Adaptive Introspection and Deployment
for Robust Long Duration Autonomy

Nathan Michael, Sebastian Scherer
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh PA, 15213-3890
phone: (412) 268-7816   fax: (412) 268-1338   email: nmichael@cmu.edu, scherer@cmu.edu

LONG-TERM GOALS

Long duration autonomy for unmanned systems is difficult to achieve as current systems are limited to anticipated exceptions and do not adapt to long-term changes in the environment. The project goal is to enable long-term operation in unpredictable environments through adaptive introspection and deployment approaches, emphasizing vehicle and system level adaptation and robustness as a robotic team cooperates toward a common and persistent mission goal. At the vehicle level, we address the problems of identifying unexpected vehicle states and developing robust mitigation policies and behaviors through introspection. At the system level, we propose long-term planning methodologies that coordinate the robot team toward a common mission objective while learning, adapting to, and anticipating changing vehicle and environment conditions. The proposed research outcomes will be broadly applicable to unmanned vehicle systems operating over long time horizons to achieve a common and persistent mission. In order to evaluate the performance of the proposed methods, we also pursue an integration and experimentation design that will provoke exceptions over long durations.

OBJECTIVES

The project is organized into three research tasks. The first two tasks pursue vehicle and system level adaptation, respectively. The third task develops an integrative experimental framework toward evaluating the approaches developed through the first two tasks.

Task 1: Introspection for Robust Autonomy: Typical robot systems are only robust to known contingencies, and therefore cannot adapt to unknown contingencies. Our goal is to be able to react to unknown exceptions through adaptation in the configuration space of sensors, processes, and actuators. Processes represent sets of viable algorithms that have varying applicability, such as a set of visual odometry algorithms for monocular and stereo situations, that will be automatically selected through introspection to enable robust autonomy and reaction to unknown contingencies. We consider two research directions to achieve this goal:

- Introspection of the effectiveness of sub-components and their parameters towards mission progress, and
- Configuration changes based on the effectiveness of sub-components of the system.

Task 2: Adaptive, Anticipative, and Resource-Aware Long-Term Deployments: Cooperative deployment planners typically generate vehicle deployments based on the current system load and individual vehicle capabilities with computational complexity that scales exponentially with the length of the planning time horizon. For this reason, these methods do not capture forecasted changes in vehicle efficacy due to physical degradation and dynamic environment conditions. We consider two research directions to address these limitations:

- Adaptive vehicle deployments based on mission-level requirements and evolving environment and system models, and
- Resource-aware deployment scheduling based on anticipated resource demands and current and forecasted environment and system models.

Task 3: Persistent Deployment in a Contingency Provoking Environment: Real-world evaluation of adaptive deployment and contingency aware methods is difficult as events can be rare. Our goal is to advance the state of the art in evaluating long-term operation in autonomous systems by constructing a challenging, but feasible, experimental design that frequently provokes contingencies and requires continuous adaptation. To achieve this goal, we are developing an experimental testbed that permits:

- Principled evaluation of long-term autonomy approaches in a controlled but unpredictable environment, and
- Benchmark comparison of various approaches to long-term operation.

APPROACH

Task 1: Introspection for Robust Autonomy
Researchers: Sebastian Scherer (PI), Kristen Holtz (CMU Graduate Student)

Autonomous aerial vehicles are often desired to perform tasks that are dangerous or even impossible for humans. From urban search and rescue missions to remote exploration of nuclear disaster sites, these tasks often take UAVs to unknown environments that introduce challenges due to their dynamic and diverse nature. Among these challenges is the likely inability of external communication, including availability of global positioning systems. This requires all perception, processing, and decision making to be made onboard. The unpredictability of the environment further contributes to the need for a more robust system, capable of recovering from unanticipated faults. It is infeasible to consider all unforeseeable events and develop a system to accommodate these events, which leads to a need for active fault-tolerant control (FTC) [Zhang, 2013].

FTC in perception can pose an especially difficult problem. It has been shown that autonomous aerial systems have difficulty navigating diverse environments when relying on a single sensor and require adaptation to move between various environments. One study has established the difficulty in perception while navigating between indoor and outdoor locations, and utilized a unique sensor and state estimation algorithm for each environment [Tomic, 2012]. The suitability of the environment to a given perception configuration is determined by evaluating the variance in state estimation for each method. Another study experimented with detecting and avoiding sensor failure by having several different algorithms, each reliant on a different subset of sensors, essentially making this a redundancy management problem [Berbra, 2008]. The most reliable subset of sensors was determined via

comparison to the expected sensor output. It has also *been found* that the reliability of sensors can be determined based upon context, using a mixture of experts framework to decide when to switch between sensors online [Ravet, 2013]. These solutions rely on running several sensors, and often several algorithms, in parallel, comparing them to determine which is most reliable.

