network's routing tree. A cluster's quality metrics are the same as

for the task, but they are defined with respect to the cluster’s root instead of the network’s root. Clusters are used for compositional analysis as explained in Section 5.

The user is interested in performance at the task level, and may set QoS constraints here, such as a minimum reporting rate in the SM scenario or a maximum delay between the detection of a target and the arrival of the report at the user for TT. Some quality metrics (e.g. network lifetime), may be maximised. The QoS constraints are considered ‘soft’ (a certain percentage of violations is acceptable), because the unpredictable nature of wireless networks makes it practically impossible to give hard guarantees. The only factors that can be directly set by the system are the hardware settings of individual nodes: the parameters. These are the ‘knobs’ that should be tuned such that the task-level goals are met. To easily compare different solutions, all quality metrics are defined such that larger values are preferred over lower ones.

### 3.2. Node-Level Trade-off Models

The left side of Figure 1 gives an overview of parameters (the rounded rectangles) and how they relate to node-level quality metrics (on the right, in the box). The figure only shows the TT task; as explained below, the parameters of the SM task are the same, while the metrics are slightly different. Rectangles represent constant values, while ovals are metrics that depend on parameters. The quality metrics are grouped in three different dimensions: *reliability*, *time* and *space*. Lines with a filled circle at the end represent positive relationships: if the incoming parameter/metric becomes larger, the other also becomes larger. Likewise, lines with an open circle are negative relationships, while lines with an arrowhead are relations that are not clearly positive or negative. The diagram shows that configuring a node means making trade-offs: adjusting parameters has a positive influence on some quality metrics and a negative influence on others.

The node-level models for the SM and TT tasks are explained below. For every metric we give an example *mapping function* in Table 1, which explicitly defines its relation to the parameters. The relations do not necessarily need to be defined analytically. Other ways to obtain mappings for a set of parameter vectors are, for example, simulations or neural networks. Here, we use explicit equations mainly for speed and ease of use. Our simulations described in Section 7 show that the mapping functions are sufficiently precise to accurately predict quality metrics for both tasks.

The first parameter of a node is its sensor’s *sample rate*  $r_s$ , the number of times per second measurements are taken and processed. The *transmit power*  $P_{tx}$  is the power level at which the node’s radio transmits data. Lastly, a node employs a periodic sleep/wake schedule, with fixed period and *duty cycle*  $f$ . When the node is in sleep mode, it does not take samples (so we should set  $f < r_s$ ) and its micro-controller (MCU) is in low-power mode. We assume the transceiver (including the MAC protocol) does its own power management. Finally, each node has a *parent*, which is the node it sends its reports to according to the routing tree. The metrics are defined below, and summarised in Table 1.

**Reliability.** When a packet is transmitted, another node can receive it with a probability depending on the received SNR and the size of the packet. According to a path-loss model, the received signal power  $P_{rx}$  depends on  $P_{tx}$  and the distance to the receiver  $d$ , according to  $k_r P_{tx} (\frac{d_0}{d})^\alpha$ , for constant gain factor  $k_r$ , path-loss coefficient  $\alpha \geq 2$  and reference distance  $d_0$ . Moreover, transmissions may interfere with transmissions of other nodes. For simplicity, we assume a constant contention-loss probability  $L$ . Then, the *communication reliability*  $I$  (the probability of correctly transferring a data packet) assuming a fixed packet size  $b$  and BPSK signalling, is expressed as (1) in Table 1, with  $Q(x) = \frac{1}{2} \operatorname{erfc}(\frac{x}{\sqrt{2}})$ , and noise level  $N$  [13].

**Time.** The node’s battery has a limited capacity  $E_{batt}$ . The *lifetime*  $T$  of a node, follows from the average power level  $P$  via (2). We define the average power  $P$  as (9), with constant energy to take and process a sample ( $E_s$ ), power of the MCU in active mode ( $P_{mcu}$ ), and energy to transmit/receive a packet ( $E_{tx}$  and  $E_{rx}$ ). We assume the transceiver uses a MAC protocol that minimises idle listening, such as B-MAC [14]. The power level depends on all three parameters, as well as on the additional metric *output traffic rate*  $r_o$ , which is the average rate of packet transmissions ((7) for SM, (8) for TT). This rate includes the node’s own generated traffic, but also traffic that is to be relayed on behalf of other nodes: the *received traffic rate*  $r_i$ . The latter is a special type of metric that depends on the  $r_o$  and  $I$  values of the child-nodes in the routing

Table 1  
Node-level mappings ( $F_n$ ) for a node  $n$

<b>Reliability</b>	
$I(n) = \left(1 - Q\left(\sqrt{2P_{\text{rx}}(P_{\text{tx}}(n))/N}\right)\right)^b (1 - L)$	(1)
<b>Time</b>	
$T(n) = \frac{E_{\text{batt}}}{P(n)}$	(2)
$r(n) = r_s(n)$ [SM]	(3)
$S(n) = \left(\frac{1}{r_s(n)} + D_s\right)^{-1}$ [TT]	(4)
<b>Space</b>	
$C(n) = f(n)$	(5)
<b>Additional metrics</b>	
$r_i(n) = \sum_{i \in \text{ch}(n)} r_o(i)I(i)$	(6)
$r_o(n) = r_i(n) + f(n)r_s(n)$ [SM]	(7)
$r_o(n) = r_i(n) + m \frac{\pi R_s^2}{A} f(n)r_s(n)$ [TT]	(8)
$P(n) = E_s r_s(n) f(n) + P_{\text{mcu}} f(n) + E_{\text{tx}} r_o(n) + k E_{\text{rx}} r_i(n)$	(9)

The set of children of  $n$  is denoted by  $\text{ch}(n)$ .

tree ((6), in which  $r_o(i)I(i)$  represents the traffic out of node  $i$  that is not lost), and hence only indirectly on the parameters, which is why it is visualised differently in Figure 1. Since TT nodes only transmit when a target is in range,  $r_o$  does not only depend on the sample rate and duty cycle, but also on the target's trajectory. The fraction of the time that a target is expected to be near is assumed to be equal to the number of targets  $m$  times the fraction of the total area  $A$  that is covered by the sensor (with sensing range  $R_s$ ).

A quality metric unique to the SM task is the *reporting rate*  $r$ , a measure of the rate at which the spatial map is updated. It is taken equal to the sample rate in (3). Since TT nodes do not continuously report information, the reporting-rate metric is not used. Instead, we are interested in the *detection speed*. We define the detection delay for a sensor node in the active mode, as the time it takes from the arrival of a nearby target until its detection. This delay depends on sample rate  $r_s$ , and (combined in the constant  $D_s$ ) the duration of sampling and detection. The (worst-case) detection speed  $S$ , given by (4), is the inverse of the detection delay.

**Space.** A sensor is said to ‘cover’ the area within its sensing range when it is active. Since a sensor node is typically asleep most of the time, it does not continuously cover this area. We therefore define the metric *coverage degree*  $C$  in (5) as the fraction of the time the sensor is switched on.

### 3.3. Task- and Cluster-Level Trade-off Models

The mapping functions for the task- and cluster-level metrics, as shown in Table 2, are explained below. The functions in the table are for a cluster  $c$ , and use the node-level metrics of Table 1. Task-level mapping functions are obtained by substituting the network  $\mathcal{N}$  for  $c$ .

**Reliability.** In both scenarios, nodes send reports to the user. However, because the communication of reports usually has a limited reliability, not all reports may reach the destination. We want to know how complete the data is that is received by the user: the *information completeness*  $I^c$  is the fraction of all generated reports that arrive. This is approximately equal to the average end-to-end communication reliability

Table 2  
Cluster-level mappings ( $G_{nc}$ ) for a cluster  $c$

---

**Reliability**

$$I^c(c) = \frac{\sum_{i \in c} \prod_{j \text{ on } \mathbf{p}^i} I(j)}{|c|} \quad (10)$$


---

**Time**

$$r^c(c) = \frac{\sum_{i \in c} r(i)}{|c|} \quad [\mathbf{SM}] \quad (11)$$

$$S^c(c) = \min_{i \in c} \left\{ \left( \frac{1}{S(i)} + |\mathbf{p}^i| D_{tx} \right)^{-1} \right\} \quad [\mathbf{TT}] \quad (12)$$

$$T^c(c) = \min_{i \in c} \{T(i)\} \quad (13)$$


---

**Space**

$$C^c(c) = \min_{i \in c} \{C(i)\} \quad (14)$$


---

**Additional metrics**

$$r_o^c(c) = r_o(rt(c)) \quad (15)$$

$$parent^c(c) = parent(rt(c)) \quad (16)$$


---

The root node of cluster  $c$  is denoted by  $rt(c)$ .

over all nodes, given by (10), where  $\mathbf{p}^i$  is the path from node  $i$  to the root node.

**Time.** For an SM task, we are interested in the *reporting rate* of the network/cluster, which is defined by (11) by the average reporting rate over all nodes. A common timeliness metric of a WSN that does target detection is the time it takes from the appearance of a target until the detection report reaches the user. For each node, this delay depends on its detection speed and the hop count  $|\mathbf{p}^i|$  to the root node. The worst-case *detection speed*  $S^c$  is given by (12), where  $D_{tx}$  is the transmission delay (including MAC delay). Further, we use definition (13) for the *lifetime*  $T^c$  that considers all nodes in the network/cluster as essential.

