On the Role of Multi-Objective Optimization in Risk Mitigation for Critical Infrastructures with Robotic Sensor Networks

Jamieson McCausland
Larus Technologies Corporation
58 Antares Drive Unit 1
Ottawa, Canada K2E 7W6
Tel: +1 (613) 244 8916 x215
jamieson.mccausland@larus.com

Rami Abielmona
Larus Technologies Corporation
58 Antares Drive Unit 1
Ottawa, Canada K2E 7W6
Tel: +1 (613) 244 8916 x202
rami.abielmona@larus.com

Rafael Falcon
Larus Technologies Corporation
58 Antares Drive Unit 1
Ottawa, Canada K2E 7W6
Tel: +1 (613) 244 8916 x214
rafael.falcon@larus.com

Ana-Maria Cretu
Université du Quebec en Outaouais
101 Saint-Jean-Bosco, C.P. 1250,
Gatineau (Québec) Canada
+1 (819) 595-3900 x1527
ana-maria.cretu@uqo.ca

Emil M. Petriu
University of Ottawa
800 King Edward Ave
Ottawa, Canada K1N 6N5
+1 (613) 562 5800 x2132
petriu@eecs.uottawa.ca

ABSTRACT
The use of robotic sensor networks (RSNs) for Territorial Security (TerrSec) applications has earned an increasing popularity in recent years. In Critical Infrastructure Protection (CIP) applications, the RSN goal is to provide the information needed to maintain a secure perimeter around the desired infrastructure and efficiently coordinate a corporate response to any event that arises in the monitored region. Such a response will only involve the most suitable robotic nodes and must successfully counter any detected vulnerability in the system. This paper is a preliminary study of the role played by multi-objective optimization (MOO) in the elicitation of responses from a risk-aware RSN that is deployed around a critical infrastructure. Contrary to previous studies showcasing a *pre-optimization* auctioning scheme, where the RSN nodes bid on the basis of their knowledge of the event, we introduce a *post-optimization* auctioning scheme in which the nodes place their bids knowing what their final positions along the perimeter will be, hence calling for a more informed decision at the node level. The impact of the pre- vs. post-optimization stage in a first-price sealed bid auction system over the risk mitigation strategies elicited by the RSN is evaluated and discussed. Empirical results reveal that the pre-optimization auctioning is more suitable for dense RSNs whereas the post-optimization one is preferred in sparse RSNs. To the best of our knowledge, this is the first attempt to assess the role of MOO in risk mitigation for CIP with RSNs.

Categories and Subject Descriptors
I.2.9 [**Robotics**]: – *autonomous vehicles, sensors*

General Terms
Algorithms, Measurement, Experimentation, Security

GECCO'14, July 12–16, 2014, Vancouver, BC, Canada.

Keywords
critical infrastructure protection; robotic sensor network; risk management; evolutionary multi-objective optimization; self-organization; auction protocols; computational intelligence

1. INTRODUCTION
A robotic sensor network (RSN) is a collection of autonomous, often wirelessly connected devices that are capable of sensing, communicating and actuating, both upon each other and on the environment. An RSN could be thought of as a more advanced mobile sensor network [14] in which nodes possess typical robotic features such as gripping or lifting objects. Because of the various benefits provided by their proactive and reactive capabilities, RSNs are increasingly being used in the TerrSec arena, for instance to maintain a secure perimeter around a critical infrastructure [6] [8] [9] [12] and to efficiently coordinate a corporate response to any event that arises in the monitored region.

Recently, risk-aware RSNs for critical infrastructure protection (CIP) that can react to single or multiple events were proposed in [8] and [9], respectively. Two innovative features stand behind this novel class of CIP solutions: (1) the self-organization of the RSN nodes, (i.e., the crafting of a risk mitigation response), is the outcome of a thorough risk analysis conducted on every robotic node as well as on the entire network; this analysis is driven by a group of risk features [3] which, in turn, are extracted from the raw data streams emanating from each node in the monitoring region, and (2) evolutionary MOO algorithms are used by the response coordinator (i.e., the RSN node that best perceives the event) to locally derive a set of promising candidate responses, each of which is judged according to several conflictive objectives, such as latency or cost.

When multiple concurrent events arise in the region of interest, the framework in [9] allows the RSN nodes to formulate their bids for each event by taking into account important aspects such as their current battery level, the distance to the event or the amount of redundant coverage in that area. The bids are submitted to the response coordinator for that event, which selects the winners following the first-price sealed bid auction rules and then proceeds to optimize their final locations along the perimeter. Subsequently, the winners are instructed to relocate to such positions.

Despite the encouraging results in terms of perimeter coverage and response efficiency obtained with the above approach, there are key factors in the entire risk mitigation process whose influence remains largely unexplored. One of them is the MOO component and its relation with the bidding phase. One could imagine that, under different bidding conditions and node-specific valuation systems, the responses could considerably vary, hence impacting the overall quality of the CIP solution.

