

This is an electronic reprint of the original article. This reprint may differ from the original in pagination and typographic detail.

Integrating Learning, Optimization, and Prediction for Efficient Navigation of Swarms of Drones

Majd, Amin; Ashraf, Adnan; Troubitsyna, Elena; Daneshtalab, Masoud

Published in:

2018 26th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP)

DOI:

[10.1109/PDP2018.2018.00022](https://doi.org/10.1109/PDP2018.2018.00022)

Published: 01/01/2018

[Link to publication](#)

Please cite the original version:

Majd, A., Ashraf, A., Troubitsyna, E., & Daneshtalab, M. (2018). Integrating Learning, Optimization, and Prediction for Efficient Navigation of Swarms of Drones. In I. Merelli, P. Lio, & I. Kotenko (Eds.), *2018 26th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP)* (pp. 101–108). IEEE. <https://doi.org/10.1109/PDP2018.2018.00022>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Copyright Notice

The document is provided by the contributing author(s) as a means to ensure timely dissemination of scholarly and technical work on a non-commercial basis. This is the author's version of the work. The final version can be found on the publisher's webpage.

This document is made available only for personal use and must abide to copyrights of the publisher. Permission to make digital or hard copies of part or all of these works for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage. This works may not be reposted without the explicit permission of the copyright holder.

Permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the corresponding copyright holders. It is understood that all persons copying this information will adhere to the terms and constraints invoked by each copyright holder.

IEEE papers: © IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The final publication is available at <http://ieeexplore.ieee.org>

ACM papers: © ACM. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The final publication is available at <http://dl.acm.org/>

Springer papers: © Springer. Pre-prints are provided only for personal use. The final publication is available at <link.springer.com>

Integrating Learning, Optimization, and Prediction for Efficient Navigation of Swarms of Drones

Amin Majd*, Adnan Ashraf*, Elena Troubitsyna*, and Masoud Daneshtalab†
*Faculty of Natural Sciences and Technology, Åbo Akademi University, Turku, Finland
Email: amajd@abo.fi, aashraf@abo.fi, etroubit@abo.fi
†Division of Intelligent Future Technologies, Mälardalen University, Västerås, Sweden
Email: masoud.daneshtalab@mdh.se

Abstract—Swarms of drones are increasingly been used in a variety of monitoring and surveillance, search and rescue, and photography and filming tasks. However, despite the growing popularity of swarm-based applications of drones, there is still a lack of approaches to generate efficient drone routes while minimizing the risks of drone collisions. In this paper, we present a novel approach that integrates learning, optimization, and prediction for generating efficient and safe routes for swarms of drones. The proposed approach comprises three main components: (1) a high-performance dynamic evolutionary algorithm for optimizing drone routes, (2) a reinforcement learning algorithm for incorporating the feedback and runtime data about the system state, and (3) a prediction approach to predict the movement of drones and moving obstacles in the flying zone. We also present a parallel implementation of the proposed approach and evaluate it against two benchmarks. The results demonstrate that the proposed approach allows to significantly reduce the route lengths and computation overhead while producing efficient and safe routes.

Index Terms—Path planning; drone; swarm; evolutionary algorithms; imperialistic competition algorithm; machine learning; prediction

I. INTRODUCTION

A drone is a semi-autonomous aircraft that can be controlled and operated remotely by using a computer along with a radio-link [1]. Commercially-available drones are increasingly been used in a variety of applications such as monitoring and surveillance, search and rescue operations, photography and filming, media coverage of public events, and aerial package delivery. However, with the growing popularity and use of drones for consumer applications, the number of incidents involving drones is also increasing dramatically. In the United States alone, the Federal Aviation Administration receives more than 100 reports every month of unauthorized and potentially hazardous drone activity reported by pilots, citizens, and law enforcement¹. Ensuring a hazard-free, safe flight is also equally important for indoor applications. Therefore, motion safety of drones is a prime concern for drone operators, which refers to the ability of the drones to detect and avoid collisions with static and moving obstacles in the environment.

For larger and more complex applications which are either beyond the capabilities of a single drone or cannot be performed efficiently if only a single drone is used, multiple

drones are used together in the form of a swarm or a fleet. In such scenarios, a safe operation can not be guaranteed without preventing the drones from colliding with one another and with static and dynamically appearing, moving obstacles in the flying zone. Therefore, in the context of a drone swarm, ensuring motion safety entails devising and implementing a motion path planning and navigation system for multiple drones with an integrated support for collision prediction and avoidance.

In this paper, we present a novel approach that integrates learning, optimization, and prediction for generating efficient and safe routes for swarms of drones. We assume that the swarm executes certain missions, in which each drone flies from a start location to a destination location. The proposed approach uses geographical locations of the drones and of the successfully detected, static and dynamically appearing, moving obstacles to predict and avoid: (1) drone-to-drone collisions, (2) drone-to-static-obstacle collisions, and (3) drone-to-moving-obstacle collisions. The proposed approach comprises three main components: (1) a high-performance dynamic evolutionary algorithm (EA) for optimizing drone routes, (2) a reinforcement learning algorithm for incorporating the feedback and runtime data about the system state, and (3) a prediction approach to predict the movement of drones and moving obstacles in the flying zone. To find efficient drone routes, we propose a parallel and dynamic implementation of the imperialistic competition algorithm (ICA) [2] that allows us to find efficient collision-free routes for the drones in the swarm. Our prediction approach uses exponential moving average to filter the monitoring data and linear regression to make predictions from the filtered data.

