

Title	A constraint programming approach to the additional relay placement problem in Wireless Sensor Networks
Authors	Sitanayah, Lanny;Brown, Kenneth N.;O'Sullivan, Barry;Sreenan, Cormac J.;Quesada, Luis
Publication date	2015-04-02
Original Citation	Quesada, L., Sitanayah, L., Brown, Kenneth N., O'Sullivan, B., Sreenan, C. J. (2015) 'A Constraint Programming Approach to the Additional Relay Placement Problem in Wireless Sensor Networks', Constraints, 20(4), pp. 433-451. doi: 10.1007/s10601-015-9188-8
Type of publication	Article (peer-reviewed)
Link to publisher's version	https://link.springer.com/article/10.1007/s10601-015-9188-8 - 10.1007/s10601-015-9188-8
Rights	© Springer Science+Business Media New York 2015. This is a post-peer-review, pre-copyedit version of an article published in Constraints. The final authenticated version is available online at: http://dx.doi.org/10.1007/s10601-015-9188-8
Download date	2025-09-07 05:34:42
Item downloaded from	https://hdl.handle.net/10468/7412



UCC

University College Cork, Ireland
Coláiste na hOllscoile Corcaigh

A Constraint Programming Approach to the Additional Relay Placement Problem in Wireless Sensor Networks

Luis Quesada · Lanny Sitanayah ·
Kenneth N. Brown · Barry O'Sullivan ·
Cormac J. Sreenan

Abstract A Wireless Sensor Network (WSN) is composed of many sensor nodes which transmit their data wirelessly over a multi-hop network to data sinks. Since WSNs are subject to node failures, the network topology should be robust, so that when a failure does occur, data delivery can continue from all surviving nodes. A WSN is k -robust if an alternate length-constrained route to a sink is available for each surviving node after the failure of up to $k-1$ nodes. A WSN is strongly k -robust if there are k disjoint length-constrained routes to a sink for each node. Determining whether a network is k -robust is polynomial. However, determining whether a network is strongly k -robust is an NP-complete problem. We develop a Constraint Programming (CP) approach for deciding strongly k -robustness that outperforms a Mixed-Integer Programming (MIP) model on larger problems. A network can be made (strongly)

Luis Quesada
CTVR and INSIGHT Centre for Data Analytics, Department of Computer Science,
University College Cork, Ireland
E-mail: l.quesada@cs.ucc.ie

Lanny Sitanayah
CTVR and Mobile & Internet Systems Laboratory, Department of Computer Science,
University College Cork, Ireland
E-mail: l.sitanayah@cs.ucc.ie

Kenneth N. Brown
CTVR and INSIGHT Centre for Data Analytics, Department of Computer Science,
University College Cork, Ireland
E-mail: k.brown@cs.ucc.ie

Barry O'Sullivan
INSIGHT Centre for Data Analytics, Department of Computer Science,
University College Cork, Ireland
E-mail: b.osullivan@cs.ucc.ie

Cormac J. Sreenan
CTVR and Mobile & Internet Systems Laboratory, Department of Computer Science,
University College Cork, Ireland
E-mail: c.sreenan@cs.ucc.ie

robust by deploying extra relay nodes. We extend our CP approach to an optimisation approach by using QuickXplain to search for a minimal set of relays, and compare it to a state-of-the-art local search approach.

Keywords Constraint Optimisation Problem · Wireless Sensor Networks · Network Deployment Planning · Relay Placement · Node Disjoint Paths

1 Introduction

Rapid improvements in wireless communication and electronics have led to the development of Wireless Sensor Networks (WSNs) for monitoring in many diverse applications, including environmental assessment, fire detection, personal health management, and surveillance. A sensor node is a device with integrated sensing, processing and communication capabilities, and is typically battery powered. A WSN is composed of many nodes, which transmit their data wirelessly over a multi-hop network to data sinks. These networks are prone to failures: the wireless devices are often unreliable, they have limited battery life, transmissions may be blocked by changes in the environment, and the devices may be damaged, e.g. by weather, wildlife or human intervention. For dealing with failures, reliable routing protocols Boukerche et al (2006); Chipara et al (2006); Gnawali et al (2009); Zeng et al (2011) have been proposed, but they rely on a network topology in which alternative routes to a sink are available.

Therefore, one key objective in the planning of a WSN is to ensure some measure of robustness in the topology, so that when failures do occur routing protocols can continue to offer reliable delivery. One standard criterion is to ensure routes to the sink are available for all remaining sensor nodes after the failure of up to $k-1$ other nodes. In addition, since there are sometimes data latency requirements, there may be a limit on the path length from sensor node to sink. Therefore every node in the initial design should have k node-disjoint paths to the sink of length less than the length bound.

The importance of node-disjoint paths in WSNs is motivated by properties that are important to sensor network applications, such as energy efficiency, resiliency and security. To be more energy-efficient, some protocols are required to provide multipath routing capability, so that less route updating is required. In addition, by having multiple disjoint paths, a network can balance its load by distributing the traffic over more sensor nodes to reduce the probability of network congestion, improve energy efficiency and achieve better energy management.

The second motivation for node-disjoint paths is to provide path resilience and data transmission reliability. Path resiliency is necessary for protecting networks against node and link failures. By transmitting multiple copies of an original data packet over different paths, a network can ensure packet recovery from several path failures. This can also give us time to go out and repair the network before a data stream is lost. Data resiliency by sending duplicate messages down each discovered path supports reliable data transmission and

improves the Quality of Service (QoS), such as network throughput (increased bandwidth), end-to-end latency and data delivery ratio.

From the network security perspective, having multiple disjoint paths can ensure that no single node sees all data and no single node can insert data into every message. In some applications, eavesdropping nodes can easily capture a message when it is routed through a single path. If this message is divided into several packets and sent through multiple disjoint paths, the task of capturing the whole message would be much more difficult.

To ensure that sensor nodes have sufficient paths, it may be necessary to add relay nodes, which do not sense, but only forward data from other nodes. The possible positions of the relay nodes may be limited, and each position can communicate with only a subset of the other nodes in the network. Finally, installing additional relays comes at a cost thus motivating solutions that minimise the number of additional relays.

