

# EUROPtus: A Mixed-Initiative Controller for Multi-vehicle Oceanographic Field Experiments

Frédéric Py<sup>1</sup>(✉), José Pinto<sup>1</sup>, Mónica A. Silva<sup>2</sup>, Tor Arne Johansen<sup>3</sup>,  
João Sousa<sup>1</sup>, and Kanna Rajan<sup>1,3</sup>(✉)

<sup>1</sup> Underwater Systems and Technology Laboratory, Faculty of Engineering,  
University of Porto, Porto, Portugal

`fredpy@gmail.com`, `{zepinto,jtasso}@fe.up.pt`

<sup>2</sup> IMAR-Açores & MARE – Marine and Environmental Sciences Center,  
Horta, Portugal

`monica.silva.imar@gmail.com`

<sup>3</sup> Department of Engineering Cybernetics,  
Center for Autonomous Marine Operations and Systems (AMOS),  
Norwegian University of Science and Technology, Trondheim, Norway  
`{tor.arne.johansen,kanna.rajan}@ntnu.no`

**Abstract.** Our research concerns the mixed-initiative coordination of air and underwater vehicles interacting over inter-operated radio and underwater communication networks for novel oceanographic field studies. In such an environment, operating multiple vehicles to observe dynamic oceanographic events such as fronts, plumes, blooms and cetaceans has required that we design, implement and operate software, methods and processes which can support ephemeral and unpredictable observations (including those of moving animals) in real-world settings with substantial constraints. We articulate an approach for coordinated measurements using such platforms, which relate directly to task outcomes. We show the use and operational value of a new Artificial Intelligence (AI) based mixed-initiative system, **EUROPtus**, for handling multiple platforms from a recent field experiment in open waters of the mid-Atlantic.

**Keywords:** Marine robotics · Oceanography · Artificial Intelligence · Mixed-initiative control

## 1 Introduction


Recent advances in robotic vehicles have made ocean observation more sustainable with the use of autonomous and semi-autonomous platforms to observe at varying spatio-temporal scales. The principal challenge however, is to observe the water column not just with point-based observations as in traditional oceanography, but across the mediums of air and water, and the air/water interface,

and doing so continuously. Such observations need not only to be *synoptic*, but also to be coordinated across space and time to requiring coordination and control of a range of robots with appropriate sensors, which necessitates the use of multi-platform systems to observe in the meso-scale ( $>50 \text{ km}^2$ ), and to follow phenomena of interest such as blooms, plumes, anoxic zones, fronts over a period of days, weeks or longer. Our experimentation [1–5] has led to the conclusion that multiple vehicles, operating in aerial, surface and underwater domains, are critical for such observational needs especially in the study of our evolving planet.


Our work therefore, deals with the operation of networked heterogenous robotic platforms while dealing with unreliable communications in the harsh conditions of the open ocean. Uncertainty in sensing, control and actuation and operational unpredictability is the norm, making operations of such platforms challenging in this environment. Further, the target applications to date, often involve dynamic and unpredictable phenomenon in unstructured environments with the use of autonomous underwater vehicles (AUVs) as well as unmanned aerial vehicles (UAVs). UAVs offer speed and the agility for synoptic observations; AUVs can then ground-truth and collect in-situ data with an array of sensors in the water column. Put together, autonomous robotic operations are key for sustained oceanographic observations in the light of such operational constraints.

This work concerns the mixed-initiative coordination of air and underwater vehicles interacting over inter-operated radio and underwater communication networks for oceanographic field studies. In such an environment, operating multiple vehicles to observe dynamic features, including motion of cetaceans or oceanographic phenomena such as fronts, plumes and blooms has required that we design, implement and operate software, methods and processes which can also support opportunistic goals amid tight operational constraints in real-world settings. Robot tasks deal with coordinating and completing observation and sampling tasks in the air, surface or underwater domains for observing the same patch of the ocean co-temporally.

We articulate an approach for coordinated measurements using such platforms, which relate directly to task outcomes. We show the use and operational value of a new Artificial Intelligence (AI) based mixed-initiative system,



**Fig. 1.** Typical setup of operators on ship/shore for vehicle control in oceanographic field operations.



**Fig. 2.** AUV (left) and UAV (right) launch operations from the NRP *Gago Coutinho* in the Açores, July 2015.

EUROptus, for handling multiple platforms (Fig. 2). Coordination in this context implies the ability to envision task completion in the light of dynamic surroundings while dealing with operational constraints. While coordination in marine robotics has typically been viewed as a means to demonstrate nominal engineering principles [6, 7], our work is focused towards solving specific oceanographic needs. Our robotic platforms are tied together with a mature set of software tools for decision support, planning, control, data visualization and archiving provided by a toolchain [8].

In operational oceanography, traditionally a graphical interface is used as a planning tool to offer decision-support capability with humans making all decisions a priori for a very uncertain, harsh and dynamic setting (Fig. 1). Our objective is to augment such methods with automated planning [9] to allow a single operator to command multiple heterogenous vehicles in real-world settings.

EUROptus is a mixed-initiative constraint-based automated planner which aids such field campaigns in coordination to ensure that operational goals are satisfied during mission operations. The planner does not model all constraints comprehensively, but is used primarily for task assignment, information gathering and situational awareness of AUVs, all the while keeping a simple resource model of UAV operations to model when tasks need to be outsourced and operators need to be engaged. The application domain involved using UAVs to spot cetaceans providing a GPS fix on their coordinates as a means for targeted in-situ measurements with AUVs. While EUROptus was designed and field tested in the specific operating environment for our experiment at sea in the waters off of the Açores<sup>1</sup>, we believe the principles behind it are general.

