MOOS-IvP-Multi: A suite of multi-agent autonomy algorithms

Preview of joint 2025 experiment with Ever Gonzalez.
Tomorrow @ 3:20

Tyler Paine: Research Scientist @ Marine Autonomy Lab

MIT Marine Autonomy Lab

MOOS-IvP-Multi: A suite multi-agent autonomy algorithms

Motivation: Complex collaborative autonomy is built from a combination of simpler algorithms in several areas:

- Safety / collision avoidance in large groups
- Communication assurance within teams
- Cooperative search
- Cooperative allocation
- Dynamic formations

We address this need with a suite of software packages:
MOOS-IvP-Multi + Swarm Toolbox (Mike Benjamin) + work by Mike Defilippo

- Multi-agent autonomy baselines for deep research activity on campus
- Uses MOOS and layered on MOOS-IvP
- Accessible to industry

MOOS-IvP-Multi: A suite multi-agent autonomy algorithms

MOOS-IvP-Multi

~ 3.1 work years, ~26000 lines of code (~10-20% LLM generated)

Additionally, 1600+ unit tests. (~1.5 work years, ~10,000 lines of code (human spec + LLM))

Safety in large groups

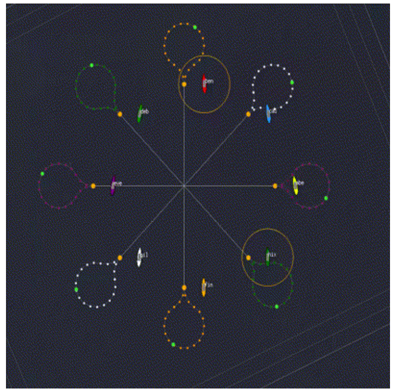
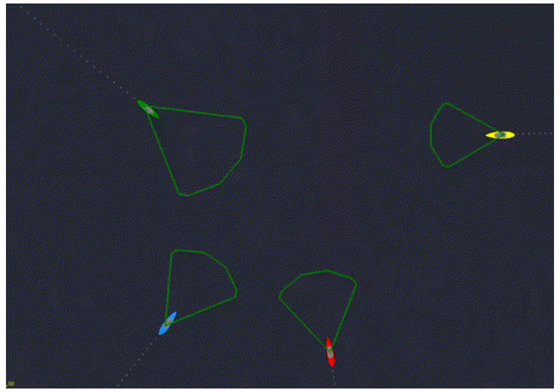

Collaborative allocation and search


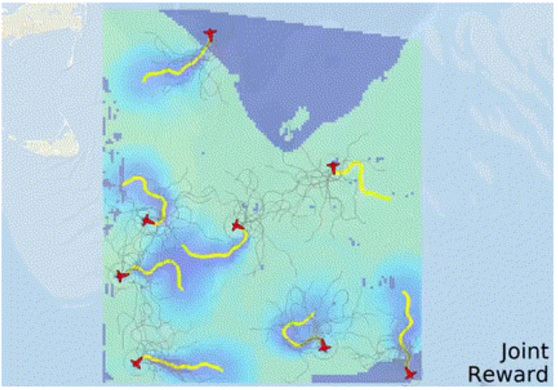

Communication maintenance and team awareness

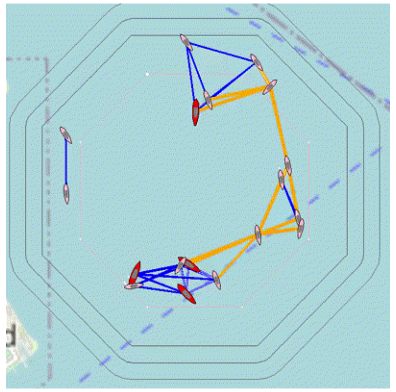
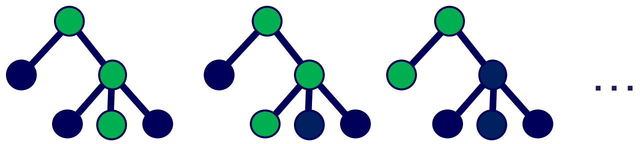