In this paper, we present constraint-based solutions to these robust WSN topology design problems. First, we consider the decision problem of whether it is possible, for a given network and a set of candidate relay positions, to find k length-bounded node-disjoint paths to the sink for a single sensor. We classify the problem as being NP-complete, and remark that the relaxed version of the problem where biconnectivity (instead of disjointness) is enforced is polynomial. We present Constraint Programming (CP) and Mixed-Integer Programming (MIP) approaches, and compare to a previously published local search method Sitanayah et al (2011). We show that the CP model solves the problems in less time than it takes to generate the MIP models. We then consider the extended decision problem, in which we must find qualifying paths for all sensors in the same network. We demonstrate that the CP solution is able to solve these problems in reasonable time, but that the MIP model does not scale. Finally, we consider the optimisation problem, in which the aim is to minimise the number of relays. The MIP model is fastest on small problems but again does not scale up to larger problems. We use QuickXplain Junker (2004) to develop an approximate CP solution, which we show is competitive in time with the local search method on the larger problems, although with lower quality solutions. We also relax the robustness requirement by enforcing biconnectivity instead of disjointness. When doing so, we observe that our QuickXplain based approach provides solutions that are not only competitive in time but also in quality.

2 Background and Related Work

A WSN can be modelled as a graph $G=(V, E)$, where V is a set of nodes and E is a set of edges. Each edge connects two nodes that are within transmission range of each other¹, and the two nodes are said to be *adjacent*. A *path* of length t between two nodes v and w is a sequence of nodes $v=v_0, v_1, \dots, v_t=w$, such

¹ For simplicity we assume bi-directional links, but this could be easily relaxed by specifying a more complex connectivity graph.

that v_i and v_{i+1} are adjacent for each $0 \leq i < t$. Two nodes are *connected* if there is a path between them. A graph is connected if every pair of nodes is connected.

The problem of deploying relay nodes for increased reliability has long been acknowledged as a significant problem Bredin et al (2005); Pu et al (2008); Han et al (2010); Misra et al (2008). Greedy Randomised Adaptive Search Procedure for Additional Relay Placement (GRASP-ARP) Sitanayah et al (2011), a recently published local search approach, has been shown to deploy fewer relay nodes with faster runtimes than the closest known approach Bredin et al (2005); Pu et al (2008). It uses the GRASP stochastic local search method Feo and Resende (1989, 1995); Resende and Ribeiro (2002) to deploy additional relay nodes for ensuring (k, l) -sink-connectivity, where all sensor nodes have k node-disjoint paths of length $\leq l$ to the sinks. GRASP-ARP works by (i) generating an initial feasible solution by adding relays until sufficient paths are found, (ii) exploring the neighbourhood of the initial solution by adding and removing relays to reduce the number of required relays, and (iii) iterating until a stopping criterion is satisfied Sitanayah et al (2011). Within each stage, GRASP-ARP uses a Counting-Paths algorithm based on the Ford Fulkerson algorithm with node splitting to find k node-disjoint paths. When searching for a new augmenting path, Counting-Paths uses a best-first approach, but with a higher priority for nodes already used in a previously obtained path, with the intention of balancing the lengths of the resultant path set. Counting-Paths is guaranteed to find k node-disjoint paths if they exist.

The time complexity of the Ford Fulkerson algorithm is $O(|E|f)$, where f is the maximum flow in the graph. Therefore, the time complexity of Counting-Paths is $O(|E|k)$, where $O(|E|)$ is the time complexity of breadth first search and k is the number of disjoint paths.

3 Complexity of the Additional Relay Placement Problem

If we restrict the Additional Relay Placement decision problem to ensuring robustness for just one of the sensors in the network, it is equivalent to the Bounded Vertex Undirected Disjoint Paths problem (BVUP), i.e., the problem of finding k vertex-disjoint paths in an undirected graph for a given pair of nodes where the paths have bound length, which is known to be NP-Complete Golovach and Thilikos (2011), although fixed parameter tractable. Additionally, the problem of minimising the number of additional nodes needed to disjointly connect two nodes, given a bound on the length of the paths, is a specific case of the Generalised Constrained Shortest Link-Disjoint Paths Selection, which is also known to be NP-Hard Xiao et al (2006).

The decision problem becomes easier though if we are requested to enforce bounded k -connectivity, i.e., is it possible to reach the target from the source respecting a bound on the length after removing $k - 1$ vertices? Even though Menger showed that biconnectivity and disjointness are equivalent properties when there is no bound on the path, it has been already observed that the




Fig. 1 s and t are biconnected subject to a length constraint of 5, but there are not two disjoint paths connecting s to t subject to the same length constraint.

Algorithm 1 CHECKBICONNECTIVITY(G, s, t, λ)

```

1.  $(V, E) = \text{DIJKSTRA}(G, s, t)$ 
2. if  $|E| > \lambda$ :
3.   return false
4. else:
5.   for  $v \in V - \{s, t\}$ :
6.      $G' = \text{REMOVEVERTEX}(G, v)$ 
7.      $(V', E') = \text{DIJKSTRA}(G', s, t)$ 
8.     if  $|E'| > \lambda$ :
9.       return false
10. return true
    
```

theorems does not hold when the lengths of the paths are bounded Lovász et al (1978).

Figure 1 presents an example showing the distinction. The length bound is 5 and k is 2. There are only two vertex disjoint paths from s to t : (s, a, b, t) and (s, c, d, e, f, g, t) . The second path violates the constraint on the bound. However, there is no cut-set of size 1. Notice that if either a or b is removed it is still possible to reach the target within 5 hops.

Indeed BVUP remains NP-Hard even for $k = 2$, but Bounded Bicconnectivity can be solved polynomially. Notice that to check the latter we just need to call Dijkstra $m + 1$ times, where m is the number of nodes involved in the shortest path from the sensor to the target, as shown in Algorithm 1. The complexity of Dijkstra is $O(V \lg V + E)$ Cormen et al (2001). As the shortest path from s to t would involve $O(V)$ vertices, the complexity of Algorithm 1 is $O(V \times (V \lg V + E))$.