This paper is organized in the following manner. We briefly place the context of this work in Sect. 2. We start by describing the scientific domain in cetacean tracking in Sect. 3 and then move to the core of the paper in Sect. 4 describing technical details of EUROptus's representation and reasoning mechanism. We describe the tools, techniques and process we use in Sect. 5 and its use in the

<sup>1</sup> <http://rep15.lsts.pt/>.

experiment in Sect. 6. We conclude with lessons learned and some ideas for future work in Sect. 7.

## 2 Related Work

AI based mixed-initiative methods for planning, continue to be novel in the oceanographic domain. In recent work [10] has used an automated planner for controlling multiple AUV; **EUROPtus** moves beyond this work in multiple ways. First, the underlying plan representation allows for replanning and reasoning about time and resources. By using the identical plan representation and reasoning formulations on ship/shore with an embedded controller, we can seamlessly transfer partial plans for commanding a vehicle. Second, **EUROPtus** deals with heterogenous vehicles and their operational complexities. Finally, we use active simulation coupled with the expected execution trace while keeping the operator situationally aware.

In **EUROPtus** we are informed by efforts on the command/control of the *Spirit* and *Opportunity* rovers on Mars [11, 12] using a similar mixed-initiative approach. The operating domain at sea however, is harsher and substantially more constrained with a more dynamic pace and operational fluidity. Using automated planning as a means to provide abstraction in control over single or multiple vehicles in the oceanographic domain, continues to be novel.

Robotic vehicles, UAVs in particular, have been used for studies in animal ecology in the open ocean [13]. However, our work here, more closely supports upper water-column biology as a side effect of cetacean transit.

## 3 The Scientific Domain

The overall scientific objectives of tracking cetaceans in the open ocean are driven primarily by the need to understand the drivers behind movement behaviour of large marine predators; in particular to collect data to interpret how ecological conditions, including inter-and intraspecific, environmental, trophic and social interactions shape their decisions. The open-ocean and specific habitat features, such as seamounts, fronts, eddies, clines, and other provide the oceanographic phenomenological context within which to study such animal behavior. The Açores, in particular, presents the ideal set of conditions to pursue this goal, bringing together multi-taxa opportunities, easy access to those habitats, access to facilities and vessels as well as a team working on top predator biotelemetry with extensive field and data analysis know-how. The long-term concept to obtain such observations, amounts to near real-time 4D (space X time) tracking of large marine predators and synoptic, dynamic prey and oceanographic sampling. The concept reverses the current paradigm on oceanic observatories: instead of having a static array of equipment fixed at a locality that is only effective when the animals wander inside its influence area, the concept would be to have a dynamic laboratory that can be deployed around and follow focal

individuals/groups to collect the relevant data in the surrounding environmental ‘bubble’.

Specific scientific objectives in this domain include, understanding the physical and biological mechanisms promoting prey abundance and aggregation, understanding the relationship between different prey field properties (e.g. overall distribution and density, patch patterns) and the cetaceans’ foraging decisions and determining whether environmental properties measured in upper water column can be useful at predicting prey properties and foraging behavior at depth.

In addition to environmental factors, foraging decisions of animals are likely to be affected by multiple ecological and intrinsic factors. Typical cetacean tracking scenarios can be combined with the collection of other data that may provide new insights into the role of some of these factors. For example, collection of biopsy, fecal, and blow samples of tagged whales to do molecular sexing [14] and to assess their reproductive and nutritional state. Still images taken by an AUV to measure body length and width of surfacing whales, could serve as a proxy for age and nutritional status. Intra and interspecific competition may also constrain foraging strategies of animals, for instance, by limiting subdominant species/individuals to habitats with lower prey availability or by changing preys behavior, distribution or aggregation patterns. A first step towards understanding the influence of density-dependent factors would be conducting surveys simultaneously with the whale tagging and environmental sampling. Doing this with traditional methods would be prohibitively expensive but AUVs and UAVs offers a unique opportunity with minimal costs and personnel, to obtain data on the total number of whales foraging on the study area, allowing us to examine if and how whale abundance influences prey patterns and whale behaviour.

Together, these provide the relevant rationale and entry point for robotic vehicles, including AUVs, UAVs as well as ASVs (autonomous surface vehicles) as a viable set of observational tools for such studies. However, the science needs and intent outstrip the available technology at this stage, especially in targeting fast moving cetaceans and at depth. This experiment was therefore conceived as a starting point for long-term research and inter-disciplinary collaboration with a unique set of constraints, driven by the key scientific objective to *characterize the upper water-column environment*, primarily to ask the question “Why are cetaceans foraging at *this* location?”. This work also builds on our previous efforts in tracking much smaller fish in space and time [3, 15].