Decentralized team rosters

Collective decision-making via consensus

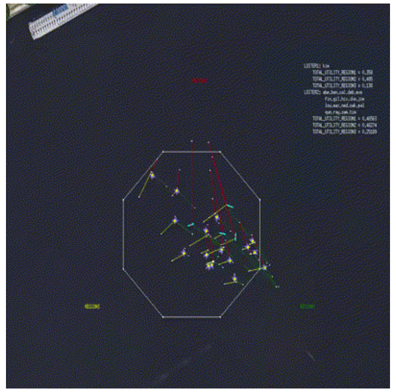
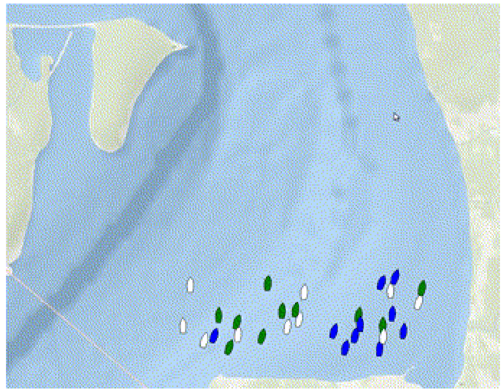

All distributed and decentralized algorithms, not centralized!

Core Area 1: Safety in large groups

Challenge: Collision avoidance in large groups is not trivial and not covered by COLREGs

Control Barrier Functions (CBFs)

