

Solution to the Fixed Airbase Problem for Autonomous URAV Site Visitation Sequencing

Amit Agarwal¹, Meng-Hiot Lim², Chan Yee Chew³, Tong Kiang Poo⁴,
Meng Joo Er¹, and Yew Kong Leong¹

¹ Intelligent Systems Center, Techno Plaza, Nanyang Technological University
Singapore 639798

{pg0212208t, emjer}@ntu.edu.sg, leongy.dynamics@stengg.com

² School of Electrical & Electronic Engineering, Nanyang Technological University
Singapore 639798

{emhlim, 810304085117, S7931481A}@ntu.edu.sg

Abstract. In [1], we modeled the problem of safe site visitation sequencing for an unmanned reconnaissance aerial vehicle (URAV) in graph-theoretic language. Our goal was to uncover several alternative flight paths along which the URAV can fly safely, no more than once, to all sites of interest. We offered a solution to a simpler abstraction of this problem using GAs – if we could not uncover any Hamiltonian cycles we reported maximally long paths between any two vertices. This work reports two orthogonal extensions of [1] – a. In addition to possibly reporting Hamiltonian cycles, the GA reports several alternative long paths that begin at the airbase. b. A randomized version of the gene migration scheme is used. The randomization does not appear to affect the quality of the best plan. However it raises the average quality of the best k plans. We give a modified gene encoding scheme and report simulation results.

1 Background

In this work, and in [1], our focus is on planning the point-to-point route of a single unmanned reconnaissance aerial vehicle (URAV) from the airbase, A , to possibly each nonoverlapping useful vicinity, $s_i \in S$ of site(s) of interest and back. The presence of $O(|S|)$ threats, T_i in the reconnaissance zone is assumed. Without loss of generality, it is also assumed that these threats are evenly distributed in the region of reconnaissance interest. As discussed in [1], the structure that represents the union of minimal risk density flight paths and S is often a multigraph with $O(|S|)$ edges. Several graphs may be defined on this multigraph, each being optimal in a certain local sense. In [1], for the lack of a better term, any such graph was referred to as an approximated Voronoi graph ($A-Vor(G)$). However, in order to avoid contributing to the rather misleading notion that this graph is an approximation of a Voronoi graph [2, 3] in any sense, we shall call it *minimal risk density target visitation* (MRDTV) graph.

Since MRDTV graphs are often sparse, suboptimal site visitation sequencing will cause the URAV to revisit a useful vicinity in order to continue flying to other useful

vicinities, strictly along minimal risk density paths. Revisitations can lead to an increase in the total time for mission completion, expose the URUV to threats for a longer duration than absolutely necessary, take away the element of surprise and may lead to decreased useful range of the URUV due to unnecessary fuel consumption. Thus the following problem, whose mathematical structure is akin to that of the NP-complete Hamiltonian cycle problem [4], gains importance –

Given an airbase, A , and a set of nonoverlapping useful vicinities of reconnaissance sites, \mathcal{S} , find orderings $\prod_{\mathcal{S}}^*$ such that a URUV can take-off from A , fly to the useful vicinities of each site, $s_i \in \mathcal{S}$, along a subset of the edges of an MRDTV graph exactly once in the order defined by $\prod_{\mathcal{S}}^* \in \prod_{\mathcal{S}}^*$ and, fly back safely to A .

If the MRDTV graph under consideration is non-Hamiltonian or, alternately, sufficient number of alternative solutions in $\prod_{\mathcal{S}}^*$ are difficult to uncover then solutions that give sufficiently long acyclic paths beginning at A can still help greatly in lowering the risk to the URUV.

We use genetic algorithms to solve this problem. Even while genetic algorithms have not been particularly popular for use on TSP related problems on sparse graphs, we found them attractive since their intrinsic parallelism is naturally suited to our need for several alternative site visitation sequences. The interested reader may look at [5, 6, 7, 8, 9] for surveys or discussions on techniques for solving the Hamiltonian cycle problem.