In this study, we look at the fault-tolerant perception method problem when limited to only one RGB-D sensor input. As explored in [Handa, 2014], different algorithms developed for an RGB-D sensor can be more suitable in some environments than others, resulting in variable performance based upon the context.

We created a system able to determine which visual odometry algorithm is most suitable given RGB-D sensory input and switch between them. Rather than manually tuning a mapping between characteristics of the environment and which visual odometry method to choose we learnt a mapping from previous data. This method is able to not only run online, but also to re-evaluate switching parameters online if confronted with poor or unexpected results. We do this first by extracting a feature vector from the available RGB and depth images as well as previous output of the visual odometry algorithms. Then we used methods described in [Sturm, 2012] to evaluate trajectory error, and train a random forest regressor to predict the trajectory error of a given visual odometry algorithm, given a feature vector (See Figure 1). The findings are submitted in [Holtz, 2015].
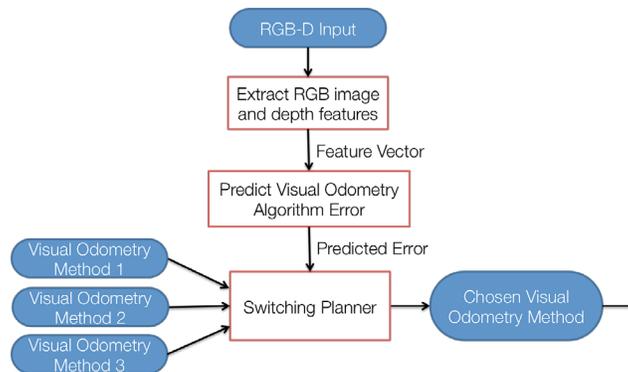
```
                    ┌─────────────┐
                    │ RGB-D Input │
                    └─────────────┘
                           │
                           ▼
              ┌──────────────────────┐
              │  Extract RGB image   │
              │  and depth features  │
              └──────────────────────┘
                           │  Feature Vector
                           ▼
              ┌──────────────────────┐
              │ Predict Visual Odometry │
              │    Algorithm Error   │
              └──────────────────────┘
                           │  Predicted Error
                           ▼
┌─────────────┐
│Visual Odometry│
│   Method 1   │──┐
└─────────────┘   │
┌─────────────┐   │   ┌──────────────────┐     ┌──────────────┐
│Visual Odometry│──┼──▶│ Switching Planner │────▶│ Chosen Visual │──▶
│   Method 2   │   │   └──────────────────┘     │Odometry Method│
└─────────────┘   │                             └──────────────┘
┌─────────────┐   │
│Visual Odometry│──┘
│   Method 3   │
└─────────────┘
```

Figure 1: Architecture for the adaptive visual odometry system.

Task 2: Adaptive, Anticipative, and Resource-Aware Long-Term Deployments
Researchers: Nathan Michael (PI), Derek Mitchell (CMU Graduate Student), Tiffany Bassey (CMU Undergraduate Student), Micah Corah (CMU/RPI Undergraduate Student), Vibhav Ganesh (CMU Undergraduate Student)

Multi-vehicle deployment algorithms that leverage optimization-based techniques provide a means to consider both system resources and capabilities while generating optimal vehicle plans toward applications such as surveillance and environmental monitoring. Through the introduction of vehicle energy constraints, these methods can be extended to consider long-duration deployments with an endurance that exceeds a single power cycle [Stump, 2011]. However, approaches that do not consider evolving task cost are unable to adapt to time-varying platform and environment models. We are developing deployment algorithms that adapt to available system resources (number of robots, power levels) and vehicle capabilities ascertained through introspection (physical health, task suitability). To

enable long-term planning, we pursue multi-resolution and approximate approaches that permit planning over long horizons (e.g., days, weeks, months). These methods will provide insight into resource sufficiency given the required load and anticipated effects from forecasted vehicle and environment changes.

The initial problem is formulated as a Multi-Robot Persistent Coverage Problem (MRPCP). The MRPCP considers a routing problem where each vehicle has a limited fuel capacity but can refuel at fuel depots placed in the environment. We introduce an objective function that minimizes the maximum period of tour cycles, as in [Turpin, 2014], placing more focus on an even distribution of target visitation so that no target is left unvisited for an extended period of time. The tour is computed online as the cost changes between cycles.