4 The Bounded Vertex Undirected Disjoint Paths Problem

The main focus of this paper is the decision problem: given a sensor and a sink in a network, do there exist k node-disjoint paths of length no greater than l to the sink? We saw above that this problem is the same problem as the Bounded Vertex Undirected Disjoint Paths problem, which is NP-complete. At first sight, GRASP-ARP offers a solution, since its initialisation phase will continue adding relays until qualifying paths are found, relying on Counting-




Fig. 2 A network for which Counting-Paths fails to find two disjoint paths between nodes 1 and 10 of length at most 4.

Paths to check the paths. However, Counting-Paths is a heuristic algorithm, and does not guarantee to meet the length bound.

Consider the problem shown in Figure 2, where the problem is to find two node-disjoint paths of length less than 5 from node 1 to node 10. Counting-Paths uses a breadth-first search to find the shortest path, with nodes ordered by smallest ID. During the first iteration, it finds path A = (1,3,9,10) of length 3. The second search then finds path B = (1,2,5,6,8,10), and the algorithm terminates with the two path lengths of 3 and 5, and thus fails to satisfy the length bound. However, there are two node-disjoint paths of length 4: (1,3,6,8,10) and (1,4,7,9,10). To obtain this solution, the second search would have had to return path C = (1,4,7,9,3,6,8,10), which would produce the correct output when A and C are merged and overlapping segments removed. Path C in this case is not returned, since the breadth-first search reaches node 6 earlier in path B and marks it as has been visited. Similar counter-examples exist for different ordering heuristics.

4.1 CP approaches

Model. As shown in Figure 3, we model the problem of finding k disjoint paths between a pair of nodes by cloning the source node k times. Finding k disjoint paths is equivalent to finding a tree whose leaf nodes are the clones of the source node, and the root is the target node. Table 1 shows the constants and variables used in the CP model. Table 2 shows a CP flow based model for the problem of finding k length-bounded disjoint paths for a sensor s .

As the source node is cloned k times, we use the z variables where z_i denotes the successor of the i -th clone. We post an ALLDIFFERENT constraint on the z_i variables to ensure that we end up with k different paths (Constraint (1)). Constraints (2) and (3) connect the successor variables to the flow variables.



Fig. 3 Cloning the source.

Table 1 Constants and variables used in the CP models

Constants	Variables
<ul style="list-style-type: none"> – V is the set of nodes in the network. – S is the set of sensors. – $R = V - S$ is the set of relays. – k is the lower bound on the connectivity. – s is the sensor (the source). – t is the target. – λ is the upper bound on the length of the paths. 	<ul style="list-style-type: none"> – $y_v \in V$ is an integer variable referring to the successor of node v. – $z_i \in V - \{s\}$ is an integer variable referring to the successor of the i-th clone of s. – $f_{uv} \in \{0, 1\}$ is a Boolean variable referring to the flow on edge $\langle u, v \rangle$. – $l_v \in \{0, \dots, \lambda\}$ is an integer variable referring to the length of the path from node s to node v.

Table 2 CP I: Flow based decomposition of the tree constraint

$$\begin{aligned}
 & \text{ALLDIFFERENT}([z_1, \dots, z_k]) & (1) \\
 & z_i = v \Rightarrow f_{sv} = 1 & \forall 1 \leq i \leq k, \forall v \in V & (2) \\
 & y_u = v \Leftrightarrow f_{uv} = 1 & \forall u \in V - \{s, t\}, v \in V \text{ if } u \neq v & (3) \\
 & \sum_{v \in \text{in}(u)} f_{vu} = \sum_{v \in \text{out}(u)} f_{uv} & \forall u \in V - \{s, t\} & (4) \\
 & f_{uv} = 1 \Rightarrow l_v \geq l_u + 1 & \forall \langle u, v \rangle \in E & (5)
 \end{aligned}$$

We remark that Constraint (2) is an implication and not a double implication since we do not want z_i to be set to v when z_j is set to v (assuming $i \neq j$). That is, having a double implication would be inconsistent with the fact that the z variables need to be set to different values. Notice that Constraint (3) is only posted in those cases where $u \neq v$ thus allowing the successor of an unused node to be bound to itself. In Constraint (4) we use $\text{in}(v)$ and $\text{out}(v)$ to denote the incoming and outgoing nodes of v . Constraint (4) ensures that the number of used incoming edges of a node is equal to the number of used outgoing edges. In order to ensure that no node is used twice we constrain the incoming and outgoing degrees of each internal node to be at most one. This is implicitly enforced by the y_v variables since a successor variable can only be bound to one value, in conjunction with Constraint (4). Constraint (5) connects the length variables to the flow variables². A solution of the decision problem at hand can be expressed in terms of the y and z variables. The

² We use \geq instead of $=$ because the paths from the source to the target may have different lengths.

determination of these variables determines the other variables so these are the decision variables in this model.

In what follows, we will call the model of Table 2 CP I. Alternatively, we can model our decision problem by using one single Tree constraint Beldiceanu et al (2008) since both the degree of the nodes and the length of the path in the tree can be constrained through the interface of the constraint. We use CP II to refer to the model obtained using the Tree constraint.

Labelling strategy. During the backtracking search process we construct the paths in a systematic way: we pick one of the clones of the source and decide its successor, then the successor of its successor and we continue this way until we reach the target. Then we pick another clone and do the same until the paths of all clones are found. When all paths have been found, the remaining successor variables are set to their self values. We remark that in constraint programming labelling is interleaved with propagation. During propagation some successor variables may get determined thus avoiding the consideration of those variables during labelling.

