

# Iterator-Based Temporal Logic Task Planning

Sebastián A. Zudaire<sup>1</sup>, Martin Garrett<sup>2</sup>, and Sebastián Uchitel<sup>3</sup>

**Abstract**—Temporal logic task planning for robotic systems suffers from state explosion when specifications involve large numbers of discrete locations. We provide a novel approach, particularly suited for tasks specifications with universally quantified locations, that has constant time with respect to the number of locations, enabling synthesis of plans for an arbitrary number of them. We propose a hybrid control framework that uses an iterator to manage the discretised workspace hiding it from a plan enacted by a discrete event controller. A downside of our approach is that it incurs in increased overhead when executing a synthesised plan. We demonstrate that the overhead is reasonable for missions of a fixed-wing Unmanned Aerial Vehicle in simulated and real scenarios for up to 700 000 locations.

## I. INTRODUCTION

Discrete event controller synthesis is receiving increased attention as a means for providing robot applications correct-by-construction task plans (e.g., [1], [2], [3]). Synthesis from temporal logic specifications requires a discrete abstraction of the environment to establish a discrete event model that can be analysed exhaustively to produce task plans.

Synthesis algorithms are computationally complex (e.g., [4] is polynomial) with respect to the number of states of the discrete model. Hence, it is crucial to establish an abstraction of the environment that is sufficiently fine grained to allow appropriately capturing task requirements but coarse enough so as to not making synthesis intractable.

A robot’s workspace may be naturally discretised to the sensors’ capabilities, e.g., land mapping with a low-autonomy Unmanned Aerial Vehicle (UAV) can require over 400 discrete locations [5]. The number of discrete locations can induce a combinatorial growth in the size of the discrete event model, which in turn can make synthesis intractable.

We provide a novel approach that allows scaling the number of locations in task planning by exploiting the following observation: Many robot tasks specifications are, or can be, expressed as a universal quantification over a set of locations (e.g., “*For all locations in the discrete workspace, if the location satisfies ... then visit it and do ... if ...*”). Examples include tasks in [1], [6], [7], [8], [9], [10], [11] and [12], where common robot task are surveyed. The common ground in these papers is the explicit management of locations, that makes synthesis intractable when increased. Indeed they do not report building plans for over 1200 locations.

Our approach uses a hybrid control framework [13], which can work with any motion planner [14], in which synthesised

task plans execute over an API that provides an iterator that manages and hides the discretised workspace, offering the plan one location at a time. Plans are synthesised from a specification that includes the task requirements and a model of iterator that abstracts the number of locations that it manages. Hence, the *synthesis time is constant* with respect to these locations.

The price to be paid for constant synthesis time with respect to locations managed by the iterator is at runtime: Plans can only make decisions and act upon the location currently offered by the iterator and cannot refer to locations explicitly. That is, a plan cannot request going to a named location  $x$ , rather it must iterate over locations asking for each one if it is location  $x$  (similar to the sensor-based approach in [6]). Thus, the order in which the iterator selects locations impacts the overall robot behaviour. A particularly bad case is if in a scenario with millions of locations, the iterator offers location  $x$  at the very end.

We show that a hybrid control layer in which location sorting uses shortest trajectory is fast enough to provide acceptable though sub-optimal flight paths for tasks involving hundreds of thousands of locations (significantly beyond what synthesis with explicit location management is capable of) for a fixed-wing UAV. The design is complementary to work on motion and trajectory planning [15]. Indeed, more sophisticated reasoning below the discrete plan layer can be modularly included into the hybrid controller and could provide enhanced performance.

In summary, we present a hybrid controller approach aimed at tasks specifications with universally quantified locations that does not suffer from synthesis scalability limitations with respect to the number of locations. Task plans are synthesised from specifications given as Labelled Transition Systems and Fluent Linear Temporal Logic. Despite using simple location motion planning and trajectory control approaches, we demonstrate by simulating and flying four tasks: search and follow [6], search and map [16], patrol [7], and cover [17], that the approach can scale to hundreds of thousands of discrete locations.

## II. PRELIMINARIES

**Labelled Transition System:** (LTS) [18] are automata where transitions are labelled with actions that constitute the interactions of the modelled system with its environment. We partition actions into controlled and uncontrolled to specify assumptions about the environment and safety requirements for a controller. Figure 2b models the assumption that `yes.fire` and `no.fire` are responses to `fire?`, and the safety property that `fire?` is not issued before the response to a previous `fire?`.

<sup>1</sup>Instituto Balseiro - Universidad Nacional de Cuyo, Argentina. [sebastian.zudaire@ib.edu.ar](mailto:sebastian.zudaire@ib.edu.ar)

<sup>2</sup>CNEA, Argentina. [medgarrett@cab.cnea.gov.ar](mailto:medgarrett@cab.cnea.gov.ar)

<sup>3</sup>Universidad de Buenos Aires, Buenos Aires, Argentina and Imperial College London, UK. [suchitel@dc.uba.ar](mailto:suchitel@dc.uba.ar)