The learning component is based on the K -nearest neighbor (KNN) learning algorithm [3]. Each placement produced by our parallel ICA for a given swarm and environment state is evaluated to train the system. Both the learning and optimization algorithms work in parallel to compute alternative routing solutions. The results are compared and a more efficient solution is chosen. Since with each run the training set increases, eventually the learning algorithm becomes capable of proposing better solutions.

We believe that our proposed approach provides a promising solution to ensure efficient, collision-free navigation of the drones in a swarm. We also present a parallel implementation

¹https://www.faa.gov/uas/resources/uas_sightings_report/

of the proposed approach and evaluate it against two benchmarks. The results demonstrate that the proposed approach allows to significantly reduce the route lengths and computation overhead while producing efficient and safe routes.

The paper is structured as follows. Section II briefly reviews EAs and ICA. Section III presents the proposed approach that integrates learning, optimization, and prediction for generating efficient and safe drone routes. In Section IV, we describe our prediction approach in detail. Section V presents some important implementation details and experimental results. Finally in Section VI, we review important related works and present our conclusions.

II. IMPERIALISTIC COMPETITION ALGORITHM

Evolutionary computing comprises a set of optimization algorithms, which are inspired by a biological or societal evolution [4]. The evolutionary algorithms (EAs) are widely used in the swarm systems due to their ability to find, in a highly performant way, near-optimal solutions for the computationally hard problems. An EA mimics the survival of the fittest principle of the nature.

There is a large variety of EAs. Some of them are inspired by natural phenomena, while others, such as Imperialist Competitive Algorithm (ICA) [2], by the social processes. The algorithm simulates a human social evolution. Its parallel implementation [5] has shown a remarkable performance in comparison with the other EAs and offers a promising solution supporting compute-intensive tasks of swarm-based systems.

Figure 1 presents a flowchart highlighting the main steps in the ICA. The algorithm starts by a random generation of a set of countries - the genotypes - in the search space of the optimization problem. The fitness function determines the power of each country. The countries with the best values of the fitness function become *imperialists*, while the other countries become *colonies*. The colonies are divided among the imperialists and hence the overall search space is divided into empires. An association of a colony with an imperialist means that only the genotype of the imperialist and its associated colonies are used for crossover. The intuition behind it as follows: since the imperialist has a higher value of the fitness function, by crossing over with an associated colony, which is known to have a lower value of the fitness function, we inherit the strongest traits of the current population.

The mutation and crossover are implemented by *assimilation* and *revolution* operators. Assimilation moves colonies closer to an imperialist in its socio-political characteristics. For instance, it can be implemented by replacing a certain bit in a colony genotype with the corresponding bit of the imperialist. Revolution results in a drastic change of a colony's characteristics. It can be implemented by a random replacement of a certain bit in the colony genotype. As a result of assimilation and revolution, a colony might reach a better position and get a chance to take over the control of the entire empire, that is, to overthrow the current imperialist. This can happen only if the evaluation of the fitness function of such a colony gives a

higher value (when solving a maximization problem) than the value of the fitness function of the current imperialist.

The next step of the algorithm computes the power of each empire and implements the imperialistic competition, which corresponds to the selection of the survivals process. The power of an empire is computed by aggregating the fitness value of the imperialist and a weighted sum of the fitness values of the colonies. The imperialists also try to take possession of colonies of other empires, that is, the weakest empire loses its weakest colony. In each step of the algorithm, based on their power, all the empires get a chance to take control of one or more of the colonies of the weakest empire. The steps of the algorithm are repeated until a termination condition is reached. As a result, the imperialist of the strongest empire produces the best solution. To improve performance of ICA, we introduce the notion of multi-population, that is, we divide the overall search space into multiple populations or clusters and perform a local search within each one of them. The best local solutions are then taken as input to perform the search in the entire search space. The multi-population based search also allows to use the inter-population *migration* operation, which migrates the best country from one population and uses it to replace the worst country in another population. Since the local search procedures are independent of each other, they can be implemented in parallel. Moreover, the multi-population based search enables a wider exploration of the search space, which helps to find high quality solutions.

III. THE PROPOSED APPROACH

A swarm of drones is a typical example of a complex distributed networked system [6]. Each drone can be seen as a mobile sensing node that is capable of collecting monitoring data and communicating with some other drones in the swarm as well as with the cloud-based navigation center. Finding an efficient route for each drone in the swarm to ensure motion safety is a complex optimization problem. Therefore, we need to rely on certain heuristics to achieve the required objectives. In this paper, we propose a dynamic EA to compute the drone routes [7]. The proposed algorithm is based on the imperialistic competition algorithm (ICA) [2].

Figure 2 presents an overview of the proposed approach called DIANA (Dynamic Intelligent Autonomous Navigation Algorithm). The *Offline Part* in the figure uses our proposed ICA-based route generation algorithm to generate drone routes before the start of the mission. Moreover, it computes and uses the shortest paths between the start and destination locations of the drones. Since a drone swarm is a highly dynamic system, we augment our offline module with an *online* approach that provides runtime means for monitoring and reconfiguration. The information obtained from the *Dynamic Monitoring* component has two main purposes. On one hand, it is a feedback mechanism. On the other hand, it allows us to detect the changes in the drone swarm and in the flying zone. Such changes may invoke swarm reconfiguration and regeneration of the drone routes. In addition, the *Prediction* module uses the runtime monitoring data to predict the movement of drones