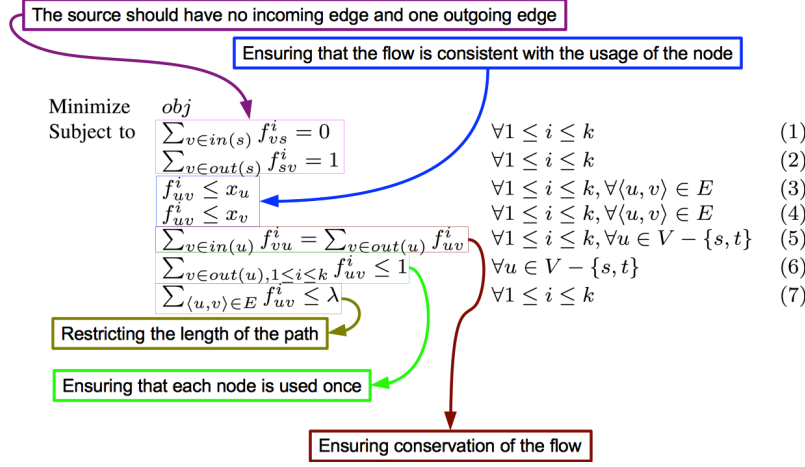
In order to implement this dynamic variable ordering, we use the lower bound of the l variables: we pick the y_i variable whose l_i has the highest lower bound. If the highest lower bound is 0, we know that the node is not participating in any path and therefore the value that we pick for y_i is i . Otherwise, the value that we pick for y_i is the one associated with the closest distance to the target aiming at minimising the length of the path³.

4.2 MIP approaches


In this section we present our MIP approaches to BVUP. We remark that in MIP we are forced to express all constraints in terms of linear equations, which leads us to models that are more verbose with respect to the CP models. The set of constants and variables of the model is contained in the set of constants and variables of CP. The only difference is that, in MIP, we replicate the set of flow variables per path. More concretely, instead of f_{uv} , we have f_{uv}^i . We also use variables x_v , which are Boolean variables referring to the usage of the nodes.

The model is presented in Table 3. Constraints (1) and (2) ensure that flow emanating from the source is 1, for each value of k (i.e., k paths start from the source). Constraints (3) and (4) disable the flow on a given arc if their nodes are not used. Constraint (5) ensures the conservation of the flow. We enforce disjointness by constraining the outgoing flow for every internal node (for every value of k) to be at most 1 (Constraint (6)). Because of Constraint (5), Constraint (6) indirectly enforces that the incoming flow for every internal node is at most one (for every value of k). As we are keeping a separate graph per path, the length constraint is enforced by imposing that the used edges in

³ The distance matrix is not updated during search so it might happen that the chosen successor is not the closest to the target.

Table 3 MIP IV: Handling the upper bound on the length of the path via replication of the graph


each graph is less than λ (see Equation (7)). As BVUP is a decision problem, the objective (obj) is 1 (i.e., we are just interested in a solution that satisfy the constraints).


Fig. 4 Splitting relay nodes.

Here we are describing the MIP approach that provided us with the best results (MIP IV). However, three other approaches were tried: MIP I, MIP II and MIP III. MIP I is actually an approach to a relaxed version of BVUP where we ignore the length constraint. In MIP II we encode the length constraint by associating an integer variable with each node representing its distance to the source and constraining those variable through conditional constraints. In MIP III we encoded disjointness by splitting the nodes (as shown in Fig 4). That is, instead of constraining the incoming degree and the outgoing degree of the node to be 1, we split each node and set the capacity of the new edge to 1 to enforce that the node only participates in one path (i.e., used once).

5 Extending to multi-sensor BVUP

The problem of finding disjoint paths for a selection of sensors to a single target (multi-sensor BVUP) can be decomposed into a set of BVUPs (one per sensor). We recall that we are still dealing with a decision problem. As such we are not concerned about the number of sensors that we use in total, therefore each BVUP can be solved independently. Both the CP and the MIP models can be extended by applying them individually to each sensor.

As mentioned in Section 4 we focus on dealing with (multi-sensor) BVUP since this problem alone is NP-complete. The approach to the optimisation problem that we discuss in the next section (Section 6) can be seen as solving a succession of BVUP problems, and therefore relies heavily on how well we tackle this decision problem. This is why in the results discussed later we take a set of 7474 instances encountered while solving the optimisation problem and elaborate on the effectiveness of our BVUP approach when solving those instances.

Results The instances used in the empirical evaluation are connectivity graphs. To generate them, firstly we generate WSN topologies similar to the technique used in Sitanayah et al (2011). The two-dimensional network area is divided into grid cells, where one sensor node is placed inside one unit grid square of $8\text{ m} \times 8\text{ m}$ and the coordinates are randomly perturbed. This is an approximation of manual deployment of sensor nodes, such as in a building or a city that has regular symmetry. In this simulation, we want the original topologies, i.e. topologies without relays, to be as sparse as possible, because sufficiently dense networks do not need additional relays to guarantee the existence of disjoint paths. In order to get sparse networks (average degree 2-3), we generate more grid points than the number of sensor nodes. For example, we use 6×6 and 11×11 grid squares to randomly deploy 25 and 100 nodes, respectively. Candidate relays are also distributed in a grid area, where a candidate occupies a unit grid square of $6\text{ m} \times 6\text{ m}$. For n25 and n100 topologies, we use 49 and 196 candidate relays, respectively. Both sensor and relay nodes use the same transmission range, i.e. 10 metres. The location of the sink was fixed at the top-left corner of the network. The maximum path length is set to 10 for n25 and 20 for n100 networks.

When it comes to the platforms used for the experiments of this section and the next one, the MIP experiments were obtained using CPLEX 12.3 and the version of Choco used in the CP experiments is 2.1.5. The experiments were run on Linux 2.6.25 x64 on a Dual Quad Core Xeon CPU machine with overall 11.76 GB of RAM and processor speed of 2.66 GHz.

The CP approach solves all the n100 instances with average solution time of 511 sec. The MIP models do not scale so well, with the time taken to generate the instances being longer than the solution time for the CP model.

Figures 5-8 present the performance of our BVUP approach measured over a set of 7474 instances generated during an execution of our QuickXplan based approach to the Additional Relay Placement problem (explained in the next




Fig. 5 Evaluation of multi-sensor BVUP - Time distribution.