This paper is a preliminary study of the role played by MOO in the elicitation of responses from a risk-aware RSN that is deployed around a critical infrastructure. Contrary to the *pre-optimization* auctioning scheme described in [8] and [9], where the RSN nodes bid on the basis of their knowledge of the event, we consider a *post-optimization* auctioning scheme in which the nodes place their bids knowing what their final positions along the perimeter will be, hence calling for a more informed decision at the node level. The impact of the pre- vs. post-optimization stage in a first-price sealed bid auction system over the risk mitigation strategies elicited by the RSN is evaluated and discussed. Empirical results reveal that the pre-optimization auctioning is more suitable for dense RSNs whereas the post-optimization one is preferred in sparse RSNs. To the best of our knowledge, this is the first attempt to assess the role of MOO in risk mitigation for CIP with RSNs.

The rest of the paper is structured as follows. Section II reviews some relevant works. Section III elaborates on the proposed post-optimization bidding scenario. Experimental results are discussed in Section IV while conclusions are outlined in Section V.

2. RELATED WORK

This section briefly reviews several works in the areas of market-based robot coordination and risk-aware robotic systems for CIP.

2.1 Market-based Robot Coordination

Market-based task allocation is a well-studied and quite successful coordination paradigm in robotics [1] [4] [5] [7-10][12]. The idea of casting different tasks in the system as auctions and letting robots bid on them to maximize their individual profit has proved valuable in improving the overall team efficiency.

Gerkey and Mataric [4] developed MURDOCH, an auction-based task allocation system which features a publish/subscribe communication model. The main goal is to create an inter-robot collaboration framework in which responses to tasks are decided in a cooperative manner. The protocol can handle robots of dissimilar types and capabilities. Only robots that are properly equipped to handle a particular task can bid in response to the corresponding broadcast auction. MURDOCH names a single winner and monitors its progress in real-time.

Mezei *et al* [10] address robot-to-robot collaboration with the goal of optimizing several objectives like communication costs, task execution costs and latency. The proposed task allocation framework relies on the aggregation of different auctions in a distributed fashion. Five auction protocols are discussed and compared to a simple auction protocol. The authors only focus on a single-robot, single-task scenario in which the cost metric is inversely proportional to the distance from the event.

Pustowka and Caicedo [12] introduce multi-robot collaboration in a surveillance application. Given a grid representation of a known environment, the goal is to have robots cover least frequently visited cells. Nearby grid cells are clustered with K-means. These clusters of cells become the tasks robots will bid for. The cost used in the utility function is the length of the path to the desired location, as computed by the A* search algorithm. The highest bidder for each task becomes the winner. This approach, however, requires full a priori knowledge of the environment and suffers from sensitivity to crucial parameters like the number of clusters (K) to be found.

Kaleci et al [5] consider a team of heterogeneous robots with three types of tasks: cleaning, carrying and monitoring. The sensing and motion models of each robot contribute to the formulation of the price and cost of a task and the authors employ two schemes to determine the auction winners. Additionally, the Hungarian algorithm is used to solve the robot-to-task assignment problem.

The above approaches illustrate the effectiveness of market-based robot cooperation. However, none of them features either MOO or risk analysis in the elicitation of candidate responses to the events. They are neither concerned with CIP.

2.2 Risk-Aware Robotic Systems for CIP

Safeguarding a critical infrastructure with an RSN typically involves appropriate situational assessment, risk detection, mitigation and prevention. The lifecycle portrayed in Fig. 1 governs the operational workflow of the RSN in this context.

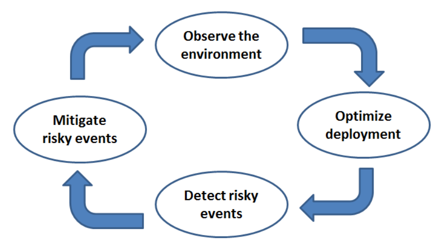

Figure 1. Lifecycle of a RSN for CIP.

An RSN for CIP is proposed in [8]. The novelty with this approach is that the RSN implements the risk management framework put forth in [2]. From the information submitted by each RSN node, three risk features are extracted: *degree of distress*, *intruder proximity factor* and *terrain maneuverability*. The overall risk posed by any RSN node is assessed on the basis of these local risk features and those nodes exceeding a permissible risk threshold are flagged as "nodes in distress" (NIDs), i.e., robotic entities likely to originate a security breach in the perimeter. The response from the RSN is to self-organize in order to maintain as much coverage as possible around the critical infrastructure while minimizing the cost of doing so. For each NID, a restricted group of candidate topologies are evolved with NSGA-II [2] as the evolutionary MOO algorithm of choice and then ranked following some network operator preferences. The network manager then decides on the most suitable response topology, which is enacted upon the environment.

The above framework could only handle a single response to a single event. That is, a network response must be effectuated before another response (such as one to another event) could be orchestrated. With this limitation in mind, the authors in [9] later augmented this approach to handle multiple concurrent events, hence creating several independent network regions which can autonomously evolve their own response sets and thus mitigate their local risks. Each response coordinator (i.e. the perimeter node that best perceives the event) initiates an auction to all RSN nodes by advertising the event location. The nodes calculate their availability by using a Sugeno fuzzy inference system (FIS) [11] and decide whether to bid or not. The response coordinator selects the winners according to the rules of the first-price sealed bid auction and then evolves their final positions with NSGA-II. Finally, the winner nodes relocate to these positions along the perimeter.