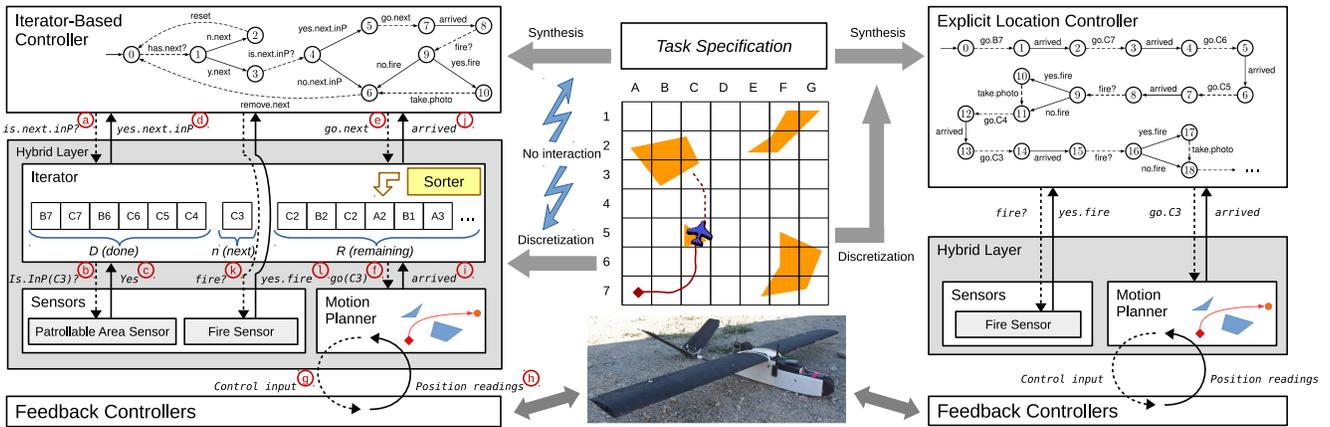


Fig. 1: System Architecture for Iterator-Based and Explicit location plans. The fixed-wing UAV has a wingspan of 1.6 m

Complex models can be constructed by LTS composition. We use a standard definition of parallel composition ( $\parallel$ ) that models the asynchronous execution of LTS, interleaving non-shared actions and forcing synchronisation of shared actions.

**Fluent Linear Temporal Logic:** (FLTL) [19] is also used to describe environment assumptions and task requirements. FLTL is a linear-time temporal logic that uses fluents to describe states over sequences of actions.

A fluent  $fl$  is defined by a set of initiating actions, a set of terminating actions, and an initial value. We may omit set notation for singletons, e.g.,  $Going = \langle go.next, arrived \rangle_{Initially \perp}$ . We may use an action label  $\ell$  for the fluent defined as  $fl = \langle \ell, Act \setminus \{\ell\} \rangle$ . Thus, the fluent  $remove.next$  is only true just after the occurrence of the action  $remove.next$ .

FLTL is defined similarly to propositional LTL but where a fluent holds at a position  $i$  in a trace  $\pi$  based on the events occurring in  $\pi$  up to  $i$ . Temporal connectives are interpreted as standard:  $\diamond\varphi$ ,  $\square\varphi$ , and  $\varphi \mathbf{W} \psi$  mean that  $\varphi$  eventually holds, always holds and holds until  $\psi$  respectively.

**Discrete Event Controller Synthesis** is defined as follows: Given an LTS  $E$  with a set of controllable actions  $L$ , assumption  $A$  and goal  $G$  expressed in FLTL, find an LTS  $C$  such that  $E \parallel C$  is deadlock free,  $C$  does not block any non-controlled actions, and for every trace of  $E \parallel C$  if the trace satisfies the assumption  $A$ , then the trace satisfies  $G$ .

When goals and assumptions are restricted to a GR(1) form [4] the control problem can be solved in polynomial time. MTSA [20] solves GR(1) control problems expressed with LTS and FLTL, requiring assumptions and goals in FLTL to be either *i*) of the form  $\bigwedge_{i=1}^n \square \diamond \varphi_i$  where  $\varphi_i$  are Boolean combinations of fluents, or *ii*) safety properties [21].

### III. DISCRETE ABSTRACTION

#### A. Iterator-Based Task Plans

Consider a task for a UAV in which various locations of a grid-based map must be patrolled and photographs must be taken if fire is detected. A plan for such a task in a temporal logic task and motion planning approach (e.g., [8]) might look like the LTS on the top right of Figure 1 (*Explicit Location Controller*) where the locations to be patrolled are

$P = \{C5, A2, B2, \dots\}$  (orange areas in Figure 1). The plan sequentially visits each location, checks for fire and takes a picture accordingly. Although the size of the plan grows linearly with  $P$ , the state space over which it is computed grows exponentially (i.e.,  $2^{\|P\|}$ ) as it must at least capture all possible orders in which locations in  $P$  could be visited.

A more compact plan for the same task may be synthesised if a richer execution environment is assumed. Consider an iterator that abstracts the size of discrete workspace and which can provide its locations, one at a time. In this case, a plan (top left of Figure 1) could consist of a loop iterating over the locations, checking for each location if it requires patrolling ( $is.next.inP?$ ) and if so ( $yes.next.inP?$ ) going to the location ( $go.next$ ) and upon arrival ( $arrived$ ) checking if the current UAV location has fire ( $fire?$ ) and taking a photo ( $take.photo$ ) if needed. As locations are not explicitly treated in the plan, its size does not depend neither on the total number of locations nor the size of  $P$ . Similarly, the state space from which the plan can be synthesised is not affected, achieving *constant synthesis time* with respect to the number of locations. In the remainder of this section we report on how to specify and synthesise iterator plans.