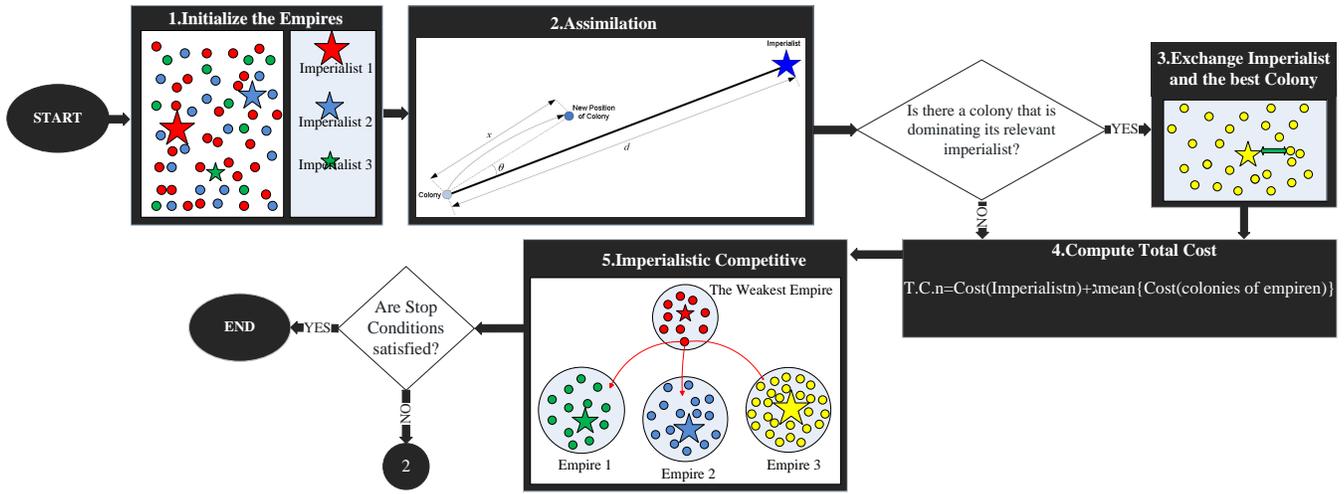


Fig. 1. A flowchart highlighting the main steps in the imperialistic competition algorithm

and moving obstacles in the flying zone. We feed the monitoring data and prediction results to both a *Dynamic EA* and a *Learning Algorithm*. Both modules compute the alternative routes. The routes are compared by the *Decision Center* that chooses the best solutions from the proposed alternatives. The *Navigation Center* then issues the corresponding commands to the drones.

The proposed DIANA approach uses the parallel implementation of the ICA [5] with an integrated migration operation (referred henceforth as the MICAP algorithm). The MICAP algorithm computes the initial drone routes as well as the subsequent new routes during the execution of the mission. For the learning module, we use a *K*-nearest neighbor (KNN)-based learning algorithm [3] that runs alongside the MICAP algorithm.

The main steps of the DIANA approach are presented in Algorithm 1. The algorithm starts with the preplanning of

the mission - the offline execution of MICAP to find the initial routes. It then obtains the current actual and predicted state of the system by invoking the dynamic monitoring and prediction modules and runs the MICAP algorithm and the KNN-based learning algorithm until the mission completes. In each iteration, it compares the results of MICAP with that of the KNN-based learning algorithm and selects the best solutions. The KNN algorithm is a popular learning method. In this paper, we use it for classification. The main idea of the algorithm is as follows. We select *K* samples in the training set. The algorithm predicts the numerical target - the nearest neighbor - based on a similarity measure, which in our case is a distance function. We compute the numerical target as the average of the Euclidian distances of the *K* nearest neighbors.

In our approach, the algorithm takes as the input *K* solutions generated for the same system state. Then it computes an average solution. For each drone in the swarm, it takes *K* routes proposed for it. Then it computes an average route in terms of the Euclidian distance. By computing such a route for each drone in the swarm, the learning component calculates the complete solution for the swarm. The decision center uses a fitness function to compute fitness values for the solutions produced by MICAP and the learning algorithm. It then chooses the solution with the highest fitness value.

IV. PREDICTION APPROACH

Our prediction approach is based on the two-step approach of Andreolini and Casolari [8] and Andreolini et al. [9]. The two-step approach allows to make predictions under real-time constraints. It is based on the rationale that periodic sampling of data offers an instantaneous view of the trends. However, raw data are of little help for distinguishing different types of trends in the data. The direct use of the measured raw data does not solve the problem, because raw data can be highly variable. Prediction based on the monitored raw data can be risky and inconvenient. Thus, it is preferable to operate on a representation of the behavior of the system. The approach

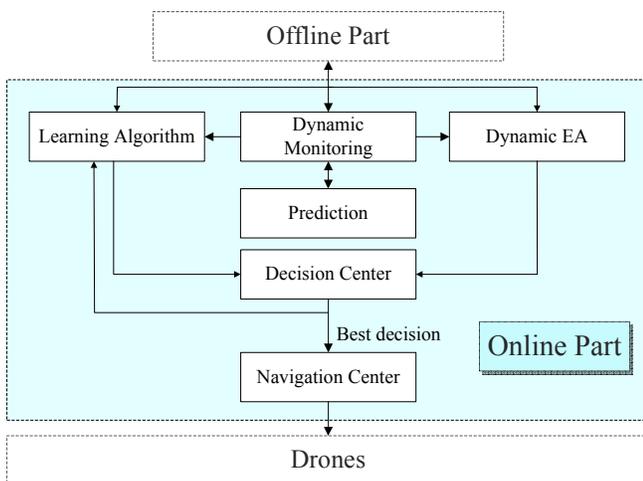


Fig. 2. Overview of the proposed approach

Algorithm 1 DIANA

```
1: {Offline part}
2: Best-Placement ← Call MICAP(Initial-State)
3: Send(Best-Placement) → Navigation Center
4: {Online part}
5: while Mission is in progress do
6:   Current-State ← Call Dynamic Monitoring and
     Prediction modules
7:   Best-EC-Result ← Call MICAP(Current-State)
8:   Best-KNN-Result ← Call KNN(Current-State)
9:   Best-Placement ← Max(Best-EC-Result,
     Best-KNN-Result)
10:  Send(Best-Placement) → Navigation Center
11:  Send(Best-Placement) → KNN-Dataset
12: end while
```