## 4 Technical Approach

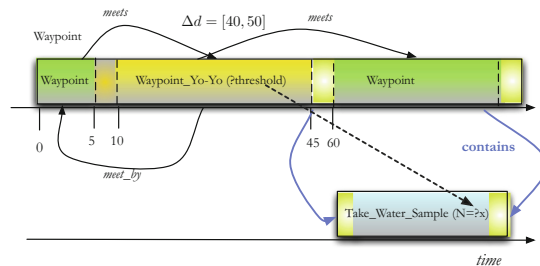
Mixed-initiative methods are used to provide a human operator situational awareness and commanding capability for robotic vehicles where the mapping is one-to-one between operator and robot. Our aim is to simplify further by offloading the operator’s cognitive burden especially in coordinating *multiple* heterogenous vehicles. We do so leveraging constrained-based temporal planning [9]. While the focus is on human-in-the-loop interaction, our system interacts with the T-REX plan-execution system embedded onboard our AUVs [16–18].

Both EUROptus and T-REX rely on the same rich representational formalism for plan synthesis; we outline this briefly. Additional details can be found in [16–19].

#### 4.1 Representational and Planning Framework

Traditionally robotic execution has relied on dispatching commands at precise times with an executive. Such linear sequences of precisely timed commands give no ability to adjust execution on the basis of sensory information. The consequences of the intrinsic inflexibility of such sequences are critical; because they are inflexible, sequences are brittle and therefore must necessarily be designed considering worst case scenarios. EUROptus’s representation significantly broadens the way robots can be commanded [20] with the interpretation of a *temporally flexible plan* which represents each start time as a flexible timepoint variable backed by an explicit network of bounded delay constraints between such temporal variables. These timepoints [21] represent temporal intervals signifying a change in state and are of the form  $[lb, ub]$  (where  $lb, ub \in \mathbb{N}$ ) are temporal lower and upper bounds respectively. Instead of specifying a fixed integer, time is represented as an interval (see Fig. 3) leaving room for adaptation at execution time while representing uncertainty (in outcomes or the environment). When the executive considers when to start a task, it propagates information through the constraint network, computes a time bound for the variable, selects an actual execution time within the bound, and starts the task at that time. Temporally flexible plans therefore, express a *range of possible outcomes* of the robots interaction with the environment within which the executive can elect at run time. The fact that constraints are explicitly represented ensures that through propagation the executive will respect global limits expressed in the plan (e.g., don’t start a task until a certain condition has been satisfied) while still satisfying some overall deadline.

A plan is composed of temporally scoped predicates called *tokens*. A token can be described as a first order logic predicate, with its associated temporal scope (*start, duration, end*) using flexible interval arithmetic [22]. All the attributes of a token are described as a domain of possible values for this token in the plan context. In order to be part of the plan, a token needs to be associated with one timeline only. A *timeline* is a sequence of tokens describing the evolution of a state variable while enforcing a strict sequence of tokens with no concurrency within the



**Fig. 3.** Tokens with flexible temporal intervals and parametric constraints between tokens. This example shows the triggering of an AUV’s water sampler based on a feature threshold while the vehicle Yo-Yo’s. The *Waypoint\_Yo-Yo* token has a flexible duration, start & end times.

only. A *timeline* is a sequence of tokens describing the evolution of a state variable while enforcing a strict sequence of tokens with no concurrency within the

timeline. Concurrency between timelines and therefore between tokens on separate timelines, is the basis for concurrent state variable evolution. Tokens are causally linked by rules in a *domain model* that describe temporal relations and/or causality links between tokens [16,23]. Finally a token can be marked optionally as either being a *Fact* or *Goal*. While a *Fact* requires no justification, a *Goal* will need not only to be inserted in the plan but necessarily have a causal chain connecting it to one or more *Facts*. The underlying planner for EUROptus, is EUROPA [19,23]; while the key concepts are not tied to this specific planner, our implementation relies heavily on its flexible temporal representation as also the basic principle of searching in plan-space [9].

The planner works by continuously repairing *flaws* in a plan until no more flaws are present; for execution, all partial plans must not have any flaws. Typically we deal with two types of flaws: an *open condition* where a token is not associated with a timeline and can be resolved by either inserting it into a specific timeline, or merging the token with one that is compatible. Or a *threat* where an inserted token may impact others indirectly through possible overlapping requirements. The plan solver then needs to enforce a scheduling constraint on those potentially conflicting tokens so they cannot overlap.

The solver resolves these flaws until either it reaches an inconsistency (i.e. a situation where plan constraints cannot be satisfied), in which case the planner will backtrack to explore alternate solutions; or a consistent solution is found and the plan presents no further flaws generating a valid solution. Additionally, while the plan might be complete, it does not have to commit to the value of its variables. For example, the start time of a token can be left to be the interval [1, 10] as long as it does not present a threat to the partial-plan. This leaves the decision of the start time to the executive which is critical for operating in uncertain real-world environments. Details of the solver and the planning engine are outside the scope of this paper and can be found in [18].

## 4.2 Planning and Execution with Asynchronicity

Typically in fully autonomous systems such as T-REX, planning and execution are intertwined. Both manipulate the same plan representation and plan execution is required to occur at every clock cycle. This influences in turn, the outcome of the planning process while ensuring that any plan produced by the agent is taking into account state evolution with the assumption that the agent is within a synchronous and fully observable world. Such a design, while appropriate for embedded systems, is incompatible with communication challenged field operations where observations are often asynchronous in nature, especially where coordination with humans on launch and recovery operations can be complex. Further, observations from the robots can arrive sporadically, as with our AUVs, and observability is limited by the availability of an acoustic channel to the vehicle.