The bidding phase in [8] and [9] takes place before the actual optimization of the responses. A more informed approach would be to evolve the target perimeter locations first, advertise them to all nodes and then conduct the bidding process. This paper is aimed at gauging the impact that pre- vs. post-optimization has over the risk mitigation strategies elicited by the RSN.

3. PROPOSED METHODOLOGY

In this section, we unfold the methodological approach involving MOO and market-based robot coordination for a risk-aware RSN in a CIP scenario.

As envisioned in [8] and [9], the robotic nodes surrounding the critical infrastructure play multiple roles in the system. First, they can *detect high-risk events* that could jeopardize the integrity of the surveillance conducted by the RSN upon the region of interest. Second, they can become *auctioneers* for a particular event, which means they are responsible for advertising some event-related information to the rest of the RSN nodes, gathering the received bids and clear off the auction, i.e., announce the winner(s). Finally, they could participate as *bidders* in any auction circulating around the network.

3.1 Risk-Driven Event Detection

The RSN nodes in our study can detect two types of events that are quite relevant to the CIP context: (1) another node's failure (e.g., because of battery depletion) or (2) an attempted perimeter intrusion. Both events demand a corporate RSN response in the form of a topological self-organization, whether to fill the coverage gap brought forth by the failed node or to increase the sampling capabilities around a portion of the perimeter where the intrusion attempt might have taken place.

The Risk Management Framework (RMF) in [3] enunciated by Falcon *et al* is the backbone of the event detection phase in our study. From the raw data stream periodically submitted by each RSN node to the operations center, a parallel risk stream is dynamically extracted. This risk stream consists of a collection of user-defined *risk features* which are modeled after different constructs. In our example, each RSN node reports: (1) its current battery level, in percentage; (2) its distance, in meters, from a potential intruder and (3) the terrain maneuverability index associated with its geographical location. The distance from a potential intruder can be gauged with the help of a laser range finder in order to provide the RSN node with depth perception. The terrain maneuverability index corresponding to the node's location can be queried from the Knowledge Base of the RSN's deployment environment. In this study, a random number in [0; 1] has been used for each location. Table 1 depicts the raw-feature-to-risk-feature mapping in our CIP case study.

Table 1. Risk feature extraction

Raw Feature	Risk Feature	Modeling Construct	Parameters / Expression
Battery Level (%)	Degree of Distress ([0;1])	Fuzzy set with a triangular membership function	A = 0 B = 0 C = 100
Distance to Potential Intruder (m)	Intrusion Risk ([0;1])	Fuzzy set with a trapezoidal membership function	A = 0 B = 0 C = 1 D = 5
Geographical position (x, y)	Terrain Risk ([0;1])	Linear mapping	1 – terrain maneuverability index of the node's position

The raw data input vector submitted by each RSN node is mapped onto an output risk vector, defined as a collection of local risk values, one per risk feature. The local risk value of any risk feature is determined by applying its modeling construct and parameters/expression to the corresponding input (raw feature data) in Table 1. As per the RMF flow in [3], the overall risk assessment for an RSN unit is obtained after the application of an FIS to the collection of local risk values. For simplicity, our Mamdani-type FIS [11] consists of a single fuzzy rule that reads as follows:

```
IF Degree of Distress is DD-HIGH OR
Intrusion Risk is IR-HIGH OR
Terrain Risk in TR-HIGH THEN
Overall Risk is OR-HIGH
```

The linguistic terms DD-HIGH, IR-HIGH, TR-HIGH and OR-HIGH are modeled as fuzzy sets by taking into account the network manager's knowledge about these local risks. If the overall risk of any RSN node exceeds a user-set threshold, then it is flagged as a "node in distress" and the pursuit of a corporate risk mitigation strategy to assist that node is initiated.

3.2 Pre-Optimization Auctioning Scheme

As explained in Section 1, a market-based approach is the mechanism through which inter-robot coordination emerges in [8] and [9]. The node that best perceives the event (NID) becomes its auctioneer and starts advertising the NID's location to all remaining RSN nodes. Upon receipt of this message, an *available* RSN node (i.e., one that is not currently part of any other risk mitigation strategy) will formulate its bid and communicate it back to the auctioneer. Once all bids have been gathered, the auctioneer determines the winner(s) and begins the optimization phase to find out which among those winning nodes will be used in the response and where they should relocate.

3.2.1 Node Bidding

Each robotic node formulates its *availability metric* (bid) to the announced task by considering three major factors: its current *battery level*, its *distance to the event* and its *coverage redundancy*. A Sugeno FIS [11] is employed to render the final bid value. Figure 2 depicts the fuzzy system embedded on each robotic node.