**Space.** The area that is covered by the nodes, if all nodes would be active, is called the covered area. However, a sensor node only covers the area in its range for a fraction of the time. We therefore introduce the metric *coverage degree*  $C^c$ . For a point in the covered area,  $C^c$  is defined as the percentage of the time that it is covered by *at least one* sensor, during a certain period. The coverage degree for the whole covered area is the minimum coverage degree over the whole area. To calculate  $C^c$  for the network/cluster, we need the locations and sensing ranges, plus the coverage degrees of all the nodes. We could approximate the target area by a grid of points and take the minimum coverage degree for each point. For this example, however, we use the form of (14), which is accurate if every sensor covers some area that cannot be covered by any other sensor.

Finally, a cluster's output traffic and parent node are defined as the output traffic and parent node of its root node.

#### 4. The Pareto-Algebra Approach

Our method to configure a WSN is rooted in the algebraic approach to Pareto analysis introduced by Geilen et al. [3]. Pareto analysis is a way to analyse trade-offs in a system in order to find optimal combinations of parameters and quality metrics, called *Pareto points*. We survey some essential concepts from [3].

#### 4.1. Configurations and Minimisation

Consider a general system with parameters and quality metrics. Each of the parameters or metrics can hold values in a specific range or domain that is determined by the characteristics of the hardware and its environment. Such a domain is called a *quantity*, which is a discrete set  $Q$  of values, with a partial order  $\preceq_Q$  (if the quantity is clear from the context, we simply write  $\preceq$ ). If  $q_1, q_2 \in Q$ , then  $q_1 \preceq_Q q_2$  means that the value  $q_1$  is considered better than  $q_2$ . The ordering of a quantity allows to express a preference of certain values over others. For example, for a WSN, we can have  $CovDegree = \{0.3, 0.5, 0.8\}$  with  $0.8 \preceq 0.5 \preceq 0.3$  ( $\preceq$  equals  $\geq$  for greater-is-better).

A configuration space  $\mathcal{S}$  is the Cartesian product  $Q_1 \times \dots \times Q_n$  of a finite number of quantities, and a configuration  $\bar{c} = (c_1, \dots, c_n)$  is an element of such a configuration space. The configuration space holds all possible configurations of a system, given a set of quantities. Since the space can be very large, it is desirable to select only potentially useful configurations for further analysis, instead of analysing all possibilities. Pareto analysis is able to make such a selection, given the preferences expressed in the ordering of quantities.

A dominance relation is used to find configurations that are clearly worse than others and do not have to be considered any further. For  $\bar{c}_1, \bar{c}_2 \in \mathcal{S}$ , configuration  $\bar{c}_1$  is said to *dominate*  $\bar{c}_2$ , denoted by  $\bar{c}_1 \preceq \bar{c}_2$ , if and only if for every quantity  $Q_k$  of  $\mathcal{S}$ ,  $\bar{c}_1(Q_k) \preceq_{Q_k} \bar{c}_2(Q_k)$ . Dominance is a partial order and hence a reflexive relation: every configuration dominates itself. The irreflexive variant, *strict dominance*, is denoted  $\prec$ . Configuration  $\bar{c}_1$  dominates an other configuration  $\bar{c}_2$ , when it is better in at least one quantity and not worse in any of the other quantities. For example, if we have  $\mathcal{S} = CovDegree \times Lifetime$ , with  $Lifetime = \mathbb{N}^+$  and  $\preceq = \geq$ , then  $(0.8, 100) \preceq (0.5, 100)$ , which means that we do not have to consider the second configuration. However,  $(0.8, 100) \not\preceq (0.5, 200)$  and also  $(0.5, 200) \not\preceq (0.8, 100)$ , implying none of the two is clearly better.

**Definition 1 (Pareto Minimal)** *A set  $\mathcal{C}$  of configurations is Pareto minimal iff for any  $\bar{c}_1, \bar{c}_2 \in \mathcal{C}$ ,  $\bar{c}_1 \not\preceq \bar{c}_2$ .*

We denote the Pareto-minimal set of an arbitrary configuration set  $\mathcal{C}$  by  $\min(\mathcal{C})$  and call the process of computing it *minimisation*. For every configuration in  $\mathcal{C}$ , there is an element of  $\min(\mathcal{C})$  that dominates it. The selected configurations are called *Pareto (optimal) configurations* or *Pareto points*. The Pareto-minimal set is unique for finite sets of configurations. Hence, when using a finite configuration set  $\mathcal{C}$ , we only need to consider the subset  $\min(\mathcal{C})$  and we can ignore all the other configurations.

#### 4.2. Building Sets of Configurations

A system often has metrics that depend on other metrics: high-level metrics could be derived from lower-level metrics, while these lower-level metrics themselves may depend on parameters. For a configuration space  $\mathcal{S}$ , we can define a function  $f : \mathcal{S} \rightarrow Q$ , where the new quantity  $Q$  is called a *derived quantity*. In this work, we call  $f$  a *mapping function*. We can extend a configuration set  $\mathcal{C}$  using  $f$ , to create  $\mathcal{C}_f = \{\bar{c} \cdot f(\bar{c}) \mid \bar{c} \in \mathcal{C}\}$ , where the dot ( $\cdot$ ) denotes concatenation. However, for Pareto algebra to be useful, we need to impose an extra restriction on mapping functions. Suppose we have two configurations  $\bar{c}_1, \bar{c}_2 \in \mathcal{C}$ , with  $\bar{c}_1 \preceq \bar{c}_2$  and  $f(\bar{c}_1) \not\preceq f(\bar{c}_2)$ . This would mean that for configurations  $\bar{c} \notin \min(\mathcal{C})$ ,  $\bar{c} \cdot f(\bar{c})$  could be in  $\min(\mathcal{C}_f)$ . This is undesirable, because when minimising before adding the new quantity (the key idea of Pareto algebra is the ability to minimise at intermediate steps of the analysis), potentially optimal configurations may get lost. Therefore, mapping functions that are applied after minimisation should be *monotone*.

**Definition 2 (Monotonicity)** *A function  $f : \mathcal{S} \rightarrow Q$  is monotone iff for any  $\bar{c}_1, \bar{c}_2 \in \mathcal{S}$  such that  $\bar{c}_1 \preceq_{\mathcal{S}} \bar{c}_2$ ,  $f(\bar{c}_1) \preceq_Q f(\bar{c}_2)$ .*

This is the generic definition of monotonicity for partial orders. A function  $h$  on real numbers (where  $\preceq$  equals  $\geq$ ), for example, is monotone if  $x \geq y$  implies  $h(x) \geq h(y)$  for all  $x, y \in \mathbb{R}$  ( $h$  is a non-decreasing function).

A configuration set can be constructed by adding derived quantities, but we can also combine two configuration sets from different spaces. For example, the configuration sets of two sensor nodes may be combined

into one joint configuration set. This is done by the *free product* operation. The free product of configuration sets  $\mathcal{C}_1$  and  $\mathcal{C}_2$  is the Cartesian product  $\mathcal{C}_1 \times \mathcal{C}_2$ . If  $\mathcal{C}_1$  and  $\mathcal{C}_2$  respectively contain  $n$  and  $m$  configurations, then  $\mathcal{C}_1 \times \mathcal{C}_2$  contains  $nm$  configurations. The free product preserves minimality:  $\min(\mathcal{C}_1 \times \mathcal{C}_2) = \min(\mathcal{C}_1) \times \min(\mathcal{C}_2)$ .

#### 4.3. Constraints and the Final Decision

Another important operation of Pareto algebra that is needed to include QoS requirements, is the ability to apply constraints to quantities. A set  $\mathcal{D}$  of configurations from configuration space  $\mathcal{S}$  is called *safe* if for all  $\bar{c}_1, \bar{c}_2 \in \mathcal{S}$ , such that  $\bar{c}_1 \preceq \bar{c}_2$ ,  $\bar{c}_2 \in \mathcal{D}$  implies that  $\bar{c}_1 \in \mathcal{D}$ . A safe set of configurations is also called a *safe constraint*. Applying a safe constraint  $\mathcal{D}$  to a configuration set  $\mathcal{C} \subseteq \mathcal{S}$  yields configuration set  $\mathcal{C} \cap \mathcal{D}$ . Applying a non-safe constraint after a minimisation step may result in the loss of Pareto points. Therefore, if we want to minimise intermediate results, only safe constraints should be used. Also, a safe constraint  $\mathcal{D}$  preserves minimality.

When a configuration set of the whole system has been constructed, one can remove configurations that do not meet the QoS constraints. The result is a set of achievable quality-metric values plus the parameters to attain these values. If this set is empty, the user's requirements cannot be met by any combination of the given sets of parameters. If there are multiple solutions, a choice between them should be made on other grounds, for example by prioritising some metrics.