Procedure 1 MICAP(Current-State)

```
1: if Processor  $P_i$  then
2:   for j=1 to Population Size do
3:     Initial-Population [j] ←  $C_{country_j}$ 
4:     Cost_Initial-Population
5:     [j] ←  $\alpha \cdot \text{Covering Quality} - \frac{\beta}{\text{CommunicationCost}}$ 
6:   end for
7:   Initial-Population ← Sort Initial-Population
8:   for j=1 to #Empire do
9:     Imperialist[j] ← Initial-Population [j]
10:  end for
11:  for j=#Empire+1 to Population Size do
12:    Assigned as a Colony to an Empire
13:  end for
14:  for All Colony do
15:    Assimilate Colony → Imperialist
16:  end for
17:  for All Colony do
18:    if Colony  $\geq$  Imperialist then
19:      Colony  $\leftrightarrow$  Imperialist
20:    end if
21:  end for
22: end if
```

Procedure 2 KNN(Current-State)

```
1: for All Instance  $\in$  Dataset do
2:   Find Nearest-Neighbor
3: end for
4: return Nearest-Neighbor
```

Procedure 3 Prediction

```
1: Invoke two parallel instances of the Tracker
2: {see Section IV}
3: Invoke two parallel instances of the Predictor
4: {see Section IV}
5: return prediction results
```

involves *trackers* that may offer a representative view of the

trends to the *predictors*, thus achieving the two-step approach.

A tracker filters out the noise and can be used to yield a more regular view of the trend of an obstacle or drone movement. It takes as input a raw measure s_i monitored at time t_i , and a set of previously collected n measures, that is $\vec{S}_n(t_i) = (s_{i-n}, \dots, s_i)$, and outputs a representation of the trends l_i at time t_i . Formally, tracker is a function $Tracker(\vec{S}_n(t_i)) : \mathbb{R}^n \rightarrow \mathbb{R}$. Multiple applications of tracker provides a sequence of values that yield a regular trend of the movement. There are different classes of linear and non-linear trackers, such as simple moving average (SMA), exponential moving average (EMA), and cubic spline (CS) [9]. SMA is a simple linear method, but has known shortcomings of increased oscillations when n is small and of significant delay when n is large. CS is a non-linear method that is more expensive to compute than SMA and EMA, but instead of returning a new tracker value for each raw measure, it returns a new tracker value after n measures. More sophisticated (time-series) models often require training periods to compute the parameters and/or off-line analyses. Similarly, the linear (auto) regressive models such as ARMA and ARIMA, usually require frequent updates to their parameters in the case of highly variable systems [8]. Therefore, the proposed approach implements a tracker based on the EMA model, which limits the delay without incurring oscillations and computes a tracker value for each measure.

EMA is the weighted mean of the n measures in the vector $\vec{S}_n(t_i)$, computed at time t_i , where $i > n$, and the weights decrease exponentially. An EMA based tracker is defined as

$$EMA(\vec{S}_n(t_i)) = \alpha \cdot s_i + (1 - \alpha) \cdot EMA(\vec{S}_n(t_{i-1}))$$

where $\alpha = \frac{2}{n+1}$. The initial value $EMA(\vec{S}_n(t_n))$ is set to the arithmetic mean of the first n values

$$EMA(\vec{S}_n(t_n)) = \frac{\sum_{j=0}^n s_j}{n}$$

The predictor takes as input a set of tracker values $\vec{L}_q(t_i) = l_{i-q}, \dots, l_i$ and outputs a future tracker value at time t_{i+k} , where $k > 0$. Formally, predictor is a function $Predictor_k(\vec{L}_q(t_i)) : \mathbb{R}^q \rightarrow \mathbb{R}$. With the use of the trackers that provide high correlation among values, even simple linear predictors are sufficient to predict the future trend of the movement. The predictor is characterized by the prediction window k and the past time window q . Using a simple linear regression model [10], the predictor uses the last q tracker values $\vec{L}_q(t_i)$ [11]. It is based on a straight line defined as

$$l = \Theta_0 + \Theta_1 \cdot t$$

where Θ_0 and Θ_1 are unknown constants, called regression coefficients, which can be estimated at runtime based on the tracker values $\vec{L}_q(t_i)$ in the past time window. One common approach to estimate these regression coefficients is to use

the least-square estimation method [10]. The least-square estimators of Θ_0 and Θ_1 , say $\hat{\Theta}_0$ and $\hat{\Theta}_1$, are computed as

$$\hat{\Theta}_0 = \bar{l} - \hat{\Theta}_1 \cdot \bar{t}$$

and

$$\hat{\Theta}_1 = \frac{\sum_{j=i-q}^i (l_j \cdot t_j) - \frac{\left(\sum_{j=i-q}^i l_j\right) \cdot \left(\sum_{j=i-q}^i t_j\right)}{q}}{\sum_{j=i-q}^i t_j^2 - \frac{\left(\sum_{j=i-q}^i t_j\right)^2}{q}}$$

where

$$\bar{l} = \frac{1}{q} \sum_{j=i-q}^i l_j \quad \text{and} \quad \bar{t} = \frac{1}{q} \sum_{j=i-q}^i t_j$$

The predictor returns a predicted future tracker value \hat{l}_{i+k} that corresponds to time t_{i+k} . It is computed as follows:

$$Predictor_k(\vec{L}_q(t_i)) = \hat{l}_{i+k} = \hat{\Theta}_0 + \hat{\Theta}_1 \cdot t_{i+k}$$

Prediction results depend upon selection of proper values for the tracker and predictor parameters. Therefore, it is necessary to find a value for n that represents a good tradeoff between a reduced delay and a reduced degree of oscillations [8]. Similarly, the values of q and k should be selected carefully.