In [1], our algorithm suffered from one major weakness: in case several different Hamiltonian cycles matching the number of plans specified by the user were not found (or, did not exist), our hybrid genetic algorithm (HGA) reported paths which though were maximally long but did not begin (or, end) in the airbase. Fig 1a shows a typical example of a long path reported. Its length is 18 (maximum possible length is 20). However, this visitation sequence plan does not begin at the airbase. And thus, its maximum useful length is only 12.

In this work we remedy this problem by changing the fitness function and making minor changes to the crossover and mutation operators [1]. Also, it was slightly more efficient to change the encoding scheme to restrict the first gene in any chromosome to corresponding to A . Fig 1b shows a typical example of the plans we now obtain – it originates at the airbase. In this example, the maximum useful length of the plan (17) and the maximal length are the same. The changes we have made to our earlier HGA do not however affect its ability to discover Hamiltonian cycles.

We had another, though minor, problem with our earlier HGA. At times, it was unable to uncover sufficient number of alternative plans that met the user's qualitative requirements even while they were known to exist. In this work, the next gene to be migrated (see Section 4) is chosen probabilistically (as opposed to sequential consideration in our former HGA) from the candidate set. This has, permitted us to obtain a greater number of alternative good plans.

The remaining paper is organized as follows. The larger focus of our work is on enabling a greater level of autonomy in remotely piloted vehicles. We briefly discuss in Section 2, the benefits of enabling intelligent onboard decision making capabilities

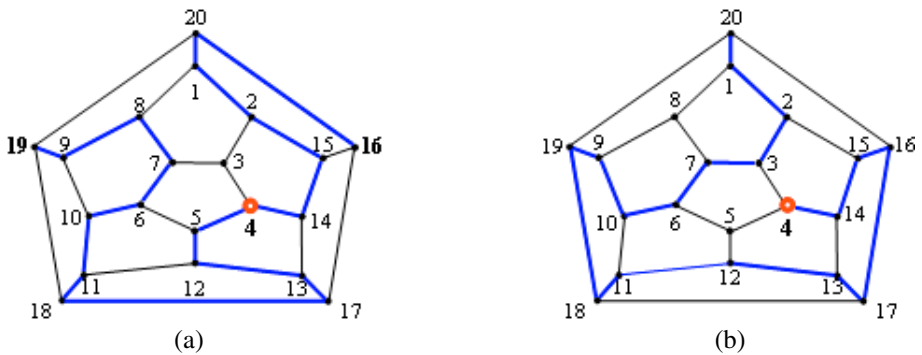


Fig. 1. An MRDTV graph with 19 useful vicinities of reconnaissance sites and an airbase (labeled 4). Thick blue lines represent a visitation plan: (a) represents a solution typically uncovered using our earlier HGA technique in [1] and (b) represents a solution typically uncovered using the technique described in this paper

in UAVs. In Section 3 we briefly discuss how an MRDTV graph may be generated. In Section 4 we discuss our HGA. We report the results in Section 6 and conclude in the following section.

2 Unmanned to Autonomous

UAVs can be loosely classified into two categories – *a.* remotely piloted *b.* autonomous. Remotely piloted vehicles still carry out tasks such as information acquisition and transmission, takeoff and landing, monitoring of onboard systems and complex flight control from outside the loop of the remote operator. However, several other critical tasks such as terrain masking, basic flight maneuvers, ammunition or sensor selection, target and role selection, decisions regarding cooperation, information processing, waypoint path planning and visitation sequencing are carried out with the help of varying degrees of remote human assistance.

Here we briefly discuss some basic motivations for the interest in going further, from unmanned, to autonomous. Some of our arguments possibly extend to other unmanned systems as well.

Bandwidth is an important resource in a digitized operation theatre [10, 11]. Due to the need for reconnaissance information in near real-time, a URUV must transmit images or videos of objects of interest while in flight. In the former case, the demand on the downlink bandwidth can be lowered by transmitting at a lower data-rate. The later case may not permit such a convenience since the video data could quickly saturate the onboard memory buffer. However, in both cases, transmissions are necessitated only by the capture of potentially useful information and thus, are intermittent and permit the use of asynchronous time division multiplexing to improve bandwidth efficiency. This bandwidth requirement can be further lowered by only transmitting extracted information upon intelligent onboard processing.