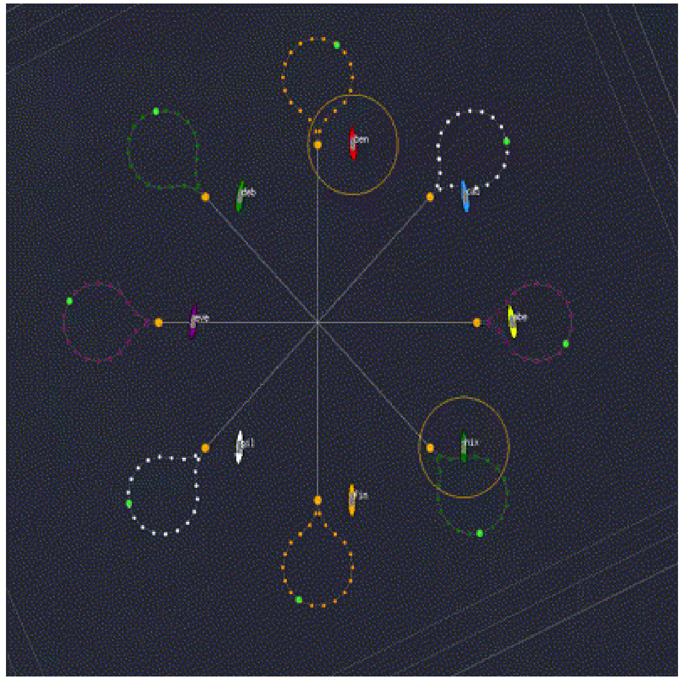

☑ Experimentally tested

Reachability

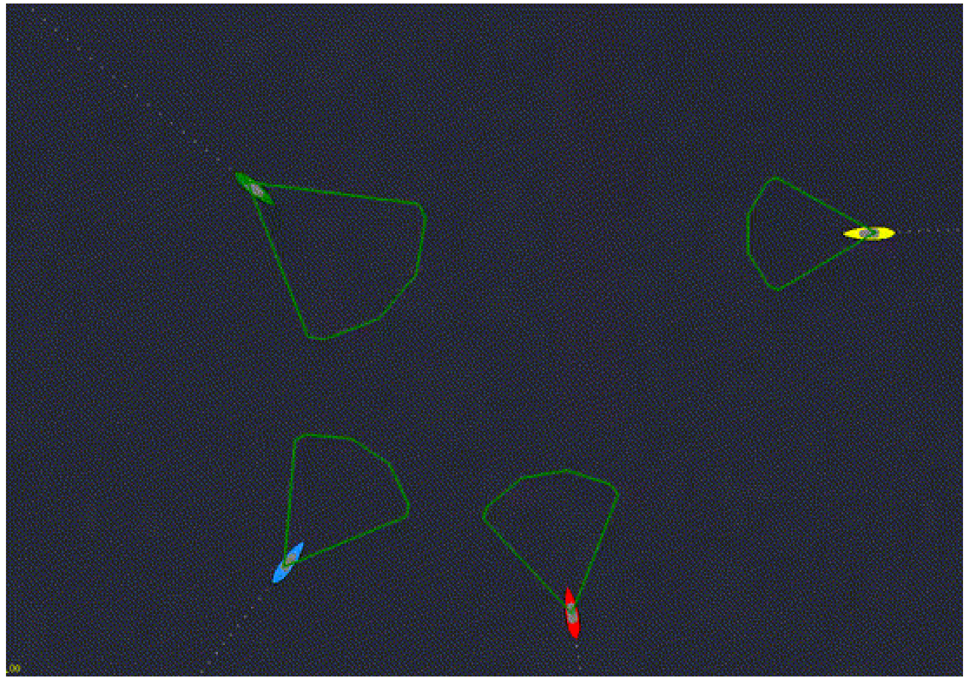

☑ Experimentally tested

- Useful for complex "in extremis" situations with multiple vehicles
- Encapsulates pHelmIvP and controller in a safety filter.
- CBF minimally modifies any dangerous control input to ensure safety.
- Intend to make extendable for your custom CBFs

- Checks for the possibility of a collision over some time horizon
- Configurable flag postings
- User configurable reachable set based on experimental testing of vessel dynamics

Core Area 2: Collaborative allocation and search

Challenge: Cooperation in groups is complicated, especially in a dynamic environment.