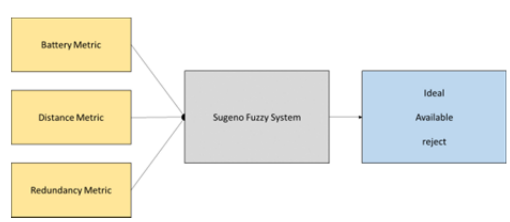

Figure 2. Sugeno fuzzy model for node bidding

The *battery metric* is an incentive to bid based on the node's remaining energy. We treat it as a linguistic variable with the following set of linguistic terms and their configurations:

- *Poor battery*: Triangular membership function with A=0.0, B=0.0 and C=0.4
- *Average battery*: Trapezoidal membership function with A=0.1, B=0.35, C=0.65 and D=0.9
- *Good battery*: Triangular membership function with A=0.6, B=1.0 and C=1.0

The *distance metric* is the node's incentive to bid based on the distance from the risk event. The closer the robotic node is to the event, the higher its motivation as it will not spend too much energy in locomotion. The normalized node-to-NID distance is the input to all the linguistic terms of this linguistic variable:

- *Far distance*: Trapezoidal membership function with A=0.0, B=0.0, C=0.65 and D = 1.0
- *Near distance*: Triangular membership function with A=0.5, B=1.0 and C=1.0

Finally, the *redundancy metric* is the node's incentive to bid given the amount of redundant coverage in its section of the perimeter, i.e., the overlap among the field of views of the RSN nodes w.r.t the set of perimeter points in that section. The more overlap, the higher the node's incentive to leave its current position and relocate somewhere else. The *redundancy metric* linguistic variable accepts a redundancy ratio as an input to its collection of linguistic terms. The redundancy ratio is the percentage of the perimeter points surveyed by this node that are also covered by nearby nodes. The linguistic terms for this fuzzy variable and their configurations are given below:

- *Low redundancy*: Triangular membership function with A=0.0, B=0.0 and C=0.4
- *Medium redundancy*: Triangular membership function with A=0.1, B=0.5 and C=0.9
- *High redundancy*: Triangular membership function with A=0.6, B=1.0 and C=1.0

The membership grade for a linguistic variable is the highest among the membership grades of its linguistic terms. Then, the availability metric of the bidder node is simply the average of the membership grades for the three linguistic variables under consideration. An ideal bidder yields 1.0 for availability while a rejection of the bid yields 0.0. All nodes not rejecting the bid submit their availabilities to the auctioneer, who then selects the winners and notifies them accordingly.

3.2.2 Multi-Objective Optimization

The auctioneer then proceeds to optimize the number of the winner nodes that are actually required to respond to the event and their target locations via NSGA-II [2].

3.2.2.1 Chromosome Encoding

Figure 3 portrays the structure of a solution (i.e. chromosome) in the MOO scheme conducted after the bidding phase has finished.

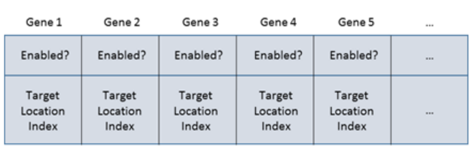

Figure 3. A two-layered chromosome encoding. The *node selection* layer decides on the inclusion of a winner node in the event response. The *target location* layer represents the index of the perimeter point where it should be relocated.

3.2.2.2 Genetic Operators

After two parent chromosomes are selected, crossover is probabilistically decided. Two crossover operators are proposed:

- *One-point crossover* for the node selection layer. This means that a single crossover point is randomly chosen and both parents exchange genetic segments in the node selection layer.
- *Uniform crossover* for the target selection layer. This means that genes in the target selection layer are drawn, one at a time, from a random parent, until a child chromosome is completed.

Mutation is also probabilistically decided for each population member. When a chromosome is mutated, either layer could be affected. A random gene is chosen in the node selection layer and its value is negated. On the target selection layer, if a gene is to be mutated, a new random target index on the interval of [0; T] will be selected, where T is the number of target perimeter points.

3.2.2.3 Objective Functions

Each population individual (risk mitigation strategy) is evaluated according to two mutually conflictive objectives: the *perimeter coverage* it provides and the *energy cost* involved in enacting its response, i.e., the energy spent in relocating all participant nodes to their target positions.

It is clear from the above description that in this pre-optimization auctioning scenario, the RSN nodes do not possess any knowledge about their final deployment along the perimeter; hence their bids are solely based on the circumstances governing their current deployment and on the location of the event.

3.3 Post-Optimization Auctioning Scheme

An interesting alternative to the aforementioned risk mitigation methodology in Section 3.2 is to conduct the optimization phase first and, based on its results, let the RSN nodes place their bids.

3.3.1 Multi-Objective Optimization

In this case, the auctioneer proceeds to optimize via NSGA-II the target locations of the hypothetical responders, i.e., where should the responders be placed along the perimeter so that the coverage is restored as much as possible and the amount of overlap among the responders is reduced as little as possible.