The general Vehicle Routing Problem (VRP) is NP-hard and so we employ (1) heuristic methods to provide suboptimal routes and (2) formulate the problem as a Mixed Integer Program (MIP) and solve to optimality using standard Branch-and-Bound techniques. MIP formulations have the advantage of providing optimal solutions, but are often unsolvable in a reasonable amount of time for any configuration of significant size. We pursue a heuristic formulation that provides a bounded suboptimal (approximate) solution with polynomial complexity. Building on the solution technique proposed by [Sundar, 2014], the heuristic solution provides an initialization point for a Branch-and-Bound Mixed Integer Linear Program (MILP) optimization formulation with mild extensions toward improved online performance.

The proposed formulation addresses dynamic and unknown costs through the dropping of targets based on current and expected traversal costs. In the MILP formulation, we solve for a set of variables that determine the minimum fuel requirement at each target necessary to continue along the tour. If the fuel level of the robot drops below this level, subsequent targets are dropped until the tour becomes feasible again. Once the cycle is complete, we re-compute given the measured cost and adjust the tour accordingly. Targets are reassigned based on individual vehicle capacity and learned transition costs. The present formulation introduces transition costs based on power consumption but is readily extended to consider alternate metrics such as perceived system health through introspection. The approach is detailed in a submitted conference publication [Mitchell, 2015]

Task 3: Persistent Deployment in a Contingency Provoking Environment
Researchers: Integrative task with researchers from Tasks 1 and 2.

We are evaluating and benchmarking our introspection and adaptive deployment algorithms in a long-term autonomy context through a novel testbed at the Field Robotics Center in an aerial vehicle persistent coordination context that requires constant adaptation and provokes exceptions. Preliminary progress toward the objectives of this task is detailed below.

WORK COMPLETED

Task 1: Introspection for Robust Autonomy
We completed the creation of a framework to enable adaptive introspection and were able to apply this framework to a set of datasets from visual odometry algorithms to create a switching paradigm. We developed a novel switching approach that resulted in significantly less errors on challenging datasets.

The results are reported in [Holtz, 2015] and correspond to the successful completion of the Task 1, Year 1 milestone. In year 2 we will continue with the approach and push towards generalization of the approach towards motion planning.

Task 2: Adaptive, Anticipative, and Resource-Aware Long-Term Deployments: From the project start date, the work completed resulted in the following contributions:

- MILP formulation of the MRPCP, a routing problem that favors even distribution of target visitation given fuel constraints;
- Heuristic algorithm to solve the MRPCP to within a proven optimality bound; and
- Method of dropping and recovering targets with repeated MRPCP solutions.

The results are reported in a submitted conference publication [Mitchell, 2015] and correspond to the successful completion of the Task 2, Year 1 milestones as noted in the project proposal.

Task 3: Persistent Deployment in a Contingency Provoking Environment: Through the use of leveraged hardware and capabilities developed through related programs, preliminary experimental evaluation of the techniques proposed in Tasks 1-2 were conducted, including:

- An online implementation and test of the algorithm on a set of datasets collected on the ex-USS SHADWELL. Results that demonstrate the feasibility of the approach.
- Persistent deployment algorithm evaluation on a team of four specialized micro air vehicles with onboard capabilties to sense power usage during operation and autonomously recharging using recharging stations. The deployment algorithm adapts to the changing traversal cost based on the observed power usage [Mitchell, 2015] and represents a preliminary version of the experiments proposed for Task 3, Year 3.

RESULTS

Task 1: Introspection for Robust Autonomy
This adaptive odometry switching method was implemented on a system and first results are shown from datasets collected in challenging environments collected onboard the ex-USS SHADWELL (See Figure 2).

We evaluated the odometry switching planner on 5 min run moving through areas with motion blur, image contrast beyond the camera limit, and dark images and the approach was able to pick the correct method. As can be seen in Figure 3 the total error at the end of the complete path is less because the method switches the algorithms in the cases were they make large errors.

Figure 2: Example images from the RGB-D dataset used for visual odometry. This data provokes frequent failures of visual odometry methods because of difficult lighting, and motion blur, for example.
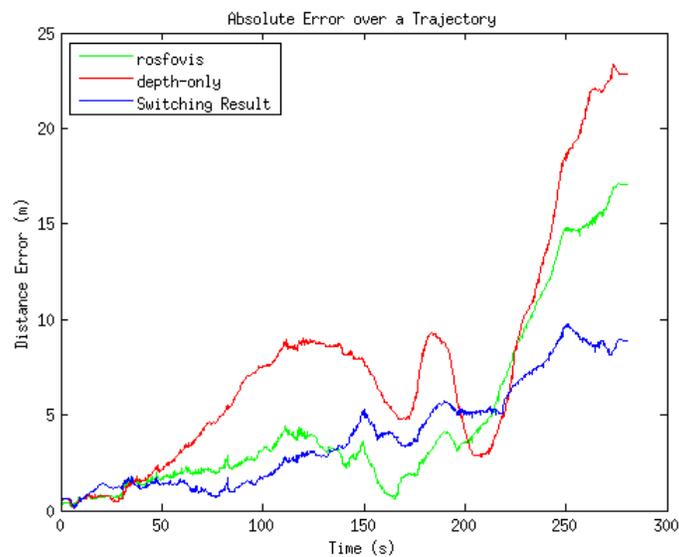