### 5. Configuring a WSN


Consider a network as a set of nodes denoted by their indices,  $\mathcal{N} = \{0, 1, \dots, |\mathcal{N}| - 1\}$ . As mentioned in Section 3, the network is organised as a tree, where the root of the tree receives reports from all nodes and delivers these to the user. Each node parameter is associated with a quantity that contains all its possible values. For example, transmit power could have a quantity  $TxPower = \{0, 5\}$  (in dBm). The *parameter space* for a sensor node, denoted  $\mathcal{S}_{P,n}$ , is the Cartesian product of all parameter quantities. In our model, we have  $\mathcal{S}_{P,n} = TxPower \times SampleRate \times DutyCycle$ . The parameter space for the whole task  $\mathcal{S}_{P,t}$  is equal to  $(\mathcal{S}_{P,n})^{|\mathcal{N}|}$ .

The end-user that is running a task on the WSN is interested in the task-level behaviour of the network, expressed in a collection of quality metrics. Each metric is also associated with a quantity. The *quality space*  $\mathcal{S}_M$  for the SM task, for example, is  $InfoCompleteness \times ReportingRate \times Lifetime \times CovDegree$ . We can see a WSN as a system that transforms a vector of parameters to a vector of quality metrics. A *configuration* is a collection of parameter values together with the collection of resulting quality-metric values.

**Definition 3 (Mapping)** *A mapping  $F : \mathcal{S}_P \rightarrow \mathcal{S}_M$ , for a parameter space  $\mathcal{S}_P$  and quality space  $\mathcal{S}_M$ , derives a vector of quality metrics from a vector of parameters (like derived quantities, Section 4). More precisely,  $F$  is a tuple of mapping functions  $f_i : \mathcal{S}_P \rightarrow Q_i$ , one for each quality metric  $i$ :  $F = (f_0, f_1, \dots, f_{k-1})$ , with  $k$  the number of quality metrics. Likewise, an incremental mapping  $G : \mathcal{S}_{M,1} \rightarrow \mathcal{S}_{M,2}$  derives a vector of quality metrics from another vector of (possibly different) metrics. The functions  $F$  and  $G$  can be lifted to sets:  $F(\mathcal{C}) = \{F(\bar{c}) \mid \bar{c} \in \mathcal{C}\}$  and similar for  $G$ .*

In the remainder, for a mapping  $F$ , we use subscripts 'n', 'c' and 't' for a mapping to node, cluster (Section 5.1) and task level respectively. For  $G$ , we use a two-letter subscript, to indicate the source and destination level. To find the set of Pareto optimal configurations using a flat system model, we would have to map *every* parameter vector in  $\mathcal{S}_{P,t}$  to a vector of quality metrics by means of a mapping  $F_t$ , and Pareto minimise the resulting set of metrics.  $F_t$  can be defined by combining the functions in Tables 1 and 2. However, the size of  $\mathcal{S}_{P,t}$  increases exponentially with the number of nodes in the network, so such an approach would not be scalable. A solution is to exploit hierarchy in the system, by building and minimising sets of configurations step by step: start with individual nodes, and then incrementally combine nodes into clusters (see Section 3). We expect that minimisation discards a significant number of dominated configurations in each step, such that only the resulting (Pareto) configurations need to be forwarded to the next level. The correctness of this approach – the resulting set of configurations from the one-step and





(a) Cluster-level quality metrics can be derived from the parameters of the nodes in the cluster by a mapping  $F_c$ . Alternatively, they can be derived from node-level quality metrics ( $G_{nc}$ ).

(b) Clusters  $A$  and  $B$  are combined into cluster  $AB$ . The cluster quality metrics of  $AB$  can always be directly derived from the parameters of the nodes in the cluster ( $F_c$ ). If they can also be derived from the cluster-level metrics of  $A$  and  $B$  and this mapping  $G_{cc}$  is monotone, the combining action is also monotone.

Fig. 2. Cluster-level mappings.

---

```

1 create initial one-node clusters
2 repeat until a single cluster remains:
3   monotonically combine  $\geq 2$  clusters
4   derive quality metrics of new cluster
5   minimise new cluster's configuration set

```

---

Algorithm 1. Combining clusters incrementally

clustered approaches should be identical – is investigated in Sections 5.1–5.3. The algorithm’s complexity is discussed in Section 5.4.

### 5.1. The Cluster Method

After mapping the parameter spaces to node-level quality metrics by functions from Table 1, and minimising the resulting sets, we obtain Pareto-optimal subsets of the configuration space  $\mathcal{S}_{M,n}$  for each node in the network. A cluster is a subset of  $\mathcal{N}$ , and the parameters of its containing nodes become the cluster’s parameters. Therefore, an  $n$ -node cluster  $i$  has parameter space  $\mathcal{S}_{P,i} = (\mathcal{S}_{P,n})^n$ . A cluster’s quality metrics are derived from its parameters by means of the mapping  $F_c : (\mathcal{S}_{P,n})^n \rightarrow \mathcal{S}_{M,c}$  (see Figure 2(a)). The cluster-level quality metrics are the same as the task-level quality metrics, only the number of nodes could be different. Thus, if we can compute the quality metrics of a cluster containing all nodes, we immediately obtain the quality metrics for the task. Note that  $F_c$  can be implemented by deriving cluster metrics from node metrics by mapping  $G_{nc}$  of Table 2; see the grey arrow in Figure 2(a). In the analysis below, we use  $F_c$  to keep the explanation clear and general.

Clusters can be combined into a larger cluster by constructing the free product of the parameter spaces of those clusters, and deriving new cluster quality metrics with  $F_c$  (see Figure 2(b)). However, it is usually more efficient to *incrementally* derive new metrics from the cluster metrics of the clusters that are combined instead. We show below that this is possible for our WSN models, and in general for many practically useful cases.

The basic structure of the algorithm is given as Algorithm 1. First, each node is converted into a one-node cluster (using  $F_c$ ). In each loop iteration, multiple clusters are combined into one large cluster. Line 3 combines two or more clusters (selected under conditions defined below) by taking the free product of the configuration sets of these clusters. Then, the set of quality-metric vectors is derived and minimised (lines 4–5). The loop stops when one cluster remains that contains all nodes in the network. The result is a set of Pareto-optimal task-level configurations.

However, it is generally not possible to combine any arbitrary set of clusters. Firstly, all mapping functions to cluster quality metrics need to be defined for the combined cluster. Some mapping functions in our model, for instance, are defined with respect to the routing tree. Therefore, each cluster that uses this function needs to be a tree: a sub-tree of the network’s routing tree.

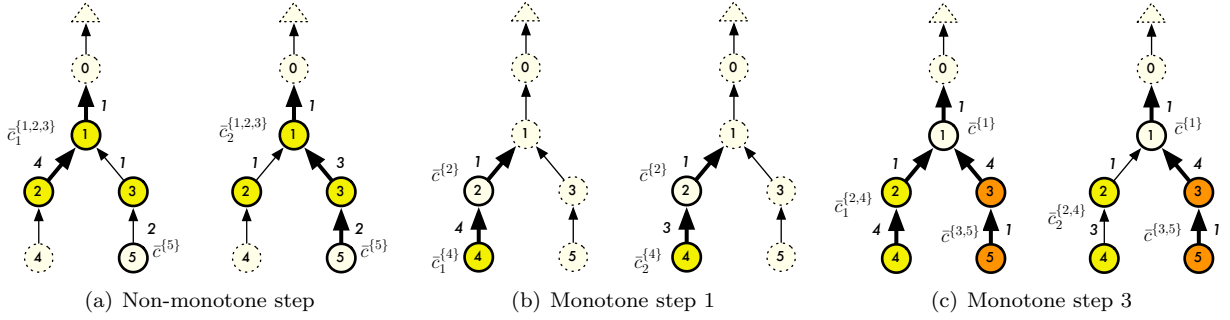


Fig. 3. Examples of non-monotone (a) and monotone clustering steps (b)–(c). A number at the arc coming out of node  $i$  is the delay  $T_i$ .

Secondly, we need to check for *monotonicity*. We need to make sure that the configurations that are removed by minimisation in earlier clustering steps, could never become optimal in later steps. This concept is related to monotonicity as defined in Definition 2.

**Definition 4 (Monotonicity of a clustering step)** Suppose we are combining clusters 0 to  $\ell - 1$  with parameter sets  $\mathcal{C}_{P,i} \subseteq \mathcal{S}_{P,i}$  for all clusters  $0 \leq i < \ell$ . The action of combining these clusters is monotone, iff for all  $\bar{c}_1^i, \bar{c}_2^i \in \mathcal{C}_{P,i}$  and  $0 \leq i < \ell$ ,  $F_c(\bar{c}_1^i) \preceq F_c(\bar{c}_2^i)$  implies  $F_c(\bar{c}_1^0 \dots \bar{c}_1^{\ell-1}) \preceq F_c(\bar{c}_2^0 \dots \bar{c}_2^{\ell-1})$ .

A monotone clustering step preserves the dominance order of cluster configurations. This implies that a cluster configuration  $\bar{c}$  that is dominated before the clustering step, can safely be removed by minimisation, because all configurations of the combined cluster that incorporate  $\bar{c}$  would be dominated by other configurations after clustering. In other words, *if all clustering steps are monotone, none of the eventual task-level Pareto configurations are lost*, and the result from the clustered algorithm would be the same as the result from an all-at-once algorithm. The following lemma is illustrated in Figure 2(b).