#### B. Specification of Iterator-Based Task Plans

Three aspects of the system must be abstracted to obtain a discrete event model from which to synthesise iterator-based task plans: the iterator, sensors and actuators.

*1) Iterator Abstraction:* The iterator is an abstract data type that manages a set of discrete locations  $\mathcal{L}$  derived through the discretization of a region. We chose for simplicity to use grid-based maps as in [2], [9], [17], [22].

An iterator is a triple  $\langle D, n, R \rangle$  where  $D$  and  $R$  are sets of locations representing those that the plan has already processed (*Done*) and those that remain (*Remaining*), and  $n$  is the next location to be processed by the plan. The iterator is initialised as follows:  $n$  is set to one element of  $\mathcal{L}$ ,  $R = \mathcal{L} \setminus \{n\}$ ,  $D = \emptyset$ . We define the following operations:

- $has.next?$ : returns true if and only if  $n$  is not null.
- $remove.next$ : adds  $n$  to  $D$  and if  $R = \emptyset$  sets  $n$  to null, otherwise sets  $n$  to a location in  $R$  and  $R = R \setminus \{n\}$ .

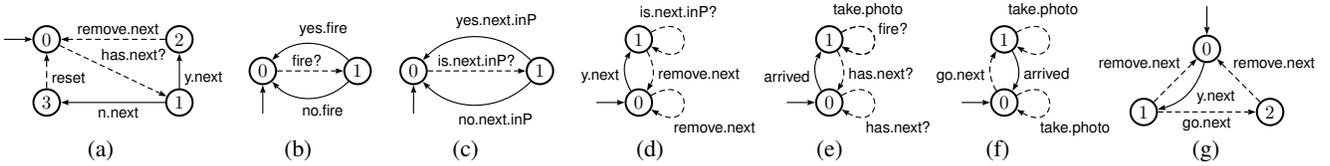


Fig. 2: Dashed (non-dashed) lines are controlled (uncontrolled) actions. (a) Iterator Model. (b) Fire sensor for current location. (c) Patrollable area sensor. (d) Constraint: `is.next.inP?` only when `y.next`. (e) Constraint: `take.photo` and `fire?` only when `arrived`. (f) Simplified UAV capabilities for Fire Patrol task. (g) Constraint: `go.next` only when `y.next`.

- `reset`: the iterator is reinitialised.

In Figure 2a we depict an LTS that models interactions with an iterator. We model the return values of `has.next?` with two different events `y.next` and `n.next`. These two events are defined to be uncontrollable, i.e., it is the iterator and not the plan that decides whether the response is `y.next` or `n.next`. A requirement such as  $\Box\Diamond\text{has.next?} \wedge \Box(\text{y.next} \Rightarrow (\neg\text{remove.next} \mathbf{W} \text{go.next}))$  will make the robot continuously visit all the locations with which the iterator was initialised.

2) *Sensors*: Similarly to [6], we introduce binary sensors to model interaction with the environment. In this iterator-based setting, sensors can answer queries regarding the location that the robot is at and/or for the next location selected by the iterator. For example, Figure 2b models a fire sensor that can be queried about the existence of fire at the current robot location. Figure 2c shows an abstraction of the sensor that responds if the next location selected by the iterator is one that must be patrolled.

For sensing over the next location, attribute  $n$  must not be null. An additional LTS is included to constrain the occurrence of `is.next.inP?` queries to between `y.next` and `remove.next` as in Figure 2d.

Additional constraints are typically needed for sensing over the current location to ensure that the plan is aware of what the current location is. For example, sensing for fire (`fire?`) and taking a photo (`take.photo`) should occur between having arrived to a particular location and starting to analyse the next (`has.next?`), Figure 2e.

3) *Primitive Capabilities*: Using a control-driven discretization [23], we define controllable/uncontrollable pairs to model the start/end of control modes [24]. For example, `go.next` commands the robot to move to the next location according to the Iterator, and the uncontrollable action `arrived` indicates that the target location has been reached. Other capabilities may be reasonably modelled as instantaneous such as `take.photo`.

In Figure 2f, we depict the minimal capability model of a robot for the Fire Patrol task. We also require that the `go.next` command only be issued when the iterator has a next location to be processed (Figure 2g).

### C. Task Specification Example

To help understand how the abstraction described above works, we elaborate on how the Fire Patrol task can be specified to obtain the plan shown in top left of Figure 1.

We build an environment model  $E$  as a parallel composition that describes assumptions and constraints on how

the infrastructure on which the plan will execute. For the Fire Patrol task the composition includes exactly all the LTS described above: the iterator, the fire and patrol sensors with associated constraints and the robot capabilities and constraints, as depicted in Figure 2. .

We structure the task specification with one property stating which location should be visited ( $\varphi_1$ ) and another one for what should be done at visited locations ( $\varphi_2$ ). We also require  $\varphi_0 = \Box\Diamond\text{has.next?}$  to ensure that the plan continually processes locations from the iterator.