In a two-dimensional flying zone, the current location of an obstacle or a drone monitored at time t_i contains two values (x, y) , which represent the horizontal and the vertical axis, respectively. Therefore, for predicting the future location of an obstacle or a drone at time t_{i+k} , we use two parallel trackers and two parallel predictors. One pair of tracker and predictor works with the x values and the other pair works with the y values.

V. IMPLEMENTATION AND EXPERIMENTAL RESULTS

In this section, we present some important implementation details and experimental results. We assume that the flying zone *AREA* is represented by a grid as shown in Figure 3(a). Our goal is to find an efficient, collision-free route for each drone from the initial start location of the drone to the destination location.

We use the example shown in Figure 3(b) to explain the principles used for defining the countries in the MICAP algorithm. For the drone $d1$, initially situated at location 20, the shortest path from the initial location to the destination is a sequence $\langle 20, 19, 18, 17, 16, 11, 6 \rangle$. We note that the path of each drone can be succinctly represented by a *turning point* – we call it *middle point*, which would be 16 for $d1$. Therefore, the proposed approach uses the middle points to optimize the drone routes. By using different values for a drone's middle point, it is possible to generate different routes for the drone. Thus, it allows to explore different alternatives in the search

space. Lets assume that the middle points for the two other drones in Figure 3(b) are 12 and 9. A country representing all drone routes for the swarm in Figure 3 can then be defined as a triple $\langle \langle 16, 12, 9 \rangle \rangle$. In general, for nd drones a country is an nd -tuple consisting of the middle points of the corresponding drones.

In the offline, planned part of the proposed approach, all locations occupied by static obstacles and the initial locations of the moving obstacles are marked as occupied or unsafe. The offline route planning module then generates all shortest paths between each pair of locations in the flying zone *AREA* while avoiding all locations marked as occupied or unsafe and stores the generated routes in a database, which is then used to compose the routes for the individual drones as a concatenation of the shortest routes from initial locations to the middle points and from the middle points to the final destinations. The shortest routes are computed using the algorithm proposed by Dijkstra [12].

The fitness function to evaluate the fitness of each country optimizes the safety/performance ratio. The first argument of our fitness function is the distance metric

$$Distance\ Metric = \sum_{i=1}^{nd} Distance_{Current_i \rightarrow Middle_i} + Distance_{Middle_i \rightarrow Destination_i}$$

It defines the total length of the drone routes according to the given solution. For our example in Figure 3(b) the distance metric of the routing defined by the country $\langle 16, 12, 9 \rangle$ is the sum of the lengths of the drone paths: $6+6+6=18$. The second argument of the fitness function defines the number of cross points associated with the given solution. For our example in Figure 3(b), the number of cross points is 3: in location 17 between the route 1 and 2, in location 12 between the route 2 and 3, and in the location 11 between route 1 and 3, correspondingly.

The third argument is the safety level of the time gap at the cross point. We introduce three safety levels: 0 if there is no cross points, 1 if the time gap at the cross point is above the safety threshold, and 2 if the time gap is below the threshold. For our example in Figure 3(b), the time gap at cross point 17 is 1, because the drones arrive at that point at times 3 and 2, the time gap for the cross point 12 is 2, because the drones arrive there at times 3 and 1, and for the cross point 11, it is 5. As a matter of illustration, we can assume that the time gaps below threshold 2 are classified as level 2, while the time gaps at and above threshold 2 as level 1. Hence, the cross point 17 obtains level 2, while the cross points 11 and 12 obtain level 1 each. The safety level of a complete routing solution for the swarm can then be computed by aggregating the individual safety levels of all cross points in the solution: $2+1+1=4$. We define our route optimization task as a minimization problem with the following fitness function:

$$Fitness\ Function = Distance\ Metric + \alpha \cdot Number\ of\ CrossPoint + \beta \cdot Level$$