**Lemma 5** A clustering step in which clusters 0 to  $\ell - 1$  are combined is monotone, if the function  $F_c$  to calculate the quality metrics of the combined cluster can be rewritten as a monotone function of only the values of the cluster-level quality metrics of clusters that are combined (metrics of individual nodes are not explicitly needed). This means that  $F_c$  can be rewritten as

$$F_c(\bar{c}^0 \dots \bar{c}^{\ell-1}) = G_{cc}(F_c(\bar{c}^0), \dots, F_c(\bar{c}^{\ell-1})) \quad (17)$$

for some monotone  $\ell$ -ary function  $G_{cc} : (\mathcal{S}_{M,c})^\ell \rightarrow \mathcal{S}_{M,c}$ . The function  $G_{cc}$  is a tuple  $(g_0, g_1, \dots, g_{\ell-1})$  of monotone functions that incrementally compute new cluster-level quality metrics ( $k$  is the number of quality metrics).

**Proof** This follows immediately from the monotonicity of  $G_{cc}$ .  $\square$

This result implies that new cluster-level quality metrics can be derived from clusters that are monotonically combined in several ways: by the function  $F_c$  from the parameters of all *nodes* inside the new cluster, or by the function  $G_{cc}$  from the quality metrics of all the *clusters*, or a hybrid form. This is an implementation choice that can be made per mapping function, based on efficiency of computation and storage. Usually, the second option would be the most efficient. Before providing  $G_{cc}$  for the models of Section 3, we first illustrate the method by means of an example.

## 5.2. The Method Applied

Consider a mapping function  $f_d$  to calculate a maximum-delay metric (captured in quantity  $Q_d$ ) for a cluster, given a parameter vector  $\bar{c}$ . Note that the WSN model for target tracking has a speed metric instead of a delay metric. For the sake of the example, however, we use the delay, for which lower values are preferred ( $\preceq_{Q_d}$  equals  $\preceq$ ). We express  $f_d$  in terms of the values  $T_i$ , the time to send a message from node  $i$  to its parent in the routing tree.

$$f_d(\bar{c}) = \max_{p \in L} \left\{ \sum_{i \text{ on } p} T_i \right\}, \quad (18)$$

where  $L$  is the set of paths from all leaf nodes in the cluster to the root node (including the root node itself). Clearly, this function does not depend only on the  $T_i$  values of the nodes in the cluster; also the way the nodes are connected in the routing tree plays an important role. Therefore, when combining clusters with maximum delay as a quality metric, it is necessary that these clusters are connected by links in the network's routing tree, such that the new cluster is a tree. Furthermore, we need to make sure that the clustering step is monotone by ensuring that (17) holds for  $f_d$ . We first give an example of a clustering strategy in which the monotonicity condition fails, and then we show a clustering strategy that is monotone. Throughout the example, we only show the delay metric, but keep in mind that there may be other quality metrics as well. This means that even if a configuration has a worse delay than another, it may or may not be a Pareto configuration, depending on the values of other metrics. We write a configuration  $\bar{c}$  as a vector of only  $T_i$  values.

Suppose we have a cluster  $\{1, 2, 3\}$ , as in Figure 3(a). The cluster has multiple possible configurations (different values of  $T_1$ ,  $T_2$  and  $T_3$ ). The figure shows two configurations:  $\bar{c}_1^{\{1,2,3\}} = (1, 4, 1)$  and  $\bar{c}_2^{\{1,2,3\}} = (1, 1, 3)$ . The cluster delays are  $f_d(\bar{c}_1^{\{1,2,3\}}) = \max\{1 + 4, 1 + 1\} = 5$  and  $f_d(\bar{c}_2^{\{1,2,3\}}) = \max\{1 + 1, 1 + 3\} = 4$  (the thick arrows in the figure show the bottleneck path). Thus,  $f_d(\bar{c}_2^{\{1,2,3\}}) \preceq_{Q_d} f_d(\bar{c}_1^{\{1,2,3\}})$  and after minimisation,  $\bar{c}_1^{\{1,2,3\}}$  may be eliminated (depending on the other metrics). Now we join this cluster with downstream cluster  $\{5\}$ , with one configuration  $\bar{c}^{\{5\}} = (2)$ . The new configurations are  $\bar{c}_1^{\{1,2,3\}} \cdot \bar{c}^{\{5\}} = (1, 4, 1, 2)$  and  $\bar{c}_2^{\{1,2,3\}} \cdot \bar{c}^{\{5\}} = (1, 1, 3, 2)$ . The delays become  $f_d(\bar{c}_1^{\{1,2,3\}} \cdot \bar{c}^{\{5\}}) = \max\{1 + 4, 1 + 1 + 2\} = 5$  and  $f_d(\bar{c}_2^{\{1,2,3\}} \cdot \bar{c}^{\{5\}}) = \max\{1 + 1, 1 + 3 + 2\} = 6$ . This implies that  $f_d(\bar{c}_1^{\{1,2,3\}} \cdot \bar{c}^{\{5\}}) \preceq_{Q_d} f_d(\bar{c}_2^{\{1,2,3\}} \cdot \bar{c}^{\{5\}})$ , but  $\bar{c}_1^{\{1,2,3\}}$  may have been discarded in the previous step! This implies that this clustering step is non-monotone. The reason is that the addition of a *downstream* cluster extends the bottleneck path in one configuration but not in the other one. Observe that this cannot occur when adding upstream clusters; the bottleneck path is then always affected.

Algorithm 1 prescribes that the procedure starts by initialising all nodes as one-node clusters. After this, clusters can be combined step by step, but we need to ensure, firstly, that the new cluster contains a tree (otherwise  $f_d$  is undefined), and secondly, that the clustering step is monotone. Consequently, a situation as in Figure 3(a) should not occur. For clusters with a maximum-delay metric as defined above, we can combine clusters with a special restriction: *the clusters that are combined form a tree with root node  $R$  that (besides  $R$ ) comprises all nodes below  $R$  in the network's routing tree.* We can thus include this condition in line 3 of Algorithm 1.

Figure 3(b) shows a possible first clustering step, after initialisation. Here one-node clusters  $\{2\}$  and  $\{4\}$  are combined. The figure shows two configurations  $\bar{c}_1^{\{4\}} = (4)$  and  $\bar{c}_2^{\{4\}} = (3)$  of cluster  $\{4\}$ , and one configuration  $\bar{c}^{\{2\}} = (1)$  of  $\{2\}$ . The mapping function for the combined cluster  $\{2, 4\}$  is simply  $T_4 + T_2$ , which is 5 and 4 for configurations  $\bar{c}_1^{\{4\}} \cdot \bar{c}^{\{2\}}$  and  $\bar{c}_2^{\{4\}} \cdot \bar{c}^{\{2\}}$  respectively. This step is clearly monotone, and this is always the case when combining clusters that contain only one path: the max-function can be left out and the remaining summation is always monotone.

The second clustering step in the example would be the combination of  $\{3\}$  and  $\{5\}$ , which goes in the same way as step 1. A possible third step is given in Figure 3(c). In this step, we are joining cluster  $\{2, 4\}$  (having configurations  $\bar{c}_1^{\{2,4\}} = (1, 4)$  and  $\bar{c}_2^{\{2,4\}} = (1, 3)$ ) with  $\{3, 5\}$  (configuration (4,1)) and  $\{1\}$  (configuration (1)). Configuration  $\bar{c}_1^{\{2,4\}}$  has a longer delay ( $1 + 4 = 5$ ) than  $\bar{c}_2^{\{2,4\}}$  ( $1 + 3 = 4$ ), so it could have been discarded in step 1. Therefore, it should not be possible that a combination of this configuration with any of the other clusters' configurations becomes a unique Pareto point. The joint cluster's delay  $f_d(\bar{c}^{\{2,4\}} \cdot \bar{c}^{\{3,5\}} \cdot \bar{c}^{\{1\}})$  is

$$\begin{aligned} & \max \{T_4 + T_2 + T_1, T_5 + T_3 + T_1\} = \\ & \max \left\{ f_d(\bar{c}^{\{2,4\}}) + f_d(\bar{c}^{\{1\}}), f_d(\bar{c}^{\{3,5\}}) + f_d(\bar{c}^{\{1\}}) \right\} = \\ & \max \left\{ f_d(\bar{c}^{\{2,4\}}), f_d(\bar{c}^{\{3,5\}}) \right\} + f_d(\bar{c}^{\{1\}}). \end{aligned} \quad (19)$$

---

```

1 function CreateCluster(root):
2   if root is leaf node:
3     set received traffic  $r_i = 0$ 
4     return one-node cluster for root
5   for each child call CreateCluster(child)
6   determine received traffic  $r_i$  for root
7   create one-node cluster for root
8   combine root and child clusters
9   derive quality metrics of new cluster
10  minimise configuration set
11  return configuration set

```

---

Algorithm 2. Cluster combining for WSN model

The last line can be seen as a function  $g(x, y, z) = \max\{x, y\} + z$ , which is monotone, and therefore this clustering step is monotone according to Lemma 5. In the example, using  $\bar{c}_1^{\{2,4\}}$  or  $\bar{c}_2^{\{2,4\}}$  will lead to a combined-cluster delay of  $1 + 4 + 1 = 6$  or  $4 + 1 + 4 = 6$  respectively. Although  $\bar{c}_1^{\{1,2,3,4,5\}}$  is not strictly worse than  $\bar{c}_2^{\{1,2,3,4,5\}}$ , it is still dominated (equal values dominate each other), so  $\bar{c}_1^{\{2,4\}}$  could have been safely removed. Configuration  $\bar{c}_1^{\{1,2,3,4,5\}}$  may be worse in dimensions other than  $Q_d$  to render it strictly dominated, or the two configurations may be identical in terms of quality metrics, in which case only one needs to be kept. But  $\bar{c}_1^{\{1,2,3,4,5\}}$  will never strictly dominate  $\bar{c}_2^{\{1,2,3,4,5\}}$ .