For property  $\varphi_1$  we need to introduce three fluents:  $\text{MustPatrol} = \langle \text{yes.next.inP}, \text{has.next?} \rangle_{\text{Initially}_\perp}$  that is true when the location selected by the iterator has been confirmed to be in the set of patrollable locations  $P$  (`yes.next.inP`).  $\text{PatrolAnswered} = \langle \{ \text{no.next.inP}, \text{yes.next.inP} \}, \text{has.next?} \rangle_{\text{Initially}_\perp}$  is true when a response to `is.next.inP?` has been received.  $\text{Arrived} = \langle \text{arrived}, \text{has.next?} \rangle_{\text{Initially}_\perp}$  is true when the UAV has arrived to the location selected by the iterator.

The patrol condition  $\text{VisitCondition} = \text{PatrolAnswered} \wedge (\text{MustPatrol} \iff \text{Arrived})$ , is that the `is.next.inP?` query has been responded, and the UAV has arrived at that location if and only if the response was `yes.next.inP`. Additionally,  $\varphi_1 = \Box(\text{y.next} \Rightarrow \neg\text{remove.next} \mathbf{W} \text{VisitCondition})$  requires, for every new location that is selected by the iterator, to not remove that location from the iterator until the  $\text{VisitCondition}$  is achieved.

The specification of what to achieve at each visited location ( $\varphi_2$ ) follows a similar pattern. We use a fluent  $\text{FireDetected} = \langle \text{yes.fire}, \text{has.next?} \rangle_{\text{Initially}_\perp}$  to model that fire has been detected at the current location, fluent  $\text{PhotoTaken} = \langle \text{take.photo}, \text{has.next?} \rangle_{\text{Initially}_\perp}$  to model that a photo has been taken and  $\text{FireAnswered} = \langle \{ \text{yes.fire}, \text{no.fire} \}, \text{has.next?} \rangle_{\text{Initially}_\perp}$  to model reception of a response to `fire?`.

The condition to be achieved once arrived at a location  $\text{ArrivedCondition} = \text{FireAnswered} \wedge (\text{FireDetected} \iff \text{PhotoTaken})$  is that a response from the fire sensor must have been received and that a photo should be taken if and only if the response is positive. Consequently, we have  $\varphi_2 = \Box(\text{arrived} \Rightarrow \neg\text{remove.next} \mathbf{W} \text{ArrivedCondition})$ .

If  $E$ ,  $\varphi_0$ ,  $\varphi_1$ , and  $\varphi_2$  as defined above are fed to MTSA [20] then the resulting controller is the one depicted in the top left of Figure 1 (*Iterator-Based Controller*).

## IV. HYBRID CONTROL LAYER

A hybrid control layer (e.g., [1], [13]) provides an interface between a discrete controller and the lower level continuous

control of the robot. Figure 1 shows an architecture both for our iterator-based approach and one that manages locations explicitly at the discrete layer (e.g., [6], [13]).

In an *iterator-based approach* the workspace is discretised independently of the synthesis procedure and fed to the Iterator module before the start of the mission. For the Fire Patrol task, the discretization also feeds the Patrollable Area Sensor with the locations that appear in orange ( $P = \{C5, A2, B2, \dots\}$ ) in the map of Figure 1.

At runtime, `has.next?`, `y.next` and `n.next` are used to loop over the discretised locations. In Figure 1 the next location in the Iterator is C3. When the plan executes `is.next.inP?`, it produces a call to the Iterator (see red ③ in Figure 1) that then forwards the request `is.inP(C3)?` to the Patrollable Area sensor ④. The sensor confirms that  $C3 \in P$  ⑤ and the plan receives event `yes.next.inP` ⑥.

Similarly, when `go.next` is issued ⑦, the Iterator makes a `go(C3)` call to the motion planner ⑧. The motion planner generates the control inputs ⑨⑩ to reach C3 possibly also performing static and dynamic obstacle avoidance.

Once the target location is reached, the arrived event is propagated upwards ⑪⑫ to the plan which then queries the existence of fire? ⑬. Note that as this query involves sensing the current robot location (and not the Iterator’s next location), the event is sent directly through to the Fire Sensor.

In an *explicit location approach*, synthesis requires information of the discretization to determine the order in which locations are to be visited. In the Fire Patrol example, the order in which locations in  $P$  are to be patrolled is decided by the synthesis procedure (instead of the Iterator).

Furthermore, the explicit location plan controls the path that must be followed by the robot, i.e. while the iterator-based approach set C5 as the first patrol location to visit, the explicit-location plan sets the path to be followed to reach C5: B7, C7, C6, C5. This allows some static obstacle avoidance manoeuvres (e.g., [1], [2]). Nonetheless, the motion planner must still generate the control inputs between adjacent discrete locations, deal with fine grained static obstacle avoidance and also dynamic obstacle avoidance (e.g., [11]).

**Location Sorting:** The order in which the Iterator offers locations to the discrete event controller can have significant impact. Consider a Fire Patrol mission shown on the left side of Figure 1. The UAV started in location A7 and was offered B7 as the first location. As it is not an area to be patrolled, it was removed from the Iterator. This also occurred for C7, B6, and C6. Only when C5 was selected as the next element did the UAV `go.next` to that location. Many more locations that do not correspond to patrol areas could have been offered thus delaying the first `go.next` command. In addition, a much more distant patrol location (e.g., G1) could have been selected, forcing possibly a less efficient patrol strategy.