Operator-at-the-joystick requires continuous, real-time reception of high definition video at the control station. In certain scenarios, especially those involving low flying

vehicles, operations in civilian airspace or the presence of hostile units the demand on situation awareness can be very high. This can greatly increase the remote operator's workload and require a downlink bandwidth ranging from a few tens of MHz to about 10 GHz [12]. Energy and time efficient onboard algorithms for terrain masking, obstacle recognition and avoidance and, map updating can thus greatly reduce the overall demand on the downlink bandwidth.

Recently Veverka and Campbell [13] report some experiments simulated to study the impact of complexity of operation, number of UAVs, desired speed of response and the nature and time duration of tasks on operator performance. Their results suggest that appropriate automation of vehicular activities help in increasing operator performance in situations that involve complex decision making or attention to 'several' vehicles.

The last issue we discuss relating to the need for increased autonomy is cooperative waypoint path planning. Site visitation sequencing is a sub problem of this problem and the focus of interest of this work. Often, due to sensor errors, changes in wind conditions, non-availability of complete or high definition maps of the terrain, hostile units or sites of interest, replanning at various levels is necessary. Replanning can often be as difficult as planning and the problem of finding a good plan that meets the constraints is NP complete in general. Also, a remote station doing the job of replanning, specially based on inputs due to changes in the perception of the inflight vehicles themselves of their environment, can require the transmission of high definition digital data. Thus, centralized planning at the base station, especially for several UAVs, can create an avoidable computing bottleneck. Moreover, in the absence of onboard replanning capabilities, the UAV may need to pass sensitive information not available to the base station and thereby permitting the non benign units a greater opportunity for jamming or even eavesdropping. Beyond these considerations, two-way propagation delay, limits on maximum downlink signal strength, possible loss of contact for short intervals of time and other factors could limit the radius of operation of a UAV. This can be particularly limiting for an unmanned combat aerial vehicle (UCAV) in hot pursuit or an 'escaping' URUV.

A transition from manned to unmanned systems promises to liberate man from dangerous, time-consuming and repetitive tasks. Further movement, in the direction of onboard autonomy, is likely to relegate several complex tasks concerning the operation, task execution and cooperative planning to onboard decision support systems [14, 15, 16].

3 Generating MRDTV Graphs

A basic element in our solution methodology for obtaining several alternative site visitation plans for a URUV is a *minimal risk density target visitation*, or the MRDTV, graph. This data structure was introduced in [1] as $(A - Vor(G))$. It is a discreet connected subset of a suitably bounded 3D space over S . Under the assumptions stated in [1] and an expectation that only those hostile ground units will likely target the URUV that are closest to it, its edges represent the set of all possible minimal risk density flight paths for the URUV.

The algorithm for generating an MRDTV graph takes the following inputs –

- a. A – the airbase
- b. S – the set of all nonintersecting useful vicinities of reconnaissance sites that have been assigned to the URAV or acquired by the URAV.
- c. T – the set of threats

The algorithm is outlined as follows. These steps yield an MRDTV multigraph. Further processing using information related to user-requirements and vehicular limitations generates an MRDTV graph. Details regarding this step are not included here.

STEP 1: Generate the Voronoi graph on the set T using the algorithm outlined in [17] or any other optimal algorithm. Elements in T are assumed, without loss of generality, to be points.

STEP 2: Suitably limit the Voronoi graph using a sufficiently large bounding box.

STEP 3: Compute a minimal risk density multigraph as shown in [1].

Our HGA works with an MRDTV graph as input. Specifically, the vertices of the graph bijectively map to the genes in a chromosome. Any two genes in a particular chromosome are said to be *connected* iff in the corresponding MRDTV graph, there exists an edge between the useful vicinities identified by those genes. The HGA uncovers several alternative site visitation sequences of maximum or maximal length (Hamiltonian cycles or acyclic paths) in the MRDTV graph. In case the maximal length paths are acyclic, they are guaranteed to begin at the airbase.