We also see in (19) how the delay mapping function for step 3 can be rewritten from a function that operates on node-level metrics to a function that uses cluster-level metrics. The implementation of the latter is clearly the most efficient, since the computation is easier and fewer values need to be stored.

The final step to complete the network is to combine cluster  $\{1, 2, 3, 4, 5\}$  with cluster  $\{0\}$ , with  $T_0 = 2$ . This step is straightforward: the delay is equal to  $f_d(\bar{c}^{\{1,2,3,4,5\}} \cdot \bar{c}^{\{0\}}) = f_d(\bar{c}^{\{1,2,3,4,5\}}) + f_d(\bar{c}^{\{0\}})$  and this step is thus monotone. The delays are  $6 + 2 = 8$  for both configurations.

**Proposition 6 (Monotonicity of Algorithm 1 for  $f_d$ )** *All clustering steps of Algorithm 1 are monotone for  $f_d$ , if in each clustering step, the clusters that are combined form a tree with root node  $R$  that (besides  $R$ ) comprises all nodes below  $R$  in the network's routing tree.*

**Proof** For each step of the algorithm, we need to ensure that  $f_d$  is defined for the combined cluster, and that the step is monotone. The algorithm maintains two invariants:

- (i) Each cluster is a tree. This ensures that  $f_d$  is defined for each cluster.
- (ii) A cluster either contains exactly one node, or it contains all nodes that are lower in the network's routing tree than the cluster's root node (it is a *leaf cluster*).

Line 1 initialises both invariants, while line 3 plus the selection condition trivially maintains them, whichever cluster is chosen for combination. Invariant (ii) ensures that, when combining clusters, the root of a multiple-node cluster  $M$  is always connected to the root  $R$  of the new cluster via a path containing only one-node clusters. If not, a node belonging to another multiple-node cluster would be on this path, but this implies the existence of a multiple-node cluster that is not a leaf cluster, which contradicts invariant (ii). The maximum delay from any leaf node in cluster  $M$  to  $R$  is equal to the maximum delay of cluster  $M$  as a whole, plus the delays of the extra nodes on the path (including  $R$ ). This is a summation of only cluster quality metrics, which is monotone. Thus, function  $f_d$  applied to a combination of one-node and multiple-node clusters is equivalent to the maximum over the maximum delays for all leaf clusters (either multiple-node or one-node). This is a monotone function using only metrics of the combined clusters, which ensures that the clustering step itself is monotone for the delay metric (Lemma 5).  $\square$

### 5.3. Correctness for the WSN Models

To apply the cluster method to the WSN models of Section 3, we need to slightly reorganise the algorithm. The node model contains the metric received-traffic rate, which depends on the output-traffic rates of nodes

Table 3  
Incremental mappings ( $G_{cc}$ ) for a cluster  $c$

<b>Reliability</b>	
$I_{\Sigma}^c(c) = I_{\Sigma}^c(rt(c)) \cdot \left( 1 + \sum_{i \in ch(c)} I_{\Sigma}^c(i) \right)$	(20)
<b>Time</b>	
$r_{\Sigma}^c(c) = \sum_{i \in sub(c)} r_{\Sigma}^c(i) \quad [\mathbf{SM}]$	(21)
$S^c(c) = \min \left\{ S^c(rt(c)), \min_{i \in ch(c)} \left\{ \left( \frac{1}{S^c(i)} + D_{tx} \right)^{-1} \right\} \right\} \quad [\mathbf{TT}]$	(22)
$T^c(c) = \min_{i \in sub(c)} \{T^c(i)\}$	(23)
<b>Space</b>	
$C^c(c) = \min_{i \in sub(c)} \{C^c(i)\}$	(24)
<b>Additional metrics</b>	
$r_o^c(c) = r_o^c(rt(c))$	(25)
$parent^c(c) = parent^c(rt(c))$	(26)

*Note:* (20), (22), (25) and (26) depend on a tree; the others do not. For combined cluster  $c$ , the root cluster is denoted  $rt(c)$ , the set of child clusters  $ch(c)$ ;  $sub(c) = \{rt(c)\} \cup ch(c)$ .

lower in the tree. Since the algorithm works from leaves to root, and it is most efficient to climb up the tree node by node (see Section 5.4), we can interleave the computation of received rate – and thus the nodes’ initialisation – with the clustering steps, instead of computing all node-level configurations beforehand. Algorithm 2 is the adapted algorithm, written in a recursive form that should be called with the network’s root node as argument. To prove the correctness of the cluster method for the WSN model, we need to provide a monotone mapping  $G_{cc}$  (Lemma 5). Monotonicity should be verified for each mapping function  $g_i$  separately, and a clustering step should only combine clusters that can be monotonically combined (note that the functions in mappings  $F$  need *not* be monotone). Table 3 gives  $G_{cc}$ .  $I_{\Sigma}^c$  and  $r_{\Sigma}^c$  are cumulative metrics that should be divided by  $|c|$  to get the actual quality metrics  $I^c$  and  $r^c$  of Table 2. The incremental computation of these cumulative metrics is more efficient than the computation of the actual metrics.

**Theorem 7 (Monotonicity of Algorithm 2)** *The cluster method (Alg. 2) is monotone for the WSN model.*

**Proof** It can be shown by similar arguments as in Proposition 6 (plus the fact that minimum, addition and multiplication are monotone) that, given the clustering strategy in Algorithm 2, all mapping functions in the model can be written as monotone functions from cluster to cluster metrics as in Table 3. Therefore, by Lemma 5, Algorithm 2 is monotone for the WSN model.  $\square$

#### 5.4. Complexity of the Cluster Method

The operations of Pareto algebra mostly have a polynomial time complexity which is at most quadratic (a basic minimisation algorithm) in the number of configurations  $n$  [15]. A crucial operation is combining two sets of configurations of size  $n$  and  $m$  by a free product, which has complexity  $\mathcal{O}(n \cdot m)$ , and increases the number of configurations from  $n + m$  to  $n \cdot m$ . The free product increases the number of configurations, while minimisation and applying constraints usually reduce this number.

The efficiency of Algs. 1 and 2 mainly depends on the number of clusters  $\ell$  that are combined per step, and the number of configurations  $|\mathcal{C}_i|$  in each cluster  $i$ . Line 3 of Algorithm 1 (line 8 of Algorithm 2) combines configuration sets with a free product. The size of the resulting set  $\mathcal{C}_{\text{prod}}$  is equal to  $|\mathcal{C}_0| \cdot \dots \cdot |\mathcal{C}_{\ell-1}|$ , and the time complexity of the free-product operation is  $\mathcal{O}(|\mathcal{C}_{\text{prod}}|)$ . The complexity of the derivation step in the following line also depends on  $|\mathcal{C}_{\text{prod}}|$  (metrics need to be derived for each new configuration), but on the complexity of the mapping functions as well. Finally, the complexity of the minimisation operation also depends on  $|\mathcal{C}_{\text{prod}}|$ , as said above.

If we consider the number of configurations per node as given and at most  $n$ , mapping functions can be evaluated in constant time, and assuming a quadratic minimisation algorithm is used, the complexity of joining all nodes in one step is  $\mathcal{O}(n^{2|\mathcal{N}|})$ . This is obviously not scalable and therefore not useful for WSN in general. Therefore, it makes sense to join as few clusters as possible in each step and to rely on minimisation to keep the configuration sets small. The algorithm's complexity could even become linear, if each step is able to reduce the set size to roughly  $m$  when cluster configuration sets of size  $m$  are combined. On the other hand, if the minimisation operation does not manage to significantly reduce the size of  $\mathcal{C}_{\text{prod}}$ , the complexity would again be exponential. This number of configurations that can be minimised away greatly depends on the mapping functions and the configurations' values, and it is hard to give general bounds. However, from experiments we see that in practical cases, the algorithm does behave about linearly in the number of nodes (see Section 7).

An extra constraint comes from the fact that clusters can only be combined if this action is monotone. In case we have mapping functions that are defined with respect to a routing tree, like the delay function discussed earlier, and we use Algorithm 2, the number of clusters that we need to combine in one step depends on the number of immediate child nodes of the parent cluster. Therefore, the complexity of the algorithm would also depend on the node degree in this case: the algorithm performs best for a routing tree with a low node degree. Since for WSN a low node degree is generally preferred, because of the better distribution of traffic, this is not a severe limitation. We come back to this in Section 7.

## 6. Implementation

Section 5 gives an overview of the cluster algorithm for configuring a WSN. To efficiently implement this algorithm we suggest optimisations that significantly improve its memory complexity and therefore also its timing. We give some further implementation details as well, and explain that the algorithm can be implemented and executed in a distributed fashion.