Consequently, an important component of the hybrid layer is the Sorter. Our hybrid layer design works on the assumption that the best next location to offer is a function of the distance from the current robot location. This is a challenge as sorting must be done over a large set of locations regularly. Sorting is performed while the robot is travelling between

the requested location (`go.next`) and the moment it reaches it (`arrived`). Distances are computed with respect to the location that the robot will have once arrived occurs.

The sorting criteria must be simple enough to allow fast computing of each location’s priority but not oversimplified, in order to produce acceptable overall task trajectories. In the next section we demonstrate experimentally that sorting over trajectory length using a simplified robot dynamic behaviour model allows fast enough sorting while providing reasonable trajectories for a fixed-wing UAV travelling at 17m/s.

## V. VALIDATION

We first show applicability to tasks taken from [6], [12] and [16], and we analyse scalability. Task specifications and results are available at [25].

### A. Experimental Configuration

All experiments were run on either a simulated or real fixed-wing battery-powered UAV.

We used the robot in Figure 1, with low-level control provided by an off-the-shelf Pixhawk autopilot loaded with Ardupilot firmware [26] ArduPlane, and sensors providing information of the system’s environment (e.g., Raspberry Pi Camera Module V2 for capturing ground images).

We built the discrete plan interpreter and hybrid control layer by extending the Ground Control Station (GCS) software MAVProxy [16] with custom Python modules. The hybrid control is run on an onboard Raspberry Pi 3B+ and communicates with the autopilot via the telemetry serial port. We also used an instance of MAVProxy to allow human monitoring from the ground on a laptop which communicates with the autopilot via a SiK Telemetry Radio. Note, however, that mission execution is entirely run on onboard.

For simulations, we replaced the plane and autopilot with the ArduPilot Software In The Loop (SITL) simulator that simulates the UAV’s dynamics, autopilot and physical environment. This allows keeping the exact same onboard computer and hybrid control software as in the real flights. In the simulations we feature automatic takeoff and landing, while for safety reasons in the real flights a remote control (RC) radio system was connected to the Pixhawk to perform manual takeoff and landing.

The robot capability model used is an extension of Figure 2f to support taking-off and landing. Discrete event controllers were synthesised using MTSA [20] and loaded onto the onboard computer before starting each mission.

**Motion Planning and Iterator Sorting.** Although motion planning has been a greatly researched for both online and offline computation (e.g., [14], [15]), we implemented a fairly simple scheme, sufficient for our experimental goals, that does not consider static or dynamic obstacles.

The motion planner generates sequences of control inputs for the autopilot based on trajectories that are computed by concatenating straight paths and turns for a given maximum turn radius, similar to [27]. The planner finds (assuming constant speed and a maximum of two turns) a trajectory for a given arrival direction at location  $n$  from the current

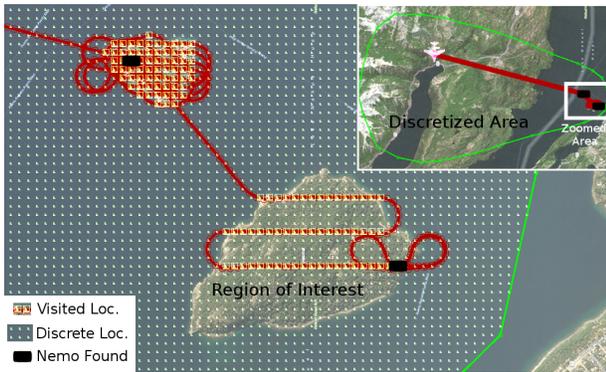


Fig. 3: Simulation of the Find Nemo task.

location and velocity vector of the UAV. We force the arrival direction to be parallel to a given fixed direction (e.g., grid axis to favour straight orderly grid coverage or perpendicular to the wind direction to increase flight stability).

The Sorter uses the same trajectory computation. With a  $50\text{ m} \times 50\text{ m}$  discretization and a UAV that flies at  $17\text{ m/s}$ , the minimum flight time between a `go.next` command and an arrived event is  $2.9\text{ s}$ . In this time, at least 40 000 locations can be sorted on the onboard Raspberry Pi 3B+.

## B. Tasks

1) *Find Nemo*: The Find Nemo task [6] requires a robot with a Nemo sensor and camera to search for Nemo in 4 regions of interest out of a total of 12. The task is to continuously, for all regions of interest, go, sense for Nemo, and if found stay and photograph.

We synthesised and ran a simulated task for 437 regions of interest over a total of 102 307 regions. We used as locations of interest the two islands that can be seen in Figure 3. We randomised the appearance and disappearance of Nemo. Figure 3 shows the UAV’s path while searching and finding Nemo in two locations (one in each island), visiting this location until Nemo disappears and then resuming the search.

2) *Search and Map Target*: Inspired on [16], we specified and flew a task requiring to find a red target in a field and to map out all locations from which the target is visible and then land. The task specification can be structured as two modes. The first is a search, very much like Nemo, using a target sensor that captures an image and processes it to search for a red  $2\text{ m} \times 2\text{ m}$  target. If the target is found, the location is stored in a shared data structure. The second is a map mode in which all cells adjacent to one from which the target was seen must be visited, for which a second sensor responds `y.adjacent.next` if the location selected by the Iterator is adjacent to one in the shared data structure.