Execution can still be integrated into a deliberation process such as planning; but we go a step further in considering that *execution is a part of planning*. Execution feedback is integrated into deliberation by taking advantage of the

way the planner works; the planner will stop searching as soon as no more flaws are found and the introduction of a new token (including *Facts*) in the plan create new flaws in the plan.

*Example:* Consider that the planner has no more flaws and that the next command to be executed has been already dispatched (see below). Because of communication latencies, which may result from intermittent connectivity, we receive feedback from AUV<sub>a</sub> that indicates that its state changed from *Inactive* to *Operating* an hour ago. Our approach is then to create the *Fact* token *Operating* for AUV<sub>a</sub> starting at the time corresponding to an hour ago. This token is added to the plan generating a new *Open condition* flaw that the solver needs to resolve.

The resolution can be either as simple as a token merge, if the token reflects exactly what was planned, or it is conflicting with the partial plan, requiring in turn, the planner to backtrack. The insertion of such a token as a *Fact* however, is akin to plan recognition, i.e. an agreement with the currently maintained partial plan structure. As the planner operates continuously keeping its search state alive, this recognition can impact the search by forcing the planner to backtrack over past decisions until it finds an alternative new solution given the injected *Fact*. As long as we assume that the decisions impacted by the new observation are not close to the root of the search tree it allows for a plan resolution in few steps without an adverse impact to performance. Further, such an assumption is reasonable as often past observations arrive in relative chronological order rarely impacting the distant past in plan history. These steps also highlight how execution tracking can be a pure deliberative task within the same planning engine.

A remaining problem is to decide which part of the plan is ready to be executed and sent to a vehicle via NEPTUS, a visualization, planning and situational awareness tool. Many actions in a partial plan can be conditioned by the need to observe a situation. Actions are specific tokens introduced in EUROPA, that instantiate a causal relationship between their *condition* and *effect* tokens. It reflects the classical approach to describe a plan domain [9] but is substantially more expressive and allows its use for additional semantics. An action is a special kind of token with temporal relations expressed either as *conditions* necessary for its execution, or expected *effects* of this actions [18].

To command AUVs for a survey we needed EUROPTUS to observe that: an AUV is ready and in the *Inactive* state and that we have a cetacean position update that is at most 30 min old; this heuristic was imposed to ensure “freshness” of cetacean tracks. When deciding on dispatching a partial plan, EUROPTUS does an analysis of the causality structure of an action with its tokens. It does so by introducing the notion of a *Justified* token as: a token is *Justified* if either it is a *Fact*, the condition of a *Justified* action or it is an action for which all the effects are *Justified*.

In a complete plan, an action in the plan can be dispatched for execution when all its conditions are *Justified* and its start time interval contains the current time. While action justification is reasonable, we need to ensure that we do not



dispatch the action before its valid start time. However, we could potentially be in a situation where an action does not have all its justified conditions due to one or more missing messages and yet its start time is before the current time (i.e. it *should have* been dispatched for execution in the past). Dispatch is then postponed and we address this as a *Pending action* flaw, which applies to any action of the plan that could start at the current time but does not have all of its conditions *Justified* nor is it *Justified* itself. Its default resolution is to restrict the start time to be postponed. Consequently the planner needs to postpone dispatching this action until either new observations justify the action or the start time can no longer be pushed; the latter triggers a backtrack for an alternate solution.

Any pending action that has all of its conditions *Justified* is dispatched to NEPTUS for execution which eventually will receive the observation of its completion (from the vehicle) and report it to EUROPtus. This results in a control loop that is managed as a pure continuous deliberation process governed by the principles just described.

The position update then comes into play for the AUV operations model. As the AUVs are driven by an embedded T-REX agent, EUROPtus could be further extended to directly leverage such a positional update for “direct” control of the vehicles. Yet the limitation in terms of communication had to be taken into account. Our AUVs can communicate with the ship only if they are at the surface and either in Wifi range (which might also be not desirable as the ship can present a threat if operated too close to the vehicle) or long enough to the surface so it can initiate a satellite connection. When the vehicle is underwater, it was assumed that there is no means to communicate<sup>2</sup>.

### 4.3 Mixed-Initiative Interaction

EUROPtus deals with *operational* constraints for field experiments. It leverages and builds on existing work in fully-autonomous AUV operations. It has a simple resource model which is integrated into the planner to enable as complex a coordination model as necessary. By using EUROPtus, the operator offloads a portion of the planning task. This is important when considering networks with very high variability of vehicle configurations and capabilities. However, the operator is still in charge of providing high-level goals, and supervising the plans sent to the vehicles.

When EUROPtus determines so, it can request new objectives for vehicles deployed in the field or prompt humans behind an operations console. For instance, operators are expected to execute some task and provide input (e.g. inspect collected data and determine a list of waypoints to be visited) or, EUROPtus generates a high-level objective that is sent to an autonomous vehicle to replan in-situ accordingly. If a new plan is found, this is reported to the consoles from which operators can both provide new objectives and inputs, and/or

---

<sup>2</sup> Note: acoustic communication was only used to track vehicle location.

recall existing objectives. This is the focus of such mixed-initiative interaction between vehicle(s) and operator.