### 6.1. Interleaved Combining, Deriving and Minimising

Since a product set can be very large, an implementation that first computes the whole product set and then derives quality metrics and minimises the resulting set would need an excessive amount of memory. Therefore, lines 8–10 of Algorithm 2 are implemented in an interleaved fashion: when an element of the product set is computed, metrics are immediately derived, and the resulting configuration is then integrated in a configuration set that is kept minimal. This is illustrated in the code fragment in Algorithm 3. The incremental minimisation function in Algorithm 4 is derived from the Simple Cull algorithm [15]. The asymptotic time complexity of this approach is the same as before, but the memory demand is now in the order of the size of the minimised configuration sets, instead of the product sets, and thus a lot smaller.

### 6.2. Quantisation

The metric values that are obtained from the mapping functions are of limited accuracy, and we can take this into account while performing minimisation. Given a configuration set with two example configurations  $\bar{a} = (20, 0.1)$  and  $\bar{b} = (2, 0.1000001)$ , we see that  $\bar{a} \not\preceq \bar{b}$  and  $\bar{b} \not\preceq \bar{a}$ , and hence both are Pareto optimal. However, we may consider the difference between the values in the second quantity to be not significant.

---

```

1  $C_{\min} \leftarrow \emptyset$ 
2 for all  $\bar{c}_0 \in \mathcal{C}_0, \bar{c}_1 \in \mathcal{C}_1, \dots, \bar{c}_{\ell-1} \in \mathcal{C}_{\ell-1}$ :
3    $\bar{m} \leftarrow G_{cc}(\bar{c}_0, \bar{c}_1, \dots, \bar{c}_{\ell-1})$ 
4    $C_{\min} \leftarrow \text{AddAndMin}(C_{\min}, \bar{m})$ 

```

---

Algorithm 3. Interleaved combine/derive/minimise step

---

```

1 function AddAndMin( $\mathcal{C}, \bar{c}$ ):
2   for all  $\bar{a} \in \mathcal{C}$ :
3     if  $\bar{a} \preceq \bar{c}$ :
4       return  $\mathcal{C}$ 
5     else if  $\bar{c} \preceq \bar{a}$ :
6        $\mathcal{C} \leftarrow \mathcal{C} \setminus \{\bar{a}\}$ 
7   return  $\mathcal{C} \cup \{\bar{c}\}$ 

```

---

Algorithm 4. Incremental minimisation function

Looking at only the first quantity,  $\bar{a}$  is clearly better. We decide to treat both second-quantity values as being equal to 0.1, such that  $\bar{a} \preceq \bar{b}$ , and  $\bar{b}$  will be removed by minimisation.

We capture the accuracy for each quantity in a vector  $\bar{q}$ , and define a function  $R_{\bar{q}} : \mathcal{S} \rightarrow \mathcal{S}$  that quantises a configuration.  $R_{\bar{q}}$  rounds every value in a configuration to the nearest multiple of the corresponding value in  $\bar{q}$ . This kind of quantisation is trivially monotone (rounding does not change the order of configurations in a quantity). The comparison in line 3 of the minimisation function in Algorithm 4 now becomes  $R_{\bar{q}}(\bar{a}) \preceq R_{\bar{q}}(\bar{c})$ , and likewise in line 5. Note that we only quantise in comparisons, but we use the unquantised configurations for further processing to prevent the propagation of rounding errors.

### 6.3. Indexing

Recall that a cluster configuration is a vector of parameter values (for all nodes in the cluster) plus a vector of quality metrics. A straightforward implementation of the algorithm would store each configuration exactly in this format. However, as illustrated in Figure 2(b), the metrics of a cluster can be directly computed from the metrics of lower-level clusters; the parameters are not needed. Our next optimisation measure uses this opportunity to save on precious storage space and time-consuming memory accesses.

First of all, in the case that multiple nodes have the same parameter vector sets, we need to store this set only once and in the configurations we can use indices to parameter vectors. Still, a cluster configuration would contain a parameter index for each node contained in the cluster, and therefore the size of configurations grows when the cluster algorithm progresses and clusters become larger. As a result, the memory demand of the algorithm for large networks may be prohibitive. We therefore extend the use of indexing in the cluster algorithm, as shown in Figure 4. In each step of the algorithm, a (one-node) root cluster is combined with its child clusters. For each of these clusters we need a configuration set with metric vectors. Each metric vector in a configuration set has an index, and when combining metric vectors into new configurations, their indices are stored with them. After deriving quality metrics and minimising the new configuration set, this set is split into a set that holds only the metric vectors and a linked set that contains vectors of indices, called the *indexing table*. Subsequently, only the table of metric vectors is used in the next cluster step.

Thus, besides the parameter set, for every node in the network only an indexing table is stored. Memory for intermediate cluster configuration sets can be freed immediately after usage in a cluster step. When the final set of Pareto points (for the whole network) has been computed, one metric vector will be selected, and we need to reconstruct the parameter vector that gives rise to it by tracing back through the indexing tables. Algorithm 5 shows how this is done by a recursive function that is called with the ID of the network's root and the index of the selected configuration.

### 6.4. Computation of Received Traffic

As usual in sensor networks, the nodes in our model need to forward data that they receive from other nodes. This amount of data, called the received traffic  $r_i$  in Section 3, needs to be known in order to compute node-level quality metrics. This is what happens in line 6 of Algorithm 2. As stated in Table 1,  $r_i$  is the sum of the output-traffic values ( $r_o$ ) of a node's children. A complication is that we typically have a whole

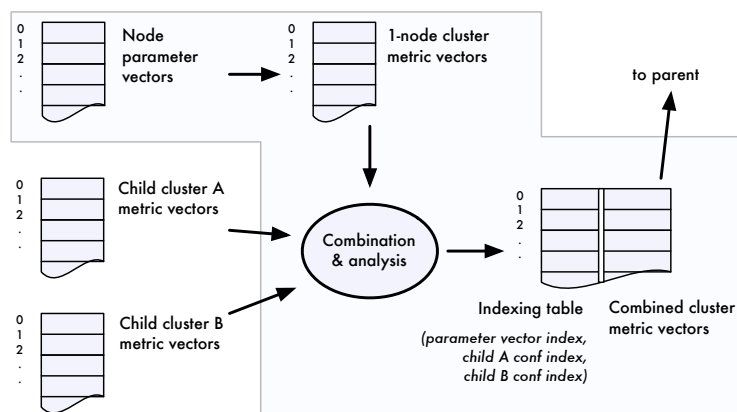


Fig. 4. Implementation of one algorithm step (node degree 3). The figure shows the tables that are needed for configuration sets and indexing, which constitute almost all of the memory needed. The numbers shown to the left of the tables are the indices.

---

```

1 function reconstructParams(n, i):
2   k ← Index[n][i]
3   if n is a leaf node:
4     return P[n][k[0]]
5    $\bar{p} \leftarrow P[n][k[0]]$ 
6   for all child nodes c of n, enumerated by j:
7      $\bar{p} \leftarrow \bar{p} \cdot \text{reconstructParams}(c, k[j+1])$ 
8   return  $\bar{p}$ 

```

---

Algorithm 5. Reconstructing a parameter vector.  $\text{Index}[n][i]$  selects the  $i^{\text{th}}$  index in the indexing table of node  $n$ , while  $P[n][i]$  is the  $i^{\text{th}}$  parameter vector in its parameter set. The first column in an indexing table is always the parameter index for the node; the remaining columns are indices of configurations of child clusters.

set of  $r_o$  values for each child: one for each Pareto point. Only after completion of the cluster algorithm and selection of the final configuration, we know the final  $r_o$  values for all nodes.

Deciding on the  $r_i$  value to use, we can choose to be conservative. There exists a negative relation between  $r_i$  and node lifetime, so if we use the highest  $r_o$  values for each child, we would end up with the worst-case node lifetime. Another option is to use average values, such that there is a higher chance to be close to the eventual value, or any other choice that is found to be useful. This choice does not affect the monotonicity of the cluster algorithm: for any  $r_i$ , lifetime values may be different, but the dominance order of configurations is the same. However, whichever option we choose, it remains an estimation of  $r_i$ , and therefore we should recompute the metrics of the final task-level configuration set (all the way from the parameters of each node) to obtain the real lifetime values. Since there are typically just a few Pareto configurations, this can be done quickly.

If we are recomputing the metrics of the Pareto set after running the cluster algorithm, it does not really matter how we compute  $r_i$  at intermediate steps. However, if we do something with the lifetime values during steps of the algorithm, the choice becomes important. We could for example apply constraints, as explained in Section 8.2 below, in which case we should use the conservative approach.