Figure 3: Absolute odometry error for two different visual odometry methods (Fovis, and Depth-only odometry) and the result of switching the methods in the environment shown above. The adaptive method is more robust and can therefore outperform the other methods.

Task 2: Adaptive, Anticipative, and Resource-Aware Long-Term Deployments: The key result of the proposed approach is the successfully execution of deployment plans that adapt to changing cost over multiple system power cycles in simulation and partially through experimentation.
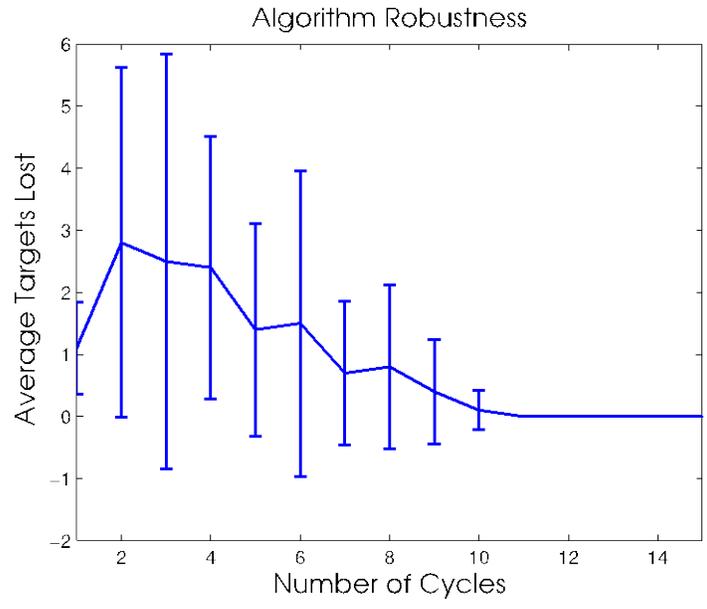
6

Figure 4: Number of dropped targets as a function of cycles completed

In Figure 4, we show that the initial number of dropped targets can be significant when the travel costs are unknown and that online cost learning and adaptation results in fewer dropped targets.
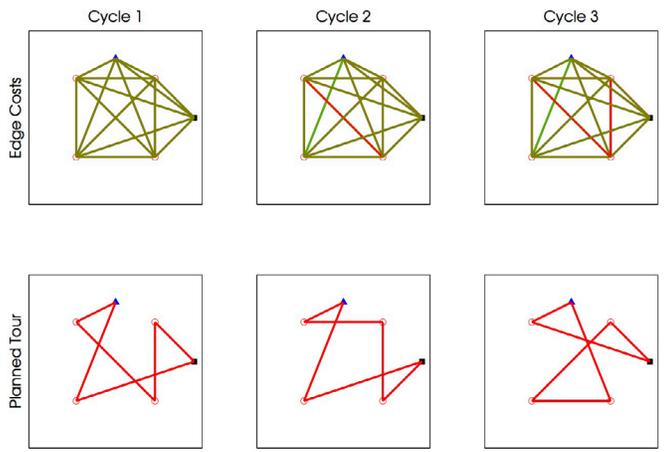

Figure 5: Planner reconfigures tours with appearance of lower cost edges (cycle 1, left; cycle 2, middle; cycle 3, right).

As the transition costs adapt and approach the true values, edges with lower traversal cost result in reconfigured tours, balancing the resource consumption between multiple systems (Fig. 5). As the costs stabilize, the solution approaches the optimal solution through formulation. If costs do not change, the time reserved for computing each plan in subsequent cycles is reallocated to refine previous solutions. The solution will evolve over time through continual refinement of task modeling.

Task 3: Persistent Deployment in a Contingency Provoking Environment: Several capabilities were developed this year toward the planned integrative evaluation.