4 Genetic Algorithm

Often in applications of genetic algorithms to vehicle routing problems, researchers do not explicitly consider the issue of starting point of the route [18, 19] since either a complete graph representation is assumed or the vehicles have the possibility of beginning its route at any location. In our application, the URAV must takeoff from a particular site; its airbase. Also, MRDTV graphs are not complete (if fact, they are usually, $O(|S| + |T|)$) and thereby, the Hamiltonian cycles are not a given or trivial to uncover. Thus, in case the HGA is not able to uncover sufficient number of different visitation sequences for the URAV that form a Hamiltonian cycle in the MRDTV graph, it must report the best maximally long paths that originate at the airbase.

In this section, in order to keep this paper as self-contained as possible, we reproduce elements of our HGA from [1] which remain the same in this paper. In addition, we report changes to the HGA which enable us to take into account the problem of starting from a fixed airbase and address the minor problem of generating sufficient number of good alternative plans. Though we report our simulations in the single URAV case, our algorithm can be used in certain cases involving multiple URAVs. The flowchart of the HGA is given in Fig 2.

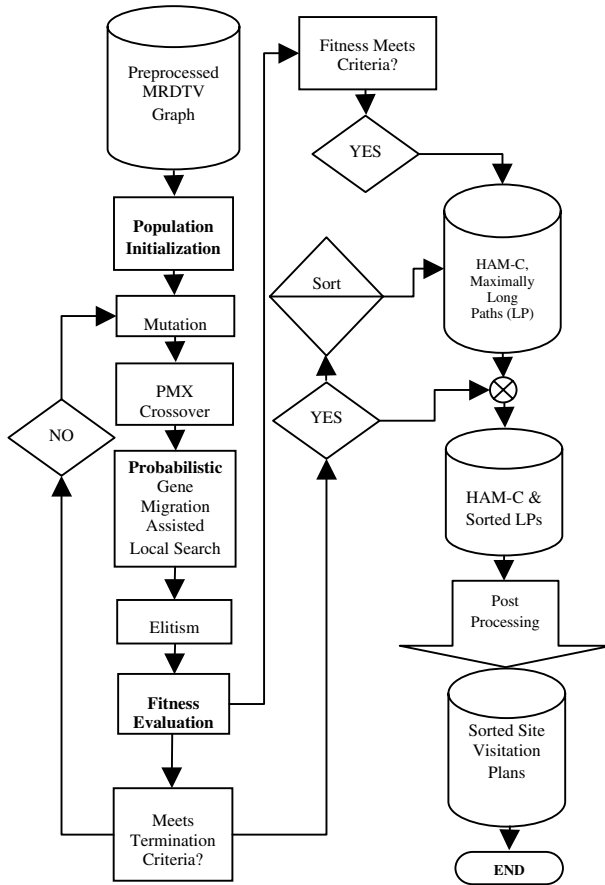


Fig. 2. HGA flowchart. Portions of this algorithm which differ from the one discussed in [1] are in bold

Gene Encoding. The mapping from the set of vertices of the MRDTV graph and the set of genes in a chromosome is bijective. Thus, each chromosome has $|S + A|$ genes. Also, the allelic value of any gene is the same as the numeric label of the corresponding vertex. One of the differences in the HGA used here is that the first gene in all chromosomes is fixed and always has the allelic value of the label for A .

Population Initialization. The population of chromosomes is randomly initialized. However, the allelic value of the first gene in all chromosomes is fixed.

Mutation. Chromosomes are probabilistically selected for mutation. Once selected, a multiple site, paired mutation is performed. Mutation must occur in pairs else they will result in the violation of the bijective mapping scheme. A restriction is that no gene can mutate to the allelic value of the first gene.

PMX Crossover. We select chromosomes in pairs for crossover using a roulette wheel scheme. We used the PMX operator [20] which is a popular scheme in tour-related problems. It guarantees feasible children chromosomes.

Fitness Function. The idea of *connected* genes is mentioned in Section 3. A chromosome's fitness is the length of the uninterrupted sequence of connected genes beginning at the first gene minus one. The objective of our HGA is to maximize this length as opposed to maximize the length of a connected sequence beginning at any gene in [1]. In the limit that one is able to uncover a Hamiltonian cycle in the process, the two objectives lead to the same solution topology – a ring. However, our current method is superior as discussed in Section 1.