### 6.5. Distributed Execution

The cluster algorithm as described in this paper is given as a centralised algorithm that is run separately from the WSN, before starting the WSN's task. However, Algorithm 2 treats nodes in a leaf-to-root fashion. A node only depends on information from downstream nodes to compute configurations for the whole cluster with itself as root. It is therefore also possible to execute the algorithm in a distributed way, in which each



Table 4  
Model constants for TelosB nodes, TT task

Reliability		Time		Space	
$k_r$	1000	$D_{tx}$	120 ms	$R_s$	10 m
$\alpha$	3	$D_s$	35 ms	$A$	$100 \cdot  \mathcal{N}  \text{ m}^2$
$d_0$	1 m	$E_s$	0.15 mJ	$m$	$ \mathcal{N} /100$
$N$	-18 dBm	$E_{batt}$	12 Wh		
$b$	288 bit	$P_{mcu}$	5.4 mW		
$L$	0.015	$E_{rx}$	8.2 mJ		

node passes the optimal configurations on to its parent, after computing the Pareto set for its cluster. When the network’s root is reached and a final task-level configuration has been chosen, this configuration is communicated back down the tree. The use of indexing as described above is important to enable running the algorithm on resource-constrained nodes, since it minimises the communication overhead as well as the memory usage. The distributed execution will be further developed in future work.


## 7. Experimental Results

We implemented Algorithm 2 with the optimisations discussed in the previous section and ran it for networks of different sizes, for both the SM and TT models (on a PC with a 1.86GHz Core 2 Duo processor). To gain insight in the scalability of the approach, we recorded the number of configurations considered in each step when the clusters are combined ( $|\mathcal{C}_{prod}|$ ), as well as the run time. Constraints were not used in these experiments; the goal was to find *all* Pareto points. See Section 8.2 for extra experiments involving constraints. The results in this paper are much better than our previous results [4] (about two orders of magnitude), because we included the optimisations described in Section 6, and we are now using an implementation in C++ instead of Python (a compiled vs. an interpreted language).


For each network size, we randomly distributed sensor nodes in a square area. To ensure an even distribution across the area, we placed the nodes with a certain variance around fixed grid points. While scaling the number of nodes, the area was scaled accordingly, such that the node density was equal for all networks. For each network, the transmission range was set to 20m, and a shortest-path spanning tree (SPST) was created. To ensure a fair comparison between the results for all networks, we only used SPSTs in which each node has at most 4 immediate child nodes (algorithm complexity depends on node degree, see Sections 5.4 and 7.5). We first set the quantities for the node-level parameters as follows:  $TxPower = \{0, -5, -10\}$  (dBm),  $SampleRate = \{0.5, 0.3, 0.1\}$  (Hz),  $DutyCycle = \{0.2, 0.4, 0.6\}$ . This leads to  $3^3 = 27$  possible configurations per node. The constants in the nodes’ mapping functions were chosen to match Crossbow TelosB [16] sensor nodes (power usage, transceiver parameters); the constants  $L$  and  $D_{tx}$  were estimated by simulation (it is reasonable to assume that certain constants can be determined empirically at WSN deployment time). See Table 4 for an overview. Subsequently, we did the same tests for just 8 configurations per node, by omitting the last parameter value in each quantity. To achieve some robustness in the measured run times and configuration counts, for each network size and number of configurations, we analysed 10 different networks.

### 7.1. Run Time and Number of Configurations


It turned out that the maximum number of configurations simultaneously considered (over all steps) stays limited, even when the network size increases. For the tests with 8 configurations per node, this maximum was roughly  $60 \cdot 10^3$ , for both SM and TT (see Figure 5). For the case of 27 configurations per node,  $|\mathcal{C}_{prod}|$  increased to about  $2.1 \cdot 10^6$ . Moreover, judging from Figure 5, the average run time of the algorithm increases roughly proportionally with the number of nodes in all scenarios, which is good considering that




(a) Run time, 8 configurations/node



(b) Run time, 27 configurations/node



(c) Max. set size, 8 configurations/node



(d) Max. set size, 27 configurations/node

Fig. 5. Run time and number of configurations. For increasing network size, the run time (a and b) does not grow more than linearly for both the SM (circle markers) and TT (cross markers) scenarios and two different sizes of the parameter space. The reason for this linear growth is that the maximum number of simultaneously considered configurations at any clustering step stays limited (c and d).

the underlying configuration space grows exponentially. Therefore, we may conclude that the algorithm is very well scalable and thus suitable for the configuration of a WSN.

For example, for an instance of a 900-node network, the TT scenario, and 27 configurations per node, the algorithm took 45.9 seconds to complete. The total number of possible configurations is  $27^{900}$ , but there were never more than 1,258,712 configurations to be considered at the same time, during any of the algorithm's iterations. There are 9 Pareto-optimal configurations, as given in Table 5. Each of these solutions has a corresponding set of parameter values for each node. We see that there are clear trade-offs between most quality metrics.

## 7.2. Memory Usage

Further, we recorded the average memory usage (over all steps) of the cluster algorithm, in the same experiments as above. Figure 6 shows the results for the old and new implementations of the algorithm (before and after the optimisations given in Sections 6.1–6.3). It is clear that the average memory usage of the optimised implementation is nearly independent of the network size, while the old implementation needs more memory for larger networks.


## 7.3. Model Accuracy

The results of Table 5 were tested in a network simulator based on OMNeT++ [17]. The nodes in the simulator were modelled after the TelosB as well and use B-MAC [14] to communicate. Energy usage, data


Table 5

Analysis results for 900-node example network. Differences with simulation in parentheses.

Information Completeness	Detection Speed	Lifetime	Coverage Degree
$I^c$ (%)	$S^c$ ( $\frac{1}{s} \cdot 10^3$ )	$T^c$ (h)	$C^c$
84 (-3)	41 (-0.4)	3481 (+8)	0.2
63 (-2)	41 (-0.4)	3773 (+120)	0.2
2.0 (+0.5)	41 (-0.4)	4073 (+388)	0.2
78 (-1)	41 (-0.3)	1821 (+3)	0.4
59 (+1)	41 (-0.4)	1970 (+54)	0.4
2.0 (+0.3)	41 (-0.4)	2123 (+195)	0.4
78 (+5)	41 (-0.2)	1214 (-76)	0.6
59 (+6)	41 (-0.1)	1313 (-40)	0.6
2.0 (+0.4)	41 (-0.4)	1415 (+55)	0.6



(a) TT task (post-optimisation max: 51800)



(b) SM task (post-optimisation max: 43600)

Fig. 6. Comparison of the average memory usage before (dashed lines) and after (solid lines) the optimisations described in Section 6. Measured is the average number of values that need to be stored over all steps of the cluster algorithm. The average memory usage of the optimised algorithm is nearly independent of the network size. The results are for networks with 27 configurations per node.

loss and delay were determined by simulating each configuration, after which the quality-metrics information completeness, detection speed and lifetime were computed (the mapping function for coverage degree is accurate by definition, and therefore not tested). The difference between computed and simulated values is given in brackets in the table (e.g. ‘-10’ indicates that the computed value is 10 lower than the simulated value). The deviations are relatively small in general (within 10%). They are caused by inaccuracies in the model and not in the way Pareto points are computed, and can be further reduced by fine-tuning the model. The same tests were done for 18 other random networks of different sizes, for both SM and TT (5 random configurations per network). On average, the deviations were 1.5% (percentage point), 0.9% and 5.4%, for completeness, speed and lifetime respectively. For information-completeness, the deviation is reported in percentage points, because the completeness percentages may become very small (see Table 5). In such a case, an acceptable small absolute deviation in percentage points would result in a misleadingly high relative deviation.




Fig. 7. The effect of node degrees on the run time of the cluster algorithm. The x-axis specifies the highest node degree in the network. The results are for 100-node networks with 27 configurations per node.

#### 7.4. Comparison with a Genetic Algorithm

We also explored the configuration space of the example network of Table 5 via a genetic algorithm [11]. Even after running for three days, it returned 6 configurations (among others (0.8, 12, 3869, 0.2)), which were all strictly dominated by at least one configuration in Table 5. It turns out that configurations with the best metric values are so rare and isolated in the total space of size  $27^{900}$ , that the genetic algorithm is doomed to fail: the probability of finding the Pareto points goes to zero. The best metric values found were 8.4, 12, 3869 and 0.2 (order as in the table), which is 90%, 70%, 5% and 67% lower than the best values found by our method. This result confirms the expected result that a search space of  $27^{900}$  is too large to search efficiently and accurately via a randomised approach, and it emphasises the strength of our exact algebraic approach.

#### 7.5. The Impact of the Node Degree

We mentioned in Section 5.4 that the complexity of the cluster algorithm heavily depends on the degree of the nodes in the routing tree. This is because in one iteration of the algorithm, a node (one-node cluster) needs to be combined with all of its child clusters at once. To illustrate this, we constructed various 100-node networks with different maximum node degrees, and fed them to the cluster algorithm as before. The resulting run times are displayed in Figure 6 (averages over 10 networks of the same maximum degree).

The figure shows that the impact of the node degree is significant. This implies that if we have the freedom to choose a routing tree, we could select one that minimises the maximum node degree, in order to reduce the complexity of the QoS optimiser. We will look at this in future work.

## 8. Heterogeneity and Constraints

### 8.1. The Cluster Method for Heterogeneous Networks

We have given two realistic WSN tasks and have shown that the cluster method is able to find the optimal solutions very quickly, and scales up to very large networks. However, the method is not restricted to these two models. Other models can be used, provided the monotonicity requirements are met. Moreover, it is possible to analyse *heterogeneous* networks with the same method. We could for example introduce different types of nodes, or even different tasks running on the same network.