Dynamic group allocation via combinatorial optimization

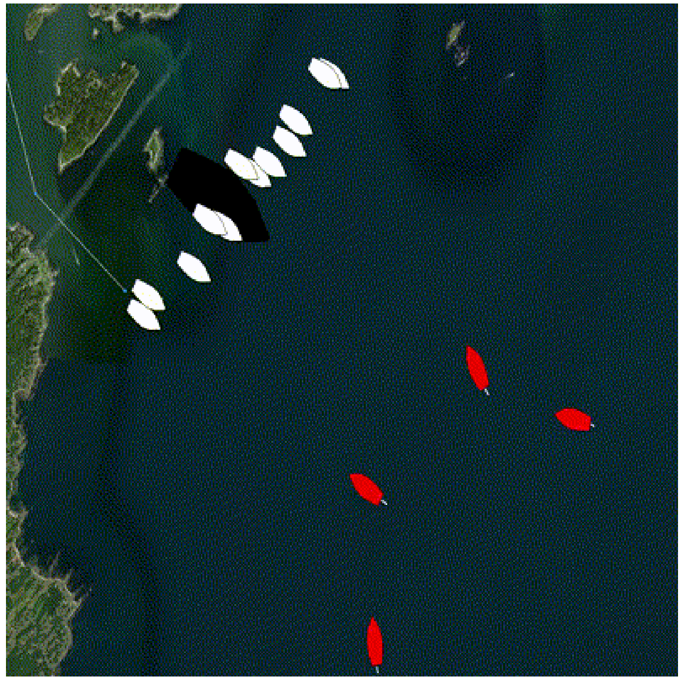

☑ Experimentally tested

Cooperative adaptive search and sampling

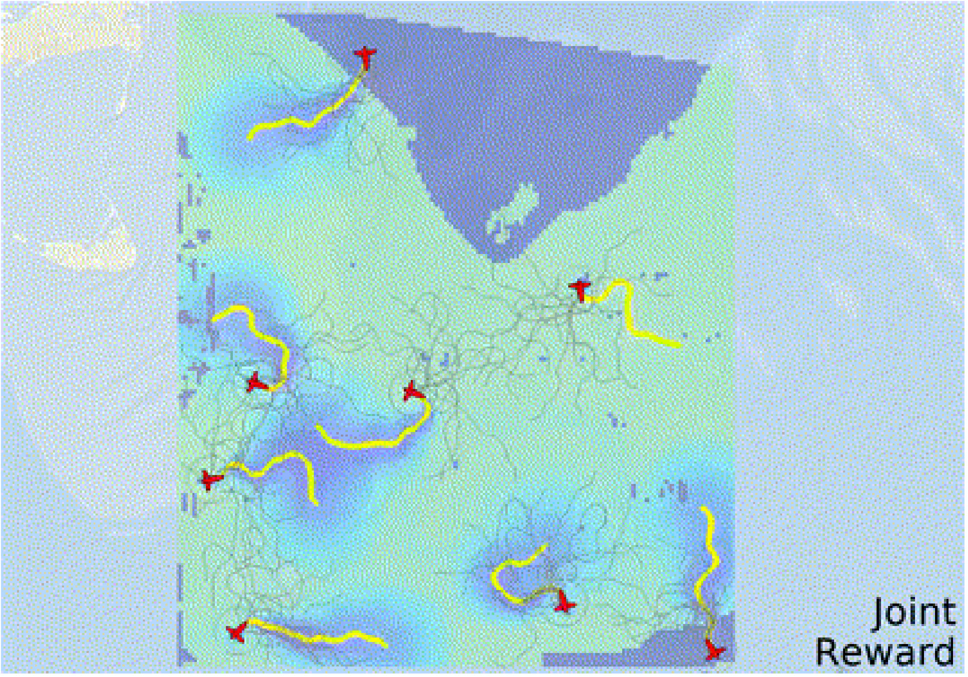

- Solves the problem of matching N teammates with M target locations
- Collectively minimize the total team cost
- Adjusts to time varying conditions

- Enables efficient cooperative search for features
- Each agent builds a global map using input from neighbors
- Path planning solved via a robust, scalable Monte-Carlo Tree Search

Core Area 3: Communication maintenance and team awareness

Challenge: Maintain communication networks and understand who is on my team

Network communication maintenance

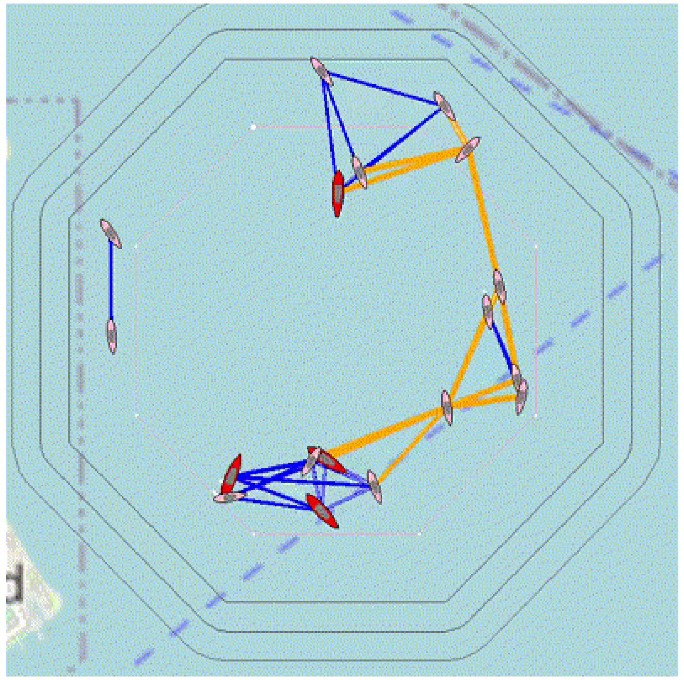

Decentralized team roster

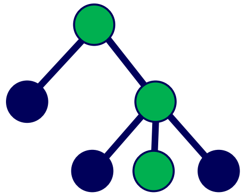

Configurable tree sets own team assignment
Ex: TEAM=TEAM1:ALPHA_PLATOON:ELEMENT_1

App assembles and shares a TEAM_REPORT
Ex: TEAM_REPORT=abe@TEAM1:ALPHA_PLATOON:...