3.3.1.1 Chromosome Encoding

A variable-length chromosome encoding is considered, where each gene contains the $<x, y>$ coordinates of the target location for a hypothetical responder. The cardinality of the chromosome (number of genes) varies from 1 to the maximum number of possible responders in the RSN. Figure 4 displays a response orchestrated by four RSN nodes (not decided yet).

Gene 1	Gene 2	Gene 3	Gene 4
x_1	x_2	x_3	x_4
y_1	y_2	y_3	y_4

Figure 4. A variable-length chromosome encoding the target locations of four hypothetical responder nodes

When the population is first initialized, solutions of different cardinality are injected in the population, with their coordinates within the $<x, y>$ bounding box of the perimeter section subject to intrusion/node failure. This bounding box can be easily computed by the auctioneer prior to the execution of the NSGA-II scheme.

3.3.1.2 Genetic Operators

A single-point crossover operator is applied probabilistically. A single-gene mutation operator is also probabilistically enforced on each population member. The mutation is interpreted as the random generation of another <x, y> pair within the vicinity of the existing one and inside the bounding box.

3.3.1.3 Objective Functions

The maximization of the *perimeter coverage* achieved by the chromosome (response) and the minimization of the *overlap* among the nodes in the response stand as the two decision objectives to be evolved via NSGA-II.

3.3.2 Node Bidding

The bidding process in this scenario follows the same procedure outline in Section 3.2.1 with the addition of a fourth indicator: the *delta overage metric,* i.e., the difference between robot coverage at the optimized position minus the coverage of the robot at its original position. $DeltaCoverage = \{Coverage\_Optimized\_Position\} - \{Coverage\_Current\_Position\}$.

It is easy to notice that the coverage gain ratio is a more informative measure for the bidder as it takes into account the consequences of leaving the node's current location and the benefits (in terms of coverage restored) of doing so. Hence, the bidder's utility function (availability metric) is more reflective of the true internal valuation of a node in a multi-agent system. The downside is that nodes behave more selfishly and this implies that the high-risk events the network is currently experiencing might not receive a suitable degree of responsiveness.

4. EXPERIMENTAL VALIDATION

This section is devoted to the empirical validation of the proposed pre- and post-optimization auctioning schemes. We simulate a CIP scenario in Microsoft Robotics Developer Studio 4 (MRDS)[1]. The setting is a military compound with two access points at paved roads. Two different RSN deployments are investigated: a *sparse* scenario of 36 nodes and a *dense* scenario of 57 nodes, each deployed along the perimeter of the compound, thus creating a virtual fence.

The experiment simulates a single node being at risk caused by a nearly depleted battery. In the experiment, the pre- and post-optimization auctioning schemes will be compared based on: ***coverage recovery***, ***response cost***, and ***response time***.

4.1 Sparse Scenario: Pre-Optimization Case

A loss of functionality of an RSN node will most likely induce coverage gaps, thus reducing the chance of a detected intrusion. This case was explored by randomly initializing the battery levels of the robotic nodes using a uniform distribution between 20% and 100%. In the case of the *sparse* scenario, robot #23 became a NID with battery level 10% with distress risk 96.6%. A risk threshold of 90% was used, which triggered the risk event. Node #22 handles the response generation and will assume the role of *auctioneer*. Node #22 takes the role of auctioneer because of it was the closest to the NID.

[1] http://www.microsoft.com/robotics Accessed 2014-04-02.

Due to the risky state of node #23, portions of the security perimeter under its jurisdiction are at stake, i.e. $-70.119 \leq x \leq -69.817$ and $33.577 \leq y \leq 48.019$, as shown in Fig. 5. The risky region compromises the risk spatial information broadcast by robot #22 (auctioneer) to all other RSN nodes in a *task announcement message*.

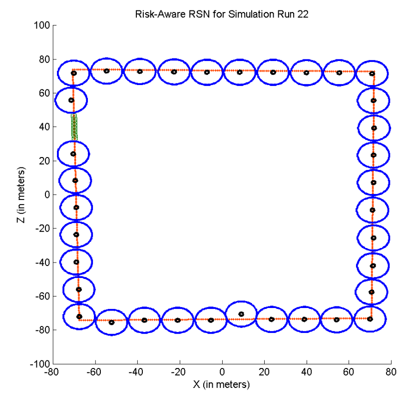

Figure 5. A two-dimensional representation of the RSN for the *sparse* scenario. Node locations are represented by small black circles. Each node's field of view is visualized by the blue circular region. Risky perimeter points are circles in green at the top left of the figure.

The RSN nodes receive the task announcement message and formulate their bids as described in Section 3.2.1. Table 2 lists the RSN nodes that did not reject the auction.

Table 2. Available bidders for the scenario in section 4.1

Node	$x_{battery}$	$x_{coverage}$	$x_{distance}$	$y_{availability}$
20	0.971	0.0	43.157	0.295
21	0.997	0.0	39.191	0.458
22	0.999	0.0	23.260	1.000
24	0.989	0.0	22.961	1.000
25	0.998	0.0	38.701	0.478

The five bidders form a candidate response group, which will be optimized using NSGA-II. The crossover and mutation probabilities were set to 0.85 and 0.1, respectively. A maximum Pareto Archive Set (PAS) size of 10 is used and the initial population contains 50 individuals. The stop criterion for the optimization is reaching 100 generations. Table 3 reports the non-dominated solutions found in the PAS once the optimization finished.