As an example, we study a network running both SM and TT at the same time. We have three node types: nodes that do either SM or TT, and nodes that do both. The network can contain any distribution and any number of these node types. One practical scenario to use this could be a disaster scene in which

Table 6  
Metrics for combined SM/TT clusters.

Metric	Task
Information Completeness	SM
Information Completeness	TT
Reporting Rate	SM
Detection Speed	TT
Lifetime	SM/TT
Coverage Degree	SM
Coverage Degree	TT
Output Traffic (additional metric)	SM/TT

we constantly need to observe the temperature across the region to be aware of fire, while at the same time we want to track people walking around. We may specify that the tracking information needs to be more fine-grained than the temperature mapping task, and hence we could deploy a large number of TT nodes plus a smaller number of SM+TT nodes (assuming that using combined nodes is more cost effective than using separate SM nodes).

We let both tasks share the same routing tree, which means that TT nodes also relay traffic for the SM task and vice versa. We assume that SM and TT need different sensors, so the combined SM+TT nodes need to have two sensors, and hence two sample-rate parameters. It is therefore needed to create a new node-level model (mapping functions) for the combined SM+TT nodes. The quality metrics in this model are the union of the metrics of the SM and TT models. Further, a new cluster model is needed. Note that for the cluster method to work, all clusters need to have the same metrics. And since a cluster can now have both the SM and TT tasks, we need the quality metrics for both tasks incorporated in the configuration space; see Table 6 for an overview of the metrics. Note that some metrics are shared by both tasks, while others are not. For example, we are interested in the individual coverage of the tasks, but we define a shared lifetime metric because both tasks share the same network.

The mapping functions can be easily derived from the functions in Table 3. However, they do depend on the type of the root node of the cluster. For instance, if the root is an SM node, it only forwards TT traffic and it does not add a new detection-delay term. Therefore, (22) will just be

$$S^c(c) = \min_{i \in ch(c)} \left\{ \left( \frac{1}{S^c(i)} + D_{tx} \right)^{-1} \right\}.$$

There are similar considerations for the other metrics. Most importantly, all mapping functions are still monotone, so the cluster algorithm will return all Pareto-optimal configurations.

We tested the model for different combinations of the three node types in a similar set-up as in Section 7, with averages over 10 random networks of 900 nodes per case. The results are summarised in Table 7. Note that the run-times vary quite a bit: in some cases the run-times are shorter than in the homogeneous case, sometimes longer. The differences arise from a varying number of Pareto points, as seen in the last two columns of Table 7. These numbers are quite unpredictable, since they depend on many factors. The number of Pareto points may change a lot even when small changes to model constants are made. Most importantly, also in the heterogeneous networks we analysed, the run-times do not explode, but remain very short given the complexity of the problem. The last two rows in the table are for cases with only one task, as done before, and listed for comparison. Note that the run times are slightly higher than in Figure 5, due to a larger overhead (we are still using cluster configurations for the SM+TT model).

Table 7  
Experimental results for heterogeneous networks.

SM (%)	TT (%)	SM+TT (%)	Run-time (s)	Pareto points	
				1	2
20	20	60	14	3	8
33	33	33	14	3	8
40	40	20	18	3	8
45	45	10	29	3	9
50	50	0	138	17	14
0	90	10	18	4	8
0	0	100	20	3	8
100	0	0	42	10	11
0	100	0	49	12	11

<sup>1</sup>number of Pareto points for the whole network

<sup>2</sup>average number of Pareto points over all clusters

## 8.2. Applying Constraints On The Fly

The approach until now was to first find all Pareto points, and subsequently check them against the given QoS constraints. However, it is possible to apply constraints earlier, at each step of the cluster algorithm, in order to reduce the configuration sets further, and thus improve the analysis time. Pareto algebra allows this, as long as the constraints are *safe* (see Section 4), which is the case for all constraints that pose a lower-bound on the quality metrics we use.

If we place a completeness constraint of 50% and a coverage constraint of 0.4 on the 900-node example network with Pareto points in Table 5 (which originally took 45.9 seconds to compute), we are left with just four complying configurations for the complete network, and the run time reduces to just 0.5 seconds. The maximum number of simultaneously considered configurations in any of the clustering steps reduces from 1,258,712 to 1,080. Obviously, if the constraints are stricter, more configurations can be removed at each step, and the run time may improve dramatically.

## 9. Conclusion

We have introduced a method for determining the node configurations of a WSN such that high-level QoS constraints are met. The method is based on Pareto algebra, which is used to reason about QoS trade-offs. It features an algorithm that keeps the working set of possible configurations small, by analysing parts of the network one by one, and meanwhile discarding configurations that are Pareto dominated by others. We gave accurate models for WSNs doing target tracking and spatial mapping, in which the available trade-offs are made explicit. We showed how to analyse the Pareto optimal configurations of such WSNs with the proposed method, how to efficiently implement the algorithm, and presented test results for various network sizes. These results show that the method is scalable – both in terms of run time and memory usage – and therefore practically usable for WSNs, even with large numbers of nodes. Furthermore, the nature of the algorithm makes it very well suitable for a distributed implementation, which scales even better and is suitable for run-time application. This is an area that we will further explore in future work. We also showed that a traditional genetic algorithm approach does not scale. The WSN models and configuration scheme can easily be extended for more complicated applications and networks, including heterogeneous networks with various types of nodes. Future work will focus on expanding the method with middleware features, such as adaptation to run-time changes. Another interesting aspect for future work is complexity control. The algorithm’s complexity is dominated by the size of the working set of configurations. This set can be kept small by limiting the number of Pareto points at intermediate steps, e.g. via approximation techniques or

simply via selection of the most interesting configurations from the quality perspective. Furthermore, we will investigate how to construct routing trees with reduced node degrees, such that the complexity of the cluster algorithm is reduced, without sacrificing performance in terms of the QoS metrics. We could even trade-off some application quality for optimisation speed, if we maintain compliance with the QoS constraints.

## References

- [1] D. Chen, P. K. Varshney, QoS support in wireless sensor networks: A survey, in: *Int. Conference on Wireless Networks (ICWN 2004)*, CSREA Press, 2004.
- [2] Y. Yu, B. Krishnamachari, V. Prasanna, Issues in designing middleware for wireless sensor networks, *IEEE Network* 18 (1) (2004) 15–21.
- [3] M. Geilen, T. Basten, B. Theelen, R. Otten, An algebra of Pareto points, *Fundamenta Informaticae* 78 (1) (2007) 35–74.
- [4] R. Hoes, T. Basten, C.-K. Tham, M. Geilen, H. Corporaal, Analysing QoS trade-offs in wireless sensor networks, in: *10<sup>th</sup> ACM Symposium on Modeling, Analysis, and Simulation of Wireless and Mobile Systems (MSWiM)*, Proc., ACM Press, New York, NY, USA, 2007, pp. 60–69.
- [5] T. He, J. Stankovic, C. Lu, T. Abdelzaher, SPEED: A stateless protocol for real-time communication in sensor networks, in: *ICDCS 2003*, Proc., IEEE, 2003.
- [6] K. Akkaya, M. Younis, An energy-aware QoS routing protocol for wireless sensor networks, in: *ICDCSW 2003*, Proc., IEEE, 2003, pp. 710–715.
- [7] M. Perillo, W. B. Heinzelman, Providing application QoS through intelligent sensor management, in: *Int. Workshop on Sensor Network Protocols and Applications (SNPA '03)*, IEEE, 2003.
- [8] K. Römer, O. Kasten, F. Mattern, Middleware challenges for wireless sensor networks, *SIGMOBILE Mob. Comput. Commun. Rev.* 6 (4) (2002) 59–61.
- [9] L. Thiele, S. Chakraborty, M. Gries, S. Künzli, A framework for evaluating design tradeoffs in packet processing architectures, in: *39th Design Automation Conference (DAC 2002)*, ACM Press, New Orleans LA, USA, 2002, pp. 880–885.
- [10] G. Palermo, C. Silvano, V. Zaccaria, Multi-objective design space exploration of embedded systems, *Journal of Embedded Computing* 1 (3).
- [11] E. Zitzler, L. Thiele, Multiobjective evolutionary algorithms: a comparative case study and the strength pareto approach, *IEEE Transactions on Evolutionary Computation* 3 (4) (1999) 257–271.
- [12] C. Lee, J. Lehoczky, R. Rajkumar, D. Siewiorek, On quality of service optimization with discrete QoS options, in: *Real-Time Technology and Applications Symposium*, Proc., IEEE, 1998.
- [13] S. Haykin, *A Introduction to Analog and Digital Communications*, John Wiley & Sons, 1989.
- [14] J. Polastre, J. Hill, D. Culler, Versatile low power media access for wireless sensor networks, in: *Embedded networked sensor systems (SenSys '04)*, Proc., ACM Press, New York, NY, USA, 2004, pp. 95–107.
- [15] M. Geilen, T. Basten, A calculator for Pareto points, in: *Design, Automation and Test in Europe (DATE)*, Proc., IEEE Computer Society Press, Los Alamitos, CA, USA, 2007, pp. 285–291.
- [16] Crossbow technology website, <http://www.xbow.com>.  
URL <http://www.xbow.com>
- [17] OMNeT++ website, <http://www.omnetpp.org>.  
URL <http://www.omnetpp.org>