## 5 Experiment Infrastructure

Our robotic hardware consisted of multiple Light AUVs [24] and customized Skywalker X-8 UAVs. The architecture for command and control is built on top of a mature toolchain [8], which has a back-seat driver API to the DUNE software which is in charge of navigation, logging and management of all communications onboard the vehicle. DUNE allows external controllers to provide desired poses for the platform while receiving progress updates on their attainability. This allows for the development of external controllers which are not tied to specific vehicle hardware allowing DUNE to use provided slack to improve vehicle safety, navigation or battery optimization. Using this API, embedded on the AUVs is T-REX [16, 18]; when it receives a goal, its objective is to synthesize partial plans and execute these while simultaneously monitoring execution onboard the vehicle. We use a publish-subscribe message-centric system, IMC [25], that is used for state updates all throughout the system, and to send commands and high-level objectives to the platforms. NEPTUS provides visualization, situational awareness and commanding with human operators. NEPTUS consoles consist of an empty canvas populated with visual widgets and map layers that reflect the data received and allow interaction with the network over IMC. These together are the components of the toolchain; EUROptus augments them with a shore/ship side component.

For this experiment we consider two types of assets; a UAV that is operated with human-in-the-loop waypoint-based control and AUVs running T-REX which can receive survey objectives via a timeline goal from EUROptus. The UAV launch and recovery operations involves substantial human involvement and its operation currently requires close monitoring. Consequently, the basis of interaction with EUROptus is simplified as follows:

- EUROptus can request a new deployment through NEPTUS directed to the UAV console operator
- the maximum UAV operation time is assumed to be 30 min given typical battery life
- replacing a UAV battery on shore/ship including recovery and new launch will take approximately 15 min between 2 surveys<sup>3</sup>
- a cetacean position update is expected to be the observed outcome of a successful UAV survey and when available generates a goal within EUROptus which is instantiated on a timeline

---

<sup>3</sup> Both the launch/recovery operation and charging times where approximate yet reasonable estimates.

EUROPtus approximately models UAV operational constraints and its interaction is indirect as it sends a message to the UAV operators console to trigger the operators response. Its observations similarly are driven by a human operator’s event when s/he manually identifies and “marks” a cetacean position on NEPTUS.

The EUROPtus AUV model is based on:

- A timeline representing the vehicle position updated whenever a position update is received and placed in the timeline according to its observation timestamp
- A set of possible surveys both parameterized with their scale in meters (representing the outer box surrounding the survey), its centroid (represented by a latitude and longitude) and its orientation (a rotation angle) [26]. From this information and the AUVs known nominal speed and location, EUROPtus is able to identify both the entry and exit location along with the completion time of this survey
- The high level operational state of the vehicle being either *Inactive*, *Operating* and *Survey* the latter which takes as an argument a fully instantiated survey as above

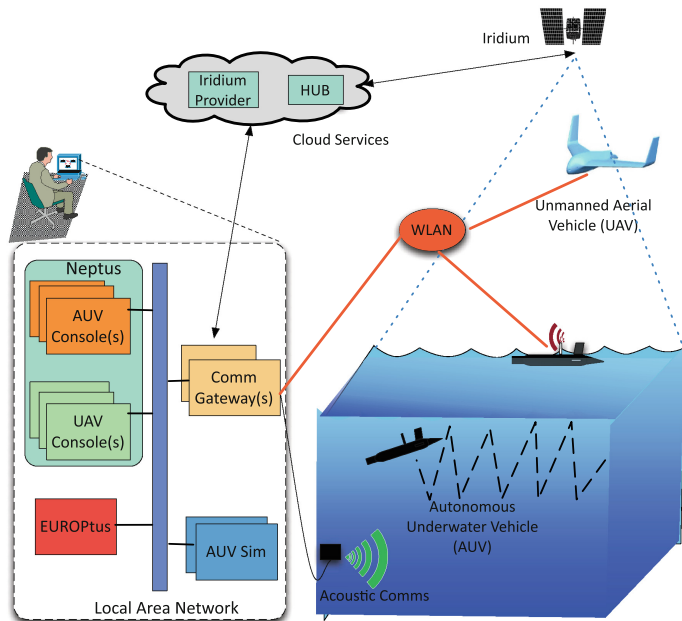
Typically within EUROPtus the AUV’s overall state cycles between *Inactive*, *Operating* and *Survey* in executing a survey. The *Operating* duration depends on the scale of the survey and the distance from the survey start point (which should be the last position observed when the vehicle is *Inactive* and on the surface. *Survey* being a goal state, its duration is short since it is used as a feedback confirming the successful completion of the survey request. This timeline – along with the AUV position – effectively produced by T-REX onboard, is the means of interaction between the AUV and EUROPtus on ship. Time-stamped messages from the AUVs, as noted, come with significant time delays. This means that the executive associated with EUROPtus should not only be able to dispatch the actions in a timely manner, but also to integrate observations from the past and, when required by the model, delay the execution of a specific action until its conditions are effectively observed.

Our experiment required having multiple vehicles disconnected for substantial (15–30 min) periods of time. In order to improve the operators’ situational awareness of vehicle positions and progress, we used a number of simulated vehicles running off commands sent to the actual platforms. As a consequence, the simulated vehicles execute the *identical* set of commands albeit in an idealized environment. While doing so and in periods of loss of contact, operators can determine with clarity, what each asset is *expected* to be doing. In the case of AUVs, any updates received over any available communication channel, are used to reset the simulated vehicle’s localization filter. Operators interact through different NEPTUS consoles which are adapted to mission and operator-specific needs.