Table 3. Candidate response set for section 4.1

Response	Nr. Responders	Coverage (%)	Energy Cost (%)
1	0	83.16	0.00
2	**1**	**84.21**	**7.90**

Response 2 is selected by the network manager as the topology of choice since it has the highest coverage of 84.21% (recovery of 1.05%) at an energy cost of 0.07% (based on the battery usage model in the simulator). This solution is illustrated in Fig. 6.

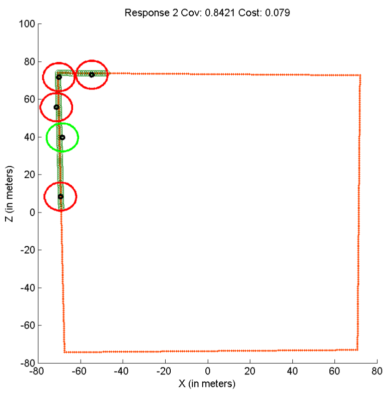

Figure 6. The RSN after response 2 in Table 3 was enforced. Only the 5 bidding nodes are displayed. Red robots remain at their original position while green robots moved to the optimized position.

The response took 50.76 seconds to generate, recovering 1.06% coverage initially lost by the NID. The responding robot executed the response with an energy cost of 7.90%.

4.2 Sparse Scenario: Post-Optimization Case

The post-optimization auctioning scheme was carried out on the exact same scenario. This time, the auctioneer optimizes the target locations of hypothetical responders first and then runs the auction. The PAS after 10 generation is shown in Table 4, which describes the optimal position, the coverage of the risk region and the number of redundant grid points surveyed by the responders.

Table 4. Candidate response set for section 4.2

Response	Solution	Coverage Risk Region (%)	Redundant Sensor Coverage (%)
1	**(-70.011, 39.667)**	**100.0**	**0**
2	(-69.810, 39.945)	100.0	0
3	(-70.090, 40.051)	100.0	0
4	(-69.985, 39.838)	100.0	0
5	(-69.919, 39.854)	100.0	0
6	(-69.896, 39.563)	100.0	0

All solutions provide the maximum coverage in the risky region with no overlap. Response 1 (in bold) is arbitrarily selected. The auction session revealed no bidders in the auctioned position, hence the nodes decided not to respond to this event, for otherwise a bigger coverage gap would have been created if they abandoned their original locations. The time taken to generate the response was 26.93 seconds.

4.3 Dense Scenario: Pre-Optimization Case

In the *dense* scenario, node #32 is flagged as a NID due to a low battery of 10% resulting in a distress risk of 96.6%. Node #33 handles the response generation and assumes the role of *auctioneer*. Node #33 takes the role of auctioneer because it is the closest node to the NID. The grid points at risk are within $-33.107 \leq x \leq -19.626$ and $73.576 \leq y \leq 73.692$ (shown in Fig. 7).

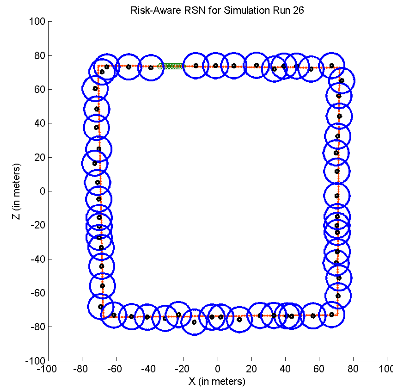

Figure 7. A two-dimensional representation of the RSN for the *dense* scenario. Risky coverage perimeter points are circles in green at the top-left of the figure.

The RSN nodes receive the task announcement message from the auctioneer and formulate their bids as portrayed in Table 5.

Table 5. Available bidders for the scenario in section 4.3

Node	$x_{battery}$	$x_{coverage}$	$x_{distance}$	$y_{availability}$
34	0.999	0.267	31.933	0.699
33	0.999	0.313	18.818	1.000
36	0.999	1.000	47.798	0.283
35	0.999	0.900	44.823	0.395
30	0.999	0.267	32.878	0.657
31	0.999	0.200	21.181	1.000
29	0.999	0.375	43.620	0.500
37	0.999	0.533	53.376	0.067

The eight bidders form a response group. Table 6 shows the PAS once the optimization has completed.

Table 6. Candidate response set for section 4.3

Response	Nr. Responders	Coverage (%)	Energy Cost (%)
1	0	89.38	0.00
2	1	90.27	0.87
3	1	92.04	2.43
4	**2**	**99.12**	**18.60**
5	2	92.92	2.88
6	1	91.15	1.08

Response 4 (in bold) offers the maximum coverage recovery and is selected as the desired response, depicted in Fig. 8.