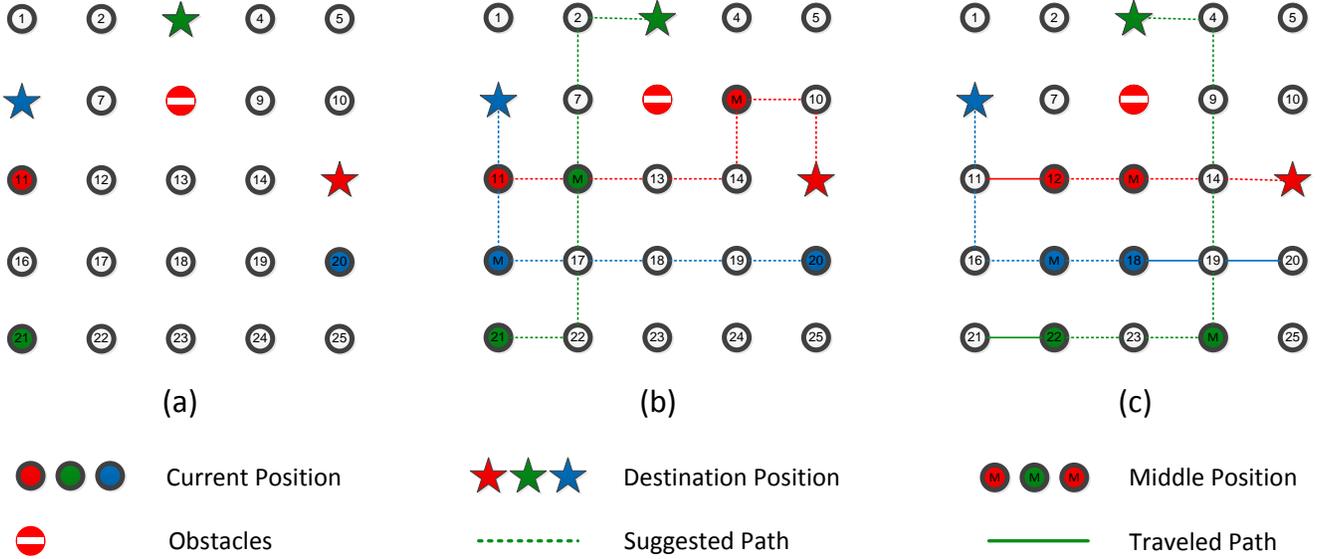


Fig. 3. An illustrative example of drone route planning

Here α and β are the weight coefficients defined as

$$1 \leq \alpha \leq \frac{nd}{2} \quad 1 \leq \beta \leq \sqrt{np} \times nd$$

where nd is the number of drones and np is the total number of points. These values allow us to adapt the fitness function evaluation based on the level of complexity of the flying zone and the number of drones. For our example in Figure 3(b), the value of the fitness function is computed as $18 + 1.5 \times 3 + 5 \times 4 = 42.5$.

A. Experiment Design and Setup

We have implemented the proposed DIANA approach on a shared memory model. We used the message passing interface (MPI) to parallelize the proposed algorithm and MPICH² to run the algorithm. To implement DIANA, we used four processors in a ring topology. Our algorithm was tested on Intel[®] Xeon[®] E5-1620 v3 @ 3.50 GHz processors with 16 GB memory and NVIDIA[®] GeForce[®] GTX 1080 graphics processing units.

We present results from two benchmark implementations. Both benchmarks are based on a 50×50 flying zone with 18 static obstacles, 1 moving obstacle, and 8 drones. We also compare the results with two alternative approaches:

- A well-known approach called Dynamic Genetic Algorithm (DGA) [13], which addresses a similar problem.
- A baseline approach called DANA (Dynamic Autonomous Navigation Algorithm), which is similar to DIANA except that it does not use learning and prediction.

B. Results and Analysis

Figure 4 and 6 present the actual and predicted movement of the moving obstacle in Benchmark 1 and 2, respectively.

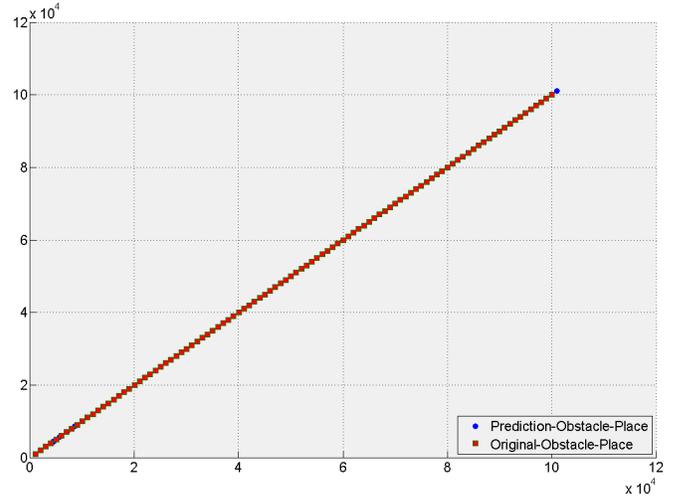


Fig. 4. Benchmark 1: actual and predicted movement of the moving obstacle

In Benchmark 1, the moving obstacle moved on a straight line while in Benchmark 2 the obstacle moved randomly. The results show that as soon as there were enough data points to make good predictions, our prediction approach started to provide precise estimates of the obstacle movement with a low prediction error. The prediction results were used by our KNN-based learning algorithm and our MICAP algorithm to generate new routes for the drones in an online manner.

Figure 5 and 7 present the flying zones from Benchmark 1 and 2, respectively. They depict the static obstacles, the actual movement of the moving obstacle, and the drone routes. The results show that all drones successfully completed their maneuvers while avoiding collisions with the static and moving obstacles and with the other drones in the flying zone.