EUROPtus was implemented to orchestrate the operational setting involving multiple assets. For this purpose, a special NEPTUS plug-in acts as a routing device for incoming and outgoing messages. The plug-in redirects data from EUROPtus to the controlled/simulated AUVs and, at the same time reports all received updates to EUROPtus; conversely this plug-in also notifies UAV operators about requests.

## 6 Experiment Setup


Our experiment was in open waters south of Pico Island in the mid-Atlantic Açorian archipelago (Fig. 5). All operations were carried out on board the Portuguese Navy research vessel, the NRP *Almirante Gago Coutinho* with AUV and UAV launch and recoveries from the aft deck; all operators and pilots were onboard the vessel. Multiple (sectorial and omnidirectional) antennas for providing 802.11 WiFi and Airmax coverage via ubiquity radios



**Fig. 4.** The operational architecture of EUROPtus used in the experiment in comparison to more typical approaches shown in Fig. 1.

on UAVs were provisioned with multiple communications gateways via a local area network onboard the vessel, allowing operators to be simultaneously aware of AUV and UAV operations (Fig. 4). The AUVs were equipped with RBR XR620 CTD (Conductivity, temperature and density) probe, WHOI acoustic modems, an Iridium SBD modem for satellite communication, an experimental Holographic imager and a Turner Designs Cyclops-7 wet-probe with a fluorometer. Skywalker X-8 UAVs at our disposal were equipped with Far-IR cameras, with one vehicle equipped with a new light-weight hyper-spectral imager.

Our scenario called for either an experienced human observer on a separate boat, or airborne UAVs with real-time imagers, to detect foraging cetaceans at the surface. As animals were moving, a reference point was determined either



**Fig. 5.** The open water location of the experiment 15 Nautical miles south of the island of Pico in the Açores.


from imagery or from the visual observer, which was communicated to the operators on the vessel. If AUVs were not already in the water, they were launched. Repetitive AUV surveys around this targeted spot were to commence, using two AUVs to measure spatial variations in the water column. Both T-REX-enabled AUVs were tasked by EUROPtus to synthesize yo-yo based survey patterns [26] in two concentric square patterns of  $400 \times 400 \text{ m}^2$  and  $800 \times 800 \text{ m}^2$  diving to a depth of 50 m. The vehicles also surfaced on the corners of the squares in order to localize with a GPS fix. The larger pattern was used to sample *outside* the foraging area, as a measure of understanding the variability in the upper water-column. The smaller survey was expected to take 25 min, while the larger 50 min at about 2.5 knots speed over ground for the vehicles. Consequently two inner surveys provided a dense coverage co-temporal to the single outer survey, (Fig. 6(a)). At the end of the survey, the next cetacean target was to be determined either via the visual observer or a UAV and the vehicles re-tasked by EUROPtus.

One extended objective was to attempt not only co-temporal AUV surveys, but to coordinate the survey of the sampling area with UAV overflights with Near-IR and hyper-spectral sensors. In doing so, it was thought, we could obtain additional science data of the ocean surface to be merged and subsequently studied to understand the bio-geochemical composition in the light of any effluents from the passing cetaceans. This was successfully achieved by sharing of mission plans between NEPTUS consoles and by coordinating the execution of plans via EUROPtus. Figure 6(b), shows vehicles in the same operating area, one airborne, one at the surface and one AUV underwater (connected acoustically). At the same time, multiple NEPTUS consoles were receiving and controlling the different vehicles from the ship. In EUROPtus we tie the AUVs operational model with that of the UAV to ensure coordination; UAVs if not in the air were launched from the aft deck, to be surveying the same area overhead.


Figure 7 shows EUROptus’s deliberation steps and depth along the first hour of the July 19th mission. This figure shows the number of *steps* increasing monotonically reflecting the integration of new observations as they were received from NEPTUS. The *depth* grows slower showing that backtracking occurs during execution. Except for the small downward spike, the *depth* tends to either remain flat or climb, indicating that in general backtracking was only impacted by recent decision and hence the planner recovered gracefully. The downward spike was due to erroneous timestamps on some AUV location messages before they reached EUROptus. Despite the large jump in steps ( $\approx 700$ ) the planner recovered in less than 3 s.

## 7 Lessons Learned and Conclusions

In this paper we show how EUROptus, a mixed-initiative planner aids coordinated operation of multiple heterogenous assets with unpredictability of the environment. This work is informed from years of iterating toward a balanced approach between autonomous decision-making and adequate operator supervision. The multi-vehicle operation was carried out by a mix of operators and automated planners onboard the vehicles and at the control station. EUROptus continues to be an experimental system in our attempt to understand the boundaries of autonomy and autonomous operations in the context of oceanographic field deployments. It shows promise in relieving the operator of the intricacies of synchronization between asset operation and focuses on higher-level coordination including scientific goal-achievement and operational safety. A more traditional waypoint-based planning tool, NEPTUS, was not only extended to dispatch commands from EUROptus but to provide enhanced situational awareness with an overview of their expected behaviour to users.



(a) Two AUV trajectories with CTD profiles (red: warmer surface waters, blue: colder deeper waters) executing consecutive co-temporal surveys around a targeted cetacean position.



(b) UAV trajectory (green) superimposed on AUV tracklines (red and blue) from 6(a). UAV tracks show several search trajectories and long-distance tests.