We ran our UAV for this task for 29 502 locations. The estimated flight time of covering all locations optimally at  $17\text{ m/s}$  is over 24 hours, significantly over the flight autonomy of the UAV (40 min or  $\sim 800$  locations). For this reason the red target was set relatively close to the launchpad. The video in [25] shows the flight of the UAV including both

the trajectory as displayed by the monitoring ground station, the photos taken at locations and the mapped area.

3) *Ordered Patrol*: In [12], a study of common UAV task requirements taken from over robotics papers is presented. One common requirement is that of an ordered patrol which requires visiting a set of locations in a particular order and at each one performing some task. In essence, this is the Fire Patrol task (which is another common pattern [12], the unordered patrol), but with an order for which the patrol locations must be visited. This task requires explicit treatment of locations at the plan level. A sensor for each location is needed to include in the requirements the sequence in which these sensors must be used to find the places to visit.

We ran simulated versions of this task which we report on in Section V-D to analyse the scalability of our approach.

4) *Cover*: Another requirement referred to in [12] is that of covering (i.e., visiting) a statically defined region. The specification of such tasks in an iterator-based fashion does not differ significantly from the Fire Patrol task. The set of locations to be covered can be determined at runtime by the plan by using a sensor like `is.next.inP?`. We ran simulated versions of this task which we report on in Section V-D.

## C. Synthesis Time Scalability

The synthesis time for the tasks discussed in this paper was lower than  $5\text{ s}$  on a laptop with a Intel i7 3.5GHz processor and 12GB of RAM. As discussed previously, this time is independent of the size of the locations over which the Iterator operates, which means that, except for Ordered Patrol, the number of locations can be scaled indefinitely. This includes both total locations and the locations to-patrol, of-interest, to-cover, and red locations in the Fire Patrol, Find Nemo, Cover, and Search and Map tasks. However, our approach is not independent of the number of sensors required to specify the task. Thus, for the Ordered Patrol task, although the total number of locations can be scaled, the number of locations to be patrolled in a specific order cannot. Indeed, for  $n$  locations to be orderly visited, the state space for synthesis will grow  $2^n$  as in any synthesis approach [7].

## D. Runtime Overhead

To analyse the overhead (in terms of flight time) of iterating increasingly large location sets at runtime we selected, based on our understanding of our approach, best and worst case tasks. To understand the impact of the Sorting component, we ran these tasks for three sorting strategies: the one described in Section IV (distance), a highly inefficient one that puts the interesting elements to visit at the end of the iterator (last), and a random ordering (random).

Our *worst case task* is an Ordered Patrol (Section V-B.3) because at runtime the Iterator may continuously offer last the next patrol location to be visited according to plan. This forces the plan to iterate over all locations every time before `go.next`. We simulated the Ordered Patrol of 3 locations as in [28] with variable amount of discrete locations.

Our *best case task* is a Cover task as one iteration over the whole location universe suffices to complete the task.

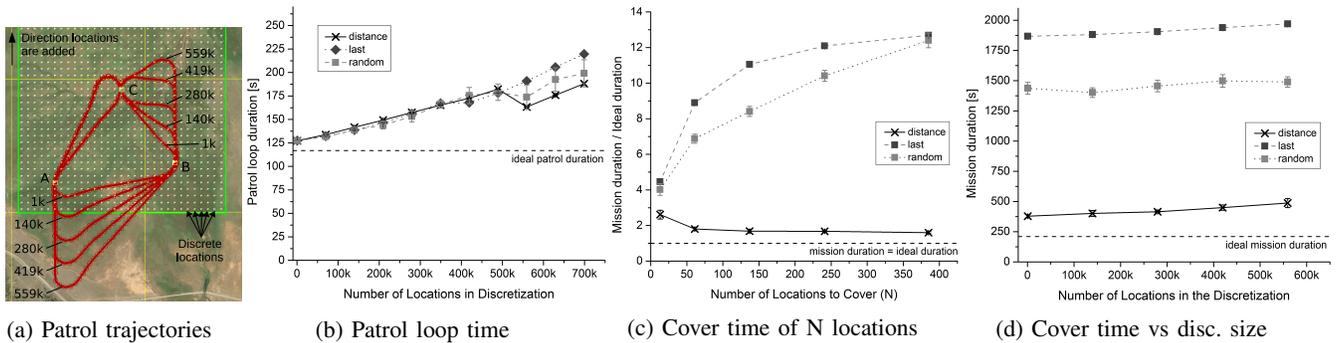


Fig. 4: (a) Patrol trajectories for different discretization sizes using *last* sorter. (b) Loop duration for Ordered Patrol varying total number of locations and sorting criteria. (c) Proportional overhead for Cover task varying number of locations to cover, universe of locations fixed at 713. (d) Duration of covering 61 locations varying total locations. *Note*: The error bars are three times the estimated standard error of the mean, and can be smaller than the symbol. Ideal duration is calculated using the minimum flight distance needed for the mission using constant speed, ignoring the UAV’s movement restrictions.

Of course, the order in which locations to be covered are offered may produce more or less efficient coverage paths. We simulated a cover task with variable amount of discrete locations and contiguous locations to be covered.