Fig. 6 Evaluation of multi-sensor BVUP - Time vs Failures.

section), when solving the optimisation problem associated with an instance of 100 sensor nodes and 196 candidate relays. As it can be observed in Figure 5, close to 99% of the instances are solved with 10 or fewer failures. In Figure 6 we are comparing the time distribution with the failure distribution. As expected, there is almost a direct correlation between time and failures (ignoring the few




Fig. 7 Evaluation of multi-sensor BVUP - Time vs Cardinality.




Fig. 8 Evaluation of multi-sensor BVUP - Failures vs Cardinality.

cases where we time out). Figure 7 relates the time to the cardinality of the set of candidate relays. When focussing on the hard cases, we can observe that there is a high variance with respect to the cardinality thus suggesting that the difficulty of the instances is not correlated to the cardinality of the set. The relation between failures and cardinality (shown in 8) suggests that the

speed (i.e., number of failures per unit of time) reduces when we have more relays, which is not surprising since the time spent in propagation increases with the number of relays.

6 Dealing with the Minimisation of Relays

Modelling the minimisation of relays in MIP is straightforward: we just set the objective to the sum of used relays ($\sum_{v \in V - (S \cup \{t\})} x_v$) and replicate the constraints per sensor. In CP we could do something similar: replicate the constraints per sensor, set the objective as done in MIP and use Branch and Bound. However, neither of the two approaches work in practice due to the size complexity of the models. For this reason, we explore another alternative in CP: to compute an approximation to an optimal set of relays by mapping a solution of the multi-sensor BVUP problem to conflicts and computing minimal conflicts using QuickXplain.

QuickXplain is an algorithm designed to compute preferred explanations and relaxations for over-constrained problems Junker (2004). In its basic form the algorithm receives a set of inconsistent constraints S and returns a subset C , which corresponds to a conflict of S . C is a minimal set in the sense that C is still inconsistent but the removal of any of its constraints will make it consistent. QuickXplain outperforms related approaches for computing minimal conflicts by applying a divide-and-conquer strategy.

We approximate the minimisation of the number of used relays by mapping solutions of the optimisation problem to conflicts. We say that *a conflict is a set of relays that allows us to find disjoint paths for all the sensors*.

Under this interpretation, we have as many constraints in S as candidate relays. That is, for each relay r we have a constraint in S stating that r is *available* to be used in a path from a sensor to the target. Relays that are not constrained are *unavailable*. Notice that the more constraints we add to the set the more likely it is to have a conflict. To put it differently, if a subset C' of S is consistent any subset of C' is consistent too thus showing that our model holds monotonicity, which is required to use a minimal conflict approach. Following this interpretation, we have that a minimal conflict corresponds to a minimal set of relays needed to ensure disjointness.

We remark that the solution that we get from our QuickXplain based method is an approximation to the objective of finding a minimum set of relays since the solution is a minimal set and not a minimum set. That is, it could be very well the case that there is another minimal set of relays of smaller cardinality. Moreover, as we are using a timeout for the satisfiability checks, minimality is also approximated since a set of relays might be discarded (i.e., flagged as a non conflict) due to the timeout.

Results Tables 4 and 5 present the results that we have obtained with the different approaches on the n25 and n100 instances introduced in Section 5. For each approach we report the best cost found (the number of chosen relays)

and the time (in seconds) spent in finding it. The results of the local search approach (LS), already published in Sitanayah et al (2011), were reproduced using GCC 4.4.4.

Our first observation is that the problem becomes very easy when the length constraint is disregarded. Indeed, BVUP without the length constraint is a mere flow problem. We suspect that the problem is still polynomial when minimising the number of relays (see column MIP I in Table 4) but this still remains to be proved. On the n25 instances Sitanayah et al (2011), MIPs II and III failed almost completely because of the encoding of the length constraint, timing out after 30 minutes. MIP IV solved almost all instances in less than 10 seconds, faster than the local search, but on a few instances was significantly slower. CP I outperformed CP II.

As mentioned before, in CP II we use the Tree constraint Beldiceanu et al (2008). The Tree constraint (as presented in Beldiceanu et al (2008)) is in the process of being migrated to the latest version of Choco⁴. The experiments of the CP II approach were carried out using a preliminary version kindly provided by Jean-Guillaume Fages. In this version, the constraints on the degree of the nodes are handled independently using occurrence constraints. There is no pruning taking advantage of the structure of the graph to filter the degree of each node. This is important in our problem since we can early detect failures by discovering articulation points.

We noticed that one issue with our CP approach is the high number of BVUP instances that need to be solved during the execution of QuickXplain. In general, we spend very little time but the number of instances is high. In CP I and CP II we are recreating and reading the models associated with the sensors every time. CP III is an optimisation of CP I where we keep the models in memory to avoid their recreation. In CP IV we enhance CP III by ordering the relays for QuickXplain by increasing distance from the sink, and reducing the timeout for solving each BVUP instance.

In order to appreciate the impact of ordering the relays in CP IV, in Figure ?? we compare this heuristic with its corresponding anti-heuristic on the n25 instances. That is, in the anti-heuristic we order the relays by decreasing distance from the sink. As can be observed in Figure ??, the choice of heuristic makes a significant difference to the performance of the model. The closest-first heuristic reduces the runtime over the anti-heuristic by more than 80% (199 sec to 1072 sec), and yet generates solutions with fewer relays (7.95 to 9.25).