Probabilistic Gene Migration Assisted Local Search. The gene migration scheme is a local search scheme, without which, the performance of the HGA is extremely poor. It works as follows. Define *nonconnected* genes as an idempotent negation of *connected* genes in a chromosome. Select a *nonconnected* gene in a chromosome and check if its insertion in any position in the *connected* sequence equals or improves the fitness of the corresponding chromosome; if it does then preserve the migration. This procedure should be repeated for all *nonconnected* genes until all of them are moved to the connected set or, alternately, no such movement is possible within a given chromosome. Currently, we use a uniform probabilistic scheme for selecting the gene that is the next candidate for potential migration as opposed to a sequential selection scheme in [1]. While this doesn't seem to affect the quality of the solutions in our simulations, it does appear to help generating a greater variety of good alternative solutions.

The GA outputs alternative plans in a nonincreasing order of their fitness values.

5 Dataset and Results

The dataset used here are simulated MRDTV graphs which are topologically isomorphic to the dataset for [1]. However, in this work, we generated 3 graphs from each of the 14 graphs in [1]. Each of these triplets differs from each other in the location of the airbase even while being topologically isomorphic to one another. Thus, in all, our data comprises 42 labeled graphs. It is available at [21].

The algorithm was coded in Java++. 15 simulation runs were conducted for each graph on a Pentium IV, 2.4 GHz PC. We generated up to 15 alternative site visitation plans in all simulations discussed here. Some results are tabulated in Table 1. For each triplet of graphs whose unlabeled graphs are isomorphic, the simulation results did not differ much. This suggests that at least for the dataset considered here, the algorithm's performance does not depend much on where the airbase is located. Instead, the overall graph topology is far more important. Thus, in order to conserve space, in each row of results in Table 1, we have given the average values obtained from simulations on three graphs that differ only in the node which represents the airbase.

Table 1. Simulation Results. The dataset was generated from 14 nonisomorphic graphs. It comprises 3 graphs generated from each of these 14 graphs that differ in the vertex which is used to represent the airbase.

Simulated MRDTV GRAPHS	Max. known Path Length (PL)	Results with Graph Compression Preprocessing		
		Avg. Best PL	Max. PL	Avg.CPU Time (s)
V_{CG} = number of vertices in compressed graph				
G01(15, 20), $V_{G01_{CG}} = 13$	14	14.0	14	0.040
G02(20, 30), $V_{G02_{CG}} = 20$	20	20.0	20	0.063
G03(20, 37), $V_{G03_{CG}} = 19$	20	20.0	20	0.066
G04(46, 69), $V_{G04_{CG}} = 46$	45	45.0	45	1.04
G05(46, 71), $V_{G05_{CG}} = 46$	46	45.1	46	0.56
G06(46, 93), $V_{G06_{CG}} = 46$	46	45.7	46	0.781
G07(49, 96), $V_{G07_{CG}} = 46$	48	48.0	48	0.461
G08(49, 142), $V_{G08_{CG}} = 49$	48	48.0	48	0.443
G09(100, 197), $V_{G09_{CG}} = 61$	99	95.6	98	4.458
G10(100, 197), $V_{G10_{CG}} = 60$	100	98.1	100	5.986
G11(100, 203), $V_{G11_{CG}} = 60$	100	96.4	99	4.324
G12(100, 207), $V_{G12_{CG}} = 77$	99	97.5	98	7.053
G13(100, 217), $V_{G13_{CG}} = 86$	99	97.6	99	10.913
G14(100, 217), $V_{G14_{CG}} = 85$	100	98.5	99	6.647

6 Conclusion

We have given a hybrid genetic algorithm which finds up to a user-specified number of alternative reconnaissance site visitation plans for an unmanned reconnaissance aerial vehicle. The plans are Hamiltonian cycles and acyclic paths beginning at the airbase in a *minimal risk density target visitation* graph along which the URAV can possibly fly to a large number of useful vicinities of reconnaissance sites without any revisitation. This helps lower the risk to the URAV operating in a hostile zone. A probabilistic gene migration scheme helps us obtain a greater diversity in visitation orderings while not affecting the fitness. Thus a larger number of alternative good site visitation sequencing plans are made available to the user. This can be particularly useful in cases involving large uncertainties or requiring further decision refinements. Our work helps in building onboard decision making capabilities in unmanned reconnaissance aerial vehicles using genetic algorithms. We believe that our general solution methodology is applicable to any unmanned vehicle system (ground, or marine) operating in threat zones. However, the utility of such solutions to other classes of vehicles probably needs further investigations.