**Fig. 6.** Coordinated observations of a UAV with two AUVs in the water-column on July 19<sup>th</sup> 2015.


The addition of EUROPtus relieves operators of the intricacies of synchronization between assets and focuses more on abstraction such as the pursuance of scientific goals and the safety of the assets. To address the latter, NEPTUS was extended to not only dispatch commands from EUROPtus to the assets but also to intercept them in order to give an overview of expected behavior to the users. Its ability to handle observations with temporal delays and ability to dispatch commands only when all their conditions are observed proved to be effective. Operators consequently, could focus on high level operational concerns and rely on NEPTUS to maintain situational awareness for understanding the current status of the mission.

Among its shortcomings however, is the way *Pending action* flaws are resolved. The only way the system can rectify them is by postponing the action until it has all of its conditions *Justified*. An interesting extension would be for the planner to actively enrich the plan by proactively *asking* for operator input. This 'inquisitive' approach would then prevent the system from replanning just because there was no feedback before a certain deadline.

Another aspect we would like to explore is allowing EUROPtus to forget part of its search when it no longer impacts the plan. A novelty of EUROPtus, when compared to classical approaches, is that it keeps all of its search history. One benefit is that EUROPtus can revisit the past and restore a partial plan fragment in the light of a delayed observation or justification. A side effect is that the receding planning horizon can result in computational search being slower as time advances. A solution to this potential issue would be to prune nodes from the search tree as they are justified by observations. This would require replacing the chronological backtracking search allowing the system to run for sustained periods without a performance impact.

One key insight has been that such a tool is *critical and effective* for operational oceanography even if more work needs to be done. The criticality has to do with the complexity of ship-based operations and with multiple assets with different operational envelopes. Coordinating a team of researchers dispersed in the ship dealing with launch, recovery and operational awareness requires a fine level of coordination which a tool such as EUROPtus can aid with and augment.

A major take-away from this experiment overall, has been that by using autonomous systems to study the oceanographic environmental context would require greater coordination between vehicle operations and marine scientists,



**Fig. 7.** Number of deliberation steps in EUROPtus and its search tree depth for a subset of July 19<sup>th</sup> operations. EUROPtus was running on a Linux Virtual Machine on a 2.8GHz Intel Core i7.

greater capacity of the operators and vehicles to quickly respond to the dynamics of the targets (whether they be animals or fronts) and vehicles and support vessels that can be quickly dispatched and recovered. To be effective, joint field work has to be simplified, operations must be based on small platforms, that can be easily maneuvered, ensuring real-time interaction between technologists and oceanographers.

**Acknowledgements.** Johansen was partly sponsored by the Research Council of Norway through the Centres of Excellence funding scheme, grant number 223254 – NTNU AMOS, and by grant number 235348. Rajan and Py were supported by the US Office of Naval Research, ONR Grant # N00014-14-1-0536. Silva was supported by FEDER and COMPETE funds and by FCT grant # IF/00943/2013. Silva was also supported by FRCT from the Government of the Azores. Authors are grateful to the Cetacean Ecology Research Group including Pablo Chellavard Navarro, Rui Prieto, Cláudia Oliveira, Marta Tobena, and skipper Renato Bettencourt, and to NTNUs UAV team, including Lars Semb, Krzysztof Cisek, Frederik Leira and João Fortuna. This work was partially supported by the SUNRISE - “Sensing, monitoring and actuating on the UNDERwater world through a federated Research InfraStructure Extending the Future Internet”, project #611449, funded by the European Union’s Seventh Framework Programme for Research and Technological Development - Large scale integrating project (IP). Finally, we are grateful to the entire LSTS team that participated in the REP-15 exercise in the Açores.

## References

1. Ryan, J.P., Johnson, S., Sherman, A., Rajan, K., Py, F., Thomas, H., Harvey, J., Bird, L., Paduan, J., Vrijenhoek, R.: Mobile autonomous process sampling within coastal ocean observing systems. *Limnol. Oceanogr. Methods* **8**, 394–402 (2010)
2. Ryan, J.P., McManus, M.A., Kudela, R.M., Artigas, M.L., Bellingham, J.G., Chavez, F.P., Doucette, G., Foley, D., Godin, M., Harvey, J.B.J., Marin III, R., Messié, M., Mikulski, C., Pennington, T., Py, F., Rajan, K., Shulman, I., Wang, Z., Zhang, Y.: Boundary influences on HAB phytoplankton ecology in a stratification-enhanced upwelling shadow. *Deep Sea Res. Part II* **101**, 63–79 (2014)
3. Pinto, J., Faria, M., Fortuna, J., Martins, R., Sousa, J., Queiroz, N., Py, F., Rajan, K.: Chasing fish: tracking and control in a autonomous multi-vehicle real-world experiment. In: *MTS/IEEE Oceans*, San Diego, California (2013)
4. Faria, M., Pinto, J., Py, F., Fortuna, J., Dias, H., Leira, R.M.F., Johansen, T.A., Sousa, J.B., Rajan, K.: Coordinating UAVs and AUVs for oceanographic field experiments: challenges and lessons learned. In: *IEEE International Conference on Robotics and Automation (ICRA)*, Hong Kong (2014)
5. Das, J., Py, F., Harvey, J., Ryan, J., Gellene, A., Graham, R., Caron, D.A., Rajan, K., Sukhatme, G.S.: Data-driven robotic sampling for marine ecosystem monitoring. *Int. J. Robot. Res.* **34**(12), 1435–1452 (2015)
6. Encarnação, P., Pascoal, A.: Combined trajectory tracking and path following: an application to the coordinated control of autonomous marine craft. In: *Decision and Control*, vol. 1, pp. 964–969. IEEE (2001)