7

- Autonomous flight capabilities (perception, planning, and control) for micro air vehicles operating in general indoor environments.
- Hardware and software capabilities to enable online deployment of multiple systems and online recharging and redeployment toward continuous experimentation.
- A novel micro air vehicle (hover mass of 350 g) with onboard perception, computation, and recharging capabilities amenable to large scale, many robot experiments.
- A set of algorithms to enable redundant architectures.

IMPACT/APPLICATIONS

There is a great push towards unmanned vehicles of all kinds (water, air, and land) in the Navy and their greatest benefit is in their potential advantage of not requiring human supervision for long durations to take over tasks such as patrol and monitoring. Current systems are frequently close to the operator and require frequent intervention during off-nominal or complex situations. However, if manning of these assets is necessary, the benefits of autonomous systems cannot be reaped.

The methods we are developing in this project will enable robust, efficient, and scalable autonomous systems that can operate on the order of months or years by going beyond the state of the art in todays autonomy methods. This will be a force multiplier to the Navy as it will enable the operation of many more assets than operators.

Our contributions to the fundamental way in which robots adapt themselves to unforeseen circumstances will enable unmanned vehicles to operate over periods of months by adapting to failures in the lower-level processes and learning the capabilities of the assets. This will enable faster deployment and reaction to new situations for which the unmanned assets were not initially developed.

RELATED PROJECTS

Capabilities developed toward Task 3 are largely funded through other programs sponsored and in collaboration with the Army Research Laboratory and Office of Naval Research (in addition to this effort). An brief overview of each program follows.

- Micro Autonomous Systems and Technology (MAST) CTA program funded by ARL develops autonomous, multifunctional, collaborative ensembles of agile, mobile microsystems to enhance tactical situational awareness in urban and complex terrain for small unit operations (see mast-cta.org).
- Absolue Localization in GPS-Denied Environments for Autonomous Unmanned Ground and Micro-Air Vehicle Systems (N121-101-0336, PM: Tom McKenna, Prime: Sensible Machines Inc.). We leverage the micro air vehicle developed in this project and datasets collected during challenging missions to give motivating scenarios for the techniques developed in this project.

REFERENCES

[Berba, 2008] Berbra, C., S. Lesecq, and J. J. Martinez. "A multi-observer switching strategy for fault-tolerant control of a quadrotor helicopter." *Control and Automation, 2008 16th Mediterranean Conference on*. IEEE, 2008.

[Handa, 2014] Handa, A, et al. "A Benchmark for RGB-D Visual Odometry, 3D Reconstruction and SLAM". *ICRA*, 2014

[Ravet, 2013] Ravet, Alexandre, et al. "Learning to combine multi-sensor information for context dependent state estimation." *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*. IEEE, 2013.

[Stump, 2011] E. Stump and N. Michael, "Multi-Robot Persistent Surveillance Planning as a Vehicle Routing Problem," presented at the Proc. of the IEEE Conf. on Autom. Sci. and Eng., Trieste, Italy, 2011, pp. 569–575.

[Sturm, 2012] Sturm, Jürgen, et al. "A benchmark for the evaluation of RGB-D SLAM systems." *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*. IEEE, 2012.

[Sundar, 2014] K. Sundar and S. Rathinam, "Algorithms for Routing an Unmanned Aerial Vehicle in the Presence of Refueling Depots," *TASE*, vol. 11, no. 1, pp. 287–294, Jan. 2014.

[Tomic, 2012] Tomic, Teodor, et al. "Toward a fully autonomous uav: Research platform for indoor and outdoor urban search and rescue." *Robotics & Automation Magazine, IEEE* 19.3 (2012): 46-56.

[Zhang, 2013] Zhang, Y. M., et al. "Development of advanced FDD and FTC techniques with application to an unmanned quadrotor helicopter testbed." *Journal of the Franklin Institute* 350.9 (2013): 2396-2422.

PUBLICATIONS

[Holtz, 2015] K. Holtz, S. Scherer, "Learning Metohd Context to Enable Robust RGB-D Odometry," submitted to the IEEE Intl. Conf. on Robot. and Autom., Seattle, WA, 2015. [submitted, refereed]

[Mitchell, 2015] D. Mitchell, M. Corah, N. Chakraborty, K. Sycara, and N. Michael, "Long-Term Persistent Coverage with Micro Air Vehicles in Dynamic Environments," submitted to the IEEE Intl. Conf. on Robot. and Autom., Seattle, WA, 2015. [submitted, refereed]

[Turpin, 2014] M. Turpin, N. Michael, and V. Kumar, "An Approximation Algorithm for Time Optimal Multi-Robot Routing," presented at the Proc. of the Intl. Workshop on the Algorithmic Foundations of Robot., Istanbul, Turkey, 2014. [published, refereed]