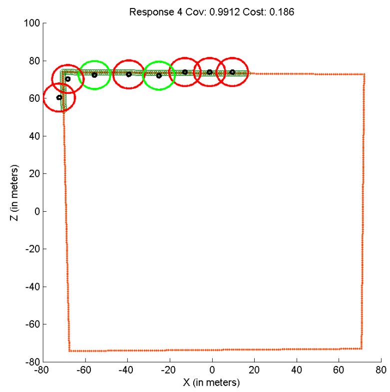

Figure 8. The RSN after response 4 in Table 6 was enforced. Only the eight bidding nodes are displayed. Robots in red remain at their original positions. Those in green move to their optimized locations.

This response took 60.50 seconds, recovered 9.74% coverage initially lost by the NID. The responding robots executed the response with an energy cost of 18.60%.

4.4 Dense Scenario: Post-Optimization Case

In this case, the response generation technique is the post-optimization auctioning for the same NID as in Section 4.3. The NSGA-II is executed over the risky region defined in the previous section as depicted in Fig. 7. The optimal position is <-23.734, 73.636>, which would recover the entire risky region with no overlap in sensor coverage. This position is auctioned by Node #33 to the entire RSN. The bids placed are tabulated in Table 7.

Table 7. Available bidders for the scenario in section 4.4

Node	$x_{battery}$	$x_{coverage}$	$x_{distance}$	$x_{deltaCov}$	$y_{availability}$
47	0.998	0.812	115.931	0.376	0.148
12	0.999	0.800	144.415	0.376	0.099
27	0.998	0.866	57.178	0.418	0.147
36	**0.999**	**0.238**	**44.664**	**0.502**	**0.409**
15	0.998	0.812	129.489	0.376	0.121
55	0.998	0.813	146.424	0.376	0.098
57	0.998	0.867	149.300	0.418	0.056
1	0.998	0.875	150.105	0.418	0.042

The auction session revealed that node #36 would be the best responder for the optimized position (bold entry in Table 7). This response is schematically depicted in Fig. 9.

The response achieved perimeter coverage of 100% (10.62%) or full recovery. This response was generated in 64.03 seconds with an energy cost of 22.33%.

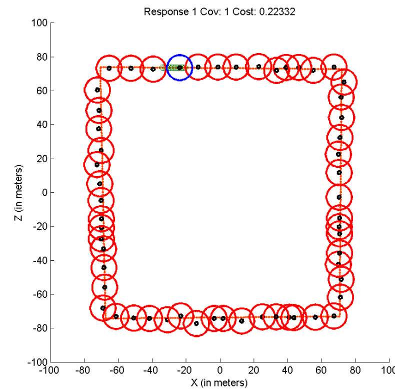

Figure 9. The RSN after node #36 responds by relocating to the target position. Robots in red remain at their original positions. The robot in blue is Node #36 at auctioned position.

4.5 Discussion

The two auctioning techniques approach the same problem from different standpoints, both using the NSGA-II evolutionary multi-objective optimizer and a fuzzy-based auction technique. In the case of the pre-optimization auctioning, a subset of the RSN nodes is gathered prior to the optimization. This set of potential responders is then evolved to determine the actual number of nodes needed to provide adequate coverage as well as their perimeter positions. On the other hand, with the post-optimization auctioning, we allow the system to determine nearly-optimal perimeter positions for hypothetical responders and their number. This information elicited by the NSGA-II is then bid out to the RSN members. The advantage witnessed with the post-optimization case is that a robot is able to use the new coverage gain versus coverage lost ratio to affect the bid process.

The response generation time in the pre-optimization case is usually longer than in the post-optimization scheme given that the former must explore all possible locations within the spatial range of all selected responder robots whereas the latter is confined to exploring the spatial region around the high-risk perimeter segment. We can expect higher response costs in the post-optimization case since the current geographical locations of the actual responders are not taken into consideration during the evolutionary process, as they are in the pre-optimization case.

The post-optimization auctioning scheme can offer superior advantages to its pre-optimization counterpart, but this depends on the density of the deployed RSN (i.e., number of nodes) and the amount of perimeter coverage overlap among them. In an RSN with fewer nodes and scarce overlap, an RSN member under the pre-optimization auctioning scheme will gladly leave its position to relocate to the event area, likely causing a coverage gap at its current location. Nodes under the post-optimization auctioning scheme will be more thoughtful about this reality and may decide not to bid at all (or bid very poorly) hence preserving the quality of the surveillance operation in their area of responsibility.

5. CONCLUSIONS

This paper has elaborated on the role played by MOO in the elicitation of different risk mitigation responses from an RSN in a CIP scenario. Two schemes have been considered. The first one establishes that the RSN nodes bid on the basis of event-related information and then the number of responders and their final locations along the perimeter are optimized. In the second one, the final locations and number of hypothetical responders are optimized first and then the bidding process takes place.