For the n100 instances (Table 5), the MIP model did not scale up, failing to find a feasible solution in 1 hour. The CP approach is competitive in time with the local search method, although with lower quality solutions. It is important to remark that most of the time is spent re-reading models. In Table 5 we show both the time spent by Choco solving the models and the total time. We attempted to keep the solvers in memory to avoid the re-reading of the

⁴ <http://www.emn.fr/z-info/choco-solver/>

Table 4 Results with the different approaches for 25 sensor nodes and 49 candidate relays

instance	LS		MIP I		MIP II		MIP III		MIP IV		CP I		CP II	
	cost	time	cost	time	cost	time	cost	time	cost	time	cost	time	cost	time
scen-n25-01	6	29.48	4	0.63	4	407	4	608	4	1.56	4	433	6	673
scen-n25-02	7	10.74	4	0.47	-	1800	5	1800	5	3.07	5	106	6	301
scen-n25-03	8	20.29	4	0.58	-	1800	-	1800	5	14.68	8	505	9	651
scen-n25-04	6	18.76	5	0.64	5	49	5	577	5	1.78	9	145	11	666
scen-n25-05	9	11.73	3	0.77	-	1800	-	1800	4	5.82	7	558	8	841
scen-n25-06	10	14.56	5	1.06	-	1800	24	1800	8	359.90	9	298	9	428
scen-n25-07	7	47.25	6	1.02	-	1800	-	1800	7	9.15	7	732	7	320
scen-n25-08	11	19.87	5	0.70	-	1800	30	1800	7	117.36	9	466	9	779
scen-n25-09	9	23.72	6	1.36	-	1800	-	1800	7	33.44	7	229	7	235
scen-n25-10	7	17.02	5	0.53	-	1800	-	1800	6	9.46	6	107	6	199
scen-n25-11	9	12.86	6	1.10	-	1800	-	1800	7	19.04	10	240	11	575
scen-n25-12	3	19.57	3	0.52	-	1800	3	1255	3	3.07	8	432	5	474
scen-n25-13	10	20.81	6	0.94	-	1800	-	1800	8	205.16	10	215	10	634
scen-n25-14	7	18.52	6	0.96	-	1800	-	1800	6	3.08	10	562	11	1172
scen-n25-15	7	1.83	5	0.48	-	1800	6	1640	6	2.02	8	154	6	158
scen-n25-16	6	40.24	5	0.51	-	1800	-	1800	5	2.12	7	173	7	174
scen-n25-17	9	14.79	5	0.92	-	1800	-	1800	6	23.83	9	96	9	373
scen-n25-18	9	10.74	7	0.84	-	1800	-	1800	7	6.76	7	295	7	259
scen-n25-19	8	16.84	5	1.07	-	1800	-	1800	6	4.68	7	537	7	636
scen-n25-20	10	30.08	5	0.65	-	1800	-	1800	7	150.98	10	592	12	1532

Table 5 Results with the different approaches for 100 sensor nodes and 196 candidate relays

instance	LS		CP III			CP IV		
	cost	time	cost	solve time	total time	cost	solve time	total time
scen-n100-01	9	22108.70	19	4038.98	6723.31	17	1411.00	3845.96
scen-n100-02	11	13144.90	20	4541.52	6787.96	18	1244.68	3486.35
scen-n100-03	5	997.63	19	3357.86	6014.40	15	1447.18	4214.37
scen-n100-04	5	174.95	15	3273.79	6268.68	13	1515.62	4421.60
scen-n100-05	12	15699.70	20	2896.29	5876.92	20	1532.72	4451.92
scen-n100-06	7	4905.40	20	5042.72	8241.95	19	1473.99	4122.45
scen-n100-07	6	1974.81	21	4255.08	7169.11	16	1653.99	4725.40
scen-n100-08	11	14462.00	17	3092.81	5523.46	19	1697.22	4816.48
scen-n100-09	8	1751.50	14	2534.82	5013.07	13	1383.42	4042.70
scen-n100-10	8	17362.70	20	4133.57	7122.75	18	1278.03	3872.07
scen-n100-11	6	12237.10	14	3762.80	6394.02	12	1376.92	4029.24
scen-n100-12	9	2043.61	17	3305.84	6377.01	18	1582.52	4574.32
scen-n100-13	9	1623.28	15	3964.06	6763.79	19	1563.62	4523.04
scen-n100-14	6	14761.90	15	4080.36	7045.69	16	1521.96	4283.17
scen-n100-15	7	15098.70	11	2101.22	4129.03	10	883.64	2909.79
scen-n100-16	5	14958.10	11	1930.22	3884.59	11	871.12	2921.89
scen-n100-17	8	2291.31	21	3290.02	5959.84	16	1524.06	4466.98
scen-n100-18	8	33180.30	21	3828.45	6463.51	14	1164.11	3527.96
scen-n100-19	7	15476.90	16	3129.26	6001.84	14	1365.71	3972.77
scen-n100-20	10	16516.90	14	2985.34	5776.93	15	1287.88	3877.77

models using the notion of worlds in Choco, but run out of memory for the big instances.




Fig. 9 Comparing two heuristics for selecting the next relay during the execution of Quick-Xplain – Comparing time. In *closest* we select the relay that is closest to the target while in *furthest* we do the opposite

7 Enforcing Biconnectivity instead of Disjointness

As shown in Section 3, checking bounded disjointness is NP-complete. However, checking bounded biconnectivity is polynomial. In this section we relax the Additional Relay Placement Problem by enforcing biconnectivity instead of disjointness. In the LS approach, Counting-Paths is replaced with Algorithm 1. In the CP approach, as the decision problem is polynomial, Algorithm 1 is used to check consistency. Even though Constraint Programming is not used for checking the decision problem, we still use the concept of minimal conflict to address the minimisation of relays therefore our approach is still CP based.

We summarise the results and compare them to those of obtained when disjointness is enforced in tables 6 and 7. As we can observe, the gap in quality between the two approaches is reduced for bounded biconnectivity. Indeed we see that the quality of the cost obtained with the CP approach is comparable to those obtained by the LS approach. We also observe that the time spent by the CP approach is less.