App collects TEAM_REPORTs from others and posts rosters
Ex: TEAM1_LIST = abe,ben,cal,deb,eve,fin, ….
ALPHA_PLATOON_LIST = abe,deb,cal,eve, …
BRAVO_PLATOON_LIST = ben,fin, ...
ELEMENT_1_LIST = abe,eve

- Reconstructs the network topology
- Identifies weak links critical to maintaining communication
- Behavior aims to maneuver vehicles to maintain network connections

- User specified logic tree to determine own team membership
- Shares team structure with teammates
- Reconstructs team rosters from neighbor reports.

Core Area 4: Collective decision-making with a family of consensus protocols

Challenge: Consensus about collective decisions in groups with limited communication

Continuous nonlinear coupled consensus/dissensus

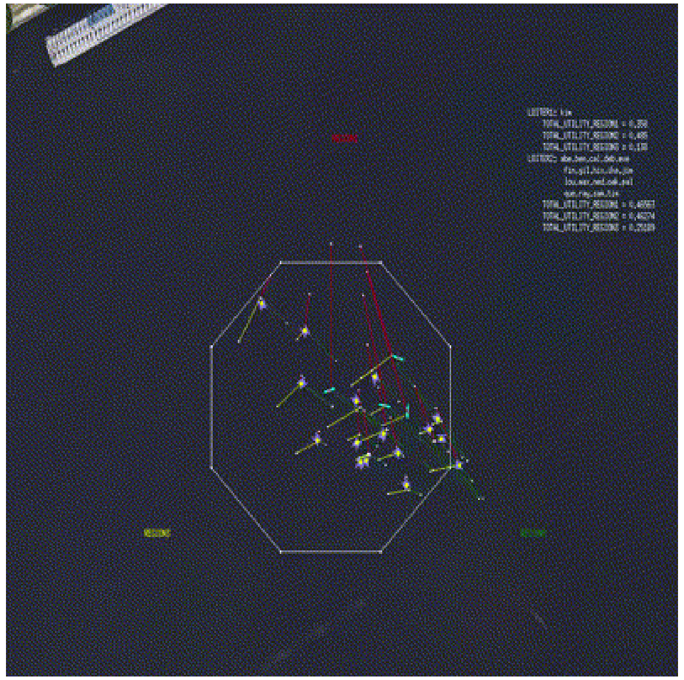

☑ Experimentally tested

- Rapid distributed decision-making over coupled/ranked options
- Works even if some agents cannot observe all costs, reducing communication
- BHV_Flock for special and common case of consensus on speed and heading.