References

1. Agarwal A., Lim, M.H., Xu Y.: Graph Compression via Compaction of degree-2 nodes for URUV Visitation Sequencing. UVS Tech 2003. Brussels, Belgium, December 2003
2. Vleugels J., Overmars, M.: Approximating Generalized Voronoi Diagrams in Any Dimension. Intl. J. on Computational Geom. and Applications. 8:201-221, 1998.
3. Arya S., Vigneron, A.: Approximating a Voronoi Cell. Research Report HKUST-TCSC-2003-10, 2003
4. Cormen, T.H., Leiserson, C.E., Rivest, R.L., Stein, C.: Introduction to Algorithms, 2nd edn. MIT Press, 2001
5. Wagner, I.A., Bruckstein, A.M.: Hamiltonian(t) – An Ant Inspired Heuristic for Recognizing Hamiltonian Graphs. Special Session on Ant Colony Methods, 1999 Congress on Evolutionary Computation. Washington DC, USA, July 1999
6. Rubin, F.: A Search Procedure for Hamiltonian Paths and Circuits. J. of the Assoc. of Computational Machinery. Vol. 21, No. 4, October 1974
7. Frank, Martel, J.C.: Phase Transitions in Random Graphs. 1st Intl. Conf. on Principles and Practice of Constraint Programming, Cassis, France, September 1995
8. Frieze, A. Yukich, J.: Probabilistic Analysis of the Traveling Salesman Problem. Book Chapter, June 2001
9. Vandegriend, B.: Finding Hamiltonian Cycles: Algorithms, Graphs and Performance. MS Thesis, University of Alberta, Spring 1998.
10. Mazade, L. (ed.): The Army's Bandwidth Bottleneck, Congressional Budget Office (USA), August 2003
11. Klausner K.A.: Command and Control of Air and space Forces Require Significant Attention to Bandwidth. Air & Space Power J., November 2002
12. Kopp, C: The UCAV Ascendancy: What are the Problem Issues? UAV Australia Conference. Melbourne, Australia, February 2001
13. Veverka, J., Campbell M.: Experimental Study of Information Load on Operators of Semi-Autonomous Systems. AIAA Guidance, Navigation and Control Conf. and Exhibit. Texas, USA, August 2003
14. CD ROM, Proceedings. UVS Tech 2003. Brussels, Belgium, December 2003
15. Unmanned Systems 2003 Proceedings. Association of Unmanned Vehicle Systems. Maryland, USA, July 2003
16. Unmanned Aerial Vehicles Roadmap 2002 – 2027. US DoD, December 2002
17. Van Kreveld, M., Overmars, M., Schwarzkopf, O., De Berg, M. (ed.): Computational Geometry. Algorithms and Applications. Springer-Verlag, 2nd Revised Edn. 2000
18. Braysy, O.: Genetic Algorithms for the Vehicle routing Problem with Time Windows. Arpakannus 1/2001 - Special Issue on Bioinformatics and Genetic Algorithms. 2001
19. Pereira, B.F., Tavares, J., Machado, P., Costa, E.: GVR: a New Genetic Representation for the Vehicle Routing Problem. 13th Irish Conference on Artificial Intelligence and Cognitive Science. Limerick, Ireland, September 2002
20. Goldberg D.E., Lingle, R.: Alleles, Loci, and the TSP. Proc. of the 1st Intl. Conf. on Genetic Algorithms, New Jersey, USA, 1985
21. http://www.ntu.edu.sg/home/emhlim/dataset/ALCT_SFAPAURAVSVS_GECCO04/