We noticed that in average the consistency checks for biconnectivity used to take more time than the consistency checks for disjointness. For instance, for the case of scen-n100-1 there were 93 consistency checks which spent in total 3845.96 seconds when disjointness was enforced. When biconnectivity was enforced, there were only 55 checks with a total time of 10620.27 seconds. This was somehow surprising since the first problem is NP-complete while the second one was polynomial. However, when we look carefully at the instances




Fig. 10 Comparing two heuristics for selecting the next relay during the execution of QuickXplain – Comparing quality of the solution. In *closest* we select the relay that is closest to the target while in *furthest* we do the opposite

of bounded disjointness we see that most of them are easy instances which can be solved almost without failing (see Figure 8). Indeed, when there is no failure the CP approach is much faster than Algorithm 1. As the price of constructing the models is paid only once, the price of finding a solution with few failures is pretty much linear with respect size of the graph.

The situation of the LS approach when comparing disjointness enforcement and biconnectivity enforcement is similar. Counting-Paths spent much less time than Algorithm 1. However it is important to remark that Counting-Paths is still an heuristic when it comes to checking biconnectivity. Indeed, the same example presented in Figure could be used to show this fact. Notice that the graph is biconnected since it was possible to find two disjoint paths and disjointness is a stronger property.

8 Conclusion and Future Work

Designing wireless sensor network topologies that are robust to failures is an important task for a wide range of different monitoring applications, and is solved in the literature by local search approaches that add additional relay nodes to achieve robustness. We have identified that the underlying decision problem – do there exist k node-disjoint paths of length at most l from a sensor to a sink – is known to be NP-complete. We remarked that the relaxed version of the problem where biconnectivity (instead of disjointness) is enforced is polynomial. We have developed complete constraint programming

Table 6 Results with the different approaches for 25 sensor nodes and 49 candidate relays

instance	LS				CP			
	Biconnected		Disjoint		Biconnected		Disjoint	
	cost	time	cost	time	cost	time	cost	time
scen-n25-01	5	54.12	6	29.48	4	17.32	4	35.16
scen-n25-02	7	17.91	7	10.74	9	21.27	9	37.77
scen-n25-03	11	36.83	8	20.29	9	12.71	10	33.29
scen-n25-04	5	24.65	6	18.76	7	23.47	7	40.96
scen-n25-05	8	37.64	9	11.73	7	25.75	8	82.09
scen-n25-06	7	36.50	10	14.56	8	14.62	8	23.22
scen-n25-07	7	62.76	7	47.25	9	21.45	8	75.47
scen-n25-08	8	13.87	11	19.87	8	18.70	8	37.41
scen-n25-09	8	35.22	9	23.72	8	20.00	8	95.84
scen-n25-10	7	29.33	7	17.02	7	17.76	8	64.53
scen-n25-11	9	8.63	9	12.86	8	22.40	8	31.98
scen-n25-12	5	39.85	3	19.57	5	12.62	6	47.82
scen-n25-13	9	45.40	10	20.81	10	24.84	11	49.33
scen-n25-14	7	9.40	7	18.52	8	21.03	9	35.94
scen-n25-15	6	18.61	7	1.83	8	23.68	8	33.46
scen-n25-16	6	50.62	6	40.24	6	18.25	7	42.35
scen-n25-17	8	34.66	9	14.79	9	16.70	9	26.90
scen-n25-18	10	53.35	9	10.74	9	16.82	9	65.47
scen-n25-19	7	36.95	8	16.84	8	21.30	9	78.71
scen-n25-20	9	12.99	10	30.08	7	18.18	9	71.37

and mixed integer programming models for solving the problem. For larger problems, the CP model, implemented in Choco, outperforms the MIP approach, implemented in CPLEX, returning the solutions while the MIP model is still generating the problems. This provides an important tool for network operators, enabling them to evaluate proposed deployments in real time. We then presented the first approach to extending the CP solution to the minimisation problem, where the aim is to minimise the number of additional relays to ensure robustness. We implemented an approach based on QuickXplain, which searches for a minimal set of relays, producing solutions in faster time than the local search, but with higher cost. We also relaxed the robustness requirement by enforcing biconnectivity instead of disjointness. When doing so, we observed that our QuickXplain based approach provides solutions that are not only competitive in time but also in quality.

Future work will focus on both the decision problem and the optimisation problem. For the decision problem, we will study the performance of the CP model for a wider range of networks. For the optimisation problem, we will implement the following Large Neighbourhood Search (LNS) approach.

A LNS approach to the Additional Relay Placement problem. Once we have a solution of multi-sensor BVUP, we can use CP for implementing moves that iteratively improve the quality of the solution. More precisely, suppose that the set of relays used by the current multi-sensor BVUP solution is $R_1 \subseteq R$. We want to find an $R_2 \subseteq R_1$ that can be replaced with a set R_3 (subset of R) such that $(R_1 - R_2) \cup R_3$ is still a solution and $|R_3| < |R_2|$.

Table 7 Results with the different approaches for 100 sensor nodes and 196 candidate relays

instance	LS				CP			
	Biconnected		Disjoint		Biconnected		Disjoint	
	cost	time	cost	time	cost	time	cost	time
scen-n100-01	7	63925.80	9	22108.70	7	10620.27	17	3845.96
scen-n100-02	7	78736.10	11	13144.90	9	9000.49	18	3486.35
scen-n100-03	5	14690.80	5	997.63	8	17742.10	15	4214.37
scen-n100-04	5	9330.01	5	174.95	5	15668.00	13	4421.60
scen-n100-05	12	12531.60	12	15699.70	13	21409.07	20	4451.92
scen-n100-06	5	36585.00	7	4905.40	5	8935.33	19	4122.45
scen-n100-07	5	17316.50	6	1974.81	5	10450.57	16	4725.40
scen-n100-08	10	62401.30	11	14462.00	9	13890.45	19	4816.48
scen-n100-09	5	23083.20	8	1751.50	8	17058.07	13	4042.70
scen-n100-10	8	33678.40	8	17362.70	13	16633.87	18	3872.07
scen-n100-11	6	25495.50	6	12237.10	8	14398.69	12	4029.24
scen-n100-12	9	29646.20	9	2043.61	10	16505.11	18	4574.32
scen-n100-13	9	24286.60	9	1623.28	9	16635.71	19	4523.04
scen-n100-14	6	17614.30	6	14761.90	5	13209.81	16	4283.17
scen-n100-15	8	23288.30	7	15098.70	7	13589.61	10	2909.79
scen-n100-16	7	70312.50	5	14958.10	6	11315.53	11	2921.89
scen-n100-17	7	30987.20	8	2291.31	8	15061.85	16	4466.98
scen-n100-18	7	81552.10	8	33180.30	7	8848.40	14	3527.96
scen-n100-19	9	66798.60	7	15476.90	7	12118.61	14	3972.77
scen-n100-20	9	47651.50	10	16516.90	10	16781.80	15	3877.77