For the Ordered Patrol task, we depict in Figure 4b the time to visit all locations once for an increasing location universe size. Each data point is the average over 30 simulations for a particular sorting criteria. For the Cover task, we show in Figure 4c how cover time increases as the number of locations to cover does. Mission duration (up to 3 hours each) required limiting experimentation to 6 samples per sorting criteria and size. Finally, we depict how the size of location universe impacts mission duration when requiring to cover 61 locations (Figure 4d).

Figure 4b shows that degradation seems linear in the number of total locations and that the three sorting strategies make little difference in relative terms when comparing the overall mission duration to the ideal mission duration. The overhead for 700 000 locations is at most 88% over an ideal patrol mission. Figure 4a exemplifies why the the UAV’s trajectory between the three locations to patrol degrades. As the universe of locations increases, so does the distance the UAV is flying in a straight line while iterating over all these locations to get the next patrollable location.

In contrast, Figure 4c shows that sorting based on trajectory distance makes a big difference and provides near constant proportional overhead (60%) when increasing the size of region to cover. Figure 4d also shows relevance of the sorting criteria and near constant duration covering a 61 location size while increasing the universe of locations.

## VI. RELATED WORK

The state explosion problem in temporal mission planning has been addressed in various ways. *i*) Improved online/offline motion planning (e.g., [9], [29], [30], [31]) is orthogonal to our approach and can be introduced within our hybrid layer replacing the Sorter and Motion Planner components. We do not compare empirically against these approaches, rather we show that even simple sorting and

motion strategies already yield reasonable results. *ii*) Advances in synthesis efficiency (e.g., [32], [33], [34]) are also orthogonal. An iterator abstraction can be used with a variety of approaches to synthesis to scale task planning orders of magnitude beyond what can be achieved when all universally quantified locations must be explicitly referred to. *iii*) Alternative strategies for integrating task and motion planning. The distinctive feature of the strategy presented in this paper is the use of an Iterator coupled with runtime motion planning. In [10] a plan outline is manually constructed without explicit naming of locations, later to be filled offline by an SMT solver. Scale is limited by the solver, which depends on the number of locations. Indeed, as in another SMT-based approach [22] reported cases are below 800 discrete locations. In [7] robot’s paths are produced for LTL specifications by solving constrained reachability problems, but total complexity depends on the number of locations. [1], [3] combine task and motion requirements in a GR(1) specification. Complexity of GR(1) is polynomial respect to the state space which grows combinatorially to the number of locations. [35] propose a highly hierarchical approach which is able to solve large complex workspaces, but requires domain-dependent choices for the hierarchy.

In all these approaches, increasing the number of universally quantified locations beyond a couple of thousand implies not being able to compute a plan. In contrast, in our approach, the increase in locations does not impede producing a plan, albeit with degraded mission trajectories due to runtime motion planning.

## VII. CONCLUSIONS

We propose iterator-based task planning to provide constant synthesis complexity with respect to the number of discrete locations universally quantified in a task specification. Iterator-based plans run on a hybrid control layer that performs runtime motion planning. We show that simple location prioritisation and motion planning strategies suffice to provide adequate mission behaviour for iterator-based plans both in simulated and real UAV missions.