Table I presents a comparison of the results of DIANA,

²<https://www.mpich.org/>

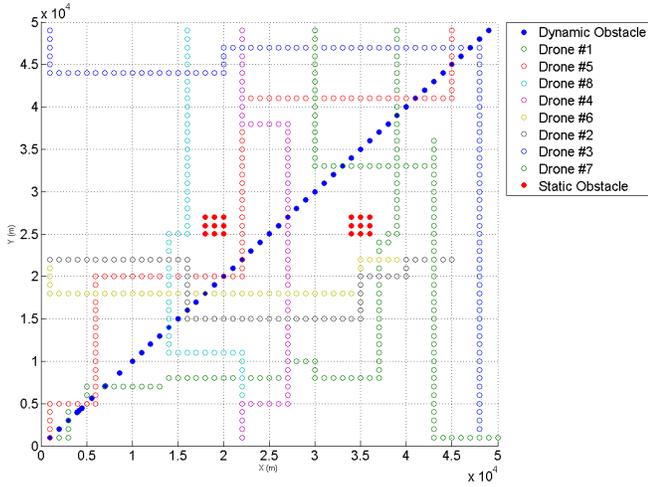


Fig. 5. A snapshot of the flying zone from Benchmark 1

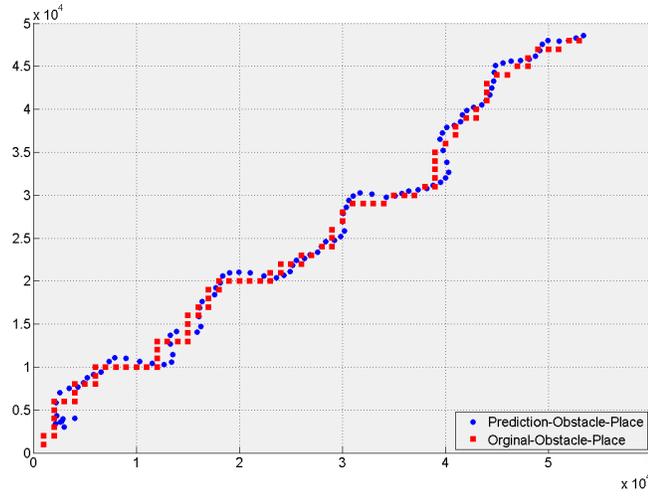


Fig. 6. Benchmark 2: actual and predicted movement of the moving obstacle

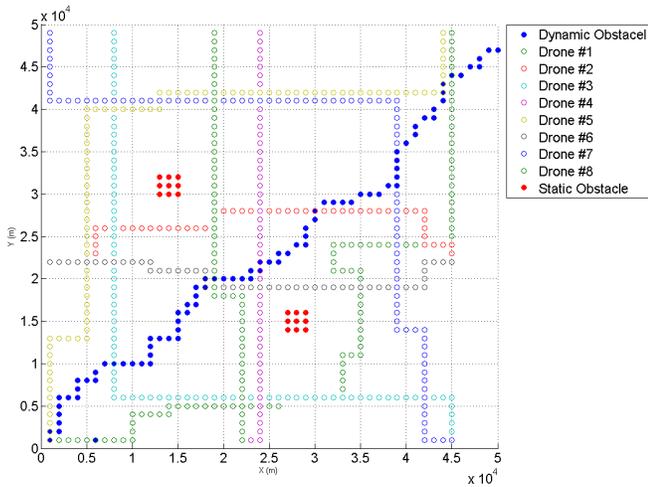


Fig. 7. A snapshot of the flying zone from Benchmark 2

DGA, and DANA. The comparison is based on three metrics:

- 1) *Route length*: the length of a drone route measured as the number of steps in the drone route. To be minimized to generate shorter routes.
- 2) *Minimum distance*: the minimum distance between a drone and a moving obstacle. To be maximized to generate safer routes.
- 3) *Frequency of route regeneration*: the number of times the drone routes are regenerated. To be minimized for reducing the re-computation overhead.

The results show that DIANA outperformed DGA and DANA in all three metrics. It minimized the route lengths and the frequency of route regeneration more efficiently than DGA and DANA. Moreover, it generated drone routes with a higher minimum distance. In Benchmark 1, DIANA produced 10.13% and 13.03% shorter routes than DANA and DGA, respectively. Similarly, in Benchmark 2, it generated 17.65% and 22.23% shorter routes than DANA and DGA, respectively. Therefore, DIANA generated safer and shorter drone routes while minimizing the re-computation overhead.

VI. RELATED WORK AND CONCLUSIONS

The problem of motion safety of semi-autonomous robotic systems is currently attracting significant research attention. A comprehensive overview of the problems associated with autonomous mobile robots is given in [14]. The analysis carried out in [15] shows that the most prominent routing schemes do not guarantee motion safety. Our approach resolves this issue and ensures not only safety but also provides efficient online routing.

Macek et al. [16] proposed a layered architectural solution for robot navigation. However, in their work, they focused on the safety issues associated with the navigation of a single vehicle and did not consider the problem of collision-free navigation in the context of swarms of robots. Aniculaesei et al. [17] presented a formal approach that employs formal verification to ensure motion safety. Petti and Fraichard [18] proposed an approach that relies on partial motion planning to ensure safety. Their solution supports navigation of a single vehicle. In our work, we have discretized the flying zone and have developed a highly efficient approach that computes the next safe states for an entire swarm and provides a mechanism for online route regeneration and collision avoidance.

Olivieri and Endler [19] presented an approach for movement coordination of swarms of drones using smart phones and mobile communication networks. Their work focuses on the internal communication of the swarm and does not provide a solution for collision-free route generation. Barry and Tedrake [20] proposed an obstacle detection algorithm for drones that allows to detect and avoid collisions in realtime. Similarly, Lin [21] presented a realtime path planner for drones that detects and avoids moving obstacles. These approaches are only applicable for individual drones and they do not provide support for a swarm of drones. In our work, we focused on collision prediction and avoidance and efficient navigation of swarms of drones.

TABLE I
A COMPARISON OF THE PROPOSED APPROACH WITH TWO ALTERNATIVE APPROACHES

	Benchmark 1			Benchmark 2		
	DIANA	DANA	DGA	DIANA	DANA	DGA
Route length	648	721	745	616	784	792
Minimum distance	5	3	2	4	2	2
Frequency of route regeneration	17	28	39	21	38	37

A comprehensive literature review on motion planning algorithms for drones can be found in [22]. The approaches reviewed in [22] are applicable to a preliminary, offline motion planning phase to plan and produce an efficient path or trajectory for a drone before the start of the mission. A more recent survey on motion planning of drones can be found in [23]. Augugliaro et al. [24] also presented a planned approach for generating collision-free trajectories for a drone fleet. In contrast to these approaches, our proposed approach combines offline motion planning with a more realistic online route generation approach to produce efficient collision-free routes.