Let us assume that we set the cardinalities of R_2 and R_3 to α and β (which are parameters to be tuned). The task now is to find a set R_2 of cardinality α to be replaced with an R_3 of cardinality β such that $(R_1 - R_2) \cup R_3$ is still a solution. That is, we have $\alpha + \beta$ decision variables where the domain of those variables in R_2 is R_1 and the domain of those variables in R_3 is $R - R_1$.

Our plan is to use CP for solving this decision problem. It is easy to show that the decision problem to be solved by CP is NP-complete. We prove this by reducing multi-sensor BVUP to this problem. Let \mathcal{R} be the universe of relays for the LNS decision problem. We set \mathcal{R} to $R \cup R_1$ where R_1 (in the reduction) is a set of dummy relays ensuring connectivity such that $|R_1| = |R| + 1$. We constraint $|R_2|$ to be equal to $|R_1|$ and $|R_3|$ to be equal to $|R|$. Finding R_2 and R_3 would be equivalent to solving the given multi-sensor BVUP problem.

The intuition behind the reduction is that checking whether the newly added relays can reconnect the sensors can be complex since, in the worst case, it can be as complex as solving a multi-sensor BVUP instance from scratch. Even though the LNS decision problem is NP-complete, the practicality of the approach would rely on the fact that α and β will be set to small numbers.

Acknowledgements This work was supported by Science Foundation Ireland Grant No. 08/CE/1523

References

- Beldiceanu N, Flener P, Lorca X (2008) Combining Tree Partitioning, Precedence, and Incomparability Constraints. *Constraints* 13(4):459–489
- Boukerche A, Pazzi RWN, Araujo RB (2006) Fault-Tolerant Wireless Sensor Network Routing Protocols for the Supervision of Context-Aware Physical Environments. *Journal of Parallel and Distributed Computing* 66(4):586–599
- Bredin JL, Demaine ED, Hajiaghayi M, Rus D (2005) Deploying Sensor Networks with Guaranteed Capacity and Fault Tolerance. In: *Proc. 6th ACM Int'l Symp. Mobile Ad Hoc Networking and Computing (MobiHoc'05)*, pp 309–319
- Chipara O, He Z, Xing G, Chen Q, Wang X, Lu C, Stankovic J, Abdelzaher T (2006) Real-time Power-Aware Routing in Wireless Sensor Networks. In: *Proc. 14th IEEE Workshop Quality of Service (IWQoS'06)*, pp 83–92
- Cormen TH, Stein C, Rivest RL, Leiserson CE (2001) *Introduction to Algorithms*, 2nd edn. McGraw-Hill Higher Education
- Feo TA, Resende MGC (1989) A Probabilistic Heuristic for a Computationally Difficult Set Covering Problem. *Operations Research Letters* 8:67–71
- Feo TA, Resende MGC (1995) Greedy Randomized Adaptive Search Procedures. *Journal of Global Optimization* 6:109–133
- Gnawali O, Fonseca R, Jamieson K, Moss D, Levis P (2009) Collection Tree Protocol. In: *Proc. 7th ACM Conf. Embedded Networked Sensor Systems (SenSys'09)*
- Golovach PA, Thilikos DM (2011) Paths of Bounded Length and Their Cuts: Parameterized Complexity and Algorithms. *Discrete Optimization* 8(1):72–86
- Han X, Cao X, Lloyd EL, Shen CC (2010) Fault-tolerant Relay Node Placement in Heterogeneous Wireless Sensor Networks. *IEEE Transactions on Mobile Computing* 9(5):643–656
- Junker U (2004) QUICKXPLAIN: Preferred Explanations and Relaxations for Over-Constrained Problems. In: McGuinness DL, Ferguson G (eds) *AAAI, AAAI Press/The MIT Press*, pp 167–172
- Lovász L, Neumann-Lara V, Plummer M (1978) Mengerian Theorems for Paths of Bounded Length. *Periodica Mathematica Hungarica* 9(4):269–276
- Misra S, Hong SD, Xue G, Tang J (2008) Constrained Relay Node Placement in Wireless Sensor Networks to Meet Connectivity and Survivability Requirements. In: *Proc. 27th Ann. IEEE Conf. Computer Communications (INFOCOM'08)*, pp 281–285
- Pu J, Xiong Z, Lu X (2008) Fault-Tolerant Deployment with k -connectivity and Partial k -connectivity in Sensor Networks. *Wireless Communications and Mobile Computing* 9(7):909–919
- Resende MGC, Ribeiro CC (2002) Greedy Randomized Adaptive Search Procedures. In: Glover F, Kochenberger G (eds) *State of the Art Handbook in Metaheuristics*, Kluwer Academic Publishers, pp 219–249

- Sitanayah L, Brown KN, Sreenan CJ (2011) Fault-Tolerant Relay Deployment for k Node-Disjoint Paths in Wireless Sensor Networks. In: Proc. 4th Int'l Conf. IFIP Wireless Days (WD'11), IEEE, pp 1–6
- Xiao Y, Thulasiraman K, Xue G (2006) Constrained Shortest Link-Disjoint Paths Selection: A Network Programming Based Approach. IEEE Trans Circuits and Systems I: Regular Papers 53(5):1174–1187, DOI 10.1109/TCSI.2006.869907
- Zeng Y, Sreenan CJ, Sitanayah L, Xiong N, Park JH, Zheng G (2011) An Emergency-Adaptive Routing Scheme for Wireless Sensor Networks for Building Fire Hazard Monitoring. Sensors 11(3):2899–2919