The empirical evidence seems to indicate that the pre-optimization auctioning method is preferred when the robots have low sensing ranges or there is enough coverage redundancy in the perimeter, e.g., in the case of a dense RSN. The post-optimization auctioning technique is better at handling larger robot response ranges since it does not require exploring the response region of each individual robot. It also gives RSN bidders enough coverage grounds to decide whether it is worth relocating or not, which is very helpful in the case of a sparse RSN with low coverage redundancy in order to avoid gaps due to the relocation process.

As a future research direction, we would like to measure how influential the underlying auction protocol is in the elicitation of risk mitigation strategies by an RSN. Thus far, only the first price sealed bid auction protocol has been studied. It would be interesting to shed light on what the RSN response landscape looks like when other well-known auction protocols like English, Dutch or Vickrey are enforced.

6. REFERENCES

[1] Dias, M.B., Zlot, R., Kalra, N., and Stentz, A. 2006. Market-based multirobot coordination: a survey and analysis. *Proceedings of the IEEE*. 94, 7 (July 2006), 1257-1270. DOI= http://dx.doi.org/10.1109/JPROC.2006.876939

[2] Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans. Evolutionary Computation*. 6, 2 (April 2002), 182-197. DOI= http://dx.doi.org/10.1109/4235.996017

[3] Falcon, R., Abielmona, R., and Nayak, A. 2011. An evolving risk management framework for wireless sensor networks. In *Proceedings of the IEEE Int'l Conference on Computational Intelligence for Measurement Systems and Applications* (Ottawa, Canada, Sept 19-21, 2011). CIMSA'11, 1-6. DOI= http://dx.doi.org/10.1109/CIMSA.2011.6059924

[4] Gerkey, B.P., and Mataric, M.J. 2002. Sold!: auction methods for multirobot coordination. *IEEE Trans. Robotics and Automation*. 18, 5 (October 2002), 758-768. DOI= http://dx.doi.org/10.1109/TRA.2002.803462

[5] Kaleci, B. 2010. Market-based task allocation by using assignment problem. In *Proceedings of the IEEE International Conference on Systems Man and Cybernetics* (Istanbul, Turkey, October 10-13, 2010). SMC'10, 135-141. DOI= http://dx.doi.org/10.1109/ICSMC.2010.5642222

[6] Langendoerfer, P. Wireless sensor and actuator networks for critical infrastructure protection. 2011. URL: http://www.wsan4cip.eu/

[7] Li, X., Falcon, R., Nayak, A., and Stojmenovic, I. 2012. Servicing wireless sensor networks by mobile robots. *IEEE Communications Magazine*. 50, 7 (July 2012), 147-154. DOI= http://dx.doi.org/10.1109/MCOM.2012.6231291

[8] McCausland, J., DiNardo, G., Falcon, R., Abielmona, R., Groza, V., and Petriu, E. 2013. A proactive risk-aware robotic sensor network for critical infrastructure protection. In Proceedings of the *IEEE International Conference on Computational Intelligence and Virtual Environments for Measurement Systems and Applications* (Milan, Italy, July 15-17, 2013). CIVEMSA'13, 132-137. DOI= http://dx.doi.org/10.1109/CIVEMSA.2013.6617409

[9] McCausland, J., Abielmona, R., Falcon, R., Cretu, A.M., and Petriu, E. 2013. Auction-based node selection of optimal and concurrent responses for a risk-aware robotic sensor network. In *Proceedings of the IEEE International Symposium on Robotic and Sensor Environments (*Washington, USA, October 21-23, 2013). ROSE'13, 136-141. DOI= http://dx.doi.org/10.1109/ROSE.2013.6698432

[10] Mezei, I., Malbasa, V., and Stojmenovic, I. 2009. Auction aggregation protocols for wireless robot-robot coordination. In *Ad Hoc, Mobile and Wireless Networks*. LNCS 5793, 180-193. DOI= http://dx.doi.org/10.1007/978-3-642-04383-3_14

[11] Negnevitsky, M. 2005. Artificial Intelligence: a guide to intelligent systems. Pearson Education Ltd. Essex, England

[12] Pustowka, A., and Caicedo, E.F. 2012. Market-based task allocation in a multi-robot surveillance system. In *Proceedings of the 2012 Brazilian Robotics Symposium and Latin American Robotics Symposium (*Fortaleza, Brazil, October 16-19, 2012). SBR-LARS'12, 185-189. DOI= http://dx.doi.org/10.1109/SBR-LARS.2012.37

[13] Rodriguez-Canosa, G., del-Cerro G., and Cruz A.B. 2014. Detection and tracking of dynamic objects by using a multirobot system: application to critical infrastructures surveillance. *Sensors* 14, 2 (Feb 2014), 2911-2943. DOI = http://dx.crossref.org/10.3390%2Fs140202911

[14] Wang, G, Cao, G., La Porta, T., and Zhang, W. 2005. Sensor relocation in mobile sensor networks. In *Proceedings of the 24th Annual Joint Conference of the IEEE Computer and Communication Societies.* (Miami, USA, March 13-17, 2005). INFOCOM'05. 2302-2312 vol 4, DOI= http://dx.doi.org/10.1109/INFCOM.2005.1498517