In this paper, we presented a novel approach that integrates learning, optimization, and prediction for generating efficient and safe routes for swarms of drones. The proposed approach comprises three main components: (1) a high-performance dynamic evolutionary algorithm for optimizing drone routes, (2) a reinforcement learning algorithm for incorporating the feedback and runtime data about the system state, and (3) a prediction approach to predict the movement of drones and moving obstacles in the flying zone. We also presented a parallel implementation of the proposed approach and evaluated it against two benchmarks. The results demonstrated that the proposed approach significantly reduces the route lengths and computation overhead while producing efficient and safe routes.

ACKNOWLEDGMENT

The work was supported by the Academy of Finland projects OpenCPS: Open Integrated Framework for Accelerating Development of Resilient CPS and CoRA: Continuous Resilience Assurance of Complex Software-Intensive Systems.

REFERENCES

[1] R. Austin, *Unmanned Aircraft Systems: UAVS Design, Development and Deployment*, ser. Aerospace Series. Wiley, 2010.

[2] E. Atashpaz-Gargari and C. Lucas, "Imperialist competitive algorithm: An algorithm for optimization inspired by imperialistic competition," in *IEEE Congress on Evolutionary Computation*, 2007, pp. 4661–4667.

[3] M. Mejdoub and C. Ben Amar, "Classification improvement of local feature vectors over the KNN algorithm," *Multimedia Tools and Applications*, vol. 64, no. 1, pp. 197–218, 2013.

[4] C. Guestrin, A. Krause, and A. P. Singh, "Near-optimal sensor placements in gaussian processes," in *Proceedings of the 22nd International Conference on Machine Learning*, ser. ICML '05. ACM, 2005, pp. 265–272.

[5] K. Kar and S. Banerjee, "Node placement for connected coverage in sensor networks," *WiOpt'03: Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, 2003. [Online]. Available: <https://hal.inria.fr/inria-00466114>

[6] N. Heo and P. K. Varshney, "Energy-efficient deployment of intelligent mobile sensor networks," *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, vol. 35, no. 1, pp. 78–92, 2005.

[7] A. Majd, G. Sahebi, M. Daneshdalan, J. Plosila, and H. Tenhunen, "Placement of smart mobile access points in wireless sensor networks and cyber-physical systems using fog computing," in *IEEE International Conference on Scalable Computing and Communications (ScalCom)*, July 2016, pp. 680–689.

[8] M. Andreolini and S. Casolari, "Load prediction models in web-based systems," in *Proceedings of the 1st International Conference on Performance Evaluation Methodologies and Tools*. ACM, 2006.

[9] M. Andreolini, S. Casolari, and M. Colajanni, "Models and framework for supporting runtime decisions in web-based systems," *ACM Transactions on the Web*, vol. 2, no. 3, pp. 17:1–17:43, 2008.

[10] D. Montgomery, E. Peck, and G. Vining, *Introduction to Linear Regression Analysis*, ser. Wiley Series in Probability and Statistics. John Wiley & Sons, 2012.

[11] A. Ashraf, B. Byholm, and I. Porres, "A session-based adaptive admission control approach for virtualized application servers," in *2012 IEEE Fifth International Conference on Utility and Cloud Computing*, 2012, pp. 65–72.

[12] E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische Mathematik*, vol. 1, no. 1, pp. 269–271, 1959.

[13] S. Indhumathi and D. Venkatesan, "Improving coverage deployment for dynamic nodes using genetic algorithm in wireless sensor networks," *Indian Journal of Science and Technology*, vol. 8, no. 16, 2015.

[14] R. Siegwart, I. R. Nourbakhsh, and D. Scaramuzza, *Introduction to Autonomous Mobile Robots*, ser. Intelligent robotics and autonomous agents. MIT Press, 2011.

[15] T. Fraichard, "A short paper about motion safety," in *Proceedings 2007 IEEE International Conference on Robotics and Automation*, 2007, pp. 1140–1145.

[16] K. Macek, D. Vasquez, T. Fraichard, and R. Siegwart, "Safe vehicle navigation in dynamic urban scenarios," in *2008 11th International IEEE Conference on Intelligent Transportation Systems*, 2008, pp. 482–489.

[17] A. Aniculaesei, D. Arnsberger, F. Howar, and A. Rausch, "Towards the verification of safety-critical autonomous systems in dynamic environments," in *Proceedings of the The First Workshop on Verification and Validation of Cyber-Physical Systems*, 2016, pp. 79–90.

[18] S. Petti and T. Fraichard, "Partial motion planning framework for reactive planning within dynamic environments," in *Proceedings of the IFAC/AAAI International Conference on Informatics in Control, Automation and Robotics*, 2005.

[19] B. J. O. de Souza and M. Endler, "Coordinating movement within swarms of UAVs through mobile networks," in *2015 IEEE International Conference on Pervasive Computing and Communication Workshops (PerCom Workshops)*, 2015, pp. 154–159.

[20] A. J. Barry and R. Tedrake, "Pushbroom stereo for high-speed navigation in cluttered environments," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 3046–3052.

[21] Y. Lin, "Moving obstacle avoidance for unmanned aerial vehicles," Ph.D. dissertation, Arizona State University, 2015.

[22] C. Goerzen, Z. Kong, and B. Mettler, "A survey of motion planning algorithms from the perspective of autonomous UAV guidance," *Journal of Intelligent and Robotic Systems*, vol. 57, no. 1, p. 65, 2009.

[23] F. Kendoul, "Survey of advances in guidance, navigation, and control of unmanned rotorcraft systems," *Journal of Field Robotics*, vol. 29, no. 2, pp. 315–378, 2012.

[24] F. Augugliaro, A. P. Schoellig, and R. D'Andrea, "Generation of collision-free trajectories for a quadcopter fleet: A sequential convex programming approach," *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1917–1922, 2012.