Discrete consensus protocols

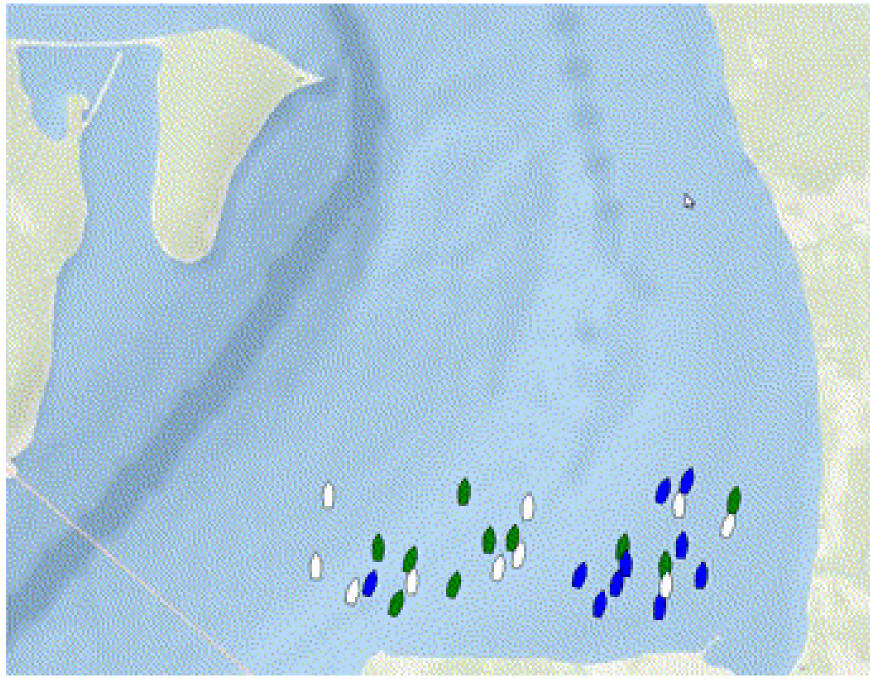

- Execute protocols like “Best of N” and minimum cost.
- Configurable run conditions, Active/Runflags, and conditional end flags.
- Versatile, handles multiple simultaneous protocols.
- Intend to make extendable for your custom protocols

Future directions for MOOS-IvP-Multi

In industry:

Build capability by combining elements

Build out more baseline methods +

Safety

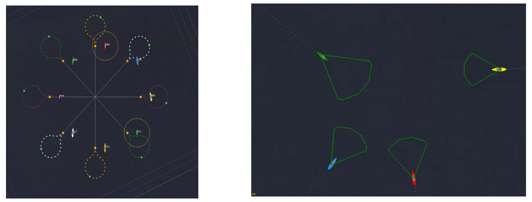

Allocation and search

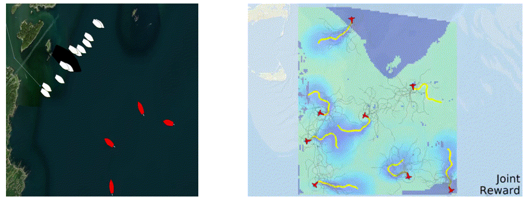

Communication

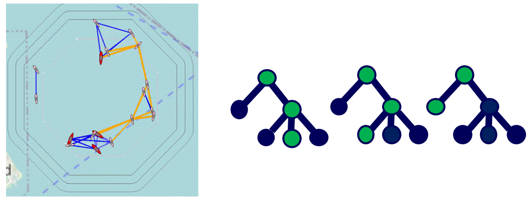

Consensus

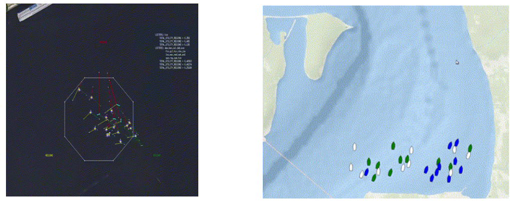

In academia: Use as baseline for deeper research questions

T. M. Paine and M. R. Benjamin, "A Model for Multi-Agent Autonomy That Uses Opinion Dynamics and Multi-Objective Behavior Optimization," 2024 IEEE International Conference on Robotics and Automation (ICRA), Yokohama, Japan, 2024, pp. 8305-8311
T. M. Paine, A. Bizyaeva, M. R. Benjamin, “Adaptive bias for dissensus in nonlinear opinion dynamics with application to evolutionary division of labor games,” 2024 IEEE 63rd Conference on Decision and Control (CDC) 2024 Dec 16. pp. 2443-2450
T. Paine, B. Long, J. Wenger, M. DeFilippo, J. Usevitch, M. Benjamin, “Distributed Control Barrier Functions for Safe Multi-Vehicle Navigation in Heterogeneous USV Fleets,” arXiv preprint arXiv:2601.11335. 2026 Jan 16.
T. M. Paine, A. Bizyaeva, M. R. Benjamin, “Census-Based Population Autonomy For Distributed Robotic Teaming,” arXiv preprint arXiv:2511.02147. 2025 Nov 4.

MIT Marine Autonomy Lab