## REFERENCES

- [1] G. E. Fainekos, H. Kress-Gazit, and G. J. Pappas, "Temporal logic motion planning for mobile robots," in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, 2005, pp. 2020–2025.
- [2] E. M. Wolff, U. Topcu, and R. M. Murray, "Efficient reactive controller synthesis for a fragment of linear temporal logic," in *2013 IEEE International Conference on Robotics and Automation*, May 2013, pp. 5033–5040.
- [3] J. A. DeCastro, V. Raman, and H. Kress-Gazit, "Dynamics-driven adaptive abstraction for reactive high-level mission and motion planning," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 369–376.
- [4] N. Piterman, A. Pnueli, and Y. Sa'ar, "Synthesis of reactive (1) designs," *Lecture notes in computer science*, vol. 3855, pp. 364–380, 2006.
- [5] A. Tariq, S. M. Osama, and A. Gillani, "Development of a low cost and light weight uav for photogrammetry and precision land mapping using aerial imagery," in *2016 International Conference on Frontiers of Information Technology (FIT)*, 2016, pp. 360–364.
- [6] H. Kress-Gazit, G. E. Fainekos, and G. J. Pappas, "Temporal-logic-based reactive mission and motion planning," *IEEE Transactions on Robotics*, vol. 25, no. 6, pp. 1370–1381, 2009.
- [7] E. M. Wolff, U. Topcu, and R. M. Murray, "Automaton-guided controller synthesis for nonlinear systems with temporal logic," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2013, pp. 4332–4339.
- [8] H. Kress-Gazit, G. E. Fainekos, and G. J. Pappas, "Where's waldo? sensor-based temporal logic motion planning," in *Proceedings 2007 IEEE International Conference on Robotics and Automation*, 2007, pp. 3116–3121.
- [9] S. C. Livingston and R. M. Murray, "Just-in-time synthesis for reactive motion planning with temporal logic," in *2013 IEEE International Conference on Robotics and Automation*, 2013, pp. 5048–5053.
- [10] S. Nedunuri, S. Prabhu, M. Moll, S. Chaudhuri, and L. E. Kavraki, "Smt-based synthesis of integrated task and motion plans from plan outlines," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, May 2014, pp. 655–662.
- [11] C. I. Vasile and C. Belta, "Reactive sampling-based temporal logic path planning," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 4310–4315.
- [12] C. Menghi, C. Tsigkanos, T. Berger, and P. Pelliccione, "Psalm: Specification of dependable robotic missions," in *2019 IEEE/ACM 41st International Conference on Software Engineering: Companion Proceedings (ICSE-Companion)*, May 2019, pp. 99–102.
- [13] H. Kress-Gazit, G. Fainekos, and G. Pappas, "Translating structured english to robot controllers," *Advanced Robotics*, vol. 22, pp. 1343–1359, 10 2008.
- [14] X. Ji and J. Li, "Online motion planning for uav under uncertain environment," in *2015 8th International Symposium on Computational Intelligence and Design (ISCID)*, vol. 2, Dec 2015, pp. 514–517.
- [15] N. Bezzo, K. Mohta, C. Nowzari, I. Lee, V. Kumar, and G. Pappas, "Online planning for energy-efficient and disturbance-aware uav operations," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct 2016, pp. 5027–5033.
- [16] H. Choi, M. Geeves, B. Alsalam, and F. Gonzalez, "Open source computer-vision based guidance system for uavs on-board decision making," in *2016 IEEE Aerospace Conference*, 2016, pp. 1–5.
- [17] M. Wei and V. Isler, "Coverage path planning under the energy constraint," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, May 2018, pp. 368–373.
- [18] R. Keller, "Formal verification of parallel programs." *Communications of the ACM*, vol. 19, pp. 371–384, 07 1976.
- [19] D. Giannakopoulou and J. Magee, "Fluent model checking for event-based systems," in *Proceedings of the 9th European Software Engineering Conference Held Jointly with 11th ACM SIGSOFT International Symposium on Foundations of Software Engineering*, ser. ESEC/FSE-11. New York, NY, USA: ACM, 2003, pp. 257–266.
- [20] N. R. D'Ippolito, V. Braberman, N. Piterman, and S. Uchitel, "Synthesis of live behaviour models," in *Proceedings of the eighteenth ACM SIGSOFT international symposium on Foundations of software engineering*, ser. FSE '10. New York, NY, USA: ACM, 2010, pp. 77–86.
- [21] L. Lamport, "Proving the correctness of multiprocess programs," *IEEE Transactions on Software Engineering*, vol. 3, no. 2, pp. 125–143, Mar. 1977.
- [22] T. Kundu and I. Saha, "Energy-aware temporal logic motion planning for mobile robots," in *2019 International Conference on Robotics and Automation (ICRA)*, May 2019, pp. 8599–8605.
- [23] C. Belta, A. Bicchi, M. Egerstedt, E. Frazzoli, E. Klavins, and G. J. Pappas, "Symbolic planning and control of robot motion [grand challenges of robotics]," *IEEE Robotics Automation Magazine*, vol. 14, no. 1, pp. 61–70, 2007.
- [24] A. Balkan, M. Vardi, and P. Tabuada, "Mode-target games: Reactive synthesis for control applications," *IEEE Transactions on Automatic Control*, vol. 63, no. 1, pp. 196–202, 2018.
- [25] Repository Package, accessed 2019-09-15. [Online]. Available: <https://bitbucket.org/szudaire/iterator-based-planning>
- [26] A. Sinisterra, M. Dhanak, and N. Kouvaras, "A usv platform for surface autonomy," in *OCEANS 2017 - Anchorage*, 2017, pp. 1–8.
- [27] S. Hota and D. Ghose, "Optimal transition trajectory for waypoint following," in *2013 IEEE International Conference on Control Applications (CCA)*, 2013, pp. 1030–1035.
- [28] E. M. Wolff and R. M. Murray, *Optimal Control of Nonlinear Systems with Temporal Logic Specifications*. Springer International Publishing, 2016, pp. 21–37.
- [29] T. M. Cabreira, L. Brisolara, and P. Ferreira Jr, "Survey on coverage path planning with unmanned aerial vehicles," *Drones*, vol. 3, p. 4, 01 2019.
- [30] Z. Yan, J. He, and J. Li, "An improved multi-auv patrol path planning method," in *2017 IEEE International Conference on Mechatronics and Automation (ICMA)*, Aug 2017, pp. 1930–1936.
- [31] M. Ghosh, S. Thomas, M. Morales, S. Rodriguez, and N. M. Amato, "Motion planning using hierarchical aggregation of workspace obstacles," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct 2016, pp. 5716–5721.
- [32] T. Wongpiromsarn, U. Topcu, N. Ozay, H. Xu, and R. M. Murray, "Tulip: A software toolbox for receding horizon temporal logic planning," in *Proceedings of the 14th International Conference on Hybrid Systems: Computation and Control*, ser. HSCC '11. New York, NY, USA: ACM, 2011, pp. 313–314.
- [33] C. Finucane, Gangyuan Jing, and H. Kress-Gazit, "Ltlmop: Experimenting with language, temporal logic and robot control," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct 2010, pp. 1988–1993.
- [34] S. Dathathri and R. M. Murray, "Decomposing gr(1) games with singleton liveness guarantees for efficient synthesis," in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, 2017, pp. 911–917.
- [35] L. P. Kaelbling and T. Lozano-Prez, "Hierarchical task and motion planning in the now," in *2011 IEEE International Conference on Robotics and Automation*, May 2011, pp. 1470–1477.