7. Arrichiello, F., Das, J., Heidarsson, H.K., Sukhatme, G.S., Chiaverini, S.: Experiments in autonomous navigation with an under-actuated surface vessel via the null-space based behavioral control. In: Conference on Advanced Intelligent Mechatronics, pp. 362–367, July 2009
8. Pinto, J., Martins, P.S.D.R., Fortuna, J., Marques, E., Sousa, J.: The LSTS tool-chain for networked vehicle systems. In: MTS/IEEE Oceans, pp. 1–9. IEEE (2013)
9. Ghallab, M., Nau, D., Traverso, P., Planning, A.: Theory and Practice. Elsevier Science, San Francisco (2004)
10. Chrapa, L., Pinto, J., Ribeiro, M., Py, F., Sousa, J., Rajan, K.: On mixed-initiative planning and control for autonomous underwater vehicles. In: Proceedings of the Intelligent Robots and Systems (IROS), Hamburg, Germany (2015)
11. Ai-Chang, M., Bresina, J., Charest, L., Chase, A., Hsu, J., Jonsson, A., Kanefsky, B., Morris, P., Rajan, K., Yglesias, J., Chafin, B., Dias, W., Maldague, P.: MAP-GEN: mixed initiative planning and scheduling for the Mars'03 MER mission. *IEEE Intell. Syst.* **19**(1), 8–12 (2004)
12. Bresina, J., Jonsson, A., Morris, P., Rajan, K.: Activity planning for the mars exploration rovers. In: International Conference on Automated Planning and Scheduling (ICAPS), Monterey, California (2005)
13. Hodgson, A., Kelly, N., Peel, D.: Unmanned aerial vehicles (UAVs) for surveying marine fauna: a dugong case study. *PLoS ONE* **8** (2013)
14. Cunha, H.A., Sole-Cava, A.M.: Molecular sexing of tucuxi dolphins (*Sotalia guianensis* and *Sotalia fluviatilis*) using samples from biopsy darting and decomposed carcasses. *Genet. Mol. Biol.* **30**, 1186–1188 (2007)
15. Sousa, L.L., López-Castejón, F., Gilabert, J., Relvas, P., Couto, A., Queiroz, N., Caldas, R., Dias, P.S., Dias, H., Faria, M., Ferreira, F., Ferreira, A.S., Fortuna, J., Gomes, R.J., Loureiro, B., Martins, R., Madureira, L., Neiva, J., Oliveira, M., Pereira, J., Pinto, J., Py, F., Queiro, H., Silva, D., Sujit, P.B., Zolich, A., Johansen, T.A., Sousa, J., Rajan, K.: Integrated monitoring of mola mola behaviour in space and time. *PLoS ONE* (2016). Accepted July
16. Py, F., Rajan, K., McGann, C.: A systematic agent framework for situated autonomous systems. In: 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), Toronto, Canada, May 2010
17. Rajan, K., Py, F.: T-REX: partitioned inference for AUV mission control. In: Roberts, G.N., Sutton, R. (eds.) Further Advances in Unmanned Marine Vehicles. The Institution of Engineering and Technology (IET), August 2012
18. Rajan, K., Py, F., Berreiro, J.: Towards deliberative control in marine robotics. In: Seto, M.L. (ed.) *Autonomy in Marine Robots*, pp. 91–175. Springer, New York (2012). doi:[10.1007/978-1-4614-5659-9\\_3](https://doi.org/10.1007/978-1-4614-5659-9_3)
19. Jónsson, A., Morris, P., Muscettola, N., Rajan, K., Smith, B.: Planning in inter-planetary space: theory and practice. In: *Artificial Intelligence Planning and Scheduling (AIPS)* (2000)
20. Muscettola, N., Nayak, P., Pell, B., Williams, B.: Remote agent: to boldly go where no AI system has gone before. *Artif. Intell.* **103**, 5–48 (1998)
21. Dean, T., Boddy, M.: Reasoning about partially ordered events. *Artif. Intell.* **36**(3), 375–399 (1988)
22. Allen, J.: Towards a general theory of action and time. *Artif. Intell.* **23**(2), 123–154 (1984)
23. Frank, J., Jónsson, A.: Constraint-based attribute and interval planning. *Constraints* **8**(4), 339–364 (2003)

24. Sousa, A., Madureira, L., Coelho, J., Pinto, J., Pereira, J., Sousa, J., Dias, P.: LAUV: the man-portable autonomous underwater vehicle. *Navig. Guidance Control Underwater Veh.* **3**(1), 268–274 (2012)
25. Martins, R., Dias, P., Marques, E., Pinto, J., Sousa, J., Pereira, F.: IMC: a communication protocol for networked vehicles and sensors. In: *OCEANS 2009 - EUROPE*, pp. 1–6 (2009)
26. Das, J., Py, F., Maughan, T., Messie, M., O'Reilly, T., Ryan, J., Sukhatme, G.S., Rajan, K.: Coordinated sampling of dynamic oceanographic features with AUVs and drifters. *Int. J. Robot. Res.* **31**, 626–646 (2012)